February 1, 2021

Reading Summary: MobilityFirst

Alexander McRae, CSC 466

Paper: *Arun Venkataramani, James F. Kurose, Dipankar Raychaudhuri, Kiran Nagaraja, Morley Mao, and Suman Banerjee. 2014. MobilityFirst: a mobility-centric and trustworthy internet architecture. SIGCOMM Comput. Commun. Rev. 44, 3 (July 2014), 74–80. DOI:https://doi.org/10.1145/2656877.2656888*

The paper MobilityFirst identifies two missing aspects of the current internet stack which are becoming more and more important for modern workloads. First that TCP/IP fails to maintain connections when a peer switches networks. Second that IP conflates both identification and routing as the same thing. These flaws are problematic as modern workloads consist of long-running connections like VoIP and streaming primarily done on mobile devices such as phones and laptops.

The main idea of the paper is to split IP into the identification, a self-signing GUID, and routing information, a network address (NA). A global name service (GNS) would be used to map GUIDs to NAs. This GNS would also map human-readable names to GUIDs. Routing would be similar to current routing where the intranetwork routing it up to the network itself and the internetwork routing was standardized. A NA is effectively an autonomous system (AS).

I am convinced that splitting IP into GUIDs and NAs is a correct abstraction for computer networks in general. Its ability to do multicast and multi-home with ease in the author's examples shows how flexible and extensible it is for modern and future workloads.

However, I am not sure it solves the initial problem of mobility at scale. When peers with a connection move it signals to the other peer by updating its GUID mapping in the GNS. As the GNS service scales larger and larger I am not convinced by the paper that it can handle the excessive writes when updating. I believe that in the general case peers can tell if they are likely to move between networks and using this we could come up with more optimized cases for mobility. For example, if both peers are likely to move a rendezvous approach is taken to maintaining the connection. In every other case, the self-signing GUID allows us to verify the sender is who they say they are without needing to call back to the GNS. Although the GNS still needs the new mapping for the GUID it is not essential to maintaining a connection. This reduces the load on the GNS to only be read from on new connections.

Another aspect of the paper I felt was overlooked was the complexity of the GNS. With the added complexity of flat mapping and generally not battle hardened implementations it's likely a lot of the security benefits mentioned in the paper could be undermined by bugs in the implementations. I think that the GNS server should limit its scope more and not map human-readable names to GUIDs.

Lastly, the paper's implementation of the network infrastructure was impressive as a demonstration that the network could function within the current internet. I felt that the thing it lacked was any mention of how it could better integrate with current network technologies and how to move forward with greater adoption of the GNS. One could argue that IP is only routing information that the layers above it, UDP and TCP, have just failed to implement identification.

Overall MobilityFirst solves a good portion of the current internet’s issues and I believe it should be tested at scale as a possible alternative to the current internet, however, given IP’s broad usage I think a path to integrating with IP is more likely.