IEEE distributed systems ONLINE
Expert-authored articles and resources

IEEE DISTRIBUTED SYSTEMS ONLINE 1541-4922 © 2004 Published by the IEEE Computer Society
Vol. 5, No. 4; April 2004

Features Editor: Rebecca Deuel

Investigating Wireless Networks with WHYNET

Adam Stone

Sometimes, size matters.

At the University of California's Santa Barbara campus, computer science professor Elizabeth Belding-Royer has been studying intrusion detection mechanisms. So far, she's been able to run trials on only an eight-node network, but that's about to change. As part of a multicampus project funded by the National Science Foundation (NSF), she'll soon be expanding her test mechanism to between 20 and 40 nodes. What's even more significant is that she'll then be able to link the expanded network to similar networks at more than half a dozen other universities.

"It will better enable us to do global intrusion detection," she explains. With this vastly expanded network, she'll be able to push her technology to new limits, to see how it performs under real-world conditions.

Belding-Royer's efforts represent just a small part of WHYNET, a multi-institutional effort to build an interconnected test bed of wireless networks for use by the academic community. Funded by a US$5.5 million NSF grant, the program encompasses test beds at the University of California's San Diego, Los Angeles, Santa Barbara, Davis, and Riverside campuses, as well as at the University of Southern California and the University of Delaware. As the project evolves in the coming months, it will connect experimental networks at these schools via the Internet, letting researchers test new components on the entire system as they create and evaluate new technologies from the radio layer to protocol stacks.

EXPANDED R&D

Researchers hope the WHYNET infrastructure will let them test a broad range of evolving wireless technologies.

At UC San Diego for example, Ramesh Rao, Director of the San Diego Division of the California Institute for Telecommunications and Information Technology, has his eye on the cellular telephone network, which he says could play an ever-increasing role in the transmission of Internet data. Thus far, it's been difficult for Rao to test his ideas. "Cellular base stations are owned by cellular providers who have paid large sums of money for the right to use the spectrum. These are in the commercial world and as such they do not lend themselves to university-based research," he explains. Although his research team does hold an experimental Federal Communications Commission license to transmit in the commercial band, "this is nonetheless a world that has been largely closed to students and researchers in the university."

WHYNET could help open that up by bringing together base stations at several campuses, thereby creating a large enough cellular footprint to run trials of a significant scale. It's not the real world yet, but it's getting closer. "We would like to see more experiments run over this infrastructure—in particular, geographically diverse experiments running over this infrastructure—and here we have a larger research community eager to do the same thing," Rao says.

At USC, meanwhile, ultra-wideband (UWB) studies are the topic of the day for electrical engineering associate professor Urbashi Mitra. She's been testing diverse higher-level activities, such as security schemes, in a modular fashion to ensure that future wireless systems would always have the necessary assets available. According to Mitra, the best test methodology would be to use various different physical layers, but this can be challenging in a radio frequency setup, where it's not easy to modify the network.

To this end, the extended access WHYNET offers could be a boon, especially in exploring issues surrounding synchronization. UCLA already has a narrowband system up and running, and a wideband system in development at UCLA should be completed under the WHYNET grant. That grant will also speed development of the UWB work at USC. Taken together, these test beds should shed light on UWB networks' future requirements.

"Right now," says Mitra, "we don't completely understand the channels we are going to experience in the UWB system. The field is just very new, and so trying to build a real UWB radio system should give us a lot of insight in the unique features of such a system," especially as it interacts with narrow- and wideband infrastructures.

THE BIG PICTURE

Clearly, WHYNET's interconnected test beds will support research in a diverse array of disciplines. This isn't a coincidence, says Joseph B. Evans, program director of the Computer and Network Systems Division in the Directorate for Computer and Information Science and Engineering at NSF.

"The interesting thing in WHYNET is that there are a number of different types of wireless technologies that are being investigated, and this offers an opportunity to look at them both individually and together," he says.

In addition to the examples already mentioned, the NSF is interested in MIMO (multiple-input, multiple-output) wireless communication, software radios, and other cutting-edge wireless issues. "A multiorganizational effort like this is absolutely required when there are this many disparate technologies involved," Evans says. "No one campus has expertise in, say, mobile ad hoc routing and the low-level details of ultra-wideband technology." It takes a collective effort paired with collective resources, and that's just what WHYNET offers, as it combines the realism of physical testing with the scalability and flexibility of simulations. As planned, WHYNET will use devices and systems developed by participating researchers in conjunction with off-the-shelf components.

At UCLA, computer science professor Rajive Bagrodia gives an overview of WHYNET's projected scope of activities. He says researchers will use the network to explore, for example, software-defined radios, which let you change the wave form of the radio based on how you configure the protocol stack in software. Some will also study new protocols now in development at the MAC (medium access control)and routing layers, including protocols meant to exploit the characteristics of these emerging technologies as well as others being designed for more specific ends, such as improved quality of service or mobility.

The common theme here is to develop a whole that is greater than the sum of its parts.

"When people design these things, they typically do an excellent job of trying to design the best technology they can at that specific layer, but they are not necessarily looking at how they might interact with the other wonderful technologies that other people are designing at other layers," Bagrodia says.

The simplest examples come from the world of consumer technology. "Many of us have a cell phone, a PDA, a laptop, and a desktop. During the course of a day, you probably use all these multiple devices. Now, many people look toward the convergence of these devices, at least in terms of form factor," Bagrodia says. But that's not quite enough.

Your cell phone uses cellular technology, while your PDA taps into a wireless local area network. Might it be possible to have a converged device that not only incorporates all functions of both devices but can also select the best connectivity channel at any given moment? To find out, researchers will need an interconnected, intermodal test bed. "I'd really like to be able to do the switching automatically in a seamless manner, and I think that much of the research we are doing would hopefully aid in the development of the networking technology that would facilitate such a convergence," Bagrodia says.

It's all about getting beyond the hypothetical, says Rao. "If you want to do controlled experiments that allow you to understand what happens when information switches systems from one to another, you can theorize and calculate, but with a test bed you can actually make real observations." When WHYNET is completed, participating researchers should be able to upload models to the test bed remotely. They will thus plug into the larger system and, hopefully, be able to observe their new systems and devices at work in a real-world setting.

BIGGER TESTING

This is not NSF's first foray into the world of interconnected research beds.

At Rutgers University, NSF helped fund the wireless network test bed known as ORBIT (Open Access Research Testbed for Next-Generation Wireless Networks). A two-tier laboratory emulator and field trial network test bed, ORBIT aims to realize reproducible experiments, while also supporting the evaluation of protocols and applications in real-world settings. The project includes a local, real-world component, reaching out across campus and into the surrounding towns of central New Jersey, including about 50 nodes running a configurable mix of third-generation, high-speed cellular, along with 802.11 access. The infrastructure offers both short- and long-range wireless capabilities with an open application interface.

As with the Rutgers networks, the largely West Coast-based WHYNET should give researchers a chance to get down to the nitty-gritty of their ideas. For network engineers, this could open up new worlds of possibility.

Let's take the 802.11a card, for example. It ought to transmit at 54 megabits per second. But, if you put two computers side by side, transmitting over 802.11a to one another, they don't get anywhere near 54 mbps. "So the issue is, how do we better exploit this bandwidth at the application level?" Bagrodia says.

To understand the answers, you have to get outside the application level and into a wider testing arena. "Part of the reason we cannot fully exploit the capabilities is because we do not understand fully the interactions on all levels," Bagrodia explains. "Looking at each layer of the stack, we do not understand how they cooperate or compete to maximize that potential."

So that's one issue worth looking at in a bigger test network, but there's more.

THE GREAT UNKNOWN

Rao wants to know how you might configure a cellular network to handle Internet data with maximum efficiency. He's got a sensor network to analyze test results, and he hopes soon to turn those sensors toward every scientist's dream: The great unknown. The fact is, Rao has little idea what kind of information he is apt to glean from the WHYNET network—but that's just the point, he says. Until you put these cutting-edge projects to the test in a real-world environment (or at least a close approximation), you just can't know how they're going to perform.

You can test other science in the lab with great accuracy. If chemistry works in a test tube, it will work on the manufacturing floor too. Why doesn't this hold true for wireless? It's because wireless involves so many variables. Rain will affect performance, as will landscape interruptions. Too many people on the line can diminish quality. Small variations in the antenna's shape can change the scene dramatically. Add to this an untried technology—a new idea or device—and suddenly a complex picture becomes significantly more difficult to assess.

CONCLUSION

That's why these researchers are excited about WHYNET. In this particular case, size does matter. Researchers need a large-scale simulation. They need to set their creations loose, observe them in their native habitat. That's the only way to know whether, and how, these technologies will function as they mature into viable, marketable ideas.