Scale-Up Community-Level Measurements of Internet Adoption and Performance

Nick Feamster, Kyle MacMillan, Nicole Marwell, Tarun Mangla, Jamie Saxon
University of Chicago

For years, the Federal Communications Commission (FCC) has pointed to the "homework gap" as evidence of how disparities in broadband Internet access result in disparate outcomes across communities. In light of the COVID-19 pandemic and the immediate public emergency surrounding online education, many cities have prioritized the deployment of broadband Internet access to underserved communities.

In Chicago, the current program is a subsidy called "Chicago Connected". The program funds ISP subscriptions for three to four years. In the immediate term, there is a need to understand (1) which households need connectivity; (2) the extent of uptake (and, when uptake fails, the social, economic, and cultural reasons for uptake failure); (3) the ultimate performance of the deployed service, and the extent to which the provided service allows households to perform vital functions, such as remote learning. Over the longer term, cities such as Chicago need more sustainable solutions, including community-based networks (e.g., rooftop radio networks, community WiFi, mmWave 5G).

Long-term sustainable solutions will require bringing affordable, reliable, performance connectivity to these communities. Yet doing so starts with *holistic measurement* that provide a thorough, accurate, and precise understanding of where the problems lie. In particular, cities and civic organizations are seeking answers to technical, social, political, and cultural questions to help address these problems. Technical questions include:

- What is the nature and extent of deployed physical infrastructure, from underground fiber to cell towards to cable buildout to poles?
- What access link performance is needed to support and sustain the applications that households rely on, from remote education to work to telehealth?
- How do various wireless radio technologies (or combinations thereof) perform in real deployments, both in terms of baseline network performance and in terms of application quality?
- How do performance factors vary with the number of residents and usage patterns in a household?
- To what extent are performance bottlenecks located in the access ISP, and to what extent are they located in other parts of the end-to-end path, including device endpoints (e.g., Chromebooks), home networking equipment (e.g., WiFi routers), and other parts of the network?

In this regard, existing datasets tend to shed light only on parts of the problem (e.g., Internet "speed"); some of these datasets, although widely used, do not accurately reflect true gaps in connectivity because the measurement techniques themselves are not designed to measure the throughput of an ISP access link.

A critical component of our project thus involves the development of a technical framework for gathering a more complete set of measurements concerning the above questions—at a far finer granularity than is available in existing datasets. Doing so will be challenging because some datasets (e.g., application quality, home wireless performance) ultimately require instrumentation in home networks, which can be difficult in normal circumstances and is even more challenging in underserved communities. Other datasets will require completely new approaches to measurement, including computer vision applied to a combination of open street view data and crowdsourced images of deployed infrastructure.

Ultimately, bridging connectivity gaps in underserved communities goes beyond technical deployments. Other important aspects include community engagement (e.g., the presence of active community-based organizations), residents who recognize the need for broadband connectivity, and affordable access. Understanding these aspects of the problem can be as important as the technical deployments, because often such deployments succeed or fail based on perceived need and engagement from the community. For example, an interesting initial data point from the Chicago Connected program is that many eligible households do not enroll in the program, even though it is free. Anecdotally, reasons for lack of uptake range from lack of perceived need to fear of government intervention. Understanding the nature of these barriers, as well as how they vary by community—will be a critical aspect of this dataset.

This newly launched project involves a broad array of partners, including the City of Chicago, Chicago Public Schools, University of Chicago Medicine, City Tech Collaborative, Kids First Chicago, and Ookla. Our goal is to pilot the above approach—and curate the corresponding datasets—for two or three neighborhoods in Chicago, with the ultimate goal that the frameworks we develop can serve as a rubric and template for replication in other cities and communities, and ultimately nationwide.