RIK FARROW

WISPER: open source, long-distance wireless

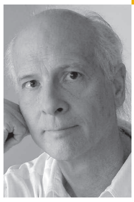

rik@usenix.org

AS THE DEVELOPED WORLD BECOMES ever more connected, the developing world falls ever farther behind. Projects such as One Laptop per Child (OLPC) seek to bring low-cost technology to aid in education to all parts of the world, but you might have noticed that there is a problem with wireless laptops in, say, central Africa. Where is the equally low-cost networking infrastructure required to connect these laptops to the global Internet?

I had considered this question, even as I mused about the reality of a $100 laptop (laptop.org). But it wasn't until Teus Hagen, Director of NLnet, caught up with me during USENIX Annual Tech '06 in Boston that I began to research the issue. Hagen had become infused with missionary zeal to create a new open source project that will, at the least, look into the prospects of creating software and related hardware platforms that will support low-cost networking.

Although the biggest focus on wireless technology has been for use in cities, wireless means that less densely populated areas might quickly become connected without large expenditures. And the research into adding free (or nearly free) wireless in large cities (e.g., projects such as RoofNet in Cambridge, MA) provide a giant step forward into creating actual software and hardware designs for expanding connectivity and communications into far-flung locations.

The Big Idea

Hagen has labeled his nascent project WISPER, an obvious play on the English word "whisper." Just like the low-powered radios used in Wi-Fi access, a network of low-powered devices that can be mass manufactured just might support the development of the communications infrastructures in much of the world. The OLPC project relies on the same economies of scale, that is, creation of a single design that can be mass-produced, bringing down the cost through economies of scale. The OLPC design does include Wi-Fi that will always be on (as long as the batteries provide power), holding out the possibility of a continuous, but purely local, network. But this type of network scales poorly. Clearly, something else is needed.

Hagen believes that the WISPER project could provide the next level of connectivity. While I can imagine that not all governments are interested in unfettered Internet connectivity, the ability to move beyond a local-only information base to a global one seems a logical next step—not an easy step, but certainly one worth exploring.

RoofNet

Perhaps the closest research project to WISPER that I've found so far is the RoofNet research network [1]. RoofNet involved a small number of systems (37) set up in a relatively small (and affluent by world standards) part of Cambridge, MA. Students installed Dell PCs running the RoofNet software package in the homes or apartments of participants. Installation included, in most cases, running a coax cable up to an omnidirectional roof-mounted antenna. In an area of mostly three- to four-story buildings, these antennas provided good coverage and were much easier to set up than directional antennas. Most systems were less than half a kilometer away from other systems, although there were outliers over a kilometer away.

RoofNet connected to the Internet via five gateways. As with the selection of RoofNet participants, the locations of gateways were more random than planned, so as a research platform, RoofNet better resembled an organically grown, rather than a carefully laid-out, wireless network. RoofNet not only worked but provided a mean bandwidth of 627 kbs and a median of 400 kbs, values that are not bad for an unplanned network using 802.11b as its wireless base.

The big story in RoofNet was not the hardware, or even that it was done, but the research that came out of the project. Some findings were not surprising, given prior work. For example, as the number of routing hops in a mesh network increases, the available bandwidth decreases. If you consider that someone who is four hops away from a gateway must share the bandwidth of intermediate systems with an ever-increasing amount of traffic, this finding just makes sense. But it makes the notion of large, unplanned mesh networks sound unworkable if it means that the number of hops between a client and a gateway must be, say, four or fewer [2].

RoofNet simply avoided the problem of the decrease in bandwidth with the increase in hop count, because no node was more than five hops from a gateway. RoofNet did, however, successfully deal with related issues. For example, traditional routing is based on following the best calculated routing metrics. Routers typically exchange reachability and response-time data among themselves and use those metrics to choose the next router to forward packets to. In RoofNet, a hybrid system, ExOR, was used instead for large file transfers.

All RoofNet Web traffic went through proxies, and these proxies would use ExOR for transferring the first 90% of large files and traditional routing for the last 10%. And using ExOR resulted in increases in transfer rates from 40 to 300%, even over single-hop routes. The key to ExOR was that it broadcast collections of packets, rather than unicasting the packets to the next hop, and included in this broadcast a manifest of the packets sent [3]. This approach takes advantage of the vagaries of Wi-Fi, where some packets might skip what would in traditional routing be the first or even later forwarders, and thus reach the destination system more quickly. ExOR also increases throughput in single hops, because it does not use RTS/CTS and has a different method for handling retransmissions of lost or damaged packets.

RoofNet also used Click [4] for assignment of addresses. Each PC was assigned an address based on the lower 24 bits of the wireless card's MAC address, with the higher 8 bits of the IP address remaining fixed. But each PC could also act as a router with NAT, so that locally connected systems would be given an address using DHCP. Because local systems were assigned private network addresses, participants had access to the Internet, but not to each other's systems.

The success of RoofNet is a relevant one, as it used innovative techniques to get beyond some of the limitations found in 802.11 hardware and standard IP software. But RoofNet itself is not a model for WISPER, where there might be many hops between a client and an Internet gateway. RoofNet was an experimental testbed.

Out of RoofNet

RoofNet has already produced a spin-off, a new company called Meraki Networks (meraki.net), which has already designed and produced the Meraki Mini, a $100 device designed for relaying Wi-Fi traffic. The Mini includes an Atheros chip that supports 802.11b/g Wi-Fi, an Ethernet port, a USB port, built-in encryption support, and a MIPs processor that runs a 2.4 Linux kernel. With the addition of some flash memory and SDRAM, the Mini is a self-contained system (just requiring a power supply) that can serve as a development platform for more mesh-style networking research.

I talked to John Bicket, the CTO and co-founder of Meraki Networks, and one of the students who was involved in the RoofNet research, about the Mini. Although the Mini sounds like the perfect hardware platform for WISPER, the match is actually far from perfect. Mini is designed to be used indoors, for example, making it a poor candidate for the environments found in developing countries. Meraki has aimed the Mini at commercial really low-cost wireless deployment in a building, such as an apartment building. Bicket said that the Mini is also designed with use by researchers in mind, as full documentation and sample device drivers are available. At $100 per box, the Mini costs just 10% of what each RoofNet node had cost, and less than many less capable devices used in mesh-network research.

There have been other network devices, such as the Meshcube (meshcube.org), intended for use in urban mesh networks, but these devices cost considerably more than the Mini. And none of the existing devices, or software, really satisfies Teus Hagen's list of goals.

WISPER

I must confess that it took me a while to understand what Hagen wanted from a WISPER project. At first, it seemed like all the parts were already in place—but they really are not, at least not yet. And Hagen has higher ambitions than a deployment in places where a RoofNet is just an alternative to several other means of having Internet connectivity. WISPER is really intended for places where the Internet has not yet penetrated—and not just physical locations.

Hagen wants WISPER to be a project not only for developing wireless networking but also for advancing IPv6. The very addressing issues that RoofNet neatly dodges need to be attacked head-on, Hagen believes, and a project like WISPER might be just the ticket. IPv6 was developed to deal with the shortcomings of IPv4, including the limited address space. But large-scale deployments of IPv6 remain scarce. And unplanned deployments, such as an ad hoc arrangement of wireless access points in a type of a mesh, really require a new way of thinking about addressing.

Hagen's complete list of goals for WISPER includes:

- IP address space solution, IPV6

- Multi-mesh scaling and answers to throughput issues
- Stability
- Operating systems other than 2.4 Linux (OpenWRT), such as 2.6, BSD, or even Minix 3
- Configuration and simplicity
- Autodynamic configuration
- Wi-Fi long distance: a cheaper WIMAX
- Standard hardware
- Open source code, to allow cooperative development
- Free access and free availability
- VoIP integration (Asterisk)
- Misuse measurements
- Security aspects (node and user authentication)
- Privacy (existing mesh solutions offer no privacy)
- A mobile mesh to allow Wi-Fi VoIP as an alternative to cellular phones

Hagen's wish list is certainly optimistic. But some items, such as auto-configuration, are a must for a device that will be installed and used by a very nontechnical population. Others, such as VoIP, may at first seem whimsical, but for parts of the world with little or no communications infrastructure and high illiteracy rates, the capacity to support voice communications will actually be very important.

Hagen turns out to be ideally placed for supporting an optimistic and open-ended project such as WISPER. Hagen is the Director of NLnet, an organization that came about when early Internet infrastructure in The Netherlands was sold to a commercial company. NLnet can invest millions of U.S. dollars per year in launching network projects, projects designed to become part of the public domain and be totally open source.

At one point in our many email exchanges, Hagen suggested that I liken WISPER to the old UUCP mail exchange networks. These networks were very ad hoc, relying on people using their own phone lines and systems to relay other people's email messages, files, and USENET postings. The Internet itself sprang from this very primitive dial-up network. In the same way, Hagen hopes that WISPER will help spawn the next level of global connectivity.

If WISPER excites you and you want to participate (or perhaps just comment on what may appear to you to be a "wild-assed scheme"), you can visit the Wiki at http://www.nlnet.nl/wifi.

REFERENCES

[1] Papers about the MIT CSAIL RoofNet research project: http://pdos.csail.mit.edu/roofnet/doku.php?id=publications.

[2] Some comments about mesh networks, including a geometric decrease in bandwidth as hop counts increase: http://www.oreillynet.com/pub/a/wireless/2004/01/22/wirelessmesh.html.

[3] Sanjit Biswas and Robert Morris, "Opportunistic Routing in Multi-Hop Wireless Networks": http://pdos.csail.mit.edu/papers/roofnet:exor-sigcomm05/.

[4] Eddie Kohler, Robert Morris, and Massimiliano Poletto, "Modular Components for Network Address Translation": http://pdos.csail.mit.edu/papers/click-rewriter/.

See also "Wireless Networking in the Developing World," with handy how-to info on boosting a Wi-Fi device: http://www.wndw.net/.