

Towards a Quantitative Comparison of Location-Independent Network Architectures

Zhaoyu Gao Arun Venkataramani James Kurose Simon Heimlicher
{gaozy, arun, kurose, heimlicher}@cs.umass.edu
School of Computer Science, University of Massachusetts Amherst

ABSTRACT

This paper presents a quantitative methodology and results comparing different approaches for *location-independent* communication. Our approach is empirical and is based on real Internet topologies, routing tables from real routers, and a measured workload of the mobility of devices and content across network addresses today. We measure the extent of network mobility exhibited by mobile devices with a home-brewed Android app deployed on hundreds of smartphones, and measure the network mobility of Internet content from distributed vantage points. We combine this measured data with our quantitative methodology to analyze the different cost-benefit tradeoffs struck by location-independent network architectures with respect to routing update cost, path stretch, and forwarding table size. We find that more than 20% of users change over 10 IP addresses a day, suggesting that mobility is the norm rather than the exception, so intrinsic and efficient network support for mobility is critical. We also find that with purely name-based routing approaches, each event involving the mobility of a device or popular content may result in an update at up to 14% of Internet routers; but, the fraction of impacted routers is much smaller for the long tail of unpopular content. These results suggest that recent proposals for *pure* name-based networking may be suitable for highly aggregateable content that moves infrequently but may need to be augmented with addressing-assisted approaches to handle device mobility.

Categories and Subject Descriptors

C.4.3 [Computer Systems Organization]: COMPUTER-COMMUNICATION NETWORKS—*Network Architecture and Design*

Keywords

Location-independence; mobility; network architecture

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
SIGCOMM'14, August 17–22, 2014, Chicago, IL, USA.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

ACM 978-1-4503-2836-4/14/08 ...\$15.00.

<http://dx.doi.org/10.1145/2619239.2626333>.

1. INTRODUCTION

Providing an abstraction of *location-independent communication*—enabling communication using fixed names without concern for changing network locations—has been a long-time goal of networking research. For example, proposals designed to provide seamless host mobility [36] target an abstraction of the form `connect(host_id)`; a number of proposals for information-centric networking [28, 13, 22, 16, 19] target an abstraction of the form `get(content_name)`.

Why is the current TCP/IP Internet seen as falling short of this goal? A common criticism is the so-called location-identity conflation problem [42]. The Internet uses an IP address to identify an interface as well as its network location. As a result, connections break when an endpoint changes network addresses, requiring application-layer workarounds to provide a semblance of seamless mobility. Advocates of information-centric networking argue that the Internet, having inherited a century-old, tethered, device-to-device communication abstraction from the wired telephony world, is poorly-suited for an Internet dominated by content traffic and a communication abstraction requiring endpoints to first obtain the network location of a host serving the requested content instead of simply procuring a copy from *any* convenient location. In response, researchers have proposed a number of designs to refactor naming, addressing, and routing in order to enable location-independent communication [22, 12, 36, 11, 28, 13, 22, 16, 19].

Our work is motivated by the observation that, although many architectural proposals share location-independence as a key design goal, there has been little prior research quantitatively comparing the different cost-benefit tradeoffs struck by these architectures in achieving that goal. One reason for the paucity of cross-architectural comparisons may be that network architecture is considered by some as part science and in good part *art* [51]. Another is that until recently, most Internet architecture efforts rarely went beyond paper designs, so a lack of a reasonably complete design specification and protocol-level implementation made it hard to justify investing research effort comparing them; however programs such as GENI[4], FIA[6], and FIRE [3] are changing this state of affairs. We believe that a quantitative comparison of different approaches for location-independent communication is timely and indeed critical for gaining a deeper understanding of cross-cutting architectural principles.

Our contribution is a quantitative methodology and empirical results comparing different location-independent architectures based on a common set of metrics, namely, routing update cost, path stretch, and forwarding table size, in