UCL DEPARTMENT OF
COMPUTER SCIENCE

Research Note

RN/14/17

Wi-Fi Goes to Town: Seamless Internet Connectivity in Metropolitan Underground Transport

8 January 2015

Kyle Jamieson

Abstract

In recent years we have seen Wi-Fi move from a “hotspot” model to fuller coverage patterns where connectivity is available everywhere, such as across university and enterprise campuses (recall, for example, the Eduroam service). As people have come to expect more from Wi-Fi of late, there has been concurrent interest in blanketing underground transit systems with wireless coverage, from Transport for London's involvement with Virgin Wireless to the New York Metropolitan Transport Authority's business relationship with the four US national mobile carriers. But these latter systems are focused on providing service to metropolitan underground platforms, frustratingly disconnecting and reconnecting users on their journey through the transit network.

We describe work in progress aimed at truly seamless Internet connectivity for the hundreds of thousands of daily commuters utilising a densely-populated urban area’s metropolitan underground transport system, light rail, or surface roads. We envision hundreds of London commuters on a single train simultaneously watching YouTube videos, surfing the web, chatting with friends and family over video, and downloading files while moving through an underground transit tunnel in a train or car at 30-40 miles per hour.

Wi-Fi Goes to Town

Kyle Jamieson
8 January 2015
Department of Computer Science
University College London
Gower Street, London WC1E 6BT, United Kingdom
Direct dial: +44 (0)20 7679 1390, Email: k.jamieson@cs.ucl.ac.uk

Abstract

We are building truly seamless Internet connectivity for the thousands of commuters utilizing a densely-populated urban area's metropolitan underground transport system. We will look at systems that can handle hundreds of commuters on a single train are simultaneously watching YouTube videos, surfing the web, chatting with friends over video, and downloading files while moving through an underground transit tunnel at up to 30 mph. The work will consider a rich design space consisting of varying antenna, inter-AP handoff, transport-, and physical-layer technologies. One of the key points in this design space is a system of large numbers of inexpensive APs along the ceiling of a tunnel, connected by backhaul Ethernet and controlled by predictive algorithms that anticipate where the user will be as the train moves at high speed.

Problem Setting

In recent years we have seen Wi-Fi move from a "hotspot" model to fuller coverage patterns where connectivity is available everywhere, such as across university and enterprise campuses.

As people have come to expect more from Wi-Fi of late, there has been concurrent interest in blanketing underground transit systems with wireless coverage, from Transport for London's involvement with Virgin Wireless to the New York Metropolitan Transport Authority's business relationship with the four US national mobile carriers, and the New York Metropolitan Transport Authority. But these systems are focused on providing service to platforms, frustratingly disconnecting and reconnecting users on their journey through the transit network. The result is a net win for all parties involved. Network providers gain advertising opportunities, the transit system gains revenue, either directly through customer billing or indirectly through deals, and riders reap all the benefits of being constantly connected.

The default approach to wireless coverage of long and narrow spaces is the *leaky feeder coaxial* cabling shown in Figure 1(a). Unlike conventional coaxial cable, the leaky feeder cable radiates energy along its length. Runs can reach 200 meters or more when amplifiers are used at regular intervals (40–50 meters for typical Wi-Fi power levels and leaky feeder dissipation) along the run.

But the leaky feeder faces three fundamental problems in its application to our problem:

1. The leaky feeder emits the same signal along its length, preventing *spatial reuse* of the shared wireless medium between two or more regions covered by the same leaky feeder cable/amplifier run.
2. Advances in single-user MIMO (standardized in 802.11n) cut against the grain of leaky feeder deployments. To see why, consider again the deployment shown in Figure 1(a). Using lightweight channel state information (CSI) measurements from a tool by Halperin *et al.* [1], we measure the wireless channel several hundred times per second as a user moves down the stairwell. Using standard methods [2], we then estimate wireless channel capacity as an upper bound on achievable throughput. Figure 1(b) shows this capacity in the presence of mobility with a leaky-feeder based access point (AP): notice the deeper fades that come about from using the leaky feeder, as compared to the more predictable fading pattern of the AP handover shown in Figure 1(c).
3. Advances is multi-user MIMO (standardized in 802.11ac) also cut across the grain of leaky feeder deployments. 802.11ac divides up the shared wireless medium spatially, but when the signal travels along the entire length of the leaky feeder, the ability to divide the medium spatially becomes degraded for the same reason as (1).

In contrast to the leaky feeder, conventional Wi-Fi APs are designed to radiate power roughly equally in all directions, even those using the newest 802.11ac access points. Deploying one or more Wi-Fi APs in a long, often curved tunnel yields diminishing returns, as most of the signal bounces off the walls of the tunnel instead of propagating along its length.

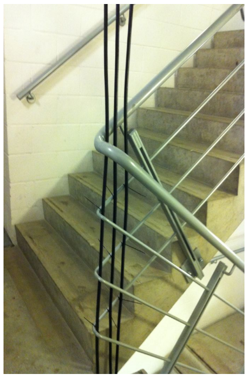

(a) A leaky feeder deployment at UCL in a stairwell.

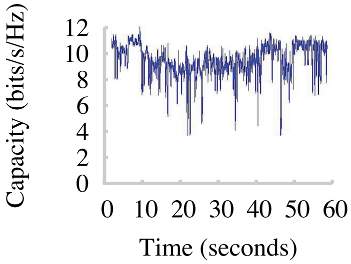

(b) Wireless capacity (bits per second, per Hz) for a user following the same path as in Figure 1(c) but served by *leaky feeder* cables shown in Figure 1(a).

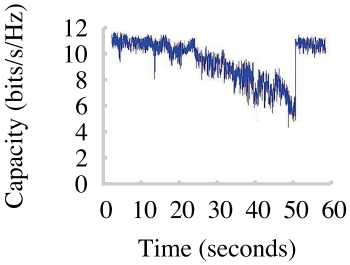

(c) Wireless capacity (bits per second, per Hz) during a handoff between two low-powered APs as a user descends the stairs.

Figure 1: Comparing wireless capacity during AP handoff with a leaky feeder in a test deployment at UCL.

Description of Proposed Work

We are interested in exploring the design space of this problem, starting from an array of low-cost, low-power APs connected by an Ethernet backhaul. The design space is rich, and roughly breaks down into the following factors, with a high degree of interaction between the different factors:

Cell size. By tuning transmit power, we can adjust the cell size, trading off the overhead associated with roaming between APs for increased amounts of spatial reuse down the length of the tunnel. We anticipate experimenting with cell sizes of 5–15 m as a starting point.

Handoff protocol. The research challenge that goes hand-in-hand with tuning down the cell size is to develop new algorithms to coordinate the backplane connecting the APs to the rest of the Internet.

While there are many existing handoff protocols for both WiFi and mobile cellular data networks, we find ourselves in an extreme position in the design space when an underground train passes by a series (hundreds) of five-meter cells at 30 mph. Imagine TCP flows traversing between an Internet server and a user's mobile with inter-packet times in the 100s of milliseconds to seconds. In this situation, traditional roaming protocols will either incur unnecessary overhead in the form of control packets to signal transitions between APs, or fail. We intend to investigate "predictive" handoff protocols that send downstream packets to the AP the user will be near once the packet reaches our system's backhaul.

This decision will likely involve estimating the train's speed either indirectly or directly, estimating the strength of each AP's wireless channel based on the estimated speed, and then making a forwarding decision based on these estimates. A certain amount of "fudging" is possible by allowing two or more APs to transmit downlink traffic together, but spatial reuse then suffers, so the challenge in this part of the work is to make the best forwarding decisions based on available CSI and user inertial sensor data (accelerometer and gyroscope) measurements.

Antenna design. While leaky feeder technology on its own has the problems noted above, there is reason to believe that relatively short leaky feeder runs may be advantageous, as the shape of the transmitting antenna (the leaky feeder) more closely matches the shape of the radio environment (the tunnel). Short leaky feeder runs would also reduce the amount of handoff needed, possibly improving AP prediction accuracy, because the signal strength a mobile client experiences as a function of horizontal displacement from the AP is flatter for a leaky feeder-equipped AP than it is for a conventional antenna-based AP, as the data above shows.

Methodology

Our approach will be prototype-driven: we intend to build and if possible deploy our designs in a real tunnel, leveraging ties we have with colleagues at UCL who have connections to Transport for London as well as other relevant research projects.

The rough work plan will involve prototyping designs that explore key points in the above design space. We will look at the following designs at the outset:

1. A single amplified leaky-feeder run that is as long as practicable.
2. A small cell-size (5 m) conventional AP distribution network connected with an Ethernet backhaul. We are considering the Raspberry Pi for this design point.
3. A medium cell-size (10-20 m) leaky-feeder augmented AP distribution network connected with an Ethernet backhaul.

Once connected to the rest of the Internet, these systems will provide fertile ground for exploring the handoff protocols sketched above. We will then be able to augment our theoretical capacity results with empirical TCP and RTP-based throughput readings from the testbed.

Data Policy

We intend to disseminate the output of the work in open academic venues (conferences, journals, and open-access publications). We hope to build relationships with the transport organizations mentioned in the proposal so that the public will directly benefit from the networking designs we build. Transport for London in particular has a history of working with universities to provide access to some data about their network, and we hope to continue in this vein.

References

[1] D. Halperin, W. Hu, A. Sheth, and D. Wetherall. Tool release: Gathering 802.11n traces with channel state information. *ACM SIGCOMM CCR*, 41(1):53, Jan 2011.

[2] D. Tse and P. Viswanath. *Fundamentals of Wireless Communication*. Cambridge Univ. Press, 2005.