



Folly Consulting

Internet Research and Development

The Challenge of Mobility in Low Power Networks

Elwyn Davies, Folly Consulting Ltd

Version 0.0 – 11 February 2011

Background

For the last 3 years I have been involved in an EU Framework Programme experiment-based research project called Networking for Communications Challenged Communities (N4C – n4c.eu). This project is intended to investigate methods of bringing Internet-style communications to areas where there is little or no infrastructure especially for power and wired or wireless communication. Our technology of choice for the experiments that we have carried out is Delay- and Disruption-Tolerant Networking (DTN - [DTNRG wiki](http://DTNRG.wiki)) using primarily the DTN2 reference implementation that is coordinated by the IRTF DTN Research Group.

The Challenge

We have been pretty successful in deploying very low power 'village routers' in relatively static locations using solar energy to overcome the lack of power infrastructure. The DTN technology allows data to be transferred using 'data mules' that pick up and drop off data during limited time encounters. The data exchange uses Wi-Fi connectivity. However the system has had to rely on an asymmetrical system, whereby there are well-known fixed points such as the village routers and gateways to the wired Internet which provide Wi-Fi access points while the data mules act as clients of these access points when they come into range. In practice the data mules are used to transfer data one hop between the static systems.

Our original intention was also to develop a system whereby 'the user is the network' in which fully portable devices such as smartphones and Internet tablets would act as data mules exchanging data on a relatively ad hoc basis with devices that they happened to encounter. Thus every node would be both a client workstation and a router/information repository.

The environment that we wish to target is such that at any given time a node is mostly able to communicate with at most one other node and for a large part of the time the mobile nodes will be out of range of any other node, whether mobile or static. Unlike conventional wired networks or the sort of networks envisaged in ROLL, multi-hop connectivity at one time will (almost) never be possible.

At the outset of the project we expected the major challenge to be getting the routing of data bundles (as the unit of data transfer in DTN is called) to work in a system where there is no stable connectivity pattern, but rather only a probabilistic chance of knowing whether a bundle should be passed to another node that is encountered as an effective way to get it to its destination. A number of the project participants had worked on a 'routing protocol' known as



Folly Consulting

Internet Research and Development

PROPHET [[PROPHET Internet draft](#)] that addresses this problem, creating a set of delivery predictabilities that mirrors the expected encounter patterns and adapts dynamically to changes in these patterns, relying on the idea that humans don't move about entirely randomly.

In practice there is another blockage on the road to making a 'symmetric' DTN system in which nodes are peers which all play similar roles in the network as opposed to the asymmetric system that we have demonstrated to work. However, before we look at the problem, it is important to note that at the level of DTN bundles and the exchange of these bundles there is no difference between the symmetric and the asymmetric case. The differences lie in how the nodes discover each other and the change from a static routing setup to a dynamic routing scheme.

Our attempts to create a dynamic network of portable devices was thoroughly stymied by the lack of a symmetric wireless system that (firstly) would work with nodes that were out of contact with any other node for a large part of the time but could quickly and automatically discover other nodes coming within wireless range without draining the portable unit's batteries in a matter of a few minutes and (secondly) could be relied on to establish contact with an arbitrary other piece of hardware nominally running the same 'standard'. The only system that would provide the system requirements and be currently available on the majority of portable hardware is Wi-Fi in ad hoc mode.

Issues with Ad Hoc Wi-Fi

The main inherent problem with ad hoc Wi-Fi is that the discovery mechanism involves active transmission of beacon signals at relatively frequent intervals. Unlike the situation with infrastructure mode where clients are essentially passive and only need to listen occasionally, the hardware cannot go into low power or sleep mode for a large fraction of the time, and the Wi-Fi transmitter needs to be turned on for a significant fraction of the time even when there is nothing going on and there are no prospective partners in range. Consequently when the Wi-Fi interface on a typical mobile phone or tablet is set to ad-hoc mode, power consumption is much greater than as an infrastructure client and battery lifetime is reduced from perhaps days to tens of minutes.

As an example, I have what is now a fairly old Nokia E61 phone that lasts several days between charges in normal usage. I have the JoikuSpot ([JoikuSpot Shop](#)) application installed which turns my mobile phone into what looks like a Wi-Fi access point but actually uses ad hoc mode. This is really useful as a quick and dirty way of connecting a standard laptop to the Internet without the need for a dongle, but don't expect the battery to last long! Maybe 30 minutes at best.

The other problem is that ad hoc Wi-Fi may be a standard but, unlike infrastructure mode, it has not been tested to destruction in heterogeneous environments. As a result although it is likely that two similar pieces of hardware will correctly negotiate an ad hoc connection, the same is not true for any given pair of dissimilar devices. Likewise there is very high risk that additional devices will not merge into an existing linkage as they are supposed to. But, let us give credit where it is due: different types of Nokia devices will generally talk to one another.

Consequently it is extremely difficult to carry out any meaningful real world tests of the 'network is the user' scenario with the commercial hardware available today. At best the user would have to



Folly Consulting

Internet Research and Development

enable ad hoc Wi-Fi for a brief period when a prospective partner hove into sight and even the number of contacts between battery charges that could be made would be very small. In the sort of power challenged areas where we are trying to experiment this is extremely limiting.

A Solution Comes into Sight

The Bluetooth organization has recently standardized Low Energy Bluetooth in v4.0 of the Core specification([Bluetooth Core Specification v4.0.](#)) In many ways the specification is similar to ZigBee in terms of wake up time (30ms), power consumption (as low as 1% of conventional Bluetooth) and range (up to 200m) but it has the advantages that it is closely related to existing Bluetooth technology, with dual-mode chips in progress, and is capable of symmetric operation whereas ZigBee requires a network coordinator.

Low Energy Bluetooth also offers much faster discovery and connection setup than either conventional Bluetooth or Wi-Fi. This is an important characteristic for the sort of ad hoc connectivity that we are interested in for DTN, as nodes may only be in range for a few tens of seconds and it is important to make the most of the available connection opportunities. With Wi-Fi we find that a significant part of the opportunity is often taken up in making the connection rather than providing goodput for bundles.

The Way Forwards

Folly Consulting is investigating the use of Low Energy Bluetooth in the sort of scenario that we originally wished to study in N4C but have not been able to do because of the hardware limitations. We currently have a small number of prototype devices for experimentation. We expect devices to be available commercially during the current year.

The original challenge for dynamic routing is indeed significant but we have made good progress in improving the original algorithms which were proposed for PRoPHET during N4C. We hope that we will be able to demonstrate the improved algorithms shortly and progress the specification as an Experimental RFC during 2011.

Overall we see the combination of Low Energy Bluetooth, PRoPHET routing and DTN as a way of providing connectivity in 'networks' where there is no well defined topology and connectivity is only available on a local one-to-one basis over short periods. End-to-end data transmission is achieved by sequence of such local connections spread over a period time that may last hours or even days. The advent of Low Energy Bluetooth appears to make it possible to use conventional smartphone type devices as data mules in a symmetric network while maintaining useful battery lifetime.