Research Statement

Mario A. Sánchez
Department of Electrical Engineering and Computer Science
Northwestern University
msanchez@eecs.northwestern.edu

My research interests focus on experimental systems deployment, performance analysis, network protocols and distributed Internet services such as DNS and content distribution networks (CDNs). I am interested in the application and implementation of low-level systems to analyze the performance of underlying protocols and their properties. I am currently working on several large research projects that entail evaluating the performance of layered network services from the perspective of individual peers in P2P systems located at the edge of the network.

Detecting ISP Interference from the Edge: Network traffic management has traditionally allowed ISPs to control infrastructure load and operational costs. Recent reports showing that such management interferes with P2P traffic, however, have sparked a heated debate regarding acceptable management policies. Beyond some highly publicized instances of interference and the increasing availability of diagnostic tools to detect them, it is yet unknown how common and widespread these policies are. This project focuses on how to identify and quantify ISP interference exclusively through natural traffic generated by thousands of real P2P users worldwide. While this approach prevents ISPs from masking their behavior, it introduces several challenges for isolating the ISP (as opposed to the application) as the likely cause for performance changes. Through the use of multiple passively gathered performance signals available from the application layer (i.e., without root privileges) we are able to confidently identify ISPs currently interfering with P2P traffic.

DNS Performance and its impact on CDN services: The Domain Name Systems originally envisioned purpose of delegating the task of Internet name resolution among different distributed systems has expanded

dramatically. Many new services have emerged that leverage the existing ubiquitous deployment of the DNS hierarchy to offer other functionality; one such service is offered by Content Distribution Network (CDN) systems. CDNs attempt to reduce latency in the delivery of content (improving overall performance) by relying on DNS name resolution to evaluate the clients proximity to the closest content replica server that will serve the request.

The emergence of such services has stressed the importance of DNS robustness and has redefined the term DNS reliability as well as the impact of configuration errors. The impact on a clients perceived quality-of-service from a timed-out request directed to a primary DNS server is no longer limited to an increased resolution delay while the request is redirected to a secondary DNS server (possibly located in a different network or even country). It also translates into an increased overall delay as the client is directed to a content replica sever that is as far away from it as the responding DNS server.

Previous DNS studies focus on lookup measurements, success rate, consistency and overall response time but due to the inherent difficulty of accessing a representative set of end nodes, analyzing the system from the edge of the network has proven to be a challenge. Taking advantage of a deployment of measurement software located on thousands of peers in a large-scale P2P network we are performing an extensive analysis on the reliability of the DNS hierarchy, evaluating the proximity of clients to their configured DNS servers as well as analyzing the effect of DNS configuration errors to ultimately weight-in the impact on these overlay networks.