UNIVERSITY OF CALIFORNIA
Santa Barbara

Re-architecting Internet Access and Wireless Networks for Rural Developing Regions

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Computer Science

by

David Lloyd Johnson

Committee in Charge:

Professor Elizabeth Belding, Chair

Professor Kevin Almeroth

Professor Heather Zheng

March 2013

The Dissertation of
David Lloyd Johnson is approved:

Professor Kevin Almeroth

Professor Heather Zheng

Professor Elizabeth Belding, Committee Chairperson

March 2013

Re-architecting Internet Access and Wireless Networks for Rural Developing Regions

by

David Lloyd Johnson

Acknowledgements

Funding for this work was in part through NSF NetSE award CNS-1064821 and a grant from Google. Visits back to Africa during my PhD studies were funded by the Council for Scientific and Industrial Research, Pretoria, South Africa. The dissertation work was also supported by the University of California Santa Barbara Office of International Students and Scholars Presidents Work-Study Award.

This PhD would not be possible without the support I had from a wide range of people. Firstly, I would like to thank my incredible wife, Nula, who has been at my side through this whole journey. The many coffees talking at Zizos in Goleta about my PhD work and her growing vision to use her legal degree to address human rights through the Internet helped spur me on. We have three beautiful boys, Samuel, Luke and Benjamin (born during my PhD) and being able to navigate them through all the changes while moving back and forth between South Africa and the USA was tough. Nula home schooled them during our first year in America to help them through the transition. While I was buried in my studies, Nula often had to bear the brunt of their emotional stress. Samuel, Luke and Benjamin, I can't imagine life without you guys, you have kept me youthful, less self-absorbed, creative and helped me to see things through fresh young eyes.

To my Mom and Dad, you are where this PhD really began. You instilled a love for learning in me from when I was just beginning to crawl. Mom, for all the times you spent encouraging me to build puzzles when I was only three and taking me to museums and science centres. Dad, you have been pivotal in teaching me to think and thoroughly investigate a topic before accepting anything as truth, both in the areas of science and faith. A big thanks to my Dad for also doing the final detailed editing on this thesis.

My advisor, Elizabeth, has played a critical part guiding me through this PhD. From my first day in the USA she helped my family and me settle into the local environment. Within my first year at UCSB, she had helped me to formulate my research area with a perfect balance between allowing me to explore new areas that I am passionate about and clear guidance on what areas are likely to provide a solid platform for my PhD. Elizabeth set a very high standard for my research and writing, always pushing me to do better and aim higher. She provided expert advice from the conceptual phase of my papers until the last few adjustments. When I look back at the researcher I was 4 years ago and the person I am today, I can see significant changes in the way I think critically as a researcher and Elizabeth had a strong part to play in that. I would also like to thank my other Committee members: Kevin Almeroth for opening the gateway for my interests in ICT for development with his technology and society seminar and solid advice on formulating a research area and Heather Zheng for her inputs on cognitive radios and dynamic spectrum access.

I have had four of the most interesting and rewarding years of my life working with students in the moment lab at UCSB. A special thanks to Veljko for all the work we did together on analysing rural Internet traffic and running white space experiments. I will never forget our drives up into the misty mountains above Goleta to set up battery powered white space transmitters or long hours around the white board thinking about some measurement or networking problem. Thank you Mariya for roughing it with me in Zambia, hanging off lattice masts to set up VillageCell base stations and staying strong in the midst off the massive challenges that are unique to rural Africa. Prashanth, thanks for your great work on the mesh gateway protocol. Mike, for your abstract bigger-picture thinking. Ashish for your deep insights into all things MAC. Stephan, for helping me understand Teo Gonzales complex course on NP-completeness. Sumit, without you around much of the digital communication world would have remained a mystery, Daniel, for helping me understand all things American from corn dogs to football and a whole bunch of TCP insights thrown in. Lara, for agreeing to work with me in the very first quarter of our PhDs on my crazy mesh gateway idea in Ben's class — you were calm even in the final hour when I was freaking out. Thank you to the great masters students who agreed to work with me on my rural networking project ideas: Stratos, Ana, Maggie, Sheena, Abhinav, Arghyadip, Rohit, Sathiish and Devdeep — without you much of this work would have remained on the list

of things to do. Thank you to Sarah for taking over the reigns of the VillageShare project — I look forward to watching your unique insights that you will add to this project in the coming months. The diversity of people in this lab who represented India, Bulgaria, Serbia, Lebanon, Poland and the USA have provided me with deep insights that go far beyond my field of interest in my PhD.

This PhD embraces a field that involves far more than technology or computer science or a set of quantitative measurements. It involves thinking about human needs and involves love, caring and self-sacrifice for the marginalised of society. Gertjan has been a mentor on every level in this key part the work, as somebody who can bridge and interpret rural Africa, as an ambassador for ICT in rural Africa and a Christian with a heart for giving people hope. Gertjan may be one of the most important voices in the field of ICT4D after living for a decade in rural Africa. In the past few years, he has helped me constantly adjust my assumptions and even my innate prejudices about rural Africa and been a good friend to have around even if it was often over Skype — our Skype conversations about our struggles and hopes could probably fill a thesis chapter and more.

Much of this work was done in collaboration with LinkNet in Macha, Zambia. A particular thanks goes to Fred Mweetwa, CEO MachaWorks, for facilitating many of our interactions with local people living in Macha. Fred has been an inspiration to me — his unwavering commitment to transforming rural communities through technology such as community radio and the Internet and inspiring local talent in the midst of massive challenges in rural Zambia have challenged my definition of what is possible in rural areas.

The CSIR Meraka institute provided additional financial support to make it possible for me to study towards a PhD in the USA while having my family with me; without them this PhD would simply have been impossible. Before my PhD work I had many rewarding years working on rural wireless networking at the CSIR with some colleagues who continue to do great work in the field: Kobus Roux was key to our original vision of reaching rural areas with the Internet back in 2004 and has continued to do excellent work in the field, initiating the "Broadband for All" project and other projects to reach rural schools with the Internet — many that will form great vehicles for reinvesting my PhD work back into South Africa. John Issac helped formulate our vision for reaching rural Africa and was always a good person to digest ideas with in the Meraka tea room. Ntsibane Ntlatlapa, now head of the Networks and Media group at the CSIR, was very supportive throughout my PhD, facilitating my return visits to CSIR and spearheading some of the new white space initiatives in South Africa that will also form part of my post-PhD work. Albert Lysko, has always been a great scientist to collaborate with on anything to do with electromagnetics and we have continued to collaborate on antenna aspects of wireless communication that computer scientists like me find

very daunting to understand. Kagiso Chikane (then head of Meraka) agreed to let me go for four years and support me during my studies. Laurens Cloete (now head of Meraka), has always been one of the best sounding boards for ideas - especially my early thoughts on what I will do after my PhD. To my good friend Kim Gush, I missed our fun creative times in our shared office which took the sting out of the sometimes rather dull world of engineering and computer science; it was good to enjoy some of this again on my visits to South Africa and thanks for hosting that crazy Serbian, Veljko, over the summer break. I also wish to thank many of my other CSIR collaborators: Fisseha Mekuria, Moshe Masonta, Thomas Olwal, John Hay, Johann Hugo, and Ajay Makan.

Ultimately the purpose and vision for my life, of which this PhD is part, and the very breath that I breath can only be found in HaMashiach Yeshua — Jesus the Messiah, the King of Kings, Lord of Lords, the light of the world. My safety and salvation, my light in the darkness when all other hope is faded and gone, in him I live and move and have my being. To him be the glory for ever and ever Amen.

Curriculum Vitæ
David Lloyd Johnson

Education

1995	Bachelor of Science in Electronic Engineering, University of Cape Town
2007	Master of Science in Computer Engineering, University of Pretoria
2012	Master of Science in Computer Science, University of California, Santa Barbara

Experience

1996-1998	Software engineer, Information Delivery Software, Icomtek, CSIR; Pretoria, South Africa
1998-2003	Telecommunications engineer, IST centre, CSIR, Pretoria, South Africa
2003-2005	Technology group leader, Mobile platform technologies, Meraka, CSIR, Pretoria, South Africa
2005-2007	Acting research group leader, Wireless Africa research group, Meraka, CSIR, Pretoria, South Africa
2008-2010	Teaching Assistant, University of California, Santa Barbara
2010-2013	Research Assistant, University of California, Santa Barbara

Selected Publications

D.L. Johnson, "Evaluation of a single radio rural mesh network in South Africa", *ICTD'07*, Bangalore, India, December 2007

D.L. Johnson and A. Lysko, "Comparison of MANET Routing Protocols Using a Scaled Indoor Wireless Grid. Mobile Networks and Applications", *Mobile Networks and Applications*, 13 (2007), pp. 82–96

D.L. Johnson, N. Ntlatlapa and C. Aichele, "A simple pragmatic approach to mesh routing using BATMAN", *IFIP International Symposium on Wireless Communications and Information Technology in Developing Countries*, Pretoria, South Africa, September, 2008

D.L. Johnson and K. Roux, “Building rural wireless networks: lessons learnt and future directions”, *WINS-DR’08*, San Francisco, CA, September, 2008

D.L. Johnson and G. Hancke, “Comparison of two routing metrics in OLSR on a grid based mesh network”, *Ad Hoc Networks*, 7 (2008), pp. 374–387

D.L. Johnson, E.M. Belding, K. Almeroth, and G. van Stam, “Internet usage and performance analysis of a rural wireless network in Macha, Zambia”, *NSDR’10*, San Francisco, CA, June, 2010

P.A.K. Acharya, D.L. Johnson, and E.M. Belding, “Gateway-aware Routing for Wireless Mesh Networks”, *Meshtech’10*, San Francisco, November, 2010

D.L. Johnson, V. Pejovic, E.M. Belding, and G. van Stam “Traffic characterization and internet usage in rural Africa”, *WWW’11*, Hyperabad, India, March, 2011

D.L. Johnson, V. Pejovic, E.M. Belding, and G. van Stam, “Network traffic locality in a rural African village”, *ICTD’12*, Atlanta, GA, March, 2012

A. Anand, V. Pejovic, E.M. Belding, and D.L. Johnson, “VillageCell: cost effective cellular connectivity in rural areas”, *ICTD’12*, Atlanta, GA, March, 2012

D.L. Johnson, V. Pejovic, E.M. Belding, and G. van Stam, “VillageShare: Facilitating content generation and sharing in rural networks”, *ACMDEV’12*, Atlanta, GA, March, 2012

M.T. Masonta, D.L. Johnson, and M. Mzyece, “The White Space Opportunity in Southern Africa: Measurements with Meraka Cognitive Radio Platform”, *e-Infrastructure and e-Services for Developing Countries*, 2012, pp. 64–73

V. Pejovic, D.L. Johnson, M. Zheleva, E.M. Belding, L. Parks, and G. van Stam, “The Bandwidth Divide: Obstacles to Efficient Broadband Adoption in Rural Sub-Saharan Africa”, *International Journal of Communication*, 6 (2012), pp. 2467–2491

V. Pejovic, D.L. Johnson, M. Zheleva, E.M. Belding, and A. Lysko, "VillageLink: Wide-area Wireless Coverage", *Mobihoc'13*, Bangalore, India, July, 2013

M. Zheleva, A. Paul, D.L. Johnson, and E.M. Belding, "Kwiizya: Local cellular network services in remote areas", *Mobisys'13*, Taipei, Taiwan, June, 2013

Abstract

Re-architecting Internet Access and Wireless Networks for Rural Developing Regions

David Lloyd Johnson

Internet connectivity is a key enabler of the socio-economic development of any nation. There are very few aspects of life that have not been impacted positively by access to the Internet; areas ranging from education to health and even the democratic fabric of society. Many countries such as France and Greece have recognized the Internet's intrinsic egalitarian nature and enshrined it as a basic human right. A 2012 International Telecommunication Union survey of Internet penetration [12] revealed that 34.1% of households in the world have Internet access. However, a strong access divide exists across the world. Developed countries like the USA have 76% of households connected to the Internet whereas many African countries such as Zambia and Mozambique have less than 3% of households connected. This disparity has created a "digital divide" that separates the affluent and developed nations from the developing and under-developed regions of the world.

In this work we focus on unique connectivity challenges and opportunities in rural areas, especially in developing regions. Providing access in rural areas is

particularly challenging due to its unique set of social and technical constraints. In order to quantify the challenges being faced by Internet users in rural areas, we conducted on-site and on-line interviews and analysed network traces from a rural network in Macha, Zambia. Our findings reveal severe local and global connectivity limitations that are unique to rural areas. Local wireless networks experience interference and packet loss due to poor network designs and limitations of WiFi in rural areas. Bandwidth-restricted Internet gateways are constantly congested during normal daytime usage periods. Users describe cost, limited availability and unreliability of Internet access as key barriers to on-line interaction. These obstacles prevent leisurely access to the Internet, including, amongst others, users generating and sharing media. This leads to a more transaction-like or "deliberate interaction" model [172] for many rural users.

Usage analysis reveals unique behaviour for users in rural areas. Most traffic is web-based traffic as opposed to peer-to-peer traffic in developed countries. Social media features even more prominently than in western countries, with Facebook being the most visited website; users are twice as likely to access Facebook than Google search. Further analysis of Facebook traffic shows a high degree of traffic between local users in the village. These unique patterns and constraints form the basis for ICT solutions we propose for rural regions. In order to avoid using congested rural Internet gateways, we propose a set of techniques to localize traffic.

VillageShare, a Facebook-based localization application, facilitates file sharing amongst users in the village without the need to send media over the bandwidth-constrained gateway. VillageShare is also capable of time-shifting uploads to off-peak usage periods in order to avoid upload failures. In order to exploit locality of interest in mobile phone usage, we designed VillageCell, an open-source, low-cost pico-cell-like base station, that allows users to make free local cellular calls in a village. It takes advantage of the very high penetration rate of mobile phones in rural Africa [11], which occurs even though many villages lack cellular coverage. VillageCell is also able to support SMS to instant message client exchange as well as routing calls between VillageCell phones and phones on the public switched telephone network.

The low population density and large village diameters in rural areas of Africa imply a need for novel solutions to spread wireless connectivity to individual homes. Current solutions, based on 802.11 require clear line-of-sight and have limited range. Those based on WiMax are not suitable due to high licence and deployment costs. We propose to use the recently freed TV spectrum bands known as white spaces, encompassing frequencies from 52MHz to 698MHz, to cover vast distances in rural areas. Our solution, VillageLink, builds on the existing 802.22 white space standard to optimally utilize white space spectrum. We add a feature that allows base stations to allocate optimal channels using inter-cell probing across all available TV channels. Channel probing is critical as frequency selectivity is more dependent on antenna characteristics and non-linearity of RF components in the system than on the free-space propagation laws in the white space band.

In order to evaluate our interventions, we deployed VillageShare and VillageCell in Macha, Zambia and evaluated VillageLink using simulations built on a real 3 km white space link in South Africa. We collected usage logs of these systems both through quantitative and qualitative studies. For qualitative studies we involved users in an iterative design process and made use of on-line interviews to understand user perception of our solutions. Ultimately, we hope that this research will lead to better penetration of wireless networks and improved network performance for users in rural villages, culminating in a more inclusive and representative Internet that truly reflects all languages and cultures in the world.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

Internet connectivity is a key enabler for socio-economic development of any nation [25]. It has become an assumed commodity for almost all spheres of life from public services to education to business-related activities in the modern industrialized world. Recent events in the Arab world have also demonstrated the democratizing force of the Internet. This contrasts sharply with the abuse of historical one-way media such as television and radio to broadcast government propaganda. As a result, many countries, such as France, Greece and Spain, have recognized the intrinsic egalitarian nature of the Internet in a modern society and enshrined it as a basic human right.

At the World Summit on the Information Society in 2005, the Association for Progressive Communication proposed that Internet access should be universal and affordable; they argued [157]:

> The Internet is a global public space that should be open and accessible to all on a non-discriminatory basis. The Internet, therefore, must be seen as a global public infrastructure. In this regard we recognize the Internet to be a global public good related to the concept of the common heritage of humanity and access to it is in the public interest, and must be provided as a global public commitment to equality.

According an International Telecommunication Union report [11], two-thirds of the world still don't have access to the Internet creating a "digital divide" between developed[1] and developing countries. In Africa the situation is even worse where broadband connectivity is almost completely missing. Only 4% of the population have mobile broadband subscriptions and 1% of the population have fixed (wired) broadband subscriptions [11]. Almost all of these subscriptions are in urban areas, leaving rural areas, which account for 70% of the population, mostly without any Internet connectivity. Sub-Saharan African countries like Liberia, Niger and Ethiopia have the lowest penetration rates in the world with less than 1% of the population being connected to the Internet. With the "digital

[1]The term "developed countries" is used to denote 47 countries that according to the November 2011 release of the United Nations Human Development Index (HDI) are awarded a very high HDI value `http://hdr.undp.org/en/media/HDR_2011_EN_Table1.pdf`. All other countries are termed developing countries. To avoid broad generalization, a group of emerging national economies: Brazil, Russia, India, China and South Africa, known as BRICS, distinguish themselves by their large, fast-growing economies.

divide" being so pronounced in the African continent and Sub-Saharan Africa in particular, we focus our attention on this region although our solutions are equally applicable to other developing regions of the world.

Although the current trend of increasing available bandwidth in Africa through undersea cables helps boost the available aggregate capacity in Africa, it does not do much for many of the landlocked countries and rural regions in Africa. In most cases a rural village is up to a thousand kilometres away and the cost of covering this distance to reach scattered rural villages is too high. As a result many rural deployments rely on expensive satellite links to provide Internet access to a village with low-cost licence-free commodity hardware such as WiFi to distribute the signal. The cost of Internet connectivity remains a major constraint to Internet connectivity in rural areas where people often live on less than a dollar per day.

There are many obstacles to overcome beyond cost when providing Internet access to rural villages. Energy supply is the most essential part of any communication infrastructure, and yet in many rural regions there is either no electricity grid or electricity is unreliable. Fibre back-haul is usually completely absent in rural regions and this can only be overcome with long-range wireless links to nearby cities or expensive satellite links. Spreading the signal from a well-provisioned Internet "point of presence" to a rural population is difficult as population densities are low. Once infrastructure is put in place to overcome these obstacles, well-trained personnel are often not available to maintain it. Finally, regulation in Africa is often either ambiguous or restrictive making it difficult for communities to self-provision or for small businesses to build communications infrastructure.

In spite of these obstacles, there have been a number of projects which attempt to provide access infrastructure in challenging rural regions around the world by NGOs and University research groups [20, 160, 89, 117]. Most of these projects have used modified commodity hardware, such as WiFi, to reach distant communities in difficult terrains such as the Airjaldi network in the mountains of Dharamsala, India [160], the semi-rural community of Peebles Valley in the east of South Africa [89] or the deep-rural community in Macha, Zambia [117]. Although essentially the same technology was used, the success of these networks has been mixed. Successes were almost exclusively due to training and involvement of local community members or nearby support structures .

Obstacles to widespread adoption of Internet in peoples homes, in many cases, are technical in nature with line-of-sight technology such as WiFi not being well suited to areas with dense vegetation or low population densities. Severely bandwidth-limited and expensive connectivity available to users remains another key constraint to full digital inclusion. Although some users have access to a computer connected to the Internet, even if not in their home, they now form part of a “bandwidth divide” with users who are in well-provisioned networks in developed nations. This leads to less locally-generated content on the Internet from rural regions of the world and difficulties when trying to collaborate internationally due to issues such as email attachments failing to upload or poor quality VoIP calls. From these observations, it is clear that much remains to be done to make rural communities global citizens of the digital age.

1.1.1 Problem statement

The obstacles to ICT adoption in rural areas include technical, environmental and cultural barriers. There is a general lack of substantive knowledge from long-term contextual analysis of ICT use in rural areas and research and development of ICT needs to be informed by such studies.

1.2 Dissertation overview

Addressing the unique challenges of rural or developing regions using ICT requires cross-disciplinary research that sensitizes the researcher to the unique technical and anthropological aspects of a rural community. There are many unique aspects of rural areas, highlighted in Section 1.1, that need to be assimilated into designing solutions for rural areas. Collaborative research is essential

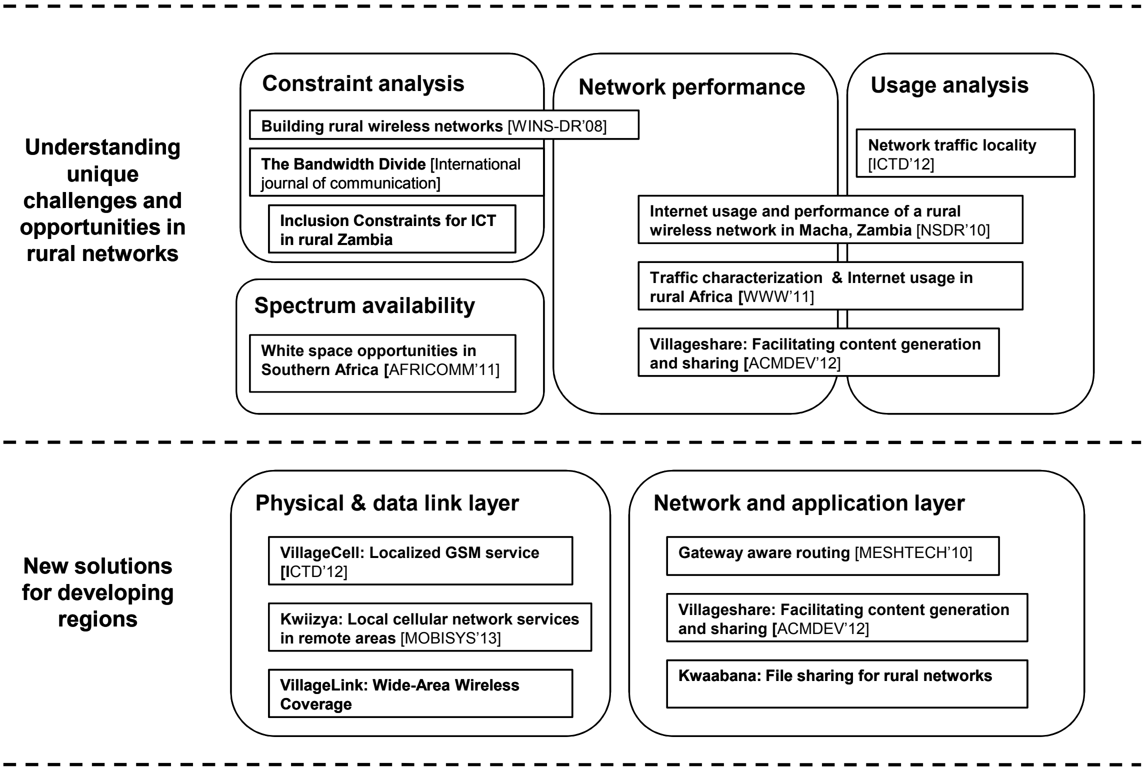

Figure 1.1: Solutions are carefully informed by constraint analysis as well as in-depth quantitative and qualitative network performance and usage analysis of a network being used by a rural community. The solutions for rural regions address both a physical layer design customized for sparse rural populations as well as network and application layer changes to optimize the utilization of scarce Internet bandwidth at the gateway.

to the analysis and design process and our work is done in close partnership with collaborators living in an African village.

Figure 1.1 gives a broad outline of the structure of this dissertation. In this dissertation we approach the problem in two phases. During the first phase we make use of network traffic analysis, and on-line and on-site interviews to understand the unique challenges and opportunities in rural networks. This analysis helps to suggest a number of key new unique research areas to address when building solutions for developing regions. At the physical and data-link layer we develop a number of novel solutions for voice and data access in rural areas. At the network and application layer we develop a series of tools which make use of the unique traffic patterns in rural areas to optimize Internet access in rural areas. Ultimately we want community members to become empowered to make use of the outputs of this research to enhance the utility of their ICT services. In the spirit of "participatory action research" [28] we see community members as research partners rather than objects of research. Inquiry is not simply a dispassionate undertaking that only seeks objective knowledge but rather includes a form of social action. We avoid rash commitments to providing the "ultimate solution" and rather see these ICT interventions as stepping stones on a long-term path towards complete inclusiveness of rural areas in the digital age.

1.2.1 Unique challenges and opportunities in rural networks

There is a growing awareness by researchers, governments and practitioners of the unique challenges that exist when implementing ICT in developing regions. The United Nations started a series of conferences called the World Summit on the Information Society (WSIS) in 2003 in response to the growing digital divide that separates developing and developed countries. In the research community, the use of ICT in developing regions is a growing but small discipline. A number of conferences in the field such as ICTD, NSDR and ACM DEV are focussed on research related to ICT for development. There are also a few multidisciplinary focus groups specializing in different aspects of the use of ICT in developing regions at Georgia Tech, University of Washington, UC Berkeley and University of Cape Town, for example.

There are a few sporadic case studies of ICT deployments in rural areas[150, 160, 127, 89, 117] but not enough to formulate generalizable usage behaviour or characteristics. Broad-based conclusions for different developing regions may not be useful due to the cultural and environmental differences between countries or even within countries. India has a far higher population density than Africa, for example. South Africa has a far higher degree of electricity and mobile phone penetration than Zambia. Even within a country like South Africa, there are large economic and cultural divides between urban and rural populations. Generalizable results may only be useful in some limited areas such as dealing with capacity constraints. However, cultural, economic, infrastructural and political specifics of different rural areas should influence choice of solutions in each area.

This warrants the need for detailed analysis of user needs and usage behaviour in the community under study. This work begins with a study of the unique technical and social constraints in rural regions in Southern Africa as well as opportunities that are unique to these rural regions. This is followed by detailed analysis of an existing wide-scale deployment of Internet connectivity in an African village. Many aspects of this analysis are generalizable to other developing regions beyond Macha. For example a comparative study of Cambodia, Iraq and Zambia using our data traces found a similar high degree of locality in Facebook traces [53].

Constraint analysis

Rural networks, especially those in developing countries, have unique challenges that set them apart from urban areas and many developed countries. Some key challenges are long distances between wireless devices, single low-bandwidth gateways to the Internet, high cost of Internet connectivity and lack of reliable power. There are also non-technical challenges when installing and maintaining technology in developing regions [26, 117]. These include transportation issues, tampering and theft (in some areas), low technical skill level of people in rural areas and political, language and cultural barriers.

We provide a set of observations from three connectivity projects in Africa. One which had limited success in Peebles Valley in the east of South Africa, a community that formed part of the author's previous research, and two which continue to provide connectivity in Dwesa, South Africa and Macha, Zambia. The Macha community was the only rural network wholly maintained and run by local community members due to an ICT training facility housed within the village. The Peebles Valley network eventually ceased operation due to a lack of a local champion and regulation barriers to running community-owned networks. Reflecting on the network in Macha after 5 years of operation exposed similar challenges faced by other rural networks but also revealed a number of unique technical issues: Windows outgoing throughput was far worse than Unix-based operating systems due to its delay-sensitive TCP protocol while power failures created network downtime often twice as long as the power outage itself due to lack of device immunity to constant power cycling. Unique social issues were also uncovered: elders expressed concern about the threat of modern technology on their culture, prioritizing services such as electricity supply to ICT facilities sometimes created tension and jealousy in communities and oral cultures in rural

Zambia prevented users from receiving the full benefit of Internet. Some of these issues inform our solutions whereas others are presented to help future researchers and practitioners adapt their solutions to local rural conditions.

Spectrum availability

"White spaces" promise to revolutionize the way connectivity is provided to rural areas. "White spaces" encompass frequencies reserved for television broadcasting from 52 MHz to 698 MHz in the USA and potentially 52 MHz to 800 MHz in ITU region 2 areas such as Africa. This spectrum provides excellent propagation properties and range in challenging rural terrain where vegetation or mountains often make "line of site" communication difficult. In 2008 the FCC opened white spaces for unlicensed use and similar efforts are beginning to occur in other parts of the world. In urban areas of the USA, such as San Francisco, only 37% of the TV spectrum is vacant whereas in rural towns, such as Juneau, Alaska, 74% of the TV spectrum is vacant [30]. In Africa the amount of available spectrum is even more attractive. In urban areas of South Africa, such as Pretoria 70% of the TV spectrum is vacant whereas in rural towns, such as Phillipstown in the Northern Cape, 97% of the TV spectrum is vacant [116]. This abundance of white space spectrum in rural areas of Africa presents a unique opportunity to provide rural connectivity solutions in this spectrum.

We carried out detailed spectrum scans in urban and rural areas of the USA as well as urban and rural areas in Africa and found many rural areas in Africa with almost no TV spectrum usage. This provides a very compelling case for white space technologies. White spaces are very well suited to rural areas in Africa that have very sparse population density and large distances between clusters of people in villages. Current wireless connectivity options, such as WiMax, 802.11 and 802.22, do not meet all the requirements of a viable communication solution in rural networks; these are: low cost, license-free, long range, non-line-of-sight, easily deployable, spectrum agile and energy efficient.

Network performance

There are a number of rural wireless networks that have been operating since 2005. Some key examples are Airjaldi in India [160], the Dwesa network in South Africa [126] and the LinkNet wireless network in Zambia [117]. These networks are unique in that they need to overcome challenges such as long distances between wireless nodes, low-bandwidth gateways to the Internet, lack of reliable power and high cost of Internet connectivity. Studies of 802.11 networks [13, 21, 33, 140] have shown a number of unique performance issues for wireless deployments in urban areas such as route flapping, co-interference, problems with default auto-rate selection and performance degradation due to the requirement for layer 2 acknowledgements. However, less is understood about wireless networks in rural areas.

In this work we analyse a set of Internet gateway trace logs, proxy access logs and OLSR-based mesh network logs from the LinkNet wireless network in Zambia over a period of 2 weeks in February 2010. The wireless network provides Internet access to approximately 300 residents of a rural village, as well as numerous international visitors. A number of wireless network anomalies were uncovered such as severe route flapping when the network was congested, unwanted packets with low TTL values flooding the network as well as interference between mesh routers and access points. These network anomalies form a strong argument for the use of white space networks in rural areas to reach distant dwellings. At the network gateway we analysed network traces to extract network performance issues that could have a significant affect on users. We found that the mix of long- and short-lived flows on the high-latency bandwidth-constrained satellite gateway caused significant performance degradation. Many daytime round-trip times of packets were between 3 and 10 seconds, sometimes reaching 60 seconds. This severely impacted web browsing interactivity and real-time applications. Malware traffic also had a significant impact on the network as many computers were infected due to lack of up-to-date virus signatures.

Usage analysis

There are a number of studies of Internet usage trends in developed regions. The Ipoque study [151] has looked at usage trends from 2007 to 2009 in different geographical areas across the globe and found that peer-to-peer traffic dominated Internet traffic. Very little is known about Internet usage trends in rural areas due to the small sample of networks that have carried out detailed traffic analysis. Studies in Ghana and Cambodia [51] revealed that portals such as Yahoo! and MSN dominated web traffic. Advertising, most which is not locally relevant, was placed second in terms of site requests. Web-based email also constituted a sizeable portion of traffic by bytes. These sites were Internet Cafes and resulted in usage patterns that followed the operating hours of the Internet Cafe.

In our work, we uncover recent trends in Internet usage in a rural village in Zambia. Our analysis, carried out in 2011, reflects a shift towards web 2.0 with social networking and content-creation becoming more prevalent worldwide. Our analysis showed dominance of web-based traffic similar to Ghana and Cambodia but key differences in usage patterns. A large portion of the traffic is not cached by the proxy server as much of the content is now served by dynamic web systems such as content distribution networks. Facebook dominated web site visits, with Facebook capturing three times more visits than the next most popular site, Google. This was a significant shift from the dominance of portals. Due to the prevalence of Facebook in our traces, we also investigated the amount of local user communication within an African village. This was an important consideration as the current hub and spoke architecture of the Internet forces traffic between local users in the village to traverse the bandwidth-constrained satellite network leading to large inefficiencies in the utilization of the gateway. Our analysis revealed that 54% of the instant messages in Facebook were between local users. This created strong supporting motivation for a novel architecture that reroutes traffic destined for local users via the local network in networks connecting to the Internet through bandwidth-constrained gateways

1.2.2 New solutions for developing regions

In this section we highlight a number of solutions to improve the performance of networks in developing regions based on findings from our analysis in Section 1.2.1. These solutions are both at the network and application layer and at the physical and data link layer. At the network and application layer we use novel routing and localization to deal with traffic patterns that caused poor network performance. At the physical and data link layer we employed new novel MAC and modulation techniques in white spaces to achieve better wireless coverage and performance to bring the Internet to remote rural populations.

Network and application layer

Wireless mesh networks (WMNs) provide an attractive method to provide Internet connectivity in developing regions. Traditional mesh routing protocols are designed to find high quality/throughput multi-hop routes in the network and greatly simplify network planning as wireless routers can often be dropped in place without any configuration. In our analysis of the wireless mesh network in Macha, Zambia as well as previous analysis of the mesh network in Peebles valley, we found that as the network began to grow, single bandwidth-constrained gateways caused congestion and high round trip times when browsing the Internet. This was especially true for users on the edge of the network. A solution to deal with the single-constrained gateway is to place extra gateways in the mesh to provide more capacity in the network and alleviate congestions at the single gateway. This scenario is possible in cases where the network begins to reach other "points of presence" at its edges, such as places where fibre, microwave backhaul or DSL is available or due to an operator of the network strategically adding new gateways in the mesh network. Existing routing protocols are not well suited to making an intelligent choice on which gateway to select. We designed a new Gateway-Aware Routing Metric (GARM) that selects high-throughput routes in the presence of heterogeneous gateways. The first part of the metric accounts for bottleneck capacity and the second models the total delay of the path. Our evaluation in simulations as well as on a testbed shows significant increase in network throughput compared to a common existing metric, Expected Transmission Time (ETT), to compute the best path to a gateway based only on path quality.

Apart from making an intelligent gateway selection, we also investigate techniques to reduce the amount of traffic passing through the gateway — usually the bottleneck in most rural networks. Our previous analysis in Section 1.2.1 revealed a high degree of local traffic in the network as well as a large amount of content generation and sharing even though many attempts to upload or share content failed. When rural users do not upload content to social media and content hosting sites, it becomes difficult for them to be active participants in the social networking revolution as well as contribute knowledge about local history, culture and customs. It is clear from the lack of Wikipedia articles in local African languages, such as Tsonga in Zambia, that there are severe obstacles to capturing local knowledge. We also found that content shared between local users in the village often traverses the bandwidth-constrained satellite link twice in order to share a file. Once, as the file is sent to a central content hosting server and again when the file is retrieved from the hosting server by the other local user. In order to address this issue we have developed VillageShare, an integrated time-delayed proxy server and locally-hosted content-sharing application. Through these two components, VillageShare facilitates localization of traffic, protecting the bandwidth-limited Internet link from content shared among local users, and minimizes upload failures by time-shifting large uploads to periods when the gateway link is under-utilized.

Physical and data link layer

VillageCell With the discovery of strong locality of interest for instant messaging in the village, we extend our argument to locality of interest for mobile devices in the village. Mobile phones are pervasive in Africa, showing impressive growth rates over the past 10 years, growing from 5% to 70% between 2000 and 2010. In our visits to Macha, Zambia and Dwesa, South Africa we found that 100% of residents owned mobile phones. However, many of the rural mobile subscribers are located in areas without mobile phone coverage and use their phones when they travel to towns with a cellular service.

Cellular communication is expensive, particularly in the developing world [12, 32] and there is a clear need to provide widespread, low-cost cellular coverage within rural communities. To meet this need, and in keeping with our localized communication model, we developed VillageCell, a picocell-inspired architecture designed specifically for developing regions. VillageCell is a novel cellphone infrastructure based on OpenBTS, an open source implementation of the GSM air interface that runs on low-cost software-defined radio boards. VillageCell can be installed at locations throughout a village and connected to existing WiFi or new white space networks for wireless backhaul. VillageCell is capable of translating GSM-based voice calls to VoIP and uses the SIP messaging protocol for the control plane. This makes it possible to establish calls between mobile phones and VoIP clients on any end-user device. In addition instant messages can be sent between mobile phones and various instant messaging clients supporting the SIP protocol. Due to our localized network architecture, calls originating and terminating within the local rural network are routed locally, avoiding the gateway link. Locally routed traffic does not use any commercial operator, thereby making local calls free. Our design also allows users to connect to the VillageCell system using their existing handsets without any modifications.

We performed lab-based experiments of VillageCell and deployed VillageCell in Macha, Zambia. In our lab-based experiments, calls were routed between multiple OpenBTSs over an indoor WiFi backhaul network. We loaded the WiFi backhaul network with background traffic from our Macha, Zambia network trace in order to simulate a more realistic scenario, where VoIP calls are sharing the capacity of the backhaul with data traffic. We found that the background traffic had very little effect on call quality and conclude that VillageCell performs adequately in a rural village using WiFi backhaul for up to 7 calls per cell for which it was designed. In the Macha deployment, we installed two VillageCell base stations connected with a WiFi backhaul. Our measurements revealed that 99% of the packets in the calls experienced delays less than 30 ms and jitter less than 3 ms — values well below thresholds that threaten performance.

VillageLink The sparse population density and large village diameters imply a need for novel technological solutions to bring Internet connectivity to individual homes. Current solutions for bringing Internet connectivity to a rural region include 802.11 modified for long distances [156] and WiMax [126]. Neither is applicable for covering an entire village. WiMax and 802.11 require clear line-of-sight for long-range connectivity. WiMax is not suitable due to the high equipment and deployment costs and its operation in licensed spectrum. WiFi often suffers from co-interference with other WiFi and non-WiFi wireless devices due to the small number of channels available and its lack of sensing capability. Cellular systems, on the other hand, have good penetration properties when deployed in sub-1000 MHz bands and their deployment in rural regions is rapidly growing. However, rural systems are designed for low-bandwidth applications and have per minute or per byte usage fees, and hence are not economically practical as the sole means of Internet access, particularly in regions where residents earn less than $1 USD per day.

As discussed in Section 1.2.1, the abundant amount of white space spectrum in rural areas is ideal for unlicensed coverage of low-density rural populations.

In addition, white space protocols standards are not finalized, allowing us to reconsider existing spectrum access schemes and fully utilize the PHY parameters flexibility to increase communication efficiency [173, 167]. IEEE 802.22, recently proposed for operation in white spaces spectrum, does not have the required flexibility due to the strict 6 MHz spectrum block requirement, lack of support for optimal choice of operating channel and lack of support for a distributed mode.

In response to these shortcomings, we developed VillageLink - a method for communication adaptation designed for the inherent challenges of a wide area setting in rural areas. We extended the current 802.22 specification to carry out channel probing in order to select optimal frequencies for a base station. After examination of the wireless propagation in a long-distance outdoor white space test bed in South Africa, we concluded that frequency selectivity is more dependent on antenna characteristics and linearity of RF components in the system than on the free-space propagation laws. With a non-uniform set of antennas being deployed at base stations, each with its own frequency selectivity and antenna radiation pattern characteristics dependent on the antenna design as well as antenna placement, it becomes essential to carry out channel probing to calculate the nature of the channel at different white space frequencies. We developed a channel probing system that is able to calculate the interference between base stations at different operating frequencies in order to choose optimal channels for base stations.

Once channels are allocated we carry out channel sounding to clients that wish to associate with base stations. Channel sounding is used to determine the delay spread between the base stations and clients which in turn can be used to calculate the frequency domain response of the channel. The channel properties between the base station and the client can then be used to carry out optimal OFDM sub-carrier allocation to clients.

1.3 Thesis statement

Our analysis of the unique challenges and opportunities in rural networks and the new solutions that adapt Internet access and wireless networks to these regions can be summarized by this thesis statement:

- Developing regions are often isolated from the digital age due to limited, costly bandwidth and unique environmental and cultural challenges. However, strong social cohesion in sparsely-populated rural areas we have studied warrants the use of novel bandwidth-saving voice and data localization techniques and long-range white space technology.

1.4 Contributions

Our key contributions towards adapting wireless access and Internet connectivity in rural developing regions are as follows:

- **Network Data collection in rural networks.** We have collected an extensive data set of rural network traffic in the research domain. The data set spans 2 years of usage from January 2011 with some outages due to server or power failures. Monitoring continues and further data sets continue to be retrieved from the village in Macha, Zambia. Monitoring traffic in rural regions is challenging due to their remote location, poor connectivity and difficulty in communicating technical concepts to support personnel in a remote village [160]. The traces captured all internal network (headers and data) traffic in the village as well as all traffic passing through the gateway to the Internet. This traffic was crucial to revealing new insights into network usage patterns and conditions unique to rural areas as well as forming the basis for the design of new network solutions that adapt rural Internet connectivity in order to improve reliability and performance.

- **Network usage patterns in rural networks.** Network usage patterns have been extensively studied in developed countries and trends are well understood. However, very little work has been carried out to study rural area network usage patterns from a live network. The usage patterns from our traces reveal communities that are active denizens of the Internet, using online social networks, studying online and using the Internet to improve farming practices. Traffic analysis reveals that the severely constrained Internet gateway experiences severe congestion in the day but is completely underutilized at night. This causes many upload and download failures which we quantify and provides strong motivation for time-shift proxies to make optimal use of the limited bandwidth of the gateway.

- **Social network analysis of rural networks.** Online social networks have typically been studied in developed world contexts [170, 171, 131]. Our work shows a number of key insights into the social connections between people in a rural village and to people outside the village. This work is the first to reveal a high degree of local connectivity between people in a rural village and allowed us to design novel localization techniques to provide bandwidth saving in the rural village of Macha, Zambia. On-site interviews were also used to try and understand the relationships between the quantitative network measurements and the qualitative information from the surveys.

- **Unique constraints of rural regions.** Many researchers working in the field of ICT for development have discovered that carrying out research in developing regions involves constraints that go far beyond adapting technology to rural regions [26, 160, 163, 162]. However, much of this work reflects short-term trends and perspectives often only derived from quantitative data. This work involves a number of local users in the village, some who have been in communities we studied for a decade, as members of the team to collect quantitative and qualitative data on issues faced when using new technological interventions. This has allowed us to report on constraints that range from environmental challenges to power quality issues to unique nuances of local cultures.

- **TV Spectrum availability in Africa.** We carried out an extensive set of spectrum scans of the television bands in urban and rural areas of the USA as well as cities and rural areas in South Africa and Zambia. The spectrum scans reveal an abundance of wireless spectrum unused by television in rural areas in Africa — 98% availability in rural Zambia — and provides a very strong case for regulation in Africa that opens up white space TV channels for licence-free opportunistic access. This is the first known set of cross-continent comparative data of white space spectrum availability using spectrum scans.

- **White spaces for rural regions.** White space networks provide a perfect fit for connecting low-density populations in Africa. However, current white space protocols such as 802.22 only provide guidance on avoiding interference with primary users and provide no solutions for selection of channels that maximize the channel quality to connected clients and minimize interference between white space cells. Our work addresses this by extending the 802.22 MAC protocol to include a channel sounding procedure to make an intelligent channel choice for each base station.

- **Cross-disciplinary and participatory action research** Our work involves strong partnerships with local champions in Macha, Zambia, embodying the concept of participatory action research [23]. We co-author our papers with local people from Macha village to ensure that our work is based on ground truth that correlates with anecdotal evidence from people who live there. New ontological understanding of many African perspectives on new media such as the Internet combined with our network traffic analysis and social network analysis reveal new insights which are applicable to designing appropriate technology in the developing regions of Africa. These also carry lessons in appropriation of technology in other developing regions around the world both in the processes to discover this knowledge as well as some of the discovered outcomes.

1.5 Dissertation outline

The remainder of this dissertation is organized as follows. In Chapter 2 we discuss challenges when deploying rural networks. In Chapter 3 we analyse the performance of a rural network in an African village. In Chapter 4 we carry out on-site analysis of Internet usage in rural Africa. Strong interest in social networks in Chapter 4 leads us to investigate the degree of locality of interest in Chapter 5. Discoveries from these chapters serve as a guideline for our network solutions in rural areas. These solutions deal with relieving congestion using gateway-aware routing in Chapter 6, providing a localized file sharing application to deal with locality of interest in Chapter 7, providing optimal channel selection for white space-based wireless connectivity in Chapter 8 and localized cellular access for rural areas in Chapter 9. We then conclude with scalability and generalizability of our solutions in the discussion section in Chapter 10 and then conclude the dissertation with Chapter 11.

Chapter 2

Rural Network Challenges

Being challenged in life is inevitable, being defeated is optional.

Roger Crawford

2.1 Introduction

Internet access is unevenly distributed across the globe and global statistics show that in 2011 developed countries had higher than 70% of individuals using the Internet whereas only 26% of individuals were using the Internet in the developing world [12]. Moreover, the variation between regions can be quite drastic. Sub-Saharan Africa, for example, lags significantly behind other developing regions; Internet penetration rates in countries such as the Democratic Republic of Congo, Liberia, Niger, and Ethiopia are less than 1% of the population [12]. Some African coastal countries such as Kenya have improved Internet connectivity over the past decade due to increased deployment of undersea cables. However, many land-locked African countries such as Zambia have very limited Internet capacity for the mostly rural population as the nearest sea-cable connection is over a thousand kilometres away.

Internet penetration and usage in the developing world lags behind the developed world for many reasons: lack of supporting infrastructure (roads and electricity), lack of supportive regulatory frameworks, and affordability, to name a few [26]. Supporting infrastructure is usually abundant in urban areas, even in some developing countries. For example, Internet penetration in South Africa is 4.6% in rural areas and 21.8% in urban areas [12]. In 2005 the developed world was predominantly urban, with three-quarters of its population living in cities. The much larger developing world, on the other hand, was predominantly rural. Urbanization rates are especially low in the most economically disadvantaged countries. For example, countries in the United Nations least developed group (by their HDI) are 70.5% rural [7].

There are also a number of less quantifiable causes of the lack of or low usage of the Internet which are not exposed [58] as work often lacks long-term contextual evidence [4]. Many projects are not extending over multiple years, or beyond the project implementation phase. As a result current literature focuses on short-term trends or broad general theory, often omitting unique cultural or environmental contexts such as African relationship economics [155] or "nondiscursive"[1] realities [31]. Challenges affecting rural areas include, amongst others, shortage of local research participation leading to inappropriate solutions, language barriers and lack of appropriate skills. Rural areas require skilled personnel and access to resources such as nearby electronics stores to maintain computer and network infrastructure. Rural areas often suffer from loss of skilled community members to urban areas as they pursue work or further study opportunities, while replacing equipment usually requires equipment to be transported from urban areas or other countries.

This chapter summarizes our experience investigating Internet usage in the rural villages of Macha, Zambia, and Dwesa, South Africa. Both of these villages have a stable wireless network that provides basic connectivity to the local population. We collected a full traffic trace for more than 2 months from a deployed network in Macha and performed an Internet usage survey in Macha and Dwesa. In a separate effort, we participated in deploying a wireless network in rural Peebles Valley, South Africa[89]. The majority of our analysis takes place in Macha, Zambia — the case study site chosen for most of the analysis in this dissertation

[1] Opposite to proceeding from topic to topic by argument or reasoning

as well as for trials of our solutions. However, similar findings at the other two sites on some of the issues discussed lend additional support for these solutions.

We combine the results of network monitoring, interview analysis, and anecdotal evidence obtained in these three locations to describe unique constraints when deploying Internet in rural Africa. Specifically, we find that constraints can be grouped into:

- **Environmental constraints** determined by environmental challenges unique to rural areas such as unreliable electricity or supply chain logistics. Rural inhabitants have little control over these challenges.

- **Skills constraints** caused by the lack of availability of trained ICT practitioners as well as the process of training and equipping these practitioners in ICT.

- **Cultural constraints** that deal with the complexity of using, installing, and maintaining ICTs in the context of African culture. Examples include perceptions of time and resources, roles and authority and the importance of oral tradition versus written contracts.

- **Economic constraints** focussed on the high costs involved in both procurement, installation, maintenance and monthly connectivity when rolling out ICT access.

- **Political constraints** as a result of poor national policies that could drive governments to remove barriers to bridging the digital divide.
- **Technology constraints** that are unique to rural areas such as bandwidth constraints and poor performance of certain operating in rural networks.

With 3 billion people living in rural areas of the developing world, where the connectivity is poor, careful examination of problems that prevent further connectivity expansion is critical to achieving universal access.

2.2 Background

Wireless networks based on WiFi technology have emerged as a viable solution for connecting previously disconnected communities in remote regions. Unlike alternatives such as fibre optics and cell phone towers, wireless networks can be built using cheap commodity hardware and do not incur an additional cost of licensing, and they allow collaborative and inclusive activities that facilitate self-management and appropriation by local communities. In recent years, isolated attempts to bridge the digital divide have been made by university research groups and non-governmental organizations [160, 127, 16, 89, 126]. A model that many of these projects follow is to bring wireless Internet connectivity (through satellite or other long-distance wireless links) to central points within a rural community, for example, to community centres, schools, or hospitals. This is commonly called the kiosk model, whereby citizens travel, often by foot, to these central areas to access the Internet. While clearly Internet access through this model is much better than no access at all, it is not a satisfactory end solution. In some cases, WiFi-based local networks are then spawned from the central points of connectivity to nearby regions to provide wider network coverage. The networks we analysed in Macha, Dwesa, and Peebles Valley are constructed in this way. Figure 2.1 shows the location of these sites that are the focus of this chapter.

2.2.1 Wireless Network in Macha, Zambia

Macha, Zambia is a typical resource-limited rural area in Africa, with scattered homesteads, very little infrastructure, and people living a subsistence lifestyle; the primary livelihood is maize farming. Like many sub-Saharan rural communities, Macha has a concentrated central area and a large, geographically dispersed rural community with a sparse population. Clusters of homes house members of a single family are generally separated from other clusters by one or more kilometers. Macha contains about 135,000 residents spread out over a 35-kilometer radius around the village center. The overall population density is 25 per square kilometer. In the middle of the community center are the facilities and housing for a mission hospital, a medical research institute, and schools.

Figure 2.1: Map of Southern Africa showing study sites: Macha in Zambia and Dwesa and Peebles valley in South Africa.

The Macha Works organization, through the LinkNet project, has deployed a network of long-distance WiFi wireless links and mesh networks that provides connectivity to about 300 community workers and visitors using satellite-based Internet and in more recent years a 2 Mbps terrestrial wireless link. Figure 2.2 shows a scaled view of the core of the LinkNet network with the position of each wireless router represented by a symbol.

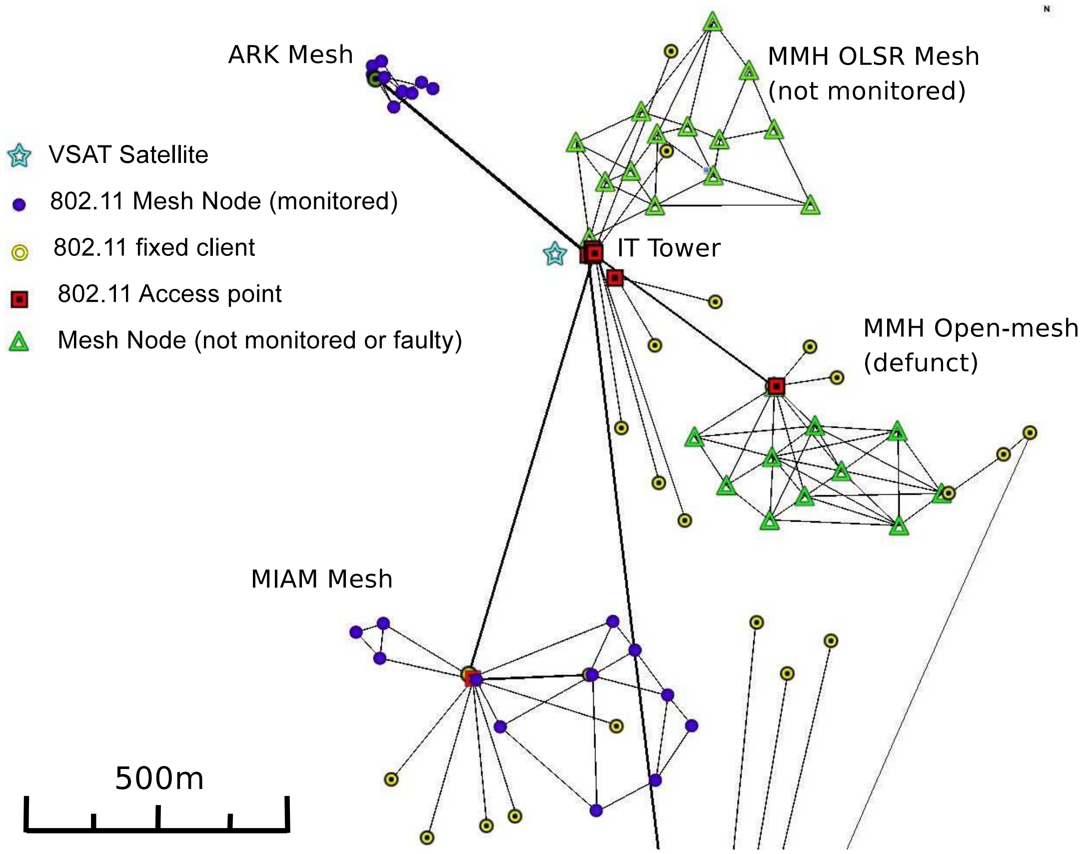

Figure 2.2: Macha wireless network in Zambia.

The original satellite link provided a committed download speed of 128 kbps (bursting up to 1 Mbps) and a committed upload speed of 64 kbps (bursting up to 256 kbps). The total monthly cost of the satellite connection was $1,200 USD per month and was covered through government subsidies as well as through Internet vouchers sold to visitors and locals. Our analysis of Macha Internet usage took place when satellite-based Internet was still in place. Most users access the

Internet at work and through community terminals, although a few people do have WiFi connectivity in their homes.

2.2.2 Wireless Network in Dwesa, South Africa

The Dwesa region is located in Eastern Cape Province, one of the poorest regions of South Africa. Similar to Macha, Dwesa is characterized by severe resource limitations, stressed infrastructure, a weak subsistence economy, and a sparse population (15,000 residents in an area of 150 square kilometers). The Dwesa community is affected by migration of young people and high crime rates. The telephone service that once existed in the area, for example, fell into disrepair after the copper telephone cables were stolen.

The Siyakhula project, led by the University of Fort Hare and Rhodes University in South Africa, has established Internet connectivity among local schools via WiMax links shown in Figure 2.3 that are several kilometers in range (Mandioma, 2007). The license for WiMax operation was provided through a local telecom, that is also a sponsor of the project. One of the schools connects to a VSAT satellite, thus serving as the Internet gateway. The satellite delivers 512 kbps download speed and 128 kbps upload speed. The connectivity is mainly used for student education purposes, school record keeping, and inter-school communication. Besides being available to students, the connectivity is offered to members

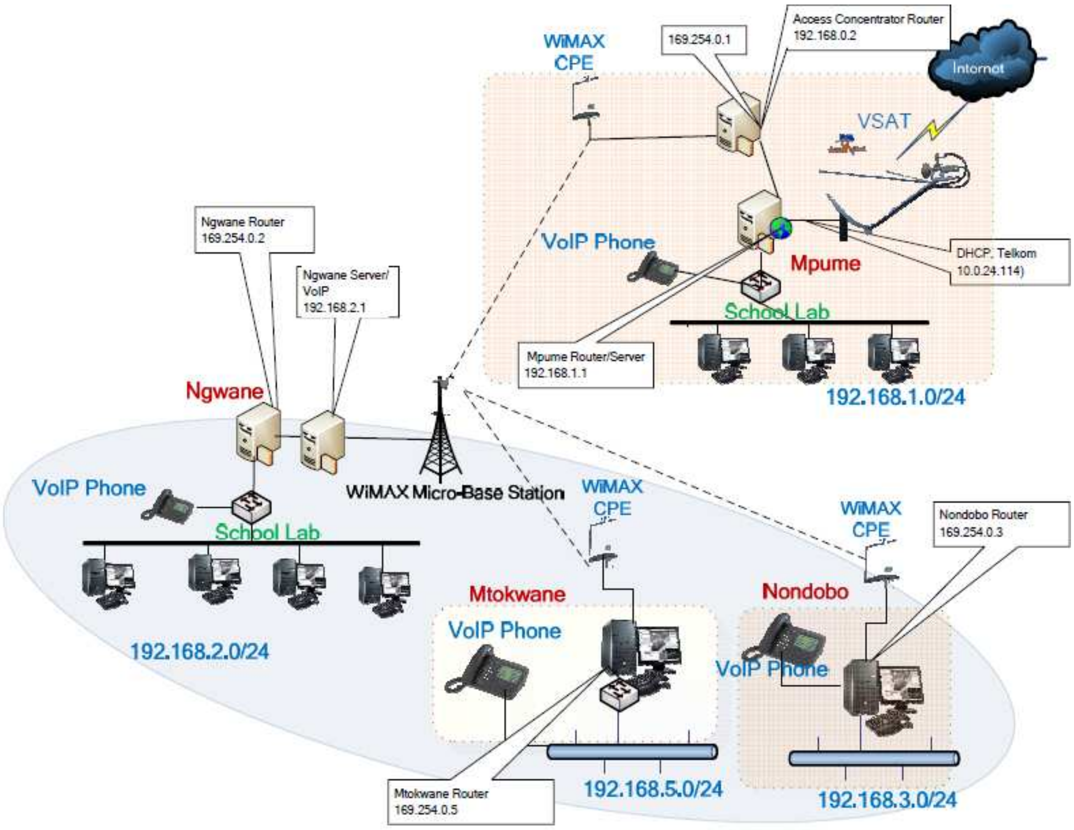

Figure 2.3: Dwesa wireless network in South Africa (Image courtesy of Martin Mandioma [112]).

of the community after school hours, when computer literacy training courses are offered to local residents.

2.2.3 Wireless Network in Peebles Valley, South Africa

Peebles Valley and the Masoyi tribal land are located in a hilly area in the north-eastern part of South Africa. The Masoyi community is under-serviced, with most roads remaining unpaved and most houses lacking running water. The community is poor and has been hugely impacted by HIV/AIDS.

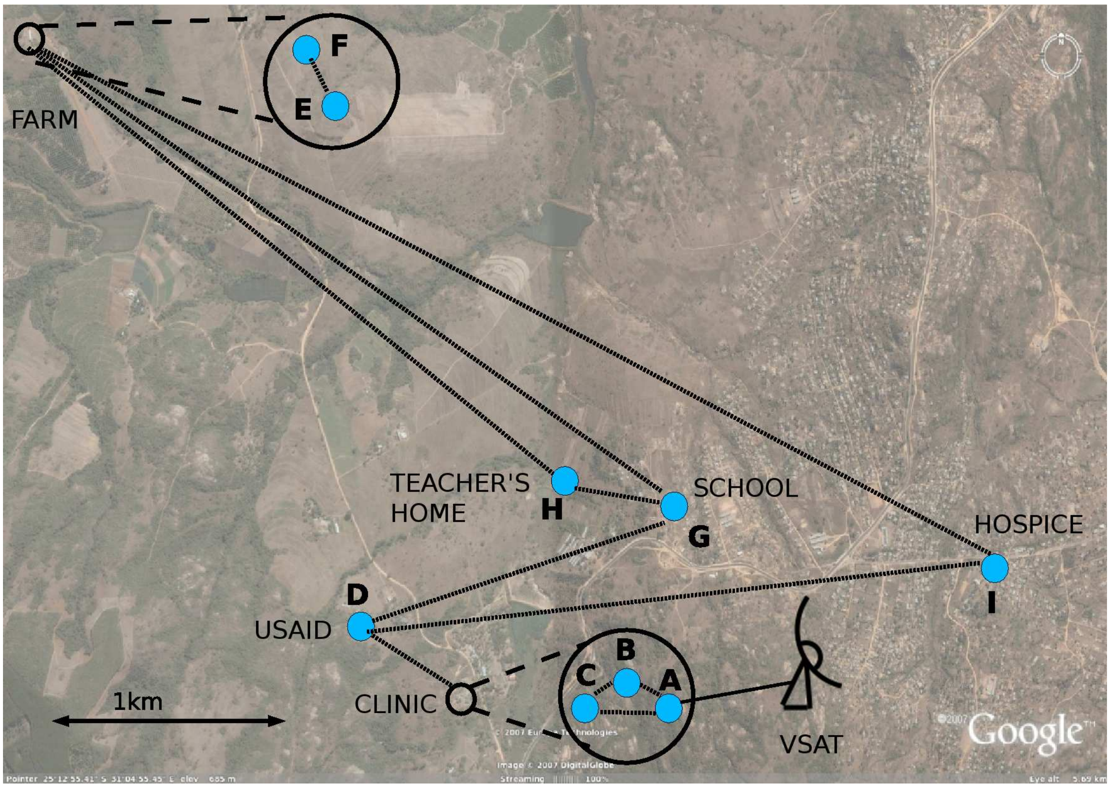

Figure 2.4: Peebles valley mesh network in South Africa.

The Peebles Valley mesh network shown in Figure 2.4, consisting of nine nodes, was deployed in 2007 over an area of about 15 square kilometers. The key user of the network was a local HIV/AIDS clinic. The clinic connected to surrounding schools, homes, farms, and other clinic infrastructure through a WiFi mesh network. A VSAT satellite Internet connectivity, which was provided free of charge by the HIV/AIDS clinic sponsor, provisioned a total of 2 GB per month at a download speed of 256 kbps and an upload speed of 64 kbps. Once the 2-GB capacity limit was reached, the Internet connection would be cut off until the beginning of the following month. This satellite link was usually underutilized every month, with clinic staff using about 60% of the available bandwidth, and no spare capacity could be carried over to the next month. This spare capacity was shared with users in the mesh network free of charge, but it had to be carefully managed by a firewall to ensure that their usage did not affect the clinic's Internet availability.

2.3 Methodology

We investigate ICT adoption in rural Zambia and South Africa through mixed-method research in order to address multi-disciplinary research questions. This is accomplished by combining quantitative network traffic monitoring in Zambia and a set of qualitative interview and survey data from Zambia and South Africa in order to extract usage trends and obstacles to ICT inclusion.

2.3.1 Network traffic monitoring

In Macha, Zambia we installed a network traffic monitoring system on the village network gateway to capture packet headers of all network packets that traverse the satellite link. The network gateway also features a web proxy server used for traffic caching. The proxy access logs allow analysis of HTTP traffic. All IP addresses were anonymised in order to protect the privacy of the users. Initially we captured 14 days of traffic from midnight, Sunday 31 January to midnight Sunday 14 February in 2010. In a follow up, we collected two months of network traffic in February, March and April 2011. Approximately 450 GB of packets were captured and analysed to assess the use and locality of traffic. Analysis of the output of gateway management software, installed in April 2010, reveals that local residents generate about 90% of the network traffic. Visitors to the community generate the other 10% of network traffic.

In Peebles Valley South Africa, we installed a bandwidth monitoring system at the gateway. This allowed us to measure the amount of traffic being used by each user in the network. The system was also used to allocate a specific amount of bandwidth to each user and block usage once the capacity limit was reached.

2.3.2 Interviews and surveys in Macha and Dwesa

On-site interviews were conducted in July and August 2010 in Macha, Zambia and Dwesa, South Africa by Pejovic [92]. A total of 37 interviews were conducted: 23 in Macha and 14 in Dwesa. The participants' ages range from 18 to 57; 15 of them are female, 22 are male, all are literate and have at least some high school education, and income ranges from zero to more than US$300 per month.

These interviews were in the form of a directive, structured questionnaire in the first phase of the conversation, in order to obtain highly quantifiable data that could be correlated with the results of our network trace analysis from Macha. In the second part of the interview, the subjects were asked less formal questions and were able to engage in a discussion with the interviewer. In African culture, narrative communication is common; thus, these open questions revealed a number of unforeseen issues.

User surveys.

In June/July 2011, we followed up these interviews with an online survey of Internet users in Macha. The aim was to investigate use of Web 2.0 applications and services [94]. The questionnaire consisted of 89 questions, implemented on the SurveyMonkey tool. Invitations were sent via email, and Facebook links, inviting people for participation on a voluntary basis. All the questions were completed by 41 out of 66 participants.

2.3.3 Long-term reflection

As part of his previous research, the author participated in sporadically administered unstructured interviews among network users in Peebles Valley. While lacking any quantifiable data, this anecdotal information provides qualitative insight into Internet usage in Peebles Valley, as the data were acquired over a three-year period.

Another member of our team provided deep qualitative insights in the Macha community by reflecting findings with qualitative, observational, longitudinal, mixed method research in the community. The research was conducted over nine years, during which that member resided in Macha and was highly respected by members of the community.

2.4 Environmental and infrastructural constraints

The constraints described in this section highlight the effect of the rural environment and lack of infrastructure on ICT roll-out and maintenance. Local users have little control over rural environmental constraints. Infrastructural constraints, on the other hand, may improve over time with sufficient economic investment. There has been a fair amount of literature on this topic. Work in Cambodia and India described challenges due to transportation issues, supply logistics, natural disasters and heat and dust [26]. Work AirJaldi in Northern India and the Aravind Eye Hospital in southern India reported on failures from poor power quality and lightning [160]. Our experiences have echoed many of these challenges and we now describe new insights, gleaned from the rural sites we have studied.

2.4.1 Geographical constraints

In Zambia, distances between towns are far. Zambia has a population density of 17 per square kilometre, a stark contrast to a developing country such as India with a population density of 371 per square kilometre. This drastically increases the time taken to transport equipment for installation or repair and presents extreme cases of the transportation challenges reported in areas such as India or Cambodia [26]. One way trips by road are expensive and typically span a number of days with frequent breakdowns, delays and accidents. There are numerous bumps on dust roads that can easily damage sensitive ICT equipment. Smaller amounts of equipment can be brought in by light aircraft using the local airstrip in Macha but costs of flying equipment into Macha are high for the small amount of cargo that can be carried.

Much of the equipment required for the VillageNet project, such as lead acid batteries, GSM antennas and computers were not available in Zambia or neighbouring countries and were driven into Zambia from South Africa by the author. This involved a 3-day journey via Botswana with a 4-wheel-drive vehicle and long delays at multiple border posts on route as equipment required temporary import permits.

2.4.2 Natural dangers

Dust is pervasive in rural Zambia, with many homes and buildings being surrounded by a bare earth yard that is swept every day. Most of Zambia has high summer temperatures above 30 degrees Celsius and buildings are open to allow movement of air as air-conditioning is costly in rural areas. IT staff report cleaning fans every three months due to dust. Within 4 weeks of living in Macha, one of the author's laptops was damaged due to heat and later diagnosis revealed dust trapped in the tiny vanes of the laptop fan. Clearly, fan-less computer hardware that can tolerate the level of heat in areas such as Zambia are the optimal solution. However, cost is a crucial factor as most of the equipment used in projects in Macha is donated older-generation equipment with fans.

Lightning damage is common in the summer months both in Macha and Peebles Valley. In Peebles Valley, our wireless mesh routers were powered using power-over-Ethernet and 10 to 15 meters of CAT5 cable were run from the router to the power-over-Ethernet injector. Two out of our ten wireless routers were damaged due to lightning. The Ethernet port was always damaged as EMF from a near strike induced high voltages in the CAT5 cable. In Macha our VillageCell device mounted on a water tower met a similar fate in the summer. A section of CAT5 cable connected VillageCell to one of the wireless routers and again lightning damaged the Ethernet port. Replacing runs of CAT5 cable with wireless adapters on the outdoor device for local access could prevent this problem in the future.

2.4.3 Electricity supply

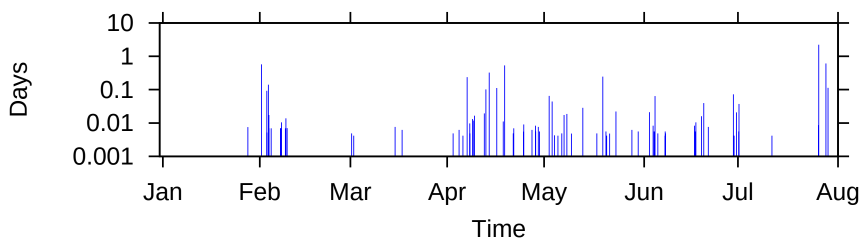

Figure 2.5: Timeline of power failures in Macha.

In Zambia, electricity is a vital constraint in rural areas. National electricity supply is either unavailable or unreliable, with power failures and voltage surges or brown-outs being common. We carried out long-term coarse-grain power availability analysis as well as detailed power quality monitoring for a month. Figure 2.5 shows our long-term power analysis; a time line of power failures in Macha during 2011. The y-axis shows the duration of the power failure. Figure 2.6 shows that, on average, the period between power failures is about 1 day, but this is masked by the fact that when power is restored there are often repeated failures. The cumulative distribution function (CDF) of duration of power failures in shown in Figure 2.7. The average duration of a power failure is approximately 1 hour but 10% of power failures can last longer than 12 hours. The longest power failure in the data set was 2 days.

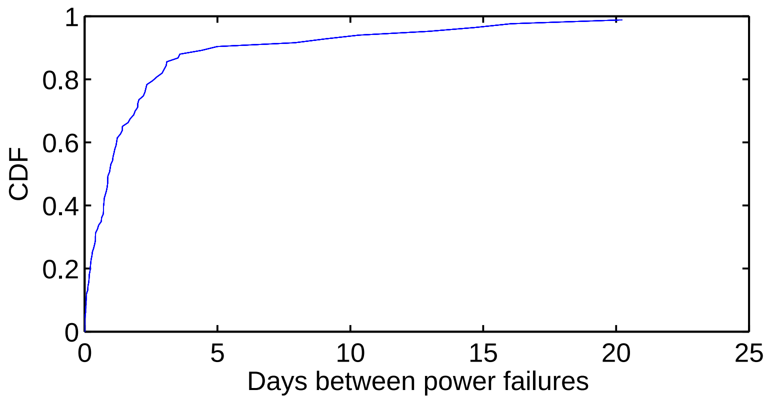

Figure 2.6: CDF of time elapsed between power failures in Macha.

Figure 2.8 shows our fine-grained power analysis where we monitored the grid power and network availability in Macha over four weeks (only 2 of these 4 weeks are shown). We find that Internet connectivity is not available to equipment

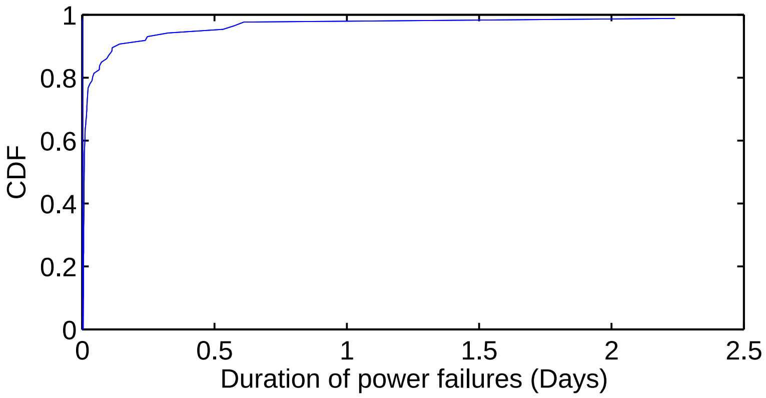

Figure 2.7: CDF of duration of power failures in Macha.

approximately 40% of the time and power is not available approximately 15% of the time. The lack of Internet connectivity is often due to upstream problems in the grid or connectivity infrastructure in Zambia or due to local equipment needing to be reset after becoming unstable when there are power failures or brown-outs (voltage drops by less than 20% of the standard voltage, 220V). The measured voltage on a raw power outlet in Macha is shown in Figure 2.8a. Three brown-outs are clearly seen in this diagram and have the potential to put equipment in an unstable state. Internet availability is shown in Figure 2.8b. It is clear that there is often no connectivity to the outside world even when power is available.

It is not only the availability of the electricity source, but also the issue of sharing of electricity that plays a significant role. When introducing electricity for ICT equipment in an area without electricity, questions are asked about the priorities for allocation and use. Certain stakeholders may request to be connected

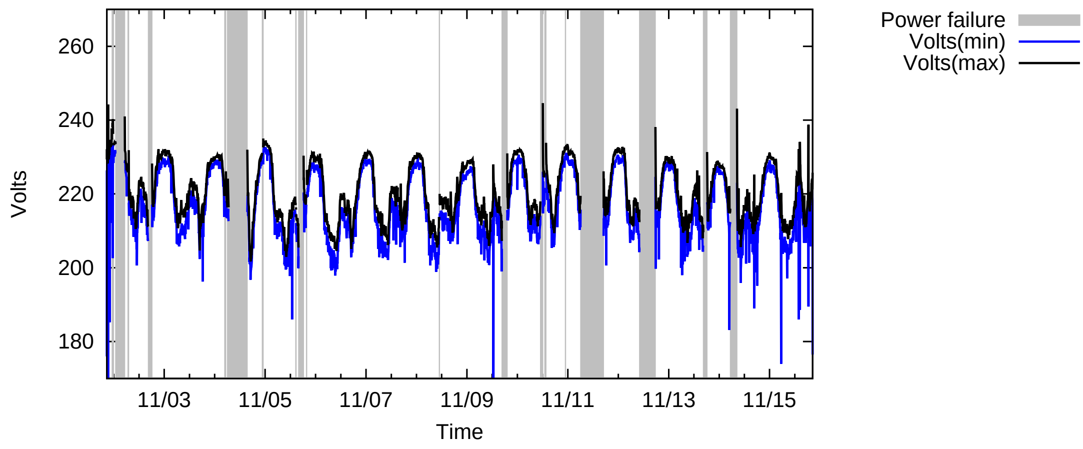

(a)

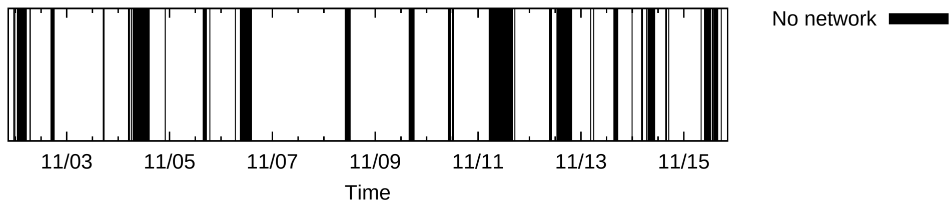

(b)

Figure 2.8: Power quality and availability in Macha over a 2 week period. (a) Gaps in the voltage reading depict the power failures. Dips below 200V represent brown-outs, (b) The shaded regions show times when no Internet connectivity was available.

first even through the electricity was primarily for the ICT equipment. In one specific location in Zambia, a local chief connected his home to an electricity source intended for ICT equipment. At another site, the local hospital management took over a solar system supplying a satellite-based Internet system and linked it with hospital resources. Careful forward planning that takes into account surrounding community members is required before extending grid-based power to electrify an ICT installation.

In practice, there are many different levels of quality of equipment arriving in rural areas. Assorted batches of donated equipment arrive with all kinds of levels of ruggedness and susceptibility to dirty power. Standard Uninterrupted Power Supplies (UPS), meant to protect equipment from energy disruptions, get easily damaged as they are the first line of defence. In Macha, out of over 40 high quality UPS systems, only 5 remained operational after one year with most failing within the first month of use.

Clearly, designing equipment that is rugged enough to withstand 'dirty' power often found in rural areas is a key issue that warrants further work. The alternative is to disconnect from the local power-grid and make use of renewable energy sources such as solar power, particularly well suited to the high percentage of sunshine days in sub-Saharan Africa. However, solar technology is plagued with issues such as batteries faltering due to high temperatures or abuse. In addition, deep-cycle batteries, required for solar installations, are costly to purchase and transport. Solar installations also typically involve equipment that requires training of specialist engineers and substantial initial investment and planning

(Bernardi et al., 2008). Further, solar panels are vulnerable to physical damage or theft.

2.4.4 Locality of usage

Somewhat related to geographical constraints is the fact that most Internet users in rural Africa are restricted to public terminals and access at school or work. Only about 10% of Internet users in the African continent access the Internet from their homes (ITU, 2011), and one can surmise that this figure is even lower for rural users. Public access comes with limited availability, greater cost, and long walking distances to points of access. We also find that at-home access allows users to enjoy the Internet in a leisurely way, using advanced applications such as online social networks and blogs. Many technical obstacles stand in the way of more prevalent at-home access in rural Africa (Surana et al., 2008). We have described the challenges in reliable grid power, but even if grid power is available it is usually restricted to schools, hospitals, and community centres. Long distances between households are another important obstacle for rural area connectivity. Typical low-cost wireless technologies such as WiFi are restricted to line-of-site links with high-power antennas to reach distant sites. Cell phone base stations can provide some Internet connectivity to wide areas but these technologies are normally not available in rural areas in Africa, use GPRS or EDGE links with very limited capacity or are too expensive to be used by the average rural Internet user. We discuss the use of a white space solution to increase Internet penetration to people's homes in Chapter 8.

2.5 Skill constraints

The knowledge base for rural ICT engineering solutions is mostly outside of Africa where research and development takes place outside the context of rural Africa. There are many challenges carrying out research in resource-limited environments [26, 76] or with African universities. We now discuss the challenges involved in obtaining sufficient skills to carry out research, development and installation or maintenance of ICTs within a rural African context.

2.5.1 Situation in Zambia

A total of only 300 people possessed graduate qualification in ICTs in Zambia in 2008 [64]. Uneven distribution of this workforce compounds the problem as the majority of the population live in rural areas and most personnel with ICT qualifications live in urban areas. This is a serious challenge as health and education facilities in rural areas begin to receive computer facilities while support staff that deal with ICT issues typically reside in urban areas [95, 145]. Rural institutions often assign ICT tasks on a task-shift basis to unqualified staff with little or no training resulting in low quality engineering efforts. As a result, faulty equipment ends up standing idle for long periods of time or is discarded.

In contrast with efforts in health and education, there are no national programmes, nor public-private partnerships, specifically designed for training of ICT practitioners in rural areas. There are no vocational training centres located in rural areas, and training is mostly left to the unregulated, commercial markets in major towns.

2.5.2 Situation in South Africa

The South African ICT sector, in contrast, is more mature with many homegrown industries ranging from cell phone companies, with a significant footprint in Africa, to mobile phone application development companies with an International footprint such as MXIT, a social network service for mobile phones.

South Africa has a smaller rural population than Zambia (38%) but a similar critical ICT skill shortage exists in rural areas. At a national scale, there are also ICT skills shortages. Between 2005 and 2010, South Africa conferred an average of 1 700 university degrees in computer science and electrical and electronic engineering, and 3,130 Technikon diplomas a year [69]. However, the 2011

ITWeb-JCSE Skills Survey suggests that current demand amounts to 20,000 to 30,000 job opportunities, or 10% to 15% of the total ICT workforce.

In both South African networks we studied, local rural ICT skills were non-existent. In the Peebles Valley network, the network was designed and built by the author who did not live in the community. In-situ training on wireless network installations was given to a local young high school graduate during each visit. However, at the end of the project, he obtained a job at an IT store in a nearby town. Although this was a positive outcome for the local community member, it had no benefit for the community as a whole. This contrasts sharply with the model used in Macha where 'local talent' or 'visionary leaders' are fostered. In the Dwesa project, ad-hoc ICT training was given on Office, Internet and network installation to some teachers and local youth. The concept was for these trainees to train other community members. However, with no clear training structure or certification, the effectiveness of this approach was limited. Other initiatives in South Africa have had more success in building and retaining skilled ICT practitioners. One such initiative is the CSIR BB4ll project [150] which trains Village Operators to become self-sustained rural ICT enterprises. Another is the Infopreneur model [114] that creates a network of local champions supported by a micro-franchise for IT support and financing.

2.5.3 Support issues

Mesh technology employed in Peebles Valley and other rural projects in South Africa such as the BB4all project is designed to ensure that very little network configuration is needed to complete an installation. Still, some basic troubleshooting knowledge is required when a link fails to determine whether the problem is, for instance, an antenna problem, a cable problem, or a problem with the wireless router. This level of expertise can be obtained through a multi-week training program with in-situ style teaching. Our interviews reveal that Macha and Dwesa differ significantly in terms of what people do when the network misbehaves. People in Dwesa often do not take any action. Users in Macha, on the other hand, are willing to organize their time around Internet availability; for example, 57% try to go online at a less congested time of the day, and 24% of users in Macha attempt to fix network problems themselves. We suspect that the incentive to take action comes from a longer history of Internet usage in Macha and the focus of Macha Works on providing local, vocational training.

2.5.4 Local champions

Wireless networks in the rural developing world are often deployed by foreigners. Once those who built them leave, or when the foreign funding is exhausted, the networks typically falter. Hence fostering local champions is key to the future sustainability of a network. Using Interview data from Macha and Dwesa, Pejovic reports that local champions can be significant in motivating other people from the community to adopt technology [92]. For example, in Macha, men are more likely than women to perform computer maintenance and virus scans themselves. In Dwesa, this discrepancy doesn't exist. Pejovic posits that this is due to Dwesa having a number of project leaders and local champions that are women, helping women to not see such tasks as "men only."

The rural Zambian co-operative social enterprise, Macha Works strives to empower 'local talent' to lead in their own communities using a wide, horizontal approach that aims to keep, and even attract, talented people to stay, or return to rural areas. Fostering a champion or visionary leader role is echoed by authors such as Unwin [163] and Toyama [162]. Macha Works carries out on-site training within the rural community of Macha in, amongst others, basic computer literacy, computer technician courses and entrepreneurship. Staff also team up with national institutes such as the University of Zambia, and international organizations in a collaborative approach [117]. Scaling up this effort, however, requires investment in similar training programmes at a national government level.

2.6 Cultural considerations

In Southern Africa, Ubuntu culture forms a integral part of the life of community members within rural areas in particular. There are different interpretations of what Ubuntu is. Liberian peace activist Leymah Gbowee defined Ubuntu as "I am what I am because of who we all are." Nobel peace prize laureate, Archbishop Desmond Tutu, offered this definition:

> A person with Ubuntu is open and available to others, affirming of others, does not feel threatened that others are able and good, based from a proper self-assurance that comes from knowing that he or she belongs in a greater whole and is diminished when others are humiliated or diminished, when others are tortured or oppressed.

African cultures embody a strong desire for progress as a community and peaceful coexistence. Resources are often communal with less emphasis on individual ownership. Rural communities build relationships carefully and slowly to allow for testing of commitments, character, and appreciation of relational aspects of new community members.

Indigenous cultural heritage expresses itself primarily through oral tradition [130] or audiovisual means such as dance or music [165]. In general, rural African culture is a society without written records, plans, and filled-up diaries. Interactions are focussed on the immediate with little planning beyond a week. Researchers from western cultures working in rural communities take time, often multiple years, to adjust to indigenous systems, beliefs and traditions that hold rural communities together [67, 115, 135].

2.6.1 Challenges introducing ICTs

The customary authority for the local, rural community is the chief. The chief must be consulted for any intervention in a rural community. The chief often refers issues or new interventions to a meeting of his (senior) headmen. Mission stations are common in many tribal communities in Africa and have their own decision-making processes that involve various layers of approval. Meetings where decision are made may be annual or bi-annual. These decisions require approval processes involving government authorities like district commissioners, regulatory authorities, ministers and other government officials. These multiple layers of authority can severely delay ICT interventions and frustrate partner institutions with tight budgetary year-based cycles.

Most technologies emerging from ICT research and development and their prime users are from within industrialized economies. Resulting ICT solutions do not necessarily align with non-technocratic cultures or non-industrialized circumstances. Solutions and processes for incorporating these solutions need to embody or reflect local traditions and practices or risk being culturally unsustainable [147]. Some examples of mistakes that could be made are: providing a new

Internet service with an online manual instead of face-to-face training, provision of Internet access in a pilot program to an individual rather than the surrounding community, lack of consultation with the chief of a village before implementing a project, trying to prove that a technology intervention will improve peoples lives without sensitivity to cultural heritage or designing a project with an expectation of results within one year.

Most rural areas have seen no major interventions beyond primary health care or some basic mobile telephony services. Thus the community can be 'closed' and attempts to make new interventions fit the values and beliefs it defends. Most people in rural areas work in different roles and capacities at the same time. Subsistence farming, practised by most rural people, with its seasonal activities and responsibilities assures food for the family. The community expects responsible individuals to participate at community meetings, weddings and funerals, and activities organised by local and traditional leadership. As part of its institutional responsibility, activities must align with these seasonal and incidental events.

Introduction of ICT access in rural communities constitutes a considerable change for a community. These changes make for periods of uncertainty and power struggles between the authorities involved can develop. The process of change is also effected by spiritual considerations influenced by an array of Western, syncretistic or traditional religious beliefs [41]. The rural communities judge activities 'sustainable' when they are welcomed by all, are comprehended and can be vocalised by all members of the community, and when all persons are included and have a sense of partaking in the development.

2.6.2 Locally relevant content

The Internet, in its current state, can be perceived as a threat among indigenous communities or simply lack relevance with its almost exclusively foreign content [17]. Conversely digital content generation and sharing can be an important avenue for cultural exchange and preservation and can even spark development activity from within the society [158]. We investigate current content generation behaviour, issues that discourage content generation and solutions that encourage users to generate and share content in Chapter 7.

2.7 Economic constraints

High cost is one of the major barriers to further Internet penetration. For example, satellite access in landlocked countries such as Zambia costs $1,400 USD per month, while the average monthly income is about $30 USD. However, quantifying the economics of countries such as Zambia is challenging as the majority of the economy is in the informal, unregulated sector [111]. Upgrading to terrestrial connectivity in rural Zambia in order to supply sufficient capacity to a large user base costs approximately $3-4,000 USD per Mbps (committed symmetric rate) with a minimum contract commitment of 2 Mbps. This is approximately 100 times the cost of a typical broadband link in urban South Africa.

Recent statistics show that connectivity remains out of reach for many in the developing world: monthly access fees in Africa require 291.3% of the average monthly income, compared to only 1.4% of the average monthly income in Europe (ITU, 2011). Even mobile phones which have been hailed as a success story in bringing some degree of connectivity to Africa may at times be implicated in the production of poverty. For example, in Ethiopia, the poorest 75% of the population who use mobile phones spend 27% of their income on them and reflect the high cost of mobile phone services in Africa [162].

Efforts to bridge the digital divide usually require significant investments from the developed world, and a sustainable solution seems to be out of reach with the current pricing schemes. Apart from high monthly connection fees, costs involved in the procurement of equipment and installation create barriers to entry. Finding suppliers of equipment in rural Africa is challenging as there is little desire by many suppliers to provide a service in rural areas due to remoteness, high transportation costs, low volumes, and communication challenges. Pricing is often an order of magnitude higher than those encountered in western countries.

Procuring equipment also comes with a multitude of challenges such as lack of acceptance of local credit cards and import constraints and delays.

Most common sources of income in rural areas (farming and fishing) only generate seasonal income. While mobile phone providers have addressed the issue of cost through their prepaid billing schemes, which allow underprivileged communities without bank accounts to tailor cell phone usage to their own ways of life (Horst, Miller, 2006), satellite Internet providers demand regular monthly payments [55]. In the LinkNet network in Zambia, an Internet voucher system was used for users in order to attempt to recoup the monthly Internet costs but this system had significant challenges of its own.

2.7.1 Voucher system

LinkNet recoups costs of connectivity using a voucher system. This allows users to buy fixed amounts of bandwidth per month. LinkNet provides a user code and pin number printed on a standard sheet of A4 paper shown in Figure 2.9. This system provides no security against unauthorized use of somebody else's code and pin combination but a possible solution using a specialized secure scratch card printing machine costs approximately $5000 USD and requires advanced maintenance skills. The LinkNet voucher system is part of a proprietary user access control gateway with monthly support fees designed for typical urban temporary Internet access scenarios such as airports or coffee shops. An improved access control system adapted for typical Internet capacity-constrained environment such as Macha would, for example, allow users to configure profiles that block automatic updates, automatic file-synchronization and save users from unwanted depletion of their Internet voucher. At the time of study, users were able to purchase capacity-based vouchers of $30 for 1 GB with a 30-day expiry period and time-based vouchers of $7 for 3 hours although time-based vouchers were almost never used.

In Macha, a 256 kbps VSAT-based connection was used before April 2011. This cost $1,400 USD per month and LinkNet was able to recoup connectivity costs through a user base that reached close to 300 users per month at times. However, Internet responsiveness became very poor during peak usage times to the point where a single web page request could take more than a minute to load. During 2011, the Internet link was upgraded to a committed 2 Mbps line at a cost of approximately $8,000 USD per month. This improved the throughput and allowed expansion of the wireless connectivity to a wider geographical area to absorb more users. However, at least 260 regular monthly users purchasing 1 GB per month were required to cover the costs of the terrestrial Internet link.

Figure 2.10 shows the number of new vouchers purchased per month as well as the number of unique machines and unique voucher logins per month – unique

MACHA WORKS
inspiring local talents

LINKNET
connecting my neighbour!

31 Days or 1 GB
Multi-Use Internet Voucher

ZMK 150,000

Card Number: **M7PEM**
PIN: **X4A79**

How to use:
1. Select "LinkNet-HotSpot' from your list of available Networks
2. Open an Internet Browser and go to a website
3. A login screen will appear automatically
4. Enter your unique Username and Password on the login screen
5. Click Login and enjoy the Internet!

This Internet Voucher expires automatically 31 Days after first login, or upon use of 1GB of data. This Internet Voucher allows access on any LinkNet network[1]

Date of First Login: ________________

LinkNet Support: web: www.link.net.zm/support, tel: 0977184615, mail: support@link.net.zm

Voucher: M7PEMX4A79
Serial Number: 2701
Expiry date: 1 Jul 2011

1 LinkNet Networks currently operational at Chitokoloki, Chikanta, Chilonga, Kalene, Macha, Minga, and Mukinge

Figure 2.9: Macha Works Internet voucher used to gain access to LinkNet.

voucher logins can exceed new vouchers purchased as 30-day voucher periods often span two months. When the number of vouchers used per month exceeds the number of unique machines, it shows that users are sharing machines. Although there was some increased uptake after the Internet connection was upgraded to 2 Mbps,

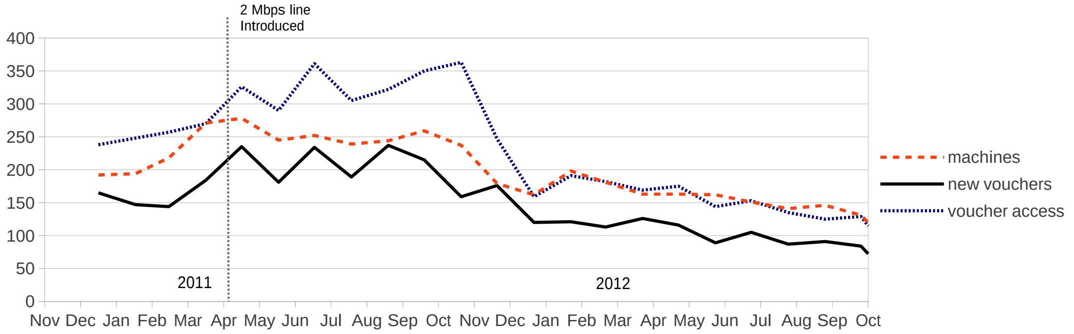

Figure 2.10: Voucher use in the LinkNet network showing number of new vouchers purchased per month as well as the number of unique machines that connected to the network per month.

a break-even point, where number of new vouchers was greater than 260, was never achieved. The highest number of new vouchers purchased within a month was 237. Most of these vouchers were purchased by people who were formally employed by Macha mission hospital, the malaria research institute or Macha Works, a sub-group forming a small percentage of the general Macha population that are mostly in the informal economic sector. With many formally employed paying users leaving in 2012, large losses were incurred by the organisation. High costs, inter-organisational tensions as well as alternative access possibilities from a mobile operator in the region led to a loss of user-base.

E-commerce challenges

The value of the Internet depends on the supporting national structures. For example, online payments are possible only if a well-developed banking system exists within a country. In Macha, 28% of the people we interviewed tried some form of online trading, often with unsatisfactory results. In Macha, a few users who did manage to gain access to a credit card were usually not able to make use of the credit card online as many e-commerce services did not accept Zambian credit cards as a means of payment. In order to fully support e-commerce, courier services are required; however, due to challenging roads and lack of named streets, courier services are only available in towns that are often hundreds of kilometres away.

2.8 Political barriers

The value of the Internet depends on supporting national structures. For example, online payments are possible only if a well-developed banking system exists within a country. Providing e-government services is only worthwhile if a nationwide Internet adoption plan is put into place to eventually provide all citizens with on-line access.

In some countries in Africa, the WiFi frequency range is subject to licensing. In such cases, regulations on wireless access are extremely important for Internet adoption, and relaxing licensing requirements for WiFi networks directly leads to higher overall Internet penetration[19] and innovation [107].

Our efforts in Peebles Valley, South Africa are just one of many examples where the network ceased to exist in part due to licensing issues. A local champion was sought to continue growing the network beyond the pilot phase. Although a local computer store owner in a nearby town was interested, this never materialized because of the uncertain regulatory conditions in South Africa at the time. During 2007 and 2008, very high license fees had to be paid by a service provider to deploy and maintain any access infrastructure, creating unfavourable conditions for local network ownership. Exemptions were made in South Africa for WiFi beyond the confines of a single premises in 2008 [9]. This allowed for small-scale wireless operators without a national telecommunication licence to operate.

In Zambia, regulatory frameworks for ICT are not necessarily conducive to widespread scale-up of ICT in rural areas. National issues, like parliamentary approval of procedures or amendments of regulations restrict progress [64]. There may be unclear executive powers, with enactment constrained by frequent changes in political environments or by legal positioning of multinational corporations such as mobile phone companies. The status and implications of regulations, licensing procedures and progress, are not clearly gazetted by the government. As a consequence frequent and personal interactions with authorities (in towns) are required to understand legal boundaries.

Attempts to install an OpenBTS-based low-cost GSM system such as VillageCell in Zambia have been challenging. Zambia does not have a legal framework for allocating test frequencies for research, unlike the FCC system of granting experimental licences for radiated power levels lower than 8 Watt. The Zambia Information and Communication Technology Authority(ZICTA) has been apprehensive to allocate GSM frequencies for VillageCell tests due to the commercial nature of GSM, even though VillageCell offers a good solution for universal service programs that ZICTA manages.

2.9 Technology constraints

There has been a considerable amount of work looking at technical obstacles to ICT inclusion in rural areas [26, 51, 160, 134, 36]. These highlight issues such as limited bandwidth, remote management challenges, low cache rates for proxy servers and dealing with large amounts of virus traffic. We now highlight some of the technical challenges we have uncovered that are not covered in previous work.

Some of these challenges receive more detailed attention in later chapters of this dissertation.

2.9.1 Linux vs Windows

Windows users in Macha were experiencing poorer performance in the LinkNet network than their Linux or Mac OS counterparts. This was especially the case with traffic requiring more upload capacity, such as Skype. Thorough network analysis of satellite Internet traffic captured over 2 months in early 2011 revealed a large disparity in performance between Windows and Linux/Mac users. Aggregate traffic was broken into hourly bins and then normalized to throughput per IP address per hour. Windows users were separated from Linux/Mac users using the Time to Live (TTL) field in the TCP header; Windows uses a TTL of 64 and Linux/MAC uses a TTL of 128. Table 2.1 shows the results.

During this measurement period, there were almost double the number of Windows users compared to Linux/Mac users logged in on average. However, Windows outgoing normalized aggregate traffic was three times worse than Linux/Mac during the phase of slow satellite connectivity. Incoming Windows normalized traffic was only marginally worse than Linux. After upgrade of the Internet connection to 2 Mbps the Windows performance improved. However, outgoing traffic speed was still two thirds that of Linux. Upgrade of the Internet link also improved the

Satellite connectivity phase 2011/01/28 to 2011/04/09	
Linux/Mac avg MB/IP/hour in	5.53
Linux/Mac avg MB/IP/hour out	1.55
Linux/Mac avg IPs/hour	9.75
Windows avg MB/IP/hour in	4.19
Windows avg MB/IP/hour out	0.55
Windows avg IPs/hour	19.24
E1 2Mbps line connectivity phase 2011/04/09 to 2011/05/19	
Linux/Mac avg MB/IP/hour in	6.33
Linux/Mac avg MB/IP/hour out	1.50
Linux/Mac avg IPs/hour	10.07
Windows avg MB/IP/hour in	7.93
Windows avg MB/IP/hour out	1.00
Windows avg IPs/hour	17.12

Table 2.1: Windows vs Linux/Mac normalized incoming and outgoing throughput.

round trip time for a packet. A satellite link induces an extra 600 ms round trip delay for the ground-to-satellite portion of the path, this delay was reduced to approximately 100 ms after the 2 Mbps Internet connection upgrade.

Further, to confirm that this scenario could be replicated in a lab environment, a Linux and Windows 7 machine were connected to a Linux server over a 1 Mbps line and an artificial delay of between 10 ms and 1 s was introduced to see the effect of increasing the delay on these operating systems. Table 2.2 shows the results of this simulation.

No delay introduced (10 ms systemic delay present)	
Linux TCP flow only	892 kbps
Linux TCP flow (Windows flow added)	822 kbps
Windows TCP flow only	968 kbps
Windows TCP flow (Linux flow added)	151 kbps
1s delay introduced	
Linux TCP flow only	860 kbps
Linux TCP flow (Windows flow added)	858 kbps
Windows TCP flow only	110 kbps
Windows TCP flow (Linux flow added)	57 kbps

Table 2.2: Windows vs Linux simulation of increasing network delay on outbound flows. Throughput measurement was carried out using the iperf tool for 20 seconds.

The simulation confirms the significant negative effect of delay on Windows machines. It also confirms that Windows is unfairly disadvantaged when there is a mix of Windows and Linux flows present and delay is high. The key reason for this discrepancy lies in the different TCP protocols used by Linux and Windows. Linux uses TCP CUBIC which changes its congestion window on the basis of the last occurring congestion event. TCP Compound, Windows' implementation, uses a delay-based window which is adjusted according to the Round Trip Time (RTT) of the last TCP packet.

There are a number of solutions to improve TCP performance over high-delay links [100, 133, 52] but none are available as a simple ready-to-use package or add-on application for a network gateway service such as clearOS used in Macha.

Installing these solutions takes considerable expertise in operating systems and advanced networks — skills that are not available in rural areas in Africa.

Although there was a sound technical reason for the difference in performance between Windows and Linux, it is challenging to share these findings with frustrated users. Users remained convinced that the performance issue was due to a poorly designed network or, worse, that Windows users were being deliberately disadvantaged.

2.9.2 Unfamiliarity with New Concepts

The African villages we visited feature Internet connections that are sometimes amongst people who have never had a fixed telephone network or even television access. This is not uncommon for infrastructure-deprived rural areas, where broadband Internet is often the principal means of mass communication. With the Internet, however, also arrive new foreign concepts that may not be intuitive to rural denizens.

In the Peebles Valley deployment, we observed unique behavioural patterns for first-time users of the wireless mesh-based Internet in rural areas. Wireless routers were powered down at night to save power even though they had little impact on total household energy usage. In wireless mesh architecture, user devices are connected to the Internet via each other; thus, powering down ones own device often causes the disconnection of other users in the network.

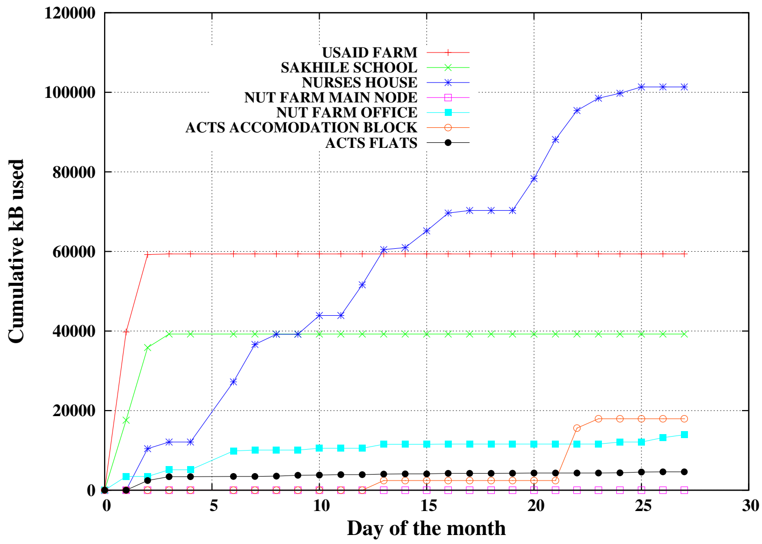

Figure 2.11: Bandwidth consumption in Peebles Valley for different users. Users are unfamiliar with the concept of bandwidth and often reach their allocated limit early in the month.

Similarly, we observed misinterpretation of the concept of network bandwidth. Figure 2.11 shows bandwidth usage in Peebles Valley over one month. The users have well-specified pre-allocated amounts of Internet traffic that they may download. We can see that some users, such as the Sakhile School, reached their limit quickly. There was little conceptual understanding of what type of Internet usage consumes large amounts of bandwidth. A graph similar to the one in the figure was available to the users, but it was of little use because the Internet users did not have sufficient educational background to interpret graphs. This includes concepts such as rate of usage and remaining capacity. A more tangible mechanism is needed on personal computers, such as a visible counter with remaining capacity. This is a more familiar concept because users often check remaining air time on their mobile phones. We observed a more sophisticated regulation of the use of the limited Internet capacity as users' online experience grew.

First-time Internet users are prone to online scams. Users in Peebles Valley were easily fooled into believing they had won huge sums of money and tricked into sharing their personal information online. Phishing scams through rogue websites were also noted. In general, it is hard for an inexperienced user to discern between a valid and a fake website or e-mail.

Computer viruses represent a significant hurdle to better-performing broadband in rural areas. We analyse the prevalence of computer viruses in Chapter 4. Infections can be prevented through up-to-date antivirus software and cautious online behavior. In rural area networks connectivity is slow, and, as our traffic analysis shows, antivirus updates often fail, leading to a higher level of infection susceptibility. Our interviews in Macha and Dwesa show that because of the slow connection, people share large files via USB keys. This method of sharing facilitates virus spread, which in turn creates more unwanted traffic in the network and consequently reduces the connectivity quality even more.

Machines running a Windows operating system are highly vulnerable to infections. Using Linux-based systems minimizes the risk of virus attacks. But many users, such as teachers at the school in Peebles Valley, were afraid to embrace an unknown operating system because they saw this as an extra hurdle to the already difficult task of becoming IT literate in Windows.

2.9.3 Network Growth versus Limited Capacity

To offset the high cost, users in rural regions often share a single satellite connection among tens or hundreds of people. Satellite connectivity is already associated with communication delay and throughput significantly lower than what can be provided with fibre optics. When a large number of users access the same satellite link, a very small amount of bandwidth is allocated to a single user. Moreover, a single satellite can be shared over multiple locations through WiFi mesh networks. The performance of these networks degrades with the number of network links that packets have to traverse on the way from the satellite gateway to the end user. This further limits link quality when sharing a single Internet connection in rural areas.

Additionally, users in remote regions are more likely to experience suboptimal online performance because of the way Internet applications are designed. The Internet is still largely centralized, and many popular applications such as

Facebook are hosted in the developed world [171]. Rural areas of the developing world are located far from data centres and servers that host popular websites. Even if the same quality of connectivity is to be provided worldwide, these remote locations will experience lag due to their physical distance from the application servers. Our network trace analysis in Macha reveals that the round-trip time of a network packet sent from the village, via the satellite link, to an application server and then back the same way can often exceed tens of seconds. Further, current centralized access models are extremely inefficient for local communication, which is dominant in rural Africa. In the current model, each instant message and voice-over-Internet-protocol (VoIP) call has to traverse a slow satellite link twice: once on the way from the sender to the central application server in the developed world and once on the way back to the receiver located in the same village as the sender. In our interviews, common complaints were a failure of instant messaging and dropped calls in VoIP applications, even when the communication is local. We examine this issue in more detail in Chapter 5.

Finally, while in the developed world dynamic ICT markets push service providers to upgrade the infrastructure and constantly increase broadband speeds, this is not the case in economically unattractive rural developing regions. At the same time, the World Wide Web is changing. From 1995 [49] to 2013 [5], the size of the average Web page has grown 90 times; the Internet has evolved from a strictly textual form into a media-rich environment with complex online applications. Connectivity in rural areas has yet to catch up with this high growth, and, unless the rate of bandwidth expansion increases faster than the rate of Web growth, predictions are that access in developing regions will effectively become worse than it is today [34].

2.10 Discussion

The environmental, skills and cultural constraints in rural areas emerge from day-to-day realities that are complex and multi-faceted. They affect the inclusion of ICT access in rural areas of the developing world. A broad, multi- and trans-disciplinary approach to ICT inclusion is necessary for a rural community that is very diverse and layered, incorporating engagement of rural communities on their own terms.

Many social and cultural assumptions are embedded in technology solutions. Their cultural specificity and their connection to power relationships are significant. These qualitative issues influence practical implementation of ICT in rural Africa. Thus, implementation of ICT in rural Africa is not only a matter of technological determination. Including ICT in rural Africa involves a holistic array of issues, informed by local context and culture.

Engineering and science mostly omits, or even excludes, the potential and usefulness of oral data. Cultural specificity of text-based, English-language and mostly western perspectives are often ill-suited for interpreting realities in rural areas of Africa. For instance, the inhabitants in rural areas that do not interact with texts often have high regard for the necessity and value of personal contact.

The constraints exposed in this chapter result in hampered access to ICT resources needed for the community to thrive. A lack of long-term, longitudinal research on use of technologies in disenfranchised areas affects the scientific knowledge base. In contrast with health science, in engineering there is no long-term scientific presence in the most disenfranchised regions of the world; there are no rural laboratories, or places equipped with engineering tools and workspaces where people can meet, share skills and work on collaborative projects within rural communities, aiming to understand and alleviate constraints. A vibrant, local community of inventors, engineers, scientists and entrepreneurs would facilitate learning from and teaching each other the practical and conceptual skills necessary to overcome the constraints for ICT access in rural Africa.

2.11 Conclusion

High-bandwidth Internet access is not present in many rural areas of the developing world. Sporadic attempts to bring broadband connectivity to isolated areas have been made, but a comprehensive evaluation of the challenges of bringing connectivity and the quality and impact of such connectivity is often lacking. This evaluation is important as it points out key technical and social challenges, in rural sub-Saharan Africa in particular, that need to be carefully considered when designing and deploying rural area networks. Due to the complex interplay of numerous technical and social issues that affect how people access the Internet and what they do online, it is inappropriate to evaluate the level of connectivity by measuring how many people have access to the Internet. The level of utility of the Internet is connected to many other factors such as locality of access and degree of locally relevant content.

A thorough examination of Internet usage at three locations in rural Africa consisting of network traffic analysis and Internet user interviews revealed several major obstacles to efficient broadband usage. These obstacles include unfamiliarity with new concepts, shortage of trained personnel, high cost of Internet access, limited Internet capacity and slow decision-making processes when initiating ICT projects across multiple local authorities amongst many others. This analysis is part of a large picture of ethical, conceptual and pragmatic issues. Constraints heavily influence practice and affect all efforts and activities aimed at bringing communication technology to rural Africa. Inclusion of ICT in rural Africa takes time and presence. Addressing local constraints within a local cultural setting is needed to assure ICT interventions are attainable, sustainable and empower the rural community.

Chapter 3

Rural Network Performance

3.1 Introduction

Rural wireless networks usually share a low-bandwidth, costly link to the Internet amongst a large user base. This means that any inefficiencies in the network can render a slow shared Internet link almost unusable. Analysing and understanding the traffic distribution, web usage patterns and source of bottlenecks can facilitate network designs that are optimized to give rural users a better Internet experience and bring down usage costs.

There has been a significant amount of work that has tracked Internet usage behaviour in urban areas over the past decade, but there is a large gap in analysis of Internet usage in the small set of rural wireless networks that are now in existence. For example, most recent Internet usage studies show that over half the Internet traffic is peer-to-peer (P2P) traffic. However, P2P traffic over a satellite link is costly and inefficient, and hence is likely to be less prevalent in a rural network.

In this work we analyse a set of Internet gateway trace logs, proxy access logs and mesh network logs from the LinkNet wireless network in Zambia over a period of 2 weeks in February 2010. The wireless network provides Internet access to approximately 300 residents of a rural village, as well as numerous international visitors. The results show that the network is heavily biased towards social networking-based web traffic and much of the potentially cacheable traffic, like Youtube videos, are not cached. A number of networking anomalies are also uncovered such as a large portion of packets with low TTL values that cause "TTL expired" responses and route flapping in the mesh network.

Based on observations from the traffic traces, we make suggestions to deal with the unique challenges of rural networks using slow shared Internet links. Some of the suggestions involve using data ferries, intelligent proxy caching, time-shifting large file downloads and packet filtering. We also suggest improvements for the mesh networks to prevent excessive route flapping.

3.2 Related work

The most up-to-date multi-year study of the distribution of Internet traffic in the developed world was carried out by Ipoque in 2008/2009 [151]. It showed that P2P traffic had decreased significantly and Web traffic had increased due to the extensive use of content servers. In Germany, for example, P2P traffic decreased from 69% to 52% and Web traffic rose from 14% to 26%.

The work that is the closest to ours is one that examines Web traffic usage in Internet Cafes and community centres in Cambodia and Ghana [51]. This paper only studied HTTP traffic and there was no wireless network aggregating traffic to the Internet connection. The traffic classification revealed that portal sites like Yahoo and MSN, advertisement sites and multimedia content accounted for the bulk of the traffic.

Wireless network performance in rural networks has been analysed in [20] and [160] but the Peebles Valley mesh network in South Africa [89] had the most similar network configuration. Analysis showed that in an unloaded network, the mesh would become the bottleneck, rather than the satellite, after 7 hops.

3.3 Background and network architecture

The LinkNet network in Macha, Zambia is described in Section 2.2.1. During the analysis period for this Chapter — early February 2010 — an Internet gateway server was configured with *ClarkConnect*[1], a pre-configured firewall and Internet gateway, to control access to the Internet. It had no user-management functionality, but some rudimentary user-management was done by checking host-names and their associated IP addresses in cases where bandwidth was used excessively.

3.4 Goals and measurement process

Our goals are three-fold: The first goal is to understand the usage patterns of users in a rural context to evaluate differences from an urban setting and gain insight into the needs of users in a rural setting. The next goal is to understand the performance of the network from the application level down to the physical level to find out what unique challenges are prevalent in rural networks. Our final goal is to make use of the learning from the first two goals and suggest ways in which the performance can be improved.

To meet these goals, two measurements points were required. The first measurement point was located at the gateway and captured all Internet traffic on the

[1]http://www.clarkconnect.com/info/

interface to the satellite and to the wireless network. The packets were captured in pcap format and a capture length (snaplen) of 96 bytes was used to minimize the size of the log file. This snaplen size was chosen to capture enough information from all the headers in the network packet. In order to analyse the HTTP traffic, the squid proxy access logs were also archived for analysis. All IP addresses were anonymized in order to protect the privacy of the users.

The second measurement point was located at the mesh gateway nodes to monitor the quality of the links in the mesh network. The monitoring daemon sent the ETX[2] values of every link in the mesh back to the gateway server every minute. The number of hops taken by mesh network packets was calculated by examining the TTL values of packets entering the gateway from the mesh.

14 days of traffic were captured from midnight, Sunday 31 January to midnight Sunday 14 February in 2010. Approximately 50 GB of packets were captured, consisting of about 6 million packet flows. Captured traffic was compressed and sent to the USA for analysis during off-peak hours. User-management software installed in April 2010 established that 10% of the traffic was from International visitors. A similar traffic distribution between the local population and visitors was likely for February but this could not be used to establish the type of traffic visitors were generating in the original data set.

[2]Expected Transmission Count (ETX) measures packet loss and is used as a routing metric.

3.5 Traffic usage analysis

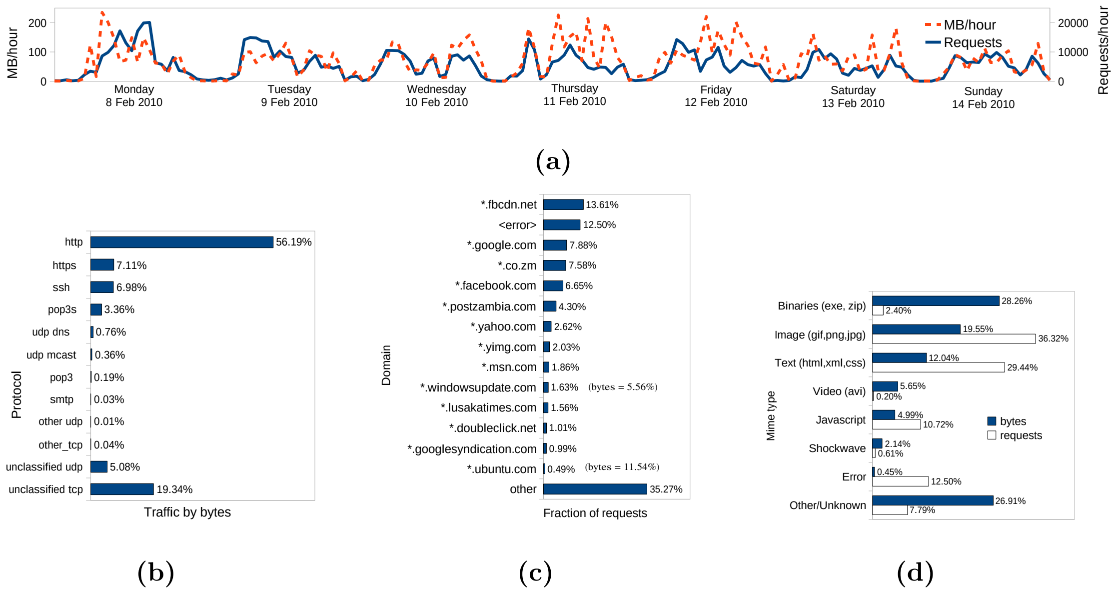

Figure 3.1: Usage analysis over 14 days: (a) Plot of the total web traffic for the satellite link over the second week; (b) Traffic classification by protocol; (c) Web traffic classification by domain visited; (d) Breakdown of traffic by MIME Type.

In this section we present an analysis of the usage patterns from squid proxy logs over 14 days organised into 1 hour bins. The plot in Figure 3.1a shows the total number of requests per hour as well as the total megabytes sent and received per hour. It follows a typical diurnal pattern but the off-peak period is very short due to users staying up late or waking up early to make full use of the extra available bandwidth during these hours. The number of requests generally tracks the amount of data downloaded except, for example, early Friday and Saturday morning where a single user was very active.

The proxy had a cache hit rate of 43% with an actual bandwidth saving of 19.59%. This low fraction of bandwidth saved is fairly common due to the dynamic nature of the Internet today; we offer some recommendations in Section 3.7 to improve this. The cache size was set to 1 gigabyte and studies have shown there is very little gain from increasing the cache size beyond 1 gigabyte [51].

3.5.1 Content distribution

To help understand the spread of traffic types, we classify packets as either TCP or UDP and then further classify these by their known Internet Assigned Numbers Authority (IANA) port numbers. Combined outbound and inbound traffic from the gateway to the satellite link was analysed.

Figure 3.1b shows the results of our classification. TCP traffic accounts for 93.24% of the traffic, which is consistent with most modern Internet usage trends. Web traffic accounts for 68.45% when standard HTTP and secure HTTP are combined and is clearly the dominant protocol. This is in sharp contrast to developed countries in which 2008/2009 studies show that Web traffic accounts for between 16% and 34% of Internet traffic due to the high prevalence of P2P traffic [151]. The large portion of ssh traffic was due to our traffic downloads.

There was also a significant amount of email traffic using pop3 store-and-forward delivery models. This traffic primarily occurred early in the morning before people arrived at work and made optimal use of off-peak hours. 26.47% of traffic could not be classified with simple port based techniques. This most likely consisted of Skype traffic, which is extensively used in Macha, as well as applications not using known assigned IANA ports. In future work we will use more advanced packet inspection tools to understand the breakdown of this unclassified traffic.

3.5.2 URL analysis

As most of the traffic is Web based, studying the sites that are visited as well as the traffic types will provide useful insights into user behaviour. Figure 3.1c classifies web traffic into the top 15 site domains. The most startling pattern that emerges is the dominance of Facebook. The Facebook host site and CDN make up 20.26% of the total requests. This is close to 3 times greater than the next most visited site, which is Google. A news article reported that Facebook overtook Google in the US in terms of number of hits in March 2010 [40]. From this data it appears that the crossover point in Macha occurred many months ago. Portals such as Google, Yahoo (*.yahoo.com, *.yimg.com) and MSN are clearly the next dominant site category (14.39% combined).

Local web sites in Zambia form the third most dominant category. This is encouraging in the sense that local relevant content and language is available to Zambians. Zambian news sites are also in the top 15 web sites visited, showing that digital news is a viable alternative when newspapers and radio are not available locally. Downloading software packages and updates from operating system sites such as Windows and Ubuntu are also in the top 15. Although they show a relatively low hit rate, the Ubuntu site, for example, has the highest share of traffic by bytes. Web site access follows a common pattern of a relatively small number of popular URLs and a long tail of other URLs. This long tail is the cause of a large pool of URLs in the "other" category.

2% of the total web requests are to known advertising domains and there are most likely many others that could be extracted from the "other" category by using pattern matching lists. The key issue, which is also raised in [51], is that the target population in developing regions is highly unlikely to be a customer of the advertised services. This essentially constitutes wasted bandwidth, which is particularly detrimental over a slow satellite link.

3.5.3 Traffic type and sizes

Because certain visited sites warrant large file downloads, we now move our focus to correlating types of web traffic with their size distributions. Figure 3.1d shows the MIME types of the traffic by bytes. There is a clear penchant towards binary downloads with large file sizes. The URL analysis in the previous section revealed that the sites that consume the most bandwidth are Windows update sites and package distribution sites for Linux. This corroborates with the evidence of binary files contributing to the largest fraction of bandwidth consumed. Large file sizes make up the tail end of the distribution, but they contribute to a large fraction of the traffic. For example, 99.9% of the traffic contains objects that are less than a megabyte, but the remaining 0.1% of the objects contribute to 50% of the total bytes downloaded.

Images and text form the bulk of the traffic types due to the dominance of web traffic in the network. The large amount of Javascript can be attributed to an increasing use of web applications when using web based email, for example. Multimedia also plays a significant role in the LinkNet network as can be seen from the Video and Shockwave traffic, which together make up 7.79% of the traffic. Inspection of the squid proxy logs showed that a significant amount of Youtube requests were made but all resulted in a cache miss. Youtube videos are served by a CDN and cannot be cached using a traditional proxy caching application like squid, which depends on a unique IP address and URI combination for its lookup.

We rank the 3162 Youtube hits over the two week measurement period according to total number of requests to the same video. The results of this analysis

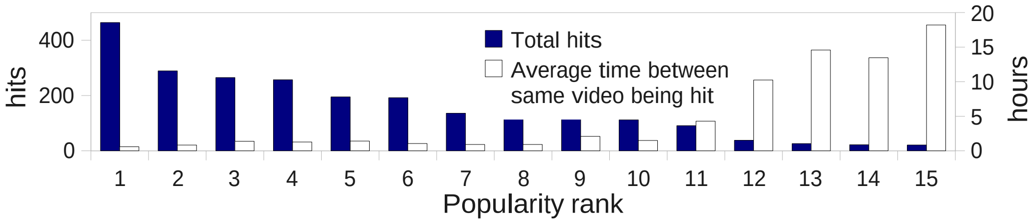

Figure 3.2: Youtube popularity.

are shown in Figure 3.2 for the 15 most popular videos. There were 451 unique videos, but the top 15 ranked videos contained 75% of the total requests. For the top ranked video, there were instances where the video was requested by six different users within a 10 minute interval. This would saturate the satellite link and the downloads would fail or be abandoned by the users. Anecdotal evidence from users in the Zambian network confirms this.

Youtube could have potentially consumed 60% of the satellite bandwidth if all 3162 Youtube videos were downloaded. However, Figure 3.1d shows that only 2.14% of the traffic by bytes is shockwave traffic. This is a result of about 87% of streaming media being abandoned by users in the first 10 seconds [60] as well as potentially slow or non-responsive Youtube downloads due to congestion. If a more advanced caching server was used and every requested Youtube video was cached in full, 86% of the bandwidth used for Youtube requests could have been saved. This is calculated by noting that there were 451 unique Youtube videos and 3162 Youtube video hits. Some of the repeated hits were due to users visiting

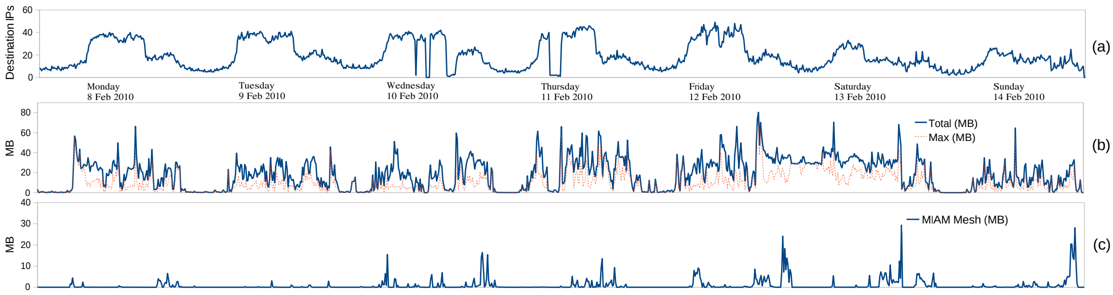

Figure 3.3: The breakdown of all inbound traffic on the satellite link in the final week of the 14 day measurement period using 10 minute bins: (a) Total number of unique IP addresses detected; (b) Total traffic vs maximum traffic for a single IP address; (c) Total bandwidth used in the MIAM mesh.

a web site containing an embedded Youtube video with the auto-play option set. Future work will quantify the extent of this problem as well as ways to limit its impact.

3.6 Network performance and troubleshooting

This section seeks to understand how the limited satellite connection bandwidth is distributed amongst the potential 300 users in Macha and how the mesh network performs under load. To do so we use traffic traces divided into 10 minute bins from the gateway and make use of mesh network traces from the mesh gateway nodes.

Figure 3.3a shows the number of unique IP addresses that were present in each 10 minute interval. This gives some indication of the number of concurrent users that were active in the network and ranged between 40 on weekdays to 25 on weekends. The four dips on Wednesday and Thursday of between 10 minutes and 2 hours were due to power failures. Some of these dips do not drop to zero as a few wireless routers are on UPS backup and laptops can also absorb short power failures. The mesh networks were only reflected as a single IP address and they would likely yield up to 5 additional concurrent users in a 10 minute window.

Figure 3.3b shows the total throughput on the gateway interfaces to the wireless network and the maximum throughput to a single IP destination in the network. This shows very clearly that there were single users that consume a large fraction of the bandwidth (50% on average). A single user would most likely be downloading large files while the balance of users were involved in more interactive activities such as web browsing.

Only the MIAM mesh was analysed as the ARK mesh gateway node was typically powered off or not as active as the MIAM mesh. Figure 3.3c confirms that the MIAM mesh was not a large contributor to the overall bandwidth consumption. Further analysis shows that it consumed the maximum bandwidth in the

LinkNet network only 4.74% of the time. The peak amount of traffic from the gateway to the MIAM mesh was 30 MB in 10 minutes, whereas single machines could achieve up to 80 MB. This is most likely due to a combination of a poor quality wireless link from the mesh network to the gateway and route flapping, which is discussed in Section 3.6.2.

There is a clear diurnal pattern in the traffic traces where almost no traffic was present between midnight and 6:00 AM. This pattern is broken on Friday night and early Saturday morning, where there was a continuous stream of activity from a low number of users. Analysis of this period of high activity showed that 52.86% of this traffic was from unclassified TCP ports compared to the 19.34% average on unclassified TCP ports over the 14 day measurement period. This may be torrent traffic for weekend entertainment. Note that there is no TV or radio reception in this area and no DVD rental store nearby.

The flow distribution over the 14 day measurement window exhibits a long tailed distribution with the most demanding flow using 22% of the bandwidth on average but as much as 90% of the bandwidth during quieter periods. The average flow size is 8.4 kB and the average flow length lasts 70 seconds. The top 0.1% of the largest flows contain 50% of the data; this correlates with earlier analysis of traffic types and sizes.

3.6.1 Performance analysis

To determine the source of networking problems, ICMP messages and HTTP messages were analysed. There were 828 "TTL expired in transit" messages *per hour* and 3900 "destination host unreachable" messages *per hour*. Further analysis showed that "TTL expired" messages were caused by the same number of packets *per hour* leaving the mesh network with TTL set to 3, 4 and 5. This made up 7% of the total traffic by packet count leaving the mesh networks. These were UDP DNS requests and were most likely caused by a malfunctioning DNS server or possibly a computer virus that wastes satellite bandwidth. Either the packets with TTLs below 32 should be discarded or the TTL should be reset before leaving the gateway.

Looking at the HTTP messages, there are 723 "503 Service not available" messages *per hour* making up 9.75% of all HTTP message responses to an HTTP GET. This is a very large fraction of HTTP requests not successfully completed. There are two reasons that this could occur. The first is that the DNS request could not be completed by the gateway server. The second is that there was no response to the HTTP GET from the upstream server. Both would be due to congestion in the network and confirm that interactive browsing is being hindered by a fully saturated satellite link often due to large downloads. Another possible cause of network congestion is a large amount of computer virus traffic. Many computers in Macha do not have updated virus scanners and it is highly likely that they are infected.

3.6.2 Mesh analysis

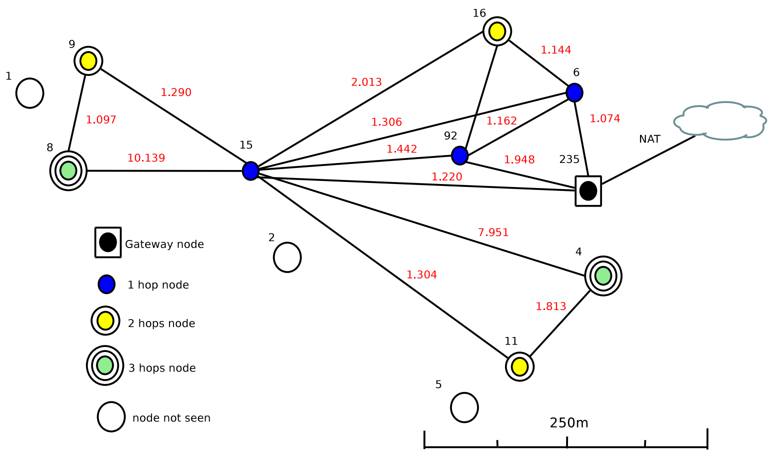

Figure 3.4: Scaled snapshot of the MIAM mesh during the 2 week measurement period with ETX values indicated per link. Note that the number hops is determined by the route with the lowest sum of ETX values.

The two mesh networks described in Section 3.3 make use of the Optimized Link State Routing (OLSR) protocol with the ETX link metric. An in depth study of the latency and throughput performance of the MIAM mesh network was carried out using active measurement techniques [16]. The 1 hop routers could achieve an average of 5 Mbps to the gateway router, but 2 and 3 hop neighbours could not achieve over 500 kbps. As a result, when the satellite link is not saturated, the bottleneck could become the mesh for 2 and 3 hop neighbours, but this was never observed in the 14 day measurement period.

Figure 3.4 shows the connectivity between all nodes in the mesh network and a snapshot of the ETX values present in the mesh at an instant in time. The average node degree in the network is 3.33 and no node has less than two potential links. However, some of the links are very poor with ETX values above 5. As an example, if node 9 or 11 go down, node 4 and 8 would need to route through very poor links. To understand the stability of the mesh network, the ETX values were tracked over 2 weeks. A Friday and Saturday portion of this is shown in Figure 3.5 together with the load on the mesh network.

It is clear that as the load on the network increases, the ETX of the weaker links tends to increase rapidly. Another effect of the changing ETX values is route flapping; this is also shown in Figure 3.5. A small amount of route changing can usually benefit a network as the routing protocol tries to discover a more optimal route. However, route flapping that approaches two route changes per second, as seen here, is detrimental to network flows.

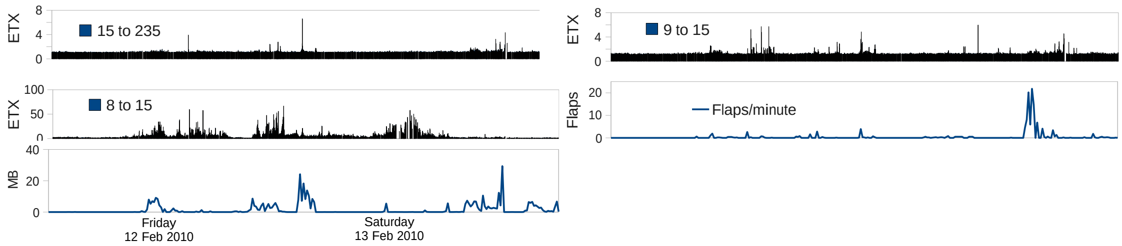

Figure 3.5: Graph showing how load on the mesh network affects the ETX link metric as well as route flapping.

3.7 Recommendations

There are two key components of rural wireless networks that need optimization to connect to the Internet through low-bandwidth links. The first is the gateway server and the second is the wireless mesh network.

3.7.1 Gateway server

Due to the limited amount of expert networking skills in the area, a *ClarkConnect* pre-configured gateway server was used in Macha. However, this and other similar pre-configured gateways are not optimized for severely constrained Internet up-links and the following customizations would have a significant positive impact on the network:

- Changing the caching behaviour of squid to be able to identify when the same content is being served by a different URL when accessing content from a CDN like Youtube. This can be done by rewriting the URL request to only preserve the static video ID. Other techniques, such as value-based web caching, deal with dynamic URLs by generating indices based on document content [51].

- Making use of cache optimizations such as HTTP chunking, time shifting and data compression discussed in [51], which would improve the throughput and response times for users. Time shifting would disallow large file downloads during peak hours and delay the download to off-peak hours.
- Filtering requests to advertising domains such as *doubleclick.net*.
- Filtering requests that are guaranteed to fail, such as packets with small TTL values.
- Installing a network monitoring tool, together with traffic measurement daemons at strategic congestion points, to help a network administrator understand network usage and performance as well as diagnose network problems. For per-user monitoring, the network subnets will need to be re-designed to avoid using NAT at the mesh gateways.
- Making use of a data-ferry synchronization station using a dedicated small-footprint PC such as "PC in a plug" and external hard drive. This will

be used to leverage frequent travellers visiting cities with low-cost high-capacity Internet links as data-ferries to fetch large popular repositories such as Ubuntu.

The Macha network is in the process of extending bandwidth and installing user authentication systems with preset and varied "service level agreements". This should help alleviate some congestion, especially once monetary value is associated with the number of bytes downloaded and the unit value increases during busy times. User management is a sensitive issue and it is highly likely that the set of rules will evolve as the balance between what is best for the community and individual users is accommodated.

3.7.2 Mesh network

In a similar fashion to the gateway server, the OLSR mesh network utilized in Macha was chosen based on being widely adopted and well documented. However it has been shown to be a sub-optimal solution and the following possible improvements could be made:

- Replace the OLSR routing protocol with a protocol that has proven better performance; protocols such as ExOR [22] or B.A.T.M.A.N. [91] perform better in the presence of poor links.

- If there is still a desire to continue to use OLSR, then replace the ETX link metric with Expected Transmission Time (ETT) to improve the overall throughput.
- Install mesh network monitoring daemons that make use of differential monitoring [138] on the wireless routers and send this information back to the gateway server.

3.8 Conclusion

It is clear from these results that the behaviour of the Internet in a rural wireless network connecting through a satellite is very different from that of an urban network in a developed region. From the usage analysis, it is apparent that the traffic is mostly web based; however, most of the traffic is not cached even though some simple proxy changes could make this traffic cacheable. A significant portion of the traffic consists of large binary files to operating system repositories. Most of this traffic could be eliminated by using people who travel to cities as data-ferries. Bandwidth is also wasted on advertisement servers, which have a very small likelihood of being relevant, as well as outgoing packets with small TTLs, which are likely to fail. The mesh network routes flap severely under load.

Other routing strategies such as stigmergic or opportunistic routing will most likely be better suited to this environment.

More attention should be given to building pre-packaged networking solutions for rural wireless networks that are cognizant of the characteristics that have been highlighted in this paper. Some of the problems were addressed in the early years of the Internet when last-mile access was similar to what is currently experienced in rural networks, but many of the problems are new as the Internet has become more dynamic and media-rich. As the average web page size continues to grow, the digital divide will widen unless innovative networking techniques, which mitigate these increasing bandwidth demands, are employed.

Chapter 4

Rural Network Internet Usage

4.1 Introduction

Bringing reliable, usable Internet connectivity to remote regions is typically plagued with problems. As discussed in Chapter 2, satellite connections are slow, often with bandwidth of only a few hundreds of kbps or 1 Mbps [117]; power sources are unreliable and devices are frequently unavailable [160]; networks are managed remotely or by poorly trained local staff [26]; and public Internet cafés have limited availability and high per minute usage costs [172]. However, despite these problems, Internet access has already revolutionized the lives of rural residents. Access to food production and health care information, distance learning programs, and global and local business opportunities has led to vast improvements in health care, quality of life, and economic earning potential. The better the quality of Internet access, in terms of availability, reliability and performance, the more residents stand to gain from their online activity.

Understanding the impact of technology in a developing rural region is a complex undertaking; it requires a multi-faceted approach that consists of both detailed traffic analysis similar to that carried out in Chapter 3 and social engagement. Local social customs discussed in 2.6 strongly influence Internet viewpoints and cannot be ignored when trying to interpret observed usage. Unfortunately, there are multiple challenges in technical and social analysis. Remote management and lack of local skilled technical staff render these networks extremely hard to monitor. Partnership with local technical organizations that help deploy and maintain infrastructure is essential. From the social science point of view, on-site data gathering is a very tedious process in remote areas. Villages are often dispersed over a large area with poor roads and lack of communication infrastructure. Language barriers pose additional challenges; often translators are required. Again, partnership with local organizations is critical for obtaining meaningful data.

In this Chapter, we focus entirely on Internet access issues in rural Africa rather than performance issues within a local wireless network and Internet gateway highlighted in Chapter 3. We examine Internet access availability and reliability, as well as Internet usage and obstacles to a fully engaged Internet experience.

To do so, we conduct a holistic study that encompasses two remote villages and both technical and social investigation of network usage. We carried out a more in-depth analysis of the same data set used in Chapter 3 to analyse Internet usage and performance. We complement our data collection with a series of interviews in Macha, Zambia in order to gain in-depth understanding of the usage patterns of local constituents. We performed the same survey in Dwesa, a remote village in South Africa, that obtained Internet access via a satellite link and a local WiMAX network in 2005. In total, we conducted 37 interviews.

The analysis of our two-week trace data reveals unique access periods for different categories of web sites, a small fraction of long-lived flows consuming the majority of the bandwidth, a constant stream of malware traffic and poor network performance with large round trip times. We observe a clear influence of the network performance on the user behaviour. In addition, Internet usage habits change once access is available at home, as opposed to being limited to public terminals with restricted usage hours, such as schools and Internet cafés. These findings suggest that rural area Internet penetration should not be evaluated through simple binary "have" and "have nots". We show that only through a comprehensive socio-technical study of a network can we obtain a full picture of the Internet usage in rural areas of the developing world. Such an outcome is crucial for a finer synergy between those who design computer networks and those who use them.

This Chapter is structured as follows. Section 4.2 provides necessary background information. In Section 4.3 we analyze the traffic characteristics in the monitored network in Macha, while in Section 4.4 we discuss Internet usage trends with an emphasis on time variation. Section 4.5 examines the presence of virus and other malware traffic in the network. Deeper socio-economic analysis of the interview results is presented in Section 4.6. Based on our analysis, in Section 4.7 we recommend a set of improvements for the Macha network in particular, and possibly rural area wireless networks in general. We complete the paper with the overview of related work in Section 7.7 and our conclusions in Section 7.9.

4.2 Methodology

The local environment as well as the LinkNet network, providing Internet to approximately 300 users in Macha, Zambia, is described in Section 2.2.1. Background on the Dwesa community in South Africa providing Internet to schools and households using satellite and WiMax in given in Section 2.2.2. The process used to gather interview data in both Macha and Dwesa is given in Section 2.3.2.

We extracted two types of statistics from the data: descriptive and mean/category-comparing. With the former we try to provide a clear picture of the Internet usage in Macha and Dwesa, while with the latter we provide possible explanations for certain observations. Due to the small number of samples we did not perform regression analysis but concentrated on the independent samples t-test for comparing means and χ^2 test for comparing categories. We report test results for which the two-tailed significance was lower than .1. We feel that this slightly looser requirement can be justified for the sample size and the domain in which we are working.[1]

4.3 Traffic Characterization

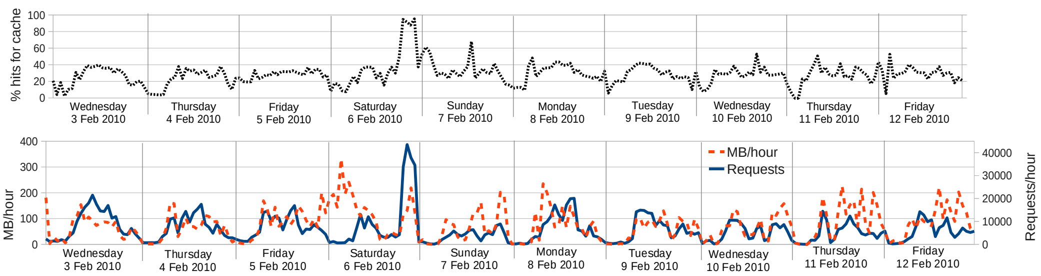

Figure 4.1: Usage analysis over 10 days.

[1]We report statistics according to the American Psychological Association standards: χ^2 statistics are reported with degrees of freedom and sample size in parentheses, the Pearson chi-square value, and the significance level; T-tests are reported like χ^2, but only the degrees of freedom are in parentheses; mean values are labeled with M_i.

In this section we seek to understand the high level usage characteristics of the network, and the network's ability to support the offered load. We are particularly interested in how the bandwidth constrained satellite connection affects network performance. We analyze throughput over time to understand the day/night cycle of usage and then proceed to study the distribution of flow sizes and lifetimes to determine the impact on the user experience. TCP round trip times (RTTs) are used to understand the delays users experience during interactive browsing and real-time activities such as instant messaging and VoIP. The performance of the cache also gives clues into network behaviour.

The proxy had a cache hit rate of 43% with an actual bandwidth saving of 19.59%. This low fraction of bandwidth saved is fairly common in a standard unmodified squid proxy server due to the dynamic nature of the Internet today [51]. The cache size was set to 1 gigabyte; studies have shown there is very little gain from cache sizes beyond 1 GB [51].

Figure 4.1 shows the traffic load, number of web requests and cache hit rate over the 10 day measurement period. What emerges is a clear, typical diurnal usage pattern, with the exception of Friday night and Saturday evening. The lack of usage during off-peak hours is due to the inaccessibility of public Internet terminals during this time. Not surprisingly, our interview data shows that those who have access at home are more likely to use it after-hours ($\chi^2(1, N = 28) = 5.2, p = .041$).

Friday night and early Saturday morning traffic displayed a sudden increase in aggregate download rate caused by a small number of requests per hour (approximately 700 or 10% of the average request rate). Inspection of the trace files reveals that these large downloads were requested from two single machines. Examination of the proxy logs, during this same period, indicates this was mainly due software updates and some Facebook and sporting web site accesses. A more detailed discussion of web usage behaviour over time is given in Section 4.4. On Saturday evening a satellite failure resulted in anomalous proxy behavior; we discuss this event in Section 4.3.1 when we analyse HTTP response codes.

Three power failures occurred during our monitoring interval: two on Wednesday 10 February and one on Thursday 11 February. These power failures caused corresponding dips in network usage. Power failures in Macha generally last anywhere from an hour to a few days. Our interviews revealed that on the average six such failures happen in a month.

Given the cyclical usage pattern and the bandwidth constraints of the satellite link, it is clear that the available bandwidth could be better utilized. In particular, users and/or administrators should be trained to set up systems that time-shift large downloads to periods when the network is quiet, using cron jobs, for exam-

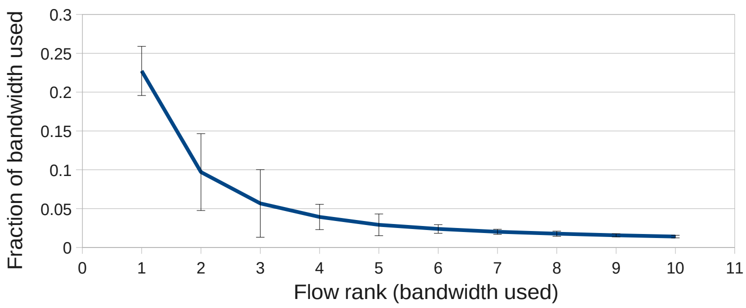

Figure 4.2: Flow distribution, organized from largest to smallest flow by bytes in one hour bins, over the 10 day measurement window.

ple. This would offload some of the peak-hour traffic, providing more capacity for real-time and interactive traffic. We discuss this further in Section 4.4.

4.3.1 Web object and flow analysis

Studies have shown that traffic in the Internet is typically characterized by many small short flows, such as web requests, and a few large flows, such as file downloads/sharing [27]. Understanding the distribution of flows becomes critical in a network connected over a slow satellite link, as the presence of large flows will likely make interactive activities, such as instant messaging and web browsing, very slow, and possibly even unusable. To determine the traffic composition in the Macha network, we look at flow distribution.

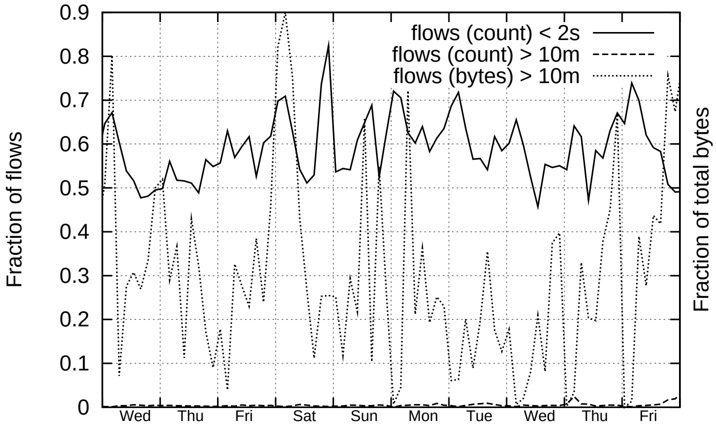

Figure 4.3: Fraction of short and long-running flows and fraction of bytes used by long-running flows over 10 day measurement window.

To understand the general behaviour of all the traffic in the network we plot the flow distribution of the top 10 flows, shown in Figure 4.2, ranked by size of flow. The figure shows the average size of the flow in one hour bins together with standard deviation over the 10 day measurement window. The largest flow, in terms of size, consumed about 22% of the bandwidth while the largest ten flows account for approximately 75% of the bandwidth demand. Thousands of smaller flows form the rest of the long tail distribution not shown in this figure. A flow is always established between two single hosts and thus a small subset of users are consuming the majority of the traffic in this network.

These smaller flows make up many of the interactive applications on the Internet such as web browsing and instant messaging. To understand the balance of small and large flows both in time and in size, the flows were categorized into short flows lasting less than 2 seconds and long flows lasting more than 10 minutes. This is plotted in Figure 4.3.

Typically, about 60% of the flows are short-lived (less than 2 seconds in duration). However, there is also a consistent fraction of long-lived flows (over 10 minutes), about 0.47% of the total. The fraction of long-lived flows increases during the weekend when more large files are downloaded. Though a small fraction of the total number of flows, the long-lived flows often consume a large fraction of the bandwidth, as much as 90% of the traffic on one occasion. A system designed for rural areas needs to cope with this high volume of short flows. Policies could be developed that place a higher priority on short flows to ensure that interactive activities such as instant messaging do not suffer from long delays.

The impact of large flows, as well as latency from the satellite link, is reflected in the RTT of the TCP traffic. Figure 4.4 shows the RTT of TCP traffic measured in one minute intervals, averaged over all flows within each one minute window measured at the outgoing interface of the gateway. A typical GEO satellite will incur a round trip delay of approximately 560 milliseconds in an uncongested network. Figure 4.4 indicates that congestion causes delays far beyond those due to the GEO satellite. The 100 second RTT on Wednesday 02/10 was due to a satellite fault; however other large delays of up to 60 seconds represent real delays experienced by users when the network was congested. Many daytime RTT averages are as high as 3 to 10 seconds, making web browsing and real-time applications difficult, if not impossible.

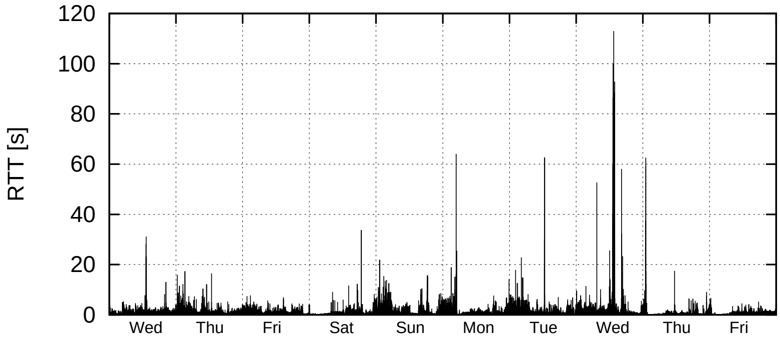

Figure 4.4: Average RTT measured in one minute bins.

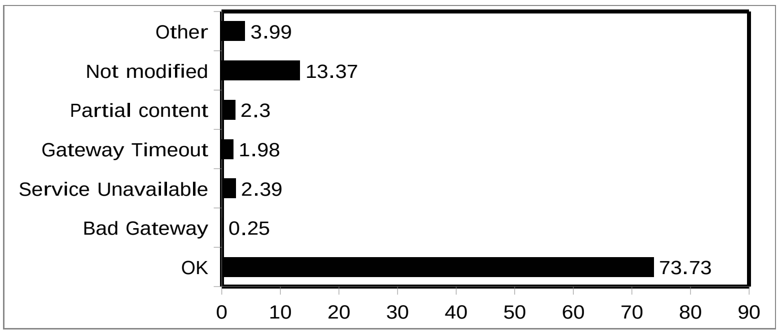

Figure 4.5: HTTP response codes for proxy server.

Another effect of large flows and congestion is session time-outs when the proxy is trying to retrieve web content. We analyze HTTP proxy responses to better understand web performance and plot the results in Figure 4.5. We find that during the 10 day period, approximately 86% of the HTTP requests were serviced successfully (found by summing the “OK” and “Not modified” responses). Approximately 4% of the requests resulted in a “Gateway timeout” or “Service unavailable” response, while another 2% received only partial content. To better understand web performance, we analyzed the HTTP responses as a time series (not shown due to space limitations). We find that the “service unavailable” errors occur only when the network is under heavy load, whereas the “gateway timeout” response consistently occurs at any time of the day or night. For example, 90% of HTTP requests received “service unavailable” errors on Saturday night; at the same time the overall number of requests peaked (Figure 4.1). We suspect that it is an automated application, rather than user behaviour that resulted in this phenomenon.

We discuss some ideas from existing and future work to mitigate observed problems in the next steps and related work sections.

4.4 Internet Usage Characterization

In this section we more closely analyze the offered load of the Macha network to better understand network usage in terms of application breakdown and web traffic classification.

In Chapter 3 we found that Web traffic accounts for 68.45% of total Internet traffic. Web traffic is far more pronounced than in the developed world, where recent studies show only 16% to 34% Web traffic [151]. In our analysis in Chapter 3, we did not specifically identify any peer-to-peer (P2P) traffic. Since P2P rarely uses known ports we left the possibility that it could be hidden in the 26.47% of traffic that we could not classify. However, our interviews reveal that P2P is indeed not popular in Macha. None of the interview participants reported using this application. Instead, 78% of interviewees transfer large files, such as movies and music downloads, from hand to hand via USB keys. This way of sharing alleviates the problem of the gateway bottleneck in a P2P system. Unfortunately, it has detrimental consequences on network security as it facilitates virus spreading.

The most popular online applications are web browsing and email. Our interviews reveal that 100% of those who use the Internet in Macha use both applications. The next most popular applications are VoIP and instant messaging via Skype and gTalk - 73% of the interviewees use these applications. This suggests that a large part of the unclassified traffic most probably belongs to these applications. A comparison of PC users in Macha and Dwesa yields interesting results. Macha users are more experienced and are more likely to use web browsing ($\chi^2(1, N = 37) = 8.07, p = .014$) and email ($\chi^2(1, N = 37) = 5.99, p = .035$).

Examination of web proxy logs reveals that web site accesses can be roughly grouped in the following categories: “Facebook”; “Web portals” - including Yahoo, Google and Live.com; “News sites and Youtube” - including Post Zambia and Lusaka Times[2] as well as youtube.com; and “Software updates” - including OS and anti-virus updates. Note that the proxy log entries obscure the difference between the user intended behavior (e.g. accessing a news site), and the automatic hits (e.g. scheduled software updates). We extend our analysis by considering Web access patterns over time. In figures 4.6a and 4.6b, we break down the total web requests and total web traffic load (in bytes) of each domain type over time, respectively. The values are shown averaged over hour intervals.

From Figure 4.6a we see that, at almost any time, “Facebook” and “Web portals” account for the largest number of HTTP requests. This is not surprising as 77% of interviewees in Macha reported using Facebook and 96% report “Googles” for work or school related material and practical problem solutions. However, the pattern by which the web accesses occur is quite different for the two categories. Web portal access frequency does not change much based on the time of day - the percentage of traffic to portals is roughly constant. Facebook, on the other hand, shows distinctive peaks late at the night. Since our interview data reveals that the late night users are most likely accessing the Internet from their homes,

[2]www.postzambia.com, www.lusakatimes.com

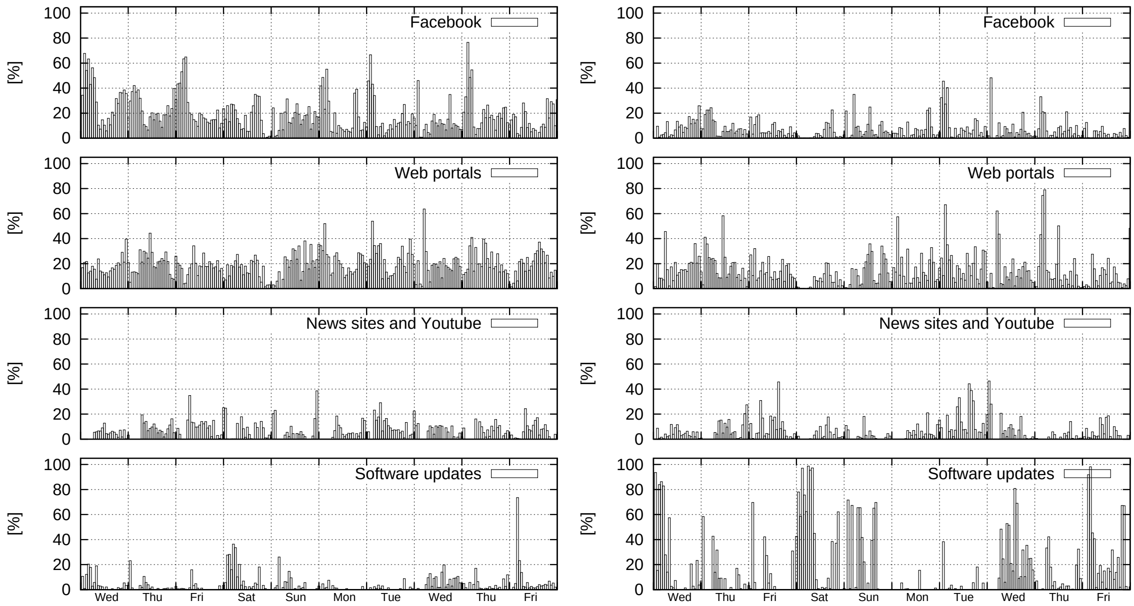

(a) HTTP category requests over time

(b) HTTP category traffic load over time

Figure 4.6: HTTP traffic requests

we deduct that restricting connectivity to public locations within a rural area, such as Internet cafés and schools, severely limits the usage of OSNs and leisure applications.

The traffic load distribution in Figure 4.6b shows that the majority of bytes transferred often belongs to a small number of flows, which is a similar finding to prior studies [27]. These are typically due to software updates, which can frequently consume the majority of the available bandwidth. Burstiness of these

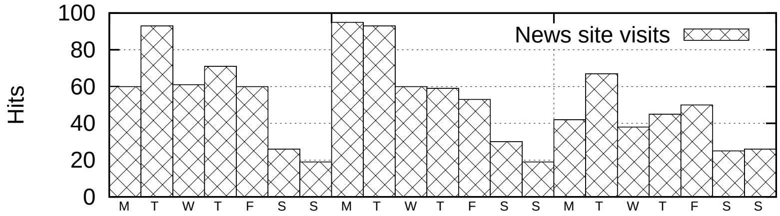

Figure 4.7: News site visits per day.

heavy-weight requests leads us to believe that they represent automatic, scheduled OS and anti-virus software updates.

4.4.1 Media access

Access to news sources is often limited in rural areas. There are only three TV stations in Dwesa and one in Macha. Newspapers from larger cities arrive highly irregularly. Interviewees praised the Internet for its provision of multiple news sources. All of the interviewees in Macha and 79% in Dwesa consider the Internet to be the preferred source of information over television and newspapers. Prior work has shown that Internet usage impacts legacy communication sources in the developed world [97]. We see the same trend in the two African villages. People who extensively use the Internet (five or more hours per day) devote drastically less time to television (M_1 =11 hours, 45 minutes per week; M_2 =3 hours 15 minutes per week; $t(28) = -3.08, p = .005$) and printed media (M_1 =2 hours 30 minutes per week; M_2 =0 hours 0 minutes per week; $t(27) = -5.52, p = .000$).

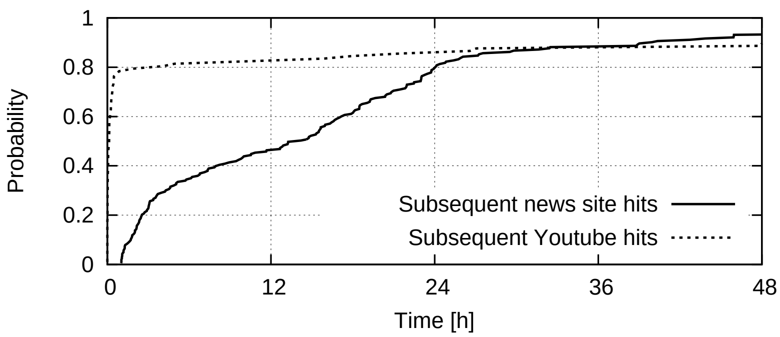

Figure 4.8: CDF of media site visit interarrival times.

That the Internet is heavily used as a news source can be observed in Figure 4.6a. We see that news site access quickly peaks in the mornings when people arrive to work. The usage then diminishes until late in the evening, when almost no accesses are observed. We extend our analysis over a three week proxy log that overlaps with our ten day measurement period and plot the news site hit distribution over that period in Figure 4.7. We see that the news sites are up to four times more popular during week days than weekends. Interestingly, we observe that the number of news site visits decreases as the week progresses, a phenomenon that is not uncommon for the developed world either[3].

To isolate individual user behavior we calculate inter-visit times for individual IPs and plot the CDF of news websites and Youtube visits for all the unique IPs in Figure 4.8. We observe that most visits to the news sites are made with an interarrival time of less than a day and that the distribution is rather uniform,

[3]http://www.alexa.com/siteinfo/nytimes.com

indicating periodic checking for new information every few hours or each working day. Youtube visits, on the other hand are bursty - almost 79% of the visits by the same user happen within the same hour. We intuit that this can be explained by the Youtube website organization where a visitor is often presented with multiple related videos. In addition, a number of interview participants complained about the unpredictability of Youtube performance. We suspect that in this case the network limitations dictate the user behavior: once the Internet connectivity provides a usable Youtube quality of service, users seize the opportunity to watch multiple videos they were unable to download earlier.

4.5 Malware traffic

Keeping a computer free from malware in a developing country is a challenge due to the need for up-to-date virus signatures and the increasing rate of new forms of attacks. These signatures can be made available as often as twice per day due to the speed at which new attacks appear. The challenge in a rural area is particularly great because virus signature downloads consume a large amount of the scarce satellite bandwidth, particularly when each computer must individually download its own patches. There is also a lack of specialist knowledge to analyse intrusion detection system logs, enforce timely software patching, and assist with removing malware when machines become infected. In addition, users have a perception that network administrators are solely responsible for protection of the network.

Approximately 60% of machines in the Macha network run Microsoft Windows; the remainder of the machines run Linux or MacOS. Windows machines are particularly vulnerable to Malware attacks, especially when the latest security patches have not been installed.

4.5.1 Detection method

We utilized the *snort* intrusion detection tool for detecting malware. Snort makes use of a set of rules to detect the presence of malware on the network. It is capable of detecting a virus payload on a network as well as the presence of suspicious traffic once the payload has infected a machine and malware traffic is being generated. Our analysis was carried out in non-real-time on truncated packets; it was therefore not possible to detect virus payloads, which require deep packet inspection. We looked for the presence of all known malware type attacks. These are:

- ICMP scans: malware searches for a host to attack

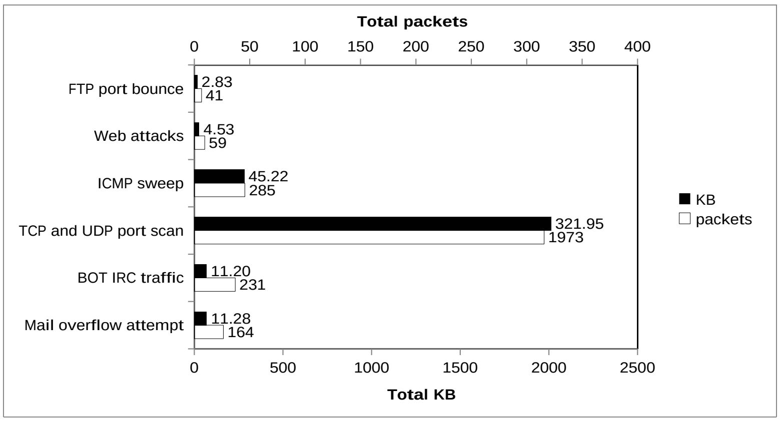

Figure 4.9: Types of attacks on the Macha network.

- UDP and TCP port scans: malware has found a host and is looking for open ports on this host to attack
- Bot traffic to IRC server: malware contacts a public IRC server to send messages and receive commands
- Web attacks: malware tries known web site attacks such as buffer overflows

4.5.2 Analysis

Figure 4.9 shows the type of attacks present in the network, broken down by total bytes consumed per attack. This is typical behaviour of a bot, in which an IRC server is contacted as a command and control center and the bot scans for potential machines to attack. Out of the 201 unique IP addresses seen over the 10 day measurement window, 9 machines were infected by bots trying to contact an

IRC server. 118 machines were carrying out some form of port scanning, however this may have some false positives due to other software that carries out port scanning such as *nmap*. Figure 4.10 shows the regularity of bot and port scan attack traffic over time. There is a clear pattern of the virus contacting the bot IRC command and control center at certain times of the day. The port scanning traffic is fairly consistent over the the full measurement window except for some brief periods late at night and early in the morning where activity decreases due to some machines being powered down.

The total impact on the satellite network cannot be accurately quantified as many of these malware bots transmit large amounts of traffic, such as mail spam, on well-known Internet ports. However, with such a high percentage of infected machines generating unwanted malware traffic in a network, we surmise that the traffic volume due to malware will be significant, negatively impacting an already strained satellite link.

4.6 Social Considerations

Internet connectivity has great potential for transforming a society. On one hand it can create multiple opportunities, while on another it can polarize the society and create a digital divide between those who do and those who do not

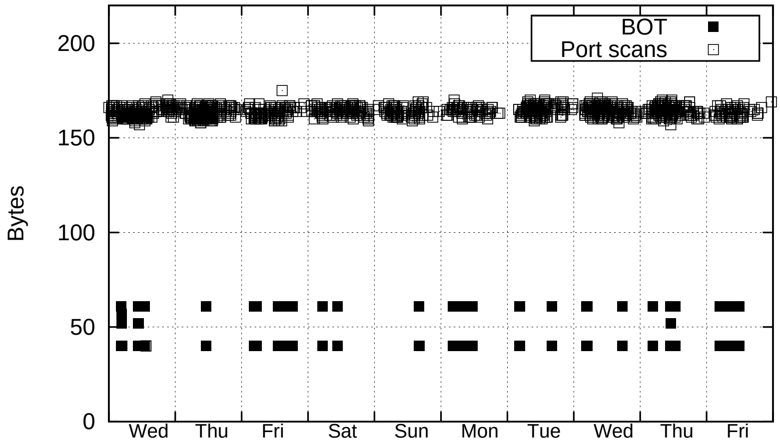

Figure 4.10: Time series of virus traffic.

have access. In this section we examine the societal impact of Internet provisioning in Macha and Dwesa. Our analysis is based on the interviews we performed in both villages. Unless otherwise specified, the aggregate results are presented.

4.6.1 Economic impact

Benefits of Internet access have been quickly realized by the rural Africans. In Macha, for example, local farmers have used the Internet to gain expertise on crop rotations, a move which completely revitalized the local agriculture [117]. Similarly, one of the interviewees in Dwesa trades local arts and crafts online. Unfortunately, a lot of the opportunities are missed because of the lack of large scale plans and business infrastructure. In many rural areas, such as Macha and Dwesa, banks are not present. In fact, obtaining a credit card is often impossible.

While micro-financing solutions help with the local economy, they cannot be used to compete in the global economy. Online trading has been attempted by 28% of the interviewees, often with unsatisfactory results. The lack of nation-wide Internet adoption plans severely hampers many of its useful aspects. For example, e-government services are not available in Zambia.

The Internet can create new local jobs, but it can also promote migrations as the local population can apply for distant jobs online. We observe a striking difference between Macha and Dwesa in that sense. In Dwesa, people are much more likely to use the Internet for job searches ($\chi^2(1, N = 30) = 11.32; p = .001$). South Africa has the second highest GINI coefficient in the world[4] and migrations to affluent areas are very common. Thus, our results show a particular case of the technology adaptation to the local context, not a case of technological determinism.

4.6.2 Social capital

Online social networks (OSNs) are highly popular in rural Africa, with 69% of interviewees being active OSN users. While more preferred among the younger population[5], they are much more than a leisure activity. In Macha 29% of the

[4] GINI coefficient is a measure of income inequality.

[5] The average age of an OSN user is 26, while the non-user average age is 36 years ($t(30) = -2.63, p = .024$).

interviewees use Facebook for some sort of business correspondence. For example, a local pastor uses Facebook to send spiritual messages to his church followers. OSNs can serve as a great cultural bridge - 82% of the interviewees have remote online friends whom they had not yet met in the real world.

Social ties are established over OSNs with both local and remote friends and relatives (77% have OSN friends within their village and 91% have remote OSN friends). To a lesser extent email is also used for local communication (47% use it for local and 91% for remote correspondence). Of all the means of online communication instant messaging has the highest level of usage locality (80% use it for local and 75% for remote correspondence). The results are not surprising as asynchronous communication provided by email and Facebook messages is more suitable for cases where personal contact does not happen often. However, the insights should be carefully considered from the network systems point of view as the locality of interaction can be used to save the satellite bandwidth [171].

4.6.3 Gender roles

When the Internet revolution started in the developed world, a digital divide appeared as the first appropriators of the new technology were mostly young white males [68]. Since the relative uniformity of the demographic in Macha and

Dwesa and our sample size do not allow race and age bisection, we concentrate on differences in Internet usage among men and women.

While both men and women use computers, women we interviewed are less likely to own a computer ($\chi^2(1, 37) = 8.29, p = .007$). Women spend significantly less time using the computer: on the average men spend 22 hours per week using a computer, while women spend 8 hours ($t(34) = 3.21, p = .003$). When online, both genders are equally likely to use email, OSNs and Youtube.

In Macha, men are more likely to perform computer maintenance themselves than women ($\chi^2(1, 22) = 6.14, p = .023$). They are also more likely to do virus scans ($\chi^2(1, 22) = 4.62, p = .054$). In Dwesa we did not observe such a discrepancy. The reason is, we believe, that one of the project leaders in Dwesa is a woman and a number of local champions are women. As a consequence, the outreach is stronger and women are not likely to see such tasks as "men only".

4.7 Next Steps

It is clear from both our traffic analysis and our interview data that the network in its current configuration undergoes periods of unusability; extremely slow response times, web page request timeouts, and instant message losses are some of the key problems observed. At the core of these problems is over-saturation of the slow satellite link due to the large number of users. In Macha, there are IT support staff who have achieved an impressive skill set in network engineering, considering that many of the administrators who have participated in the network deployment and maintenance do not have a college education. However, the problems experienced in this network require advanced skills in networking engineering. Due to this limited skill set, a pre-configured gateway, *ClarkConnect*, was used. Such a preconfigured gateway is not optimized for rural bandwidth constrained links; there are some fairly simple improvements that can be made using known solutions. There is also scope for appropriate technology research to mitigate these problems. For example on the client side, a different abstraction of web search with underlying intelligent proxies could be provided. Some good examples are the TEK search engine [108] that delivers a low-bandwidth copy of a web page to a user using email and Ruralcafe [38] that provides an expanded search query interface, allowing users to enter additional search terms and maximize the utility of their search results.

We focus on interventions in the network and from our analysis, we identify the need for four key areas:

- The caching behaviour of squid must be changed to identify when the same content is being served by different URLs, as in the case of CDNs.

- Content servers, to mirror services such as Wikipedia, operating system updates and virus updates, need to be installed in Macha with traffic redirected to these servers.

- Traffic destined for local recipients, such as shared pictures and local voice calls, needs to avoid traversal of the satellite link.

- Internet access should be extended to more homes so that off-peak hours of the satellite link can be more optimally used.

There are additional actions that, if enforced, would notably improve the performance of the network. For instance, when a machine with malware is detected, it should be disconnected from the network and only be reconnected once it has been disinfected.

The first two interventions have some known solutions [6], however these need to be constantly adapted as the web evolves. The third intervention is an area that will require new solutions; determining which traffic is of local interest is not a trivial problem. This is especially true of web-based social interaction [171]. The fourth intervention requires solutions to more widely cover rural houses. This is a challenge because homes are often many kilometres away from the community center where Internet is available. There are some new options, using white spaces

[6] http://drupal.airjaldi.com/node/264

spectrum, which can achieve longer distances and non-line-of-sight links that go beyond the current gambit of WiFi solutions for rural areas.

4.8 Related Work

The performance of rural area wireless networks has been investigated since their inception in the early 2000s. Work by Chebrolu et al. [33] and by Sheth et al. [156] deals with wireless propagation problems over long distance links. A comprehensive study of rural area network problems is presented in [160]. In this paper, technical issues were juxtaposed with social obstacles of deploying networks in the developing world. Our work is similar in a sense that it also considers both perspectives. However, Surana et al. concentrate on system troubleshooting with an accent on the energy problems. In addition they provide only anecdotal evidence of the social problems. Our study is geared towards understanding of user behaviour and provides a quantifiable description of social occurrences. Specifics of network usage in the developing world is also the subject of [152] and [104]. The former investigates a specific application (VoIP) and its economic feasibility in rural areas, while the latter includes interviews of telecentre users in rural India. In this paper, we do not focus on a single application and supplement our on-site interview data with a comprehensive network trace.

From a large body of work that analyses the behaviour of the Internet traffic, [51] is the closest to ours; it examines web traffic usage in Internet Cafés and community centres in Cambodia and Ghana. While it serves as a good starting point, this paper only studied HTTP traffic, and there was not a wireless network aggregating traffic to the Internet connection. In our analysis, we analysed full TCP dumps and provide insights on a number of problems, including virus and bot presence. Moreover, the profile of Web traffic changed drastically. The popularity of Web 2.0 sites, such as Facebook and Youtube, that we observe in rural Africa could not be envisioned even in the developed world up to a few years ago.

Estimates, in 2007, of the 600 million machines connected to the Internet infected with botnets range from 150 million which are vulnerable to infection to a conservative estimate of 12 million currently infected by bots [57]. A study at Aachen University in Germany in 2007 showed that in 8 weeks, 13.4 million successful exploits were discovered due to 2034 unique malware binaries circulating amongst 16,000 unique IP addresses [57]. Botnests, which are the key platform for most Internet attacks, are surveyed in [59]. This work highlights that botnet command and control has moved beyond the common methods such as contacting an IRC server. Methods used today make use of HTTP and P2P traffic and are thus harder to detect with common intrusion detection systems, such as snort. This emphasizes the fact that the malware we detected may only be the tip of the iceberg. While we are not aware of any previous work studying the problem of viruses in rural network deployments, the importance of user education in rural networks is highlighted in [84].

4.9 Conclusion

The results from our analysis, as well as from previous studies [26], highlight the fact that rural networks need to be treated as a special class of network due to their unique set of challenges. We find that while traffic is primarily web based, a large portion of cache-able traffic from CDNs is not cached. Further, there are many long-lived, non-real-time flows during peak usage times of the network that negatively affect the interactivity of web browsing and often cause instant messaging to fail. These large flows are typically due to automated file downloads, for example when an operating system requests an update. Very high TCP round trip times during the day, sometimes over 10 seconds, with high probability lead to a change in user behaviour; a highly interactive multi-search web experience is likely to become far more singular and deliberate in purpose, almost approaching the environment of a DTN network.

Chapter 5

Rural Network Locality of Interest

5.1 Introduction

The Internet has evolved both in terms of size and application since its birth in the early 1990s. There is an increasing amount of off-PC storage and processing using cloud computing for services such as navigation, photo sharing and file hosting. Many applications that in the past were run on a user's operating system, such as email, word processors and instant message clients, are now run on web browsers. These features have brought users in developed countries closer to the vision of "anywhere any-time" computing, where devices connected to high speed

Internet connections delegate computing power and storage capacity to cloud computing services. However, what is not clear are the implications for users in developing regions where Internet access speeds of between 128 MHz and 256 MHz are common, as was typical of dial-up users of the 1990s.

One of the key consequences of web-based computing is an increase in traffic load. In the rural villages we have studied in Zambia [90], there can be as many as 60 concurrent users sharing a single relatively slow satellite link. As a result, web-applications become increasingly slow, to the point where they are unusable.

We have seen in Chapter 3 that social networking is the most popular web application in the village of Macha, Zambia and users share messages, pictures, music and software with each other, where the sender and recipient are often in the same village. When a user shares an object with another user in the same village using an off-PC storage server such as Facebook or Dropbox, the same file traverses the slow satellite gateway twice: once as it is uploaded to the server, and again as it is downloaded to the recipient. This leads to congestion of an already constrained satellite link.

The most widely used solution to deal with inefficiencies in this centralized model is to use peer-to-peer (P2P) networking. Bittorrent is an example of a popular P2P networking protocol and accounts for between 27% to 55% of Internet traffic, depending on geographic location [151]. Bittorrent works on the principle of "tit for tat" in which users may not download content faster than they can upload content. Peers are always located outside the local network, rendering Bittorrent unusable or extremely slow due to the limited capability of the satellite link. Although we found only traces of P2P traffic (1% of total aggregate traffic passing through gateway) it did account for a surprising portion of traffic with large outbound flows (35% of all outbound flows greater than 100KB). Many users in Macha use Skype, which also uses P2P networking. A large portion of these large P2P outgoing flows may be Skype traffic, which is difficult to distinguish from Bittorrent traffic. We found no evidence of P2P traffic being directly routed between two local users as most local users have no direct network routes to each other and super-peers — well provisioned nodes for routing P2P traffic — are always located non-locally. Clearly traditional P2P networking has not been the panacea to save Internet gateway bandwidth in rural networks.

In order to find a more well suited solution to the inefficiencies of centralized web access, we analyse the locality of interest in the Macha network and its potential to save gateway bandwidth. We carry out two tasks in this study. The first task is to measure the strength of social connections in the village of Macha by extracting a social graph from Facebook instant message chat. This helps us to understand the potential of users to share information and files. The second task is to measure the fraction of local traffic sent directly between users without the use of the Internet, or the fraction of traffic being used to upload content to Internet file-sharing services in order to determine how effectively or ineffectively users are utilizing the local network.

We extended our collected traffic traces used in Chapter 3 and Chapter 4 to a period of two months in Macha, Zambia. This traffic includes all local traffic within the network and all traffic to and from the Internet. Facebook instant messaging was common in our trace; however there was no direct local file sharing and a very small fraction of traffic using file synchronizations (0.94%) or file sharing services (0.65%). We believe this is due to the poor nature of the network and the relatively high cost of bandwidth ($30/GB) rather than a desire not to share these objects; on-line interviews substantiate this hypothesis.

Our interaction graphs of Facebook instant message chats reveal that 54% of chats are between local users in the village, even though only 35% of the observed users were local to the village. We also find that people who travel to areas outside the village are strong sources and sinks for local information exchange. We extract other statistics, such as social degree and clustering coefficients, for the social graph. Packet flows are analysed to understand the percentage of traffic sent directly between users in the village and via central web servers. Finally, we correlate some of our findings with on-line interviews that were collected from 77 users in the village from a wide range of age groups.

Other than interesting anthropological inferences from this analysis, there are many technical implications. The high degree of local conversation provides strong motivation for an automated localization engine that is able to intercept content for local users and re-route it directly to a local user rather than traversing the satellite link. Relocating services that enable user-to-user interaction from the Internet to servers in the local village is also promising for improving local performance.

5.2 Related Work

Improving Internet usability in rural regions has been studied by Chen *et al* [35, 36]. In [35], an asynchronous queuing model was used where users can queue web requests as well as a cache search feature with predictive text. Users responded positively even though a custom web frame was used to search a local cache or queue content. In [36], a browser plugin was used to pre-fetch pages and serve stale cached pages. This achieved an average acceleration of 2.8x for users browsing non-video web pages. These results bode well for interventions that may need the web access paradigm to be modified somewhat. Our study looks at traffic locality in a rural network and its implications on improving network performance and, as such, represents another tool in the toolbox of techniques to improve network performance in rural regions.

A recent media and society study of the localization in the Internet highlights the fact that as the Internet continues to grow, it is becoming "more local" [136]. This phenomena is beginning to blur the boundaries between online and offline social domains, and it is this trend that demands a localization approach to network design, especially in isolated rural communities.

There are many studies on social network interactions, both at a structural level using friend lists and at an interaction level using wall posts [105, 121]. Wilson *et al* argue that social links created by friend lists are not valid indicators of user interactions [170]. This is shown by the fact that the number of "friend adds" account for 45% of the activity per day whereas comments only account for 10% of the activity. Interestingly the common notion of small-world clustering, which is present in a social graph derived from "friend adds", is absent from the interaction graph. Our work captures the physical locality of the users, which adds a new dimension to interaction graph analysis.

Locality of interest has been primarily studied in the domain of peer-to-peer networks. For example, a semantic clustering technique is employed by Handurukande *et al* in [66]. The semantic relationship is either implicit, using information such as peer-history, or explicit, using meta-information about the file, such as whether it is music or video. In a rural village, a P2P mechanism that first looks for a peer in a local subnet may be a possible solution to making file sharing more efficient.

5.3 Data collection and processing

5.3.1 Data collection

A simplified model of the network in the Macha community in Zambia is depicted in Figure 5.1. Computers are connected through a bridged wireless network to the gateway, which is connected to the Internet through a slow satellite link. All traffic in the network passes through the bridge that is monitored passively by a data collector. This allows us to capture all local traffic as well as traffic to the Internet.

Local or non-local direct links between machines using services such as secure copy are very rare in networks without advanced computer users and we found no evidence of this traffic. P2P traffic also creates direct links between machines and is common amongst many users in the developed world (between 42.51% and 69.95% in 2008/2009 [151]), but because of the high cost of data and low bandwidth capacity, it makes up a small portion of traffic in Macha (1% of total aggregate traffic). The majority of traffic is either to and from an Internet web server or between two local clients through an Internet web service. When using a web service like Facebook, a content delivery network (CDN) is typically employed to receive and deliver content. As shown in Figure 5.1, it is fairly common for a user in the village (Client 1) to send a packet to a service like Facebook through a CDN (CDN 1) and for a different CDN (CDN 2) to deliver a packet to another user in the village (Client 2).

Over the course of two months (February and March 2011), we captured traces of all traffic passing through the bridge, which connects all machines in the network, using tcpdump. The total traffic collected is 250 GB and is stored as multiple pcap formatted files.

5.3.2 Data characteristics

As described earlier, one of our main goals is to determine the quantity of traffic, from the Macha trace, between two local machines in the village. This is trivial to determine for traffic sent directly between two machines by checking whether the source and destination IP address are both in the local village subnet. However, determining this for traffic that passes through an Internet server is far more difficult. Hence, we need to compare uplink and downlink flows to look for similar patterns leaving and entering the network. These patterns are characterized by features such as IP addresses, port numbers flow sizes, content

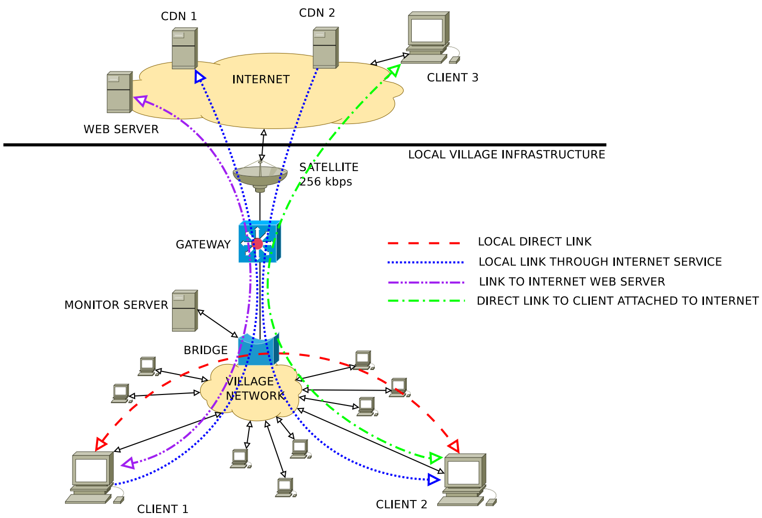

Figure 5.1: A simplified model of the network architecture in the Macha network. All traffic in the village passes through a bridge. Four types of possible traffic are highlighted. All traffic types are captured by our monitor server.

headers and, when these are not sufficient, features unique to specific objects, such as colour distribution in images. This traffic may be real-time, as in the case of VOIP calls or instant messaging, or it may be non-real time when using a file sharing service, for example.

Real-time traffic between two local machines routed through an Internet server is relatively simple to detect using IP address pairs for flow classification. Returning to Figure 5.1, if Client 1 routes a real-time connection to Client 2 through CDN 1, there will be a steady stream of packets from Client 1 to CDN 1 through the gateway. This same stream of packets will appear from CDN 1 to Client 2 within a time-window typical of the round trip time of the satellite link. If two different CDN servers are used, as is also shown in Figure 5.1, the CDN domain, rather than IP addresses can be used as a flow feature. For example, two Facebook CDN's for IM traffic could be "im-a.ak.fbcdn.net" and "im-b.ak.fbcdn.net"; we would use the domain "fbcdn.net" rather than the IP address to define the Internet server being used in the flow.

The task becomes more challenging with non real-time traffic where features such as IP addresses and flow lengths may not be enough to correlate an upload with a download. For example, when using services that host pictures or videos such as Picasa and Youtube, the hosting server will compress or manipulate the object in some way, such that the uploaded object does not match the downloaded object. Fingerprinting can help solve this problem. For example, an image can be fingerprinted using a histogram of colours and will be immune to compression.

A further complication is introduced when encryption is employed — common on many web mail and file sharing sites — rendering even fingerprinting techniques useless. In this study we were only able to analyse flows between local machines, real-time Instant Message Facebook flows and Facebook image sharing using unique ID information inside Facebook filenames. Extracting locality in flows that have been manipulated or which employ encryption is left as an exer-

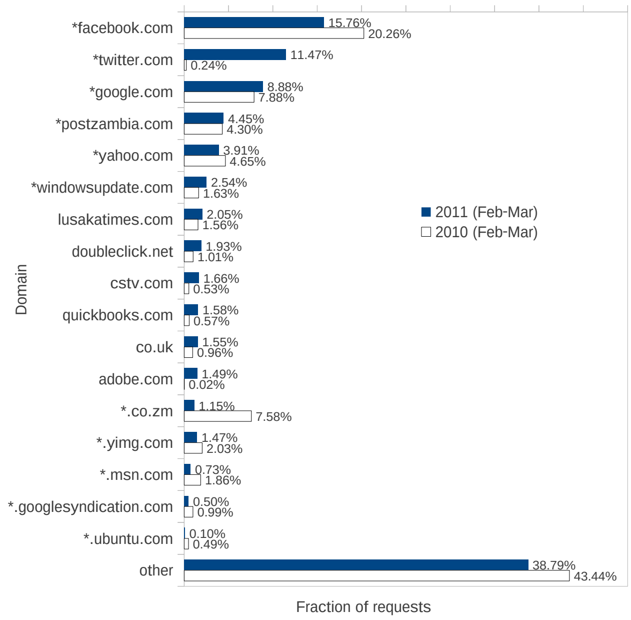

Figure 5.2: Internet usage analysis in Macha, Zambia for February and March 2010 and 2011.

cise for future work. However, we believe that the social structure of the online community in Macha, extracted using Facebook analysis, is applicable in other messaging and file sharing applications.

5.3.3 Identifying local Facebook traffic

In Chapter 3, we analysed Internet usage statistics in 2010 and found that Facebook traffic accounted for the majority of traffic in the network. We carried out the same analysis using our new 2 month 2011 network trace and found that this trend continued; however, some new interesting usage trends appeared. We compare the traffic analysis from 2010 and 2011 traces in Figure 5.2. Although Facebook declined slightly from 20.26% to 15.76%, Twitter, another social networking site, gained a strong following. While constituting only .24% of the network traffic in 2010, it has become the second most popular web service in 2011. Considering that both Facebook and Twitter are social networking services, social networking now accounts for almost three times the number of page visits compared with Google. The continued dominance of Facebook gives credence to our notion that Facebook traffic is a good representation of social connections in Macha, Zambia.

Facebook instant message traffic was extracted by searching for a header, unique to Facebook, in the HTML body of the packets in our network trace. Although Facebook uses a different schema for sending and receiving a packet during an IM conversation, we were able to simplify the analysis by only analysing the incoming packets due to the inherent behaviour shown in Figure 5.3. For every message transmitted, the same message is sent back to both the sending user and receiving user as it is displayed on the IM web client. To determine whether a conversation is local, we check whether we receive two packets with the same user pair on two different local machines (machines with different IP addresses).

In order to evaluate locality of interest for more bandwidth-intensive shared content, we also investigate the locality of Facebook photo sharing. Facebook

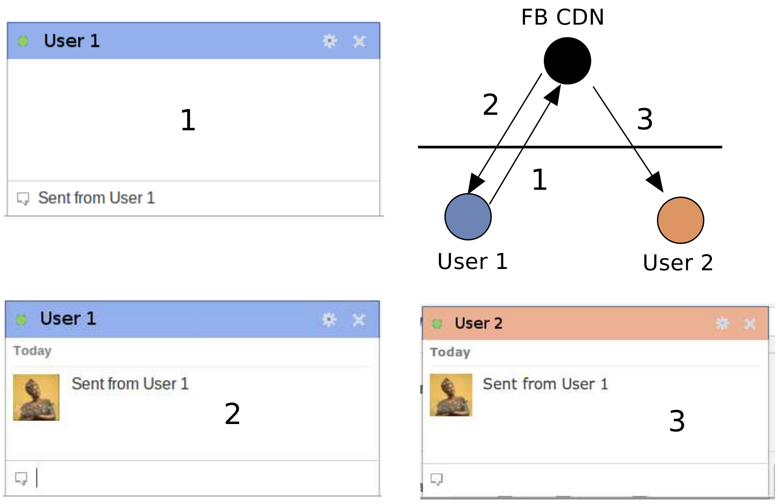

Figure 5.3: Outgoing and incoming packets when sending an IM message in Facebook. (a) A message is typed on user A's Facebook IM console. (b) The message is displayed on user A's Facebook IM dialogue area. (c) The message is displayed on user B's Facebook IM dialogue area. (d) A flow diagram of packets sent between client A and B and the Facebook CDN for each of the 3 events in (a),(b) and (c).

photos are downloaded via an HTTP GET request to Facebook CDN servers. Facebook encodes a user's unique Facebook ID into the filename of any user-generated image. The ID used in the filename will always be the Facebook ID of the user who uploaded the image. We exploit this feature to establish the quantity of Facebook image content that was locally generated and viewed, filtering out profile thumbnails as they have no bearing on locality of interest. This is done using a list of local user IDs extracted extracted through our Facebook instant message analysis and comparing these to the originator IDs of each Facebook user generated image.

5.4 Social Graph Analysis

In this section, we present an analysis of a social graph built from instant messages sent in Facebook. We call this social graph an "interaction graph" as it captures the actual interactivity between users rather than passive links. Instant messages represent the strongest possible relationship linkage between users in Facebook as they indicate a one-to-one mapping between users, and they capture true user interaction rather than passive relationships represented by "friend adds" and wall posts. Previous studies support the notion of different levels of relationships in Facebook [170]. They showed that wall posts and photo comments represent a much stronger indication of relationship bonds between Facebook users than "friend adds".

5.4.1 Social graph

In order to understand the relationships between users within and outside the village, we create a social graph that captures users as nodes and conversations as edges. The weight of the edge depicts the number of conversations in the two month measurement period between two specific users. Figure 5.4 shows the central section of the Facebook interaction graph for the village. The graph has been arranged so that users who are well connected with each other are closer to the centre. The color, as shown by the key, depicts whether the user was always in the village, always outside the village, or a traveller. Users were detected as travellers when they originated a message both from within and from outside the village during the trace. This was possible to detect as users were tracked using Facebook IDs rather than associated IP addresses, which are not persistent. Figure 5.5 shows the edge of the interaction graph with isolated communities of users, often having one to six local village users, connecting up to 20 outside users. A complete version of this social graph is shown in Appendix A.

There are a few key conclusions we can draw from these interaction graphs:

- There are key users in the village that act as strong links to the outside world. Often these key people are travellers. These users are easily identified by the fan shape motifs in the graph.
- There are a number of isolated communities where the majority of nodes are external, often linking to only one or two local users. From personal observations made in the village, these are likely researchers who visit the

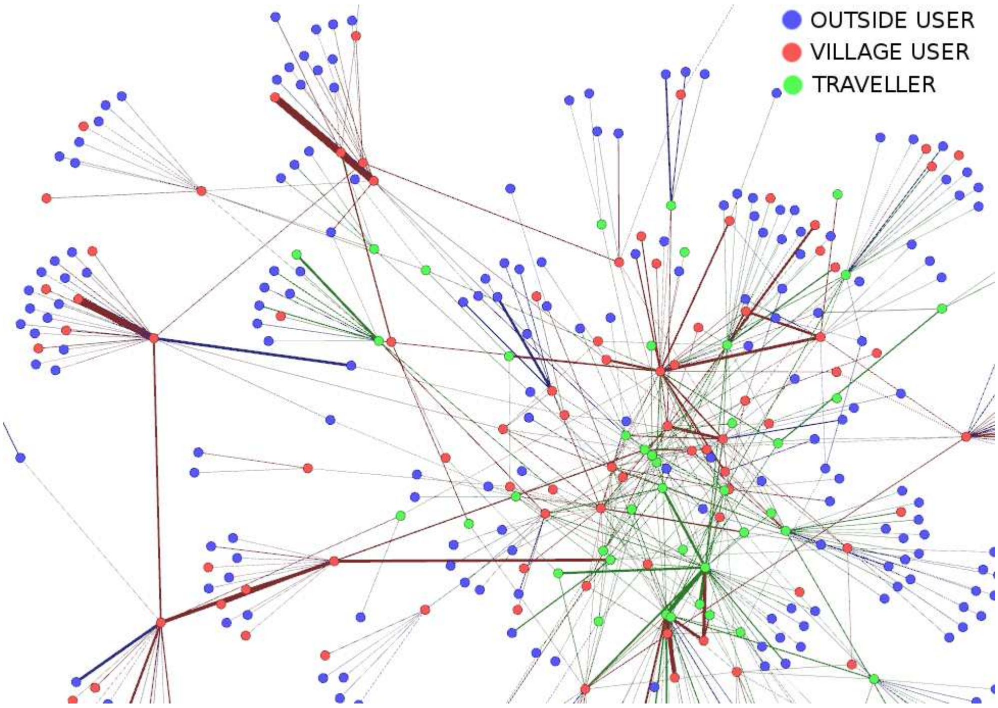

Figure 5.4: Interaction graph for Facebook IM traffic showing users as nodes who either remain in the village, travel but return to the village or who are always outside the village. The edges denote conversations. Thicker edges indicate that more instant messages were sent between users.

village to carry out their research and collaborate with overseas colleagues and have very little online interaction with local people.

- The strongest bonds (highest number of conversations) are between local users, especially those that travel and key contact people. These users most

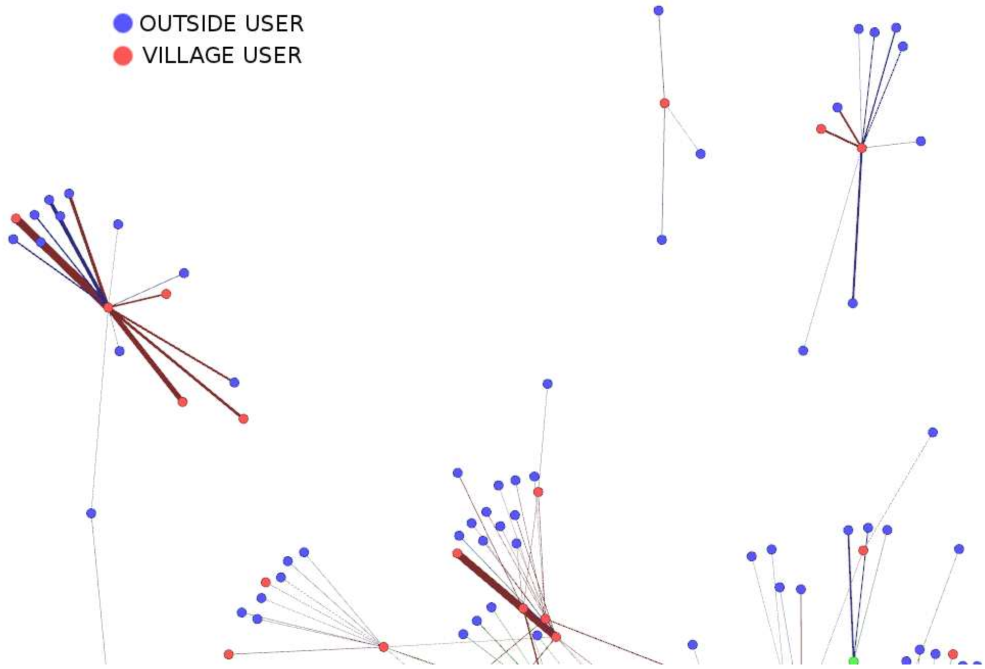

Figure 5.5: A section on the periphery of the interaction graph showing isolated communities of users.

likely form information brokers enabling the flow of information between users outside and inside the village.

- There are surprisingly few cases of one external user connecting to many users in the village. One interpretation is that users who arrive in the village come from a diverse set of communities, where people in each community only know one person in the village. Another is that outside users who meet travellers from Macha tend to use only one person as their key contact or gatekeeper in the village.

It is clear from this interaction graph that users in the village form a very close knit community with stronger links between other local users than to users outside the village. However, the graph also reveals that Macha is not a homogeneous community. There is a small set of isolated individuals and communities; on-the-ground observations reveal that these are often non-locals living in the village. The presence of non-locals in the community is common due to international visitors and relocations and rotation of Zambian health and education personnel nationwide.

In order to understand the characteristics of this interaction graph, such as the fraction of local messaging and clustering, we now perform more detailed statistical analysis.

5.4.2 Statistical analysis

Over the two month measurement period 573 unique Facebook users were identified, of which 140 were local users who never left the village and 43 were users who travelled. There were 14,217 unique instant messages sent between 726 unique user pairs. Our analysis reveals that 54% of the instant messages were between local users in the village even though only 35% of the users were in the village, with 7.5% of these local users occasionally travelling. This shows that instant messaging is used extensively for intra-village communication.

To understand the statistical distribution of these relationships and messages among the users in the village, we plot the cumulative distribution function (CDF) of the social degree between local users, and from local to external users in Figure 5.6. Social degree is a measure of the number of edges connected to a node; this correlates to the number of users with which a specific user has communicated. Only a local user's perspective is shown as the social degree of the external users cannot be fully known. We also plot the CDF of messages sent between local users, and from local to external users, in Figure 5.7.

Node degree between local users reveals a large number of small cliques and a small set of well connected community members who act as messenger hubs. Local users who have connections with external users generally tend to have a large number of connections. 70% of local users have a degree of 3 or less, compared to 7 or less for local to external users. The average degree of a local user is 3.6 and the average degree from a local user to an external user is 5.3. However, Figure 5.7 shows that local users send more messages to each other than to external users even though they have a lower node degree. Local to local user communication also has a heavier tail revealing a small subset of very active local users who message each other. Looking at the social graph, we see that well connected users are often also responsible for high message counts, confirming their role as information brokers. These results confirm that it is not enough to look at the connections between users in a social graph. Interaction between users reveals the true behavioural characteristics of users, and in our specific domain of interest, potential of users to interact locally.

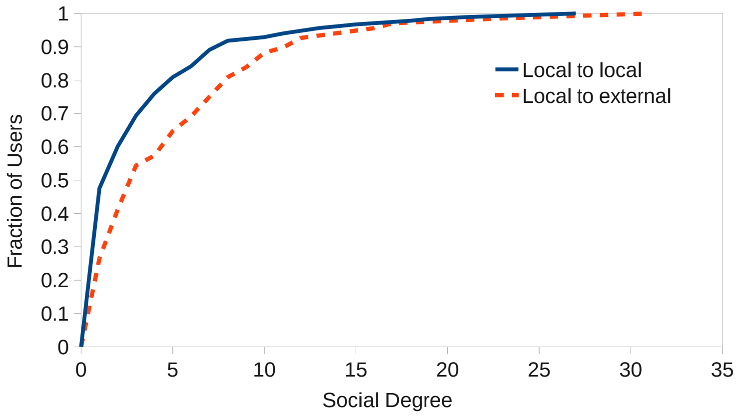

Figure 5.6: CDF of social degree between local Macha users in Facebook and from local users to users outside Macha.

.

In order to understand the level of cohesiveness of users in the social graph, a number of social graph related metrics are employed:

- **Clustering coefficient** measures the tendency of nodes in a graph to cluster together. For a node with N neighbours and E edges between these neighbours, the clustering coefficient is (2E)/(N(N-1)). A higher clustering coefficient can be interpreted as nodes forming tightly connected localized cliques with their direct neighbours.

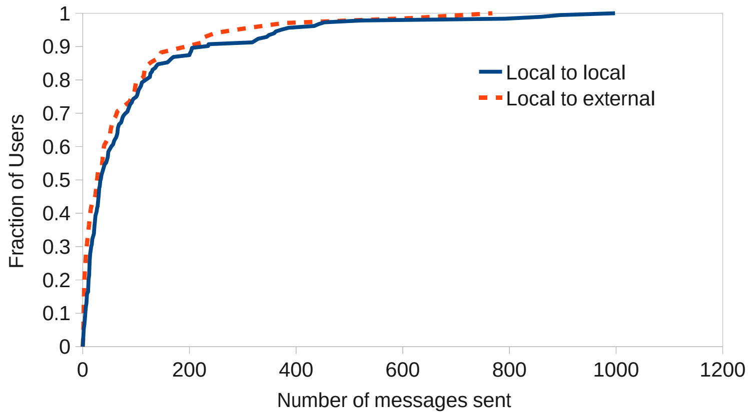

Figure 5.7: CDF of instant messages between local Macha users in Facebook and from local users to users outside Macha.

- **Average path length** defines the number of edges along the shortest path between all possible pairs of nodes in the network. This measures the degree of separation between users. A lower average path length indicates a community that is closely connected and through which information can spread quickly.

- **Eccentricity** of a graph is the maximum distance between any two nodes.

- **Diameter** is the maximum of all eccentricities.

The clustering coefficient is only measured for local users in the network as we do not have a complete picture of Facebook interaction for external users. This value effectively measures the cohesiveness of the local community rather than checking for a "small-world effect" in the larger Facebook community. The average clustering coefficient is 0.1 (average on facebook is 0.164 [170]) for all the local nodes. Taking into account that, on average, regular interactivity only occurs with one fourth of a user's Facebook friend list [170], this value still represents a local community that is strongly connected.

The average path length between all nodes in the graph is 3.798 (average on facebook is 4.8 [170]). This is far less than the six-degrees of separation hypothesis for social graphs [119] as we have an incomplete graph, with edges only extending to the first tier of external users. This, together with the clustering coefficient, provides further proof for a strongly connected local community.

The diameter of the local graph is 8 (average on facebook is 9.8 [170] and average on Orkit is 9 [170]). Considering that this is a localized graph, the value is surprisingly high. This is due to a set of outlier users who are weakly connected to other users in the network (low social degree). They form a a category of users who do not use Facebook IM as a regular means of communication and most likely use other channels such as different IM clients, SMS or email.

From this social graph analysis, it is clear that the potential for local interaction, other than instant message interaction captured in the social graph, is very high. If the infrastructure was supportive of local connectivity, users would share music, pictures, videos and software as a natural consequence of these relationships as is seen in well provisioned networks in developed countries. In the next section, we determine if indeed there are any attempts to send data between local users.

5.5 Locality of Facebook images

In this section, we present an analysis of Facebook images shared between Facebook friends. We are particularly interested in the fraction of local-user Facebook images that are viewed by local users in the village in order to extract the degree of locality of interest.

The results are summarized in Table 5.1. From our data set, we extract 6066 unique Facebook IDs from user generated images, of which 182, or 3%, are from local users. Although only a small fraction of Facebook users are local, these users generate a significant portion of image content observed in the trace. Our analysis reveals that 9% of the unique user generated images are local. However, 24% of the images viewed, including some viewed repeatedly, are local images. Further analysis shows that a local user views a local image 5.2 times on average, as opposed to 1.5 times on average for an externally generated image. This could be interpreted as just over three times more interest in locally generated images demonstrating strong interest in local content among Facebook users in Macha. However, we cannot make any claims about whether or not the same content from

Macha is also popular outside of the village. Thus, any localization solution also has to allow distribution of local content to the global Internet.

Table 5.1: Locality of Facebook image generation and sharing in Macha, Zambia. A large part of the actual content observed in the trace comes from the village itself, indicating high locality of interest.

Sample information	
Total unique Facebook IDs	6066
Total local Facebook IDs	182 (3%)
Image counts	
Total unique user images	27403
Total local unique user images	2370 (9%)
Total images viewed	50308
Total local images viewed	12273 (24%)

Section 7.4 reveals that Facebook image uploads represent 39% of the total upload traffic volume (Figure 7.7b). Much of this traffic could be saved by using a localised Facebook image and file sharing application; we propose such a solution in Chapter 7.

5.6 Local traffic patterns

Information can be exchanged between users using a deliberate or passive action. A deliberate action includes activities such as making a Skype call, transferring a file directly between two computers, or sending an email. A passive action involves a user uploading a file to a server, such as an image upload on Facebook, and another user downloading this file at a later point in time. In this section we analyse our traffic traces to detect deliberate actions to share content with other users in the village or use of services which may lead to passive content sharing.

We begin with an analysis of all traffic transmitted directly between local machines in the network.

5.6.1 Traffic sent directly between machines

The local wireless network makes use of approximately 100 802.11b/g wireless routers to connect close to 300 users in the village. These routers are typically able to operate between 1 Mbps and 5 Mbps depending on the distance between them and level of interference. Hence the local network has substantially higher capacity than the satellite gateway (128 MHz) and, as such, represents a large amount of unused spare capacity. In order to understand whether this local capacity is utilized we now analyse the amount of traffic transmitted between users in the local network.

We summarize the high level findings of local traffic in Macha in Table 5.2. Traffic, including gateway servers (DNS server, capture portal server), consists primarily of HTTP traffic as users sign onto the Internet using the local web

Table 5.2: Statistics for traffic between local machines.

Description	Quantity
Total flows	50,355,759
Total Data (GB)	249,338
Including gateway servers	
Total local flows	13,546,582 (26%)
Total local data (GB)	12,548 (5%)
Excluding gateway servers	
Total local flows	6,726 (0.013%)
Total local data (GB)	7 (0.0029%)

"capture portal". There were also many DNS requests to the local DNS server to resolve host names. This traffic constitutes 26% of the total aggregate traffic passing through the gateway. Due to the local network having between five and twenty times more bandwidth than the satellite link and the opportunity for local connections to create isolated flows between local users in different parts of the network, we conclude that the local network will have a large amount of residual capacity.

In order to understand whether users connect directly to each other in the network, we exclude traffic to the gateway servers, which perform common tasks of Internet login and DNS resolution, and calculate the remaining residual traffic transmitted directly between client machines. This makes up a very small portion of the traffic, constituting only 0.0029% of the total traffic seen during the measurement period. Further investigation of this traffic revealed that 44% is Netbios traffic, a Windows networking protocol, and 2.4% is samba traffic, a Windows network and printer sharing service. These are services which are automatically activated by Windows and do not represent an explicit desire by a user to connect to another user. No evidence was found of files being transferred using Samba or any other protocol such as the "File Transfer Protocol" or "Secure Copy". There was also no evidence of any VoIP calls being made directly between two local machines without the use of the Internet.

Clearly the best solution to conserve the bandwidth-limited and costly Internet gateway is to make use of services that can enable direct user-to-user communication or host content on local servers in Macha. For example, a simple FTP server could host music files that users wish to share or a Asterisk VoIP server could enable local calls between users in Macha using soft-phones. However, this requires advanced computer skills and breaks the current paradigm of web-based computer usage seen in Chapter 3. As a result, we primarily observe flow patterns where users make use of web services to share content or make phone calls. We discuss this scenario in the following section.

5.6.2 Traffic sent locally through Internet servers

As highlighted in Section 5.3.2, the process of extracting local traffic routed via servers on the Internet is a complex process. In order to determine the quantity of potential local traffic, we first analyse the destination domains of large outgoing flows (flows greater than 100KB). These flows have the potential to contain objects that may be downloaded later or streamed in real-time to local users in the village. We found that 15% of the outgoing traffic consisted of large outgoing flows. We divide this portion of traffic into domain categories in Figure 5.8. We are also able to check whether any of the Skype flows embedded in the "Skype and Bittorrent" category are routed between local users through an Internet super-peer. We leave the task of matching up and down flows in manipulated objects, such as photos and images, and encrypted flows for future work.

Skype and Bittorrent applications create the largest amount of upload activity. It was not possible to separate these two traffic types, as they both behave in a similar manner. They use a mix of UDP and TCP and random source and destination ports, and they exhibit balanced throughput for incoming and outgoing streams. Bittorrent over a satellite gateway with cap-based-Internet access is a costly option as users pay for somebody else downloading fragments of the file. This is due to the 'tit-for-tat' nature of Bittorrent. Hence we expect that most of this traffic is Skype rather than Bittorrent (strong evidence for Skype

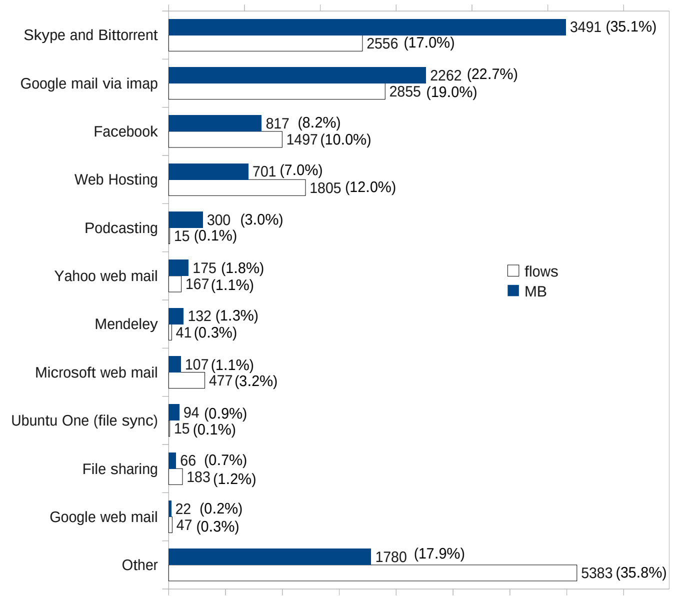

Figure 5.8: Domains visited by outgoing flows greater than 100KB.

dominance is also provided in our interviews in Section 5.7), but further advanced data mining techniques will be required to extract the exact breakdown of these two traffic types. Skype and Bittorrent have the potential to be used to make calls between users in the village or exchange files. However this is far more likely to occur with Skype as Bittorrent is mostly used to download software, music or video that was not sourced by local friends.

The “Google mail via imap” category is for users who connect through an IMAP email client and send file attachments. This traffic type is encrypted using SSL. Analysis of the Facebook traffic category shows primarily image uploads which are usually compressed by the Facebook server. Matching up and down flows for both these forms of traffic, containing encryption and image manipulation, is left as a future exercise. There were a number of web-hosting sites used in Macha, such as softlayer and freewebs; these are grouped together to form one category called “Web hosting”. Although no Youtube uploads were found, users in Macha regularly upload podcasts to Podomatic, which accepts short audio or video clips that users have generated using mobile phones or web cams. Only a small portion of file sharing was found as well as some file synchronization services through Ubuntu One. Mendeley is a program for managing and sharing research papers, and some academic paper uploads occurred. There was a surprisingly small fraction of uploads using web-based email clients; most of the mail attachments occur through non web-based email clients.

Using a set of basic heuristics, it was possible to identify two local Skype calls in the measurement period that were routed between two local users via an international super-peer. Each flow is time-stamped and a simple check was made to see whether a local user connects to an external super-peer, and this same super-peer connects back to a different local user. The second heuristic is to check that the time-stamp between the outgoing and incoming flows is not more than a few seconds. These two skype calls lasted 10 minutes, and 3 minutes respectively, and make up an insignificant proportion of the total number of Skype and Bittorrent traffic seen.

Although it is difficult to conclusively determine the amount of local content sharing in Macha due to much of the traffic being encrypted or manipulated, we conclude that there is most likely only a small portion of overall uplink network traffic that is being used to share large objects between local users. One of the key reasons for this is the high cost of bandwidth, together with the Internet-cap-model being used. This is in stark contrast to the large amount of instant messaging between users. For example, sharing a 500M video costs a user $15. It is far more cost effective to put the video on a flash drive and walk to another local user. However, if the architecture of the network made it possible to exchange these files for free using a local web-based application or through intelligent routing, local file sharing would be sure to increase as it has in developed regions.

5.7 On-line Interviews

To understand whether the demographics and qualitative usage behaviour of the users in the village have any bearing on the behaviour we observed in our trace, we conducted an on-line survey in Macha. This survey was conducted during June and July 2011 and collected data broadly focused on access and usage of Web 2.0 applications and services. The survey was implemented on the SurveyMonkey tool, and was based upon an open source example. It was the first time that an extended questionnaire was offered in Macha with only an online version available. Users were invited to participate via email and Facebook links. Local support to help respondents in interacting with the online survey was provided upon request. 77 users living in Macha participated in the survey consisting of 89 questions.

Some demographic findings from the survey:

- 69% of respondents were between 20 and 30 years old.
- 34% of the respondents were female, and 66% male.
- 88% of respondents were able to use computers and the Internet to achieve most of their objectives. 12% regarded themselves as novice to computers.
- 67% of respondents use the Internet more than 3 hours a day.
- 49% of respondents have Internet connectivity at home.
- 87% of respondents use the Internet at work.
- 71% of respondents use the Internet for learning.
- 51% of respondents use the Internet for entertainment.
- 91% of respondents wish to access the Internet more frequently. 34% are prevented from doing so because of Internet connectivity costs, 33% because of bandwidth limitations, and 25% because of the (institutional) regulations for use.

The survey examined current online activity in 16 categories. Key findings of the survey relevant to social networking are:

- In online engagement, the use of photo sharing stands out, with 60% of the respondents sharing pictures on Facebook and 52% commenting on them, at least several times per month or more;
- 34% of respondents do watch videos on a video sharing website several times per month or more.
- 53% of respondents indicate that they work collaboratively online using tools such as Google Docs.
- 54% of respondents interact on social networks like Facebook several times a week or more. A mere 9% of respondents never interact on social networks, while 24% never used Facebook. 72% use instant messaging.
- All respondents use e-mail. 59% of respondents make phone calls over the Internet several times per month or more, and 59% of respondents search for news online at least several times a week.
- 94% of respondents never used the Internet for an online business. 73% have never used the Internet to purchase goods.

In conclusion, the survey and measurements give indications of an Internet community of 300 users with 200 regular users in Macha. Most people desire more access and interaction, but costs and bandwidth limitations restrain them from doing so. Many of these findings correlate well with what we have observed in our analysis. The 140 local Facebook users active on instant messaging correlates with 72% of the 200 regular users who claim to actively use instant messaging. The well connected interaction graph in Facebook is supported by the fact that only 9% of the community does not use social networking. The dominance of email traffic in our trace is also well supported by the fact that all respondents claim to use email.

The lower proportion of users who have Internet access at home versus users who have access at work will most likely have a negative effective on local user connections as activities like local VoIP and file sharing often occur after work. Internet users in Macha are primarily under 30 and represent a new burgeoning group of trend setters for the rest of the surrounding rural community. As these users, and others that will join, begin to generate and share content and become active members of the global digital village, every effort should be made to mask the crippling effect of the slow Internet gateway. Building a novel new localized network software architecture, that takes full advantage the unique strong clustering revealed in our traffic analysis, is one such mechanism to do this.

5.8 Conclusion

Social networking has grown in Macha and the network in Macha is used extensively for intra-village conversation in the form of instant messaging. For Facebook instant messaging, our statistical analysis revealed that 35% of the users were local to the village, with 7% of these being travellers, and 54% of the instant messages sent were between local users. Analysis of Facebook image sharing, showed that in Macha 24% of images viewed were locally generated and there was approximately 3 times views of a local user's image compared to a non-local user's image.

Interviews also revealed a large proportion of users who make extensive use of social networking with instant messaging being especially popular. Although there is strong social cohesion in this community, no direct local networking traffic related to sharing files or making voice or video calls was found. There was also very little use of Skype between users in the village via Internet super-peers (2 calls within the two month measurement window) or file sharing (0.65% of large outgoing flows). File synchronization services on the Internet were also minimal (0.94% of large outgoing flows) due to the high cost of bandwidth ($32/GB) and bandwidth limitations of the satellite gateway.

It is clear that the current hub and spoke architecture of the web, with most services running as monolithic servers, is not well suited to rural networks with low-bandwidth gateways. However, the strong social cohesion and interdependence often found in rural communities is a key differentiator which, if taken advantage of, can make substantial improvements to network performance when bandwidth-limited Internet gateways are used. A new localized network software architecture is needed to take full advantage of this strong clustering found in rural communities. This software should place services that enable user-to-user interaction and file sharing in the village. Once in place, rural communities will disseminate information amongst each other more efficiently and Internet responsiveness will improve due to offloading "local" traffic from the gateway.

Chapter 6

Gateway-aware Routing: Load Balancing in Rural Networks

6.1 Introduction

IEEE 802.11-based wireless mesh networks (WMNs) have been used as an economic and convenient technology to provide connectivity to rural areas, especially in developing regions[160, 20, 117, 89]. The network in Macha, Zambia has used mesh networking extensively to extend access to community workers and in Chapter 3 we analysed the performance of this network. Even in developed countries, community-wide WMNs have provided Internet connectivity to a large geographic area. We envision the following architecture to use WMNs for Internet access in these areas. Wireless routers are placed in homes to provide in-residence Internet access. A small number of these routers are connected to the Internet and function as gateway nodes for the network. These gateway nodes may use technologies such as DSL, WiMAX, 1xEVDO, cellular, long distance WiFi or satellite as means for capacity injection/backhaul connection. Backhaul connections may be hosted and shared by local businesses or by members of the community.

Traditional mesh routing solutions have focused on simply finding the best route to the gateway to reach the Internet. Two important underlying assumptions of these routing solutions are: 1) all gateway nodes are equally capable in terms of resources such as bandwidth capacity and delay to connect to the Internet; and/or 2) the capacity *bottleneck* is in the wireless multihop portion of the WMN. We now examine each of these assumptions in the context of rural and community WMNs.

In some WMN deployments, all the gateway nodes are similarly provisioned. However, this is not the case in many other scenarios. In rural/developing regions, cost and availability considerations influence the uplink connectivity options that can be used. Another scenario is a neighborhood community WMN wherein volunteers share a portion of their existing Internet connections. These may be DSL, cable-modem, WiMAX or GPRS/EDGE and leads to a heterogeneous mix of gateway uplinks for the WMN. Many Internet connections do not provide guaranteed bandwidth, but instead the capacity varies with network load. Clearly, the assumption of homogeneous gateways does not always hold.

The second assumption in traditional routing is that the capacity of the multihop mesh route is less than that of a gateway uplink. This is more common in a pure IEEE-802.11b network with a raw capacity of 11Mbit/s (and typical multihop capacity of the order of 1Mbit/s [21]). However with the advent of newer 802.11 radio technologies, the capacity of mesh networks has increased by orders of magnitude. (54 Mbit/s for 802.11g or 802.11a and up to 300 Mbit/s for 802.11n). These capacities may be significantly higher than those of the common uplink technologies, e.g. the download capacity for DSL ranges from 1.5Mbit/s to 24Mbit/s, the typical capacity of a WiMAX client is 4Mbit/s [120], a VSAT connection is 1024kbit/s [117]. Therefore, in many scenarios, the capacity of the mesh route may be more than that of the gateway.

The throughput performance of clients connected to a mesh router, intuitively, is influenced by the routing path to the Internet, including the chosen gateway node. Traditional routing solutions, however, have dealt primarily with routing inside the mesh with wireless capacity maximization and interference minimization as prominent objectives. Clearly, in WMNs with varied gateway capacities, the choice of gateway also has a large influence on the performance of the network. Therefore, we believe the correct question to consider in routing is the following:

"at each mesh router, what is the best route to access the Internet?" This problem statement captures the constraints introduced by the gateway capabilities in addition to the problems of wireless capacity and interference.

In this chapter, we propose a gateway-aware mesh routing solution that intelligently selects gateways for each mesh router based on the multihop route in the mesh as well as the capability of the gateway. We develop a new composite routing metric called Gateway-aware Routing Metric (GARM) that captures these aspects of routing in WMNs. We implement our routing metric for the upcoming IEEE-802.11s (mesh networking) standard in a simulator. Further, we evaluate this metric on a testbed network based on an existing routing solution [21]. Our evaluations demonstrate the importance of intelligent gateway selection in order to improve network performance. We show that the GARM metric can increase the overall network throughput, in some cases by 22%, through better utilization of existing resources.

6.2 Related Work

There is a significant body of research work that focuses on improving the performance of WMNs. These have addressed various aspects of WMN operations including routing, channel assignment, and interference management. However, most of these solutions either consider the network to have a single gateway, or multiple homogeneous gateways.

Multi-gateway systems have been studied in the context of load balancing among the gateways. The Hyacinth architecture considers multiple gateways and suggests gateway selection based on available bandwidth, to achieve load balancing [143]. This architecture requires accurate knowledge of the flow bandwidth requirements at each mesh router, which may be non-trivial to estimate. Nandiraju et al. propose dynamic gateway switching in a WMN to achieve load balancing: the gateway node monitors congestion levels on its uplink, and at the onset of congestion sends a message to an associated mesh router to choose a different gateway [125]. In contrast, our approach is proactive and does not require frequent congestion measurement. Lakshmanan et al. suggest the simultaneous use of multiple gateways by striping packets of a flow to maximize uplink capacity utilization [106]. Their architecture, however, requires a super-gateway that handles packet re-ordering caused by packet striping. Such an architecture may not be possible in rural WMNs because the uplinks may be provided by different ISPs. Other systems, as proposed in [161], formulate the problem as a linear optimization problem. This approach, however, requires global knowledge of gateway and link capacities, and is centralized. On the other hand, our system is distributed and builds upon the existing routing framework.

6.3 Background

We now outline our assumptions and considerations for the design of a gateway-aware routing protocol in WMNs.

Assumptions: We assume the mesh network consists of a multi-tiered architecture with a mesh backhaul layer responsible for communication among the mesh routers (including the gateway node), and a client access layer that communicates with end-user devices. This architecture is commonly used by several real-world deployments [20, 21]. Our protocol is currently designed for a mesh backhaul layer that uses one 802.11 radio at each node. We assume that only mesh nodes use our routing protocol and select one gateway at a time, i.e., multi-gateway associations are not permitted.

Design Preliminaries: We consider the problem of finding the best route to the Internet from any node, given a choice of multiple gateway nodes and several candidate multi-hop routes to reach these gateways. The gateways are characterized by different capacities, and the candidate routes by different path qualities. A naïve approach would be to include an additional virtual node connected to each gateway with a link capacity corresponding to each gateway's capacity. The problem is now reduced to finding the best route from each mesh node to this

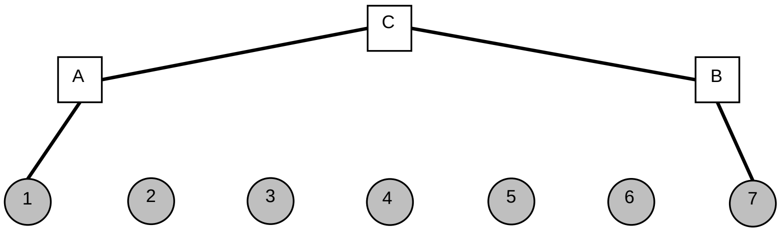

Figure 6.1: Simple Line Topology: The gray circles are mesh nodes. Gateway nodes 1 and 7 are connected to routers A and B via Ethernet-like connections. We measure the throughput to node 5.

virtual node. The uplink is mapped as an additional wireless link and existing routing algorithms are used.

We conduct simulation experiments using Qualnet [1] to understand the feasibility and/or pitfalls of this simple approach. In particular, we study the effects of gateway capacity on the throughput performance of the mesh network.

We consider the line topology of Figure 6.1 with seven equidistant mesh nodes. This topology also approximates the decision process in a more complex complete graph where at some point two paths extend to either gateway. In each experiment trial, the distance between the nodes is varied to ensure that only adjacent nodes can communicate directly. All the mesh nodes use the IEEE 802.11s mesh networking extension with HWMP routing protocol to form the mesh topology. The radios operate using the 802.11b/g MAC and PHY standard and bit-rate is fixed uniformly across the network. All other parameters are

[1]Qualnet Network Simulator, Version 4.5. Available: http://www.scalable-networks.com

set to the default values provided by the simulator. For each experiment trial, we use a different bit-rate in order to vary the wireless mesh capacity. Nodes 1 and 7 are the gateway nodes and are connected to the external nodes A and B via wired-links. The capacities of these links are varied to simulate different capacities. Nodes A and B are connected to node C with a high capacity link. Node C forms one end-point of all communication and represents the larger Internet.

In order to study the impact of gateway choice, only one of the uplinks (Link 1A or 7B) is active at any time. We initiate one UDP flow from node C to node 5 and measure the throughput for this flow through each gateway. Table 6.1 shows the average of five trials of these experiments.

Gateway Node	Gateway Capacity (Mbps)	Wireless bit-rate (Mbps)	UDP throughput (Mbps)
1	1.5	36	1.443
7	0.5	36	0.493
1	1.5	2	0.367
7	0.5	2	0.458

Table 6.1: Throughput with two gateways in a line topology.

We make two key observations from these results. First, using a gateway with a high capacity uplink can lead to higher throughput, even though it may take more hops (or mesh air-time) to communicate with the gateway. Therefore, an intelligent choice of gateway can indeed lead to better network performance.

Second, using a lower wireless bit-rate, the throughput while using the closer gateway with lower uplink capacity is higher than that using a high capacity distant gateway. We conclude that the total end-end throughput is dominated by one of either the gateway capacity or the capacity of the mesh path, i.e., one of these acts as a bottleneck for the path and effectively determines the maximum achievable throughput.

The above observations, imply that the protocol should treat the two parts of the path (i.e., multihop wireless path and the gateway uplink) independently. The metric should be able to identify the bottleneck portion of each path, and choose an appropriate gateway. The challenge, therefore, is to design a routing metric that simultaneously considers the gateway capacity as well as the quality of the multihop route from a given node to the gateway. The design of such a metric necessitates a mapping of gateway capacity to the path quality, in order to compare and identify the bottleneck. We next present the design of the GARM metric that enables joint selection of gateways and routes to these gateways.

6.4 Design

Based on the observations in Section 6.3, we can consider a gateway-aware route in a WMN to consist of two parts: first, the multihop wireless path; second, the uplink at the gateway node. We present the metric design for these components, and then describe the unified metric.

Metric for the wireless path: Yang et al. identify four fundamental characteristics that a WMN routing metric should capture [174]. The routing metric should increase with the length of path, capture packet loss ratios of the links, consider link capacities, and help reduce interference in the channel shared by the wireless links. Early research on routing for multihop wireless networks commonly used hop count as the routing metric. This metric incorporates only the first characteristic. The Expected Transmission Count (ETX) [46], defined as the expected number of MAC layer transmissions to deliver a packet, only captures the first two characteristics.

The Expected Transmission Time (ETT) metric improves upon ETX by considering the differences in link capacities [50]. The ETT of a link i is defined as the expected duration of a successful MAC layer transmission over the link i. The ETT of a link l is defined by the following relationship:

$$ETT_i = ETX_i \cdot \frac{S}{B_i}$$

where S is the packet size and B_i is the bit-rate for link i. The weight of a path p is the sum of the individual ETTs of the links along the path:

$$mETT = \sum_{l \in p} ETT_l$$

The ETT metric captures the impact of the link capacities on the path performance. For single radio networks, the ETT metric also considers the impact of interference. Therefore, ETT has all the desirable properties for the routing metric of WMN. Routing using the ETT metric has been shown to provide high throughput paths, compared to other previous metrics such as ETX [21, 50]. We therefore choose to use ETT as the metric for the wireless portion of the gateway-aware metric. We also note that the Airtime metric, one of the proposed routing candidate metrics in the IEEE 802.11s standard, is similar to the ETT metric.

Gateway capacity metric: We define the gateway capacity metric $gwETT$ as the time required to transmit a packet of size S on the uplink, and is given by

$$gwETT = ETX_{gw} \cdot \frac{S}{B_{gw}}$$

where B_{gw} is the capacity of the gateway and ETX_{gw} is the expected transmission count for the uplink. In the simple case, we assume that the uplink is a reliable medium or has negligible loss rates, and ETX_{gw} is one. If the uplink is an unreliable medium, we assume that the ETX_{gw} is provided to the routing protocol by some external module (e.g., a link quality measurement tool). Such an ETT-like design of the gateway capacity metric enables a simple mapping and direct comparison to the wireless path ETT-metric.

Combining the two metrics: In Section 6.3, we noted that the throughput of a path is determined by the bottleneck portion of the path. A comparison of the two ETT metrics, $mETT$ and $gwETT$, identifies this bottleneck; the larger of the two indicates that more time is required for transmission on the corresponding portion, and therefore is the bottleneck on the path.

$$M_i = max(mETT, gwETT)$$

This metric, however, is not *isotonic* [174]. Isotonicity is an important property for routing metrics to ensure that efficient algorithms such as Dijkstra can be used to calculate paths with minimum weight. In addition, isotonicity also ensures loop-free forwarding in hop-by-hop routing.

Therefore, we use the following metric for routing to a gateway:

$$GARM = \beta \cdot M_i + (1 - \beta) \cdot (mETT + gwETT)$$

This Gateway-aware Routing Metric has two parts. The first part of the metric accounts for bottleneck capacity. The second factor models the total delay of the path, including the uplink. β is used to control the balance between these two factors. The gateway with the least GARM value is chosen as the default gateway. We note that $mETT$ increases at each hop, and since $gwETT$ is constant for a given gateway, M_i either stays the same or increases at each hop. Therefore $GARM$ is isotonic. This metric requires the propagation of the $gwETT$ metric in addition to the path metric, i.e., each node needs to know of the gateway capacity metric for its routing calculations.

6.4.1 Design Discussion

Capacity measurement: One of the inputs to the GARM metric is the gateway capacity. We assume that this value is input to the routing software by an external entity. This capacity input is straight-forward on uplinks with guaranteed/fixed bandwidth, e.g., Ethernet, in a community WMN wherein a volunteer wants to share only a fixed portion of his Internet connection. On uplinks with time-varying capacity, (e.g. wireless uplinks), we assume the presence of capacity-measurement tools such as *pathrate*. In response to a change in uplink capacity, gateway selection may change at the same time as required to propagate the routing information.

Load-sensitivity: GARM does not consider the traffic load on the gateway for gateway selection. Load-based routing leads to routing instability and route flapping [174]. A similar effect would occur for gateway selection as well. However, in networks where the mesh paths are of similar qualities, GARM achieves implicit load-balancing among gateways. Due to the formulation of the $gwETT$ metric and its comparison with $mETT$, GARM associates mesh routers to gateways approximately in proportion to the gateway capacities.

Multi-radio WMNs: The GARM metric is designed for operation on single radio/channel networks. As noted before, operation of the gateway uplink is independent of the wireless path. This behavior can be considered analogous to a wireless link operating on a different channel. Previous research has designed routing protocols such as MR-LQSR with the WCETT metric [50]. We believe such solutions can be extended to be gateway-aware by considering the uplink as a wireless link operating on an orthogonal channel.

6.5 Implementation and Performance Evaluation

We implement the GARM routing metric in two environments. First, we extend the 802.11s-HWMP implementation in Qualnet 4.5 to include the composite GARM metric. The 802.11s protocol provides a framework for WMN operations including neighbor discovery, link quality estimation, route discovery and route updates. The 802.11s routing beacons that originate from gateway/portal nodes include additional information about the $gwETT$ value corresponding to its uplink capacity. Intermediate routers repeat the $gwETT$ value when they forward the routing beacons, thereby indicating the $gwETT$ value of the path in which the update propagated. Our second implementation of the GARM metric is for the Click-based MIT RoofNet software [21]. Here too, the gateway routing advertisements of the SRCR routing protocol are extended to include the $gwETT$ value, and intermediate routers rebroadcast this information. We use this implementation for our experiments on the UCSB MeshNet testbed (see Sections 6.5.2 and 6.5.3).

To evaluate the performance of the GARM metric in WMN routing, we first study the accuracy of the metric in selecting gateways. Subsequently, we characterize the throughput gains obtained by intelligent gateway selection. We compare the performance of our solution with that of routing protocols that use the ETT metric. We evaluate GARM in simulation-based experiments as well on a testbed network. The simulation study allows fine-grained control of network parameters such as topology, link quality, bit-rate, etc. to enable better understanding of the metric. The testbed represents an uncontrolled environment with variable and heterogeneous link qualities, dynamic bit-rate, external interference, etc. and provides a realistic WMN environment.

In the evaluation of GARM we use different combinations of gateway uplink capacities. We choose three gateway capacities that represent the uplink technologies that are used in rural WMNs: 4Mbit/s to represent ADSL, long-distance Wifi, WipLL (Wireless IP Local Loop), WiMAX [160, 120, 2]; 1.5Mbit/s for CDMA450, T1, 1xEVDO [120, 2]; 0.5Mbit/s for a VSAT link or EDGE [117, 2]. The absolute values of the uplink capacities can vary due to the level of service purchased. For example service providers will often provide a guaranteed rate and a maximum burst rate which varies based on the number of users. A possible solution is to advertise a typical mean for a specific day and hour using historic measurements. Calculating actual instantaneous capacity is a non-trivial problem and for now we assume that the gateway capacity is constant for the duration of the experiments. However, it is important to note that there is an order of magnitude difference in their relative capacities. As we demonstrate later, the larger the difference in gateway capacities, the more critical is the role of gateway-aware routing in increasing network performance.

6.5.1 Gateway Selection Accuracy

We first use simulation-based experiments to study GARM's accuracy in selecting gateways under different operating environments. Therefore, we conduct experiments with various combinations of uplink and wireless path capacities.

Our experiments use twenty-five 802.11b/g nodes in a 5x5 uniform grid topology. For each experiment, we fix the link bit-rate to be either 2Mbit/s, 11Mbit/s, or 36Mbit/s; all nodes use the same bit-rate. The distance between the nodes is varied accordingly to ensure that only 1-hop neighbors have direct communication. Two diagonally opposite corners of the grid are chosen as gateway nodes. Similar to the topology shown in Figure 6.1, the gateway nodes are connected to external nodes A and B via wired-links. The capacities of these links are varied. Nodes A and B are connected to node C with a high capacity link. Node C is the end-point of all communication and represents the Internet. The two gateway capacities are chosen as a combination of 4Mbit/s, 1.5Mbit/s and 0.5Mbit/s, leading to three distinct uplink capacity pairs. These uplink capacity pairs, together with the link bit-rates, lead to nine different combinations of gateway and wireless path capacities. Through simulations, we verify that GARM selects the appropriate gateway in each scenario.

For each scenario, we perform five trials of the following experiments. We first determine the best gateway for each wireless node in the network. This is achieved by comparing the throughput of a TCP flow from node C to each wireless node, and evaluating the throughput through each gateway. We call this method the Oracle approach because it determines the maximum achievable throughput. Next, we repeat the TCP throughput experiment and let the routing protocol choose the gateway. We record the throughput as well as the choice of gateway. The experiment is repeated for ETT and GARM with β values of 0.0, 0.5 and 0.9. Note that $\beta = 0$ represents the case where the metric does not identify the bottleneck, instead it only considers the sum of $mETT$ and $gwETT$. For every node, we determine whether the routing protocol chooses the best gateway, and calculate the difference in the achieved throughput from that of the Oracle approach.

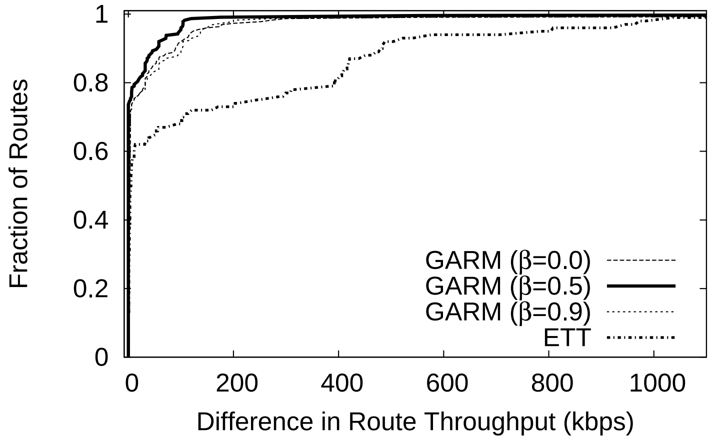

Figure 6.2: Accuracy of gateway selection: CDF of difference in throughput using the GARM and ETT routing metrics compared to that using the Oracle approach (best gateway). A perfect routing metric has zero difference.

Figure 6.2 plots a CDF of the difference in throughput for each route using the GARM and ETT metrics compared to the Oracle approach over all nine capacity combinations. A perfect routing metric would always choose the best gateway (similar to the Oracle approach), and therefore have zero difference in throughput for all routes. From the graph, we see that ETT selects the best gateway for only 60% of the nodes. Further, for 20% of the routes, the loss of throughput due to poor gateway choice is significant (more than 400kbit/s). GARM (with β=0.5), on the other hand, has a very small fraction of routes (<3%) wherein the loss of throughput is greater than 150kbit/s. With GARM, 76% of the nodes select the best gateway. For the nodes that do not select the best gateway, we see that the difference in throughput is less than 150kbit/s. This shows that the GARM metric is unable to distinguish between routes that differ by less than 150kbit/s.

The graph shows that the performance of GARM with β values of 0, 0.5, and 0.9 is quite similar. This leads us to believe that GARM is not very sensitive to the value of β. However, we recommend the use of $\beta = 0.5$, where bottleneck capacity and path delay including the uplink are perfectly balanced, since it provides the best results among the three values we tested.

6.5.2 Throughput Performance

We next evaluate GARM on the UCSB MeshNet indoor testbed [109]. The testbed consists of 15 nodes deployed in offices on different floors of a building. Each node consists of two Atheros-based 802.11a/b/g radios, and uses the MadWifi driver (v0.9.4) on Linux (kernel v.2.6-15). We use only one of the radios at each node and operate it in 802.11b/g mode. Each node is also connected to a LAN via Ethernet. This provides a control interface to manage the nodes and conduct experiments. This setup also allows us to provision any testbed node to be a gateway node with the Ethernet link as uplink, thereby enabling evaluation of different topologies. We use the MIT RoofNet software for routing in this network. The SRCR routing protocol [21], implemented in Click, uses the

ETT metric for routing. Our implementation incorporates GARM in the SRCR protocol. For our experiments, we use a dedicated host machine on the LAN to act as an endpoint for TCP tests and a Web server. In the context of an actual WMN, this dedicated host represents the Internet.

We now characterize the throughput gains achieved at each mesh node with intelligent gateway selection using the GARM metric. We use two gateways in the network, each with different capacity. The capacity of the gateways is again chosen among 4Mbit/s, 1.5Mbit/s and 0.5Mbit/s. We use the `tc` traffic control utility to limit the gateway capacities to these values. The two gateways are randomly chosen from nodes in the network. The selection of two random gateways with a given capacity combination (e.g., 4Mbit/s and 1.5Mbit/s) constitutes a network topology for the experiment. With each topology, we perform the following. For each node in the network, we use the `nuttcp` throughput measurement tool to measure the throughput of a 3-minute TCP stream from the dedicated LAN host to the node. This represents this maximum throughput at the node when downloading a file from the Internet. Only one stream is active at a time. We record the average throughput of each node based on three trials. We repeat the experiment for five different topologies, for both the GARM and ETT routing metrics.

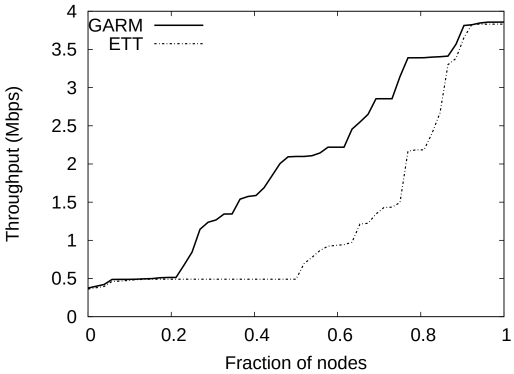

(a) Gateway Capacities: 4Mbit/s, 0.5Mbit/s

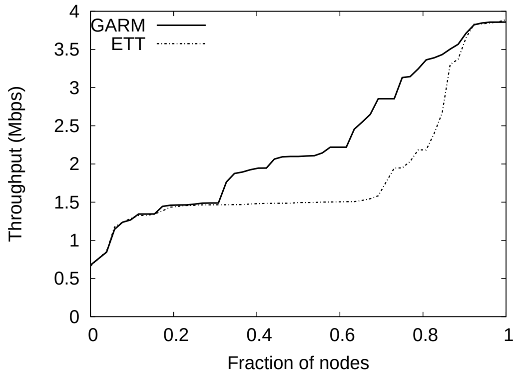

(b) Gateway Capacities: 4Mbit/s, 1.5Mbit/s

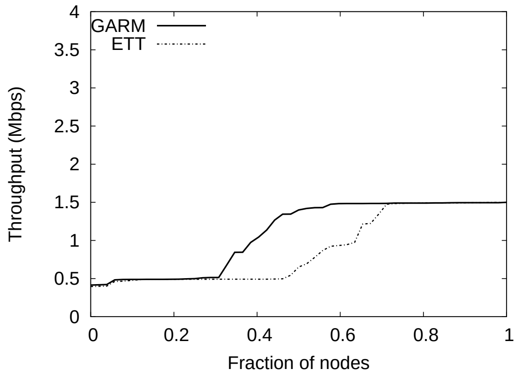

(c) Gateway Capacities: 1.5Mbit/s, 0.5Mbit/s

Figure 6.3: Distribution of node TCP throughput with different gateway capacities. GARM increases throughput for up to 67% of the nodes. The performance increases as the the difference in gateway capacities grows.

We plot the distribution (CDF) of individual node throughputs for these testbed experiments in Figure 6.3. Figures 6.3a, (b) and (c) show the distribution of throughputs when the capacities of the gateway is 4Mbit/s & 0.5Mbit/s, 4Mbit/s & 1.5Mbit/s, and 1.5Mbit/s & 0.5Mbit/s, respectively. The percentage of nodes that obtain better throughput with GARM are 67%, 53% and 36%, respectively. The average increase in throughput (and percentage increase in throughput) for these scenarios is 679kbit/s (58%), 403kbit/s (24%), and 194kbit/s (22%) respectively. Although GARM provides better throughput for many nodes in all three scenarios, we note that the absolute gain and the number of affected nodes is influenced by the difference in the gateway capacities. The larger the difference, the larger the fraction of nodes affected and greater the throughput improvement.

6.5.3 Overall Network Performance

Having evaluated the performance gains of individual nodes in the network, we now characterize the overall network throughput performance. Each node in the network simultaneously downloads a five megabyte file from the dedicated LAN host. At each node, the time required for completion of the download is measured. The experiment is repeated with different gateway capacity combinations and for five different topologies.

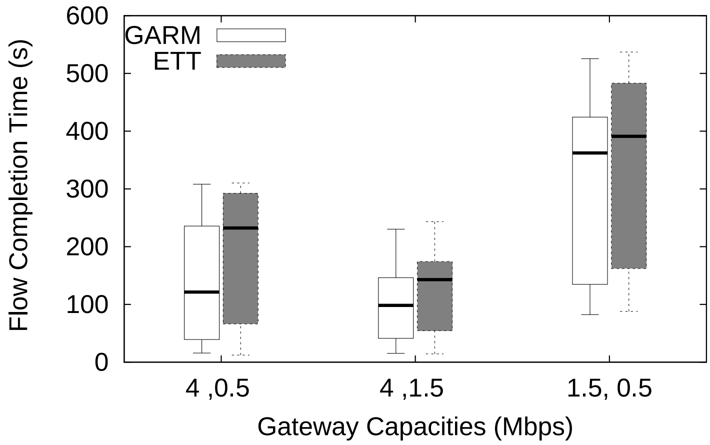

Figure 6.4: Distribution of time required to download a 5 MByte file in the UCSB MeshNet.

Figure 6.4 shows a box-and-whiskers plot depicting the quartiles of the distribution of flow completion time for the GARM and ETT metrics[2]. The median values are shown in the box and the whiskers represent the minimum and maximum values. From the graph, we observe that GARM reduces the median download time in each case. Further, the distribution of the lower and the median quartiles is skewed towards lower flow completion times, implying that a large number of flows complete sooner when using GARM than when using ETT. This effect is most significant when the difference between the gateway capacities is large. For example, with gateway capacities of 4Mbit/s and 0.5Mbit/s, the average completion time for GARM and ETT are 146.9s and 179.3s, respectively.

[2]We choose the box-and-whiskers representation instead of average and standard deviation because the distribution is approximately bimodal; this is because the flow completion time in many cases is limited by the capacity of one of the two gateways.

This represents a 22% increase in average network throughput, through intelligent gateway selection that results in better utilization of network resources.

6.6 Conclusion

Wireless mesh networks in developing regions are connected to the Internet with a variety of uplink technologies, each with different capacities. In this paper, we have shown that the selection of gateways for routing plays a very important role in determining the performance of the network. We presented the design of the GARM routing metric that can effectively choose the best gateway for each mesh node. Through performance evaluations we showed that intelligent gateway selection can increase the throughput of nodes in the WMN, and increase the effective capacity of the network. This framework could also be extended to incorporate aspects such as the different costs of gateway bandwidth, power availability at each gateway and application-aware gateway selection.

Chapter 7

VillageShare: Encouraging Content Sharing

7.1 Introduction

It is clear that online social networks (OSNs) have revolutionized communication worldwide. The statistics are astounding: As of 2012, 1.2 billion people use Facebook.[1] Twitter adds about 500,000 new users a day.[2] 800 million unique visits to YouTube occur every month.[3] Chapter 3 and Chapter 5 demonstrated that even residents of the far reaches of rural Africa use Facebook frequently to

[1] http://www.statisticbrain.com/social-networking-statistics
[2] http://www.jeffbullas.com/2011/09/02/20-stunning-social-media-statistics
[3] http://www.statisticbrain.com/youtube-statistics

communicate with others in their community as well as with friends, family, and colleagues worldwide. In these areas, online social networks (OSNs) are critical as alternate communication infrastructures (e.g. fixed telephone lines, cellular service) are infrequently available.

However, what is also clear is that access to OSNs is not uniformly available to the world's citizens. In developing regions, Internet access is often slow and unreliable; it can be orders of magnitude slower than in developed countries, as residents access the Internet through low bandwidth satellite or other long distance wireless links. Websites have also become more dynamic, and as a result, less cachable. This is particularly true of social media sites, where content can change at sub-minute granularity.

These challenges have resulted in a social media digital divide, where Internet users in developing regions with poor Internet access have difficulty actively participating in and contributing to social networks. Chapter 4 showed that it is common for 75% of uploads to fail in rural Zambia, making it difficult for Zambians to become producers of Internet content. User-generated content, such as Wikipedia articles, is virtually non-existent in local African languages[4]. Thus, on the web, rural Africans have to revert to a major language, e.g. English.

[4]For example, there are no Wikipedia articles in Chi-Tonga, Ila and Lozi languages spoken by more than two million people in (rural) Zambia, whereas approximately the same number of Slovenians enjoy a bounty of more than 100,000 articles in their native language.

Despite these challenges, however, there is a clear desire for participation. Our Macha traffic analysis has shown that Facebook is the most popular URL request, accounting for 15% of all web requests during our two month measurement window in 2011; Twitter was the second most popular URL request, accounting for 11.5% of all requests. As of 2012, Facebook users come from 213 countries worldwide[5] while tweets have appeared from 229 countries[6].

Chapter 5 studied the spatial locality of Facebook instant message (IM) and picture exchanges in Macha and discovered a number of interesting facts. Most importantly, it discovered that 54% of IMs were between local users – users within a single village. This is despite the fact that only 35% of the observed users were local to the village. In addition we found that 25% of all shared Facebook photos were generated by local users in the village. Other studies have shown similar findings, that as the Internet continues to grow, it is becoming "more local" [171, 136].

In this Chapter, we endeavour to take advantage of the spatial locality of social network traffic to provide a better OSN experience to users in developing regions. We also take advantage of off-peak periods in the network to upload content more reliably. Specifically, we develop and evaluate two novel architectures to optimize local sharing, while simultaneously facilitating streamlined content

[5] www.socialbakers.com/facebook-statistics/

[6] http://aworldoftweets.frogdesign.com

sharing with external users and content servers. The first solution makes used of a content-sharing Facebook application combined with a time-delayed proxy server and the second solution utilizes a local cloud-based application combined with traffic-shaping to facilitate file sharing and file uploads. VillageShare is called Kwaabana in Zambia; a local Tsonga meaning 'to share'.

Our Facebook-based VillageShare application, makes use of the existing Facebook social graph to facilitate file sharing between Facebook friends. It ensures that all content shared between two local Facebook friends is stored in the village and content shared between a Facebook user within the Village and outside the village only uses the Internet link when it is underutilized. We evaluate the system theoretically using existing traffic patterns to predict the amount of bandwidth that could be saved as well as the impact of time-shifting flows. Our findings show that between 11% and 22% of the large outbound flow traffic, normally utilizing the satellite gateway, can be sent between local users using a file server hosted locally by VillageShare. The time-delay proxy creates enough off-peak capacity to fulfill the current traffic load with enough capacity remaining to accommodate additional large upload requests.

Our cloud-based VillageShare application utilizes a local cloud-storage solution, providing file hosting and sharing within the village. A global VillageShare cloud storage-server is put in place to allow village users to share content with non-local users. Databases between these two servers are partially synchronized based on which content is shared between local and international users. A sophisticated synchronization engine is put in place to minimize disruption to interactive traffic, i.e. HTTP traffic. We have deployed our cloud-based VillageShare in Macha, Zambia and found that users utilize the service both as a means of locally backing up images from their mobile phones as well as for sharing images, documents, music and videos. VillageShare has also been used to share large videos. Video sharing is a slow and expensive exercise when streamed from an international file hosting service in Chapter 4 we observed that the majority of video downloads failed. Due to our smart synchronization strategy, synchronization of files between the local and global VillageShare servers shows little effect on web browsing performance in Macha and periods of low Internet gateway utilization are exploited for synchronization.

We believe that VillageShare will be key to enabling vibrant collaboration amongst users in remote communities, as well as sharing of content generated by residents of remote regions with the international community. We hope that, as a result, users who previously abandoned the idea of uploading media due to frustration will be encouraged to generate and share digital content, becoming active, full-fledged participants in social networking and content generation and distribution.

7.2 Methodology

Details about the Macha Network and background about the location are given in Section 2.2.1. Details on methods used to capture traffic traces are given in Section 2.3.1.

Content generation and sharing in rural areas can be influenced by a complex mix of technological and social factors [113]. To investigate opportunities for improved content generation and sharing, we perform a holistic investigation of Internet usage in rural Macha, Zambia. Our study incorporates both network performance profiling and traffic analysis as well as social surveys.

The method used for our first social survey in 2010 on the use of Web 2.0 applications in Macha, Zambia is described in Section 2.3.2. We carried out a 2nd survey on current sharing behaviour in 2012 in Macha as well as a rural connectivity project – Broadband 4 All[7] in South Africa. We included the South African project in order to understand if any trends can be generalized. The survey asked questions about the type of media community members produce, such as audio, images or video. It then requested the respondents to report on their methods for sharing this media with their local or non-local contacts and obstacles they experienced sharing content. rural connectivity project – Broadband 4 All[8]

[7] http://www.broadband4all.co.za
[8] http://www.broadband4all.co.za

7.3 Survey results

7.3.1 Content Generation in Rural Africa

Online presence has great potential to facilitate the preservation of intangible cultural heritage within the local community and its sharing with the wider society. If only access to web-browsing is available, but content generation is not supported, local customs and way of life threaten to be replaced by a lifestyle that is observed online, which mostly originates from urban developed areas [164]. In addition, if the content coming from a rural community is created only by a subset of users, this micro digital divide can even further polarize the community. We contrast content generation with the demographic of Internet users in Macha we surveyed. We find that neither gender, age, nor reported level of IT competence impacts content generation. This promising result emphasizes the importance of facilitating content generation as a tool to tackle the existing inequalities in the developing world [75].

The results of our social survey show that content generation is strongly associated with user-reported accessibility of Internet connection (Kendall's tau-b $(N = 41) = .249, p = .071$). This is well aligned with a previous study [44] that identifies perceived ease of use as one of the key factors for user acceptance of IT. Limited bandwidth is the main hurdle for full fledged Internet access in rural

Zambia - 41.5% of the interviewees mentioned it as one of the reasons why they do not spend more time online. Limited bandwidth was a strong differentiator between those who do and those who do not generate content in the US during the transition period from dial-up to broadband connections [77]. Thus, to stimulate content generation and sharing in rural areas, it is crucial to provide support for bandwidth-hungry applications.

7.3.2 Content sharing behaviour in Rural Africa

A total of 14 users responded to the online survey in Macha. The participants' age ranged from 18 to 35; 9 of them were female and 6 were male. The majority of users made use of mobile phones to produce content: 62% used phones to capture images, 60% used phones to capture video, and 40% for recording audio. For each type of media, we asked whether the user shared this content with others and what techniques were utilized for sharing. The results are shown in Figure 7.1. The overwhelming majority made use of external media, such as CDs and USB flash drives, to share content with others. All media was shared online to some extent except for video; video was shared exclusively via removable media. We probed this behaviour further by asking about obstacles to sharing content. 87% of the respondents selected cost (i.e. of Internet access due to per-byte usage fees) as a significant obstacle; 75% selected speed; and 50% selected reliability of the Internet as impediments to sharing.

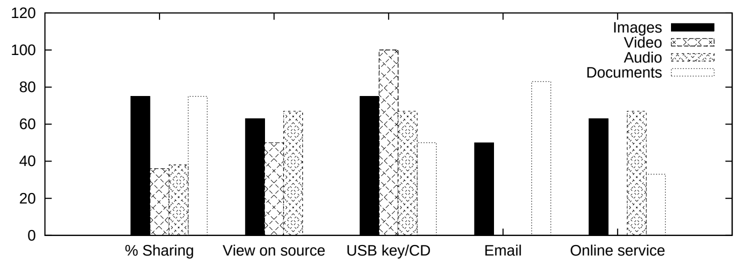

Figure 7.1: Results from online survey on current file sharing trends in Macha.

A total of 15 village operators completed the survey in South Africa. The participants' age ranged from 22 to 39 with most respondents being under 30; 11 of them were male and 4 were female. Again many users made use of a mobile phone to produce content, but higher quality recording devices featured just as strongly; 77% used mobile phones to capture images (70% used digital cameras), 62% used mobile phones to capture video (85% used video cameras, 39% used digital cameras), 72% used mobile phones for recording audio (72% used computer recoding equipment).

The results for sharing content are shown in Figure 7.2. Again many users made use of external media, such as CDs and USB flash drives, to share content with others. Many more users were viewing content with others directly on the device (approximately 80% compared to approximately 60% in Macha); this is most likely due to more users having higher quality recording devices with better screens for viewing. The fraction of video shared was higher than Macha (60% compared to 35% in Macha) with some early signs of attempting to share videos online (15% using email and 15% using online services); no video was shared online in Macha. Village Operators do not pay for their Internet access and have no bandwidth cap. They do however, have controlled access to the Internet via a white list of approved web sites. We surmise that the penchant towards more online sharing is due to the zero-cost barrier to using web sites that facilitate sharing.

In response to obstacles to sharing content with others, the most significant factor was moving the content from the device to a computer (50% of respondents). This is a perplexing results and we assume that this is due to non-usb mobile phones which need specialized cables or computers with Bluetooth to offload the content. The second most significant factor was reliability of the Internet (42%). Cost and speed were only an obstacle to 17% of respondents – an expected result considering the service is provided free of charge. The obstacles are entirely different to Macha due to the different business model in place but still reflect a unique set of issues that rural users face when attempting to share their content with other users.

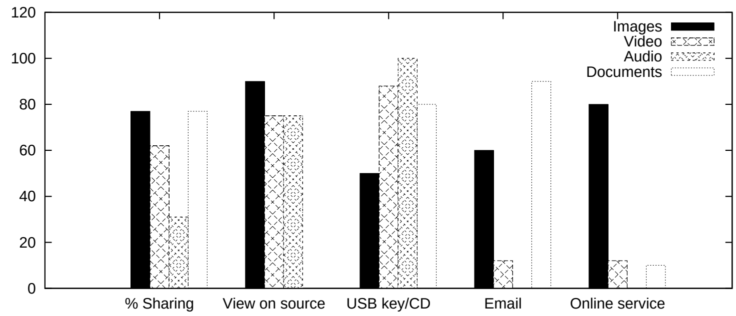

Figure 7.2: Results from online survey on current file sharing trends in the BB4all project, South Africa.

Respondents were then asked how their content sharing behaviour would change if there was a specialized service to address this need. On a graded scale with four options from 'very unlikely' to 'very likely' they were asked if they would share more content with local users if there was a free service. In Macha, 86% responded with 'very likely' and 14% responded with 'likely'. In South Africa, 58% responded with 'very likely' and 42% responded with 'likely'. We then asked whether they would share more content if the service was 10 times faster than an Internet service (an estimate of the increase in upload/download speed if a local service was used). In Macha 71% responded with 'very likely' and 29% responded with 'likely'. In South Africa 58% responded with 'very likely' and 42% responded with 'likely'

From this survey it is clear that VillageShare has great potential to increase content sharing and online collaboration within the community. The motivation in Macha is mostly related to saving on costs whereas the motivation in South Africa appears to be related to reliability of the link to the Internet. Video sharing clearly has the most to gain from a local file sharing service. The problems faced when offloading content from mobile phones without WiFi could be solved by placing a GSM-pico cell or VillageCell, described in Chapter 9, in the community to offload digital content directly to VillageShare at no cost.

7.4 Traffic analysis

A user that is limited to content consumption does not experience a full level of utility from the Internet. In rural areas, the lack of bandwidth prevents users from generating and sharing content. The essence of the problem lies in both technology that is used for access (i.e. long distance terrestrial wireless or satellite links), as well as the way online applications are designed (centralized architecture). In this section we investigate specifics of rural area network traffic and identify opportunities for improvement of connection utilization. As described in Section 2.2.1, Macha is connected to the Internet through a satellite. Hence, in the remainder of this paper we focus on a satellite gateway link. However, our work is generalizable to any low bandwidth gateway link.

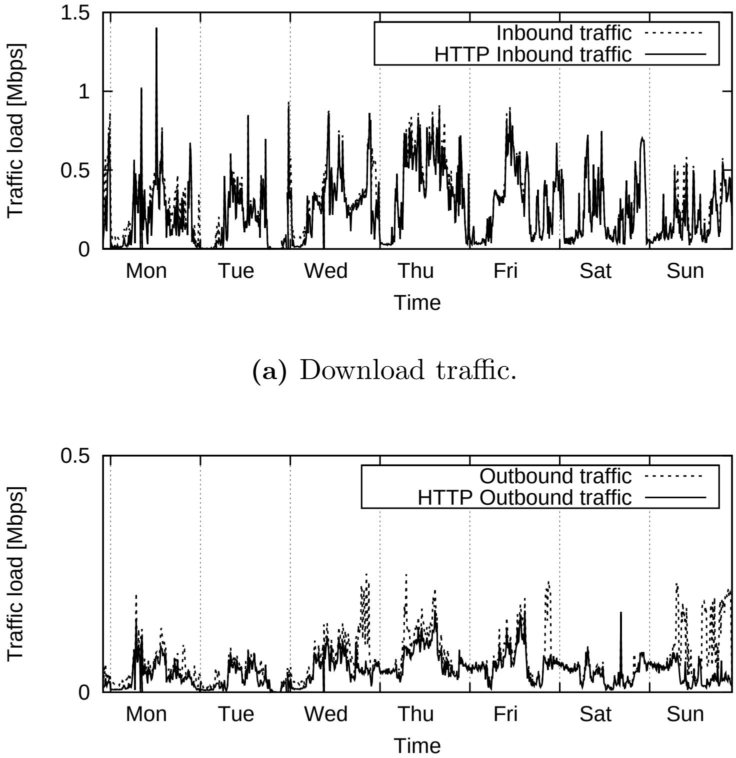

(a) Download traffic.

(b) Upload traffic.

Figure 7.3: Network usage over a week (the x-axis ticks mark starts of days). There is a clear diurnal pattern.

7.4.1 Temporal characteristics of network traffic

A diurnal pattern of web browsing has been reported in both developed and developing world networks [166, 51, 74]. Daily periodicity, however, is much more important in bandwidth limited networks where the capacity can easily be reached during peak usage times [166]. In this work we examine not only incoming, but also outgoing network traffic, since our goal is improved content generation and sharing.

Network traffic distribution in time has been explored in the context of HTTP traffic, through web proxy log analysis [166, 51]. Here we analyze all TCP flows in the full network trace. Thus, the resulting graphs (shown for a week long period in Figure 7.3a for inbound and in Figure 7.3b for outbound traffic) present, besides HTTP, other protocols that utilize TCP. Moreover, this method of analysis allows us to capture those HTTP uploads that have not been completed successfully[9]. Figures 7.3a and 7.3b show a typical daily pattern. There is little activity late in the evenings and early in the mornings. While HTTP traffic remains dominant, in the upload there are periods when non-HTTP traffic prevails, such as Wednesday, Friday and Sunday evening in the plotted sample.

A closer analysis reveals that the majority of non-HTTP traffic is peer-to-peer (P2P) traffic; in total 8.2% of all traffic is P2P. While it is less pronounced than in the developed world [151], P2P can be highly problematic in rural area networks. P2P networking is a bandwidth costly technique for sharing content, one which networks using an asymmetric satellite gateway connection can ill afford.

We investigate the difference between total Internet usage in Macha during weekdays and weekends. We take the two-month trace and average observed

[9] HTTP POST messages are sent out after the actual content is successfully transferred.

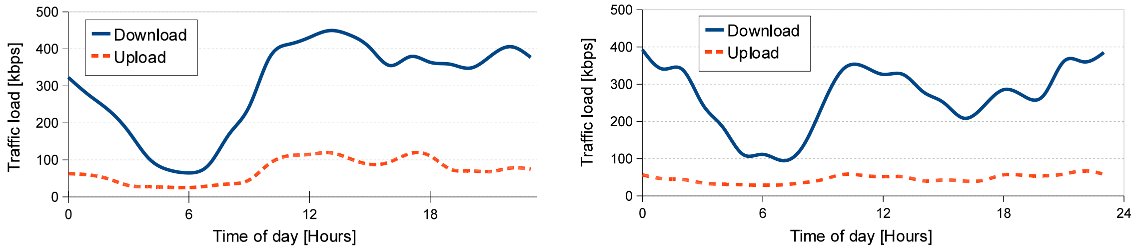

(a) Weekdays **(b)** Weekends

Figure 7.4: Network usage depending on the time of week. Traffic load differs between weekends and weekdays.

upload and download traffic over twenty four hours, separately for weekdays (Figure 7.4a) and weekends (Figure 7.4b). The usage patterns differ significantly. As described in Chapter 4, the point of access can play a major role in Internet usage, and at-home access is still rare in Macha. Thus, network load is lower in weekends. In addition, the typical diurnal pattern observed on a weekday is less pronounced on weekends. This implies that a solution that relies on low usage periods, such as our proposed proxy (Section 7.5.1), needs to adapt to changing traffic patterns.

7.4.2 Upload patterns

We now concentrate on content generation and in Figure 9.5 plot the size distribution of all TCP flows in either the upload or download direction. According to the figure, there is a significant number of small flows. These correspond to

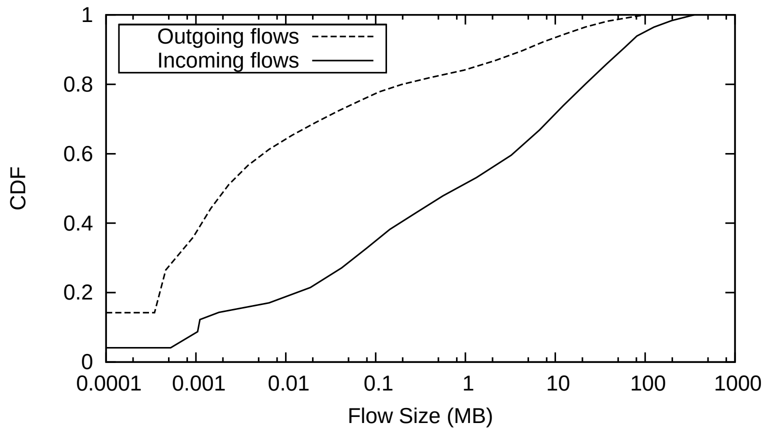

Figure 7.5: Size distribution of incoming and outgoing flows.

short, often automated messages (i.e. TCP ACKs). User-generated media, such as photos, podcasts and videos, are larger. We take 100KB as a rough size limit over which an upload flow is considered to carry a user-generated content. We first note that the fraction of outbound flows greater than our threshold of 100KB is only 22% by volume. This is in stark contrast to incoming flows where those greater than 100KB account for 80% by volume. Hence there is already a weak tendency to upload content, most likely due to both frustration and a lack of a content-generation culture. Chapter 5 did show, however, that there is an active messaging culture that is not reflected in flow volume statistics due to their small size.

We inspect HTTP uploads larger than 100KB according to the request time. We take the full two-month trace and put each of the outgoing TCP flows in one of 24 bins according to the time of day when each began. These flows are additionally differentiated based on their completion. In Figure 7.6 we show both successfully completed and aborted flows. Aborted flows may have timed out or otherwise be incomplete. During the peak usage time the number of upload attempts rises, but the success rate is still very poor.

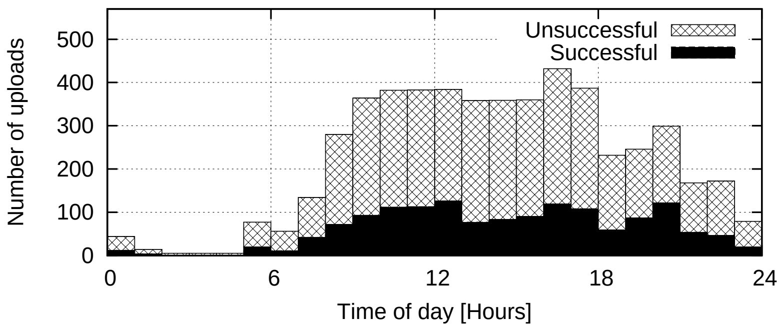

Figure 7.6: TCP upload attempts sorted in 24 hour bins.

To identify types of content that users in Macha generate, we inspect the HTTP uploads according to the application. We concentrate on those flows that are larger than 100KB; a total of 1.2GB of such flows, sent via 2260 upload requests, was detected. We extract HTTP POST requests from the trace and inspect the "host" field. We group requests by domain into eight categories based on either frequent occurrences or significant traffic load. The most popular types of traffic are: "Facebook", "Web mail" (Gmail, Yahoo mail and Hotmail), "File sharing" (4share and Skydrive), "Blogs", "Podcasts", "Education" (LinkedIn, Mendeley, JoVE), "Web design" and "Amateur radio" (Hrdlog.com). The traffic classified as "Other" corresponds to online dating services, ICQ, and automated software

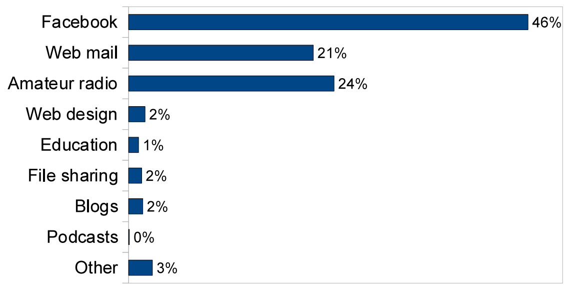

(a) By number of requests.

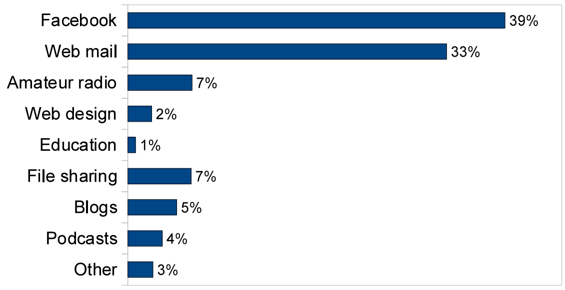

(b) By upload size.

Figure 7.7: Upload web traffic by domain. Distribution of 2260 HTTP POST requests (approx. 1.2GB) observed in the two month trace.

crash reports, among others. In figures 7.7a and 7.7b we show the distribution of upload requests and traffic volume, respectively.

More than two thirds of requests, both by number and volume, correspond to Facebook and web mail. Photo sharing via OSNs and email attachments are common, as can be seen from the trace. However, video and audio uploads are still rare in rural Zambia. We saw only two podcast uploads and no video uploads, though Chapter 4 reveals that watching (download) videos is common in Macha. Video and audio content can be significantly large; the main reason for the lack of upload media content most probably stems from poor network performance during large uploads. File sharing is present with 7% of total upload traffic, but these files are not large, with the median size slightly larger than 2MB. Blogs account for 5% of total traffic; however, these are often written by foreign visitors that are temporarily stationed in Macha. To our surprise 7% of the uploaded content corresponds to an amateur radio application. We suspect that these are automated updates.

7.5 Villageshare architecture

Two alternate VillageShare architectures are described, one that integrates with Facebook and uses a Facebook application to manage content sharing between users and another that makes use of a cloud-based content sharing solution and allows users to share content even when the Internet link is off-line.

7.5.1 Facebook-based architecture

In this section we describe the architecture for our Facebook-based content sharing system. The system makes use of a Facebook-based application and associated local and international file-storage servers to facilitate content sharing between local users and between local and non-local users. In addition it employs a mechanism for time-shifting large uploads to low-usage periods in the network. We focus on time-shifting uploads as time-shifting downloads has already been addressed in previous literature [166, 83]. In addition, supporting uploads meets our objective of facilitating content generation.

Time-shift upload proxy

The goal of the time-shift upload proxy is to “equalize” traffic burden on the satellite link. Due to congestion, uploads often fail (Figure 7.6). One way of relieving the congestion is to reschedule large uploads for periods when the network is underutilized, such as nights and weekends. Two important benefits arise from such a policy: first, uploads are likely to be completed successfully, and second, short interactive flows, that are crucial for local communication via IM/VoIP, for example, perform better with large flows out of their way.

To distinguish between short and long flows we first examine the HTTP upload behavior. A file upload using the HTTP protocol involves an ordered set of TCP messages shown in Figure 7.8. This message exchange reveals that the size of the content is only known after the content has been uploaded. The size of the content is contained in an HTTP POST message at the end of the flow. Hence the only information that can be used to determine a size of the flow, before the entire file is uploaded, is the actual real-time TCP traffic flow as the upload is in progress.

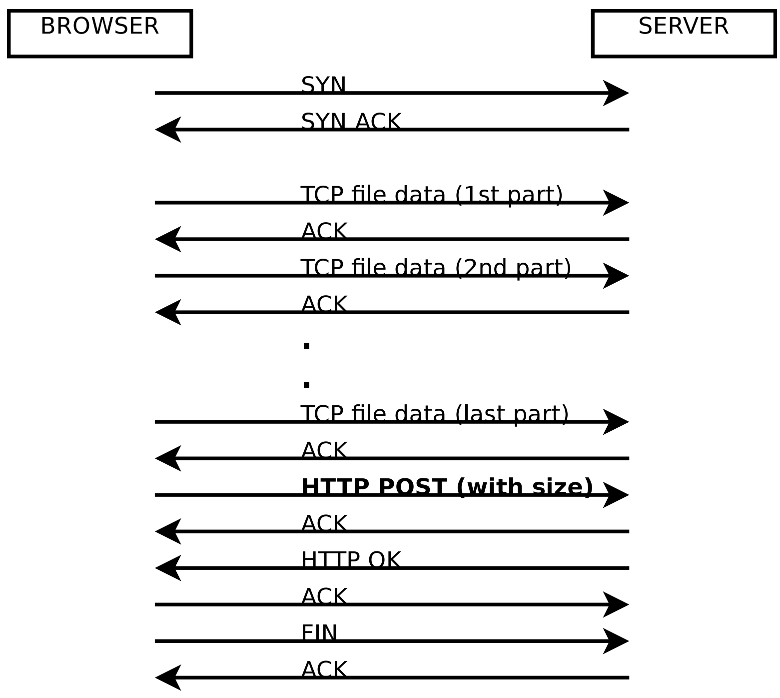

Figure 7.8: TCP flow messages during a file upload.

Another challenge for shifting large uploads is to determine a suitable time to reschedule traffic. While exact prediction of network usage at any specific moment is impossible, we can use daily and weekly trends in network usage, as they exist, to more accurately predict future utilization. A 24 hour running average of outgoing traffic is shown in Figure 7.4a for Mondays to Fridays and Figure 7.4b for Saturdays and Sundays. We estimate the maximum average capacity in the network using the maximum values in these figures. Since satellite connectivity uses a time-shared link, in which many other users within the satellite's coverage area occupy the same frequency, the capacity may fluctuate over time as users come and go. Even if that fluctuation happens, one would expect that all satellite users within the same time-zone of the satellite coverage area would exhibit the same usage pattern (similar to figures 7.4a and 7.4b) and have peak usage periods at the same time. Thus, our capacity estimation represents the lower bound, and the actual capacity during off-peak times might be even higher, which provides strong motivation for our time-shifting proxy.

We consider two approaches for identifying time periods when spare capacity is available and in which we can reschedule large uploads. The first is termed "dynamic" and attempts to scavenge any spare capacity, defined as the area between bandwidth consumed by non-rescheduled, current outgoing traffic and the capacity of the satellite. The second strategy identifies "off-peak" periods in which the traffic load is below a pre-determined threshold. This threshold value can vary such that the time-shifting proxy is able to service all the delayed uploads within 24 hours. To harness short periods of extra capacity, "dynamic" requires real-time usage estimation, for example, via TCP round trip time measurements; "off-peak" relies on a moving average of usage, and can only capture trends in capacity usage. Due to a significant variation in traffic load (Figure 7.4), the threshold can be calculated separately for weekends and weekdays.

The final piece of our proposed proxy is a set of policies that manage upload queuing and scheduling. Our scheduling policy has to balance between upload success and interactivity. When a large upload is observed during a peak period, the proxy captures the file upload in its local cache and notifies the user that the upload will be delivered to its final destination when a sufficient capacity is available. When such a moment arrives, queued upload requests are served in a FIFO manner and the cached content is transferred via the satellite link. As a part of our future work we plan to investigate more complex approaches that take into account upload size (similarly to Bassa [166]) and real-time satellite capacity measurements [39].

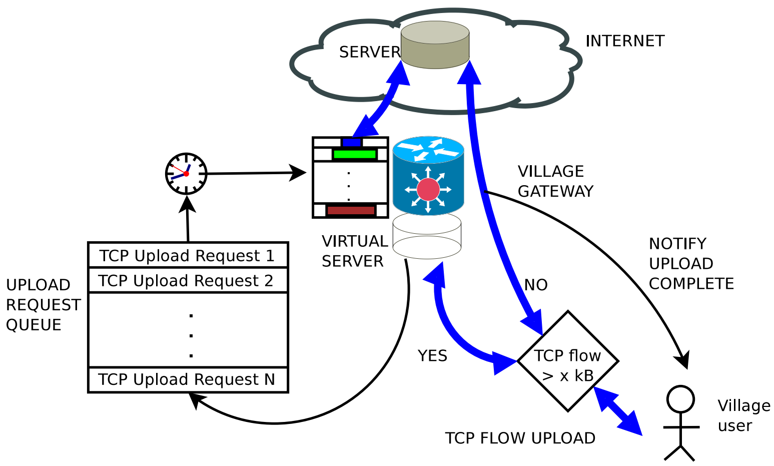

Figure 7.9: Architecture of a time-shift proxy for upload.

The full architecture of the time-shifting proxy is shown in Figure 7.9. A gateway (labeled "village gateway") is placed between the client and the Internet in order to intercept an upload. Once the outgoing flow exceeds a threshold size and the satellite link usage is high, the gateway redirects the user to a local upload page (labeled "virtual server") and the connection via the satellite link is disabled. The upload page will capture the user's authentication information required to upload the content to a content server at a later stage. An example of a common protocol used to authenticate a user with a content server, such as Facebook or YouTube, is OAuth. The gateway stores the URI details of the upload in a FIFO queue (labeled "upload request queue") and will reuse this information to upload the content when spare capacity is available.

Localized and time-shifted file sharing

Web 2.0 introduced numerous applications that allow users to generate and share content: Facebook, Picassa, Youtube, and Google documents, to name a few. However, these applications store content on servers that may be geographically distant from where the users are located. While this does not present a major issue for well connected broadband users, it significantly deteriorates Web usability for those who have low bandwidth connections [171].

To minimize local traffic on the Internet gateway and thereby improve the user experience, we develop the VillageShare Facebook application shown in Figure 7.10. VillageShare ameliorates expensive satellite link traversal by storing locally generated content on a server located within the local network. A user is able to upload a file to the file server and then share the content with their friend

Figure 7.10: Main VillageShare application screen.

list. As a Facebook application, VillageShare inherits all the useful functionality of Facebook. Figure 7.10 shows the main screen that is presented to the user, where a list of uploaded files and file operations is presented. Each uploaded file can be shared, downloaded or deleted. When the share option is selected, the user is taken to the file sharing screen where features inherited from Facebook's functionality are used to search for friends to share a file. Content can either be shared using a direct message to a friend or by posting to the owner's wall or a friends wall. The message contains a unique URI, containing the content creator's Facebook ID and the File ID, which will launch VillageShare and display the shared file to the user.

The architecture of the system is shown in Figure 7.11a. To support VillageShare functionality, a number of servers work together: (1) the Facebook server; (2) a VillageShare Facebook application server hosted externally; (3) a local file storage server hosted in the local village; and (4) a file storage server hosted ex-

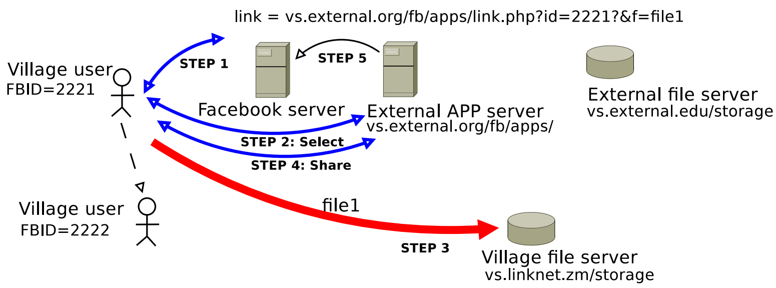

(a) Sharing between two village users.

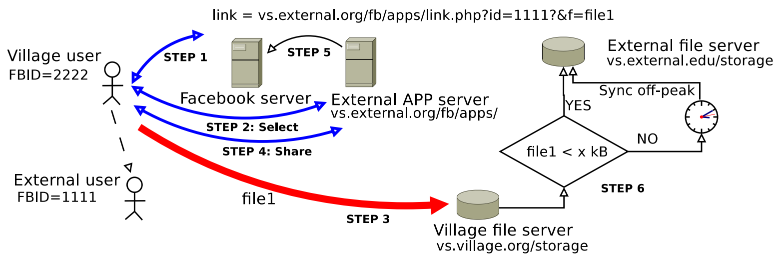

(b) Sharing between a village user and an external user.

Figure 7.11: Architecture of VillageShare Facebook application for uploading files.

ternally. A key component of VillageShare is determining whether a user is inside or outside the village. This can easily be achieved by checking the IP address of the user using the VillageShare application. All users from the village are routed through a satellite gateway using NAT, which is allocated a single IP address by the satellite service provider. When this IP address is detected, VillageShare knows that the user is in the village; all other IP addresses are outside the village.

Note that user roaming is implicitly supported, as we do not use Facebook IDs, but IP addresses to establish the locality.

The following steps, highlighted in Figure 7.11a, are followed when sharing files between local village users: (1) The user connects to Facebook and logs into their account. (2) The user launches the VillageShare application from Facebook. (3) The user uploads a file to the local file storage server through VillageShare. (4) The user selects one or more local friends with whom to share the file. (5) A message is posted on the recipient's wall or inbox with a link to the file using embedded PHP fields. All interactions with the Facebook server and the External Application server are lightweight HTTP messages; media is only sent to the Village file server.

It is also possible for a user in the village to use this application to share a file with a user who is outside the village. From a user perspective the process is the same as when sharing a file with a local user; however, in the back-end new techniques are required to prevent the satellite gateway from being used during peak usage periods. Figure 7.11b shows the case where a village user shares a file with an outside user. The first five steps are the same, with an additional synchronization step (6) added. The synchronization engine checks to see whether the file exceeds a threshold size. If so it delays synchronization between the local village file server and external file server, as described in Section 7.5.1.

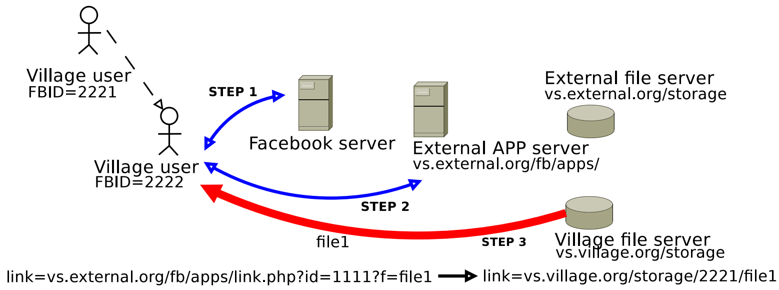

(a) Sharing between two village users.

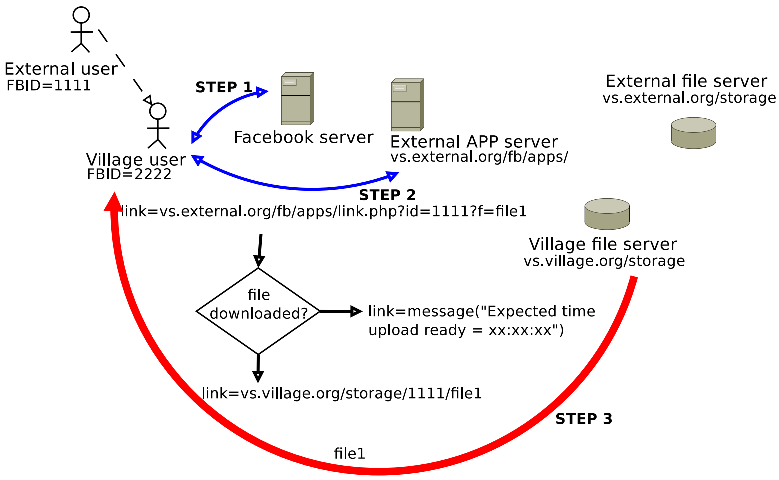

t

(b) Sharing between a village user and an external user.

Figure 7.12: Architecture of VillageShare Facebook application for downloading files.

In the download direction, when sharing a file between local village users, the following steps, shown in Figure 7.12a, takes place: (1) Recipient logs into Facebook. (2) User notices that they have a message or wall post with a link to a shared file and selects the file. VillageShare is automatically launched. This is not a hard-coded URL but rather a link to VillageShare with parameters specifying the source user and file being shared. (3) VillageShare knows that the user attempting to retrieve the file is local, using the IP address method discussed, and also knows that the file was generated by a local user using a location tag stored with the file which stores the location of the user who uploaded the file. Using this information it generates a URL pointing to the source file on the local village server and the file is downloaded. In the case of a local user downloading a file posted by an external user, shown in Figure 7.12b, the link will either resolve to a message saying that the file is not downloaded to the local store yet, or it will resolve to a URL pointing to the file on the local village sever. The synchronization status will be known as the Facebook application server will be notified when synchronization is complete.

7.5.2 Cloud-based architecture

Our initial version of VillageShare was built as a Facebook application because of Facebook's prevalence in developing regions — highlighted in Chapter 3. After deployment of the Facebook-based VillageShare, we discovered that the integration with Facebook prevented some users from participating due to Internet access costs; many users cannot afford to pay Internet connectivity charges. When a user logged in with their Facebook account, Facebook itself had to be contacted, re-

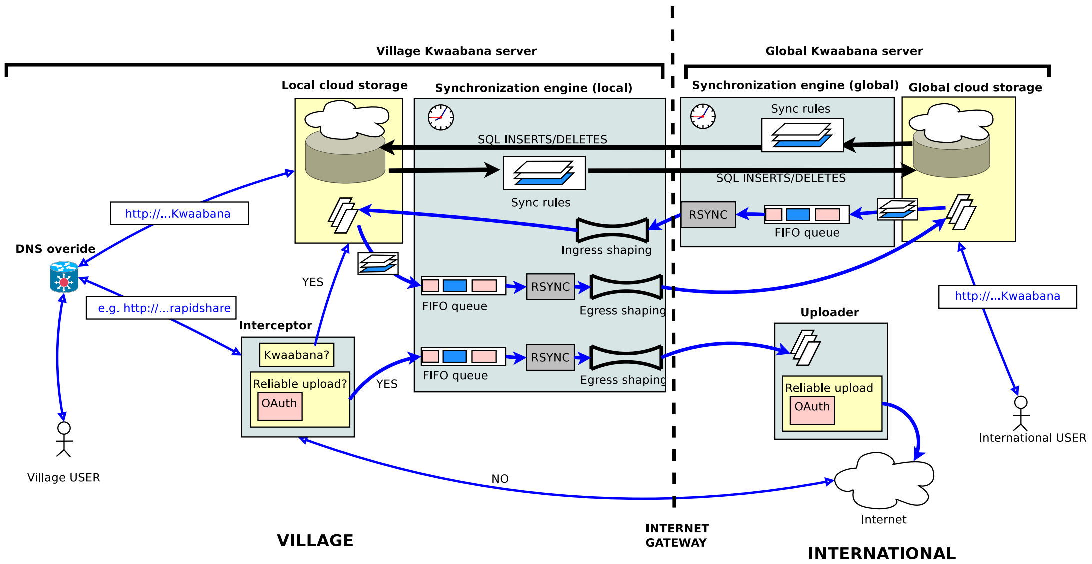

Figure 7.13: The VillageShare2 architecture.

sulting in per-byte usage fees from the local Internet provider. To provide broader access, we designed a second version of VillageShare, this time integrated with an open source file hosting service called *owncloud*. The local cloud hosting service allows users to share hosted content with any other user registered on the cloud. Users can share content with other local users for free. We now describe the latter version of VillageShare built on *owncloud* and for purposes of clarity, we call this version VillageShare2.

VillageShare2 consists of a set of networking components integrated with a cloud storage application. The architecture of the system is shown in Figure 7.13. The primary networking components include a traffic interception system (*Interceptor*) to re-direct users to VillageShare2 should they use file sharing Internet-services, and a *synchronization engine* that carries out lazy database and file synchronization. The *synchronization engine* maintains partial consistency between the local and global cloud storage server at UCSB while minimizing impact on interactive browsing traffic.

Content on the local and global cloud storage database is only partially synchronized; there is no need to have knowledge of an international user's file if it has not been shared with a local village user, and vice versa. Once files are uploaded to VillageShare2, they are stored in the user's personal storage space. The user can then share those files with other local or global users. If the files are shared between local users, the files are only stored on the local cloud storage server. Should a local user share a file with an external user, the file is placed in an upload *FIFO queue* to synchronize with the global cloud storage server. A set of *sync rules* are checked to decide which records in the local database need to be inserted into the global database. A scheduler executes the list of database updates at a regular interval and triggers the synchronization engine to upload files from the queue. We describe the database design in Section 7.5.2.

We carry out file synchronization through a service that prioritizes interactive traffic such as HTTP. Our non-interactive file uploads and downloads receive a small slice of the bandwidth during periods of high utilization, and a larger slice when utilization is low. The service also uses the Linux *rsync* tool which is able to recover from partially synchronized files should there be severe congestion or a network or power failure. This service is critical as our traffic analysis in Chapter 4 discovered that as many as 75% of file uploads are unsuccessful due to network congestion over the slow Internet access link. We describe the synchronization service in Section 7.5.2.

VillageShare2 is available to all users, regardless of whether they reside in a rural region, via a single URL. However, communication between users within the local village are routed to the village VillageShare2 service by the *DNS override* service that overrides the global Internet DNS entry. Rural users who attempt to upload files to common file hosting sites such as youtube are also intercepted to check whether the user intends to share this file with other local users. If the user does intend to share the file locally, they are redirected to the VillageShare2 service to upload their file and share with a local user. We describe this file upload interception in Section 7.5.2.

In the following sections, we describe how each of these components work together to provide the full set of VillageShare2 functionality.

DNS override

It is common practice in the Internet to ensure that requests to servers are load balanced across replicated servers. In many cases these servers are replicated worldwide and the server that is closest to the client is used. Our VillageShare2 server (www.village-net.org) is registered on a global DNS server to a machine in the USA. However, in the rural community the local DNS cache must be manually edited to point this domain to a local machine, redirecting all VillageShare2 application requests to the local server. This reduces not only delay and Internet costs, but also simplifies deployment as local and external users only need to recall a single URL.

File upload interceptions

There are a number of monolithic Internet-based file sharing services, such as youtube and yousendit, that allow a user to share files. It is, however, a waste of costly Internet bandwidth to make use of these services to share content between local users in the village. To avoid this, uploads to these services are intercepted using an *iptables* filter and redirected to the VillageShare2 service. The goal is to introduce users who may not be using VillageShare2 to the service when they use Internet services that could be performed more efficiently with VillageShare2.

Once interception has occurred, the user is asked whether the file being uploaded is destined for a local user. If so, an option is presented to use the VillageShare2 service. An additional option is presented to share files with outside users that allows the user to use a more reliable upload service. We have found in

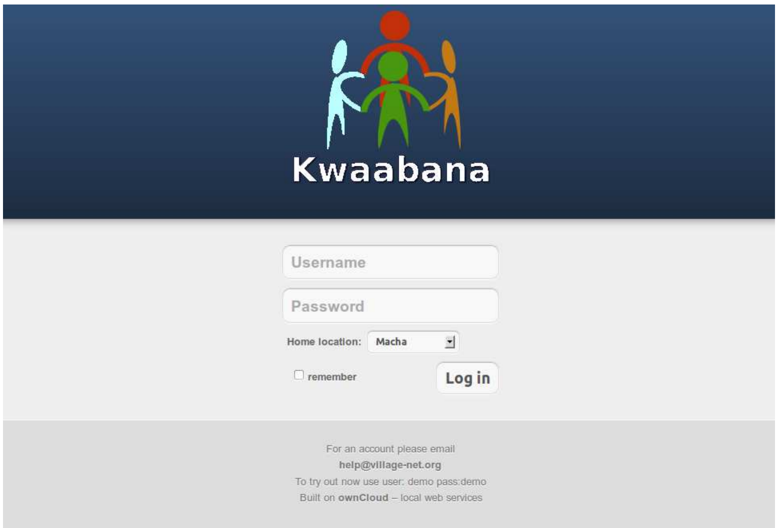

Figure 7.14: The main VillageShare2 login page. Notice the option to select location.

previous studies that uploads often time out due to network congestion or power failures, and that file hosting services do not allow users to recover from a failed upload. Use of a reliable upload service ensures that the file is first synchronized using a recoverable upload process to a server at UCSB running our uploader service. From there it is uploaded to a file sharing service. A user's OAuth[10] credentials, an authentication mechanism used by most file hosting services, are captured during the upload to re-authenticate with the file sharing service from our UCSB server. The file synchronization utilizes the Linux *rsync* tool, which is capable of recovering from a partial file upload. A file is only removed from the

[10]Open standard for authorization.

upload queue once it is successfully uploaded. When a server boots after a power failure, the synchronization service checks the upload queue and reruns the *rsync* tool on the first file in the FIFO queue, which may be partially uploaded or not uploaded at all.

User interface design

Our initial Facebook-based VillageShare application provided a very simple upload and sharing interface which allowed users to share files with their Facebook contacts. After we noticed no significant uptake of the service in the community due to exclusion of user's without Internet vouchers[11] and the lack of a rich interface which provided other features such as a media player, we modified the design to make use of a completely localized file storage and sharing solution that does not require the user to use the Internet connection.

The new VillageShare2 file sharing service utilizes a local cloud storage solution. A user registers her name on the local service and then proceeds to upload and share content with other local or external registered users. The user interface is built on an open source file hosting project called *owncloud*. We incorporate our synchronization modifications and a feature to allow users to associate their account with their Facebook profile. This enables them to receive notifications on

[11] $30 Internet vouchers are required to access the Internet for 30 days

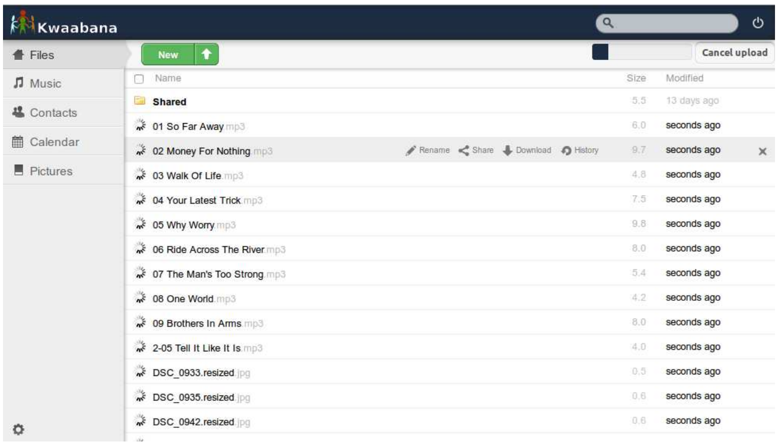

Figure 7.15: The VillageShare2 file explorer page. Selecting "share" from this page allows users to share the file with users registered in the VillageShare2 system.

their Facebook account if a user shares content with them when they are logged into the Internet.

The VillageShare2 login screen is shown in Figure 7.14. An account is requested from the administrator by sending an email request or, in the case of a user who does not have an Internet voucher, by physically visiting the administrator in the village. A user name and password is entered, followed by a location that is set to the current location by default. The purpose of the location entry is to allow visitors who have an account on a remote VillageShare2 server to access their files.

The main VillageShare2 file explorer interface is shown in Figure 7.15. The choices a user can make from this *owncloud* interface are: (a) upload a new file or set of files. Files sizes up to 2GB are supported; (b) view supported media, such as pictures or videos. Pictures can be viewed as a slide show; (c) listen to their music collection directly from the web browser; (d) download files to their computer; or (e) share a file with another registered user. When sharing a file, a user can select whether the shared file is allowed to be edited. Files that have been shared with a user are placed in the *shared* folder.

When a user logs in to an external VillageShare2 server (i.e. a server outside of her home community), a full set of file meta data is retrieved from the remote server; however, the actual files are not synchronized unless they are requested. After logging in the user is presented with a message: "Requesting your set of files from the remote VillageShare2 server." Once the file list is synchronized, the full set of files is shown in the file explorer interface. When selecting a file, a user is presented with an option to request a file from the remote server. The process of synchronizing the file list and file is presented in sections 7.5.2 and 7.5.2.

Database design

VillageShare2 stores user and file information, as well as information on sharing between users, in a *mysql* database. Files are stored in a directory structure in which a file is found by traversing folders related to the owner and owner loca-

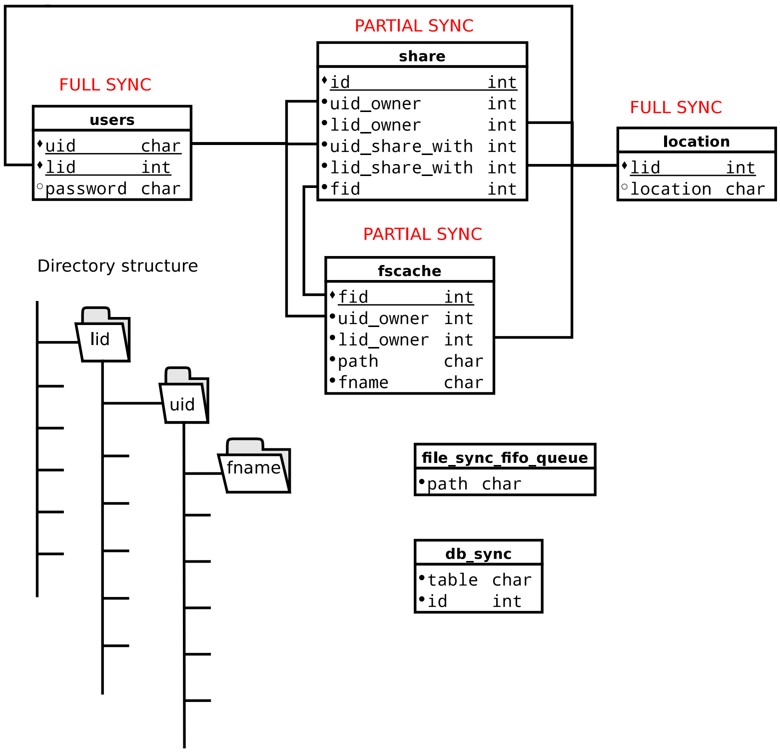

Figure 7.16: The database design, showing tables that support locality and file sharing across different VillageShare2 locations.

tion. The structure of the database showing the core tables in the VillageShare2 database as well as the directory structure is shown in Figure 7.16.[12]

To identify users, a user creates a unique user name with a location id attached to this user account. This information is stored in the *users* table. It is possible for two people with the same user name to register in different locations as the combination of user name and location id forms a primary key. As new locations are added they are stored in the *locations* table.

[12]Not all fields and tables are shown. What is shown is sufficient to illustrate the core architecture of the system.

When files are upload by a user, the user automatically becomes the owner of these files. File information such as the file name, owner of the file, location and other file meta data not shown, is stored in the *fscache* table. When a file is shared with another user, the owner id and owner location id, as well as the user id and location of the person with which the file is shared, is stored together with the file name in the *share* table.

Synchronization between databases at different locations is partial; it is not necessary to have knowledge of a set of files at another location if the files are not shared with a local user. Hence some tables require full synchronization and others do not. The *users* table is fully synchronized in order for users at different sites to have knowledge of each other and share files. The *locations* table is also fully synchronized as new locations are added to the system. From the point of view of a local user, the *fscache* table is partially synchronized as it only needs to store information for files owned by local users and for files from remote users that have shared files with local users. The *share* table is also partially synchronized as it only needs to store information about files shared between local users and files shared between local and remote users; no knowledge of files shared between two remote users is required.

If we extend this system to multiple remote sites, updates to tables requiring full synchronization would need to be propagated to all other sites and updates to tables requiring partial synchronization will only be propagated between sites with which a file share has occurred. Eventual consistency of records needing synchronization due to connectivity issues, for example, has no effect on the core operation of VillageShare2 – it simply results in no knowledge of a remote user or shared file until consistency is restored.

When a new user is added or a file is shared with a remote user, database records that need to be synchronized with the remote database are entered into the *db_sync* table. Suppose a new user *bob* is registered. This table would then contain a record {users,bob}. If a file with id *11* was shared between *bob* at local location id *1* and external user *alice* with location id *2*, two entries will be required in the *db_sync* table. Assume the id of new record in the *share* table, which captures the sharing event, is *22*. The first entry would be {fscache,11} and the second entry would be {share,22}.

As discussed in Section 7.5.2, it is also possible for a remote user logged in to a local server to request a copy of their files. In this instance, all the ids for the complete set of records for the remote user in the *fscache* table is placed in the *db_sync* table.

A scheduler in the synchronization engine runs a task regularly that converts the entries in the *db_sync* table into a set of SQL INSERT statements stored in a file and sent to the remote database to be executed. Thus, only a partial set of relevant updates are executed on the remote database. This process happens in the reverse direction as well if a new remote user is registered or a remote user shares a file with a local user. Should users or files be deleted, a set of SQL DELETE statements will be created in a similar way. Bandwidth requirements for sharing these SQL update files is low as they are text updates that only reflect differential changes.

The actual files that are shared between local and remote users are entered into a fifo queue in table *file_sync_fifo_queue*. This table also holds files that are requested from a remote user as discussed in Section 7.5.2. A synchronization service, discussed in Section 7.5.2, is used to synchronize the files in this queue between the local and remote VillageShare2 servers.

Synchronization and traffic shaping

The synchronization service is designed to synchronize files in the synchronization queue between the local VillageShare2 file store and the global VillageShare2 file store and vice versa. To optimize use of the limited-bandwidth Internet gateway that frequently exists in rural communities, we shape the synchronization traffic to minimize impact on the browsing activity of users while optimizing usage of the unused portion of bandwidth available in the Internet link.

The Traffic Control subsystem of the Linux Kernel is a sophisticated set of tools to shape traffic before it is passed to the network device. It allows the establishment of different queuing disciplines for packets based on classification according to criteria such as ports or IP addresses or any information available in the IP header (i.e. Type of Service (TOS)). These queuing disciplines, or *qdiscs*, represent the scheduling policies applied to a queue. Several types of *qdiscs* exist such as FIFO, Stochastic Fair Queuing (SFQ) and Token Bucket Filters (TBF). To shape the synchronization traffic, which is sent using the Linux *rsync* tool, we allocate a unique port to the synchronization process and use this to classify the synchronization traffic. As discussed in Section 7.5.2, the combination of a FIFO queue, in which a file entry is only removed once the file synchronization is complete, and the recovery feature of the *rsync* tool make the system robust against upload failures due to power or network outages.

In order to shape traffic on the incoming (ingress) and outgoing (egress) links, we need to take control of the Linux queue attached to the network device. To do so, the Internet gateway machine needs to become the bottleneck for the traffic flowing to the Internet. In Macha, there is a 10 Mbps Ethernet link between the gateway machine and the WiFi router; the WiFi router provides a 2.1 Mbps[13] link to the Internet. If the outgoing traffic is not shaped to just below 2.1 Mbps,

[13] 2.1Mbps is SLA agreement with the ISP.

all traffic will build up in a queue at the WiFi router which we cannot control. To control the queue, we shape the outgoing link to 2Mbps. We then classify the traffic based on any number of different criteria, such as HTTP traffic or traffic using the port allocated for file synchronization.

In a gateway serving a community where web traffic dominates, it is advantageous to divide the traffic so that most of the bandwidth is allocated to HTTP traffic. Traffic that needs low latency, such as TCP ACKs, should be pushed to the network device as quickly as possible. Because it is not interactive, synchronization traffic should receive a small slice of the bandwidth if other services are using the link but consume as much as possible when the link is dormant. To split the traffic we use a hierarchical token bucket queuing discipline. This allows us to set a token bucket-based rate when the link is under contention but allows the service to borrow tokens from other classes when capacity is available, up to a ceiling rate. HTTP traffic is sent to a stochastic fair queue to avoid one HTTP connection dominating another. ACK packets are sent to a FIFO queue to allow them to reach the network interface as fast as possible. The VillageShare2 traffic is also sent to a FIFO queue as only one file is synchronized at a time. The system design is shown in Figure 7.17

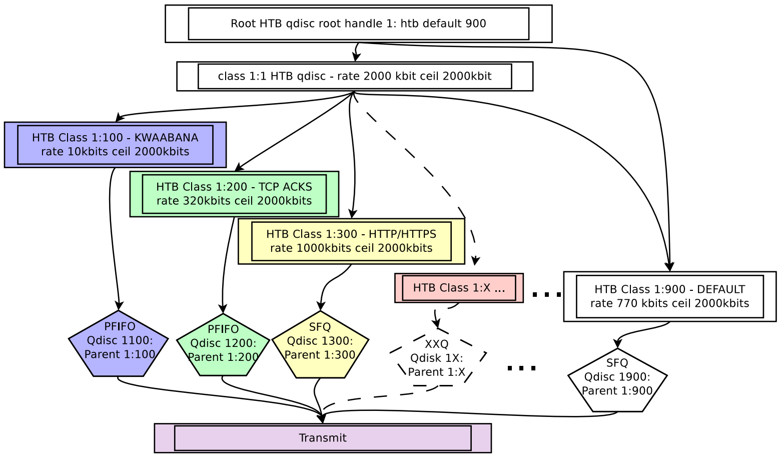

Figure 7.17: Traffic shaping design for VillageShare2 using the Traffic Control subsystem of the Linux Kernel.

7.6 Villageshare evaluation

We now evaluate both our VillageShare architectures in turn. The Facebook-based VillageShare is evaluated via a simulation as uptake in the rural village of Macha was limited; this was due to the limitations of the user interface and the exclusion of users who did not have Internet vouchers. The cloud-based VillageShare, VillageShare2, is evaluated using data from a live deployment in Macha, Zambia.

7.6.1 Facebook-based VillageShare evaluation

Our evaluation methodology utilizes intuitive calculations based on our traffic analysis to predict bounds for performance improvement that could be provided by VillageShare. We focus on outgoing traffic as our goal is to facilitate file uploads.

Traffic localization using VillageShare

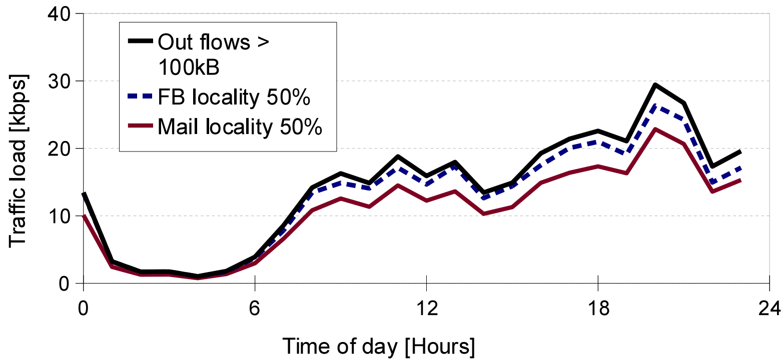

Figure 7.18: Capacity savings with large upload handling through VillageShare.

VillageShare relieves Internet link congestion by bypassing the gateway and keeping local traffic within the local network. We investigate how much new capacity would be made available based on the amount of local traffic. We perform the evaluation assuming user behavior as observed in the trace. We focus our evaluation on the dominant upload traffic from Facebook and email, as these applications are most likely to benefit from VillageShare.

We infer the amount of locally produced and locally consumed traffic from the locality of Instant message and photos exchanges on Facebook calculated in Chapter 5. 25% locality was found in Facebook photo sharing over a bandwidth constrained satellite link; less-constrained Facebook message sharing in the same network revealed 50% locality of interest. Thus, we take 25% and 50% as the lower and the upper bound on Facebook locality, respectively. If we assume that this trend is consistent in other forms of content sharing, we can apply this fraction to emails with attachments sent to local users.

Based on this assumed amount of local traffic, Figure 7.18 shows the bandwidth saved over an averaged 24 hour cycle with VillageShare. Much of the bandwidth savings occurs during the day, which helps relieve the congested gateway and improve interactive application performance. The remainder of the large uploads, not shared with local users, will utilize VillageShare's time-shift proxy. Table 7.1 summarizes the upper (50%) and lower (25%) bound of bandwidth savings due to VillageShare over a two month period. All traffic volume percentages are given as a fraction of large outgoing flows excluding P2P traffic. We exclude P2P traffic as it does not capture explicit user upload activity. We anticipate that, as upload performance improves, upload traffic volume will increase. Hence, these savings, by volume, are likely to escalate when much larger files, such as videos and software, are shared locally. Moreover, VillageShare leads to improved user experience, due to fewer failed uploads and faster upload times.

Table 7.1: *Summary of outgoing bandwidth used and potential bandwidth savings due to VillageShare.*

Traffic types	MB
Total outgoing	45287
Total outgoing ¿ 100KB (excl. P2P)	6366
email clients	1568 (24.63%)
Web-email	611 (9.60%)
Facebook	744 (11.69%)
VillageShare(VS) savings	
Facebook 25% locality	186 (2.92%)
All email 25% locality	545 (8.56%)
Facebook 50% locality	372 (5.84%)
All email 50% locality	1090 (17.11%)
Total saved due to VS (25%)	731 (11.48%)
Total saved due to VS (50%)	1462 (22.96%)

Time-shift upload proxy

We next estimate the amount of spare uplink capacity the time-shift proxy will be able to utilize in Macha. The satellite data bundle used in Macha claims to provide a 64 MHz committed rate and a 256 MHz bursting rate. From our traffic analysis, we found an average peak capacity of approximately 128 MHz (see Figure 7.4a); we use this as our baseline uplink capacity. This provides a maximum daily outgoing capacity of 1328MB. We assume that traffic between

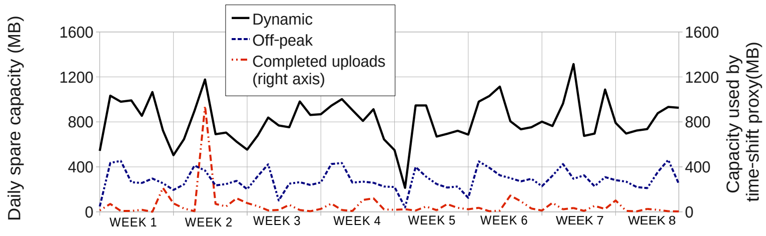

Figure 7.19: The capacity available each day for the time-shift proxy using dynamic and off-peak allocation strategy. An estimate of the bandwidth used by the proxy for completed uploads is also shown.

local users harnesses the VillageShare file-sharing application. As a result this traffic does not figure in the flows we consider for rescheduling. We concentrate on the remaining large uploads (greater than 100kB) which are rescheduled by the time-shift proxy; hence this traffic is initially removed from the current set of outgoing flows.

We evaluate the two spare capacity estimation strategies elaborated in Section 7.5.1: "dynamic" and "off-peak". In the "dynamic" strategy the proxy automatically dequeues any enqueued flows when capacity becomes available. The threshold for the "off-peak" strategy is set to a third of the satellite capacity (42.6 MHz). This results in an off-peak daily period from 3 to 8 A.M. with an additional weekend period from 2 to 5 P.M. The proxy begins to dequeue flows at the beginning of the off-peak period.

Figure 7.19 shows the spare capacity available using these two strategies over the 8 week measurement period, in intervals of days. Clearly, the dynamic strategy provides far more capacity, with an average of 700MB per day; however, it has the disadvantage of negative interaction between upload traffic (long-lived flows) and browsing traffic (short-lived flows) [61]. The off-peak strategy provides an average of 278MB extra capacity per day, which is sufficient for all but one day in our 2 months network measurement. The large dip in capacity at the beginning of week 5 is due to a full-day power failure.

We now reinsert the removed uploads, destined for external consumption, back into the network. We conservatively assume that 75% of all large email, Facebook and non-P2P large outbound flows are these externally consumed uploads. They are then reinserted using one of the two strategies, "dynamic" or "off-peak". We assume that the measured size of these flows in our traces was artificially small due to the aborted transfers. To make the sizes of these flows more realistic, we double the length of these flows. The bottom line in Figure 7.19 shows the amount of capacity required by these flows over time. Both proxy strategies can handle the majority of the current outbound flows greater than 100kB within a 24 hour period. The sudden increase in demand in the middle of week two is due to an unusual email client anomaly, where large attachments were transmitted continuously throughout the day. For this single anomaly, only the dynamic bandwidth allocation strategy would have been able to service the uploads within a single day.

It is likely that the improved network performance, due to local caching, will cause users to increase the quantity of uploads. The evaluation shows that there is enough capacity to handle a substantial increase in uploads, but we acknowledge that only an in-situ deployment will be able to evaluate the efficacy of the time-shift proxy as outgoing load increases beyond its current pattern.

7.6.2 Cloud-based VillageShare evaluation

In this section we evaluate the operation and some post-deployment restricted usage of VillageShare2. We study the type of content uploaded and shared on VillageShare2. We depict the sharing between users as a social graph to understand the distribution of sharing behaviour across the user base. We compare this social graph to previous work on the locality of interest for Facebook images and instant-messages. We also evaluate the synchronization component of VillageShare2 using simulated file shares between local and remote users. We carry out experiments during normal peak usage times in the network to ensure that our synchronization component has minimal effect on normal interactive web browsing.

Evaluation of shared content for real users

Following our deployment in Macha, we provided access to VillageShare2 to a limited number of users to test its operation "in the wild". Since the test, we have made VillageShare2 available to all village residents and usage is steadily growing. For the evaluation described in this section, we disabled synchronization between the local and global system until VillageShare2 could be integrated with the voucher based billing system. This evaluation, therefore, only presents local sharing behaviour within Macha. During these initial tests, there were 12 registered users who were encouraged to explore the use VillageShare2 for file sharing. The following results are from the initial two weeks of usage following deployment.

Table 7.2: Statistics for files stored on VillageShare2.

File Type	Quantity	Total Size(MB)
image	51(72%)	169(18%)
audio	5(7%)	71(8%)
pdf	12(17%)	6(4%)
video	2(3%)	669(73%)
executable	1(1%)	3(0.3%)

Table 7.2 shows a summary of the quantity of files uploaded over two weeks using the VillageShare2 service, while table 7.3 summarizes the quantity of files shared. Images are the dominant media type; 72% of the total files uploaded to the server. However, after only two videos were uploaded to the server, 73% of

Table 7.3: Statistics for files shared on VillageShare2.

File Type	Unique files	Total shared	Share ratio
image	6(43%)	22(46%)	3.6:1
audio	4(29%)	13(27%)	3.3:1
pdf	3(21%)	11(23%)	3.6:1
video	1(7%)	2(4%)	2:1
executable	0(0%)	0(0%)	0:0

the storage space was consumed. Only 6 of the 51 stored images were shared with other users whereas all the videos were shared with other users; one of the videos was shared with users registered on VillageShare2 and the other video was broadcast to all users in Macha by posting a URL link to the video on a Facebook page. In the audio category 4 out of the 5 audio files, which were mp3 music tracks, were shared with other users. We conclude that users have different intentions when uploading content to VillageShare2. When content is generated by the user – as most images are – VillageShare2 is used as a backup service for this content with a small fraction of these files shared with other users. Recall that using a Internet file hosting service for backup would be very costly and slow in a remote rural area like Macha. When the content is not locally generated, in the case of downloaded music or videos, the content stored on VillageShare2 is uploaded with a clear intention to share this with other users as backup of these files is not as critical.

We define the share ratio of the content as an average measure of the number of files shared with other users for a specific file type, calculated simply by dividing the total number of shared instances of a file type by the total number of unique files of that type. We find that the share ratios are fairly similar for images, audio, and pdf files – the share ratio of each is greater than 3:1. Chapter 5 analysed locality of interest in Facebook images and showed that local content was viewed 5.2 times on average by local users, compared to externally generated content, which was only viewed 1.5 times on average. We believe the share ratios are likely increase further as usage of VillageShare2 grows, particularly as it grows among users without access to Facebook.

We captured an event where one of the large 600 MB videos uploaded to VillageShare2 was broadcast to all users in Macha using a Facebook link to the content. Analyzing the apache access logs revealed that this video was viewed three times within the first day of posting. Although the video was not directly shown to users of VillageShare2, the broadcast of this video achieved a share ratio of 3:1 within the first day and saved 600 MB of upload bandwidth and 1.8 GB of download bandwidth on the Internet gateway. This is a financial saving of $72 USD, given that 1 GB vouchers cost $30 USD.

The social graph depicting sharing behaviour between users is shown in Figure 7.20. Node sizes indicate the weighted node degree and edges depict the

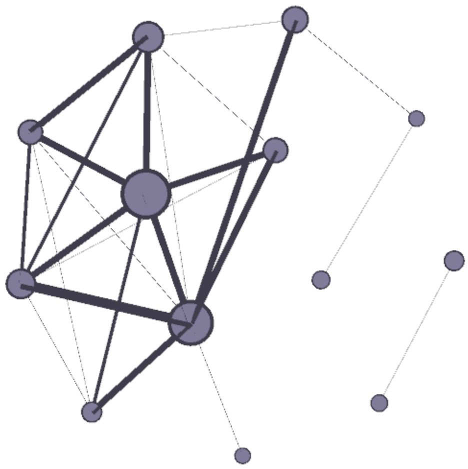

Figure 7.20: Social graph of local file sharing using VillageShare2 in Macha. Node sizes indicate the weighted node degree and edge thickness depicts the number of files shared between users.

number of files shared between users. It is clear from this graph that there is a long-tail distribution of file sharing between users – a few dominant users share the majority of files and many users only share a few files. This behaviour mimics social graph patterns we have seen in Facebook instant messaging. There are a few users who act as ‘information conduits’ or ‘information generators’ and a large number of other users who are consumers of this content.

Evaluation of synchronization

To evaluate the synchronization strategy using the Linux traffic control system, we monitored the down-link of the Internet gateway during a simulated file synchronization between the global and local VillageShare2 sever; this effectively simulates an international user sharing a file with a local village user. Figure 7.21 shows a trace of synchronization traffic and all other Internet traffic passing through the gateway; samples of the downlink throughput were taken every 2 seconds.

One 15 MB file was placed in the upload queue (from the point of view of the global VillageShare2 server). The incoming queue on the village server was restricted to 1.2 Mbps using a hierarchical token bucket (HTB). As can be seen from the plot, the HTB allows some momentary throughput to reach 2 Mbps; however, over a larger averaging window, the downlink throughput is 1.2 Mbps. The scheduler in the synchronization engine started the download at 13:47:30. At this stage the browsing traffic averaged around 800 kbps, and the synchronization service took full advantage of the spare capacity of about 400 kbps by borrowing tokens from the other *qdisc* classes. At 13:50:30, the browsing traffic rapidly increased to the maximum capacity of the link; the *qdisc* for the synchronization traffic then falls to its allocated slice of 10 kbps, with no spare tokens to borrow from other classes. The browsing traffic then eases off at 13:53:00 and the synchronization service takes advantage of the spare capacity again, with the file synchronization completing at 13:57:00. From the plot it is clear that the traffic shaping for the synchronization service is very responsive to increases and decreases in capacity and makes full use of capacity when available but backs off to its "trickle" rate when the link is saturated.

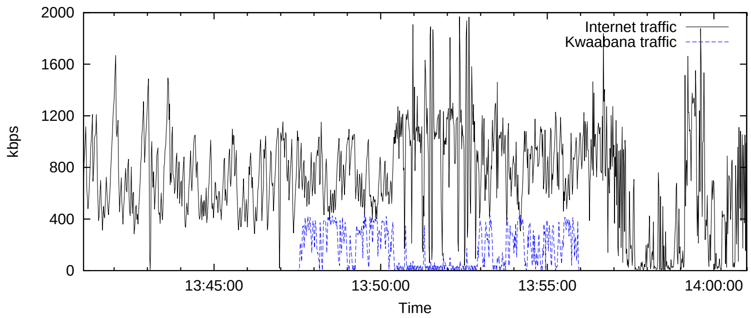

Figure 7.21: Plot of traffic passing through the gateway illustrating VillageShare2 file synchronization and all other Internet traffic.

7.7 Related Work

The lack of local content in rural Africa was noted by Van Hoorik and Mweetwa [164]. According to their observations, rural Africans do not find a representation of their customs and culture online, thus they may perceive the Internet as a "foreign body". In [113] models of online content generation are explored. The authors conclude that social and cultural factors have to be considered for successful implementation of content generation tools. In our Internet usage survey (Section 7.4), we investigate why content generation is lacking even in an existing network with experienced Internet users. We find that the lack of easy access to computer terminals and the lack of bandwidth inhibit content generation.

To cope with limited bandwidth, we propose a time-delayed proxy and an OSN-based local file sharing system. Time-delayed proxy for bandwidth-limited networks was first proposed in [51]. A Collaborative cache approach was proposed by Isaacman and Martonosi [83] in which cached content on client devices is made available to all clients in the local network. This solution acknowledges the value of local stored content as well as the importance of distributed access in unreliable networks. A system similar to the proposed one was implemented by Vithinage and Atukorale [166]. We build upon this idea and extend it so that both uploads and downloads are handled in a time-delayed manner. Locality of online interactions, on the other hand, was first observed with the advent of online social networks. Wittie *et al.* [171] exploit locality of interest in order to improve OSN usability for remote areas. Their work is geared towards larger geographic regions (e.g. whole countries). In this paper we find a micro-locality within a single village can be used as a basis for the improved file sharing scheme we propose in Section 7.5.

The value of local clouds in developing regions was explored in [53]; an order of magnitude improvement in performance was demonstrated with an application mirror installed in a local network. Collaborative caching was proposed by Isaacman and Martonosi [83] where cached content on client devices is made available to all clients in the local network. This solution acknowledges the value of local stored content as well as the importance of distributed access in unreliable networks. A system similar to the one proposed was implemented by Vithinage and Atukorale [166].

7.8 Discussion

VillageShare does not present a method that transparently solves the problem of bandwidth shortage in rural area networks. Rather, it is an intervention that facilitates content generation and sharing by providing an alternative to the existing paradigm of interactive file upload, which clearly does not work well in bandwidth-poor networks. We believe that the poor performance experienced in uploading mail attachments, podcasts, videos or other files shared with local village users will convince users to converge on usage of the VillageShare application. Although the cloud-based solution diverged from the original Facebook-based solution, in the future we would want to see both these approaches converge in order to fully exploit the existing Facebook social graph.

A tradeoff between impacting user behaviour and improving network performance exists on multiple levels. VillageShare exploits locality of interaction and

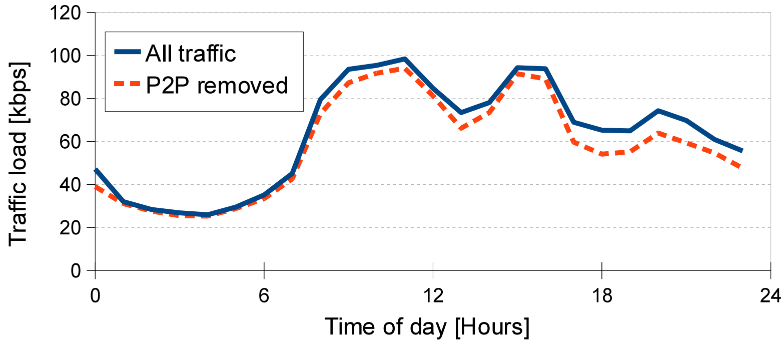

Figure 7.22: Effect of removing all P2P upload traffic on the network (averaged two week trace).

time shifting; traffic shaping is yet another aspect. For example, we find that peer-to-peer traffic utilizes 8.2% of the total bandwidth, while exhibiting very poor performance over an asymmetric satellite link. To examine the impact of this traffic, we show in Figure 7.22 the upload capacity savings when all P2P traffic is blocked at the gateway through a simple firewall rule. Although this could seem like an attractive way to relieve congestion, there are cases when content is only available through a P2P network; it would be unwise to completely remove the opportunity to use this mechanism. A tradeoff solution could be a combination of a complete blockade of P2P networking in the network combined with a service on a dedicated server outside the network to forward P2P file requests. The files can then be downloaded to the local server using the time-shift proxy method that we described earlier.

Bandwidth is not the only restriction on content generation. Cost associated with Internet access is extremely important for low-income rural residents. Our interviews reveal that when Internet usage is limited by cost, it significantly impacts the frequency of content generation ($\chi^2(N = 41) = 3.475, p = .006$). Moreover, the number of hours a user spends using the Internet (irrespective of whether one has to pay for it) determines the user's affinity towards content generation ($\chi^2(N = 41) = 5.218, p = .035$). Access at work, school or any other public location comes with restricted hours and leads to "the deliberate interaction" model [172]. Only at-home access allows leisurely content generation and sharing; we plan to tackle this aspect of the problem in our future work. In our future work we also plan to carry out a long term longitudinal study to capture the changes in content generation and content sharing behaviour of users once the VillageShare system has received considerable uptake in the community. s

7.9 Conclusion

Our traffic analysis and a social survey reveal that rural networks have a unique set of challenges and opportunities. Content generation, vital to capturing the local languages and cultures of rural areas was hindered by a number of factors. Our survey reveals that users are frustrated by costly, slow Internet services that hinder their ability to share content, specifically large files such as video When local users share content by email or Facebook, it traverses a slow Internet gateway twice and often leads to failure.

We presented two VillageShare architectures to solve this problem. Our first VillageShare system makes use of a a time-shift proxy that delays large file uploads to off-peak usage periods, and a Facebook-based file sharing application that ensures that local user traffic does not traverse the Internet gateway. Our evaluation showed that between 11% and 22% of the large outbound flow traffic, normally utilizing the satellite gateway, can be sent between local users using a Villageshare file server hosted in the village. We also noted that the time-delay proxy creates enough off-peak capacity to fulfil the current set and a large number of additional large upload requests.

Our second VillageShare system makes use of a locally hosted cloud-based solution. Early results from our deployment of our second VillageShare system show that the service facilitates easy local sharing of content as well as backup of user generated content such as photographs from mobile phones. Share ratios of over 3:1 for all local content mimic locality of interest seen in our previous Facebook image analysis. Simulations of content shared with external users using our traffic shaping engine show promising results in that a users' browsing experience is not adversely effected by the file synchronization.

As the bandwidth divide increases between developing regions connected over slow Internet and the rest of the international community, services such as VillageShare will be key to enabling vibrant collaboration amongst local users and sharing of locally generated content with the international community. Without these services, locally generated content will be trapped in digital storage islands in developing regions across the world.

Chapter 8

VillageLink: Reaching Low Density Populations

8.1 Introduction

In Chapter 2 we highlighted some of the environmental challenges that hinder Internet penetration in rural developing regions; particularly those in Africa. One of the key obstacles in these rural areas is lack of wired infrastructure, such as copper cables, fibre optic links or cell phone base stations, due to high deployment cost and low population density which renders these techniques economically infeasible. Rural areas are also hard to reach via cheap license free solutions such as

WiFi, as these technologies, operating in 2.4 or 5 GHz bands, have a very limited connectivity range.

In the 50-800 MHz band, a large block of UHF and VHF frequencies has recently been freed up due to the analog to digital TV transition. This spectrum, called white spaces, promises to deliver an affordable means of providing wide area coverage. It is extremely attractive for rural areas as the propagation range is an order of magnitude higher than in the bands used by competing technologies. However, a distributed, resource-efficient solution for network organization, especially for spectrum allocation within a network, is needed for further proliferation of rural area white space deployments.

Channel allocation aims to assign one of the available channels to each of the network nodes. Traditionally, the issue has been expressed as the graph colouring problem where a color (channel) is assigned to a node so that network interference is minimized, and consequently the capacity is maximized. In a network operating over a small set of available frequencies, such as WiFi, channels generally do not exhibit significant differences in terms of propagation properties. White space networks, however, operate over a very large span of channels, and propagation properties can vary drastically over these channels. Channel assignment in such a network has to satisfy conflicting goals: maximize useful transmission by preferring channels with superior propagation properties, and minimize interference by favouring channels that propagate over a shorter radius.

Since we propose white space connectivity for impoverished regions, we concentrate on making our solution as cost efficient as possible. Therefore, we reuse the existing TV antennas already installed in even the most remote rural areas. Unfortunately, this further complicates the problem of channel allocation as these antennas exhibit uneven and unpredictable propagation behaviour over the wide white space spectrum. Any analytical solution that provides a clear picture of frequency quality becomes impossible, and a direct inference of propagation properties is needed.

In this paper we successfully address the above challenges and deliver the following contributions:

- We design a light-weight frequency profiling methodology to evaluate channel quality. We extend the IEEE 802.22 MAC scheme to enable autonomous white space base stations to coordinate frequency probing, and exchange information used for improvement of network performance.
- We develop a novel channel allocation method that assigns operating frequencies to base stations with the aim of minimizing the impact of interference over the useful signal levels in a network. Our frequency assignment

method is fully distributed; each link needs to know only the frequency profile of itself and its immediate neighbours. This information is readily available after the light-weight probing is performed.

We compile these contributions into a practical channel profiling and allocation scheme for wide area white space networks called VillageLink. We test VillageLink frequency probing mechanism on a long-distance software-defined radio white space link we deployed. Through simulations we evaluate VillageLink's channel allocation. We show that our frequency-aware channel allocation leads to up to twice as much network capacity than an alternative heuristic based on interference avoiding. In addition, VillageLink preserves fairness among users even when the density of the network is high. With its high performance, efficient resource usage and distributed nature, VillageLink is a highly practical solution for wide area white space coverage in rural areas.

8.2 Wide-area White Space Networks

White spaces represent a historic opportunity to revolutionize wide area wireless networking. White spaces not only deliver much greater communication range than Gigahertz frequencies, they also support non-line of sight communication, including transmission through vegetation and small obstacles, which makes them highly suitable for various terrain configurations. However, white space networks have to deal with unique peculiarities of transmission over a wide band of relatively low frequencies, and should enable license-free unplanned deployments in rural developing regions.

8.2.1 TV Spectrum availability

White spaces spectrum can also be used by primary users, typically television and wireless microphones. In order to understand the amount of available white space spectrum, we conducted a number of spectrum scans in rural and urban regions in South Africa, Zambia and the USA.

South Africa has 5 TV channels utilizing VHF and UHF TV bands; all channels are available in urban centres and a portion of these channels available in rural areas. Zambia only has 2 national TV bands with very limited coverage. The USA, in contrast, has anywhere between 10 and 25 terrestrial TV channels available depending on location. Figure 8.1 shows the spectrum scan of the lower portion of the UHF band used for television broadcasting. Urban areas in the USA have very limited white space available. Rural areas of the USA are comparable to urban areas in South Africa. Rural areas in Zambia and South Africa have an abundance of white space and are thus very well suited for rural connectivity solutions.

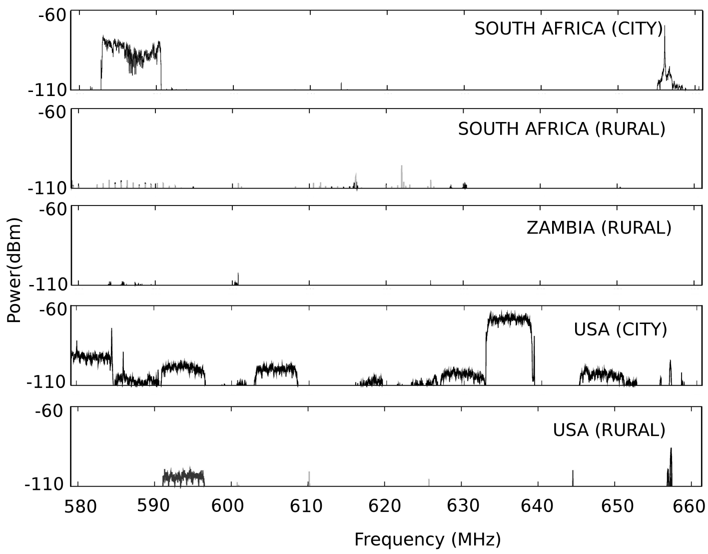

Figure 8.1: Analysis of available spectrum in the lower UHF band for urban and rural areas in South Africa, Zambia and the USA.

8.2.2 Wide band frequency selectivity

The variation in the free-space loss across a band is termed "dynamic range" and is calculated as follows:

$$D_{dB} = 20 \log \left(f_U / f_L \right)$$

where f_L and f_U are the lowest and the highest frequency in the band, respectively. In Table 8.1 we summarize the dynamic range of a number of traditional wireless systems. Free-space loss in a traditional wireless network, such as WiFi or GSM, is relatively uniform over the range of frequencies these networks operate on. The reason for low dynamic range in these networks lies in the fact that they either operate over a relatively narrow band of frequencies, such as 50 MHz for GSM and 80 MHz for 2.4 GHz WiFi, or they operate on high central frequencies where the difference between the lowest and the highest frequency diminishes, as is the case with 5 GHz WiFi. White spaces, however, operate on a wide band of low frequencies, and the difference in propagation between white space frequencies can be large. Note that the same issue does not arise in GSM (as well as 3G and 4G/LTE) networks, that can also operate on a wide range of frequencies (e.g. GSM850, GSM900, GSM1800). Unlike with white spaces, in these networks once the band selection is done the operation is restricted to a single relatively narrow range of channels.

Table 8.1: Dynamic range and fractional bandwidth of different wireless systems.

Technology	f_L**(MHz)**	f_U**(MHz)**	$\boldsymbol{D}$(dB)	$\boldsymbol{FB}$(%)
802.11 (2.4 GHz)	2412	2484	0.26	2.9
802.11 (5 GHz)	5170	5700	0.85	9.8
GSM900	935	960	0.23	2.6
White spaces	43.25	797.25	25.31	179

Besides wide dynamic range, white space links experience uneven fading due to antenna patterns. The fractional bandwidth (FB) for a frequency band, calculated as a ratio of operating bandwidth and the central frequency, determines how wideband an antenna should be in order to have the same gain over all frequencies with the band. From Table 8.1 we see that white spaces require significantly wider band antennas than GSM and WiFi. Such antennas are hard, if not impossible, to design. Consequently, white space links are highly prone to the effects of imperfect antennas.

To confirm this statement, we deployed a 3 km outdoor non-line-of-sight white space link. Each of the link nodes consists of a USRP2 radio and a dual core 2.4 GHz Pentium PC running GNUradio software. One node acts as a transmitter and sends probes at every 1 MHz over the white space spectrum. Another node, the receiver, scans the spectrum with 1 MHz spacing both with the transmitter turned off (baseline scan) and with the transmitter sending probes (signal scan).

Figure 8.2 shows the received signal strength across the UHF TV band in South Africa. Three TV stations were detected and probes did not occur at these frequencies. The received signal strength does not fall of monotonically with increasing frequency, which would be the case if only free-space loss determined the propagation loss. Instead, due to the antenna characteristics the propagation loss is non-uniform across the UHF band. Using the WIPL-D antenna modelling package we created a model of the deployed antenna. The results are shown in Figure 8.3. While an antenna with no surrounding structures has a more predictable gain patterns, when surrounding structures and antenna imperfections, such as bent elements, are introduced, the antenna gain pattern has far less pre-

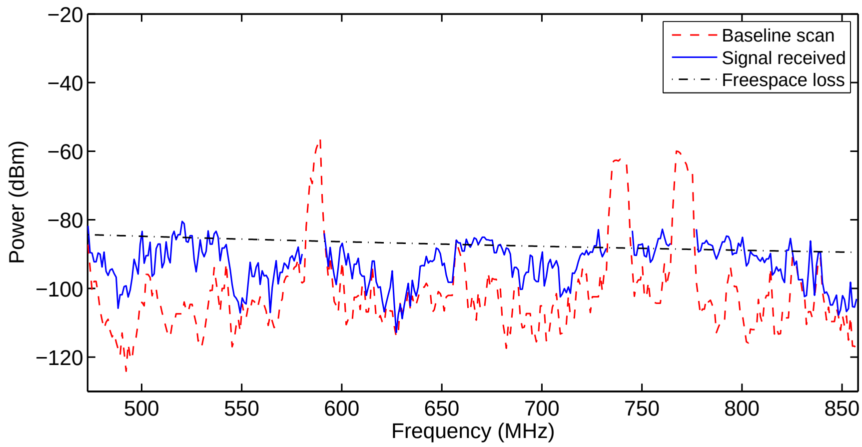

Figure 8.2: Analysis of received signal strength over the UHF band using 8 dBi yagi antennas at transmitter and receiver. The plots demonstrate that received signal strength is far more dependent on the antenna gain pattern than on attenuation due to free-space loss. This is confirmed by antenna frequency profiles of a number of TV antenna models in Figure 8.10.

dictability similar to what was seen in our received signal strength measurements. Predicting the type of TV antenna being used or the structures surrounding the antenna is not possible, which necessitates frequency probing in white spaces.

In addition to antenna effects, a part of frequency selectivity may stem from the environment and terrain effects. *Shadowing*, i.e. slow fading due to physical obstacles on the signal path would still be detected and accounted for with frequency probing. Unlike shadowing, *multipath*, which leads to rapid variation of propagation within a channel, cannot be captured through our current probing

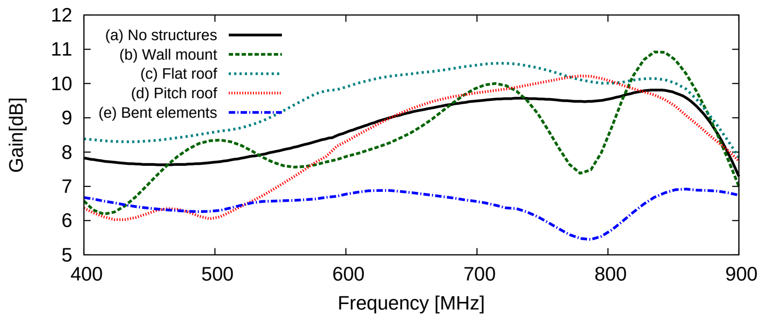

Figure 8.3: Antenna gain for yagi antenna used in white space link in the outdoor testbed. The plot shows the following scenarios: (a) Antenna with no nearby structures, (b) Antenna mounted on the side of a wall, (c) Antenna mounted on a pitch roof, (d) Antenna with imperfections due to bent elements.

method. However, channel allocation only requires knowledge of the average channel gain of the channel that is captured by frequency probing, and in Section 8.4 we devise a lightweight channel probing mechanism. (458 MHz - 860 MHz)

8.2.3 Channel assignment in white space networks

The problem of channel assignment in wireless networks is often expressed with graph coloring, where each color represents a different channel. For a link, one of the available central frequencies is assigned so that a goal, such as maximum throughput, is achieved. In the channel allocation literature on traditional wireless networks all colors are considered equal in terms of their propagation

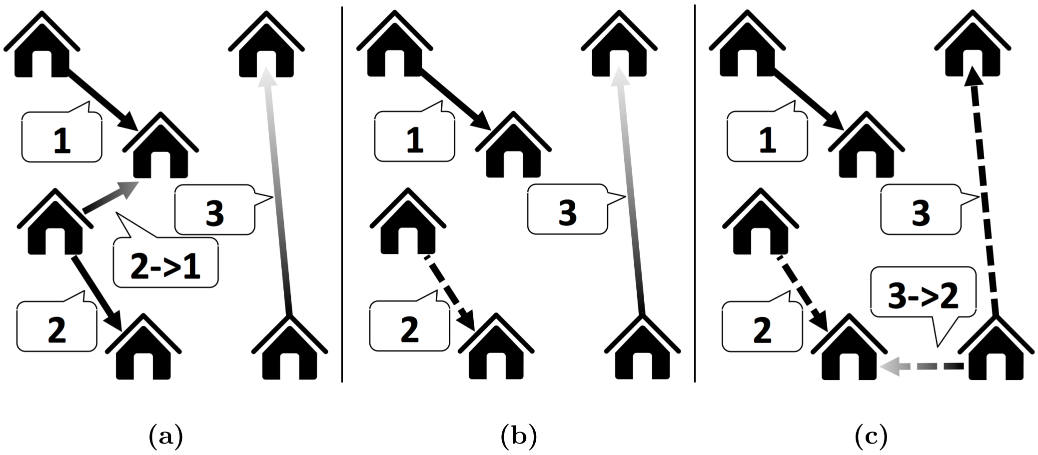

Figure 8.4: A simple example of the challenges of frequency assignment in white spaces. We want to establish links 1, 2 and 3; two channels represented by solid (inferior propagating frequency) and dashed (superior propagating frequency) colouring are given. In (a) all links operate on the inferior frequency, however there is interference from link 2 at link 1's receiver and link 3 is too long to be established on this frequency. In (b) the interference is resolved by switching the frequency for link 2. In (c) link 3 is establishment through assignment of the superior frequency. However, the superior frequency propagates further, thus interference at link 2's receiver is introduced.

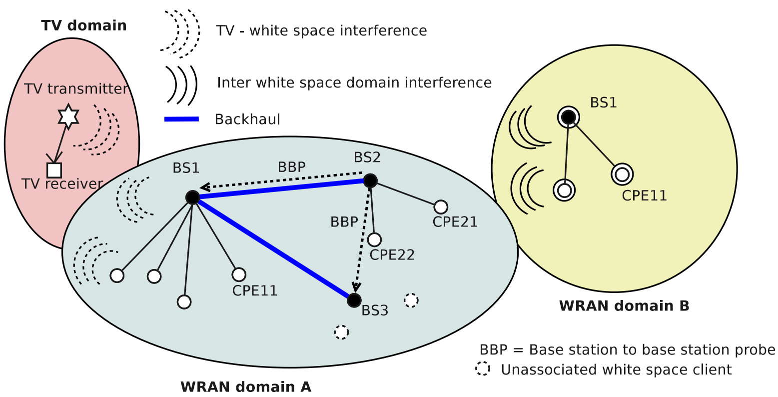

Figure 8.5: Layout of a targeted white space network showing interference scenarios between television and white spaces, and between white space networks in different domains. White space base stations within the same domain send base station to base station probes (BBPs) to calculate the channel conditions among themselves.

properties [159]. In white spaces, due to the wide dynamic range and antenna effects, the transmission range varies significantly among frequencies in the band (see Table 8.1 and Figure 8.2). Therefore, *selection of the operating frequency can impact the existence of a link itself*. This further complicates the problem of graph colouring, as now not all colours are equal. Figure 8.4 show one such example where a tradeoff between establishing links and avoiding interference is hard to achieve. In a white space network the color affects the graph structure, thus the existing approaches to frequency assignment are not directly applicable.

8.2.4 Network Architecture

The network scenario that describes the setting in which VillageLink will operate is given in Figure 8.5. In this paper, we consider wide-area white space networks that consist of individual base stations (BSs), each with a set of associated customer-premises equipment (CPE) clients. We term one such BS with its CPEs *a cell*. A BS and all the CPEs within a cell operate on the same channel[1]; thus, when considering channel allocation we use "BS" and "cell" interchangeably. All cells that are operated within the same administration are called a WRAN domain. The existence of TV transmission and other white space networks not in our service set reduces the number of channels available to the BSs within our domain. The aim of our work is to develop a channel allocation algorithm, as well as supporting structures such as a MAC layer and a frequency probing mechanism, so that the overall network performance within our white space WRAN domain is maximized.

We assume that base stations are connected with a back channel. This can be another white space radio tuned to a common channel that does not interfere with the set of channels available for base station to client connectivity. Moreover,

[1]We envision OFDMA channel sharing among the CPEs of a BS. Such an approach is mandated by IEEE 802.16 and IEEE 802.22 standards. We leave the details of subcarrier allocation as the future work, and in this paper concentrate solely on channel allocation at the BS level.

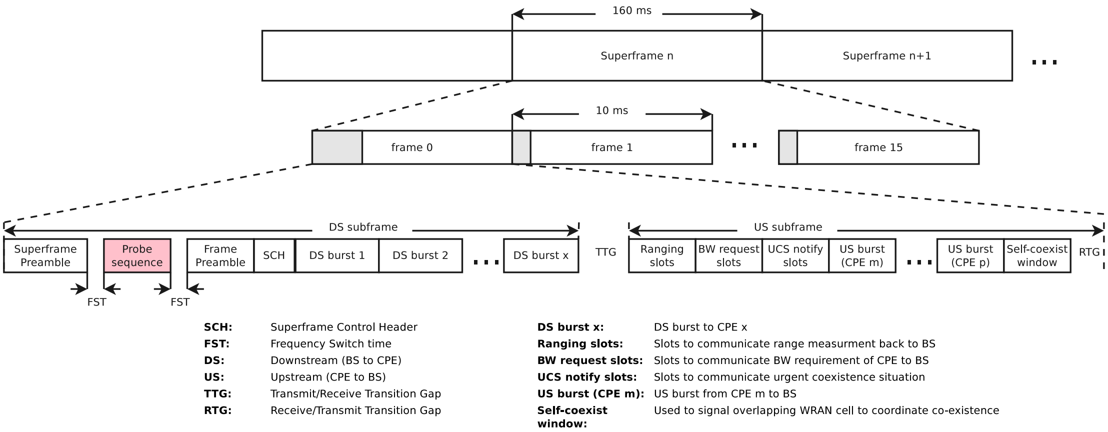

Figure 8.6: General superframe structure of 802.22 showing additional probe sequence required for frequency profiling. The 1024 bit probe sequence, in the worst case, increases overhead by 0.44%; this includes the 10μs Frequency Switch Time (FST)

because the amount of control data sent over the back channel is low, a VHF/UHF packet radio, or any other low bandwidth communication technology can be used.

8.3 Channel Probing and Medium Access for Wide-Area Networks

Channel probing is a necessary tool for propagation evaluation over a wide white space frequency range. Unfortunately, the existing MAC protocols proposed for wide area networks [141, 156, 43] do not explicitly support frequency profiling. The MAC protocol that most closely resembles our proposed system is IEEE

802.22. The protocol has built-in protection for primary users and mechanisms to move to new channels but has no built-in mechanism to choose from a set of available channels. It specifies that the channel may be chosen from the available list by an operator or by a "local routine". Thus, instead of rebuilding an entire MAC layer we propose to extend the 802.22 protocol to include a feature that performs frequency profiling on all available channels. In our frequency profiling scheme we measure the SNR value of the probe that was received at a base station using the previously measured power spectral density of channel with no probes and the power spectral density measured when a probe is present. To ensure consistency among measurements we assume that the nodes are static.

The MAC is organised into 160 ms superframes which consist of sixteen 10 ms frames. Each frame is divided into a downstream subframe and an upstream subframe; the size of these subframes depends on the amount of downstream and upstream data that needs to be sent between the base station and the clients. All base stations are equipped with a GPS and are able to synchronize their clocks to within 2 ns accuracy. Base stations start their first superframe at the start of a GPS minute cycle; this greatly simplifies inter-base station communication and scheduling. Final synchronization is carried out using the superframe preamble.

The current 802.22 MAC specification already has sophisticated mechanisms to detect other 802.22 WRAN domains using co-existence beacons, move into a time division multiplex mode in cases where base stations are forced to share the same channel, allow clients to scan for and associate with base stations, and carry out ranging between clients and base stations to account for propagation delays. However none of these mechanisms allow a base station to discover channel conditions between WRAN cells at all available channels.

Figure 8.6 shows the superframe structure of 802.22 with our modification, an additional probe sequence after the superframe preamble to allow for frequency profiling at different wireless frequencies. The probe sequence uses a 1024 bit PN sequence modulated using BPSK with 4 samples per symbol. The probe sequence incurs low overhead, using an additional 0.44% of the channel in the worst case We place the probe sequence after the superframe preamble to ensure we do not break any timing synchronization. Clients make use of the superframe preamble to synchronize any clock offsets. The probe will be transmitted on a probe channel, p_i, where $\boldsymbol{p} = (p_1, p_2...p_N)$, a set of N probe channels. The probe channel set is a subset of the complete set of available TV channels, $\boldsymbol{v} = (v_1, v_2...v_M)$ with $N < M$ after eliminating non-vacant channels. Each base station in the WRAN domain would first have consulted its spectrum database and scanned all the channels for primary users and other white space domains to ensure availability.

In order to perform frequency profiling between base stations on all available channels, a mechanism is required to coordinate probing timing, channel probe senders and listeners.When a base station is in a probing state, it sends a probe at the beginning of each superframe. It sequentially steps through the full white space TV channel set and only sends a probe if the channel is contained in the probe channel set for that base station. The entire scan takes $160M$ ms, where M is the total number of TV channels. If the probe channel is not contained in the probe set, the base station does not send any transmission in the probe sequence slot. Base stations maintain their own probe sets as they may each generate interference to primary users on different sets of frequencies. We chose a mechanism where we step through the entire set of white space channels with interspersed quiet periods on non-vacant channels to avoid needing to maintain full consistency of all probe channels amongst all base stations and associated clients.

In order to coordinate probe transmission and reception, we propose a token approach in which a base station only transmits probes when it has a probe token and listens for probes when it does not have a token. A breath-first traversal of a spanning tree of the graph is used to ensure that a token traverses the graph of base stations when the back channel forms the edges. Once the base station has finished probing across the TV channel set, it sends the token to the next base station using a traversal algorithm; the probing process is completed when all base stations have used the token. Probes can take place when base stations

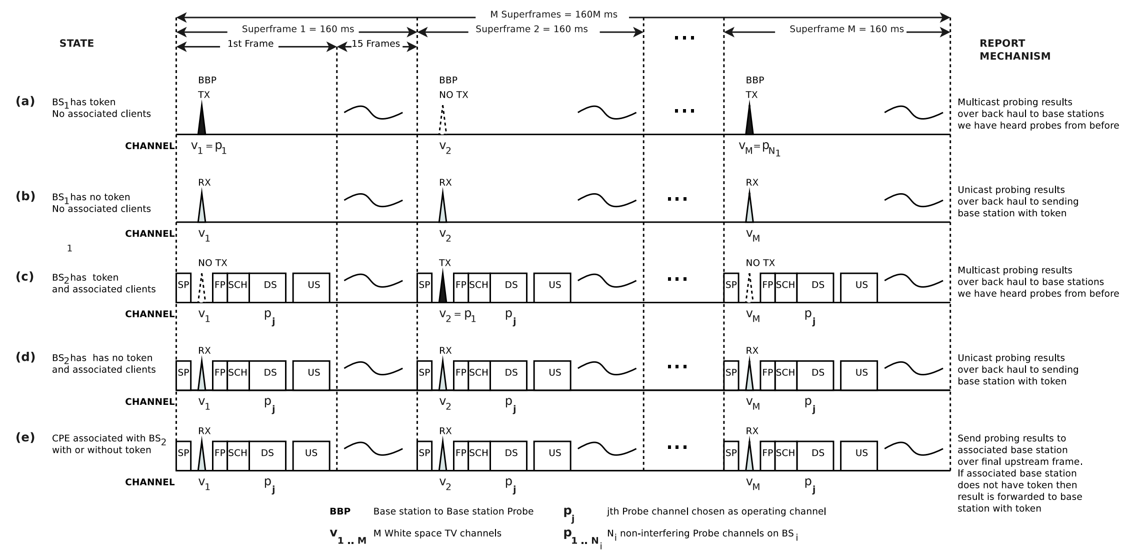

Figure 8.7: Timing diagram showing the BBP probes sent by the base station and received by other base stations and clients.

are in the initialization phase and have not chosen an operating channel, or when they are in an operating phase and are communicating with associated clients and a new base station is added, for example. A request to probe is broadcast by the base station wishing to initiate probing in the domain.

The 802.22 specification makes use of clients to sense for primary users and extend the sensing coverage area. We propose to use a similar notion when listening for probes. Clients of one base station experience interference from all other base stations. To account for this interference clients can be instructed to listen for probes from base stations with whom they are not associated. Frequency profiling results for clients are sent back to the associated base station on the final upstream frame once the client has listened on the full set of white space channels. The maximum SNR value of a received probe heard at a base station and its associated clients is used to incorporate the worst case effect on the system. These SNR values from each of the cells that received the probes are unicast on the back channel to the sending base station. Results received at sending base stations are distributed to neighbouring base stations, where two base stations are defined as neighbouring if a probe can be exchanged between them on at least one frequency.

Figure 8.7 shows a number of timing diagrams for different base station and client states. Figure 8.7(a) shows a base station in the initialization phase in which it has a token. In this case only probes are sent by the base station. Figure 8.7(c) also shows a base station with a probe token but in this case the probes are interspersed within standard 802.22 frames as the base station is actively communicating with clients. Quiet periods, where no probe is sent, are different in Figures 8.7(a) and 8.7(c) as they effect primary users on different channels. Figures 8.7(b) and (d) show a base station receiving probes and measuring the channel condition between the sending base station and itself. Figure 8.7(e) shows a client receiving probes from a base station with or without a token. In the case of a client hearing a probe from a base station it is not associated with, the frequency profile matrix is sent back to its associated base station and then forwarded on the back channel to the sending base station with the token.

Once the probing process is completed each BS i has information on signal propagation at different frequencies: 1) within its own cell, obtained through aggregation of probing results from the cell's CPEs, 2) between i and each of the BSs j that heard probes from i, and 3) within cells that are served by each of the neighbouring BSs j. We note that power failures are expected in rural areas [160], and should they be experienced, the base station can power up to its previous known state stored in non-volatile memory as the frequency profile matrix, which contains the frequency profile between all base stations at all available channels.

8.3.1 Calculating probe SNR

Each master probe is simply a pseudo-noise (PN) sequence modulated with DBPSK. The client calculates the average power measured, P_{avg} over the probe listen window.

$$P_{avg}(dB) = 10 \log \left(\frac{1}{N} \sum s(n)^2 \right) + CF$$

where $s(n)$ is the signal received, n is the sample number, N is the total number of samples and CF is the correction factor, which is calculated by calibrating the receiver. Average power is a low complexity ($O(N)$) calculation and the cognitive radio can carry this out in real time using a cumulative average.

The SNR of the probe can then be calculated by using a previously known noise level from an initial scan when no probes are present.

$$SNR(dB) = P_{avg}(dB) - N_{avg}(dB)$$

From the measured SNR we can extract the channel gain:

$$H = \frac{SNR \cdot N_0 \cdot W}{P} \tag{8.1}$$

where N_0, W and P denote the noise constant, channel width and the transmission power, respectively.

8.4 Channel Allocation

In this section we devise a distributed channel allocation algorithm that uses information obtained through frequency profiling (Section 8.3) and does not incur channel switching overhead typical for other allocation schemes. Our approach is based on the annealed Gibbs sampler, a technique that can help us minimize a target function in a distributed way. In the next subsection we present the basics of Gibbs sampling. An interested reader can find more details about Gibbs sampling in [24]. We then cast our problem to the Gibbsian framework and sketch the channel allocation algorithm.

8.4.1 Gibbs Sampling

The Gibbs sampler is a Markov chain Monte Carlo technique for obtaining random samples from a multivariate probability distribution. The sampler is useful in situations where the joint distribution is unknown or difficult to sample, but the conditional distributions of variables are known and easy to sample. In a nutshell, the Gibbs sampler that draws samples from a multivariate probability distribution $p(x_1, ..., x_N)$ works as follows:

- Initialize all variables $x_1, ..., x_N$ to (random) starting values $x_1^0, ..., x_N^0$.
- For each sample $j = 1..k$, sample each variable x_i from the conditional distribution $p(x_i | x_1^j ..., x_{i-1}^j, x_{i+1}^{j-1}, ..., x_N^{j-1})$ to obtain x_i^j.

After the above process is finished, we are left with $x_1^j, ..., x_N^j; j \in [1..k]$ samples from the joint distribution p.

We can solve the channel allocation problem through Gibbs sampling, if we obtain the samples from a multivariate probability distribution that:

1. is related to overall network performance.
2. depends on the selected operating channel of each of the base stations.
3. isolates the impact of each of the base stations on the total optimization function.

4. can be calculated in a distributed way and sampled independently at each of the base stations.

5. favors states that lead to maximum network performance.

In the following section we develop a network performance metric that can be used as a basis for a probability distribution that satisfies the above demands.

8.4.2 Network Performance Metric

Traditionally, the goal of a channel allocation protocol is to assign available channels to BSs so that the total network capacity is maximized. The capacity $C_i(c_i)$ of a single cell operating on the channel c_i is:

$$C_i(c_i) = \sum_{k \in \mathbb{K}_i} W_k \log\left(1 + SINR_{ik}(c_i)\right)$$

where $\mathbb{K}_i$ is the set of CPEs within the cell, W_k is the width of a part of the channel c_i used by CPE k, and $SINR_{ik}(c_i)$ is the signal to interference plus noise ratio at the CPE k. We approximate the presence of all clients within the cell with a single *virtual CPE* with an SINR value $SINR_i(c_i) = \sum_k SINR_{ik}(c_i)/|\mathbb{K}|$. The cell capacity is now:

$$C_i(c_i) = W \log\left(1 + SINR_i(c_i)\right)$$

where W is the full channel width, essentially a sum of all W_k as a cell operates in an OFDMA mode. This approximation hides channel distribution within the cell and helps us concentrate on inter-cell interaction.

If we consider a network with N cells, with a given channel assignment $\mathbf{c} = (c_1, c_2, ..., c_N), c_i \in \mathbb{C}$, where $\mathbb{C}$ is the set of available channels, the total network throughput is a sum of all individual capacities at their respective allocated channels:

$$C(\mathbf{c}) = \sum_i C_i(c_i) = \sum_i W \log\left(1 + SINR_i(c_i)\right) \tag{8.2}$$

A single BS's decision on the operating channel changes the interference level at all its neighboring BSs. In the above equation the interference is accounted for in the SINR, which is embedded within the logarithmic function. Thus, the impact of a single BS on the total sum is hard to isolate, and the total capacity is not a suitable metric for distributed computation using Gibbs sampling. Centralized optimization using known polynomial complexity techniques, such as linear programming, is not directly applicable either, since the target sum involves non-linear factors and discrete variables.

One of the ways to circumvent this is to revert to a tighter problem formulation that prevents interfering base stations from concurrent transmission [110]. While this can be enforced in a network that employs carrier sensing and collision avoidance, in our setting long distances between base stations render such coordination inefficient [141]. In addition, allowing some interference often yields more capacity than restricting concurrent transmissions [123]. Another approach is to modify the optimization function and instead of maximizing capacity concentrate on minimizing total network interference [96, 139]. This approach is attractive for networks, such as WiFi, where these two goals are essentially interchangeable. In a white space setting, where available channels can differ drastically in terms of their propagation properties, *a channel allocation that leads to minimal interference may not necessarily lead to maximum capacity.*

We propose a novel network performance metric – *Cumulative interference plus noise to signal ratio (CINSR)* – a sum of inverse of *SINR* experienced at each of the cells. *CINSR* can be seen as the overall ratio of the impact of harmful factors, noise and interference, to the beneficial one, received signal strength. Thus, our goal is to minimize it. Compared to metrics such as the total capacity or the overall level of interference, *CINSR* takes into account the frequency diversity that exists in white space networks, and allows distributed performance optimization with Gibbs sampling:

$$CINSR(\mathbf{c}) = \sum_{i=1}^{N} \frac{1}{SINR_i(c_i)} \tag{8.3}$$

$$= \sum_{i=1}^{N} \frac{N_0 W + \sum_{j=1..N}^{j \neq i} ch(i,j) PH_{ji}(c_i)}{PH_i(c_i)} \tag{8.4}$$

The first term in the numerator within the above sum is the thermal noise (a product of the channel width W and the noise constant N_0), whereas the second term is the sum of interference experienced at cell i, and originating from all other base stations that transmit at the same channel. Interference from a single source is a product of P - the transmission powerand $H_{ji}(c_i)$ – the propagation gain from base station j to cell i on channel c_i. The function $ch(i, j)$ is equal to 1 if i and j operate on the same channel and, otherwise it is equal to 0. The denominator in the above equation is the average signal strength received by the clients of the BS i that transmits with power P. The average channel gain towards the clients is denoted by $H_i(c_i)$.

We now isolate the impact of a single BS i on $CINSR(\mathbf{c})$ and term it *local CINSR*:

$$CINSR_i(\mathbf{c}) = \frac{N_0 W}{PH_i(c_i)} \tag{8.5}$$

$$+ \sum_{j \neq i} ch(i, j) \left(\frac{PH_{ij}(c_i)}{PH_i(c_i)} + \frac{PH_{ji}(c_i)}{PH_j(c_i)} \right) \tag{8.6}$$

Information needed for $CINSR_i(\mathbf{c})$ calculation, namely $PH_i(c_i)$, $PH_{ji}(c_i)$, $PH_{ij}(c_i)$ and $PH_j(c_i)$, is available locally at BSi, through channel probing described in Section 8.3.

8.4.3 The Gibbs distribution

The Gibbs distribution associated with the function *CINSR* and a positive temperature T is the probability distribution on c^N (the combined channel state space of all BSs) defined as:

$$\pi(\mathbf{c}) = \frac{e^{-\mathcal{CINSR}(\mathbf{c})/T}}{\sum_{\mathbf{c}' \in c^N} e^{-\mathcal{CINSR}(\mathbf{c}')/T}} \tag{8.7}$$

The above distribution is of a special interest as it favors states in which *CINSR* is low. In addition, the channel selected by BSi is independent of all non-neighboring BSs and the distribution fulfils all the conditions listed Section 8.4.1.

The Gibbs sampler draws a sequence of samples from the above distribution by having each of the BSs i independently sample its local Gibbs distribution $\pi_i(\mathbf{c})$:

$$\pi_i(\mathbf{c}) = \frac{e^{-\mathcal{CINSR}_i(c_i,(c_j)_{j\neq i})/T}}{\sum_{\mathbf{c}' \in c^N} e^{-\mathcal{CINSR}_i(c_i',(c_j)_{j\neq i})/T}} \tag{8.8}$$

and transitions to the sampled local state, converging to the stationary distribution $\pi(\mathbf{c})$ (see Section 8.4.5).

Distribution $\pi(\mathbf{c})$ favors low *CINSR* states when the temperature is low. While our goal is to minimize *CINSR*, by keeping the temperature low we risk getting stuck in a local minimum early in the process. *The annealed Gibbs sampler* introduces a slow decrease of temperature T to zero according to *a cooling schedule.* Therefore, in the beginning the probability of exploring a wide range of states is high, and as the time goes to infinity, the procedure converges to the minimum *CINSR* state (Theorem 3, Section 8.4.5). The choice of the schedule impacts the convergence speed, and we experiment with two commonly used schedules in the evaluation section.

8.4.4 Channel Allocation Algorithm

Algorithm 1 VillageLink channel allocation – distributed

{Executed at the base station i}
while $t < t_{end}$ **do**
  $T = f(T_0, t)$ {f - schedule, T_0 - starting temperature}
  for all channel $c_i' \in \mathbb{C}$ **do**
    $\mathbf{c}' = (c_1, c_2, ..., c_i', ..., c_N)$
    Calculate $CINSR_i(\mathbf{c}')$
  end for
  for all channel $c_i' \in \mathbb{C}$ **do**
    $\mathbf{c}' = (c_1, c_2, ..., c_i', ..., c_N)$
    Calculate $\pi_i(\mathbf{c}')$
  end for
  Sample a random variable according to the law π_i and choose the next channel of the BS i accordingly.
  Send information about the newly selected channel to i's neighbors.
end while
Switch the network interface to the last selected channel.

The Algorithm 1 is executed at each of the base stations. The temperature falls off with time, ensuring that the Gibbs sampler converges towards the global minimum of *CINSR*. The starting time for all the base stations has to be loosely aligned, and can be achieved through a standard synchronization scheme such as NTP.

Compared to some other distributed channel allocation schemes [96, 120], Algorithm 1 has an attractive property that no channel switching is needed until the convergence. To see why note that the calculation of the local *CINSR* is done after the probing process, and during the algorithm run the only variable parameter is $ch(i, j)$. At BS i this parameter can be updated irrespective of the actual operating channel of BS j. In every step a BS decides on its current channel and sends the decision to its neighbors, who then update their $ch(i, j)$ tables. Once the cooling schedule is completed base stations switch to their channel of choice (line 15 in Algorithm 1). This greatly speeds up the convergence, as the channel allocation process is not limited by the channel switching time.

8.4.5 Algorithm convergence

Theorem 1. *The Gibbs distribution π (equation 8.7) has the Markov property.*

Proof. A Gibbs potential V associates a real number $V_\Gamma(\mathbf{s})$ with each subset Γ of a set S. The potential is determined by the state $\mathbf{s}$ of the nodes in Γ and is defined as zero if Γ is not a clique. *An energy function* $\mathcal{E}(\mathbf{s})$ maps each of the graph states to a real number. We say that the energy function *derives from the potential V* if:

$$\mathcal{E}(\mathbf{s}) = \sum_{\Gamma} V_{\Gamma}(\mathbf{s}) \tag{8.9}$$

where the summation goes over all subsets of the set S. According to Theorem 2.1, pg. 260 in [24] the Gibbs distribution where the energy derives from a Gibbs potential is a Markov random field. This is a part of a more specific Hammersley-Clifford theorem that states necessary and sufficient conditions under which a probability distribution can be represented as a Markov random field (Theorem 2.2, pg. 262 in [24]).

The energy function that we use to construct the Gibbs distribution in equation 8.7 is *CINSR(s)*. We can represent *CINSR* as a sum of local impact of cliques of the graph of base stations $\mathcal{A}$. *CINSR* then takes the form described by equation 8.9 and can be used as the energy function for Gibbs sampling:

$$\begin{aligned} CINSR(\mathbf{c}) = & \sum_{i \in \mathcal{A}} \frac{N_0 W}{PH_i(c_i)} + \\ & + \sum_{\{i,j\} \in \mathcal{A}} ch(i,j) \left(\frac{PH_{ij}(c_i)}{PH_i(c_i)} + \frac{PH_{ji}(c_i)}{PH_j(c_i)} \right) \\ & = \sum_{\mathcal{B} \subset \mathcal{A}} V_{\mathcal{B}}(\mathbf{c}) \end{aligned}$$

Here V denotes the Gibbs potential. The potential is defined for all subsets $\mathcal{B}$ of the set of base stations $\mathcal{A}$ as:

$$V_{\mathcal{B}}(\mathbf{c}) = \begin{cases} N_0 W / PH_i(c_i) & \text{if } \mathcal{B} = \{i\} \\ ch(i,j)\left(\dfrac{PH_{ij}(c_i)}{PH_i(c_i)} + \dfrac{PH_{ji}(c_i)}{PH_j(c_i)}\right) & \text{if } \mathcal{B} = \{i,j\} \\ 0 & \text{if } |\mathcal{B}| \geq 3 \end{cases}$$

Note that the potential is non zero only for cliques of size one and two. Thus, energy $CINSR(\mathbf{c})$ derives from the Gibbs potential and, consequently π is a Markov random field. □

Theorem 2. *For a fixed network of N base stations, each running a Gibbs sampler over its local Gibbs distribution $\pi_i(\mathbf{c})$, channel allocation converges in variation*[2] *towards the Gibbs distribution π.*

Proof. The process can be described as a Gibbs sampler on a finite state homogeneous Markov chain (Theorem 1), for which the Gibbs distribution (equation 8.7) is the invariant probability measure. Example 6.5, pg. 288 in [24] proves that such a sampler converges in variation to the target distribution, and that the convergence takes place with geometric speed. □

Note that direct sampling of the capacity (equation 8.2) does not provide any guarantees on the performance as the capacity equation cannot be transformed

[2]Convergence in variation describes convergence of an array of samples to a probability distribution and is defined in [24], pg. 128.

to an energy function that derives from the Gibbs potential. Thus, we develop *CINSR*.

Theorem 3. *For a fixed network of N base stations implementing Algorithm 1, channel allocation converges in variation towards a limit distribution that only puts positive probability mass on the states of minimum global energy.*

Proof. The proof is analog to Example 8.8, pg. 311 in [24]. The annealed Gibbs sampler converges according to a strongly ergodic nonhomogeneous Markov chain: it converges in variation to a limit distribution that only puts positive probability mass on the states of minimum global energy. Conditions that the cooling schedule has to satisfy in order for convergence to happen can be found in [65]. □

8.5 Evaluation

The VillageLink system consists of our frequency profiling method built on top of the 802.22 MAC protocol, and the channel allocation algorithm based on Gibbs sampling. Experimental evaluation of such a system is challenging due to the need for a wide area outdoor deployment. In addition, off-the-shelf 802.22 equipment is not yet commercially available, and software defined radio platforms cannot support the synchronization that the MAC protocol requires [128]. Therefore, we evaluate our protocol in a simulated setting. However, the initial experimental investigation of channel probing and frequency selectivity in white spaces, presented in Section 8.2.2, was performed on a 3 km outdoor link.

8.5.1 Simulation Setup

For a comprehensive evaluation of the channel allocation algorithm, we rely on a Matlab-based custom simulator. The simulator allows us to scale our experiments over a number of cells, and to model different network layouts. We explicitly take into account high variability of signal propagation in the white space band by modeling propagation with the Friss transmission equation:

$$P_r = P_t + G_t + G_r + 20 \log \left(\frac{\lambda}{4\pi R} \right)$$

where $\lambda, R, P_r, P_t, G_t$ and G_r are the wavelength, distance between antennas, received power, transmitted power, transmitter antenna gain, and receiver antenna gain, respectively. Antenna gains depend on specific devices used and their orientations. Earlier, in our outdoor testbed, we confirmed that frequency dependence of antenna gain is the most dominant factor that leads to the frequency diversity in white spaces (Figure 8.2, Section 8.2), thus we model antenna effects in detail.

We use publicly available antenna models[3] and the Numerical Electromagnetics Code (NEC)[4] antenna modeling software to examine propagation over different

[3]http://www.hdtvprimer.com/ANTENNAS/comparing.html
[4]www.nec2.org

frequencies with different antennas. Figures 8.8b and 8.9b show the radiation patterns seen from the the center frequency (598 MHz) of the white space band for two different antennas. In Figure 8.10 we plot frequency dependence of antenna gain. We found that the shape of the antenna pattern does not change significantly for different frequencies. The gain, on the other hand, changes significantly and unpredictably, as seen in Figure 8.10. Thus, in the simulations we use the antenna pattern shape of the center frequency to account for antenna orientation, and we use the full gain over frequency diversity.

All base stations in our simulations use the Yagi antenna from Figure 8.8, as this antenna exhibits the best performance of all the antennas that were modeled. In our simulation we assume clients make use of existing TV antennas used to receive terrestrial TV broadcast signals. Operators have no control over the variety of antennas used by clients and we randomly select antennas from a set of 17 possible client TV antennas ranging from outdoor Yagi antennas with a gain of 15dBi to simple indoor loop antennas with a gain of 3dBi.

We run our experiments over a white space band from 443 MHz to 875 MHz as the antenna models we use perform reasonably well within this range. The band is divided into 36 TV channels, each 6 MHz wide, with a 6 MHz guard band between adjacent channels. In all the experiments we simulate a $100km \times 100km$ field with random BS placement and random antenna orientation. Each of the BSs has a

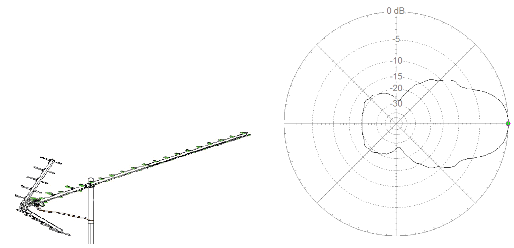

(a) (b)

Figure 8.8: Wineguard PR9032 UHF Yagi/corner reflector antenna used as a base station antenna in our evaluation. Showing (a) the antenna design and (b) its radiation pattern seen from the top of the antenna.

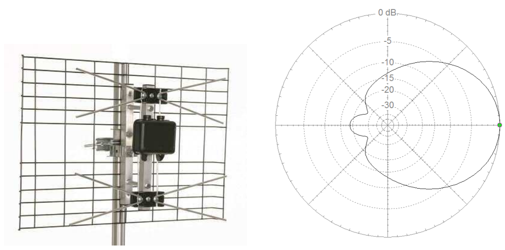

(a) (b)

Figure 8.9: AntennasDirect DB-2 2-Bay UHF antenna; one of the client antennas used for the evaluation. Showing the antenna in (a) and its radiation pattern seen from the top of the antenna in (b).

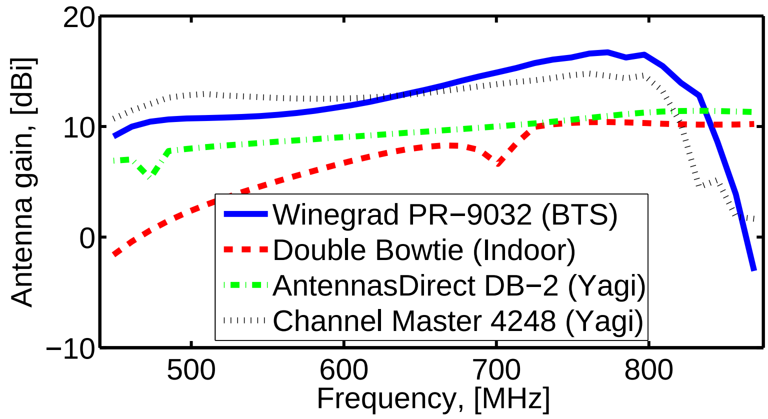

Figure 8.10: Antenna profiles of four of the antennas used in our evaluation. One of the profiles, Wineguard PR-9032, corresponds to the BS antenna and, the other three, to client antennas.

single associated virtual client at a distance uniformly picked from 0.2 km to 20 km and with its antenna pointed directly towards the BS. We also simulate a TV station that covers a part of the field with its signal and occupies two adjacent channels.

8.5.2 Channel Allocation Convergence

We simulate *Algorithm 1* behavior in a network of ten base stations and five white space channels that are available for communication. We are interested in the algorithm convergence under different Gibbs sampling parameters. We experiment with two common cooling schedules:

- logarithmic: $T = T_0 / \log(t + 2)$, proposed in [56].

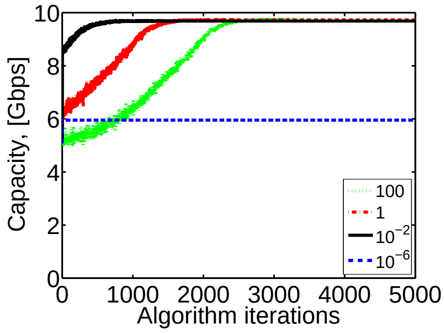

(a)

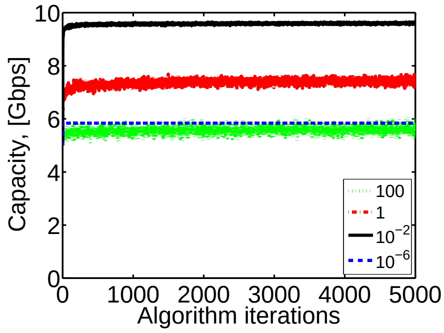

(b)

Figure 8.11: Algorithm convergence with the (a) exponential, and (b) logarithmic cooling schedule. Different lines correspond to different starting temperatures.

- exponential: $T = T_0 \alpha^t$, proposed in [99].

where T denotes the temperature at time t, T_0 is the starting temperature, and α is a real number between zero and one; we empirically find value 0.995 to work well in our experiments.

The selection of the starting temperature is important for proper annealing. In Figure 8.11 we plot total network capacity achieved with the two schedules and four different starting temperatures for each. Each point in the graphs is an average over 100 runs. The impact of the starting temperature is clearly visible: the higher T_0 is, the more time it takes for the algorithm to converge. At the same time, higher temperatures ensure exploration of a large part of the solution space, and generally lead to a better solution. We can also see that $T_0 = 10^{-6}$ does not result in any variation of capacity as the algorithm progresses – the sampler is “frozen” and BSs will stick to the initial channel allocation without exploring the full solution space. There is a trade-off, dictated by the starting temperature, between the convergence time and the assurance that the optimal value will be found. In the rest of the evaluation section we fix the starting temperature to 1, a value that allows full exploration of the solution space and converges in a reasonable amount of time.

We observe much faster convergence with the exponential schedule, that converged in all but one case ($T_0 = 10^{-6}$). The logarithmic schedule did not converge in 5000 iterations for $T_0 = 100$ and $T_0 = 1$. In the rest of this section we rely exclusively on the exponential schedule.

8.5.3 CINSR as a Performance Metric

To confirm that *CINSR* is a good choice for the network performance metric, we compare it with an alternative – overall interference and noise in the network – which is often used as a metric in channel allocation algorithms [96, 139].

The total network interference and noise is defined as:

$$I(\mathbf{c}) = \sum_{i=1}^{N} \left(N_0 W + \sum_{j=1..N}^{j \neq i} ch(i,j) P H_{ij}(c_i) \right) \tag{8.10}$$

The impact of a single BS on the sum is defined as *the local interference*:

$$I_i(\mathbf{c}) = N_0W + \sum_{j=1..N}^{j \neq i} ch(i,j)\,(PH_{ij}(c_i) + PH_{ji}(c_i)) \tag{8.11}$$

We modify the Gibbs distribution (equation 8.7) to include $I(\mathbf{c})$ instead of $CINSR(\mathbf{c})$, and the local Gibbs distribution (equation 8.8) to include $I_i(\mathbf{c})$ instead of $CINSR_i(\mathbf{c})$. The necessary conditions for the Gibbs sampler convergence still hold, and we apply an algorithm analogous to Algorithm 1.

Note that, defined this way, the interference function still uses the results of channel probing, yet it does not account for the balance between well propagating channels that are preferred by the CPEs and inferior channels that minimize inter-cell interference.

Channel under-provisioning

In the first scenario we simulate a network with a number of contending BSs higher than the number of available channels. This can be the case in the urban developed world, for example. We put 50 cells in the same $100km \times 100km$ region. We experiment with a varying number of available channels. The total network capacity is plotted in Figure 8.12a. Each point represents an average value of 20 runs of the algorithm with a different metric, *Gibbs CINSR* or *Gibbs Interference*, over the same topology.

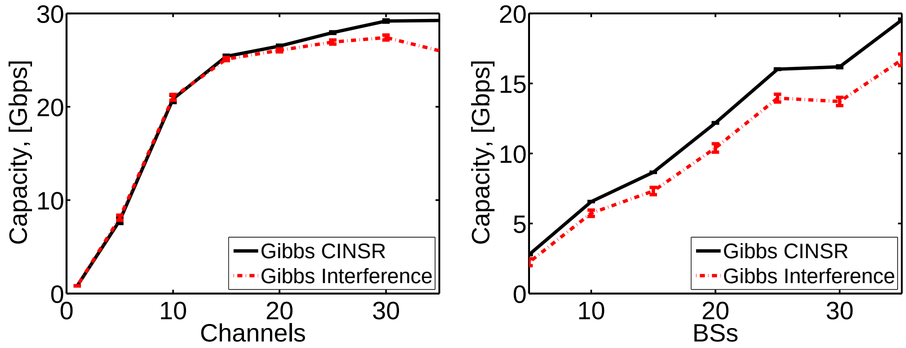

(a) Channel under-provisioning.

(b) Channel over-provisioning.

Figure 8.12: Comparison of the total network capacity achieved with CINSR and Interference metrics. We simulate under-provisioned and over-provisioned number of channels with respect to the number of base stations in the network.

When multiple cells operate on the same frequency the network is in a low SINR mode, and capacity can be increased by interference minimization. From Figure 8.12a we see that the two versions of the Gibbs sampler perform equally well with a small number of available channels. As we increase the amount of available spectrum, BSs have more freedom to operate at different channels with minimal interference. Therefore, frequency-dependent performance of CPEs associated with the BSs becomes an important factor that impacts total capacity. Since this factor is not accounted for in equation 8.10, this version of the Gibbs sampler

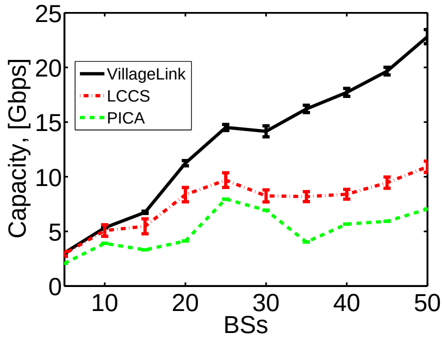

(a) 10 available channels.

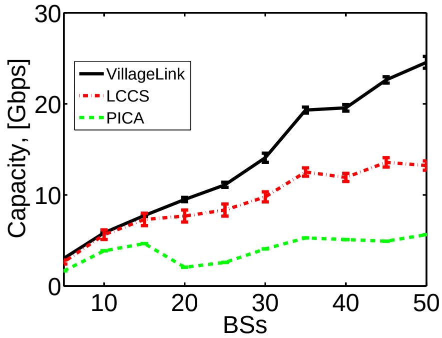

(b) 15 available channels.

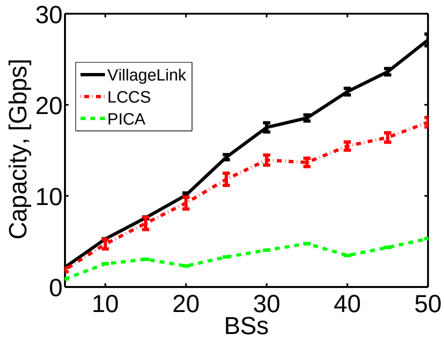

(c) 20 available channels.

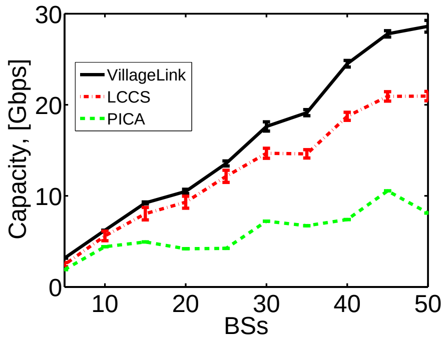

(d) 25 available channels.

Figure 8.13: Total network capacity with varying number of channels and base stations.

results in a channel allocation that delivers less capacity than the version that uses *CINSR*.

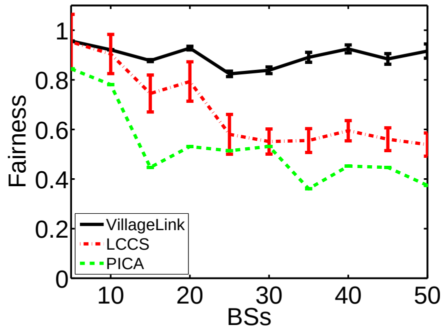

(a) 10 available channels.

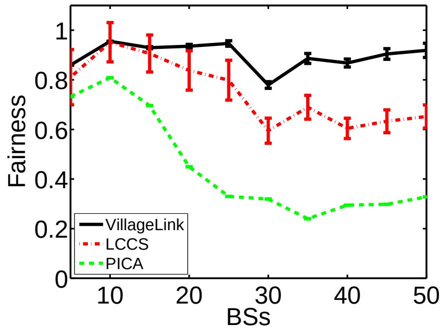

(b) 15 available channels.

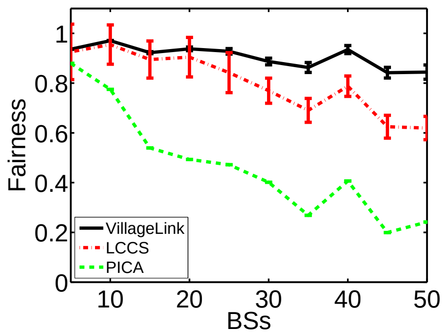

(c) 20 available channels.

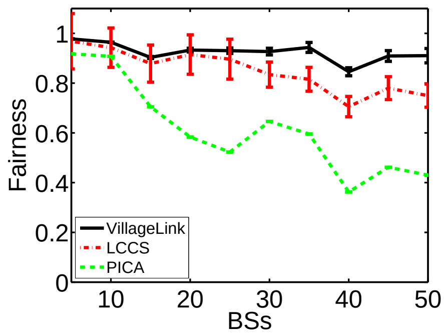

(d) 25 available channels.

Figure 8.14: Fairness with varying number of channels and base stations (the closer the fairness index value is to one - the better).

Channel over-provisioning

We now fix the number of available channels to 36 and compare the performance of the two versions of the algorithm with the number of BSs varying from 5 to 35. The total network capacity is plotted in Figure 8.12b. Each point represents an average value of 20 runs of the algorithm (*Gibbs CINSR* or *Gibbs Interference*) over the same topology.

When the number of channels is greater than the number of BSs there is more than one allocation that leads to minimal interference. However, not all of the allocations are favored by the CPEs. Through the factor $H_i(c_i)$ CINSR accounts for the frequency dependent intra-cell preferences, and assigns channels that maximize capacity within each of the cells. The results presented here point out that *channel allocation in white spaces remains important even in rural areas where the channel availability is high* [70].

8.5.4 Comparison to alternative channel allocation methods

Channel allocation is a difficult problem to solve in a distributed setting. Heuristics are often used instead of a rigorous solution and we compare our approach with:

- *Least congested channel search (LCCS)* - a heuristic where each of the BSs individually scans for a channel with the least number of other BSs assigned to it [120].

- *Preferred intra-cell channel allocation (PICA)* - in this greedy method each of the BSs selects the channel for which it observes the highest channel gain towards its own CPEs ($\arg\max_{c_i} H_i(c_i)$).

These heuristics optimize a non-submodular capacity function in a greedy manner, therefore may settle for a solution that is arbitrarily far from the optimal. VillageLink's convergence to the states of minimum *CINSR* is proven in Section 8.4.5. We compare the experimental behavior of different solutions in a number of scenarios encompassing various numbers of BSs and available white space channels. We run each of the algorithms 100 times in each of the scenarios.

Total network capacity

In Figure 8.13 we plot the total network capacity as we increase the number of cells in the system from 5 to 50. To ensure consistency among points in the graph, we do not generate a new topology every time we increase the number of cells, but add randomly placed cells to the existing topology. Each of the topology sequences are evaluated in environments with 10, 15, 20 and 25 available channels. We plot average values and two standard deviations (represented by error bars) for each data point.

VillageLink performs better or equal to the alternatives in all scenarios. The benefits of frequency-probing based channel allocation grow with the number of cells. In some scenarios, such as 50 BSs - 10 channels and 50 BSs - 15 channels, VillageLink delivers twice as much capacity as the next best alternative, LCCS. A comprehensive comparison of LCCS and PICA could unravel the importance of two conflicting goals in channel allocation: minimizing interference and maximizing intra-cell capacity in isolation, and is left for future work.

Fairness

In Figure 8.14 we plot the Jain fairness index [85] for cell capacity with channel allocations determined by VillageLink, LCCS and PICA. We plot average values and two standard deviations (represented by error bars) for each data point. Although we designed VillageLink as a method to optimize total network capacity, it also ensures a remarkably fair allocation of resources. As the number of cells grows, the fairness of VillageLink is more pronounced as it stays close to 1 while the fairness indices of PICA and LCCS drop.

8.6 Related Work

Efforts to provide broadband connectivity to remote rural regions with low-cost unlicensed options, such as modified WiFi, have been proposed in the last decade [141, 156]. While numerous rural WiFi deployments provide useful general guidelines for wide-area coverage, the propagation characteristics in white spaces are drastically different than in the WiFi bands, and networking protocols have to be reconsidered for the new spectrum. IEEE 802.22 [43] is a standardized protocol for wide area white space coverage. In VillageLink we embrace the 802.22 frame structure, and augment it with novel channel probing and operating frequency selection mechanisms.

Previous work related to channel assignment in wireless networks usually casts the problem of channel assignment as an NP-hard graph coloring problem [120]. Numerous heuristics have been proposed to provide an approximate solution ([159] and references therein). Ma and Tsang [110] recognize the channel heterogeneity in the case of wide bands and propose an integer linear programming solution for the frequency allocation problem. However, they restrict frequency reuse to well defined interference domains, thus no two BSs are allowed to transmit at the same time if they interfere. Motivated by [123], we rely on a more sophisticated representation of interference – we directly measure its impact through probing and account for any interference during the allocation process.

The Gibbs sampling, under this name, was first proposed in 1984 for image manipulation [56]. Its applicability to distributed channel allocation, client association and power control in wireless networks has been realized over the last twenty years [96, 118]. VillageLink differs from these by two important properties. First, we propose a novel network performance metric called *CINSR*, that takes into account frequency dependence of both useful signal transmissions and interference. Second, our algorithm does not require subsequent channel switching and environment sensing after each local decision is made. Rather, only control information has to be exchanged among neighboring nodes, and once the algorithm terminates only a single channel switch is made per node.

8.7 Conclusion

The heterogeneity of white space frequencies imposes unique challenges when it comes to channel allocation in a wireless network. Rather than simply minimizing interference, a channel allocation policy has to account for transmission quality over different channels as well. In this work we develop VillageLink, a channel allocation protocol that relies on the knowledge of signal propagation in the whole white space band before it performs distributed channel assignment that converges towards a network-wide optimum.

White space networks are largely unexplored, and their straightforward implementation might prove difficult due to unique characteristic they exhibit. For example, experiments we performed on VillageLink demonstrate the necessity for careful channel allocation in white spaces even when the number of devices operating in the same interference domain is low, which is a stark contrast to

WiFi networks. Our work examines only one aspect of network adaptation. The complex nature of signal propagation over a wide frequency band opens up new possibilities for protocol design and further refinement of channel access in white spaces.

Chapter 9

VillageCell: Local Cellular Access for Rural Areas

9.1 Introduction

It is clear that mobile phone usage has become one of the most prevalent means of communication worldwide. Statistics for 2011 from the International Telecommunication Union indicate that, worldwide, there are 85.7 mobile-cellular subscriptions per 100 inhabitants [12]. This number has been steadily increasing over the last 10 years, with growth over the last five years primarily driven by subscriptions in the developing world. Over the past five years developing countries annual growth has been approximately 20% compared to 5% for developed countries' [12] and the percentage of the population who owns a cell phone in the developing world increased from 23% to 68% between 2006 and 2011 [12].

While the statistics are encouraging, what they mask is the huge differential in cost and service availability and quality between the developed and the developing world. For instance, while residents of developed countries spend on average 2% of their monthly income on cellular service, the cost in developing countries is closer to 12% [12]. In addition, the ability of residents in developing regions to access cellphone technology is reduced by limited cellular deployments; while many residents own phones and buy either subscriptions or pre-paid plans, coverage may be spotty or non-existent within residential areas of rural developing regions characterized by low population densities. Further, while 4G is rapidly becoming available throughout the developed world, in developing regions coverage is often limited to 2G, or 2G + EDGE at best.

Despite these limitations and high costs, mobile phones are critical for providing communication in developing regions due to limited or non-existent telecommunications infrastructure and poor roads, making regular travel difficult. In many developing communities, cultures are oral – communication is based on oral rather than written form. Storytelling in such communities is vital to forming world views, maintaining trans-generational knowledge and teaching practical skills. Further, oral communication facilitates practical information exchange, such as crop prices, health care availability, and numerous others.

Our studies of locality of interest in Chapter 5 indicate that communication through technology largely appears between individuals in close physical proximity. Onnela *et al.* analysed 72 million calls and 17 million text messages in Europe and found that probability of communication decreases by 5 orders of magnitude when distance between communicating parties increases from 1 km to 1000 km [131]. VoIP services such as gTalk and Skype perform poorly for local voice communication due to their use of centralized services on the Internet that utilize the bandwidth-limited Internet gateway in the village. These statistics emphasize the need for solutions that support reliable local voice communication in remote areas.

Large telecom operators, however, remain reluctant to deploy cellular infrastructure in remote areas with low population densities [8]. Rural areas in both the developed and developing world typically have either limited cellular connectivity or no connectivity at all. Currently, deployment of cellular networks is complex and requires installation of Base Transceiver Stations (BTS) and supporting infrastructure. The installation cost is high, and it remains difficult for operators to establish a profitable network in areas with low income and population density. In addition, with seasonal revenues coming from subsistence agriculture, rural users often buy prepaid airtime non-uniformly throughout the year, thus leaving telecoms without a constant funding source [78, 55]. Compared to VoIP, cellphone calls can be prohibitively expensive. Finally, the BTS, which is the last hop to which handsets connect, requires a large and constant supply of energy, often unavailable in developing regions.

In this chapter, we propose a cost-effective architecture, dubbed VillageCell, for a GSM cellular network that operates in conjunction with a local rural-area network that serves as a backbone[1]. The solution uses a Software Defined Radio (SDR) controlled by a software implementation of the GSM protocol stack, called OpenBTS[2] and is able to make use of users' existing GSM phones. OpenBTS uses SDR for transmitting and receiving in the GSM bands and serves as a local cellphone base station. To extend coverage, multiple BTSs are connected through a local wireless network and calls are managed and routed via IP-based Private Branch Exchange (PBX) servers. Figure 9.1 illustrates the proposed VillageCell network architecture where cellphones, OpenBTS and PBX entities interconnect to offer widespread cellular connectivity. We describe in detail the different types of calls illustrated in the figure in Section 9.6.

[1]Note: In this Chapter, we use the term "local network" to mean the network within a rural village or community, connected to the Internet through an Internet gateway (i.e., a satellite link or a long distance WiFi link, etc.)

[2]http://openbts.sourceforge.net

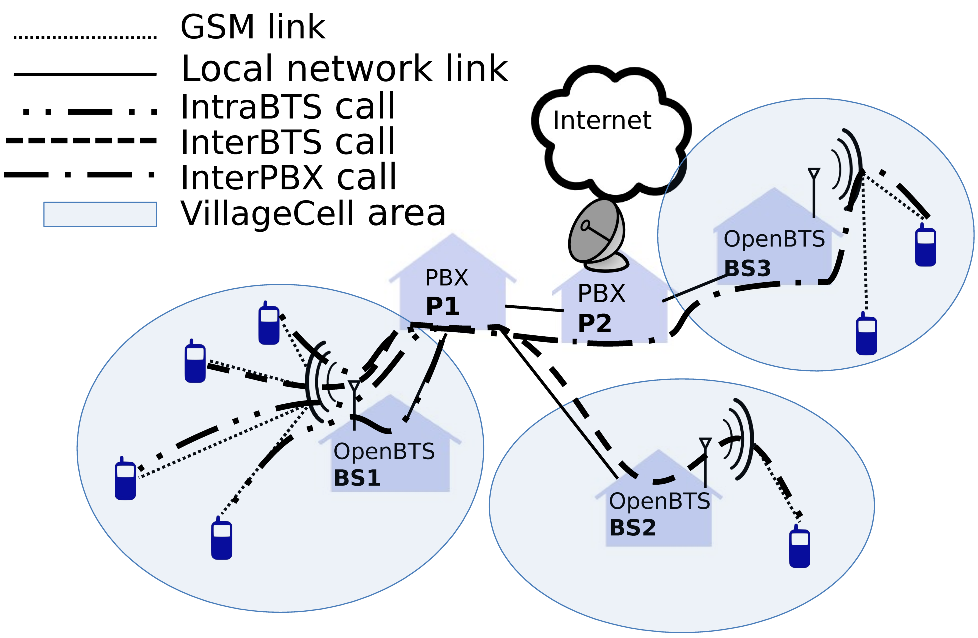

Figure 9.1: VillageCell network architecture. OpenBTS provides local coverage, while Private Branch Exchange (PBX) nodes support call management and routing. The expected number of users, their spatial layout, and the local network traffic load are taken into account for OpenBTS and Asterisk interconnection and placement.

While a single instance of OpenBTS has been proposed for rural communications before [73], to the best of our knowledge, VillageCell is the first system that provides coverage to whole villages. From that aspect, we are faced with a number of challenges. Our first challenge is related to placement and interconnection of multiple BTSs and PBX servers. VillageCell leverages any existing local wireless network. Thus, the location of BTS and PBX within the network can impact both legacy traffic as well as voice communication. The second challenge stems from the relative infancy of OpenBTS. The lack of comprehensive evaluation of OpenBTS performance as the traffic load on the wireless network and the number of users in the system change leaves us without any information on the VillageCell call quality and the system capacity. Finally, VoIP traffic is sensitive to packet delay and delivery reliability. In Chapter 4, we observed high variability of traffic load in the Macha rural network. Whether VillageCell can perform successfully in such a network is an important question we seek to answer.

To address the above challenges we evaluate VillageCell in both a lab environment with simulated traffic and a live deployment in Macha, Zambia. In our lab-based environment, we mix VillageCell traffic with our real-world wireless network trace gathered in Macha to account for realistic network conditions that inter-PBX communication faces in rural areas. The key results from our analysis, such as the call setup time, packet loss, delivery delay and jitter, demonstrate that VillageCell is indeed a viable and attractive solution for local low population density rural area communication.

We deployed two outdoor VillageCell base stations, connected via an existing WiFi link, in Macha, Zambia in June, 2012 and gave pre-registered SIM cards to a set of approximately 20 test users. In our our assessment we measured the same metrics as our lab-based experiment (call setup time, packet loss, average delay and jitter performance) for both intra- and inter-base station calls. The results in a live rural environment showed again that VillageCell can provide performance well below the thresholds that limit satisfactory call quality in an outdoor rural environment.

9.2 Voice communication in emerging regions

Voice-based applications have the potential to revolutionize developing regions. Well suited for areas with low literacy, voice delivers both global Internet content [102] and region-specific information [132] to remote communities. The range of applications span areas such as micro-payment management [103], education [169] and health care [160].

While the above benefits can be observed worldwide, the way communication tools are used often varies among different regions. Local ethnographies steer the appropriation of technology according to indigenous customs [45, 78]. In our work we concentrate on sub-Saharan Africa: a region where the narrative culture emphasizes the need for voice communication, where the lack of infrastructure is more pronounced than in the rest of the world and where the dispersion of population across a large geographic area makes the existing voice connectivity approaches challenging to implement.

9.2.1 Existing use of voice technology in Macha and Dwesa

To understand the way rural Africans appropriate VoIP communication, we use interview data described in Section 2.3.2 from Macha and Dwesa. This is also supplemented with our network trace data from Macha described in Section 2.3.1.

VoIP is highly popular in both villages, and 73% of interviewees use it through applications such as gTalk and Skype. Analysis of the traffic trace from Macha in Chapter 4 further supports this claim, with VoIP potentially contributing up to 26% of the traffic volume.

From the communication system design perspective, the locality of interactions plays a significant role in determining the most appropriate solution for a specific region. Strong locality of interest highlighted in Chapter 5 and 80% of interviewees using VoIP for intra-village interaction highlight the need for technology that enables local low-cost reliable voice communication.

9.2.2 Challenges with VoIP in rural communities

Community networks in the developing world, including those in Macha and Dwesa, often consist of a single satellite Internet gateway and a wireless network that provide connectivity to a number of end-users. In such a setting, the gateway is the bottleneck and limits the network performance. Internet-based VoIP is ill-suited for this type of a setup as voice applications such as gTalk and Skype establish a call between two nodes through a third-party Internet server[3]. In practice, this means that all VoIP communication between two persons residing in the same village has to go from the sender, over the highly congested satellite link, to the outside server, and back to the village along the same satellite link to the recipient. Our analysis in Section 4.3 of high round trip times (RTT) over the satellite links show that RTT can be in the order of tens of seconds. This makes meeting quality of service constraints exceedingly difficult, if not impossible.

The performance of VoIP can be enhanced through either reorganization of the way VoIP traffic is handled, by keeping the traffic within the village, for example, or through significant improvement of the outside Internet connection of rural villages. Cellular telephony, on the other hand, is robust with regard to the above technical issues. In addition, cell phones are far more prevalent than PCs and laptops[4]. Mobile telephony in rural developing areas faces two major problems: the coverage is often not available in sparsely populated rural areas due to high installation and operational cost, and low, seasonal income makes the price of air time out of reach for many of the residents.

[3]In the case of Skype, that server is called *supernode*, and represents a Skype user with very good connectivity, thus very likely outside of the rural area.

[4]In line with the global trends, we also find that 100% of interviewees in Macha and Dwesa own a cell phone, even though cellular coverage is sporadic.

9.3 Voice quality metrics

9.3.1 Mean opinion score

Voice call quality is often expressed in mean opinion score (MOS) and ranges from perfect (5) to impossible to communicate (1), where any score higher than 3 is considered acceptable. The E-model [82] converts packet loss and voice codec information into MOS.

- Maximum MOS for G.711 codec used in our Asterisk implementation is 4.4
- Maximum MOS for GSM 6.10 used in our FreeSwitch implementation is 3.46

9.3.2 Delay, jitter and packet loss

Delay, jitter and packet loss are three characteristics of a VoIP session that are critical to voice quality. ITU recommendation G.114 mandates that tolerable one way delay is up to 150 ms [8]. Furthermore, the theoretical minimum delay that a system can provide is dependent on the codec used. For example, GSM 6.10, used in our FreeSwitch-based VillageCell has a minimum delay of 20 ms.

Variance in the interframe arrival times at the receiver is called jitter. This is potentially more disruptive for IP telephony than the delay. In order to compensate for jitter and achieve a steady stream of packets, the receiver holds the first packet in a jitter buffer for a while before it is played out. This adds to the delay but can improve overall perceived voice quality. Typical jitter buffers in VoIP systems range from 50 ms to 100 ms and are usually adaptive to network conditions [71].

VOIP is not very tolerant of packet loss as packets containing between 40 ms and 80 ms of speech information match the duration of critical units of speech called phonemes. 2% packet loss, when using the G.711 standard without packet loss concealment[5]; G.711 with packet concealment techniques and other codecs with high compression, can tolerate 10% packet loss for a single unit reduction of MOS[146].

9.4 VillageCell

We harness the usability and prevalence of cellphones, with the affordability of VoIP communication, and propose VillageCell. VillageCell is designed with the following goals in mind:

- develop a low-cost, easy to deploy system that can be placed among groups of homes to provide localized cellular coverage.

[5] Packet loss concealment is a technique to mask the effects of packet loss in VoIP communications by using techniques such as interpolating speech gaps), results in a reduction of MOS (based on the E-model) from 4.4 (Good) to 3.4 (Fair) [146]

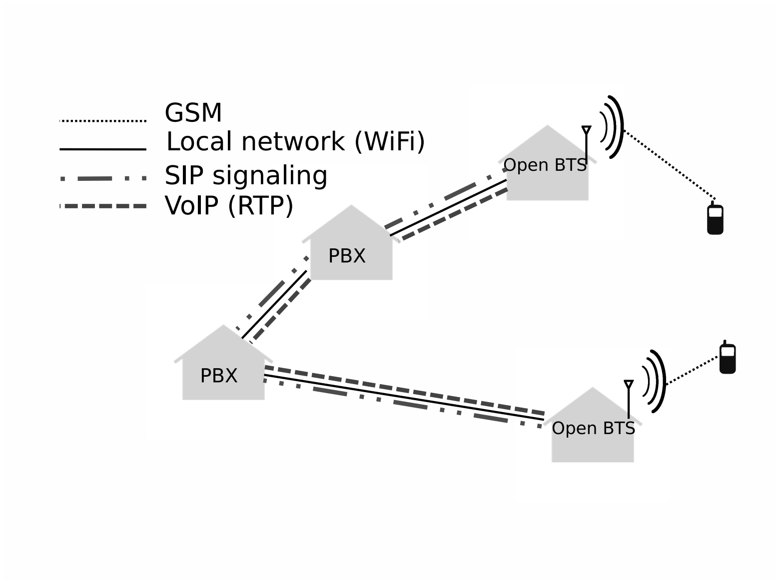

Figure 9.2: VillageCell protocols. On the MAC/PHY layer VillageCell relies on GSM and a local network protocol (usually WiFi). SIP signaling is used to establish a call, while the RTP protocol carries voice data (VoIP).

- provide free cellular calls within the local network while facilitating standard telephony connections to callers outside of the local network via VoIP.
- architect the necessary system component layout so that the call setup time and call quality are optimized.

In the following section we describe our system architecture in detail.

9.4.1 Architecture overview

VillageCell utilizes free, open-source solutions and off-the-shelf hardware to minimize the cost. Its architecture is modular and easily extensible – the VillageCell system can grow organically with the need for coverage. The main components of VillageCell are base stations and private branch exchanges.

OpenBTS is a software implementation of the complete cellular GSM protocol stack. It provides the network functionalities of *GSM registration, location updating* and *mobility management* which are, in a commercial system, distributed over multiple components such as Base Switching Centers (BSC), Mobile Switching Centers (MSC), Home Location Registers (HLR) and Visitor Location Registers (VLR). OpenBTS essentially converts GSM wireless signals into IP traffic that can be transported over a low-cost IP backbone.

OpenBTS uses SDR for the GSM air interface. SDR consists of a radio front end which transmits/receives wireless signals at the desired frequency[6] and a general purpose computer (PC) for signal processing. The OpenBTS software resides on both the front end and the PC.

One important functionality of OpenBTS is the interconversion of GSM and VoIP data. OpenBTS receives the GSM signals, demodulates them and converts them to VoIP packets that carry the call data (Figure 9.2). A call is established

[6]GSM bands are located at 850 MHz, 900 MHz, 1800 MHz or 1900 MHz.

when the signaling between the two parties is completed. This signaling is carried out by the PBX, a telephone switch system that relies on the Session Initiation Protocol (SIP) [81].

We use **Asterisk** or **Freeswitch** interchangeably as our open-source PBX implementation. The PBX works on a *client-server* model where a mobile phone in a VillageCell is presented to the PBX as a SIP client through the OpenBTS station, while the PBX acts as a SIP server. Both Asterisk and Freeswitch perform call routing and call monitoring for each of the connected SIP clients. They also maintain a database of all mobiles across the VillageCells, not only those that are directly associated with them. Finally, both Asterisk and Freeswitch allow connectivity to the public switched telephone network, and thus, integration with the global telephone system. Freeswitch, however presents a much more modular architecture that allows easy addition of voice or messaging applications, written in multiple languages, that call the Freeswitch core libraries.

The VillageCell communication range depends on the transmission power, which is limited by the specific hardware used and local regulations. In addition, villages differ in their layout. Thus, a varying number and position of cell stations is needed for different geographies. In Figure 9.3 we show an example of a typical sub-Saharan village layout from Macha, Zambia. The houses are dispersed over a wide area in small clusters with family members living in close proximity.

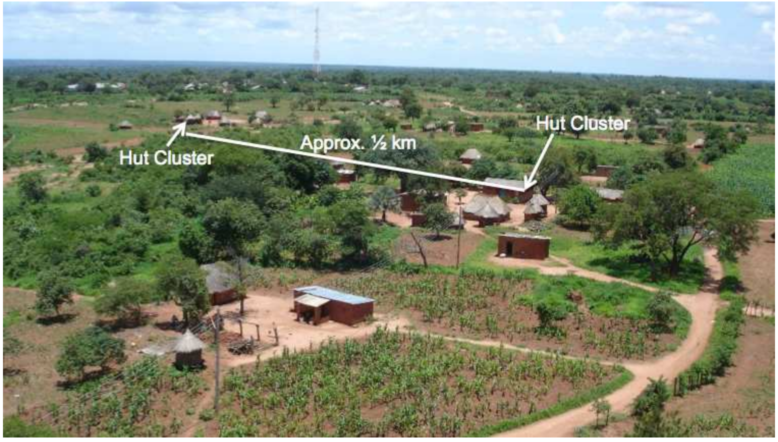

Figure 9.3: African village layout (Macha, Zambia). Clusters of houses are dispersed over a wide area. In Macha, the population density is 25 persons per km^2. Such a low population density, along with the low income, discourages large telecoms from deploying cellular networks in rural Africa.

We envision approximately one VillageCell per cluster, dependent on the distance between such clusters.

VillageCell components can be interconnected in multiple configurations; one PBX server can be common to many OpenBTS cells. Alternatively a single cell can also have a dedicated PBX server. In addition, the backhaul wireless network can carry varying quantities of non-VoIP traffic. In Section 9.6 we experimentally investigate the impact of the component layout on the call quality and the system capacity.

Connection between VillageCell base stations and PBX servers, as well as among the PBX servers themselves, can be realized with any standard IP-based technology: WiFi, WiMax, local Ethernet, 802.22. Local wireless (often WiFi-based) networks have been deployed in many isolated communities, such as Macha and Dwesa. If such a network exists, VillageCell can utilize it for call transfer. Within the underlying network, an OpenBTS or PBX server appears as just another node in the network.

9.5 Lab-based implementation

We implemented a prototype of VillageCell in a lab setting using readily available hardware components. Universal Software Radio Peripheral 2 (USRP2)[7] is a commercial SDR platform that natively supports OpenBTS software. We use a USRP2 with a general purpose PC for a VillageCell base station. The USRP2 platform hosts a powerful processing circuit (FPGA) for high bandwidth communication and a transceiver capable of operating in GSM bands. In our setup we use the 900 MHz band, as there are no interfering telecom carriers in that band in the USA. We do not amplify the USRP2 signal output, thus restricting the cellular coverage to a single indoor lab.

[7] http://www.ettus.com

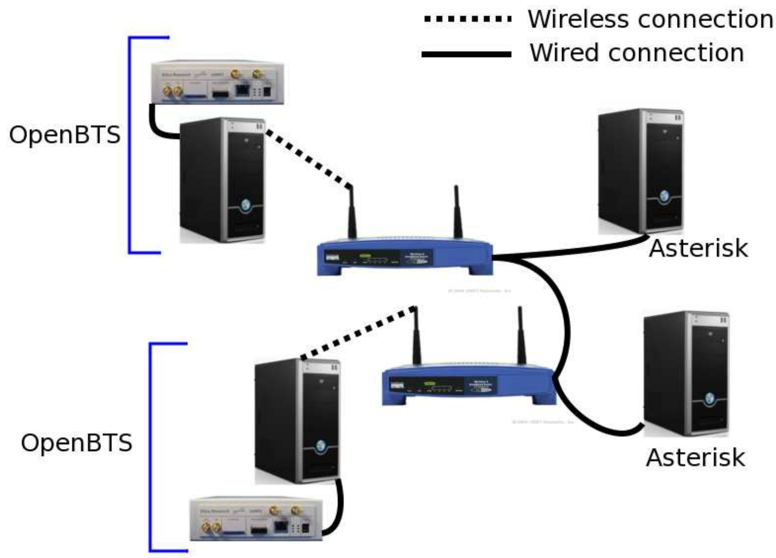

Figure 9.4: Experimental VillageCell setup. Shown is a configuration with two OpenBTS stations (each is composed of a USRP2 and a PC) and two Asterisk servers. The wireless routers ensure that the BTSs are connected via non-interfering WiFi channels. The rest of the configuration is connected via Ethernet.

For PBX, we use commodity PCs running Linux and the Asterisk software. Since Asterisk does not need a dedicated PC, it could be installed on the same machine on which OpenBTS is running. However, in order to isolate different parts of the system, we install Asterisk servers as separate entities. Connection among the components is established through two Linksys WiFi routers as per Figure 9.4. This setup represents a scaled-down version of VillageCell that would be deployed in the real world and helps us isolate the impact of individual factors, such as network layout, wireless interference, and background traffic, on the performance.

We tested our VillageCell implementation with three phone models: Nokia 3510 (from year 2002), Nokia 5300 Express Music (2006), and HTC Dream Android phone (2009). We also test the system with a range of SIM cards, with different memory sizes and belonging to different operators from both the developing and developed world, such as AT&T (USA), MTN (South Africa), Vodafone, Airtel and BSNL (India). Since we find no difference in the performance as we change the phone models and the SIM cards, we do not explicitly note these characteristics when reporting the experimental results.

9.6 Lab-based evaluation

We envision the VillageCell system on top of an existing rural area network. Thus, VillageCell voice traffic has to contend with other traffic for network resources. In this section we evaluate the capacity of the VillageCell system and the call quality in a realistic rural area network setup. Real-time voice communication has stringent packet delivery and delay requirements. While our low-cost implementation of a local cell phone architecture is not intended to compete with expensive commercial telecom equipment, VillageCell has to perform well enough so that quality local phone calls can be established.

9.6.1 Call scenarios

Three different scenarios of a VillageCell phone call can exist depending on the relationship between the call origin/destination and the architecture layout. We show these scenarios in Figure 9.1 and briefly describe them here:

- *Intra-VillageCell Call/Intra-PBX Call (IntraBTS):* The source and destination mobiles are registered as SIP clients under the same PBX server P1 and are both connected to the same OpenBTS station BS1. When a call request is made, BS1 determines the existence of the destination mobile by querying the PBX server P1 through *SIP Invite* signaling [81]. If a match is found then a communication channel is established between the station BS1 and the server P1 as *BS1 - P1 - BS1* and the call is connected.

- *Inter-VillageCell/Intra-PBX Call (InterBTS):* Here the source and destination mobiles are registered as SIP clients under the same PBX server P1 but under different OpenBTS stations. BS1 corresponding to the caller mobile contacts its controlling PBX server P1 and verifies the existence of the called mobile in a different VillageCell (with station BS2) through *SIP Invite* signaling. If a match is found then a communication channel is established between the two stations as *BS1 - P1 - BS2* and the call is made.
- *Inter-VillageCell/Inter-PBX Call (InterPBX):* In this case, each of the two communicating mobiles is registered as a SIP client with different PBX servers. OpenBTS station BS1 corresponding to the caller mobile contacts its controlling PBX server P1 and queries for the existence of the called mobile. The PBX server P1 in turn contacts PBX server P2 for the destination mobile's verification using *SIP Invite* signaling. If a match is found then a communication channel is established between the two stations as *BS1 - P1 - P2 - BS3* and the call is made.

Intuitively, the call scenario depends on the caller and callee position in the area served by VillageCell. However, with careful planning we can lay out the

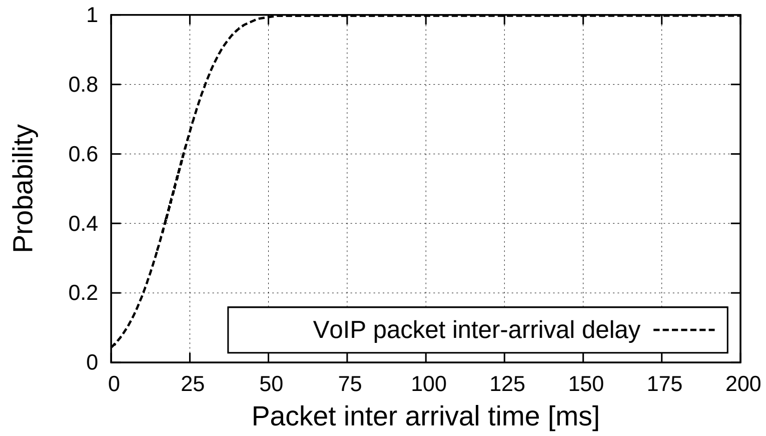

Figure 9.5: CDF of inter-arrival delay.

VillageCell components so that desirable scenarios occur more frequently than the others.

9.6.2 VillageCell call quality

To quantify the performance of our proposed architecture, we measured call setup time, maximum VoIP latency, delay jitter and VoIP packet loss for voice calls. We evaluate these parameters for each of the three call scenarios mentioned above. In a production network, the underlying wireless network will carry traffic in addition to VillageCell VoIP. To test the system under varying background load, we run a constant stream of UDP traffic with *iperf*[8] between the PBX servers, as well as between the PBX servers and BTSs, and vary the UDP traffic load

[8]http://iperf.sourceforge.net

between experiments. In each of the experiment runs we conduct a three minute long call, and for each of the data points, we average over five runs.

We measure the call setup delay as the time duration between the call initiation (*SIP Invite* signal) and call ringing notification (SIP 180 Ringing signal), both on the calling OpenBTS. In our experiments we observe call setup delay in the range of 1.5-2.0 seconds, which is an acceptable value. The default GSM voice encoding in our experiment is G.711 μ-law. This codec transmits packets every 20 ms. At the receiver, we measure the interarrival time between consecutive packets in a voice stream. In Figure 9.5 we provide the cumulative distribution of interarrival delay for the case of InterPBX scenario with 1Mbps of UDP background traffic. As observed from the figure, 85% of the VoIP packets have interarrival time of less than 25 ms, with 95% having interarrival time of less than 40 ms. The figure demonstrates that the VillageCell system is able to process and forward the packets while introducing little disturbance in the flow.

In Figure 9.6 we show end-to-end VoIP packet loss in the three scenarios. We push the UDP background traffic as high as 15 Mbps; beyond 15 Mbps network saturation occurs. The VoIP loss grows linearly with the background traffic and reaches the maximum at 15 Mbps, with 1.4% packet loss. The loss tolerance of the G.711 codec is relatively high, and as long as the packet loss stays below 10%,

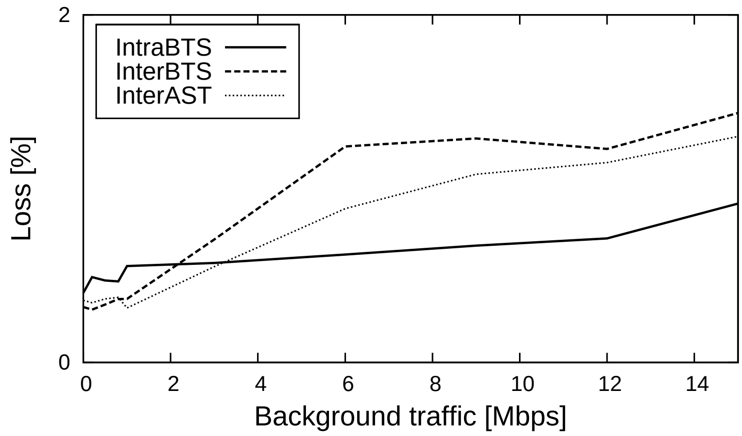

Figure 9.6: Packet loss with varying level of background traffic.

speech communication is possible. Our results show that packet loss does not limit VillageCell usability in these tests.

VoIP packets are sent at uniform 20 ms intervals; however, jitter in their inter-arrival time can impact call quality. We measure the jitter in all three test cases and show the results in Figure 9.7. We observe that the jitter increases linearly with the amount of background traffic. As discussed in Section 9.3, VoIP applications often implement receiver-side jitter buffers that store packets for some time (usually less than 100 ms) to cope with high jitter rates at the expense of increased end-to-end call delay. In our setup, the maximum jitter is always below 3 ms, thus a jitter-buffer with a short buffering period will suffice.

In Figure 9.8, we show MOS values for each scenario with increasing background traffic. We use the GSM codec G.711 μ-law which has a maximum MOS of 4.4. In all the cases call quality remains above 4, i.e. very good.

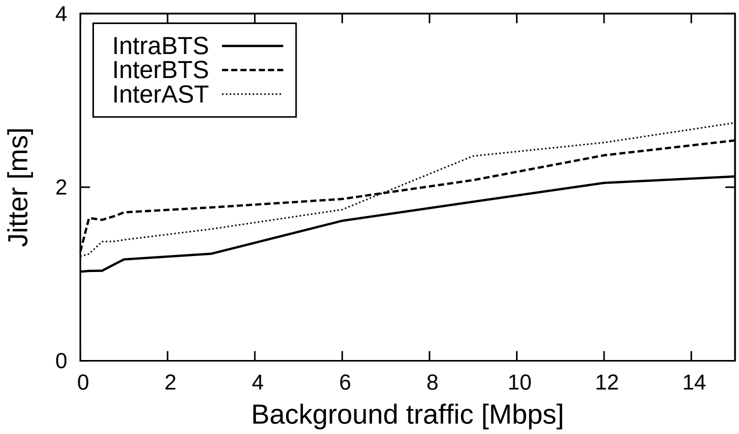

Figure 9.7: Delay jitter with varying level of background traffic.

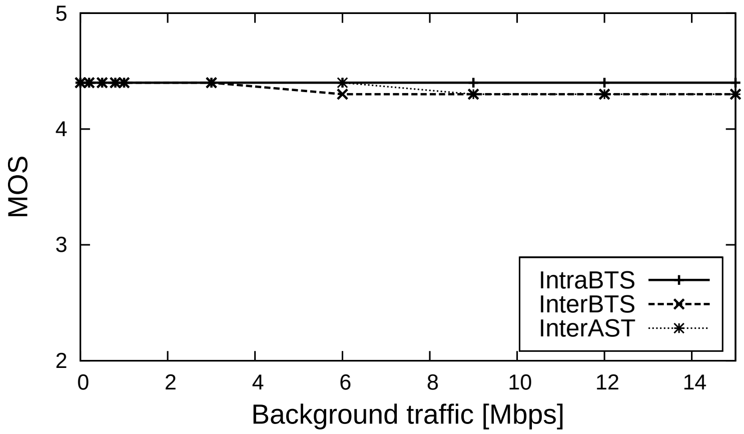

Figure 9.8: Mean opinion score (MOS) with varying level of background traffic.

	IntraBTS	InterBTS	InterPBX
Trace from Macha, Zambia	0.69%	0.82%	0.88%
iperf-generated TCP	1.00%	1.32%	1.81%

Table 9.1: Packet loss in a VillageCell system with varying background traffic types.

9.6.3 Realistic load experiments

Next we investigate VillageCell performance when the voice traffic is mixed with the traffic trace used in Chapter 3 from Macha, Zambia. This trace contains a mix of Internet protocols as the network is used for web browsing, email and non-VillageCell VoIP services, among other purposes. In our testbed we replay a randomly selected, ten-minute snippet of traffic from Macha. Similar to the UDP background traffic, we measure the packet loss that a single call experiences in each of the three configurations. For comparison, we run a separate set of experiments with iperf-generated TCP traffic as the background traffic and compare the results.

Table 9.1 summarizes the packet loss results from the experiment with the three VillageCell scenarios and two different background traffic types. The results are consistent with the earlier case of UDP background traffic. The packet loss remains below 2%, and higher loss is experienced in the InterBTS and InterPBX scenarios than in the IntraBTS scenario. This is consistent with the behavior observed under high UDP background traffic. Interestingly, iperf-TCP background traffic results in more losses than the real-world traffic from Zambia. The reason stems from the fact that iperf boosts the TCP throughput up to the limit imposed by the network conditions, which in the best case allow up to 54 Mbps. The traffic in Zambia, on the other hand, is more strictly limited by the satellite gateway capacity, which is only 1 Mbps at maximum.

9.6.4 VillageCell system capacity

We evaluate the capacity of VillageCell when it comes to multiple simultaneous calls. In our VillageCell prototype we establish a call and incrementally add more calls, up to a maximum of six calls. Once all the calls have begun, we measure the packet loss rate in each call and calculate the average value. All calls are composed of one physical phone as a receiver and one soft-phone as a caller, due to the number of devices we have at our disposal.

In Figure 9.9 we present the loss error rate for two types of configurations[9] as the number of simultaneous calls increases. We show the results with both no background UDP traffic and with 1 Mbps constant UDP traffic. In all four cases call quality experiences only a minor change in packet error rate (less than 0.3% increase) as we activate all six calls. While it is promising that we observe very little impact of the number of simultaneous calls on the call quality, in the

[9]Since we are using one soft-phone, which runs on Asterisk, there is no difference between the IntraBTS and InterBTS cases.

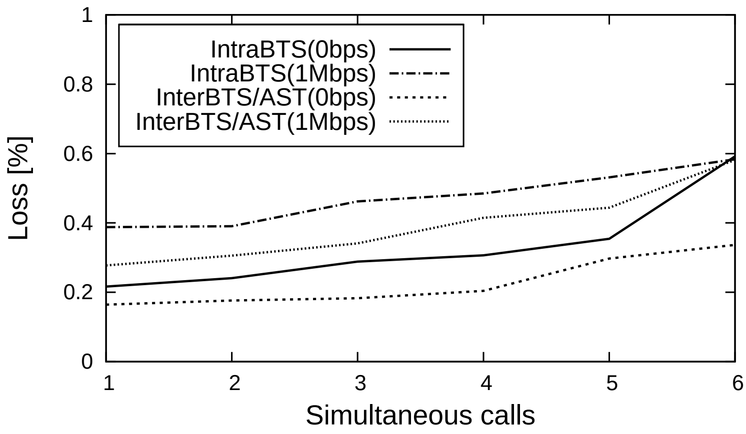

Figure 9.9: VillageCell performance with varying number of simultaneous calls.

future we plan to obtain more handsets and identify the true capacity limit of VillageCell.

9.6.5 Impact of VillageCell layout

The experimental results from the previous sections demonstrate that VillageCell provides high-quality voice communication under various network conditions. In addition, to the extent that we could test it, the VillageCell system scales well with the number of concurrent calls in the system. However, differences in the call quality can be noted among the three VillageCell scenarios: InterBTS, IntraBTS, and InterPBX setup. We analyze the three configurations with respect to the packet loss rate, delivery jitter and the number of supported calls.

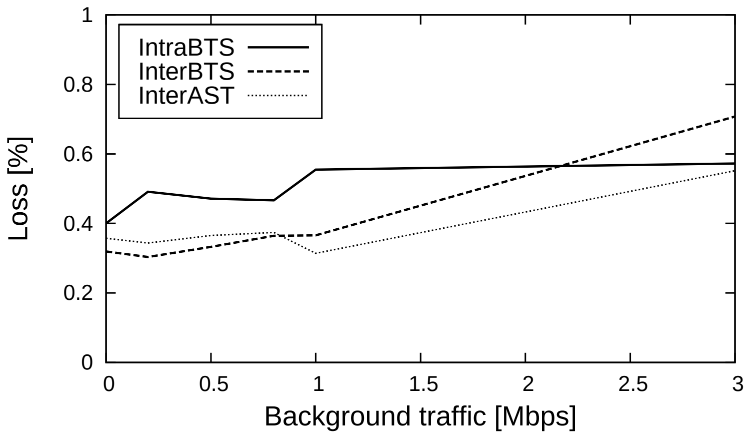

Figure 9.10: Packet loss under light background traffic load.

First, we concentrate on packet loss rate with varying levels of background UDP traffic. In Figure 9.6 we observed that the IntraBTS configuration results in lower average packet loss rate for all but low network loads. We enlarge the leftmost part of the graph in Figure 9.10. We observe that at less than 2-3 Mbps of background traffic, both alternatives (InterBTS and InterPBX) perform better than IntraBTS. The explanation stems from the distribution of losses. In the IntraBTS case, since both parties are associated with the same BTS, the same call traverses a single wireless link twice, from the BTS to the PBX server and back to the same BTS. Thus, the flow self-interference results in some dropped packets. This does not happen in the other two cases, InterBTS and InterPBX, as the flow never traverses the same link twice, nor two links in the same interference domain; the resulting loss is lower than in the IntraBTS case. When the background traffic is increased, however, the impact of uncorrelated losses on the two WiFi links (from BS1 to P1 and P1 to BS2) in the InterBTS and InterPBX configuration is more pronounced than the effect of self-interference in the IntraBTS case, thus the loss is higher. Consequently, the background traffic trace from Macha, Zambia or the TCP streaming (which is higher than 1 Mbps), is less detrimental in the case of IntraBTS configuration, as was shown in table 9.1.

The packet delivery jitter (Figure 9.7) is slightly lower in the IntraBTS case than in the other two cases. The difference is minor and can be explained by the fact that more links and PCs have to be traversed in order to establish a call in the InterBTS and InterPBX scenarios.

Finally, as shown in Figure 9.9, a higher number of simultaneous calls negatively impacts the call performance irrespective of the configuration, yet the performance of IntraBTS is worse than the performance of InterBTS/AST, regardless of the number of calls. In summary, irrespective of the number of simultaneous calls, IntraBTS calls experience more packet loss than InterBTS/AST calls as long as the background traffic remains low. We recap the findings from this section in table 9.2.

[10] For system scaling experiments we gathered results with up to 1 Mbps background traffic.

	Low background traffic	High background traffic
Packet loss	InterBTS/InterPBX	IntraBTS
Delay jitter	IntraBTS	IntraBTS
System scaling	InterBTS/InterPBX	InterBTS/InterPBX[10]

Table 9.2: Summary of VillageCell layout on the call performance. We show the optimal layout for each of the background traffic load and performance metric combinations.

9.7 Macha deployment

VillageCell was deployed in Macha in June/July 2012. Macha was chosen as the first field deployment, even though commercial cellphone coverage is available in the village center, as it is fairly well connected in terms of Internet access. This allows remote access to the system for administration and performance evaluation purposes. After consulting with the local Macha community, we named the VillageCell system, Kwiizya — a Tonga[11] word meaning "to chat."

In the following sections, we describe existing mobile phone usage in Macha, highlight design refinements we made to Kwiizya deployed in Macha as well as challenges we faced, many of which are not found in developed world deployments. We then describe how these challenges influenced our system deployment.

[11]Tonga is the native language in Zambia's Southern Province

9.7.1 Existing GSM use in Macha

Mobile phones are used extensively in Macha. For example, access to cellphones in Macha allowed farmers to overcome what they call the "briefcase buyers problem", whereby businessmen who buy maize come to the village and often attempt to exploit farmers who are unaware of market prices, buying crops at extremely low prices. Farmers in Macha can now call the Zambian Food Reserve Agency (FRA) to get information about crop prices. Furthermore, a new communication structure comprising of a combination of Internet and cellphones has emerged in Macha, whereby a farmer who has access to both the Internet and a cellphone can check crop prices online and send text messages to fellow farmers.

Cellphone coverage in Macha was first introduced in 2006 by Celtel (now Airtel). By 2012, MTN was a second active cellphone provider in the village. Residents with a cellphone subscription can use plain voice and text messaging and, where available, low data rate GPRS service for Internet access. The coverage provided by the two operators is largely available in the central part of the village; coverage is inconsistent and spotty in residential areas. Many people have subscriptions with both cellular providers to increase the likelihood of cellphone coverage at any given time as well as to make use of lower cost calls between the same operators.

9.7.2 Unique challenges

Power. The Kwiizya base stations require 12V DC and draw a maximum of 3A, which can be supplied by a 12V power transformer running off the main grid. Our deployment locations are connected to the national power grid; however, the quality of power varies drastically over time as highlighted in Section 2.4. This poor power quality is harmful to equipment and makes remote access for administration and evaluation extremely challenging.

Internet access. The remote accessibility of the system is not only influenced by power availability but by network availability as well. Power is often available in the village; however, due to an outage in the upstream link, highlighted in Section 2.4, the village gateway sometimes cannot connect to the Internet. This further limits our ability to access Kwiizya remotely.

Logging and storage. Logging is important for system administration and troubleshooting; however, special attention should be paid when enabling logging to ensure it does not deteriorate system performance, as writing to a flash-based disk can slow the system. We experienced an event where a combination of detailed logging and limited flash storage (4 GB) in the base station caused Kwiizya to malfunction.

Operation band. GSM in Africa operates in the 900/1800 band. While the 900 band provides for large coverage range compared to the 1800, one needs to be careful when selecting the operation band in order to not interfere with existing commercial operators, which typically operate in the 900 band in rural areas.

We adapt our system design to meet these challenges. The following section provides technical details about our field deployment.

(a) (b)

Figure 9.11: Our equipment in Macha: (a) the base station and (b) the power supply.

9.7.3 Technical details

Our Macha deployment operates in the GSM-1800 band. We chose this band for two reasons. First, most residents have basic dual band (900/1800) phones, so the use of GSM-1800 maximizes the number of people who will be able to use the system with their existing phones. Second, GSM-1800 does not interfere with commercial cellular providers, which operate at GSM-900 in rural areas.

For our Macha deployment we required a robust system that was fully self-contained in a weather proof enclosure and capable of a range of at least 3 km. Our USRP-based lab system was only capable of short-range communication for indoor experiments.

In order to build an outdoor long-range GSM base station, additional components were required. Figure 9.12 shows the components of the outdoor Kwiizya system. The duplexer isolates the receiver from the transmitter, permitting them to share a common GSM antenna; typically you need more than 48 dB of separation between your transmit and receive stage. It also provides rejection of transmitter noise to prevent receiver desensitization. The Low Noise Amplifier (LNA) is placed as close to the duplexer as possible to amplify the weak input signal before it experiences losses in the feedline of the USRP. We use a 1 Watt power amplifier on the transmit stage to achieve a 3 km range for the outdoor Kwiizya system.

According to the GSM standard, a clock with "absolute accuracy better than 0.05 ppm for both RF frequency generation and clocking the timebase" is required. The standard USRP has a frequency accuracy of 2.5 ppm, hence an external clock (clock tamer) is an essential requirement for a stable OpenBTS-based system.

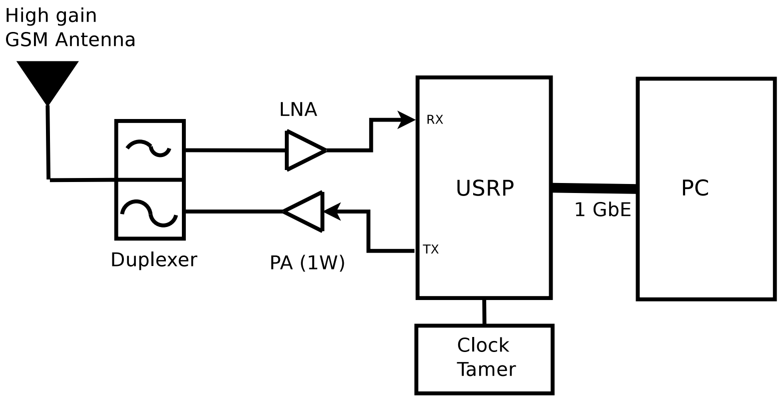

Figure 9.12: Block diagram of outdoor VillageCell(Kwiizya) system.

Instead of building the system from base components, we made use of RangeNetwork[12] Snap units — fully self-contained base stations with an embedded PC contained in a water-proof enclosure. The snap units contained all the subcomponents we required and we used this hardware to deploy VillageCell. The integrated PC contains a 4 GB flash card that runs Ubuntu 10.04.

9.7.4 Design refinements

Our initial VillageCell system made use of Asterisk. While refining our final system for deploying in Macha, we replaced Asterisk with FreeSwitch due to its modular support for voice or messaging applications.

The architecture of our Macha VillageCell deployment, Kwiizya, is depicted in Figure 9.13, To route calls within and outside Kwiizya, we use FreeSwitch. Like

[12] `http://rangenetworks.com/`

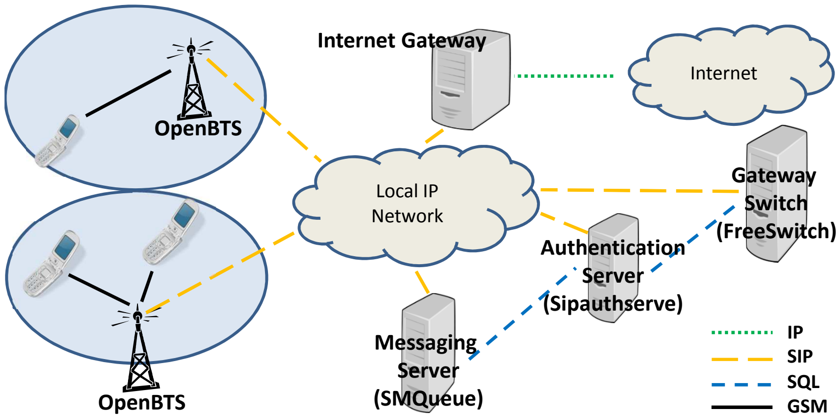

Figure 9.13: Kwiizya architecture.

Asterisk, FreeSwitch connects to OpenBTS via SIP and RTP and routes calls both in intra- and inter-BTS local scenarios. It also has the capability to route calls outside of the network to commercial cellular, fixed line and VoIP networks using SIP and SS7.

Kwiizya also utilizes Sipauthserve and SMQueue to handle user authentication and text messaging, respectively. SMQueue is the SIP-based equivalent of an SMSC (Short Text Messaging Central) in a commercial grade system. As such it interfaces with OpenBTS and makes use of commodity IP networks to transmit SMS (Short Message System). At the same time it can interface with commercial SMSCs using SS7 and SMPP. SMQueue implements a store and forward SMS queue functionality that allows messages to be delivered in a delay tolerant fashion. The latter is of great importance for areas with intermittent cellphone access and electric power availability as users are often either out of range or have their cellphone powered off. To handle user authentication and mobility, Kwiizya leverages Sipauthserve – a database server with an interface to processes SIP REGISTER messages to track mobility. Both SMQueue and Sipauthserve are queried by other network elements (e.g. FreeSwitch and OpenBTS) through SQL.

Each unit can independently run all components of a GSM network. However, to scale the network to a wider physical area, we use one of the base station PCs as a network central server, running FreeSwitch, Sipauthserve and SMQueue; the second base station runs only OpenBTS and connects to the first one (Figure 9.14) for all other services. The connection between the two base stations is facilitated through a long-distance WiFi link that was readily available through the local wireless network in Macha.

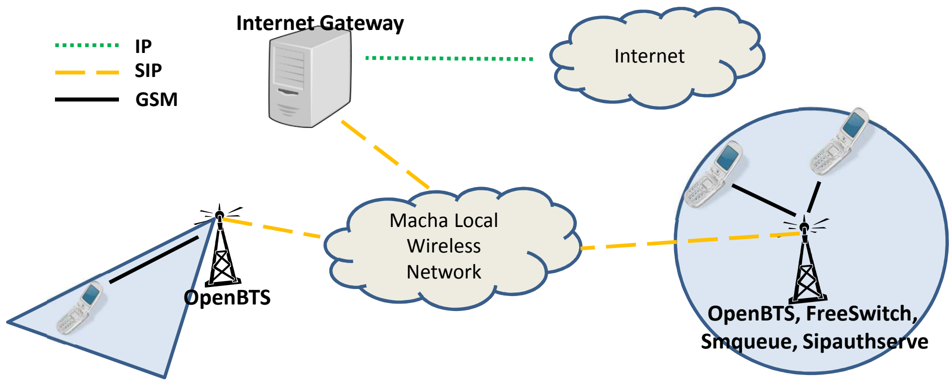

Figure 9.14: Kwiizya deployment in Macha.

To handle the problem of limited storage we attach a 1 TB external hard drive to each base station. We check the disk utilization periodically and when it reaches 90%, we offload syslogs from the base station PC onto the external hard disk to free disk space and save logs for future reference. We use simple ping-based monitoring to inform us when access to the village gateway is available from the outside, so that we can analyze network availability and know when our logs are accessible.

Determination of the most suitable power supply to resist the power fluctuations in Macha while providing stable 12V/3A to our base stations was an iterative process. We started with a combination of deep cycle car batteries and car battery chargers; however, within ten days we lost one charger due to poor power quality and two of the batteries started leaking due to overcharge. Eventually, we changed our power supply to the one shown in Figure 9.11b, which includes a UPS and a 5A/12V power supply.

The area of Macha is relatively flat with small hills, so providing coverage was a matter of installing the base stations on elevated locations. We installed one of the base stations on a 30 m water tower with a 10 m communication mast on top (Figure 9.15a). For this site we used an 11 dBi omnidirectional antenna. Our second base station was installed close to the LinkNet IT Academy (LITA) on a 12 m mast mounted on the ground (Figure 9.15b) and used an 11 dBi 90

(a)

(b)

Figure 9.15: Kwiizya sites in Macha: (a) the water tower and (b) LITA.

degree sector antenna. The distance between the two sites is 2.3 km; each site is capable of providing coverage up to 3 km within their antennas radiation pattern, depending on terrain. The water tower covered a larger population and therefore all the central services — FreeSwitch, Sipauthserve and SMQueue — were started on the water tower base station in order to minimize use of the WiFi link that connects the water tower to LITA.

While Kwiizya supports open registration with existing cellphones and SIM cards, we opted for restricted registration during our initial field testing. We manually provisioned 20 SIM cards and distributed the cards to a small set of users. This allowed us more control during the initial performance evaluation of our system, while still allowing our users unrestricted access. In the near future, we plan to deploy Kwiizya in neighbouring communities with no existing cellphone coverage and enable open registration.

We placed two GSM modems in the field to support controlled tests and experiments including text message transmission and voice calls. We chose U-Blox quadband GSM/GPRS modems that can be powered and controlled through USB. The latter is important because it allows the modems to be powered by the server to which they are connected, thereby making them independent of an individual reliable power supply. We attached the modems to an Ubuntu server and were able to access and control them remotely using AT commands, which are a suite of specialized commands for remote control of GSM modems.

Each call in Kwiizya has two associated VoIP sessions – one for the mobile call originator (MCO) and one for the mobile call terminator (MCT). A VoIP session consists of a SIP control session and forward and reverse RTP streams that carry the actual voice traffic. Kwiizya deployment in Macha supports four call scenarios depending on the locality of the MCO and the MCT. These are similar to those highlighted in our lab-based experiment in Section 9.5 with the exception of InterPBX calls. We describe these call scenarios in turn:

- **IntraBTS - no backhaul**: *Water tower — Water tower.* This is a call scenario where both the MCO and the MCT are in the vicinity of the water tower base station. This setup does not utilize the backhaul wireless link; both VoIP sessions are established only through the water tower base station.
- **InterBTS - with backhaul**: *Water tower — LITA* and *LITA – Water tower.* In the first scenario, the MCO is associated with the water tower and the MCT with LITA; the opposite holds for the second call scenario. These two scenarios are identical in terms of resource utilization. In each of the cases one of the VoIP sessions utilizes the backhaul link and the other one is local at the water tower.
- **IntraBTS - with backhaul**: *LITA — LITA.* In this case both the MCO and the MCT are attached at LITA. In this scenario, both the MCO and MCT VoIP sessions traverse the wireless backhaul.

9.8 Macha evaluation

Our evaluation of Kwiizya includes two parts — controlled experiments using test modems in Macha to check call setup times and a live deployment in Macha to evaluate voice call quality of the system with a set of test users.

9.8.1 Controlled experiments

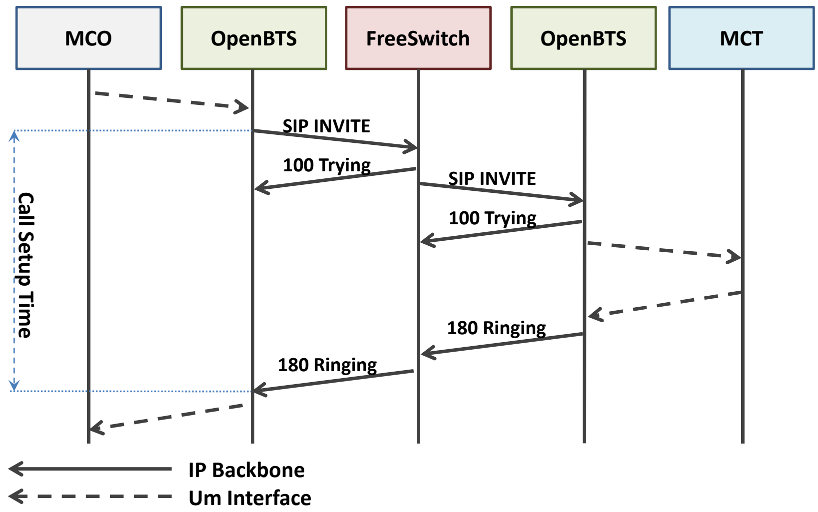

Figure 9.16: Call setup SIP message transaction.

In our controlled voice experiment, we measure the time it takes for a call to be established. As depicted in Figure 9.16, this is the time from when the first SIP INVITE message is sent by the MCO base station (informally from when the calling party hits "dial") to when the MCT starts ringing (or when the calling party starts hearing the "ring" tone). Components with the same color run on the same physical machine. We run an experiment with 15 consecutive calls. Figure 9.17 presents our results; call ID is plotted on the x-axis and call setup delay in seconds on the y-axis. The figure indicates that with one exception, it typically takes between 2 and 4 seconds for a call to initiate. We examined the delay components of the outlier call 14; the 15 second setup time is caused either by delay in the Um radio interface[13] between OpenBTS and the GSM modem or by a glitch in the GSM modem itself. We leave the evaluation of call quality to the calls placed by actual users to Section 9.8.2.

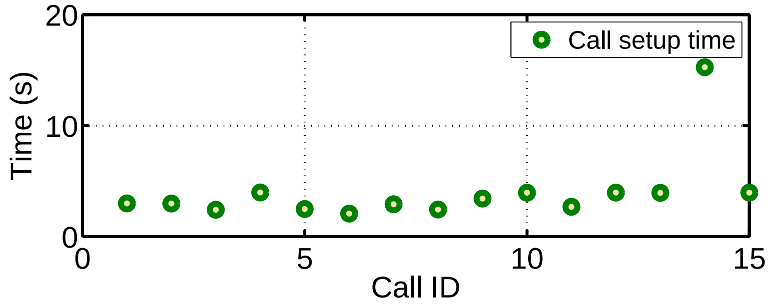

Figure 9.17: Call setup time.

9.8.2 Kwiizya field experiments

We now present results from in situ usage of our Kwiizya deployment in Macha from a period of two weeks in July 2012. We begin by describing our traffic collection system. We then evaluate call quality and compare this with our lab-based experiments.

[13]The Um interface is the air interface for the GSM mobile telephone standard.

Kwiizya monitoring

As noted in Section 9.7, our deployment includes two base stations – water tower (*WT*), which runs OpenBTS, FreeSwitch, SMQueue and Sipauthserve, and LITA (*LT*), which runs only OpenBTS and connects to *WT* to use other services. To capture all system traffic, we installed three monitoring points and run *tcpdump* at each point to capture SIP and RTP traffic:

- *lo@WT* – the loopback interface at the water tower base station that captures all internal communication in *WT* between FreeSwitch, SMQueue and Sipauthserve. This monitoring point allows evaluation of calls where either or both the MCO and the MCT are in the vicinity of *WT*.
- *eth@WT* – the Ethernet interface at the water tower base station that connects to the WiFi link from the water tower to LITA and captures all communication between *WT* and *LT*. This includes all SIP and RTP traffic related to calls where either one or both communicating parties are associated with LITA.
- *eth@LT* – the Ethernet interface at the LITA base station that connects to the WiFi link from the water tower to LITA. This monitoring point captures the same traffic as *eth@WT*, however with different timing of packets. This allows us to assess timing-related aspects of the system.

Voice call quality

We evaluate these characteristics in our system by analyzing 52 VoIP sessions from Kwiizya users in Macha that make use of the wireless backhaul to route calls. Calls using the wireless backhaul represent the worst case voice quality for calls as they can be affected by congestion or latency issues on the link.

Delay, jitter and packet Loss.

Figure 9.18 presents our results. For each VoIP session, Figures 9.18a and 9.18b plot the average and standard deviation of delay and jitter. A single point shows the average delay or jitter and standard deviation across all packets within that VoIP session. Both delay and jitter are well below the tolerated thresholds for VoIP; more than 99% of the packets in the actual call experienced delay less than 30 ms and jitter less than 3 ms. At the same time, the average delay is close to 20 ms over all VoIP sessions. Furthermore, delay and jitter do not vary much over a single VoIP session as indicated by the standard deviation bars, which shows that Kwiizya has stable performance throughout a call.

Finally, Figure 9.18c plots packet loss over the 52 VoIP sessions. Only three of all sessions suffered non-zero packet loss; however, these three were all less than 0.5%, which is within the limit for satisfactory call quality (see Section 9.3).

Mean Opinion Score (MOS).

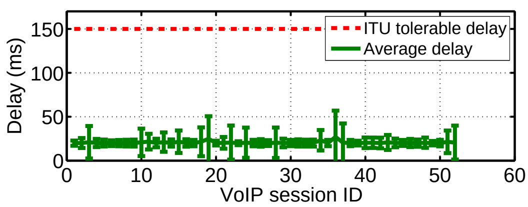

(a)

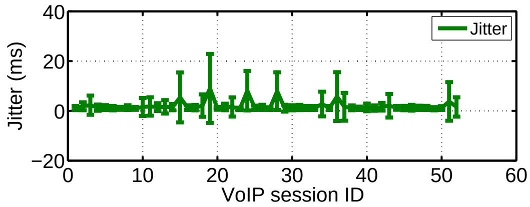

(b)

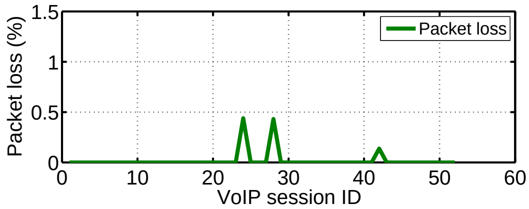

(c)

Figure 9.18: Evaluating calls: (a) average delay, (b) average jitter, and (c) packet loss.

FreeSwitch uses the GSM FR (6.10) codec, for which the maximum expected MOS is 3.46. Figure 9.19 presents our results for MOS for the 52 VoIP sessions we analyzed. On the x-axis we have session ID, while on the y-axis we have MOS. The minimum measured MOS is 3.43 and the maximum (received by 49 of the 52 sessions) is 3.46 – equal to the maximum provided by the system.

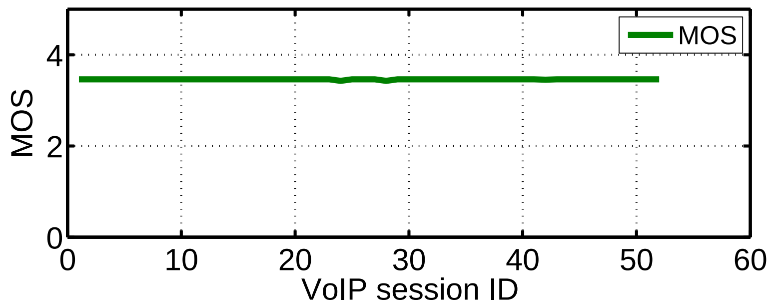

Figure 9.19: Mean Opinion Score (MOS).

9.9 Deployment planning

Our goal in designing the VillageCell architecture is for it to be flexible and adaptable to user needs. However, another set of restrictions comes from the topology of the existing community network (if any), energy resource availability, and regulatory issues. Here we discuss VillageCell planning from all of the above aspects.

9.9.1 Component layout

VillageCell can be built on top of an existing community wireless network. Because VillageCell performs differently in different configurations (IntraBTS, InterBTS and InterPBX), and with varying levels of background traffic (table 9.2) we devise a set of guidelines for VillageCell planning:

- IntraBTS performs worse than InterBTS/PBX when the background traffic is low. As a consequence, where local interaction is high, and the edge of the local network is near (where the WiFi backbone ends), rather than trying to cover a large area with one powerful BTS (IntraBTS), it is more attractive to have a few BTSs and split the load between them, as in the InterBTS architecture. Consequently, the losses will be lower, as we expect low background traffic, since there will be no aggregated traffic from nodes whose path to the gateway traverses VillageCell links.

- IntraBTS is not as sensitive to background traffic (Figure 9.6). Thus, OpenBTS-PBX communication can use a congested backbone link as long as the communication remains local (IntraBTS). If we consider a community where we expect a very high level of local interaction, we can connect the BTS and the PBX server directly to the backbone, without the need to have a separate wireless link dedicated to that connection. This reduces the planning effort and the cost of deployment.

- InterBTS and InterPBX are sensitive to high background traffic. If we have two locations where we expect a lot of mutual interaction, we should connect their BTSs to a PBX(s) with dedicated WiFi links. While this may increase the cost of deployment, it assures reliable delivery of both VoIP and existing network traffic.

- Because calls are routed through the PBX servers, we should keep the PBX servers local to the BTSs in the areas of high level of local interaction to avoid packet losses that occur in BTS - PBX dedicated WiFi links.

9.9.2 Outside connectivity

VillageCell is optimized for free local communication, though it can also connect local users to the outside world using a commercial VoIP Network. In our system, PBX machines on the edge of the local network can be connected to the outside world over the Internet via VoIP. The traffic going to a VoIP network is billed according to the VoIP operator's usage terms. Non-local call routing is performed as follows. When a PBX server at the edge of the network receives a call request to a user not present in the local network, the call request is forwarded to the database of an external VoIP provider registered on the PBX server to locate the user. If the user is found in the database, the subsequent call traffic is routed via a satellite gateway over the Internet. On the other hand, when an outside user calls a user who is located within the VillageCell system, the edge PBX server translates between a globally accessible VoIP ID and a local VillageCell phone number. In this paper, we focus on VillageCell's operation within the local network, but we note that communication outside the network is also feasible and we plan to implement it as a part of our future work.

9.9.3 Energy issues

VillageCell components, such as OpenBTS stations and PBX servers, can be built out of commodity PCs or laptops. These devices consume on the order of hundreds of watts or less. The radio front end, provided by USRP2, consumes only up to 13 watts. While this implies that VillageCell needs more than an order of magnitude less power than a commercial cellphone station, unreliability of the electrical grid in rural areas still presents a major problem. Heimerl and Brewer [73] propose powering OpenBTS base stations with wind and solar energy. This attractive alternative, however, comes with an added cost of energy harvesting equipment, which could surpass the cost of communication equipment [18]. Further investigation is needed to identify the optimal energy availability – equipment cost balance, and tackle problems of possible power shortages due to unfavourable weather conditions.

9.9.4 Licensing issues

Worldwide, operation on GSM frequencies requires a license. Usually, a license is granted on a national or a regional level to a large telecom. However, this does not necessarily prevent smaller players from deploying OpenBTS-based systems. In the United States, FCC grants experimental licenses for GSM bands as long as the irradiated power is less than 8 W. Analysis of WiFi frequency bands showed that non-restrictive licensing contributes to increased Internet connectivity [19]. The final decision, however, lies with the regulatory bodies and their assessment of local cellular coverage benefits.

9.10 Related Work

The impact of cellphone technology on residents of developing countries has been widely studied [14, 45]. Examples from the literature show that cellphones have changed the way people learn and exchange knowledge, get access to health care and handle local government activities. There is increasing effort in developing applications that use plain text messaging and voice to leverage existing feature phones. These applications can be largely divided into those that use voice and those that use text messaging. In [132] for instance, the authors deploy a social network based on voice forums for farmers to exchange knowledge about crops and crop prices. In [129] message exchanges are encoded as a sequence missed call durations. Numerous applications use text messaging as a platform. Use cases include update of remote databases through SMS [101], attendance tracking [144], web search [37] and health care [15, 47, 48, 101].

The drawback of these solutions is that they all assume existence of an underlying cellphone network. However, ITU statistics show that 48% of the rural population in sub-Saharan Africa is still disconnected; this constitutes a large fraction of a population that would benefit most from the outlined applications.

A few projects design a local cellular network solution for remote areas [73, 124]. Heimerl *et al.* design a village base station to provide off-grid deployment of local cellphone networks [73]. Mpala *et al.* study the applicability of an OpenBTS-based system in rural areas [124]. These works provide insight into design of such systems and extensive discussion of obstacles to implementation; what they lack is holistic system design and evaluation of feasibility from a system point of view.

9.11 Conclusion

In this chapter we presented VillageCell, a low-cost localized cell phone system for rural areas. We implemented VillageCell in a lab setting and evaluated it in realistic rural-area network scenarios. Through the experiments we identified technical issues that are crucial to core functionality of the VillageCell architecture: establishing local intra- and inter-VillageCell calls. We show that call quality in our system is often very good with little packet loss, fast setup time and low delay and jitter both in repeated lab experiments and in a deployment in Macha, Zambia. From the variations in performance that we observed as we modified the network layout, we derived guidelines for efficient VillageCell integration in rural-area wireless networks with a number of different user distributions and traffic load patterns.

VillageCell solves an important problem of providing localized low-cost voice connectivity. In the future, we plan to develop applications, targeting the health and education domain, specifically suited to VillageCell's unique offering. Moreover, many existing applications for developing regions that experience implementation problems due to high costs of cellphone communication or lack of sufficient coverage can benefit from VillageCell [132, 154].

Chapter 10

Discussion

This dissertation has sought to investigate and provide technical solutions for computer networks in a rural African context. We have specifically focussed on communities in which access is spread from a single bandwidth-constrained Internet gateway to a sparsely distributed population. We have provided architectural solutions which take advantage of unique usage and access aspects of these communities such as web-centric Internet access, a high prevalence of online social network usage, strong social cohesion and locality of interest, diurnal network usage patterns and a high degree of mobile phone ownership.

There are many other access modalities for rural contexts that have been explored which are either complementary or supersede our VillageNet approach.

To put our solutions into context, we now discuss how this work relates to these historic or current access techniques.

Information and Communication Technology for Development (ICT4D) has evolved through three phases over the past two decades [72]. In the first phase, from the early 1980s until about 1990, computers were deployed for improved government administration and by multinationals to foster economic growth. From the mid-1990s, a second-phase of ICT4D work began where the World Bank called for adoption of ICTs as a tool for development in response to the Millennium Development Goals (MDGs). One of the most popular choices to rapidly respond to the plight of poor, rural communities was the use of telecentres to deliver ICT services. A third phase of ICT4D is now emerging which views ICT not only as a tool for development but also as a means to an end, for example producing digital content and building services to generate income. This is a move towards active innovation within communities to achieve development goals and can be achieved for the poor, with the poor or by the poor themselves. This involves a wide range of possible access devices from mobile phones to tablets to PCs and connectivity architectures that go beyond a simple single point of access, such as a telecentre. As such, our solutions support initiatives embodied by this new third phase of ICT4D.

10.1 Alternative access modalities

Telecentres have been employed in many remote rural communities and some continue to be used [162]. Supporters of telecentres believed they would solve many challenges in the developing world: distance education would provide rural scholars with access to advanced learning, citizen services would reduce government corruption and local economies would grow, with claims of even doubling income in rural villages [88]. Many telecentre operators provide anecdotes of farmers improving crops in India and Africa [168, 149] or village operators earning good incomes from operating telecentres [88]. While there are a few successes, many more failed within a few months or years after being set up due to lack of financial sustainability [98] (typical telecentre running costs are $100 per month in Africa). The limited number of telecentres that keep operating are often run by devoted non-profit organizations and demand considerable effort or entrepreneurs, often needing to balance multiple sources of income, to keep their telecentres running [98].

Somewhat linked to the kiosk model is a model of access using ruggedized public computer terminals placed in an open public area with a large amount of off-line content and applications. This model was introduced by Mitra in the "Hole In The Wall" experiment in India in 2000. A computer was placed in a recess cut in a wall on a busy street [122] and members of the community were observed interacting with this 'high-tech device', even though they had never used a computer. Contrary to most popular skill-acquisition models, unassisted learning through trial and error allowed many users to acquire basic ICT literacy skills, especially amongst the Indian youth; this model of learning became known as 'minimally invasive education'. Building on this principle, the 'Digital Doorway' project was launched in South Africa in 2002 [63]. The project focussed on providing ruggedized kiosk-like terminals to South African youth with off-line education-targeted content — disadvantaged areas with low ICT penetration were targeted. Core to the design of the digital doorway is user interaction on multiple terminals placed on three or four facets of a kiosk containing a central computing hub. This defines the digital doorway as a social entity where a number of users congregate in one place and participate in peer learning — a concept diametric to 'one laptop per child' which focusses on individual ownership.

Although a kiosk model is a useful bridge to build IT literacy, especially amongst youth, or providing some basic computing or Internet services, it will not fully immerse users in the modern digital age and allow them to become active contributors to the World Wide Web. Our solutions are focussed on extending access beyond the kiosk model or single point of access to homes, businesses, and health and educational facilities. This allows users to avoid a deliberate interaction model [172] typical of the high cost of per-minute Internet usage in telecentres or the limited time available on a public terminal. This model also allows a village operator to recoup high Internet costs from a larger number of paying users within wireless reach of the Internet gateway.

Delay Tolerant Networks (DTNs) provide the middle ground between offline public access terminals such as the Digital Doorway and fully connected Internet solutions. Many networking applications assume end-to-end connections always exist from the client to the server. However, in many rural areas no connectivity exists. In these cases, DTN store-and-forward protocols can be utilized. Data is stored in intermediate nodes until a connection can be made to to receive data from a distant source or send data towards its destination. DTNs can also deal with networking scenarios where network disruptions occur due to frequent power failures common in many rural areas. Some examples of successfully deployed DTNs are KioskNet [62], DakNet [134] and a system deployed in Anandpuram, India, to connect Internet kiosks with buses [153]. However, DTNs cannot provide a user with a modern Internet experience as a large fraction of Internet usage we observed was often real-time, such as Facebook chat or personalized services requiring a user to login and maintain a live session or dynamic web content generated as the user interacts with a web page.

The intervention we have studied in Macha, Zambia is essentially a bottom-up approach to connectivity, in which an NGO builds capacity in a community to grow their ICT infrastructure outwards from a core nucleus. There are other interventions, which involve a top-down approach at a Government level to connect infrastructure such as schools or health facilities. Examples of such interventions are Brazil's "Broadband in Schools" program established in 2008 which has resulted in about 84% of Brazilian students having access to free broadband in urban public schools as of 2011 [86]. The Khanya project in South Africa, established in 2001, addresses the shortage of educator capacity and the need to deliver curriculum to schools through ICT [1]. It has delivered PCs, network infrastructure and training to the majority of public schools in the Western Cape province using local government funding. Affordable Internet connectivity to rural schools in South Africa, however, still remains a key challenge. The Khanya project and others like the NEPAD e-schools [54] initiative, which aims to integrate ICT in education across 16 African countries, all acknowledge the need for local community involvement and a community champion in sustaining an ICT intervention. The BB4All [150] project in South Africa, which aims primarily at reaching rural schools, embodies this principle in the village operator model. Village operators run a rural micro enterprise that builds, operates and supports localized mesh network infrastructure to reach schools initially and then businesses and homes during a later phase. The backbone networks into the rural regions, however, are funded by government. It is clear that catalysts to successful initiatives that aim to bridge technology or digital divides in developing regions require a hybrid structure of both bottom-up community-driven interventions for sustainability and top-down government-driven interventions to facilitate growing to scale. As an example the Macha Works model we have studied requires the Zambian government to make use of Universal Service Agency funding in order to scale their model to other areas in Zambia.

From this work and the author's previous research on community-wireless mesh networks built on licence-free spectrum [93], it is clear that low-cost licence-free networks are an important component in ensuring that access extends to many marginalized communities not connected by commercial operators beyond the "affordability frontier" in rural areas. The Macha community in Zambia and the BB4all project both rely heavily on license-free WiFi networks to extend the reach of the Internet. Increasing the amount of licence-free spectrum by including white spaces can only increase innovative solutions [107] for rural access and help expand connectivity to more remote communities.

Mobile networks continue to grow in developing regions at impressive rates — mobile phone subscriptions in developing regions have increased from 10% to 78% between 2001 and 2011 [12]. The development community has contemplated the possibility of the mobile phone ending global poverty [42] and there are a number of projects which have demonstrated that mobile phones can help eliminate information inefficiencies in developing-world markets and health services [87, 79]. Universities have built entire departments devoted to building mobile phone applications for developing regions [6]. One could even argue that some countries with close to 100% cell phone coverage of the population, such as South Africa, have bridged the digital divide as some degree of Internet access is available on most mobile phones. However, the limitations of mobile phone device screen sizes, limited input from mobile phone keypads, high costs and low data rates (20 MHz for GRPS) in these areas provide users with a very limited Internet experience, which at best is mostly used as a slightly more advanced messaging service. For example Mxit [3], a messaging service which runs over mobile phone data connections, is Africa's largest online social network. Some have even argued [32] that mobile phones reinforce dynamics of uneven development due to their "socially articulating" nature. Mobile phone services reflect the second phase of ICT4D services, where ICT is used as a tool for development. Digital content generation and application development reflected in the third phase of ICT4D in developing regions is still mostly carried out in the context of a PC connected to a broadband Internet connection.

10.2 Applying VillageNet to other access modalities

Although the philosophy of other access modalities in developing regions may be different to VillageNet, which aims to enhance local access and grow outwards to low-density communities from a central nucleus of Internet access, there are a number of ways in which some components of VillageNet can be used in these other contexts.

In the case of public access terminals, such as Digital Doorway or telecentres, VillageShare could be installed within a public access terminal or on a telecentre server. This would provide a local file sharing service that eliminates Internet costs and increases the speed of local content sharing amongst users who visit the facility. Once data connectivity is available on VillageCell, the large amount of educational cached content on the digital doorway could be accessed by mobile phones at no cost from nearby users. DTNs are mostly used to connect single points of access such as kiosks or telecentres and VillageShare can be used to share local content. However, sharing content with external users on a DTN network would require some modification and integration with the DTN routing system to upload or download content shared between users within the village and users outside the village. Telecentres are often the first stage of Internet access in a rural community (see Section 10.4) and VillageLink white space connectivity can be used to spread this access to the surrounding population.

Government interventions to connect schools could also benefit from VillageNet. VillageLink could be used to extend the reach of government-supplied backbone to reach remote schools and homes located in challenging terrain. VillageShare could be used by students and teachers to share education content at local schools and VillageCell could be used to provide free after-hour educational resources to learner's mobile phones within a 5 km radius of the school. South Africa has piloted a project called "Dr Math" [29] to provide mathematics instruction via a messaging service on mobile phones. This service, however, relies on students having airtime on their phones in order to use the service. VillageCell would allow students to continue using the service even when their airtime has expired.

Although many components of VillageNet will work in the other contexts highlighted, careful analysis of true user needs as well as the environmental and cultural context should be carried out before charging forward with a specific technology solution.

10.3 Moving beyond the 'haves' and 'have-nots'

Many studies on ICT access in the world reflect quantitative measurements [12], measures such as number of mobile phone subscriptions or number of Internet connections per 100 inhabitants. What these studies fail to capture are the more subtle nuances within a community — rural communities in particular — that hinder or encourage integration of ICTs into a cultural context.

In Chapters 2, 3 and 4 we capture some of the measurable obstacles that rural users experience when using the Internet: the high cost of Internet access, limited bandwidth, content uploads failing due to unreliable or congested uplinks and the high prevalence of computer viruses. Cost may be the most debilitating constraint to access; in 2010, a fixed broadband connection in a developing country cost 112% of per capita GNI and only 1.5% of per capita GNI in a developed country [10]. These costs are often due to the large investment required by a commercial operator to bring connectivity to rural areas. High costs may also be due to poor government regulation to control pricing or provide conditional licenses which provide a mandate to connect a certain fraction of users in under-serviced areas.

For a rural user paying high Internet costs, the benefits and utility of access needs to be high. Measuring the degree of economic and social benefit is a challenge without ongoing longitudinal studies. Current research often lacks long-term contextual evidence on the use of ICT in rural areas of sub-Saharan Africa in particular [4]. This is due to many projects not extending over multiple years, or beyond the project implementation phase. Some econometric studies using long-term ITU data suggest that a 10% increase in broadband penetration could result in a GDP increase of 1.38% in low and middle-income economies [137]. However, these macro-economic models do not differentiate between rural and urban communities in developing countries and with ICT penetration being higher in urban areas, they tend to be weighted towards urban development. Rural areas stand to benefit the most from from ethnographic and economic studies that seek to understand changes that occur over the long term when ICTs are introduced into a community.

Limited bandwidth is becoming an increasing concern as the average web page size continues to grow. The average size of a web page has grown more than fifty times in the past 15 year even though the average International bandwidth in Africa is only 937 bps per user, while Europe has 78,678 bps available per user [12]. We are entering an era where it may be more pertinent to talk of a “bandwidth divide” rather than a “digital divide” that focusses on access to, use of, or knowledge of ICT. The location of Internet access also plays a key role in the level of user engagement with the Internet. Wyche et al. reveal that public access, typical of the kiosk model, leads to "deliberate interaction", where online activities need to be pre-planned [172]. On the other hand, the Digital Doorway project revealed that public access amongst youth can lead to peer-learning [63]. Hence at-home and public access may be seen as playing a complimentary role.

The high prevalence of viruses detected in network traffic highlighted in Chapter 4 raises additional questions related to managing complex technology in rural regions. In Africa, in particular, decisions related to technology choices are often based on peer knowledge rather than individual choices or investigation in literature. Most community members aim for socially accepted and consented behaviour [115] with idea dissemination occurring primarily through orality [130]. This makes communicating complex concepts such as bandwidth capacity a challenge[93]. Many users choose Windows because their peers use Windows in spite of many negative consequences related to its use. Windows machines are far more susceptible to virus infections and virus protection software is challenging to keep up-to-date when connected over a slow, costly Internet link. As highlighted in Chapter 2, Windows machines also suffer from poor uplink performance on high-delay links compared to Linux or Apple operating systems due to the delay-sensitive TCP protocol used by the kernel. Hence, considerable responsibility is required when introducing technology as it acquires a much higher degree of inertia than in a developed country.

Many other constraints to digital inclusion within certain cross-sections of a community can only be discovered through unstructured interviews or discussion with community members. For example, some Zambian users perceive the Internet as a possible threat to indigenous communities [164]:

> Internet in other ways will either build or destroy our culture because of the powerful influence it has in peoples lives. Truly speaking, most of the things exposed on the Internet are from the western world and very little is from Africa. I do not want to lose my culture.

The Internet abounds with Western socially accepted norms, that may be frowned upon in other cultures. Norms such as dress-code, family structure or attitudes to elders. Culturally offensive material may appear in adverts or suggested content typical on sites such as Youtube.

This raises a key issue: how much content on the Internet is locally relevant to the average rural user? In the west, users obtain map directions, carry out e-commerce or access vast repositories of knowledge in their own home language. Many users in rural developing regions to do not enjoy these same benefits when using the Internet. Local languages are hardly represented. For example, there are no Wikipedia articles in Chi-Tonga, Ila and Lozi, languages spoken by more than two million people in (rural) Zambia, whereas approximately the same number of Slovenians enjoy a bounty of more than 100,000 articles in their native language. Information on the Internet is often tightly coupled to Western amenities such as well-marked and named roads or door-to-door courier services.

It is clear that a full picture of the utility of Internet provision, in rural areas in particular, requires more in-depth analysis beyond a simple user count of those with or without access. Available bandwidth and costs may have temporal characteristics as users make use of one or more public, home or work locations for Internet access. Device or operating system choices unsuited for rural access may negatively affect the utility of the Internet. Lack of locally relevant content, tight coupling to local amenities or economies and fear of losing cultural values, may decrease the ability of Internet to develop communities. This calls for new multifaceted metrics that measure the issues raised in this section and attempt to evaluate the true impact of broadband connectivity on rural regions.

10.4 Scaling up solutions

In this dissertation we have primarily analysed traffic in a network in a single village with approximately 100 wireless radios connecting about 300 users. This network wirelessly connects a limited number of users clustered around a central hospital. Ultimately the aim of this and many other similar networks is to expand connectivity to a much larger set of surrounding users within the village and nearby villages. Macha Works already has a model to use resource containers connected with satellite links to reach remote sites; however, spreading connectivity

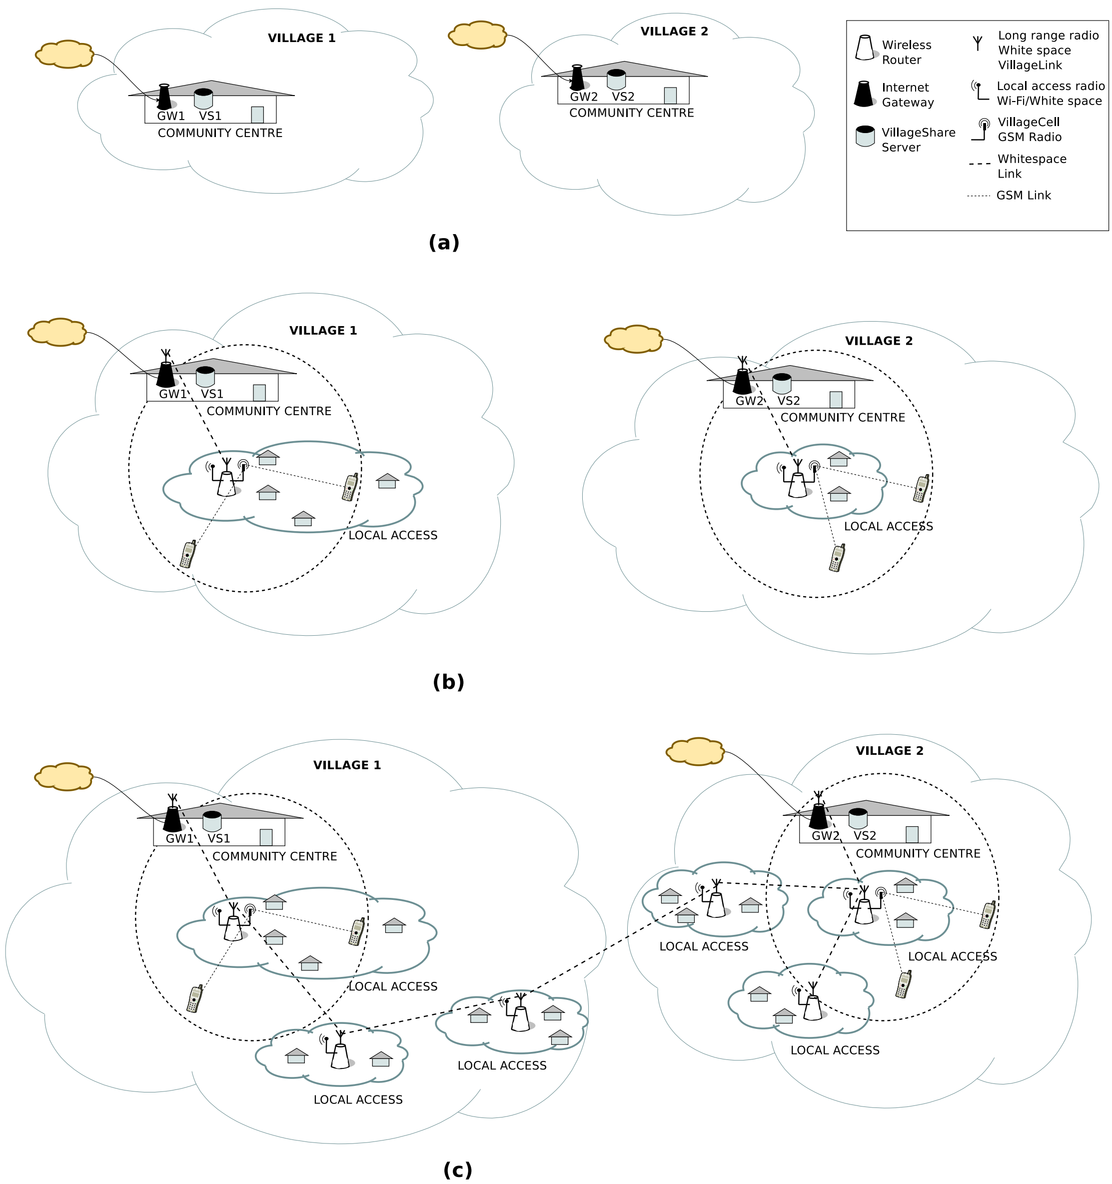

Figure 10.1: Growth of network in two villages until the point where they become linked

to a wider area around these central Internet points of presence is limited by the reach of the current WiFi system as well as the capacity and cost constraints of the Internet gateway.

The ultimate vision of our solution is to make use of white spaces to reach homes of users in the surrounding village and even adjacent villages. Figure 10.1 depicts the expansion we envisage using the VillageNet suite of technology solutions. Initial Internet provision, shown in Figure 10.1(a), would usually be to a single location, such as a community centre, where a kiosk/telecentre model is used with shared Internet access on public terminals. A VillageShare server would be placed in the community centre to facilitate free local content sharing. In Figure 10.1(b), expansion now begins to occur using White space links (VillageLink) to surrounding homes and places of work such as schools and health centres. Local access within a cluster of users is provided through WiFi or White spaces, depending on the size of the cluster needing coverage. VillageCell is installed to provide local free GSM access to the village. As more and more clusters are connected, eventually it is possible for adjacent villages, each with their own Internet gateway and VillageShare server, to be connected by long distance White space links as depicted in Figure 10.1(c). It may also be possible for a single network to expand to the point where a single Internet gateway is no longer able to provide sufficient capacity or the number of hops within the local network becomes high enough to restrict the local network capacity. In this case an additional

Internet gateway is added to the network to provide additional capacity and an additional VillageShare server is added to provide higher-speed sharing between nearby users.

In Chapter 6 we discuss a gateway-aware routing protocol that routes traffic to the optimal gateway in cases where multiple gateways are available and this system will suffice for Internet access in the scenario depicted in Figure 10.1(c). However, our VillageShare solution discussed in Chapter 7 only provides a solution for sharing between local users in a single village and between these village users and external users. In cases where multiple isolated villages utilize VillageShare, as in Figure 10.1(a) and Figure 10.1(b), a mechanism is required to facilitate efficient Inter-village sharing with minimal use of the slow Internet link. One possible solution is to cache all content shared between VillageShare servers on the global server before being sent to the VillageShare server with the destination user. This prevents resending the content from the source village VillageShare server to a user serviced by a VillageShare server in a different destination village. When Multiple VillageShare servers are connected via a local wireless network as in Figure 10.1(c), content shared with a user in an adjacent connected village should make use of the local wireless network of both villages to synchronize content. For this to occur, the VillageShare servers will need to be aware of the connection between the two local villages and route the inter-village file sharing traffic on the local network. This could be made possible by a VillageShare server discovery protocol, where all local VillageShare servers advertise their local IP addresses to the global VillageShare server. The local VillageShare servers could then probe the other local VillageShare servers on their local IP address once a day to discover if they are accessible on the local network.

In a similar way, VillageCell should also be able to create a local join between VillageCell systems that were originally located in two independent villages once a local wireless link is able to connect two VillageCell entities. Currently VillageCell requires FreeSwitch to manage the inter- and intra-base station call routing and OpenBTS for the GSM stack. In the case where there is only one VillageCell base station, both of these systems are located on a single VillageCell unit. In the case where there are multiple VillageCell base stations in a village, one base station will be designated to run the FreeSwitch service and the others will only run the OpenBTS GSM stack. When two villages are joined by a local wireless link, a trunk connection will be made between one village FreesSwitch service and the other village FreeSwitch service to allow call routing to occur between both villages without using the Internet gateway.

10.5 Generalizability of our analysis and solutions

In Section 10.2, we looked at how VillageNet can be extended to other access modalities. We now examine how widely applicable the Network performance and usage trends as well as VillageNet solutions are to other developing regions.

Our white space-based wireless network solution, VillageLink, was motivated by the fact that the spectrum allows for wide area wireless coverage in sparsely populated environments. In Chapter 8 we noted that rural regions in developing countries and even many developed countries usually have ample vacant spectrum that can be used by white space technologies. A population density map, shown in Figure 10.2, however, reveals that there are developing regions in the world where this assumption does not hold true. For example, high density rural populations can be found in India with an average population density of 371 per square kilometre or the Java province of Indonesia with a population density of 1064 per square kilometre. This is a stark contrast to countries in sub-Saharan Africa like Zambia with a population density of 17 per square kilometre. Other solutions may be more appropriate to address the high capacity and density of these areas such as WiFi mesh networks or voice and messaging systems running on Zigbee [142].

In general, however, the majority of developing countries have vast rural areas with low density populations which will be well suited to VillageLink.

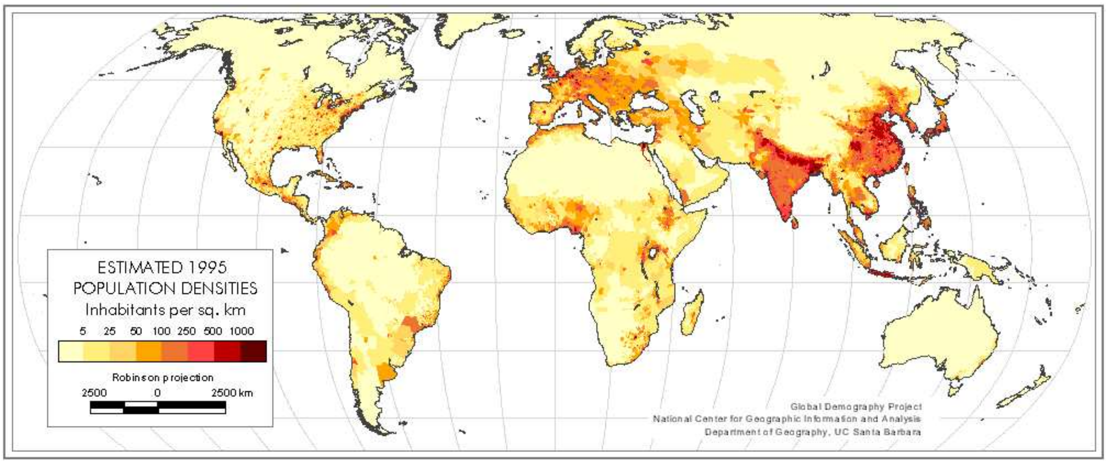

Figure 10.2: Worldwide population density showing the vast differences between developing regions in Africa and India.

Very few studies are able to gain access to data on Internet usage or performance at a worldwide level. The ITU indicators database is one such global-scale database that collects country-based information on metrics such as broadband subscriptions and available Internet bandwidth to compile the ITU report on growth of ICT globally [11, 12]. Another large-scale report at a multi-country level is the Ipoque report [151] that used data collected from 8 ISPs and 3 universities in 8 geographic regions of Africa, South America, Europe and the Middle East to analyse Internet traffic.

The ITU report releases very coarse grained data at a country level. The Ipoque report generalizes results even further and provides data at a regional level such as Southern Africa or Eastern Europe. These reports fail to reveal the large disparities that are often present within countries. For example, South Africa has a two-tier economy in which modern facilities, on-par with highly developed industrialized economies, are available in urban areas but not far away, a rural area may not have access to proper sanitation or Internet infrastructure. This is the motivation behind using the term developing region for our solutions rather than developing country.

Both reports show that generalizations about developing regions can be dangerous. For example, percentage of users with cell phone subscriptions were 125% in South Africa, 65% in Zambia and only 18% in Ethiopia in 2011. Only 2% of Ghana's population had access to an Internet connection above 2 Mbps as opposed to 28% of the population of Colombia. In the Ipoque study, 66% of the Internet traffic in Southern Africa was generated by peer-to-peer(P2P) file sharing, whereas only 44% of the Internet traffic in Northern Africa was P2P. None of these course-grain reports are sufficient to develop custom solutions that meet true user needs.

That said, the high prevalence of Facebook usage noted in Chapter 3 is also seen in other developing regions. Research on Internet usage in Cambodia and

Iraq in 2012 [53] showed that Facebook was the most visited web site. Analysis of web usage in a peri-urban primary school outside Bangalore India reveals that Facebook was the third most accessed URL [36]. Building applications for developing regions based on online social networks appears to be an overall good choice. However, in general our work has emphasized the importance of comprehensive analysis of user behaviour, current Internet usage patterns and ethnography to shape solutions for specific contexts.

10.6 Conclusion

No matter what technical solution is used, history shows us that technology is primarily a magnifier of existing institutional forces[162, 54] and is not transformative in and of itself. The Western world is heavily influenced by a pro-innovation bias [148] in which all innovation is assumed to be positive even though many innovations fail with little reported work on their failure. Successful development is best achieved when technology is combined with other basic positive economic entitlements such as education, training and access to credit instruments such as microfinancing [80]. Hence technology, ICT in our case, without the intent for true development, such as eradicating extreme poverty and hunger, achieving universal primary education or promoting gender equality, may entrench these social ills even further [32]. Technology projects should seek to amplify the impact of existing institutions already contributing successfully to development goals instead of attempting to lead with technology; technology is best employed in a supportive role to amplify intent to develop.

Chapter 11

Conclusion

Internet access has only reached one third of the world's population; the majority of the unreached population live in rural areas of the developing world. At an international level there have been concerted efforts to place the issue of reaching developing regions with the Internet on the agenda of United Nations conferences and plans. Examples of key UN initiatives are the Millenium Development Goals defined in 2000, the World Summit on the Information Society in 2005 and the World Telecommunication and Information Society Day in 2011. At a national level many countries such as Brazil, USA, South Africa, India, Australia and Canada have broadband strategies to address the access divides between urban and rural areas and increase overall Internet capacity at a national level. These efforts all recognize the powerful effect that ICTs can have on stimulating economic activity and providing improved services in domains such as health care and education.

Sporadic attempts to bring broadband connectivity to isolated areas in developing regions have been made, but a comprehensive multi-year evaluation of the challenges faced when deploying, running and maintaining ICTs in these areas, as well as analysis of the utility and usage of these services, is often lacking, especially in rural developing regions of Africa. We have carried out thorough multi-year qualitative and quantitative studies of rural networks providing Internet in rural areas of sub-Saharan Africa. Our analysis revealed unique challenges faced by these communities, ranging from concerns about the threat of the Internet to local culture to technical obstacles such as constant power failures, file upload failures and poor Windows performance compared to Linux. This analysis also revealed unique opportunities in the local community such as strong social cohesion. The analysis framework used to uncover these social-technical challenges and opportunities could also be used for commissioning future networks and assisting with iterative improvements once a network is operational.

This work provides a number of key contributions to the field of ICT for development. Our solutions address unique aspects of rural developing regions discovered through our analysis framework. We have captured a two-year data set of rural Internet usage from 300 users in a rural African village supplemented with extensive in-person and on-line interviews. This data revealed key differences with the developed world, specifically in the use of web 2.0. Further analysis of social media usage uncovered strong locality of interest amongst rural Internet denizens. The VillageShare and VillageCell solutions were developed to ensure that data traffic and voice calls sent between local users stay within the local network. Our VillageLink solution extends access to low-density areas, such as rural Africa, using white space technology. We carried out spectrum profiling in South Africa and Zambia and discovered that the complete range of white space spectrum is often available in rural African areas. Channel gain varies unpredictably across this available white space spectrum and our channel probing technique extends the 802.22 standard to select optimal white space channels.

There are a number of possible avenues for future work on VillageNet. Currently VillageCell is not robust to disconnections in the backhaul network. Once a connection to the PBX is broken, mobile phone clients connected to the OpenBTS base station will not be able to route calls. Auto-reconfiguration of VillageCell should allow the VillageCell unit closest to the base station to instantiate a new PBX instance. In addition, once data services are available on VillageCell, users should be able to upload content directly from their phone to VillageShare. VillageShare represents a localization solution for uploading or sharing static content. Localization could ultimately be extended to real-time content by sophisticated data-mining techniques that discover common motifs in outgoing and incoming data streams. The key challenge will be dealing with network encryption employed by most real-time services. VillageLink employs software-defined radios which offer a large amount of scope for modification of communication parameters to improve performance. For example, optimal sub-carrier assignment and cyclic prefix can be adjusted depending on the amount of delay spread in the channel. The modulation scheme or MAC protocol could also be modified in real time based on the sensed environment.

This research has had impact in a number of areas. Our 2010 publication on Internet usage and performance analysis of the wireless network in Macha, Zambia [90] has been cited 18 times (2 self-citations). The 2011 publication on traffic characterization and Internet usage in rural Africa [92] has been cited 15 times (5 self-citations). Our data sets have been used for comparative studies of proxy cache performance in Zambia, Cambodia and Iraq [53]. Our work has also been featured in a number of popular publications:

- UCSB Engineering Insights 2011: VillageNet poster
- UCSB Daily Nexus Oct 2011: Grant Finances Web Access in Rural Africa
- UCSB Convergence Magazine, September 2012: Bridging the Digital Divide

This work has also had social impact on the communities we have been working in. Our model of carrying out research in partnership with local Zambian researchers through co-authorship of publications and participatory research has expanded the skill-base of local Zambian ICT researchers and UCSB students alike. Our VillageShare and VillageCell localization solutions have provided bandwidth and cost-saving benefits to the community. In addition, these solutions have provided a more efficient mechanism for users to upload local content to the Internet.

This dissertation made use of elements of Paulo Freire's "participatory action research" [23] where community members were research partners. Qualitative data were gathered using on-line and in-person structured and unstructured interviews as well as co-reflection with local people living in the community by co-researchers. This process included a form of phenomenology, in which experiences are captured as knowledge, moving beyond simple positivism and recognizing the complexity of human and social problems. Kenneth Kuanda, the previous president of Zambia, reflected on the importance of knowledge gained through experience in Africa:

> Westerners have aggressive problem solving minds; Africans experience people.

This is a key lesson for future researchers in rural regions of Africa that are often culturally dissonant with Western culture. The traditional scientific method will often have limitations within deeply oral cultures in Africa in which new knowledge is not always discovered through systematic observation, measurement, and experiment, and the formulation, testing, and modification of hypotheses.

Our solutions have been tailored to the specific rural African contexts that we studied. The generalizability of our work depends on the number of common elements discovered in other communities such as bandwidth-constrained Internet gateways or locality of interest. The solutions can be used in a supportive role to efforts by communities or governments to use ICT to improve services in areas with poor Internet coverage. In countries where authoritarian governments are attempting to block or control the free flow of information on the Internet, our solutions are also capable of providing off-grid services to communities that wish to carry out voice or data communication outside of government control. However, this work carries with it the lesson that a comprehensive participatory-based analysis of the community will create a higher probability of solutions that actually match real user needs. Ultimately we hope that the lessons learnt from our analysis and solutions embodied in VillageNet will enable users in marginalized regions of the world to be more active participators in the social web and content contributors to knowledge repositories on the World Wide Web.

Bibliography

[1] *Khanya Project*. http://www.westerncape.gov.za/eng/pubs/public_info/F/148833/9 (accessed on 23 January, 2013).

[2] *List of device bandwidths*. http://en.wikipedia.org/wiki/List_of_device_bandwidths (accessed on 20 January 2013).

[3] *MXIT*. http://site.mxit.com (accessed on 23 January, 2013).

[4] *Putting Problems Before Solutions in Development*. http://www.ictworks.org/news/2011/08/03/putting-problems-solutions-development (accessed on 20 January 2013).

[5] *The HTTP archive*. http://httparchive.org (accessed on 20 February, 2013).

[6] *UCT Centre for ICT4D*. http://www.ict4d.cs.uct.ac.za (accessed on 20 February, 2013).

[7] *World Bank. 2011 World development indicators*. http://databank.worldbank.org/databank/download/WDIandGDF_excel.zip (accessed on 20 January 2013).

[8] *ITU-T Recommendation G.114, One-way transmission time*, tech. rep., International Telecommunication Union, 2003.

[9] *Regulations regarding licence exempt electronic communication networks*, tech. rep., South African Government Gazette No. 31289, July 2008.

[10] *2011 Facts and Figures*, tech. rep., International Telecommunication Union, 2011.

[11] *Measuring the Information Society*, tech. rep., International Telecommunication Union, 2011.

[12] *Measuring the information society*, tech. rep., International Telecommunications Union, 2012.

[13] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, *Link-level measurements from an 802.11b mesh network*, ACM SIGCOMM Computer Communication Review, 34 (2004), p. 121.

[14] J. C. Aker and I. M. Mbiti, *Mobile Phones and Economic Development in Africa*, Journal of Economic Perspectives, 24 (2010).

[15] R. Anderson, E. Blantz, D. Lubinski, E. O'Rourke, M. Summer, and Y. Yousoufian, *Smart Connect: last mile data connectivity for rural health facilities*, in NSDR'10, san Francisco, CA, June 2010.

[16] J. Backens, G. Mweemba, and G. van Stam, *A Rural Implementation of a 52 Node Mixed Wireless Mesh Network in Macha, Zambia*, in Africomm'09, Maputo, Mozambique, Dec. 2009.

[17] P. Ballantyne, *Collecting and propagating local development content*, tech. rep., Tanzania Commission for Science and Technology, 2002.

[18] G. Bernardi, P. Buneman, and M. K. Marina, *Tegola Tiered Mesh Network Testbed in Rural Scotland*, in WiNS-DR, San Francisco, CA, Sept. 2008.

[19] M. L. Best, *A global survey of spectrum license exemptions*, in Telecommunications Policy Research Conference, Washington, DC, Oct. 2006.

[20] P. Bhagwat, B. Raman, and D. Sanghi, *Turning 802.11 inside-out*, ACM SIGCOMM Computer Communication Review, 34 (2004), p. 33.

[21] J. Bicket, D. Aguayo, S. Biswas, and R. Morris, *Architecture and evaluation of an unplanned 802.11b mesh network*, in Mobicom'05, Cologne, Germany, Aug. 2005.

[22] S. Biswas and R. Morris, *ExOR: opportunistic multi-hop routing for wireless networks*, ACM SIGCOMM Computer Communication Review, 35 (2005).

[23] O. Borda, *Participatory (action) research in social theory: Origins and challenges*, in Handbook of action research: Participative inquiry and practice, SAGE Publications, 2001, pp. 27–37.

[24] P. Bremaud, *Markov Chains, Gibbs Fields, Monte Carlo Simulation and Queues*, Springer, 1999.

[25] E. Brewer, M. Demmer, B. Du, M. Ho, M. Kam, S. Nedevschi, J. Pal, R. Patra, S. Surana, and K. Fall, *The case for technology in developing regions*, Computer, 38 (2005), pp. 25–38.

[26] E. Brewer, M. Demmer, M. Ho, R. Honicky, J. Pal, M. Plauche, and S. Surana, *The Challenges of Technology Research for Developing Regions*, IEEE Pervasive Computing, 5 (2006), pp. 15–23.

[27] N. Brownlee and K. C. Claffy, *Understanding Internet traffic streams: Dragonflies and tortoises*, IEEE Communications Magazine, 40 (2002), pp. 110–117.

[28] I. Buskens and S. Earl, *Research for change: Outcome mapping's contribution to emancipatory action research in Africa*, Action Research, 6 (2008), pp. 171–192.

[29] L. Butgereit, *How Dr Math reaches pupils with competitions and computer games by using MXit*, in IST-Africa 2009 Conference and Exhibition, Kampala, Uganda, May 2009.

[30] M. Calabrese and B. Scott, *Measuring TV 'White Space' Available for Unlicensed Wireless Broadband*, New America Foundation and Free Press Analysis, (2006).

[31] A. S. Canagarajah, *"Nondiscursive" Requirements in Academic Publishing, Material Resources of Periphery Scholars, and the Politics of Knowledge Production*, Written Communication, 13 (1996), pp. 435–472.

[32] P. Carmody, *The Informationalization of Poverty in Africa? Mobile Phones and Economic Structure*, Information Technologies & International Development, 8 (2012), pp. 1–17.

[33] K. Chebrolu, B. Raman, and S. Sen, *Long-Distance 802.11b Links: Performance Measurements and Experience*, in MobiCom'06, Los Angeles, CA, Sept. 2006.

[34] J. Chen, *Re-architecting Web and Mobile Information Access for Emerging Regions*, PhD thesis, New York University, 2011.

[35] J. Chen, S. Amershi, A. Dhananjay, and L. Subramanian, *Comparing web interaction models in developing regions*, in ACM Symposium on Computing for Development, London, UK, Dec. 2010, ACM.

[36] J. Chen, D. Hutchful, W. Thies, and L. Subramanian, *Analyzing and accelerating web access in a school in peri-urban India*, in WWW'11, New York, NY, Mar. 2011.

[37] J. Chen, L. Subramanian, and E. Brewer, *Sms-based web search for low-end mobile devices*, in Mobicom'10, Chicago, IL, Sept. 2010.

[38] J. Chen, L. Subramanian, and J. Li, *RuralCafe: web search in the rural developing world*, in WWW'09, Madrid, Spain, Apr. 2009.

[39] L. J. Chen, T. Sun, G. Yang, M. Y. Sanadidi, and M. Gerla, *End-to-End Asymmetric Link Capacity Estimation*, in Networking'05, Waterloo, Canada, May 2005.

[40] N. Clark, *Facebook overtakes Googles.* The Independent, 18 March, 2010, http://www.independent.co.uk/life-style/gadgets-and-tech/news/facebook-overtakes-google-1923102.html (accessed on 23 January, 2013).

[41] E. Colson, *Tonga religious life in the twentieth century*, Bookworld Publishing House, 2006.

[42] S. Corbett, *Can the Cellphone Help End Global Poverty?* The New York Times, 13 April, 2008, http://www.nytimes.com/2008/04/13/magazine/13anthropology-t.htm (accessed on 20 January 2013).

[43] C. Cordeiro, K. Challapali, D. Birru, B. Manor, and S. Diego, *IEEE 802.22: an introduction to the first wireless standard based on cognitive radios*, Journal of communications, 1 (2006), pp. 38–47.

[44] F. D. Davis, *Perceived Usefulness, Perceived Ease of Use, and User Acceptance of Information Technology*, MIS Quarterly, 13 (1989), pp. 319–340.

[45] M. de Bruijn, F. Nyamnjoh, and I. Brinkman, *Mobile phones: The new talking drums of everyday Africa*, Langaa Rpcig, 2009.

[46] D. S. J. De Couto, D. Aguayo, J. Bicket, and R. Morris, *A high-throughput path metric for multi-hop wireless routing*, Wireless Networks, 11 (2005), pp. 419–434.

[47] N. Dell, S. Venkatachalam, D. Stevens, P. Yager, and G. Borriello, *Towards a point-of-care diagnostic system: automated analysis of immunoassay test data on a cell phone*, in NSDR'11, Washington, D.C, June 2011.

[48] B. DeRenzi, G. Borriello, J. Jackson, V. S. Kumar, T. S. Parikh, P. Virk, and N. Lesh, *Mobile Phone Tools for FieldBased Health care Workers in LowIncome Countries*, Mount Sinai Journal of Medicine, 78 (2011), pp. 406–418.

[49] J. Domènech, A. Pont, J. Sahuquillo, and J. A. Gil, *A user-focused evaluation of web prefetching algorithms*, Computer communications, 30 (2007), pp. 2213–2224.

[50] R. Draves, J. Padhye, and B. Zill, *Routing in multi-radio, multi-hop wireless mesh networks*, in Mobicom'04, Philadelphia, PA, Oct. 2004.

[51] B. Du, M. Demmer, and E. Brewer, *Analysis of WWW traffic in Cambodia and Ghana*, in WWW'06, Edinburgh, UK, May 2006.

[52] N. Ehsan, M. Liu, and R. J. Ragland, *Evaluation of performance enhancing proxies in internet over satellite*, International Journal of Communication Systems, 16 (2003), pp. 513–534.

[53] A. Fahad, Z. Chen, K. Shen, and J. Bigham, *An Evaluation of Web Acceleration Techniques for the Developing World*, in NSDR'12, Boston, MA, June 2012.

[54] G. Farrell, S. Isaacs, and M. Trucano, *The NEPAD e-Schools Demonstration Project: A Work in Progress: A Public Report*, tech. rep., Commonwealth of Learning; World Bank, Sept. 2007.

[55] R. Flickenger, C. Aichele, C. Fonda, J. Forster, I. Howard, T. Krag, and M. Zennaro, *Wireless Networking in the Developing World*, Limehouse Book Sprint Team, first ed., Jan. 2006.

[56] S. Geman and D. Geman, *Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 6 (1984), pp. 721–741.

[57] J. Goebel, T. Holz, and C. Willems, *Measurement and analysis of autonomous spreading malware in a university environment*, Detection of Intrusions and Malware, and Vulnerability Assessment, 4579 (2007), pp. 109–128.

[58] R. Gomez and S. Pather, *ICT Evaluation: Are We Asking the Right Questions?*, The Electronic Journal of Information Systems in Developing Countries, 50 (2011).

[59] G. Gu, R. Perdisci, J. Zhang, and W. Lee, *BotMiner: Clustering analysis of network traffic for protocol-and structure-independent botnet detection*, in USENIX Security Symposium (Security'08), San Jose, CA, July 2008.

[60] L. Guo, S. Chen, Z. Xiao, and X. Zhang, *Analysis of multimedia workloads with implications for internet streaming*, in WWW'05, Chiba, Japan, May 2005.

[61] L. Guo and I. Matta, *The war between mice and elephants*, in ICNP'01, Riverside, CA.

[62] S. Guo, M. H. Falaki, E. A. Oliver, S. U. Rahman, A. Seth, M. A. Zaharia, U. Ismail, and S. Keshav, *Design and Implementation of the KioskNet System*, in ICTD'07, Bangalore, India, Dec. 2007.

[63] K. Gush, R. de Villiers, R. Smith, and G. Cambridge, *ICTs for Global Development and Sustainability*, IGI Global, Sept. 2010.

[64] S. Habeenzu, *Zambia ICT Sector Performance review 2009/2010*, Research ICT Africa, 2 (2010).

[65] B. Hajek, *Cooling Schedules for Optimal Annealing*, Mathematics of Operations Research, 13 (1988), pp. 311–319.

[66] S. B. Handurukande, A. M. Kermarrec, F. Le Fessant, and L. Massoulié, *Exploiting semantic clustering in the eDonkey P2P network*, in Proceedings of the 11th workshop on ACM SIGOPS European workshop: beyond the PC - EW11, New York, NY, Sept. 2004.

[67] B. Harden, *Africa: dispatches from a fragile continent*, Harper Collins, 1990.

[68] E. Hargittai and S. Shafer, *Differences in Actual and Perceived Online Skills: The Role of Gender*, Social Science Quarterly, 87 (2006), pp. 432–448.

[69] L. Harris, *Mind the ICT skills gap.* ITweb, 22 October, 2013, http://www.itweb.co.za/index.php?option=com_content&view=article&id=59439 (accessed on 20 February, 2013).

[70] K. Harrison, S. M. Mishra, and A. Saha, *How Much White-Space Capacity Is There?*, in DySpan'10, Singapore, Apr. 2010.

[71] M. Hassan, A. Nayandoro, and M. Atiquzzaman, *Internet telephony: services, technical challenges, and products*, Communications Magazine, IEEE, 38 (2000), pp. 96–103.

[72] R. Heeks, *ICT4D 2.0: The Next Phase of Applying ICT for International Development*, Computer, 41 (2008), pp. 26–33.

[73] K. Heimerl and E. Brewer, *The Village Base Station*, in NSDR'10, San Francisco, CA,, June 2010.

[74] T. Henderson, D. Kotz, and I. Abyzov, *The Changing Usage of a Mature Campus-Wide Wireless Network*, in MobiCom'04, Philadelphia, PA, Sept. 2004.

[75] M. Hilbert, *Digital gender divide or technologically empowered women in developing countries? A typical case of lies, damned lies, and statistics*, Women's Studies International Forum, 34 (2011), pp. 479–489.

[76] J. D. Holm and L. Malete, *The Asymmetries of University Partnerships between Africa and the Developed World: Our experience in Botswana*, in Going Global 4 - The British Councils International Education Conference, London, UK, 2010.

[77] J. Horrigan, *Home Broadband Adoption 2006*, tech. rep., Pew Internet & American Life Project, May 2006.

[78] H. Horst and D. Miller, *The Cell Phone: an Anthropology of Communication*, Berg, 2006.

[79] T. Horvath, H. Azman, G. E. Kennedy, and G. W. Rutherford, *Mobile phone text messaging for promoting adherence to antiretroviral therapy in patients with HIV infection*, Cochrane Database of Systematic Reviews, 3 (2012).

[80] M. Hossain, *Credit for alleviation of rural poverty: The Grameen Bank in Bangladesh*, International Food Policy Research Institute, Washington, DC, 1988.

[81] IETF, *RFC 4168, The Stream Control Transmission Protocol (SCTP) as a Transport for the Session Initiation Protocol (SIP)*, 2005.

[82] International Telecommunications Union, *ITU-T Recommendation G.107, "The E-Model, a computational model for use in transmission planning"*, Dec. 1998.

[83] S. Isaacman and M. Martonosi, *Low-infrastructure methods to improve internet access for mobile users in emerging regions*, in WWW'11, New York, NY, Mar. 2011, p. 473.

[84] J. Ishmael, S. Bury, D. Pezaros, and N. Race, *Deploying Rural Community Wireless Mesh Networks*, IEEE Internet Computing, 12 (2008), pp. 22–29.

[85] R. Jain, D. M. Chiu, and W. R. Hawe, *A quantitative measure of fairness and discrimination for resource allocation in shared computer system*, Eastern Research Laboratory, Digital Equipment Corporation, 1984.

[86] M. Jensen, *Broadband in Brazil: A Multipronged Public Sector Approach to Digital Inclusion*, tech. rep., InfoDev, 2011.

[87] R. Jensen, *The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector*, The Quarterly Journal of Economics, 122 (2007), pp. 879–924.

[88] A. Jhunjhunwala, A. Ramachandran, and A. Bandyopadhyay, *n-Logue: The story of a rural service provider in India*, The Journal of Community Informatics, 1 (2004), pp. 30–38.

[89] D. L. Johnson, *Evaluation of a single radio rural mesh network in South Africa*, in ICTD'06, Bangalore, India, Dec. 2007.

[90] D. L. Johnson, E. M. Belding, K. C. Almeroth, and G. van Stam, *Internet usage and performance analysis of a rural wireless network in Macha, Zambia*, in NSDR'10, San Francisco, CA, June 2010.

[91] D. L. Johnson, N. Ntlatlapa, and C. Aichele, *A simple pragmatic approach to mesh routing using BATMAN*, in IFIP International Symposium on Wireless Communications and Information Technology in Developing Countries, Pretoria, South Africa, Oct. 2008.

[92] D. L. Johnson, V. Pejovic, E. M. Belding, and G. van Stam, *Traffic characterization and internet usage in rural Africa*, in WWW'11, Hyperabad, India, Apr. 2011, ACM, pp. 493–502.

[93] Johnson, D L, *Performance analysis of mesh networks in indoor and outdoor wireless testbeds*, PhD thesis, University of Pretoria, 2007.

[94] Johnson, D L and Pejovic, Veljko and Belding, E M and van Stam, Gertjan, *Network traffic locality in a rural African village*, in ICTD'12, Atlanta, GA, 2012.

[95] T. Karsenti, S. Collin, and T. Harper-Merrett, *Successes and Challenges from 87 African Schools Pedagogical Integration of ICT*, tech. rep., IDRC, Ottawa, 2011.

[96] B. Kauffmann, F. Baccelli, A. Chaintreau, V. Mhatre, K. Papagiannaki, and C. Diot, *Measurement-Based Self Organization of Interfering 802.11 Wireless Access Networks*, in INFOCOM'07, Anchorage, AK, May 2007.

[97] T. L. M. Kennedy, S. Aaron, A. T. Wells, and B. Wellman, *Networked Families*, tech. rep., Pew Internet and American Life Project, Oct. 2008.

[98] K. Kiri and D. Menon, *For profit rural kiosks in India: achievements and challenges*, Information for Development, 4 (2006), pp. 14–16.

[99] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, *Optimization by Simulated Annealing*, Science, 220 (1983), pp. 671–680.

[100] S. Kopparty, S. Krishnamurthy, M. Faloutsos, and S. Tripathi, *Split TCP for mobile ad hoc networks*, Globecom'02, 1 (2002), pp. 138–142.

[101] A. Kumar, J. Chen, M. Paik, and L. Subramanian, *ELMR: Efficient Lightweight Mobile Records*, in MobiHeld'09, Barcelona, Spain, Aug. 2009.

[102] A. Kumar, N. Rajput, D. Chakraborty, S. K. Agarwal, and A. A. Nanavati, *WWTW: The World Wide Telecom Web*, in NSDR'09, Kyoto, Japan, Aug. 2007.

[103] D. Kumar, D. Martin, and J. O'Neill, *The Times They Are A-Changin: Mobile Payments in India*, in CHI'11, Vancouver, Canada, May 2011.

[104] R. Kumar and M. L. Best, *Social Impact and Diffusion of Telecenter Use: A Study from the Sustainable Access in Rural India Project*, The Journal of Community Informatics, 2 (2006), pp. 1–22.

[105] R. Kumar, J. Novak, and A. Tomkins, *Structure and evolution of online social networks*, Springer, New York, NY, 2010.

[106] S. Lakshmanan, R. Sivakumar, and K. Sundaresan, *Multi-gateway association in wireless mesh networks*, Ad Hoc Networks, 7 (2009), pp. 622–637.

[107] L. Lessig, *The future of ideas: The fate of the commons in a connected world*, Vintage, 2002.

[108] L. Levison, W. Thies, and S. Amarasinghe, *Providing Web search capability for low-connectivity communities*, in International Symposium on Technology and Society, Raleigh, NC, June 2002.

[109] H. Lundgren, K. N. Ramachandran, E. M. Belding-Royer, K. C. Almeroth, M. Benny, A. Hewatt, A. Touma, and A. Jardosh, *Experiences from the design, deployment, and usage of the UCSB meshnet testbed*, IEEE Wireless Communications, 13 (2006), pp. 18–29.

[110] M. Ma and D. H. K. Tsang, *Joint Design of Spectrum Sharing and Routing with Channel Heterogeneity in Cognitive Radio Networks*, Physical Communication, 2 (2009), pp. 127–137.

[111] G. Macola, *One Zambia, Many Histories: Towards a History of Post-colonial Zambia*, Brill, 2008.

[112] M. Mandioma, *Rural Internet Connectivity: A Deployment in Dwesa-Cwebe, Eastern Cape, South Africa*, Master's thesis, University of Fort Hare, Nov. 2007.

[113] M. A. J. Manschont and C. Stroek, *Stimulation of Local Content Generation in Rural Africa*, in IST-Africa'09, Lake Victoria, Uganda, May 2009.

[114] M. Marais, *An analysis of the factors affecting the sustainability of ICT4D initiatives*, in IDIA2011, Lima, Peru, Oct. 2011.

[115] D. Maranz, *African Friends and Money Matters: Observations from Africa*, Publications in Ethongraphy, 31 (2001).

[116] M. Masonta, D. L. Johnson, and M. Mzyece, *The White Space Opportunity in Southern Africa: Measurements with Meraka Cognitive Radio Platform*, in e-Infrastructure and e-Services for Developing Countries, vol. 92

of Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, Springer, 2012, pp. 64–73.

[117] K. W. Matthee, G. Mweemba, A. V. Pais, G. van Stam, and M. Rijken, *Bringing Internet connectivity to rural Zambia using a collaborative approach*, in ICTD'07, Bangalore, India, Dec. 2007.

[118] V. Mhatre, K. Papagiannaki, and F. Baccelli, *Interfernece Mitigation through Power Control in High Density 802.11 WLANs*, in INFOCOM'07, Anchorage, AK, May 2007.

[119] S. Milgram, *The small world problem*, Psychology today, 2 (1967), pp. 60–67.

[120] S. Mishra, J. Hwang, D. Filippini, R. Moazzami, L. Subramanian, and T. Du, *Economic Analysis of Networking Technologies for Rural Developing Regions*, Lecture Notes in Computer Science, 3828 (2005), pp. 184–194.

[121] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee, *Measurement and analysis of online social networks*, in IMC '07, New York, NY, Oct. 2007, ACM Press.

[122] S. Mitra, *Minimally invasive education for mass computer literacy*, in Conference on Research in Distance and Adult Learning in Asia, Pokfulam, Hong Kong, June 2000.

[123] T. Moscibroda, R. Wattenhofer, and Y. Weber, *Protocol Design Beyond Graph-Based Models*, in HotNets'06, 2006.

[124] J. Mpala and G. van Stam, *OpenBTS, a GSM experiment in rural Zambia*, in Africomm'12, Yaounde, Cameroon, 2012.

[125] D. Nandiraju, L. Santhanam, N. Nandiraju, and D. Agrawal, *Achieving load balancing in wireless mesh networks through multiple gateways*, in IEEE International Conference on Mobile Adhoc and Sensor Systems (MASS), Vancouver, Canada, Oct. 2006.

[126] N. Ndlovu, M. Thinyane, and A. Terzoli, *Deployment and Extension of a Converged WiMAX/WiFi Network for Dwesa Community Area South Africa*, in Southern African Telecommunication Networks and Applications Conference, Ezulweni, Swaziland, Aug. 2009.

[127] A. M. Nungu, *Analysis of the Serengeti Broadband Network*, tech. rep., KTH, Stockholm, Sweden.

[128] G. Nychis, T. Hottelier, Z. Yang, S. Seshan, and P. Steenkiste, *Enabling MAC protocol implementations on software-defined radios*, in NSDI'09, Boston, MA, Apr. 2009.

[129] B. Omwenga and P. Githinkji, *Buzzenger: Asynchronous (Time-Sliced) Missed Call Duration Messaging*, in NSDR'12, Boston, MA, 2012.

[130] W. Ong, *Orality and literacy*, Routledge, 2002.

[131] J. Onnela and S. Arbesman, *Geographic constraints on social network groups*, PLoS one, 6 (2011).

[132] N. Patel, D. Chittamuru, A. Jain, P. Dave, and T. S. Parikh, *Avaaj Otalo - A Field Study of an Interactive Voice Forum for Small Farmers in Rural India*, in CHI'10, Atlanta, GA, Apr. 2010.

[133] F. Peng, L. Wu, and V. Leung, *Cross-layer enhancement of TCP split-connections over satellites links*, International Journal of Satellite Communications and Networking, 24 (2006), pp. 405–418.

[134] A. Pentland, R. Fletcher, and A. Hasson, *DakNet: rethinking connectivity in developing nations*, Computer, 37 (2004), pp. 78–83.

[135] J.-P. Platteau, *The Evolutionary Theory of Land Rights as Applied to Sub-Saharan Africa: A Critical Assessment*, Development and Change, 27 (1996), pp. 29–86.

[136] J. Postill, *Localizing the internet beyond communities and networks*, New Media & Society, 10 (2008), pp. 413–431.

[137] C. Qiang, C. Rossotto, and K. Kimura, *Economic impacts of broadband*, World Bank, 2009.

[138] R. Raghavendra, P. A. K. Acharya, E. M. Belding, and K. C. Almeroth, *MeshMon: a multi-tiered framework for wireless mesh network monitoring*, Wireless Communications and Mobile Computing, 11 (2010), pp. 1182–1196.

[139] K. N. Ramachandran, E. M. Belding, K. C. Almeroth, and M. M. Buddhikot, *Interference-Aware Channel Assignment in Multi-Radio Wireless Mesh Networks*, in INFOCOM'06, Barcelona, Spain, Apr. 2006.

[140] K. N. Ramachandran, I. Sheriff, and E. M. Belding, *Routing stability in static wireless mesh networks*, Lecture Notes in Computer Science, 4427 (2007), p. 73.

[141] B. Raman and K. Chebrolu, *Design and Evaluation of a new MAC Protocol for Long-Distance 802.11 Mesh Networks*, in MobiCom'05, Cologne, Germany, 2005.

[142] B. Raman and K. Chebrolu, *Lo3 : Low-cost, Low-power, Local Voice and Messaging for Developing Regions*, in NSDR'09, Big Sky, Montana, Oct. 2009.

[143] A. Raniwala and T.-c. Chiueh, *Architecture and Algorithms for an IEEE 802.11-Based Multi-Channel Wireless Mesh Network*, in INFOCOM'05, vol. 00, Miami, FL, Mar. 2005.

[144] A. Reda, S. Panjwani, and E. Cutrell, *Hyke: a low-cost remote attendance tracking system for developing regions*, in NSDR'11, Washington, D.C, June 2011.

[145] M. C. Reddy, S. Purao, and M. Kelly, *Developing IT Infrastructure for Rural Hospitals: A Case Study of Benefits and Challenges of Hospital-to-Hospital Partnerships.*, Journal of the American Medical Informatics Association : JAMIA, 15 (2008), pp. 554–8.

[146] A. F. Ribadeneira, *An Analysis of the MOS Under Conditions of Delay, Jitter and Packet Loss and an Analysis of the Impact of Introducing Piggybacking and Reed Solomon FEC*, PhD thesis, Georgia State University, 2007.

[147] T. Roberts, *Appropriating Technology via Rural Hackspaces.* http://laptopburns.wordpress.com/2011/08/04/appropriating-technology-via-rural-hackspaces/ (accessed on 23 February, 2013).

[148] E. M. Rogers, *Diffusion of innovations*, Simon and Schuster, 1995.

[149] W. Ross, *Kenyan farmer lauds internet as saviour of potato crop.* BBC News, 17 March, 2010, http://news.bbc.co.uk/2/hi/8569125.stm (accessed on 23 February 2013).

[150] K. Roux and M. Marais, *Design for Sustainability: Rural Connectivity with Village Operators*, in IEEE Global Humanitarian Technology Conference, Seattle, WA, Oct. 2011.

[151] H. Schulze and K. Mochalski, *Internet Study 2008/2009*, tech. rep., IPOQUE, 2009.

[152] S. Sen, S. Kole, and B. Raman, *Rural Telephony: A Socio-Economic Case Study*, in ICTD'06, Berkeley, CA, May 2006.

[153] A. Seth, D. Kroeker, M. A. Zaharia, S. Guo, and S. Keshav, *Low-cost communication for rural internet kiosks using mechanical backhaul*, in MobiCom'06, Los Angeles, CA, USA, 2006.

[154] A. Sharma Grover and E. Barnard, *The Lwazi community communication service: design and piloting of a voice-based information service*, in WWW'11, Hyperabad, India, Apr. 2011.

[155] K. Sheneberger and G. van Stam, *Relatio: An examination of the relational dimension of resource allocation*, Economics and Finance Review, 1 (2011), pp. 26–33.

[156] A. Sheth, S. Nedevschi, R. Patra, S. Surana, L. Subramanian, and E. Brewer, *WiLDNet: Design and Implementation of High Performance WiFi Based Long Distance Networks*, in NSDI, Anchorage, AK, Apr. 2007.

[157] H. Spang-Hanssen, *Public international computer network law issues*, DJOF Publishing, 2006.

[158] R. Srinivasan, *Rethinking Digital Cultures and Divides: The Case for Reflective Media*, The Information Society, 28 (2012), pp. 24–36.

[159] I. Stojmenovic, *Handbook of wireless networks and mobile computing*, Wiley-Interscience, 2002.

[160] S. Surana, R. Patra, S. Nedevschi, M. Ramos, L. Subramanian, Y. Ben-David, and E. Brewer, *Beyond pilots: Keeping rural wireless networks alive*, in NSDI, San Francisco, CA, 2008.

[161] S. Tajima, T. Higashino, N. Funabiki, and S. Yoshida, *An Internet gateway access-point selection problem for wireless infrastructure mesh networks*, in FMUIT'06, Nara, Japan, May 2006.

[162] K. Toyama, *Technology as amplifier in international development*, in iConference'11, Seattle, WA, Feb. 2011.

[163] T. Unwin, *ICT4D: Information and communication technology for development*, Cambridge University Press, 2009.

[164] P. van Hoorik and F. Mweetwa, *Use of Internet in Rural Areas of Zambia*, in IST-Africa'08, Windhoek, Namibia, May 2008.

[165] G. van Stam and F. Mweetwa, *Community Radio Provides Elderly a Platform to Have Their Voices Heard in rural Macha, Zambia*, The Journal of Community Informatics, 8 (2012).

[166] W. W. Vithanage and A. S. Atukorale, *Bassa: a time shifted web caching system for developing regions*, in NSDR'11, Washington DC, June 2011.

[167] M. Vutukuru, H. Balakrishnan, and K. Jamieson, *Cross-layer Wireless Bit Rate Adaptation*, in ACM SIGCOMM'09, Barcelona, Spain, Aug. 2009.

[168] A. Waldman, *Indian Soybean Farmers Join the Global Village.* New York Times, 1 January, 2004, http://www.nytimes.com/2004/01/01/world/india-s-soybean-farmers-join-the-global-village.html (accessed on 20 January 2013).

[169] M. Wang, R. Shen, R. Tong, F. Yang, and P. Han, *Mobile Learning with Cellphones and PocketPCs*, Advances in web-based learning - ICWL 2005, 3583/2005 (2005).

[170] C. Wilson, B. Boe, A. Sala, K. P. Puttaswamy, and B. Y. Zhao, *User interactions in social networks and their implications*, in 4th ACM European conference on Computer systems (EuroSys), Nuremberg, Germany, Apr. 2009.

[171] M. Wittie, V. Pejovic, L. Deek, K. C. Almeroth, and B. Y. Zhao, *Exploiting Locality of Interest in Online Social Network*, in CoNEXT'10, Philadelphia, PA, Dec. 2010.

[172] S. P. Wyche, T. N. Smyth, M. Chetty, P. M. Aoki, and R. E. Grinter, *Deliberate interactions: characterizing technology use in Nairobi, Kenya*, in Conference on Human Factors in Computing Systems, Atlanta, GA, Apr. 2010.

[173] L. Yang, W. Hou, L. Cao, B. Y. Zhao, and H. Zheng, *Supporting Demanding Wireless Applications with Frequency-agile Radios*, in 7th USENIX

conference on Networked systems design and implementation (NSDI), San Jose, CA, Apr. 2010.

[174] Y. Yang, J. Wang, and R. Kravets, *Designing routing metrics for mesh networks*, in WiMesh'05, Santa Clara, CA, Sept. 2005.

Appendices

Appendix A

Facebook social graph

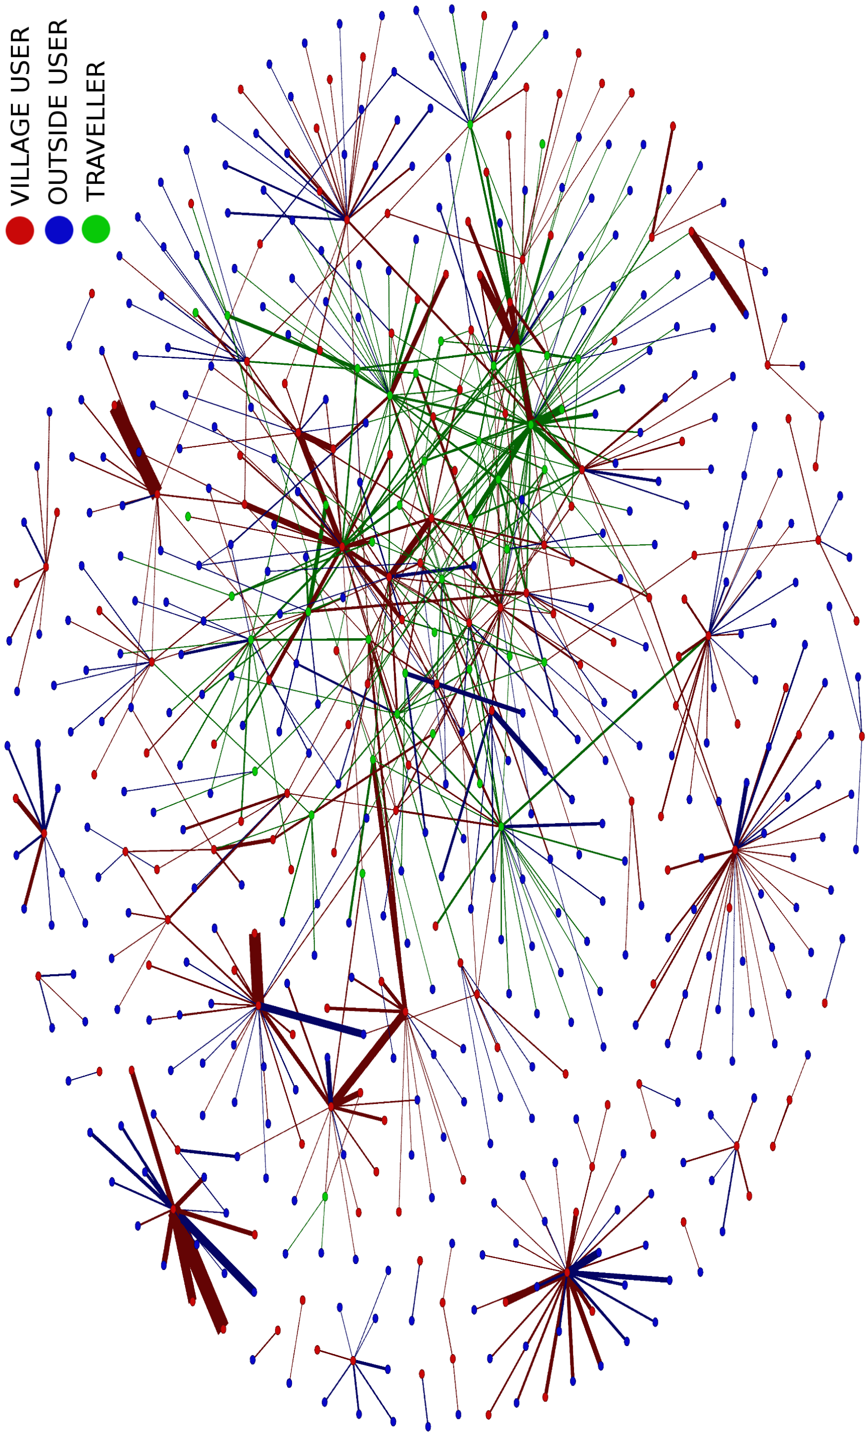

Figure A.1: Social graph of Facebook instant message exchanges.