Leah Anderson

Delay Tolerant Networking: Adapting Internet Protocols to Unfriendly Environments

With the development of all kinds of new remote data transmission devices, typical Internet connection standards are becoming unsuitable for rugged connectivity demands of modern communication and high-tech scientific exploration. Delay Tolerant Networking (DTN) attempts to meet these new demands by combining the basic principles of the standard 7-layer DARPANET architecture with the flexibility and resilience needed to function in networks with intermittent connectivity or frequent transmission delays. While still an experimental design, DTN promises many socially beneficial applications that make it worth future investigation and investment.

Unlike the Internet that was designed for persistent end-to-end connections, DTN uses a "store and forward" algorithm. This means that instead of transmitting a constant flow of packet traffic back and forth between a communicating pair of nodes, DTNs group an entire message worth of data into a single "bundle". Then during transmission, a bundle is stored persistently at each node it encounters until a "custody transfer" to the next hop node can be verified.[1]

However, DTN cannot accomplish this algorithm within the confines of the typical 7-layer Internet model; instead it requires the networks to contain an 8th "Bundle Layer" protocol on top of existing Transport protocol. This additional layer essentially consists of some sort of persistent memory storage and methods to interface with the surrounding protocols, and it must be present on every node (along with an underlying Transport layer). In some implementations, the addition of this new layer has the secondary benefit of allowing a single Internet to contain nodes running different Transport Protocols. But because DTN is not yet formalized, exact application of these general ideas vary between designers and applications.

This "store and forward" method increases a network's tolerance of disruptions or delays in two ways: first, it minimizes end-to-end communications as much as possible, and second, it prevents the need for packets to be re-sent from the source every time data gets lost during transmission.[2] Perhaps the primary benefit of DTN is the elimination of the multiple back-and-forth transmissions that are required to negotiate a connection in a traditional TCP connection. Thus in a situation where long-distance transmission times are unreasonably long, all of the information required for a desired interaction can be sent back and forth in one or two "bundles" rather than requiring lengthy pauses in transmission for the reception of connection status messages. This would be an important feature in a network designed to connect remote spacecrafts or bridge multiple planets, for example. Additionally, a DTN has a safeguard against data loss along transmission paths: a "custodian" node will not delete a bundle until the subsequent node acknowledges that it has successfully received and copied. Hence when an active connection is disrupted, the custodian can simply restart its

---

[1] Warthman, Forrest. *Delay Tolerant Networks: A Tutorial.* (2003). Available at http://www.ipnsig.org/reports/DTN_Tutorial11.pdf

[2] *ibid*

intermediate transmission without having to send notification all the way back to the bundle's original source.

Yet there are still some major implementation issues that must be overcome for DTN to become practical for many types of non-traditional networking environments. Its lack of reliability is a big concern that is inherent to the structure of the technology. Consider, for example, a “custodian” node that becomes permanently isolated from all other nodes in a path to the destination of it’s bundle: this bundle will eventually just expire and get deleted without it’s original source necessarily having any knowledge of it’s fate. Security is another concern; since every node will store an entire message in the process of transmission, if a node is captured or physically removed from the system, then it’s message would become vulnerable to interception or alteration. These aspects are currently being researched by multiple organizations contributing to the development of DTN. [3]

One of the biggest remaining challenges preventing large scale adoption of DTN is the design of an efficient algorithm for message routing in an environment where a complete end-to-end connection may take hours (or even days) to be discovered. A handful of procedures for routing have been suggested, including some methods that probabilistically estimate the future availability of “next hop” connections and others that attempt to integrate a mathematical model of relevant behaviors or environmental factors affecting connectivity.[4] Future standardization of such routing techniques remains unforeseen.

Ultimately, the applications of DTN will be limited to situations where the benefits of its structure will outweigh its potential shortcomings. Examples include remote sensing applications where the surrounding environment may leave the network prone to frequent network partitions, cases in which network nodes are constantly moving in and out of range of connectivity, and long-distance transmissions for space travel, among others. Also, many researchers working on DTN appropriately envision it as a tool for providing connectivity to remote areas in developing nations through various ad-hoc networking or message ferrying techniques; it promises the potential to spread modern communications via roving network nodes to regions where investment in widespread network infrastructure would not be profitable. [5]

---

[3] Farrell, Stephen and Cahill, Vinny. *Delay- and Disruption-Tolerant Networking*. Boston: Artech House, 2006.

[4] Warthman, 2003.

[5] Farrel and Cahill, 2006.