IEEE distributed systems ONLINE
Expert-authored articles and resources

August 2006 (vol. 7, no. 8), art. no. 0608-o8002
1541-4922
Published by the IEEE Computer Society

News

Delay-Tolerant Network Technologies Coming Together

Greg Goth

Taken separately, delay-tolerant network research efforts often reveal cutting-edge details within specific realms of communications architecture theory. As the core concepts behind DTN mature, however, these seemingly disparate components also present a mosaic of what future networks might look like.

"A large community is sort of coalescing around the Delay Tolerant Networking Research Group," says Rajesh Krishnan, senior scientist at BBN Technologies. DTNRG is an open research group of the Internet Research Task Force, the research sister organization to the Internet Engineering Task Force standards body. "Because there's going to be a model which will hopefully get standardized through the IRTF and then the IETF process," Krishnan says, "people will have a standard way of doing communications in disconnected environments." BBN is currently working on a DTN project funded by DARPA's Disruption-Tolerant Networking Program.

The DTNRG actually had its origins in discussions surrounding the inadequacies of Transmission Control Protocol and Internet Protocol for deep space and interplanetary communications. The explosion of wireless networks, however, has given the concept's terrestrial implications a much higher profile.

Building blocks come from all over

One possible piece of the larger DTN puzzle was actually conceived when wireless networks' capabilities were far less robust—and the alternatives were far fewer—than at present. Called the WiNE (Wireless Network Environment) Sensor, the technology was the brainchild of Lt. Col. Lisa Shay, now a professor at the United States Military Academy at West Point, for her PhD thesis at Rensselaer Polytechnic Institute. WiNE can discern when a link in an IEEE 802.11 network is about to fail. It then switches the connection to another link.

"It was a few years ahead of its time," Shay says. "because when we first started talking about it, we envisioned wireless devices that worked over many different bands, or platforms that worked over many different bands or interfaces."

Shay, who tested her WiNE hypothesis in the late 1990s and into the year 2000, says the promulgation of alternative wireless platforms might now provide more practical opportunities to take the technology further.

"We had 802.11," she says. "Bluetooth hadn't come out. Zigbee hadn't come out. We had no support for a cell phone data rate like GSM in Europe. So now there are actual alternatives you could switch to, whereas before we were hypothesizing. In the Army we're still in the process of developing software radios, and that's a perfect application, because with a software radio you can switch to another band."

DARPA has taken the lead in funding much of the DTN work, and front-line military networks stand out as an obvious deployment for a delay-tolerant network. Still, commercial-sector applications for the technology include sensor networks that might be affected by weather or power loss and mobile networks in regions with little infrastructure.

BBN's Krishnan says he thinks commercial and defense-related research will proceed in tandem. He sees the core DTN protocols focused on commercial application features; features that bring the protocols up to military specifications can be added later via plug-ins.

"People have been dealing with disconnections for applications like email, RFID, and meter reading—things where the device or end user is not always connected—and have come up with custom applications," he says. The development of a standard core architecture around which discrete sectors can add features would be a boon for next-generation networks.

From packets to bundles

The DTNRG Web site has links to numerous Internet drafts and tutorials explaining the differentiation between TCP/IP-based networks and those based on core DTN principles. However, group cochair Kevin Fall says we shouldn't think of DTN protocols as supplanting TCP/IP networks. Rather, they are overlays where necessary. Fall says the group's efforts within the IRTF/IETF framework will focus on getting the drafts into informational or experimental Request for Comments. The IETF maintains these RFCs as persistent documents that are suitable for formal reference. While the specifics are still immature within the formal IETF network evaluation process, the basic concepts have met with wide agreement and early experimental implementations have been promising.

Krishnan differentiates the requirements of DTNs and common end-to-end-based networks with the terms "eventual stability" and "eventual connectivity." Traditional end-to-end networks assume a stable state throughout the network from the source to the end destination, whereas DTNs assume eventual connectivity and provide store-and-forward capabilities. For example, whereas intermediate nodes in a traditional network will jettison packets for which it can't discern a complete end-to-end path, a DTN must be able to provide facilities along the route that can store data intended for a disconnected node.

Perhaps the most basic differentiator between DTN architectures and the Internet is the concept of the *bundle*, the data unit upon which DTN communication is based. A bundle travels over networks that recognize and make provisions for intermittent and discontinuous connectivity. By contrast, the Internet's packet-based communication assembles data into packets, separates them, and transmits them to an endpoint that the protocol assumes is always connected.

The bundle layer is the overlay layer that DTNRG's Fall references. It sits between the transport layer and application layer of the various discrete networks the DTN bridges.

Fall says an associated naming scheme helps provide the flexibility for internetwork communication. The naming scheme is more general than an IP address, he says. Bundle lengths are variable and based on Internet-style URIs. "We can route on those, so routing is not only forwarding a packet, say, from one place to another, but also deciding not just where it's going but also into what other next-hop protocol it's going. So we could route an email message through TCP at an open point and then through an underwater acoustic network, at least in theory, and make those decisions at each point."

Early evaluations of the DTN bundle protocol at BBN have yielded promising results. In a presentation given at the March 2006 IETF meeting, Krishnan showed that experiments with DTN algorithms over networks with 20 percent availability successfully delivered 100 percent of 264 bundles within 40 minutes. By comparison, a baseline end-to-end topology hadn't yet delivered all its bundles even after one hour.

Meeting the next round of challenges

Several significant engineering challenges remain for the DTN research teams. Among them is the processing overhead of the bundle layer, but Krishnan says he thinks the ill effect will be negligible in most cases.

"Looking at microsensors, very tiny sensors that might not even have room to hold an entire bundle, it could be an issue," he says, "but more capable nodes don't seem to have that problem. Except for the tiniest of tiny computers, I believe the bundle is a reasonable way to go, and it is possible to come up with a mapping scheme to deal with that issue for tiny computers."

Fall says that researchers studying interlayer architecture have found standard interlayer API's to be "not really that good in providing error information—like how much stuff has been sent and not acknowledged or received. And those kind of details you can get if you're down in the protocol itself."

Among the issues the BBN team will focus on in future work is developing hybrid routing algorithms that can operate under either of the eventual stability (end-to-end) or eventual connectivity (DTN) assumptions, depending on which condition is present and will lead to best performance. The researchers are also perfecting late binding, in which intermediate nodes can route and/or store information or queries based on information they, but not the originating node, possess.

However, research such as Shay's, which was done wholly unconnected to the mainstream DARPA-funded work such as BBN's, might find itself a spot in the overall DTN architecture. In fact, Shay's thesis was a continuation of work performed by RPI professors Ken Vastola, Chuanyi Ji, and some of their other graduate students under an earlier DARPA grant pertaining to wired network management.

"In wireless networks," Vastola says, "various factors—the context, traffic flow, the movement of the nodes, the number of nodes—all of that is changing all the time. You need to figure out a whole set of baselines of what might be a good link or bad link."

The research around both commercial and defense DTN architecture might yield some interesting results fairly soon, especially as the results of early experiments such as BBN's are disseminated through the community.

"We have simple prototypes up and running," says Preston Marshall, DARPA's DTN program manager. "We plan to have essential features of our next-generation DTN operating in 18 months or so. We plan to make a small accelerator that can be plugged into any network to explore how DTNs would impact its performance."

Related Links

- DS Online's Dependable Systems Community
- "Security Considerations in Space and Delay-Tolerant Networks"

Cite this article:
Greg Goth, "Delay-Tolerant Network Technologies Coming Together," *IEEE Distributed Systems Online*, vol. 7, no. 8, 2006, art. no. 0608-o8002.