

David Choffnes *Northeastern University, Michigan* Z. Morley Mao *University of Michigan*

Editors: Jacob Sorber and Lin Zhong



SUMMARY REPORT

for the First NSF Workshop on Mobile Community Measurement Infrastructure

Flexible measurement infrastructure support for experimentation are critical for enabling and supporting research work with strong experimental components to help validate ideas and evaluate the design in practice. Mobile computing research work must be evaluated in as realistic network settings as possible to help researchers understand how various factors, such as signal strength, network load affect the performance and energy metrics of interest. Currently, there is no environment that enables network visibility from an end-to-end perspective of the cellular network protocol stack: from the end-device all the way to the network server traversing through various network elements. The lack of such an integrated measurement infrastructure greatly hinders the innovation in this important research field.

In November, 2014, we held the first NSF workshop on Mobile Community Measurement Infrastructure. The goal was to identify the key requirements for designing and developing a mobile community measurement infrastructure to support cutting-edge research in mobile computing, going beyond the current research interests to support experimentation in next-generation mobile networks, such as 5G. Today's mobile systems are deployed with ad-hoc measurements providing limited visibility into network behavior; however, we need integrated, cross-layer scalable and flexible measurement infrastructure to enable transformative and bold research endeavors. In particular, most commercially deployed mobile network infrastructures are completely closed to researchers, making it difficult to experiment with or understand the impact of today's systems or future technologies. We argue that measurements for mobile wireless networks must be designed from ground up, rather than as an after-thought after the protocol and infrastructures have been built and deployed.

A goal of this report is to continue the conversation on how the community will conduct measurement in an effective, sustainable and collaborative manner. Our workshop mission aligns with previous and concurrent workshops, namely AIMS 2014 [1] and the NSF Workshop on Future Research Infrastructure held jointly with this workshop. However, our mission is more focused on measurement and monitoring support, instead of infrastructure development in general. In fact, each critical part of infrastructure should support measurement effectively, to enable real-time debugging and diagnosis, in order to evolve the wireless network for future application and user demands.

IT IS CLEAR THAT MOBILE COMMUNITY MEASUREMENT INFRASTRUCTURE IS A BROAD AREA WITH MANY CHALLENGES TO ADDRESS ACROSS MANY DISCIPLINES AND LAYERS OF THE NETWORKING STACK.

ISSUES AND CHALLENGES

The workshop focused on four high-level topics: (1) measurements and application needs, (2) testbeds, (3) operational challenges, and (4) data collection and privacy. Based on presentations, discussions and feedback from participants, we identified the following key challenges that need to be addressed in these areas.

Measurements and Application Needs

An important challenge for any measurement infrastructure is understanding what measurements to make and how they will benefit applications. We discussed a variety of current and future measurements and applications, and identified the following key challenges.

- **What data should we collect?** From fine-grained, low-level measurements to application-layer QoE metrics, the range of potential data to collect is vast. There are common challenges across all measurement domains, including coverage across time and locations, efficiency of measurements, availability of measurement data, how to use the data once it is gathered, how to crowdsource measurements, what applications will benefit and how to support infrastructures that benefit the most amount of research. It appears there is no silver bullet for determining how to instrument wireless networks; rather, flexibility, evolvability, and efficiency are essential components for infrastructure that supports measurements.
- **How do we get meaningful data?** Even if we determine what data to collect, a key challenge highlighted by participants is that today's systems are closed and may not support measurements we need. Further, even if we could get

these measurements from deployed infrastructure, it's unclear how to obtain representative datasets for user/device populations because this is often kept private by telecom operators. We believe that building open, extensible infrastructure and improving access to useful datasets are essential to enabling future innovation in wireless networks.

Testbeds

Without a large operational network to use as a playground for experiments, most researchers turn to testbeds. The research and operational communities have built a variety of testbeds each focusing on different aspects of wireless communication, but there is little work that joins the disparate testbeds into a platform that incorporates end-to-end network characteristics. Doing so is no easy task: it requires potentially covering all layers of the network stack from physical to QoE, to providing scale, diversity and measurement support for a variety of intended environments.

Another important thread of discussion is related to the data collected from such testbeds. To be information, we argue that a testbed should take into account the quality of the data collected (e.g., coverage, completeness, accuracy and integrity – all particularly relevant issues in crowdsourced data). Further, we need to augment these testbeds with data analytics that allows us to reveal deep insights and patterns.

A common pitfall is that testbeds are usually created with a target application in mind. As a result, they are carefully tuned to the PI's research needs. Adapting the testbed for shared use is often surprisingly difficult and creates in sub-optimal results. As a community we need to think about how to encourage the design of shared, flexible testbeds. A related concern is that we are sometimes replicating or performing work that might be done better by industry (or has already been done) – highlighting the need for researchers and practitioners to maintain an open dialogue to avoid redundant work.

Given the current state of testbeds, we raise the following questions for future testbed design:

- **Testbed purpose:** What is the purpose of a testbed? How can it be broadened to

support the maximum number of projects?

- **Data collection:** What is the structure that would best benefit industry, government and academia to build an integrated testbed or collect measurement data? How to deal with the noise, privacy concerns and incentives when collecting crowdsourcing data?
- **Integration with existing testbeds:** How can we best utilize existing testbeds? What new infrastructure is needed to perform exciting measurements?
- **Industry involvement:** What is the main challenge the wireless industry is facing? Do academic studies provide good feedback to the industry?

Operational Challenges

The workshop featured a panel discussion from representatives of mobile wireless providers to help understand operational challenges. We summarize the key issues and challenges identified in this panel here.

- **Instrumentation challenges:** Today there is a limited ability to perform instrumentation, e.g., phones, base stations, and doing so can be costly: monitoring baseband information requires more energy, as waking up the phone for measurement can be expensive. In general, companies are not interested in altruistic investment to make access to data easier; to address this, we need to provide direct incentives and clear use cases to convince them to invest in instrumentation.
- **Network complexity:** Today's wireless networks are complicated; the increased prevalence of virtualization only makes it harder to diagnose problems and optimize networks. Identifying how to instrument systems and use this data to master complexity is a key open challenge.
- **Innovation in areas controlled by industry:** We need infrastructure that allows us to innovate in areas currently controlled by cellular providers. In addition, a standardized, highly configurable virtual platform that supports a wide range of measurements can help us make progress in the face of limited access to operational deployments.
- **Realistic data and infrastructure:** It is important to recruit real users to use experimental infrastructure to provide realistic workloads and use cases. As



concrete action items, the research community can build our data archive and create benchmark at least for use by the research community. Another focus can be to develop experimental infrastructure and open testbeds facilitating research in this area.

Data Collection and Privacy

An important challenge for mobile wireless measurement is that the data collected can potentially reveal significant information about users, whether doing so was intentional or not. We identified the following challenges for collecting data and maintaining privacy/security.

- **Standardized mechanisms for allowing policy-based access to data:** Picking data-sharing models can be challenging and models tend to be one of: anonymize everything and share with everyone, no sharing, or share it with anyone who signs NDAs. A key open challenge is that we are not lawyers and legal implications can vary by jurisdiction.
- **Incentives for industry to share data:** Government agencies can fund industry/academia collaborations that require open datasets. Another solution is to collect some “case study” or “benchmark” data from industry. While not necessarily broadly representative, it can be used for benchmarking – similar to the Netflix challenge (but hopefully without the privacy problems). We also discussed the

utility of data sets with mobility, both for understanding human mobility and for mobile simulations and evaluations. In this context, industry has more data, as well as data processing technology, but getting data from industry is difficult.

- **A norm for data privacy/security accepted by the community:** There is much confusion in the community around IRBs, ethical standards and privacy — examples of which vary according to where you are and who you ask. One participant suggested that the community needs to develop a set of community “norms” for ethical mobile measurement when human subjects are involved. The NSF and ACM could perhaps support the community in developing and disseminating these norms among research institutions and their IRBs and brokering connections to relevant international organizations.

RECOMMENDATIONS AND FUTURE DIRECTIONS

It is clear that mobile community measurement infrastructure is a broad area with many challenges to address across many disciplines and layers of the networking stack. We heard a wide range of talks about research efforts that were interesting individually, but there was a clear need for focus, communication, coordination and collaboration across different research, industry and governmental groups.

Through discussions and survey feedback, the following common themes emerged for recommendations to make progress toward sustainable, innovative mobile community measurement infrastructure:

- NSF funding of long-lived infrastructure to host collaborative testbeds for mobile measurement. There needs to be a program that funds long-term development and maintenance of measurement infrastructure as a service to the research community.
- We recommend that researchers should work with vendors, service providers and government advocates to ensure measurements are integrated into wireless systems. There is a need for open and innovated testbeds, including low-layer tools for performing measurements, systems/architectures to simplify measurement collection and techniques to use/combine the measurements effectively. Last, we need better instrumentation to understand the performance difference for new and current spectrum uses.
- Incentives for researchers and industry to work together and share data. This can in part be solved by government policies/ programs providing “carrots” or “sticks” to encourage cooperation. Of course, it is also incumbent on researchers and industry to find mutually beneficial projects on which to collaborate.
- Instrumentation across all layers of the wireless stack, from spectrum to PHY to application layer, and we need to identify how to combine measurements across layers to address problems in today’s wireless networks and to inform future network designs.
- Address data privacy/security issues in the mobile environment, particularly due to the increased risk for leaks of subjects’ personally identifiable information (PII). We should develop community standards for gathering, securing and sharing such data, and ensure that these policies are compliant with jurisdictional restrictions. We should develop ways to “reward” those who comply with community standards for sharing data to encourage the practice.
- We should improve IRBs, e.g., creating a technology-focused IRB, and look into how social scientists (e.g., Census Bureau) deal with similar issues we are facing in the wireless measurement domain. Another

participant suggested that we should establish best practices and policies for data sharing.

- More research and policy effort should be devoted to making testbeds and datasets truly open and accessible. It is important to work closely with industry to not duplicate their efforts and instead focus on challenges that lead to addressing more forward-looking aspects of the network design across different layers for next-generation mobile networks.
- Better communication between research groups, industry, policy makers and pertinent government agencies. Many participants were unaware of salient testbeds, measurement approaches and government initiatives presented at the meeting. We suggest future workshops that focus on bringing together more focused groups of participants who share stronger interests, and use these meetings to build a community around critical mobile measurement infrastructure.

Measurement testbeds are traditionally second-class citizens in the research community – they help us produce interesting findings for publication, but take a back seat to the experiments that run over them. Based on our workshop discussion, it is clear that ad-hoc approaches to measurement infrastructure in the wireless environment will not allow us to advance the state of the art. We hope that the outcome of our discussions will help guide a new generation of testbed designs and deployments that drive experiments to inform future networked systems.

ACKNOWLEDGMENTS

This workshop was funded by the NSF (CNS-1455719). Aruna Balasubramanian, Mike Wittie, Chunyi Peng and Aaron Schulman contributed valuable scribes of the proceedings. We thank the NSF and participants for their contributions to this workshop. And in particular, we thank Thyagarajan Nandagopal for initiating the idea for the workshop and for his valuable input. ■

REFERENCES

- [1] k.c. claffy, D.D. Clark, and M. Wittie. The 6th Workshop on Active Internet Measurements (AIMS-6) Report. *SIGCOMM Comput. Commun. Rev.*, 44(5):39–44, Oct. 2014.