

Architecting Citywide Ubiquitous Wi-Fi Access

Nishanth Sastry^{*}
University of Cambridge
MIT

Jon Crowcroft
University of Cambridge

Karen Sollins
MIT

ABSTRACT

Cities around the world are currently considering building expensive Wi-Fi infrastructure. In urban areas, resident operated Wi-Fi access points (APs) are dense enough to achieve ubiquitous Internet access, provided we can induce the hosts to provide guest Wi-Fi access. However, sharing Wi-Fi involves taking on responsibility for the guest's actions. Our main contribution is a novel mechanism to handoff the host's responsibility to a trusted point by tunneling the guest's packets through it. The tunnel also guarantees that the guest's traffic cannot be subverted by malicious hosts. Using tunneling as a primitive, we show how to architect a citywide cooperative for *safely* sharing Wi-Fi with *legitimate* guests. We offer this as an economically viable alternative to investing millions in new infrastructure.

1. INTRODUCTION

People used to ubiquitous Wi-Fi at their workplaces will readily appreciate the ability to access the Internet everywhere. The economic importance of ubiquitous Wi-Fi connectivity rises dramatically in the face of new developments such as the convergence of cellular and Wi-Fi via standards such as Unlicensed Mobile Access [21]. For instance, T-Mobile USA recently introduced the ability to route cellphone calls over Wi-Fi¹.

Understanding the potential impact of ubiquitous Internet connectivity, many cities have moved to create citywide wireless access infrastructure. Usually, the solution involves Wi-Max as the backbone supporting a number of Wi-Fi APs. This is an expensive solution that can cost as much as \$150,000 per square mile². Wireless connectivity for a large city such as San Francisco will reportedly cost 15 million dollars³.

In this paper, we argue that citywide ubiquitous Wi-Fi access can be architected at near-zero cost because

^{*}Currently a visiting student at MIT CSAIL.

¹<http://www.nytimes.com/2006/07/29/technology/29phones.html>

²<http://www.jupitermedia.com/corporate/releases/05.07.06-newjupresearch.html>

³<http://www.elecdesign.com/Articles/ArticleID/12413/12413.html>



Figure 1: A map of San Francisco. Each red point represents one Wi-Fi AP (from WiGLE.net, accessed on Jul 26, 2007).

the network infrastructure is already in place: A majority of city dwellers have a broadband connection and a personal Wi-Fi AP at home. As Fig. 1 suggests, in dense metropolitan areas such as San Francisco, there is sufficient density of APs to achieve near-ubiquitous Wi-Fi by sharing access to APs amongst residents.

Our goal is to address the issue of how to induce hosts (who own the APs) to explicitly and safely share their Wi-Fi with legitimate guests, while at the same time ensuring that guests are protected from malicious hosts. We note that this is almost completely a sociological solution to a technical problem.

Our agenda is twofold. First, we examine the dangers involved in implicitly sharing Wi-Fi access, both from the perspective of a host providing an unsecured AP, and a guest hopping onto an unknown AP. Second, we question whether cities should be investing millions of tax dollars in building public Wi-Fi infrastructure and suggest a *safe, low-cost* alternative: forming a cooperative of residents to share Wi-Fi access explicitly.

Our main technical contribution is to identify tunneling as a mechanism to sidestep various legal and safety issues involved in sharing Wi-Fi access. To make Wi-Fi sharing safe for both the host and the guest, we need to remove latent trust dependencies between them.

The guest’s packets are routed through an opaque tunnel and ingress the public Internet from a third point trusted by the guest, typically the guest’s home. Since the rest of the Internet sees this the origin of the guest’s traffic, the home, rather than the host, becomes accountable for the guest’s actions (such as downloading illegal content). The home authenticates the guest before assuming this responsibility. The architecture also puts the home in control of the guest’s DHCP parameters and routing, which prevents malicious hosts from tampering.

To create a trusted network of end-points, participating city residents form a co-operative (co-op) in which members share their Wi-Fi APs with each other. The coop infrastructure helps track IP addresses assigned to members’ homes by their ISPs, and establishes tunnels by punching NAT holes when required. The co-op also enables simple optimisations for highly mobile guests (e.g. cars) who may be in contact with a given host AP for only a few seconds.

The rest of the paper is organised as follows. §2 discusses related work. §3 discusses the main obstacles that prevent or inhibit the sharing of Wi-Fi. §4 motivates tunneling as an approach to overcome these obstacles. In §5, we outline an architecture that supports the creation and maintenance of tunnels. §6 discusses optimisations for roaming and the scope and limitations of our approach. §7 concludes.

The following clarifies terminology used both in this section and in the rest of the paper: The *host* is the one sharing her Wi-Fi access point with a *guest* who is away from her *home* network where she has an always-on connection to the Internet. The tunnel gets established between the host and the home.

2. RELATED WORK

A low cost, “guerilla” approach that has previously been tried is to create a community-wide mesh network [1, 3, 18], but this requires separate network hardware setup, and can be difficult to scale beyond a core, dense facility because each node in the mesh needs to be within radio range of at least one other node.

Yet another approach is to provide Wi-Fi access in exchange for tokens which can be redeemed when the host travels to other APs in the city [12]. While such trading schemes would induce hosts with houses near popular spots like bus-stops to share Wi-Fi, home owners in more residential neighbourhoods would have little incentive to share. Since our goal is ubiquitous Wi-Fi, we opt instead for a co-op that provides equal incentive for all residents to share Wi-Fi.

Recently, a few commercial ventures such as Fon [6] and Whisher [22] have introduced different models for sharing Wi-Fi connectivity. Whisher seems to work by distributing the WEP/WPA keys to users authorised by

the host⁴, and as such, does not appear to be a viable solution for sharing Wi-Fi with guest users unknown to the host. Fon sells a custom Wi-Fi AP with firmware that requires new guest users to authenticate themselves with a Fon server using the captive portal technique. Captive portals are essentially dynamic MAC/IP Address filters containing a list of authenticated guests and therefore can easily be evaded by sniffing the address of an already authenticated guest and masquerading as that address.

To the best of our knowledge, all previous attempts to share Wi-Fi allow the guest to directly access the Internet using the host’s connection to the ISP. As we detail in §4, there are numerous legal and security related reasons why we adopt tunneling instead.

Mobile IP [14] also relies on tunnels, to make guest-node mobility transparent to external “correspondent” nodes. Mobile IP calls for triangular routing in which packets from the external node to the guest are routed through the home but packets in the reverse direction are sent directly. This is not possible in our scenario because hosts do not trust the guests to grant them direct access to external nodes. Also, mobile IP enables the external node to initiate contact with the mobile (guest) node. Consequently, the tunnel setup is more involved and requires the registration of a “care-of” address with the home agent whenever the guest moves. These choices make mobile IP difficult to use for fast-moving guests such as cars, which may be in contact with a given host AP for only a few seconds [2]. In contrast, in our system, the guest node typically initiates the contact, and we are able to incorporate simplifications aimed at enabling fast-moving guests. Traditional mobile IP has also been largely incompatible with NAT, while our design accommodates, and indeed, makes use of the reality of NATs and private IP addresses.

3. OBSTACLES

We will now review the obstacles from the host’s and the guest’s viewpoints that prevent or inhibit the sharing of personal Wi-Fi APs.

3.1 Host’s concerns

A number of APs are left unsecured either because their owners ideologically want to support free wireless access or because they lack the tech savvy to change the default factory settings of their AP. AP owners who are aware of the possible dangers and legal implications of sharing their Internet connection with potentially malicious guests nearly always secure their APs [8].

A natural concern of hosts is the possible loss of bandwidth. An inconsiderate guest should not be able to hog bandwidth to an extent that the host is unable to ade-

⁴<http://blog.whisher.com/2007/03/01/today-we-explain-sharing-your-wifi-with-whisher/#comment-46>

quately access the Internet herself. There may be a need to control guest bandwidth even when the host is not using the Internet, in order to keep the aggregate traffic envelope under limits acceptable to the host’s ISP.

Another worry, although many Wi-Fi owners are unaware of this possibility, is that a malicious guest may attack and infect other computers on the host’s network or even the wireless router itself [20].

A more difficult problem is that a guest may download illegal content such as copyrighted media files or pornography. Most ISPs’ Terms of Service explicitly ban their customers from performing or abetting such illegal acts and the host would be in violation of these terms because of the guest’s actions [9].

3.2 Guest’s concerns

Modern operating systems try to make connecting to a new wireless network painless and easy. This has led to a careless and even carefree attitude – most users do not hesitate to hop onto any foreign wireless network that is freely available. However, surreptitiously stealing bandwidth is ethically and legally questionable. Also, prudent guests should realise their susceptibility to attacks and infections by a malicious host in much the same way a malicious guest can attack a host.

Even worse, if DHCP is being used, the host’s network can assign the guest a phony DNS server IP Address, forming the basis for a sophisticated “pharming” attack [19]. By redirecting DNS requests for personal website names (web-based mail, bank websites etc.) to bogus servers that carefully reproduce the look and feel of the requested website, the guest can be conned into divulging passwords and other sensitive information.

4. TUNNELING THROUGH OBSTACLES

Our main claim is that the above obstacles can be solved if the only access granted to the guest is a tunnel to a single point in the Internet that the guest trusts. In this section, we motivate the reasons for this decision. We also briefly discuss some legal implications of sharing Wi-Fi, and overheads imposed by tunneling.

We expect that common network security practices will continue at the host end, supplementing the tunnels. For basic protection against malicious guests, firewalls should be enabled on the wireless router and/or the host’s computers. Several wireless routers as well as commercial (e.g. Fon) and open-source (e.g. DD-WRT) router firmware allow for setting bandwidth limits.

4.1 Motivation: Why tunnel?

The guest’s concerns can be solved if her traffic can be encrypted so that the host cannot snoop. However, most websites today do not support secure connections. To the guest, the tunnel provides a trusted intermediary that can encrypt *all* of its flows going through a

host that it does not necessarily trust. As we discuss below, the tunnel allows the *home* to set the guest’s DNS server, and protects against the pharming attack described above.

To resolve the host’s concerns, we need to examine the roots of malicious guest usage. The first issue is the nature of the content accessed by the guest. Guests should not be allowed to use the foreign host as an alternate Internet access point merely to evade restrictions that may be in-place at their home or other default Internet access point. These restrictions include the home ISP’s Terms of Service or Acceptable Use Policies as well as any router-based or ISP-provided parental control software⁵. A tunnel does not directly solve this problem. Instead, it hands off this responsibility to the remote tunnel endpoint. The reasoning is that if the guest has the rights to access the Internet from a point, then that point should have the responsibility and the authority to police, or limit access to content.

The second issue is the possibility of subverting ubiquitous connectivity for DDoS. Because of poor radio links (host APs could be behind walls or far away physically) and because hosts may limit the bandwidth they share, many guests could experience low-bandwidth links. Hence, legitimate guests within range of multiple host APs should be allowed to use them all simultaneously⁶, using multiple wireless cards or techniques like Multi-net [4] and multi-radio diversity [11]. However, allowing this in conjunction with ubiquitous Internet connectivity all over the city massively amplifies the bandwidth available and raises the spectre of DDoS. With our tunneling mechanism, traffic from the guest through any host eventually gets routed through a single choke point, the guest’s own home. This limits the amplification a malicious guest can achieve to the bandwidth that was already available at home.

A third, potential area of concern is that if an ISP offers value-added services such as real-time stock quotes or games-on-demand that are based on the customer at the end point, the guest, who is NAT’ed behind the same end point, will also obtain access. A similar situation arises when the host has a better point of connectivity, say a private T1 line. The host may not want to share these premium services if reciprocal access is not guaranteed. With a tunnel, the guest is only exposed to her own ISP, and access to premium services is not leaked. Furthermore, as discussed above, the bandwidth attainable by the guest is limited by the available bandwidth on her home network, which makes it safe to provide guest access even on expensive lines.

⁵We do not concern ourselves with parental control toolbars in browsers. They travel with the guest’s machine, and cannot be evaded merely by choosing alternate access points.

⁶To the best of our knowledge, Fon does not allow simultaneously connecting to multiple hosts, and this is a drawback.

In summary, most of our arguments boil down to eliminating all the reasons (other than being located away from home) that may lead a guest to use the host’s AP rather than her own connection at home.

4.2 Legal issues and accountability

Forcing the guest to create a secure tunnel removes the host’s “ability to supervise infringing activity” which can be a legal liability for an Access Point provider (A&M Records v. Napster, Inc., cited in [9]). Interestingly, courts in the US have held that a guest could be subject to the host ISP’s terms of service even if notice of the terms were not given (America Online v. LCGM, cited in [9]). Thus, it is in the guest’s interest to access the public Internet through a tunnel to her own ISP, whose terms of service are known to her.

Forcing access through the guest’s home network also improves accountability in IP traceback situations [16, 17], since the home node is now a part of any traceback path and legally binding terms of terms of service of the home’s ISP can be brought to bear upon it.

Many ISPs have terms of service that prohibit reselling Internet connectivity. By restricting guest traffic to a tunnel, the host is able to share Wi-Fi without sharing basic ISP-provided services such as DNS lookups, default routes, or IP address assignments. Also, since the tunnel is encrypted, everyone other than the guest and the home will perceive the guest’s traffic as originating from its home. The tunnel traffic itself will appear similar to peer-to-peer flows such as bittorrent and skype that already exist in the Internet.

Even so, ISPs do have a lot of control since they could contractually prohibit the tunnel traffic or affect the hosts’ willingness to share Wi-Fi by changing their pricing model. We imagine that in practical situations, co-ops could split the costs of membership with ISPs to carry the tunnel traffic. Or, municipalities could pass laws requiring the ISPs to support co-ops. The exact commercial and regulatory framework will vary from region to region and is beyond the scope of the paper.

4.3 Overheads of tunneling

The obvious worry with our approach is path length inflation, especially in cases where the host and the home are on two different ISPs that do not peer locally. However, for most purposes, guests only care about increase in latency. In broadband hosts, intracity P2P ping latencies are reported to be 30 – 60 ms [10] which is almost undetectable for most normal usage.

Bandwidth is more of a limiting factor. Notice that data downloaded by the guest is first downloaded to the home and then uploaded to the guest via the host. Most residential broadband connections are highly asymmetric. Thus the available bandwidth on home’s uplink (median 212 Kbps [10]) limits the guest’s bandwidth.

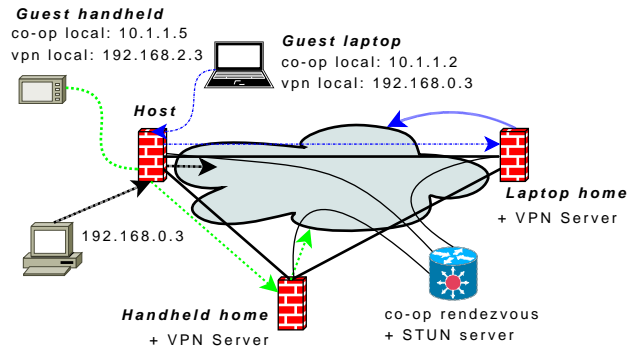


Figure 2: Architecture: Guest traffic is tunneled to home. The home end of the tunnel is a VPN server & authenticates guests. The host end is a firewall allowing only homebound packets. A STUN server helps set up tunnels when members are behind NATs. Guests are addressed by a fixed coop-local IP on the host’s network, and a vpn local IP on the home network. Both addresses may be reused outside their scope.

This is a serious limitation for some, but may be sufficient for common use cases. Studies are needed to determine whether guests can indeed achieve sufficient bandwidth in practice.

5. ARCHITECTURE

Tunneling is illustrated in Fig. 2, which shows guests with a laptop and a handheld sharing a common host AP. As shown, packets from machines belonging to the host’s network are allowed to directly access the Internet. The laptop’s and handheld’s packets, however, are tunneled through to their respective homes and ingress the Internet from there.

In this section, we discuss the underpinnings of an architecture for this kind of access: a co-op to form a trusted network of endpoints across the city and manage the guest identities, gateways at the host and home ends of the tunnel, and a STUN/rendezvous server to help setup the tunnel when the gateways are behind NATs.

5.1 Co-ops and co-op local addresses

To manage the guests’ identities, we assume that the residents of a city will form a co-op to share Wi-Fi access. Membership is voluntary (but required for guest access); but the scheme allows for incremental deployment. Even if only two residents are members, they can reciprocally share Wi-Fi at each others’ houses. The utility of the co-op grows as more residents join.

Each mobile device of a member is assigned a unique private IP address (say from the large 10.x space). We call this the *coop-local* address. Since the coop-local address is unique, members can freely use their devices for

guest access regardless of which other member’s devices are concurrently guests at the same host. Also notice that hosts can create a simple firewall rule to isolate *all* coop-local addresses from the rest of the machines on the host’s network.

One subtlety is that the coop-local address is only used to distinguish between nodes sharing a host AP. As we discuss below, the guest’s coop-local address is hidden behind the host’s NAT, so none of the other nodes in the Internet (including the home) need be aware of this address. The guest’s applications use a different “vpn-local” address (also discussed below). Furthermore, the home is configured as the default gateway of the guest. Thus, the coop-local address space may safely be reused elsewhere. In particular, the guest could log into her office VPN⁷ and access intranet machines which may also be in the coop-local address space.

5.2 NAT traversal

Many residential networks are assigned IP addresses dynamically. Some are also behind NATs. Consequently, gateways need to discover their current public IP addresses and punch NAT holes [7] to be reachable from the other ends of the tunnels. The co-op enables this through a STUN [15] server which the gateways can query to obtain their current NAT bindings (public IP address+port). As shown in Fig. 2, the server also facilitates rendezvous and tunnel setup by pushing the learned NAT bindings to all other members in a secure manner.

Periodic keepalives are used to keep NAT holes open. Also, a gateway that is behind a restricted cone or symmetric NAT may need to punch separate holes to each gateway in the co-op. In a few rare cases (e.g. NATs employing deep packet inspection), NAT holes cannot be punched, and packets between two gateways may have to be relayed through the rendezvous server.

In summary, all gateways know the current mappings between the other members’ coop-local address and the globally reachable IP address (or NAT binding) of their home gateway. We use this fact below.

5.3 Gateways: tunnel end points

As shown in Fig. 2, each member’s AP implements a gateway that performs different functions at the ends of the tunnels to support and maintain them.

5.3.1 Host gateway

At the host, where the tunnel begins, the gateway acts as a NAT and translates the coop-local addresses of guests as its own address to the external world (specifically, to the guest’s home). To send packets back to home, guests can remember the home’s IP from the

⁷This VPN would operate over the VPN tunnel between the host and the home that we describe below.

last contact. If the home’s IP has changed, and the guest sends a packet destined for the old address, the host drops it and sends the guest an ICMP “No route to host” message. The guest can then query and obtain the current address of the home from the host gateway.

5.3.2 Home gateway

At the home, where the tunnel ends, the gateway appears as a VPN server. We assume an SSL VPN-like solution, similar to OpenVPN. In OpenVPN, when the guest joins the home VPN, a virtual (TUN) device is activated on the guest machine and given a *vpn-local* address from the private IP address space. The vpn-local address is private to the *home’s* VPN and can be safely reused, for instance by other machines on the host’s network (as in Fig. 2).

The home gateway also acts as a NAT and translates the guest’s vpn-local address as its own address (the home address) to the rest of the public Internet. Thus, the guest may access the Internet through its own ISP-provided connection to the Internet. If only web access is desired, an alternate method is to run a HTTP proxy on the home’s VPN. The guest’s HTTP requests are forwarded to the proxy, which then fetches it from the WWW. If the proxy caches websites, this may be faster than the NAT option.

5.4 Security features

Malicious guests are not allowed to spray random IP addresses with traffic. As hinted in Section 5.3.1, the host gateway acts as a firewall with rules that only allow encrypted (VPN) packets between the coop-local addresses of a member and the current IP address (or NAT binding) of the member’s home. Alternately, or in addition to the firewall, the host could adopt a policy of cutting off a particular guest if no response is heard from the home within a reasonable time window – if the guest successfully authenticated with the home and opened a VPN tunnel, there should be traffic sent from the home, in response to the guest’s packets.

Coop-local addresses cannot be easily spoofed. The host end of the tunnel only allows packets destined to the home, and the VPN server at the home end authenticates the guest. Our scheme is impervious to Sybil attacks that depend upon the fact that IDs are only virtual, and one can create too many of them. In our case, since there is a real home address behind the coop-local IDs, and all packets pass through this address, Sybil attacks cannot be launched.

Notice that the ICMP message generated when a coop-local address sends packets to a destination other than its corresponding home is only sent to coop-local (10.x) addresses, and is thus contained within the host’s network. We ignore the minor risk that a malicious guest could send spurious ICMP messages to another

guest at the same host, although this could certainly be thwarted by more firewall rules.

Host firewalls cannot be compromised with false NAT bindings for malicious home gateways, because STUN requires its clients to authenticate themselves. Also, the host exchanges periodic keepalives with the home’s current address to keep NAT holes open, which would not work if the NAT binding was wrong.

Simple bandwidth limiting by hosts can also be highly effective in preventing DoS attacks by malicious guests.

6. DISCUSSION

6.1 Supporting highly mobile guests

With a very simple optimisation, highly mobile guests can preserve application sessions as they move from the range of one host AP to another. The guest’s applications use the TUN interface and its associated vpn-local address, which is assigned by the home’s VPN server. After tunnel establishment, the VPN server also pushes the VPN’s default gateway, DNS server, etc. to route all traffic over it. Since these values seldom need to change, the VPN server can be configured to always assign the same values for a given guest. In conjunction with disconnection-tolerant transports and TCP extensions [5, 13], preserving the vpn-local address allows higher-level sessions to be preserved across accesses through different host APs.

For faster link establishment, the vpn-local address to use for the TUN interface as well as the other network parameters above could also be statically configured on the guest machine. Establishing a secure VPN tunnel also involves an initial TLS handshake to exchange the random secret key used to encrypt the data. A shared static key can be used to avoid this overhead.

We carefully chose to use fixed coop-local addresses to eliminate the overheads of DHCP when a guest docks with a new host. Studies in [2] show that highly mobile clients such as cars are typically in contact with a given AP only for very brief periods (on avg. <10 secs. at 55kmph) and the time spent acquiring a new IP address can significantly affect the time available for useful data access. With fixed coop-local addresses and the above optimisations, docking with a new host only involves associating with its SSID. After this, the guest may start sending packets back to the home network.

6.2 Co-op scope and limitations

The co-operative is intended to be operated on the scale of a city. Creating tunnels over larger regions would cause unacceptable delays. Also, the host gateway’s firewall size increases roughly linearly with the size of the co-op, which places limitations on the number of members that can be supported. Larger co-ops can be supported by only creating rules for guests that

dock with the host. (Instead of creating firewall rules when the coop rendezvous server pushes a new coop-local ID to NAT-binding mapping, the host can query the rendezvous server for the current binding when a guest docks with the host. Clearly, this comes at the cost of slower link establishment for the guest.)

Finally, some countries impose restrictions on the content that can be accessed in their jurisdiction. If arbitrarily long tunnels are allowed, a guest who manages to obtain a trusted endpoint outside of this jurisdiction would be able to access “illegal” content. We do not prohibit or support such tunnels, but point to the possibility as a non-technical reason for why a co-op may need to be operated on a local, rather than global, scale.

7. CONCLUSION

In major cities of the developed world, the density of resident-operated Wi-Fi APs is sufficient to blanket the whole swathes of the city with ubiquitous Internet access. However, granting Internet access to a guest involves taking on responsibility for the guest. We have presented tunneling back to a trusted point in the Internet as a means for hosts to handoff this responsibility and thereby remove reservations about sharing their private Wi-Fi AP with the other residents of a city. Hosts further protect themselves by placing a firewall between guests and their computers. Guests can easily prevent hosts from snooping over their traffic by creating a secure tunnel. Together, the residents of a city could form a co-operative of trusted endpoints to *safely* share Internet connectivity throughout the city, rather than build expensive public Wi-Fi infrastructure.

8. ACKNOWLEDGEMENTS

We thank Rob Beverly, Ji Li, Eng-Keong Lua and Mythili Vutukuru for reading previous versions of the paper and suggesting many improvements. We also thank David Clark for clarifying some policy and legal aspects of our approach.

9. REFERENCES

- [1] BICKET, J., ET AL. Architecture and evaluation of an unplanned 802.11b mesh network. In *11th ACM MOBICOM Conf.* (2005).
- [2] BYCHKOVSKY, V., ET AL. A Measurement Study of Vehicular Internet Access Using In Situ Wi-Fi Networks. In *12th ACM MOBICOM Conf.* (2006).
- [3] CHAMPAIGN-URBANA WIRELESS. <http://cuwin.net>.
- [4] CHANDRA, R., BAHL, P., AND BAHL, P. Multinet: Connecting to multiple IEEE 802.11 networks using a single wireless card. In *IEEE Infocom* (2004).

- [5] FARRELL, S., ET AL. When TCP breaks: Delay- and disruption- tolerant networking. *IEEE Internet Computing* 10, 4 (2006), 72–78.
- [6] FON. <http://www.fon.com/>.
- [7] FORD, B., SRISURESH, P., AND KEGEL, D. Peer-to-peer communication across network address translators. In *USENIX* (2005).
- [8] GOOD, R. Should I secure my Wi-Fi access point? http://www.masternewmedia.org/news/2006/03/07/should_i_secure_my_wifi.htm.
- [9] HALE, II, R. Wi-Fi liability: Potential legal risks in accessing and operating wireless internet. *Santa Clara Computer and High Technology Law Journal* 21 (2005).
- [10] LAKSHMINARAYANAN, K., AND PADMANABHAN, V. N. Some findings on the network performance of broadband hosts. In *Internet Measurement Conference* (2003).
- [11] MIU, A. K., BALAKRISHNAN, H., AND KOKSAL, C. E. Improving Loss Resilience with Multi-Radio Diversity in Wireless Networks. In *11th ACM MOBICOM Conf.* (2005).
- [12] NIKANDER, P. Authorization and charging in public WLANs using FreeBSD and 802.1x. In *USENIX* (2002).
- [13] OTT, J., AND KUTSCHER, D. A disconnection tolerant transport for drive-thru Internet environments. In *Infocom* (2005).
- [14] PERKINS, C. IP mobility support. RFC 2002 (Standards Track), Oct. 1996.
- [15] ROSENBERG, J., WEINBERGER, J., HUITEMA, C., AND MAHY, R. STUN - Simple traversal of User Datagram Protocol (UDP) through Network Address Translators (NATs). RFC 3489 (Proposed Standard), Mar. 2003.
- [16] SAVAGE, S., ET AL. Practical network support for IP traceback. In *SIGCOMM* (2000).
- [17] SNOEREN, A., ET AL. Single-packet IP traceback. *IEEE/ACM Trans. Netw.* 10, 6 (2002).
- [18] SOCALFREENET. <http://socalfreenet.org/>.
- [19] STAMM, S., RAMZAN, Z., AND JAKOBSSON, M. Drive-by pharming. Tech. Rep. 641, Dept. of Computer Science, Indiana University, Dec 2006.
- [20] STOW, A. Can you trust a wireless router? <http://www.cs.indiana.edu/~atsow/mal-router/>. Accessed Jul 15, 2007.
- [21] UMA. <http://www.umatoday.com/umaOverview.php>. Accessed Aug 1, 2007.
- [22] WHISHER. <http://www.whisher.com/>.