

Large-scale Measurements of Wireless Network Behavior—Public Review

Kyle Jamieson
Department of Computer Science
University College London
k.jamieson@cs.ucl.ac.uk

This year marks the 30th anniversary of the F.C.C.’s decision to open up the so-called “garbage bands” at 900, 2400, and 5800 MHz to the data communications industry. When the N.C.R. Corporation subsequently introduced its WaveLAN product to the market in 1988, the world gained un-tethered Internet connectivity for mobile devices, and the research community found a significant source of inspiration for new designs and ideas.

Starting at the top of the networking stack, researchers in our community began investigating new ways of designing these wireless networks in the nineties, with the scope in the following decade widening from bits to include signals, accelerating the rate of innovation in our community.

But a skeptical narrative has run concurrently with this stream of research, rightly questioning the motivation for much of the community’s work. Questions raised include:

1. Is interference between Wi-Fi networks really a significant problem, and to what extent does the interference that impacts Wi-Fi arise from Wi-Fi devices themselves, as opposed to other types of radio transmitters operating on the same frequency bands?
2. How occupied are the frequency bands that Wi-Fi operates in, and how rapidly has their usage increased over the past one to two years?
3. What applications use most of the spectrum, and how is their traffic split between the uplink *v.* the downlink?

This paper by Biswas, Bicket, *et al.* describes large-scale measurements of wireless network behavior that begin to answer the above questions. Meraki is a cloud-based management architecture for Wi-Fi networks. Controlling hundreds of thousands of access points serving many millions of clients, the platform enables the authors to take an unprecedented number of measurements about how we are using Wi-Fi. An interesting feature of the Meraki MR18 access point is the presence of a third radio, dedicated to scanning for nearby networks while its two primary radios service clients.

The paper’s longitudinal analysis quantifies trends in the number of nearby Wi-Fi networks each Meraki-managed network sees across a six-month interval, the link delivery rates

each access point experiences across a six-month interval, and the adoption of new wireless technologies within a two-year window. The paper breaks out trends in daytime versus nighttime usage as well as analyzing the type of client radio hardware and software applications the network saw. Of note here are the two traffic classes with the outlier greatest percentages of uplink utilization: Dropcam, comprising 0.42% of total traffic, yet with the fewest clients among the top 40 applications, and online backup, comprising 0.15% of total traffic. Reflecting the state-of-the-art, the measurement methodology cuts across layers from the physical layer to the application layer, while capturing longitudinal trends.

The program committee appreciated the unique scale of the measurements and sound experimental methodology in the paper, noting that this is one of the few measurement studies from an operational carrier. Reviewers agreed that the experimental results the authors highlight are interesting, especially the lack of correlation between the number of access points overheard nearby and the wireless channel utilization. Also notable are measurements showing that the majority of Wi-Fi interference is attributable to decodable Wi-Fi transmissions.

While the focus of the paper is on data and not deployment experience, the paper does include a section highlighting these experiences. In this section the authors relate an interesting anecdote about the impact of Manhattan skyscrapers on Wi-Fi access point behavior, and describe the challenges the advent of smartphones created *ca.* 2006–2012.

Also notable is Cisco Meraki’s decision to publicly release large parts of the datasets this paper describes to the community in an anonymized form. I applaud this decision for two reasons. First, access to the paper’s data allows researchers to better focus their efforts on the most important problems. Second, open access allows the research community, industry, and government regulators to better understand what factors currently limit the performance of these networks, and gives these parties an opportunity to work together to address these factors. I hope other operators follow in Cisco Meraki’s footsteps for a more complete and open understanding of Wi-Fi performance.