Reconciling Community Resource Requirements in U-Nets

Sara Bury
sara.bury@comp.lancs.ac.uk

Johnathan Ishmael
ishmael@comp.lancs.ac.uk

Nicholas J. P. Race
race@comp.lancs.ac.uk

Mark Rouncefield
m.rouncefield@lancaster.ac.uk

Paul Smith
p.smith@comp.lancs.ac.uk

Computing Department
Lancaster University
Lancaster, LA1 4WA, United Kingdom

ABSTRACT

A challenge for future user-provided networks will be reconciling potentially conflicting demands for a finite resource, such as network bandwidth. In this position paper, we discuss how this problem is tackled in an operational community-driven wireless mesh network. A key outcome of this discussion is that, although the approach that has evolved is not ideal, it allows the reconciliation of conflicting demands for use of the shared network to reflect communal concerns; a property we believe is essential to its success, and that of user-provided networks. We argue that user-driven distributed arbitration of requests for resource is necessary in a user-provided network, and via a simple abstraction we discuss design options to enable this. The consequences of incorrect design decisions could negatively impact a network and its community of users. To help us make appropriate design decisions, we could look to scientific methods, such as game theory. However, we find them unable to model the intricacies of communal life, leading us to suggest it is necessary to stop considering users as anonymous rational persons and to start factoring in their personalities and beliefs.

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network Operations; J.4 [**Social and Behavioural Science**]: Sociology; H.5 [**Information Interfaces and Presentation**]: General

General Terms

Design, Human Factors, Management, Measurement

1. INTRODUCTION

Most commonly this is carried out using wireless technologies to route traffic via infrastructure hosted within people's homes and is typically connected to the Internet by some number of uplinks provided by individual members of the community. Such networks are not only applicable in remote areas, or where users have limited of intermittent connectivity, they are often deployed in urban centres where privately owned WLANs overlap strongly undermining the potential of the radio spectrum available. However, in this paper the focus is concentrated on networks created in areas where Internet connectivity has been unavailable, , whi

Often, the creation of such networks has been initiated by a group of people united solely by their common goal of network provision. Such a group likely contains people with a wide variety of technical knowledge and experience, and differing expectations of the amount of time and effort the project will require. User-provided networking isn't without its challenges [12], while a significant number of these are technical in nature; the noteworthy ones are those which are cross-discipline and socio-technical due to the effects of users being in or on the administrative loop. Such factors will have an impact on both the planning and deployment of the infrastructure, but also the day-to-day operation and management of the network when it is in place.

The desire to improve user-provided networks often results in a research focus centered around alterations to the underlying technologies, such as improved resource allocation at the MAC layer, or dynamic routing protocols for mobile hosts; we believe that when considering such networks this does not present a broad enough view. Based on a five year involvement with the operation of a community-based user-centric network, we present a number of further topics for investigation stemming from the realisation that the users themselves, and more importantly their actions, are a key component of any user-centric network and their involvement as consumers and governors should not be ignored.

In this position paper, we examine the challenge of resource sharing for user-provided networks, using knowledge obtained through our experiences with the members of a user-centric community network. Specifically, the paper provides the following contributions: by highlighting our own experience with the Wray Community Wireless Mesh Network (WMN), we describe the potential pitfalls of community network management, focusing on resource management, and discuss the shortcomings this has for User-provided Networks (U-Nets). Subsequently, we begin to offer a potential resolution to the issue of communal re-

U-NET'09, December 1, 2009, Rome, Italy.

source management within U-Nets, through a way of abstracting the requirements of users and providing a means of information-driven community arbitration. We then discuss the potential challenges in deploying such a system, describing how releasing the wrong information may lead to detrimental effects to both the network and the community in which its based. Finally, we discuss the pitfalls of using game theory as an avenue for design verification.

2. BACKGROUND

Approximately six years ago, Lancaster University was approached by villagers from the community of Wray (a small, relatively remote village in the North West of England). The community felt strongly that the lack of broadband availability in their village (a consequence of their remoteness) was jeopardising local businesses, education and impacting on other aspects of community life. To address this issue, the University helped to establish a WMN within the village – allowing it to investigate the resilience and capabilities of the technology, whilst providing the community with Internet access in their homes and at public locations around the village [8]. The network was offered as a shared community resource; anyone could request access (without charge) and the underlying network capacity was shared equally amongst all of the users, including those that hosted parts of the infrastructure.

Whilst the University had taken a major role in setting up the infrastructure, it was the community themselves who were to maintain and manage the network on a day-to-day basis (rather than any central authority). To facilitate this process, the community established a Wray WMN committee that was open to all villagers, offering training in the use of network administrative software and device configuration. While the Wray committee was ten strong at its creation, numbers soon diminished as people moved away from the village, apathy in dealing with users' problems set in and committee members found they had a general lack of time for management in between their day job and family responsibilities. Over time, the committee was reduced to a couple of enthusiastic members who were ultimately responsible for managing the entire network, and making policy decisions for over 500 users.

2.1 Network Management Tools in Wray

One of the initial challenges for this committee was having sufficient information available about the network to make informed decisions – they required mechanisms they could use to monitor the overall status of the network and track the usage of users in real-time. Whilst the committee had access to information on the state of each mesh node, this information was delayed by up to one hour, which led the committee members to come up with creative ways of diagnosing problems on their own: given a detailed knowledge of the infrastructure (in particular, the location of mesh nodes within the village and the users that would routinely connect to each node), they could use an instant messaging application to see which of the village users were online, thus inferring that parts of the network were operational.

Over time, it became clear that the committee required more detailed information about the usage of the network. As the user-base grew, they found an increasing number of issues associated with users connecting to the network at peak times – particularly those with weaker signal levels. As a result, they requested that the University provide them with traffic data on a per-user basis in order that they could start to introduce methods of controlling the usage of the network. The committee were provided with a breakdown of how much each user's device downloads on an hourly basis and this information was used as the starting point for policy enforcement.

2.2 Policies and Enforcement

As the community network in Wray is a shared resource, it is perhaps inevitable that the committee would start to see issues with individual users consuming significant portions of the available bandwidth, leaving others with little or no access. This became more prevalent as users grew in confidence; they would start out using the network for relatively simple e-mail and web browsing but over time many users discovered new applications (including P2P, Skype and video streaming) that were significantly more network intensive. In an attempt to alleviate this problem, the remaining committee members introduced an acceptable usage policy which stated that users must not transfer more than 50 MB/hour, and any large files (over 100MB) should not be downloaded between the hours of 9am-9pm. These figures were based upon a knowledge of the underlying capacity of the network and the average number of concurrent users.

To assist with the implementation of this policy, network usage statistics are generated for the previous 24 hours and sent to the committee members each morning. This information is then analysed by the committee and used to correlate events which may have occurred the previous day in addition to highlighting possible transgressions by users. In the event that a user appears to have exceeded their download allocation a warning is sent by the committee via email. In the case where a user ignores these warnings and repeatedly exceeds their allocation the user will be disconnected from the network and must approach the committee to explain their actions before access can be restored.

Additionally, the committee will often use contextual information to help inform decisions over warnings and disconnections. As the committee members have a personal knowledge of individuals within the village, they will often use this information to *automatically* excuse a breach of the acceptable usage policy. For example, they may consider a user returning from a long business trip, or users with visitors as mitigating factors in any breach of policy, and will often forgo contacting the individual concerned for an explanation. There are also a smaller number of policies which are enforced by the committee, such as a requirement for anti-virus software on connected devices, and an assurance that computers will be patched and kept up-to-date – these are typically requested at the time of connection, but not enforced once access has been granted.

3. U-NET RESOURCE MANAGEMENT CHALLENGES

By definition, in a user-provided network the resources that make up the infrastructure will be contributed by a number of its users. Depending on the incentive models in place, providers may wish to contribute wholly or partially the total resource they have to the community, both in terms of access (to the Internet) and transit (within the multi-hop network) resource. The resources that are allocated to the community could be provisioned in a number of different ways, arguably increasing in complexity and control:

As a best-effort service.

Our experience in Wray has shown that heavy users of the network (e.g., those using peer-to-peer file sharing services) can starve others of bandwidth, perhaps unfairly [9]. This approach, whilst straightforward to implement, is undesirable.

Allocate a fair share to users.

An alternative approach to a best-effort service would be to apportion bandwidth equally between users. In the current Wray deployment, the throughput of users is statically capped to 2Mb/sec per client. A more dynamic approach could be taken, where bandwidth is allocated equally based upon observed demand.

Based upon traffic classes.

Assigning a fair share of resources to users has merits, but lacks the flexibility to be able to prioritise certain types of traffic (e.g., media streams over email traffic) that could be desirable. This prioritisation of traffic could be achieved by classifying traffic and assigning priority to certain flows in the network. However, this has drawbacks: traffic classification is not straightforward [11], and we have observed usage of the network in Wray that suggests this policy would be too restrictive (see below). The issue of traffic management practices by traditional ISPs (as highlighted in recent press articles) also suggests a need for greater transparency between users and the management of the underlying network.

Reflecting communal requirements.

In Wray, we have observed requests for resource that transcend requirements based upon traffic class, and are more complex than simple fair-share resource allocation. Specifically, certain households have been given dispensation to break usage rules to carry out research during periods when school children's coursework is due. Also, a user of the network has asked for exemption from the rules to carry out critical activities for their work. As these are exceptional cases, the operators in Wray and the wider community have allowed these transgressions to pass, despite the detriment to their service. A simple traffic-based priority or fair-share scheme would not allow these situations to occur.

We believe a key to the success of the Wray deployment is the ability of the small management team to use context regarding communal requirements of the network to flexibly enforce fair usage policies. Without this flexibility, users of the network may have become frustrated with the level of service they received during personally critical periods, and turned to alternative provision. Furthermore, during an interview with the Wray operators where we discussed the potential for automating traffic management (e.g., based upon traffic type), they were vehemently opposed to such alterations – there is a strong desire that it should never become as restrictive as a regular ISP and it is important to them that the users feel in control of their access.

However, in a user-provided network the centralised network management approach used in Wray, which affords the desired flexibility, in all likelihood cannot be used. The introduction of multiple infrastructure (and virtual) providers, who can specify their own usage restrictions based upon an incentive model, adds complexity that could quickly become intractable. A way of managing this complexity is to have network resource allocation requirements be expressed by the users of the network (both providers and consumers) and the network attempt to fulfil these requirements on their behalf (i.e., have the network act as the user-specified policy enforcer). Demands for resource will have an affect on other users, which will need to be reconciled through trade-offs made by them. A fundamental component of such a system is its interface to the user community that shows such information as current network availability, their usage, and allows users to express how they would like to moderate their usage and that of others.

4. COMMUNITY RESOURCE CONTROL: DESIGN CONSIDERATIONS

We have argued that in a user-provided network, provisioning (or management of resources) should reflect communal requirements, and these requirements should be (requested and) arbitrated in a distributed fashion, by the community itself. As a means to simplify the problem and help us reason about how the design of a system to support this should be approached, we can express an individual's request for resource as two parameters: the request itself and a weight (or statement of importance) associated with the request. These two parameters can be submitted to the system for implementation, the outcome of which can be arbitrated by the community in various ways. Using this simple abstraction, a number of design decisions can be considered.

4.1 Resource Request Format

As mentioned above, there are two components which make up a request for network resource – the request itself and a statement of importance. An example of such a request has been seen a number of times within the current Wray deployment. A user requested unrestricted use of the network (i.e., not to be subjected to normal policy constraints), and stated this was critical for their work. As it was deemed acceptable by the community members running the network, the user was permitted to use the network without the normal policy restrictions. The question raised is: how can similar requests, such as the example above, be expressed via an on-line system?

The way in which a user requests resource from the network must be specific enough to allow the system to allocate the correct resources, while simple enough for a novice user to understand. There are several potential options as to how resources can be requested, including:

QoS parameter-based e.g., bandwidth, jitter, delay, availability.

This provides the greatest amount of control over the resource request and is perhaps the easiest to implement, yet is probably too complicated for most users.

Application-based.

A user could specify which applications are to be used, such as Web browsing, watching YouTube videos or using Skype. This would provide a rich parameter set while being simple for a user to understand. Consequently the list of applications would have to be maintained and configured by someone with relative technical knowledge.

Via an appropriate metaphor.

This would relate the resource request with something familiar to novice users, e.g., a water tap, for example – turn it up for more resource and down for less. This reduces the level of detail available to the system to tune the resource request.

Temporally.

In addition to the forms of request presented, a user could also specify their requirements in advance, e.g., "I want to watch television between 21:00 and 21:30 tomorrow."

In addition to the request for a specific resource they may also wish to place a weighting behind it, indicating how much of a necessity the resource is to them. The weighting could take one of the following potential forms:

A numeric value.

This could be to indicate how much they need the resource (e.g., -1 to +1) where negative numbers could indicate they are willing to relinquish some of their resources to other users. Positive numbers would indicate a need, or urgency about the request.

Excuses.

Similar to the Wray example mentioned previously – a user could request the resource accompanied by an excuse or comment, e.g., a piece of course work is due at school, or important videoconference is needed for work

Credits.

A user could "pay" credits to have their resource increased, similarly they could also opt to reduce their resource for some form of payment. This could form part of a wider incentive policy.

4.2 Arbitration of the Implementation

Upon receipt of a resource request, the community should (if necessary) be able to arbitrate its implementation by the network. This can be achieved in a number of different ways, discussed below. For the community to make these decisions, they need to be informed about current resource usage in the network and potentially historical data relating to the resource requestee, the precise form this should take we also discuss.

One initial design decision relates to the system behaviour in the absence of resource requests. This could include simply offering a best effort or fair-share service. This is important, as all users not making specific resource requests would fall into this category. Such users may have the consequence that they have little or no resources available. The remainder of the users' resources will be allocated based upon the arbitration carried out by the community. This can be achieved in a number of ways:

Self policing.

Users, with information about current usage and resource availability, police their own usage. For example, this could be implemented by users changing their resource request or its level of importance. Potentially, this could lead to unfair usage of the network if users try to play the system by continually requesting high resource usage.

Communal policing.

All or a subset of the community police resource requests and usage. For example, this could be implemented via a voting system – resource requests could be granted if a majority agree, or abusers of the network could be reprimanded, similarly.

As mentioned earlier, in order for a user or the community to make informed decisions regarding resource requests, they need information regarding their current network usage (or state) and that of their neighbours and possibly the entire network. This permits the community to determine if, and the consequences of a resource request being honoured. Providing current network usage to a user would allow them to assess their current usage and adjust their weight/request accordingly. Information regarding the current network usage could be presented as:

- A users own usage represented in some way
- The average usage on the network
- The highest (or lowest) usage in the network
- A cumulative total of the above (historical data)

While revealing such information would permit the community to make more informed decisions, it also raises concerns over privacy – one must consider how much data it is appropriate to reveal to the community (e.g., should individual users be named, specific protocols be highlighted or websites accesses listed). While the community may make use any information regarding the network usage, they may also glean information via off-line methods (as is currently the case in the Wray network), e.g., through gossip in the local post office.

One must also be mindful of the network's topology. A significant challenge lies in determining to which network nodes and sets of users a resource request should be propagated. In a single-hop network only one node (and its users) considers and implements a request. For a multi-hop network, a request must be propagated along the path to a required destination (e.g., to the Internet), assuming all nodes make up the path. However, for a multi-hop network with multiple paths, determining the target nodes for a resource request is less straightforward, as only a subset of the network may be affected by a given request. To resolve this issue, it is necessary to determine where resource contention could occur, as a consequence of a request. Determining where contention lies requires information from multiple network layers – while a user may use a single network-layer path, the radio devices on that route may interfere with other paths, making the contention a link-layer issue; both pieces of information are needed to determine this. This adds considerable overhead and potentially unnecessary complexity to the system. An alternative is to simply allow everyone in the community to arbitrate a request, which may have scalability issues. Managing these trade-offs is a matter for further investigation.

In summary, this discussion has highlighted a number of design decisions that relate to the expression and arbitration of resource requirements within a user-provided network. It is not immediately clear which of these designs is appropriate; we could turn to game theory to help us make appropriate design decisions.

4.3 Informing Design with Game Theory

As an interdisciplinary endeavour we have long employed variants of social science methodology in order to gain various kinds of understandings about our users and their community [5]. Not unnaturally we therefore turned to the social sciences in order to glean some predictive insights into the likely consequences of any proposed design interventions. We were interested in reasoning, in informed speculation, about practices and situations of use, following Woolgar's [13] argument, that a substantive part of any systems project involves reasoning about what users might do with the system. Quite simply, and perhaps naively, we wanted to be able to predict user's likely reactions to our design decisions in order to make the 'right' design choices. In particular we had in mind that having given users information about network use, especially about how others were using (or misusing) the network, that this would be a position and a policy it might be difficult to reverse, and whilst obviously not on a parallel with Eve eating from the tree of knowledge or Pandora opening her box, the consequences could be detrimental both to the use of a U-Net and, in the form of bitterness and recrimination, to the community itself.

Game theory was an obvious candidate approach for initial consideration, not least because the existing literature seems to relate to some relevant and similar problems. Mahajan et al. [10] discuss their experiences in using game theory to design systems that would address two, related, but specific issues, how to discourage selfishness and promote cooperative behaviour in wireless networks. We face at least one essentially similar problem in that as a general rule we also would like to encourage cooperative behaviour – and indeed who wouldn't? However, our particular concern is more of a general 'policy' issue and involves the amount and level of information we should reveal to users. How much do users need to know about how the network is being used, how other users are using or abusing the system, and so on.

Before we make a policy decision on these issues we would like to be able to predict likely reactions amongst our user group. Would having more information lead to greater levels of abuse and, potentially, community suspicion and disintegration? Like Mahajan et al. we turned to game theory in the expectation that it might enable us to model the problem we faced and make some reasonably robust predictions about possible outcomes of different policies. Similarly, we have discovered that, whilst the exercise of considering users' approach to the use of a U-Net as a 'game' has proved interesting, no obvious solution has emerged because of complications emerging from the interaction of a range of real world issues. In part, these complications have arisen from the simple fact that our involvement in Wray over a considerable period of time has meant that we doubt some of the more simplistic assumptions – notably assumptions about human behaviour, of selfishness or altruism – behind a number of the game theory models. Undoubtedly, the power of game theory comes from its ability to develop abstract models about behaviour, to enable a clarity of thinking uncluttered by mundane details. Unfortunately, the real world largely consists of exactly such mundane details and as we considered the application of game theory to this particular setting, translating from abstract solutions to implementation in the real world, we encountered a series of 'yes, but..' reactions shaped by what we already know about the idiosyncracies of our particular users.

As Crawford (1997) argues 'behavior' in games is notoriously sensitive to details of the environment, so that strategic models carry a heavy informational burden, which is often compounded in the field by an inability to observe all relevant variables'. Whilst not yet being prepared to admit that 'generally speaking, the situation is hopeless' we remain exercised by this mismatch between theory and 'reality'. We are also and obviously concerned at the impact that providing information about others activity (use of bandwidth) might have on each users own behaviour; that it might impact on 'reality' through influencing expectations about what others are doing or are likely to do [3]. At a perhaps deeper and more general level, the question is whether the users in our community are shaped by rational choice or whether other, perhaps mixed, motivations are operating, in particular motivations consequential on the view of the community as a 'moral economy' (see for example Fafchamps [6], Bowles and Gintis [4]). The fact that our WMN embraces a physical, geographical community, characterised by a range of existing relations, obligations and ties makes understanding the consequences of our actions more important and yet, paradoxically, more difficult to predict, especially using existing versions of game theory. In consequence, our approach to making design decisions about policy is likely to follow the course of Mahajan et al. who, arguing that traditional games theory did not fit well with the requirements of their problems, eventually turned to 'less formal approaches to construct and analyse solutions'.

5. IMPLEMENTATION CONSIDERATIONS

Whilst we have started to elicit design concerns associated with a system offering resource sharing and arbitration within a user-provided network, it is important that these are considered within the constraints of the underlying infrastructure upon which the network is built. As highlighted in Section 2, the Wray committee frequently requested access to fine-grained data to assist them in making informed decisions relating to network management. Our experiences in implementing monitoring systems to afford them this level of detail have highlighted the challenges when working with networks formed entirely from resource constrained devices.

The issues that we frequently encountered with this type of hardware related to their processing and storage limitations. These limitations resulted in us being unable to capture detailed traffic statistics; and having to restrict our capture to one minute averages. Additionally, we found the systems had limited abilities to perform traffic shaping and intrusion detection – leading us to develop new techniques for monitoring on these types of devices [7].

The current Wray deployment (now over six year old) is incapable of supporting the types of resource sharing and arbitration services we have outlined, and so as part of a wider renewal programme to update the infrastructure within Wray we are introducing a next-generation platform. This upgrade is based around an Alix3c hardware solution, utilising an AMD Geode 500Mhz processor and bespoke software platform built upon OpenWRT[1]. WifiDog[2] provides a captive portal system (offering user access control via a web-based login system) alongside a basic monitoring platform to capture bandwidth utilisation on a per-user level.

The upgrade provides a significant increase in processing power and available storage, but we cannot be complacent; the use of multiple radio cards and frequencies will also result in a significant increase in available bandwidth – placing

[1]http://www.openwrt.org
[2]http://dev.wifidog.org/

additional demands on any resource arbitration mechanisms. We still believe that when designing solutions for resource arbitration due consideration must be made to the capabilities of the underlying platform.

6. RELATED WORK

In this paper, we reflect on some of our wrestling with the complexities – both technical and, importantly, social – which face any effort at developing a working model of community management of user-provided networks. These complexities, and the subtle features of communal life that they point to, have generally been elided in other attempts to address network management issues (e.g., Barraca and Aguiar 2008 [2]). While admirable in at least recognising the importance of social factors in network management, what such studies generally offer is an impoverished version of social life, a model that relentlessly avoids all the intricacies, all the nuances of communal life. Of course, one can easily point to important features of a community or a network, to things like the importance of 'membership' or 'trust' and so on (as indeed Barraca and Aguiar 2008 do), but what precisely these things amount to, what they might mean in this particular community and how this might impact on design decisions cannot be assumed prior to some form of investigation. So, for example, both 'membership', 'tolerance' and 'trust' are clearly contextually determined. Knowing our own community so well enables us perhaps to anticipate some of the difficulties that might follow any attempt to embed such properties in a management system.

7. CONCLUDING REMARKS

This position paper has highlighted a number of issues in reconciling community resource requirements within user-provided networks. We argue that an understanding of users (and their behaviour) is vital when designing technical solutions for these networks. Whilst this may appear obvious, all too often technical models are designed in the complete absence of, and therefore in ignorance of, the very community and users they are supposed to serve.

In the case of the Wray WMN, we recounted some of the challenges the community faced when managing their own network. Based on their experiences we believe that the issues of resource sharing will represent a major challenge within future user-provided networks and that any solution must give due consideration to communal requirements. We suggest that these requirements should be (requested and) arbitrated in a distributed fashion, by the community itself. We proposed the use of a simple abstraction (whereby an individual's request for resource is expressed as two parameters) as a means to simplify the problem and help us reason about how the design of a system to support this could be approached.

Whilst user behaviour clearly represents a key component in the operation of a user-centric network, scientific methods allowing us to model and predict this appear to fall short of our requirements. We identified a number of limitations with game theory – it was unable to model intricate information, whereby users could enter into co-operative or anti-social activities based on their inter-personal relationships with other users. Such scenarios caused us to stop considering users as anonymous rational persons and to start factoring in their personalities and beliefs.

It is evident there is a great deal of further work required to understand both the social and technical implementation aspects of the approach we advocate to communal resource management. As part of our continuing effort in this burgeoning research area, we are using the next-generation platform within Wray to build a community arbitrated resource sharing system. Our efforts focus upon the design choices relating to the resource request format (request/weight) and the method of arbitration within the community, including how this information is presented to users.

8. REFERENCES

[1] S. M. Allen, G. Colombo, and R. M. Whitaker. Forming social networks of trust to incentivize cooperation. In *HICSS 2009*, 2009.

[2] J. Barraca and R. Aguiar. Managing community aware wireless mesh networks. In *Computers and Communications, 2008. ISCC 2008. IEEE Symposium on*, pages 278–285, July 2008.

[3] C. Bicchieri. The Grammar of Society: The Nature and Dynamics of Social Norms. *The British Journal for the Philosophy of Science*, 58(3):613–618, 2007.

[4] S. Bowles and H. Gintis. The Moral Economy of Communities: Structured Populations and the Evolution of Pro-Social Norms. *Evolution and Human Behavior*, 19(1):3–25, January 1998.

[5] S. Bury, J. Ishmael, N. J. P. Race, P. Smith, and M. Rouncefield. Towards an understanding of security concerns within communities. In *WIMOB '08: Proceedings of the 2008 IEEE International Conference on Wireless & Mobile Computing, Networking & Communication*, pages 478–483, Washington, DC, USA, 2008. IEEE Computer Society.

[6] M. Fafchamps. Solidarity Networks in Preindustrial Societies: Rational Peasants with a Moral Economy. *Economic Development and Cultural Change*, 41(1):147, 1992.

[7] F. Hugelshofer, P. Smith, D. Hutchison, and N. J. P. Race. OpenLIDS: A Lightweight Intrusion Detection System for Wireless Mesh Networks. In *15th Annual International Conference on Mobile Computing and Networking (MobiCom'2009) (to appear)*, Beijing, China, September 2009.

[8] J. Ishmael, S. Bury, D. Pezaros, and N. J. P. Race. Deploying rural community wireless mesh networks. *Internet Computing, IEEE*, 12(4):22–29, July-Aug. 2008.

[9] J. Ishmael and N. J. P. Race. Routing challenges: Peer to Peer Applications on a Community Mesh Network. In *WiMeshNets 2006*, Waterloo, Canada, August 2006.

[10] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Experiences Applying Game Theory to System Design. In *PINS '04: Proceedings of the ACM SIGCOMM Workshop on Practice and Theory of Incentives in Networked Systems*, pages 183–190, New York, NY, USA, 2004. ACM.

[11] D. Parish, K. Bharadia, A. Larkum, I. Phillips, and M. Oliver. Using packet size distributions to identify real-time networked applications. *IEEE Proceedings Communication*, 150(4):221–227, August 2003.

[12] R. Sofia and P. Mendes. User-provided networks: consumer as provider. *Communications Magazine, IEEE*, 46(12):86–91, December 2008.

[13] S. Woolgar. *Configuring the User: The Case of Usability Trials.* Routledge, New York, 1991.