

The added value of ads - Measuring the rest of the Internet

Esteban Carisimo (Universidad de Buenos Aires) and Fabián E. Bustamante (Northwestern University)

As the primary source of revenue for web content and services, Internet advertising is pervasive across networks, platforms and geographic locations. We argue that this same pervasiveness suggests ads may have another valuable role to play in the Internet - as the basis of a measurement infrastructure that can achieve universal network coverage.

While much of what we know of the Internet we have learned through measurements, the fastest growing regions of today's Internet -- from the Middle East, to Africa and Latin America -- are also the least covered by our measurement platforms. Capturing a network perspective from those barely covered regions poses a number of key challenges, from deployment incentives to logistic difficulties. First, while users lack clear incentives to share their limited bandwidth with crowdsourcing measurement projects, operators seem to find little value in publicly exposing limitations on their networks. Second, distributing hardware probes in some of these areas is at best challenging and sometimes outright impossible given countries' customs barriers for electronic equipment. Last, even if we could address these incentive and logistic challenges, the dominance of mobile over wired networks for Internet access in these regions may mean we ended up, once again, measuring what we can rather than what we want to measure.

We believe ads could be the answer to these challenges, offering a potential vector for deploying measurements at the edge of the network. Ads are pervasive in the Internet across regions, networks and platforms -- from the Web, to social networks and mobile apps. In addition, the use of an ads-based measurement framework could offer new forms of incentives, such as additional context for value to ad analytics or a form of micro-revenue for users. Furthermore, ads already incorporate functionalities that can allow us to measure users network status and performance (e.g., boomerang-based measurements¹, the Network Information API²), and new features, such as the use of micro-targeting audiences of ad campaigns, that could provide useful control knobs for Internet measurement.

Building a measurement platform based on ads presents us with a number of interesting problems, from growing concerns about user privacy to the "right" way to incorporate measurement functionality into the already complex ad ecosystem. Would the coarse granularity acceptable for Internet measurement be sufficient to address privacy concerns? Would ad-blocking interfere with measurements or could we expect them to be whitelisted given that Internet measurements typically avoid PII? Should Internet measurement be part of existing ads, as a sort of added feature ad providers choose to include, or should we have special "ads" alongside shoes and home financing ads, that content providers include? How would the different ways in which users use the various platforms, like social networks and the web, impact the measurements we are able to collect?

¹ <https://developer.akamai.com/tools/boomerang/docs/>

² https://developer.mozilla.org/en-US/docs/Web/API/Network_Information_API