International Journal of Communication 6 (2012), 2467–2491 1932–8036/20120005

Broadband Adoption
IJoC

The Bandwidth Divide: Obstacles to Efficient Broadband Adoption in Rural Sub-Saharan Africa

VELJKO PEJOVIC[*]
DAVID L. JOHNSON
MARIYA ZHELEVA
ELIZABETH BELDING
LISA PARKS
University of California, Santa Barbara

GERTJAN VAN STAM
LinkNet, Macha, Zambia

[*] Funding for this work was in part through National Science Foundation Network Science and Engineering Award CNS-1064821. The James D. Kline fellowship for work that promotes international understanding and the Council for Scientific and Industrial Research, Pretoria, South Africa, funded visits to Dwesa, South Africa, and Macha, Zambia. The authors would like to thank LinkNet, Macha, Zambia, for their cooperation on this project, specifically Fred Mweetwa, CEO Macha Works, for facilitating the interviews; the Meraka Institute and University of Fort Hare in South Africa for organizing the visit to Dwesa and facilitating the interviews; and Professor Jennifer Earl from the University of Arizona for assisting with the interview design and analysis. The authors also thank anonymous reviewers for their valuable comments. Please send e-mail correspondence to veljko@cs.ucsb.edu, davidj@cs.ucsb.edu, and mariya@cs.ucsb.edu.

Veljko Pejovic: veljko@cs.ucsb.edu
David L. Johnson: davidj@cs.ucsb.edu
Mariya Zheleva: mariya@cs.ucsb.edu
Elizabeth Belding: ebelding@cs.ucsb.edu
Lisa Parks: parks@filmandmedia.ucsb.edu
Gertjan van Stam: gertjan.vanstam@worksgroup.org
Date submitted: 2012–09–04

The Bandwidth Divide: Obstacles to Efficient Broadband Adoption in Rural Sub-Saharan Africa

Current metrics for evaluating Internet adoption capture the percentage of people with physical access to the Internet and provide a coarse understanding of actual usage. The factors for Internet adoption, however, are related not only to the provision of connectivity but also to individuals' personal experience. We concentrate on rural sub-Saharan Africa, and through network traffic analysis and social surveys we find that the location of access, connectivity speeds, and the cost of the connection together with the overall context in which the usage happens severely impact online behavior. Thus, we identify a set of metrics that describe individual perceptions of the Internet and provide an in-depth understanding of Internet usage patterns to identify obstacles to Internet adoption.

Internet connectivity is an essential factor for progress of any nation. Access to information can improve human development in areas such as education, health care, economy, and political freedom. Unfortunately, Internet access is unevenly distributed across the globe. While Internet penetration reaches staggering numbers in some areas, even basic connectivity is lacking in many developing regions. Global statistics show that in 2011, developed countries[1] had Internet penetration higher than 73% (International Telecommunication Union [ITU], 2011). In the developing world during the same year, however, only 26% of individuals were connected to the Internet.[2] Moreover, the variation between regions can be quite drastic. Sub-Saharan Africa, for example, significantly lags even behind other developing regions; penetration rates in countries such as the Democratic Republic of Congo, Liberia, Niger, and Ethiopia are less than 1% of the population.

The importance of Internet connectivity in developing regions, however, is immense. Rural areas of the developing world have few resources such as libraries, and skilled workers tend to migrate to more affluent, industrialized areas. Internet connectivity provides access to knowledge, which is crucial for social and economic development. Evidence from the fishing industry in India (Abraham, 2006) and sunflower farming in Zambia (Matthee, Mweemba, Pais, van Stam, & Rijken, 2007) shows how simple

---

[1] In this article, and in the International Telecommunication Union (ITU) statistics, 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. All other countries are termed *developing countries*. We are aware of the limitations of such a coarse-gained classification, and we use the terms strictly as a statistical convenience.

[2] The percentage of individuals using the Internet in the ITU statistics is based on results from national household surveys. For many developing countries such surveys do not exist, and the usage numbers were estimated using hot-deck imputation, where data from other countries with similar characteristics was used instead.

access to information can dramatically improve conditions in impoverished regions. Previous research shows that, everything else aside, access to information and communication technologies (ICTs) improves the gross domestic product of a country by about 1% (Waverman, Meschi, & Fuss, 2005). Access to ICTs also can have deep social impacts. For instance, Internet access enables the spread of communication and information, which can bring about political freedoms and strengthen human agency. It is not surprising that Internet censorship was a key point of struggle during the Arab Spring of 2010–2011. Finally, through their unique affordances, ICTs can be used to overcome social problems such as gender inequality in developing countries (Hilbert, 2011).

Internet penetration in the developing world lags behind the developed world for many reasons: lack of supporting infrastructure (roads and electricity), outdated regulatory frameworks, and affordability, to name a few (Brewer et al., 2005). However, a clear difference also exists between the developed and the developing world in the level of urbanization. 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 the HDI) are 70.5% rural (World Bank, 2011).

Rural areas in both developed and developing countries almost always experience lower levels of Internet penetration. The difference can be dramatic, such as in the case of South Africa (Internet penetration is 4.6% in rural areas and 21.8% in urban areas) and in Morocco (40.4% in rural areas and 75.6% in urban areas) (ITU, 2011). The fact that more than 3 billion people live in rural areas of the developing world, where the connectivity is poor, calls for careful examination of problems that prevent further connectivity expansion and for reconsideration of the means for measuring Internet penetration in these regions.

This article 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. The details of the investigation can be found in Johnson, Pejovic, Belding, and van Stam (2011, 2012) and in Johnson, Belding, and van Stam (2012). In a separate effort, we participated in deploying a wireless network in rural Peebles Valley, South Africa (Johnson, 2007). We combine the results of network monitoring, interview analysis, and anecdotal evidence obtained in these three locations to describe Internet usage in rural Africa.

The barriers that stand in the way of more ubiquitous and efficient Internet usage come from a complex interplay of both technical and social factors. For example, the location of access can be roughly categorized to access at home, at work, or in a public, commercial facility. We find that at-home connectivity generally involves a more leisurely type of Internet use with a strong emphasis on social networking. We also find that temporal correlation of usage within the village leads to congested links, which forces some to adapt their daily schedule to Internet availability. Likewise, slow upload speeds prevent substantial amounts of content generation by the locals, which, in turn, raises fears that the Internet may destroy local culture and customs, and minimizes the utility of the Internet for disenfranchised communities.

These findings suggest that rural area Internet penetration should not be evaluated through a simple binary of haves and have-nots. Rather, we propose the assessment of per-user connection capacity, location of access, perceived utility of the Internet, and cost-benefit analysis of the access to obtain a full picture of broadband adoption.

Background

Rural Area Community Networks

Wireless networks based on WiFi technology have emerged as a viable solution for connecting previously disconnected communities in remote regions. Unlike alternatives such as fiber 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 nongovernmental organizations (Bernardi, Buneman, & Marina, 2008; Brewer et al., 2005; Guo et al., 2007; Matthee et al., 2007; Sen, Kole, & Raman, 2006). 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 centers, 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 analyzed in Macha, Dwesa, and Peebles Valley are constructed in this way.

Figure 1. Map of southern Africa with locations of networks analyzed.
Highlighted are the locations of Dwesa and Peebles Valley, South Africa, and Macha, Zambia.

Wireless Network in Macha, Zambia

Macha, Zambia, highlighted in Figure 1, is a typical poor 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 and are likely separated from other clusters by 1 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.

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. Most users access the Internet at work and through community terminals, although a few people do have WiFi connectivity in their homes. The community is connected to the Internet through a VSAT satellite connection. The satellite provides 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 is US$1,200 per month and is covered through government subsidies as well as through Internet vouchers sold to visitors and locals.

Wireless Network in Dwesa, South Africa

The Dwesa region, also highlighted in Figure 1, 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 that are several kilometers in range (Mandioma, 2007). The license for WiMax operation was provided through a local telecom, which 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 interschool communication. Besides being available to students, the connectivity is offered to members of the community after school hours, when computer literacy training courses are offered to local residents.

Wireless Network in Peebles Valley, South Africa

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

The Peebles Valley mesh network, 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.

The Internet users came from a wide range of backgrounds. Regular users included a male middle-aged teacher from a local rural high school who had limited exposure to computers, a young woman who had just graduated from high school who had no exposure to the Internet, a local male middle-aged farmer who was well acquainted with the Internet, and a young boy in junior school who had only used a computer for gaming and had never used the Internet.

The network ceased to exist in 2008, and because of legal obstacles, a local entrepreneur who was interested in running the network could not lawfully inherit it from the outside organization that deployed it.

Methodology

Wireless Network Monitoring in Macha

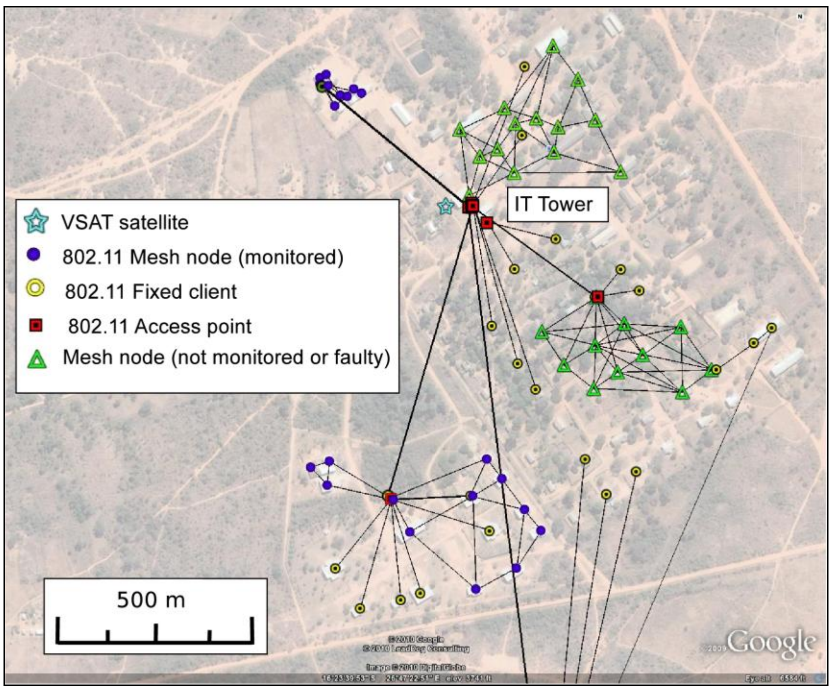

Figure 2. View of the Macha network.
Network traffic is monitored at the satellite link gateway.

Through our partnership with the Macha Works organization, we installed a network traffic monitoring system on the village network gateway as shown in Figure 2. We capture packet headers of all network packets that traverse the satellite link. This allows us to inspect traffic behavior without collecting the actual packet payload. The network gateway also features a Web proxy server used for traffic caching. To analyze the HTTP traffic, the proxy access logs were also collected. All IP addresses were anonymized to protect the privacy of the users. Initially, we captured 14 days of traffic from midnight, Sunday January 31, 2010, to midnight, Sunday February 14, 2010. In a follow-up, we collected 2 months of network traffic in February to April 2011. Approximately 450 GB of packets were captured. In 2010, captured traffic was compressed and sent to our outside location during off-peak hours for off-line analysis. In 2011, we moved the captured traffic via an external hard drive to avoid affecting the trace itself. 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 in the rest of the trace.

In-person Interviews in Macha and Dwesa

The interviews were conducted in July and August 2010, on-site in Macha and Dwesa, privately between one interviewee and one interviewer. Interview participation was on a voluntary basis, and no material awards were associated with participation. The goals of the interview and possible influence on the future quality of network service were explained to every person so they could understand the benefit of participation. In a close-knit African community, residents can be reluctant to openly talk about their habits with a complete stranger. To facilitate the interview process, we used our existing ties with local persons of authority for introduction to potential interviewees. Naturally, this method does not result in a completely random sample; however, every effort was made to ensure that different age, gender, and income groups were represented. 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.

We opted for a directive, structured questionnaire in the first phase of the conversation, because we wanted 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.

Online Survey in Macha

In June and July 2011, we administered an online survey of Internet users in Macha. The survey broadly focuses on usage of Web 2.0 applications and services. The questionnaire consisted of 89 questions and was implemented using the SurveyMonkey tool. Invitations for participation were sent via e-mail and Facebook links, and participation was voluntary. A total of 66 users responded; 41 completed all the questions. We restrict analysis to the latter group. There is significant gender, age, and Internet skills diversity within the sample.

Interviews in Peebles Valley

A member of our team who participated in network building sporadically administered unstructured interviews among network users. While lacking any quantifiable data, this anecdotal information provides deep qualitative insight into Internet usage in Peebles Valley, as the data were acquired over a long period of time by a person who earned the trust of the local community.

Reflection with Longitudinal Research in Macha

Another member of our team provided 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.

Internet Usage in Rural Africa

Traffic analysis helps us profile network usage, measure application popularity, and diagnose problems in the network. Whenever possible we juxtapose findings from the traffic analysis with the interview results.

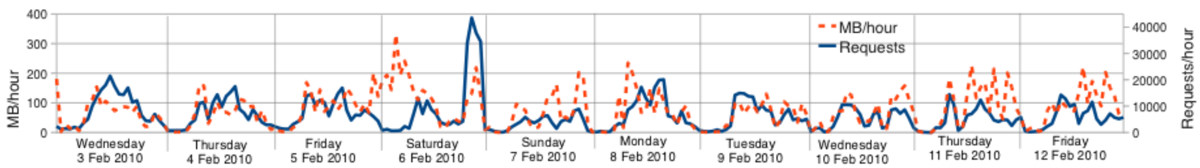

Figure 3. Internet usage in Macha over 10 days.

There is a prominent daily pattern. A low number of requests and high traffic load between Friday and Saturday indicate large file downloads. On Saturday evening, a satellite malfunction caused an unusually high amount of traffic.

Figure 3 displays the traffic load (in MB per hour) and the number of Web requests over 10 days. What emerges is a clear diurnal usage pattern, with the exception of Friday night and Saturday evening. The lack of usage during off-peak hours is partly due to the daily routine in Macha and partly due to the inaccessibility of public Internet terminals during this time. Our interview data show that users who have access at home are more likely to go online after hours than those whose only means of access is a public terminal or a work place (χ^2 (1; *N* = 28) = 5.2; *p* = .041).

Next, we concentrate on traffic types observed in the Macha trace. We observe a high prevalence of Web (HTTP) traffic (68% of the total traffic) and a lack of peer-to-peer sharing traffic, which is commonly observed in traces from developed countries (Schulze & Mochalski, 2009). Instead, our interviews reveal that most of the Internet users in Macha and Dwesa (78% of the correspondents) use USB keys to exchange large files. This method avoids the slow and congested satellite connection but also halts the distribution of content and has detrimental consequences on network security because it facilitates virus spreading.

Analyzing the HTTP traffic in the trace, we find that online social networks (OSNs) are the most popular application, comprising 20% of the Web downloads in the 2010 trace (Facebook only) and 27% of the Web downloads in the 2011 trace (Facebook and Twitter). They are even more prominent in relation to Web uploads—46% of uploads go to Facebook. The remaining uploads in the trace come from Web mail (Gmail, Yahoo mail, for example), file sharing, and so on. More information on the trace composition can be found in Johnson et al. (2011, 2012). Facebook traffic is attractive for further analysis for two reasons. First, it is used for a wide variety of applications, from instant messaging to business advertising. Second, unlike, for example, Web mail traffic, Facebook traffic is not encrypted and reveals a lot about communication patterns among its users. Therefore, we analyze Facebook interactions in detail later in this section.

Online Social Networks

Online social networks are popular in both the developing and the developed world. In 2010 Facebook surpassed even Google search and is currently the most popular online application. The interviews we conducted in Macha and Dwesa show that rural Africans use OSNs in ways typically observed in the developed world: users post photos and wall comments, send messages, and use Facebook chat. In addition, 70% of our interviewees have online-only friends, a window into distant cultures, and some of the interviewees use Facebook for business advertising.

The prominence of OSNs prompted us to investigate the traffic trace for Facebook messages exchanged in the village. Facebook includes a user ID within a message, and by connecting those IDs with the IP addresses that are local to Macha,[3] we were able to construct a graph of Facebook interactions as shown in Figure 4.

[3] All Internet users in Macha use a single satellite link, which is given a unique IP address.

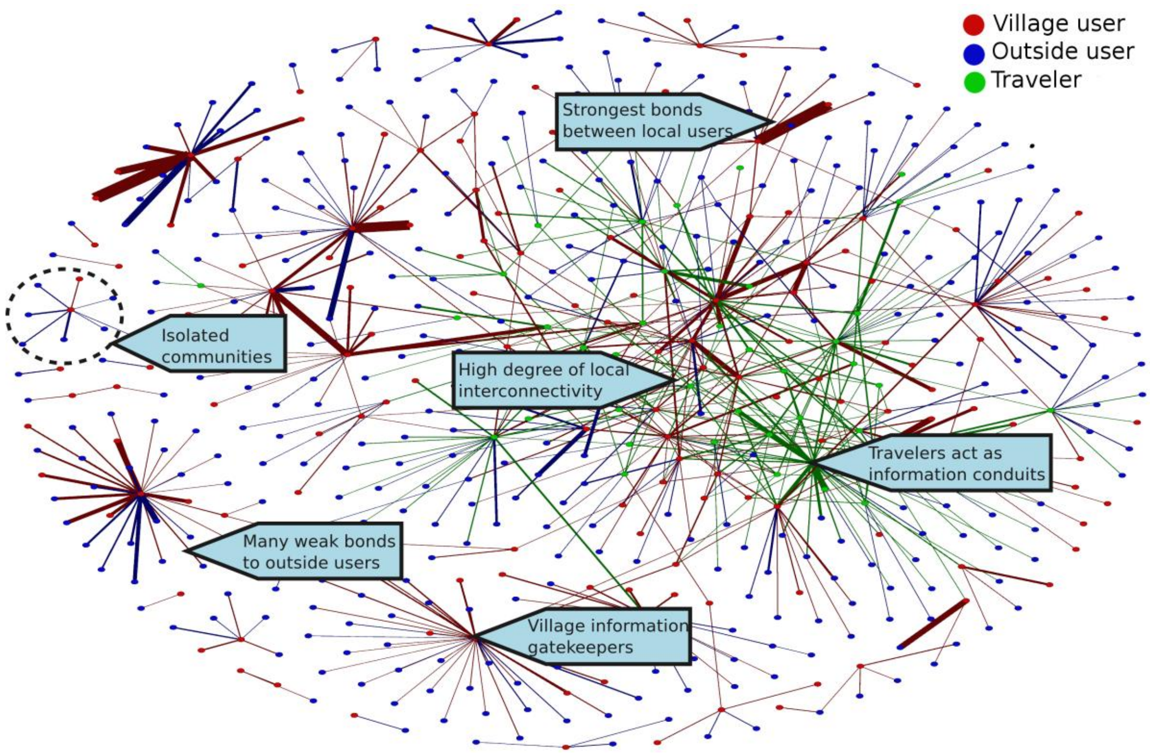

***Figure 42. Facebook interactions in Macha*.**
Users who travel outside the village are well connected and can be seen as information conduits. Interaction among local users dominates the system.

Users are classified as (a) local users, if their ID was observed only in messages with an IP address local to Macha; (b) outside users, if their ID was observed only in messages with an IP that is not within Macha; or (c) travelers, if their ID was observed both within and outside Macha. The interaction graph reveals complex patterns such as isolated cliques and highly connected users who serve as gateways among local and outside users.

We examine the locality of online interactions and find that a great deal of Facebook interactions are local, within the village. Of all the messages exchanged, 54% are among local users, although only 35% of the users in the trace are local. Similarly, interest in photos posted on Facebook reveals that pictures by local users receive four times more local views than those posted by outside users. This finding corresponds with our interview results, where 77% of correspondents reported using OSNs for local communication. The locality of Facebook interaction has been reported in the literature, albeit on a larger geographical scale (Wittie, Pejovic, Deek, Almeroth, & Zhao, 2010).

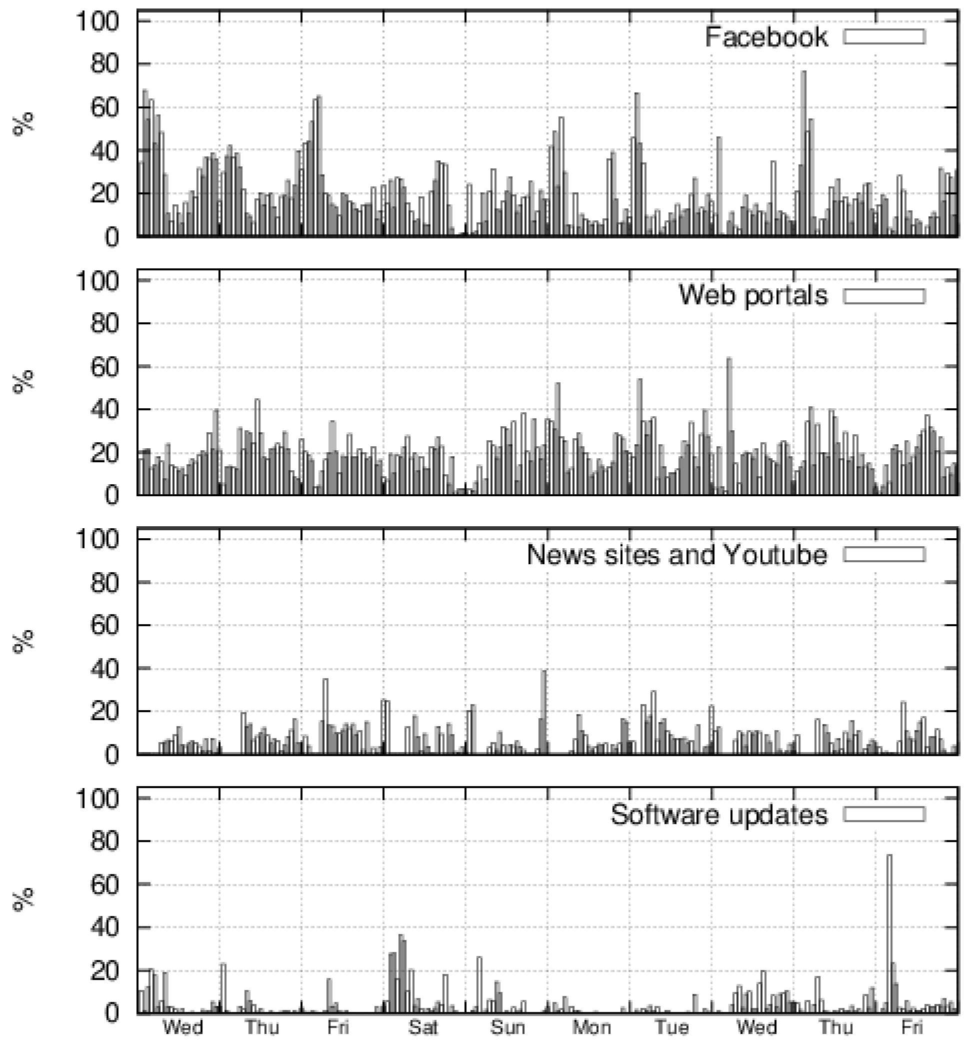

Figure 5. Categorized Web traffic in Macha.

Facebook use peaks during the night.

The popularity of Facebook within the trace shows time-varying behavior. Figure 5 shows the percentages of the total Web requests in Macha over a 10-day period. We classify the traffic as either Facebook, Web portals (Yahoo and Google), News sites (local Zambian newspapers), or Software updates (Windows, Ubuntu, and antivirus software updates). Facebook usage peaks during the night. Earlier we discovered that only those users who have Internet access at home go online during the night. We conclude that *restricting access to public terminals and workplaces severely hampers the type of applications used online.*

Content Generation

For remote rural communities, digital content generation and sharing can be an important avenue for cultural exchange and preservation. Additionally, locally generated digital content can spark development activity from within the society (Srinivasan, 2012).

To investigate content generation in Macha, we decouple upload and download traffic in the trace. The satellite that provides Internet connectivity to Macha is highly asymmetrical; the committed upload speed is half the download speed. This is one of the reasons for substantially more download traffic. The discrepancy between upload and download traffic also can be observed if we compare sizes of individual flows belonging to each of the groups (see Figure 6).

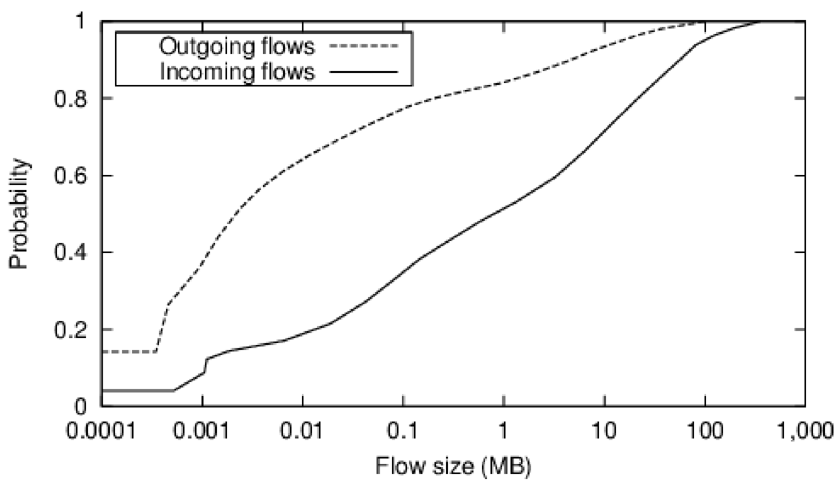

Figure 6. Size distribution of incoming and outgoing flows.
Outgoing flows correspond to file uploads, while incoming flows correspond to file downloads. The above distribution shows that users in Macha upload small files and download large files.

While 80% of the download flows are larger than 100 KB (0.1 MB), only 22% of the upload flows are larger than 100 KB, revealing a lack of large uploads. Many potential technical and nontechnical reasons exist for the lack of content generation, such as a lack of skills and incentives and oversaturation by foreign content (Ballantyne, 2002). We analyze the attempted uploads in the Macha trace, average them over the 2-month trace, and in Figure 7 plot both completed and failed uploads. With more than two-thirds of the attempted uploads failing, residents of Macha are certainly discouraged from digital content production and sharing.

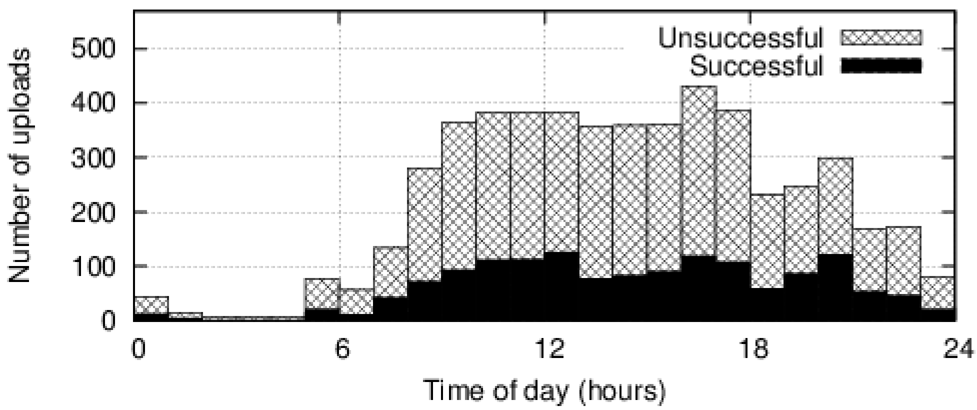

***Figure 7. Uploads in the Macha network**.*
Most attempted uploads fail, especially during peak usage times.

The online interviews we conducted in Macha show that 30% of the users generate some form of online content. Interestingly, this group is not determined by age, gender, or experience in computer usage. Online content generation in conservative societies can be a useful tool in overcoming existing social barriers that prevent certain groups from having their voices heard (Hilbert, 2011). While the simple fun of photo and video sharing represents one of the main drivers of content generation, most of the people we interviewed also report that they share online content because they like helping other people.

However, from the trace we see that content generation and sharing are not as prevalent. We found a significant association of content generation with user-reported accessibility of Internet access (Kendall's tau-b (N = 41) = 249; p = .071). This is well aligned with a previous study (Davis, 1989) that identifies perceived ease of use as one of the key factors for the acceptance of information technologies. Limited bandwidth is the main hurdle for full-fledged Internet access in rural Zambia, and 41.5% of the interviewees mentioned it as one of the reasons that 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 United States during the transition period from dial-up to broadband connections (Horrigan, 2006). Thus, to stimulate content generation and sharing in rural areas, it is crucial to provide support for bandwidth-hungry applications.

Obstacles to Efficient Usage

The obstacles to efficient Internet usage need to be understood and adopted in the design of practical solutions for Internet connectivity in rural Africa. In the course of our studies in Africa, we have identified a number of such obstacles that are caused by factors ranging from personal to governmental.

Restrictions on the Locality of Usage

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 only through at-home access can users 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). First, rural areas often lack reliable grid power. Even if the grid power is available, it is usually restricted to schools, hospitals, and community centers. Renewable energy sources such as wind and solar energy are an attractive alternative for supplying power to rural area network deployments. However, they require substantial initial investment and planning (Bernardi et al., 2008). Long distances between households are another important obstacle for rural area connectivity. Cheap wireless technologies such as WiFi are of short range—a few hundred meters or less. Cell phone base stations, especially those that support 3G and 4G LTE communication can provide high bandwidth connectivity to wide areas. However, they are expensive and not economically viable for areas with low, seasonal-income populations.

Lack of Local Content and the Language Barrier

Our traffic trace analysis shows that digital content consumption dominates over content generation. Local content generation, on the other hand, is important because it can play a key role in preserving the cultural heritage of remote communities (van Hoorik & Mweetwa, 2008). Moreover, with exclusively foreign content, the Internet can be perceived as a threat among indigenous communities:

> Internet in other ways will either build or destroy our culture because of the powerful influence it has in people's lives. Truly speaking, most of things exposed on the Internet are from the western world and very little is from Africa. I do not want to lose my culture. (van Hoorik & Mweetwa, 2008)

Unfamiliarity with New Concepts

The African villages we visited feature Internet connections, although they 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 behavioral 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 one's own device often causes the disconnection of other users in the network.

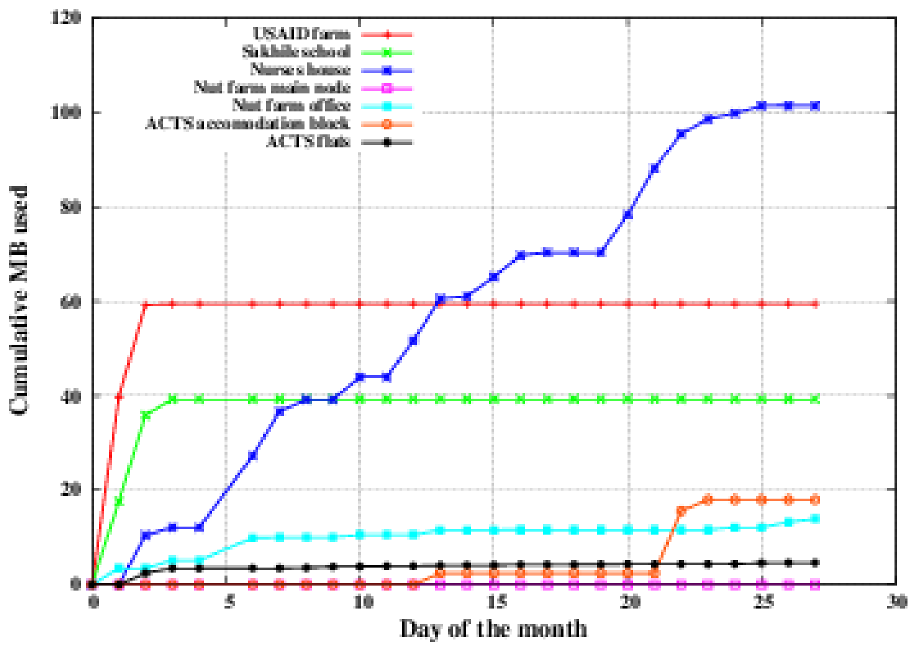

Figure 8. Bandwidth consumption in Peebles Valley for different users.
Unfamiliar with the concept of bandwidth, users often reach their allocated limit early in the month.

Similarly, we observed misinterpretation of the concept of network bandwidth. Figure 8 shows bandwidth usage in Peebles Valley over one month. The users have well-specified preallocated 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. 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. We analyzed the traffic trace from Macha and found a significant presence of botnet traffic. Bots are infected machines that are controlled remotely and carry out malicious online tasks, such as sending spam e-mails, for their botmaster. Bots infected nine out of 201 unique machines we observed in Macha, with about 100 more exhibiting suspicious behavior that might be related to botnets. It is important to note that 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.

Shortage of Trained Personnel

The mesh technology was designed to ensure that very little network configuration was 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 multiweek 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.

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. Even when the local population is involved, it is hard to attain local IT experts in rural areas, because those experts often leave for better jobs in urban areas (Flickenger et al., 2006). We observed this in Peebles Valley, where a local young high school graduate was trained to carry out wireless network installations. However, at the end of the project, he obtained a job in 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.

The impact of a local champion can be significant in motivating other people from the community to adopt technology. For example, in Macha, men are more likely than women to perform computer maintenance themselves (χ^2 (1; 22) = 6.14; p = .023). They are also more likely to do virus scans (χ^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.”

High Cost of Internet Access

High cost is one of the major barriers to further Internet penetration. For example, satellite access in Macha costs US$1,200 per month, while the average monthly income is about US$30. 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). Efforts to bridge the digital divide include significant investments from the developed world, and a sustainable solution seems to be out of reach with the current pricing schemes. In addition, the most common sources of income in rural areas (farming and fishing) generate only seasonal income. While mobile phone providers have addressed the issue of cost through their prepaid billing schemes, which allow underprivileged communities to tailor cell phone usage to their own ways of life (Horst & Miller, 2006), satellite Internet providers demand regular monthly payments (Flickenger et al., 2006).

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 fiber 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 (Wittie et al., 2010). Rural areas of the developing world are located far from data centers 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.

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 to 2010, the size of the average Web page has grown 36 times (Charzinski, 2010); the Internet 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 (Chen, 2011).

Lack of Supporting Infrastructure and Regulatory Problems

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 and Dwesa, 28% of the people we interviewed tried some form of online trading, often with unsatisfactory results. The lack of nationwide Internet adoption plans severely hampers many of its useful aspects. For example, e-government services are not available in Zambia.

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 (Best, 2006). Our efforts in Peebles Valley 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 unfavorable conditions for local network ownership.

Somewhat related to the regulatory issues is the difference that Internet access can make with respect to the existing social hierarchy. For example, Kumar and Best (2006) report that the main reason for the failure of a project that delivered e-government services to rural India stems from the resistance of low-level government clerks who perceived the project as a threat to their means of receiving bribes. From our own experience, providing Internet access to the media center teacher and students who use the school computer lab in Peebles Valley was not acceptable unless the school principal also received a computer with Internet in his or her office.

Measuring Broadband Adoption in Rural Africa

Current approaches for broadband measurement are quantitative, concentrated on estimating the number of users who have physical access to the Internet or the span of areas that are served by Internet service providers. The latter is often approximated through the number of providers serving a specific ZIP code (U.S. Government Accountability Office, 2006). Besides being quite coarse, these metrics do not capture the vast difference in the quality of access among users. While we do not argue for a different way of measuring the number of users, we claim that the quality of the Internet use experience among people in rural or remote communities requires a more elaborate metric. We argue that a comprehensive metric for measuring broadband adoption that includes all aspects of Internet access and usage requires a multidimensional approach. In Table 1 we summarize the dimensions that are, based on our analysis of Internet usage in rural Africa, the most important for efficient broadband access.

Table 1. Summary of Proposed Metric Dimensions, Their Units, and Temporal Interpretation.

Metric dimension	Units	Time variant interpretation
Per-user bandwidth	Bits-per-second	Yes
Online affordances	Complex: a function of locally relevant content, language, and supporting infrastructure	Yes
Location of access	Home/work/public terminal	No
Cost-benefit	Complex: a function of the percent of per-person gross national income (GNI) and the economic impact of connectivity	No

Metric Dimensions for Measuring Efficient Broadband Adoption

Internet connectivity is more than a question of physical access to the network. It is a matter of what the access can afford to the user that determines the utility of the Internet. Obstacles to efficient Internet utilization in rural Africa influence the way broadband adoption should be measured.

First, there is a clear impact of *per-user Internet bandwidth* that limits the type of applications that can be used. Rural areas are often deprived of access to high-speed fiber optic Internet backbone cables. Access is usually brought by satellite connectivity or other long-distance wireless links, which results in at least an order of magnitude slower connection than at-home broadband connectivity in the developed world. To offset the high cost of satellite access, the connection is shared among multiple users. In the villages we surveyed, hundreds of users share the same connection gateway. This results in an extremely low per-user bandwidth. ITU (2011) statistics, for example, show that the international Internet bandwidth in Africa reaches 937 bits per second per user, while in Europe it is 78,678 bits per second per user. Our research shows that bandwidth impacts the way files are shared and is one of the main obstacles to local content generation in rural Africa. Moreover, more Internet bandwidth is needed for an efficient online experience. The average Web page grew 36 times in just 15 years. Bandwidth-hungry applications such as video streaming services and online social networks are becoming more popular. Thus, per-user bandwidth must be considered with respect to the requirements of popular online applications at the time of the measurement.

The online realm is not of equal value to everyone. While a user in the urban developed world can engage in online commerce, obtain directions via online map services, and schedule a doctor's appointment online, users in the rural developing world generally do not enjoy these benefits. Measuring the opportunities that one can obtain online requires a complex metric that takes into account the availability of online content that is relevant to the user. The relevance is connected to the language one speaks, the culture one is a part of, and the socioeconomic and geographic context in which one resides. Measuring *online affordances* also must take into account the infrastructure that is not necessarily connected with Internet access, such as credit card banking and transport of goods, but still plays a role in the utility of some online services. Finally, online affordances should include the social aspect of using the Internet, as elaborated by the social affordances concept introduced by Wellman et al. (2006).

Location of access critically impacts Internet usage. In our research we find that at-home connectivity leads to more leisurely Internet use with an emphasis on online social networking and content generation. Wyche, Smyth, Chetty, Aoki, and Grinter (2010) show that public access leads to the so-called deliberate interaction model, where online activities must be preplanned before the access actually happens. On the other hand, Best, Kollayani, and Garg (2012) find that public access in African cybercafés has an educational note, and users engage in collaborative Web access to enhance their skills. Similarly, Rhinesmith (2012) finds that Internet users in public libraries in Philadelphia, Pennsylvania, supplement, rather than replace, at-home access with public access. The latter is seen as a more social experience.

The cost of broadband puts Internet access into perspective with other basic necessities. Broadband connectivity requires up to a few hundred times proportionally higher investment from a user who lives in the developing world than from a user who lives in the developed world. Consequently, the benefits of the access have to be high. To obtain a clear picture of the effectiveness of broadband access, we propose inclusion of a *cost-benefit* analysis in the overall metric. Measuring the affordability of broadband can be done through identification of a percentage of per capita gross national income, as in the ITU (2011) study, for example. The benefit that Internet access brings is harder to measure, and we suggest investigation of the existing economic practices and online behavior to describe the impact of the Internet on one's living standard.

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 quality and the impact of such connectivity is often lacking. This evaluation is important because it can point out, and differentiate between, good and bad practices in designing and deploying rural area networks. Because of a complex interplay of numerous technical and social issues that affect how people access the Internet and what they do online, it is inappropriate to measure broadband connectivity with only a simple metric, such as the number of users who have access to the Internet.

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 restrictions on the locality of access, a lack of locally relevant content, unfamiliarity with new concepts, shortage of trained personnel, high cost of Internet access, and limited connection capacity with respect to the Internet structure and content.

Our Internet usage analysis points out that a metric employed for measuring broadband in rural Africa has to describe the state of connectivity with respect to all the obstacles listed above. Recent research by Gomez and Pather (2012) argues that the benefits of ICTs in developing regions often come in a sophisticated form that cannot be explained through easily measurable parameters, and the authors call for a methodological shift in ICT evaluation. We propose a multidimensional metric that measures per-user Internet access bandwidth, describes personal affordances from going online, categorizes access based on the locality of connection, and examines the weight of benefits provided by going online versus the subjectively measured cost of access.

The measurement dimensions that we propose require a more in-depth approach than a simple Internet user count. Per-user bandwidth, for example, depends on the type of connectivity, the technical details of access distribution over a community that uses it, and the temporal correlation of access within the community. Similarly, measuring online affordances calls for a thorough anthropological study of the availability of digital content that the local population finds relevant and the evaluation of financial and transport infrastructure that renders the Internet more or less useful. Nevertheless, we argue that only through multifaceted metrics can we truly measure the impact of broadband connectivity on rural regions and the usage characteristics of that connectivity.

References

Abraham, R. (2006, May). Mobile phones and economic development: Evidence from the fishing industry in India. *International Conference on Information and Communication Technologies and Development (ICTD 2006)*. Berkeley, CA.

Ballantyne, P. (2002, May). *Collecting and propagating local development content* (Report of a project carried out by IICD in association with the Tanzania Commission for Science and Technology and funded by the UK Department for International Development).

Bernardi, G., Buneman, P., & Marina, M. K. (2008, September). Tegola tiered mesh network testbed in rural Scotland. *ACM Wireless Networks and Systems for Developing Regions Workshop (WiNS-DR)*. San Francisco, CA.

Best, M. (2006, September). A global survey of spectrum license exemptions. *Research Conference on Communication, Information & Internet Policy (TPRC)*. Washington, DC.

Best, M., Kollyani, B., & Garg, S. (2012, March). Sharing in public: Working with others in Ghanian cybercafes. *International Conference on Information and Communication Technologies and Development (ICTD 2012)*. Atlanta, GA.

Brewer, E., Demmer, M., Ho, M., Honicky, R. J., Pal, J., Plauche, M., & Surana, S. (2005). The challenges of technology research for developing regions. *IEEE Computer*, *38*, 15–23.

Charzinski, J. (2010). Traffic properties, client side cachability and CDN usage of popular websites. *Measuring, Modelling, and Evaluation of Computing Systems and Dependability and Fault Tolerance*, *5987*, 136–150.

Chen, J. (2011). *Re-architecting Web and mobile information access for emerging regions* (PhD thesis). Department of Mathematics, New York University, New York, NY.

Davis, F. D. (1989). Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, *13*, 319–340.

Flickenger, R., Aichele, C., Fonda, C., Forster, J., Howard, I., Krag, T., & Zennaro, M. (2006). *Wireless networking in the developing world.* Limehouse Book Sprint Team.

Gomez, R., & Pather, S. (2012). ICT evaluation: Are we asking the right questions? *Electronic Journal of Information Systems in Developing Countries, 50*.

Guo, S., Falaki, M. H., Oliver, E. A., Rahman, S. U., Seth, A., Zaharia, M. A., Ismail, U., & Keshav, S. (2007, December). Design and implementation of the KioskNet system. *International Conference on Information and Communication Technologies and Development (ICTD'07)*. Bangalore, India.

Hilbert, M. (2011). 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*, 479–489.

Horrigan, J. (2006, May). *Home broadband adoption 2006* (Technical report). Pew Internet & American Life Project. Washington, DC: Pew Research Center.

Horst, H., & Miller, D. (2006). *The cell phone: An anthropology of communication*. Oxford, UK: Berg Publishers.

International Telecommunication Union. (2011). *Measuring the Information Society.* Retrieved from www.itu.int/net/pressoffice/backgrounders/general/pdf/5.pdf

Johnson, D. L. (2007. December). Evaluation of a single radio rural mesh network in South Africa. *International Conference on Information and Communication Technologies and Development (ICTD2007)*. Bangalore, India.

Johnson, D. L., Belding E. M., & van Stam, G. (2012, March). Network traffic locality in a rural African village. *International Conference on Information and Communication Technologies and Development (ICTD 2012)*. Atlanta, GA.

Johnson, D. L., Pejovic, V., Belding E. M., & van Stam, G. (2011, March). Traffic characterization and Internet usage in rural Africa. *International World Wide Web Conference (WWW '11)*. Hyderabad, India.

Johnson, D. L., Pejovic, V., Belding E. M., & van Stam, G. (2012, March). VillageShare: Facilitating content generation and sharing in rural networks. *ACM DEV '12*. Atlanta, GA.

Kumar, R., & Best, M. L. (2006). Impact and sustainability of e-government services in developing countries: Lessons learned from Tamil Nadu, India. *The Information Society, 22*, 1–12.

Mandioma, M. (2007). *Rural Internet Connectivity: A Deployment in Dwesa-Cwebe, Eastern Cape, South Africa* (MS thesis). Computer Science Department, University of Fort Hare, South Africa.

Matthee, K. W., Mweemba, G., Pais, A. V., van Stam, G., & Rijken, M. (2007, December). Bringing Internet connectivity to rural Zambia using a collaborative approach. *International Conference on Information and Communication Technologies and Development (ICTD'07)*. Bangalore, India.

Rhinesmith, C. (2012, April). Free library hot spots: Redefining broadband adoption in Philadelphia's low-income communities. *Defining and Measuring Meaningful Broadband Adoption Workshop.* Washington, DC.

Schulze, H., & Mochalski, K. (2009). *Ipoque Internet study 2008/2009*. Leipzig, Germany, ipoque.

Sen, S., Kole, S., & Raman, B. (2006, May). Rural telephony: A socioeconomic case study. *International Conference on Information and Communication Technologies and Development (ICTD 2006)*. Berkeley, CA.

Srinivasan, R. (2012). Rethinking digital cultures and divides: The case for reflective media. *Information Society, 28*, 24–36.

Surana, S., Patra, R., Nedevschi, S., Ramos, M., Subramanian, L., Ben-David, Y. et al. (2008, April). Beyond pilots: Keeping rural wireless networks alive. *USENIX Symposium on Networked Systems Design and Implementation (NSDI '08).* San Francisco, CA.

U.S. Government Accountability Office. (2006). *Broadband deployment is extensive throughout the United States, but it is difficult to assess the extent of deployment gaps in rural areas* (Report to Congressional Committees). Retrieved from http://www.gao.gov/new.items/d06426.pdf

van Hoorik, P., & Mweetwa, F. (2008, May). Use of Internet in rural areas of Zambia. *IST–Africa '08*. Windhoek, Namibia.

Waverman, L., Meschi, M., & Fuss, M. (2005). The impact of telecoms on economic growth in developing countries. *Vodafone Policy Paper Series*, *2*, 10–23.

Wellman, B., Quan-Haase, A., Boase, J., Chen, W., Hampton, K., Diaz, I., & Miyata, K. (2006). The social affordances of the Internet for networked individualism, *Journal of Computer-Mediated Communication, 8*(3). Retrieved from http://jcmc.indiana.edu/vol8/issue3/wellman.html

Wittie, M., Pejovic, V., Deek, L., Almeroth, K., & Zhao, B. (2010, December). Exploiting locality of interest in online social networks. *ACM CoNEXT.* Philadelphia, PA.

World Bank. (2011). *World development indicators.* Retrieved from http://databank.worldbank.org/databank/download/WDIandGDF_excel.zip

Wyche, S., Smyth, T., Chetty, M., Aoki, P., & Grinter, R. (2010, April). Deliberate interactions: Characterizing technology use in Nairobi, Kenya. *ACM Conference on Human Factors in Computing Systems (CHI '10).* Atlanta, GA.