Modeling and Analysis of Human Encounters

Andrew G. Miklas† (student), Kelvin K. W. Chan∗, Stefan Saroiu†
†University of Toronto and ∗Google, Inc.

1 Motivation

A decade after its birth, the Internet continues to change. Recent news articles are reporting that the mainstream Internet user is shifting away from powerful desktop machines with always-on broadband Internet connections towards battery-powered, mobile, lightweight, cell-phone-like end-systems with sparse Internet connectivity [3, 6]. The number of new applications for the "new Internet" is quickly growing. For example, a startup is offering an Internet dating service allowing their clients to use their cell-phones' Bluetooth radios to detect when they are in the proximity of a person that matches their interests [4]. A different startup offers to re-tool billboards so that they can deliver short movie clips and ads via Bluetooth to a passerby [2]. A recent research project is bringing Internet connectivity to rural India through the use of buses and cars [7].

An important "trait" that distinguishes these new end-hosts from typical Internet hosts is that they are mobile, since they are carried by their owners in their day-to-day activities. These new systems' workloads are therefore influenced by people's mobility patterns. For example, the workload of the previously mentioned Internet dating service [4] is shaped by how their clients move, travel, and interact.

Our work presents an in-depth analysis and models of human mobility patterns. Our goals are:

1. To understand people's mobility patterns so we can anticipate the mobility patterns of tomorrow's Internet hosts.

2. To demonstrate that significant opportunity exists to optimize performance in emerging Internet applications by exploiting the Internet users' mobility characteristics.

2 Understanding People's Encounters Rather Than Their Mobility Patterns

While recent projects [1, 9] have proposed models that capture people's mobility patterns, validating these models is challenging because little experimental data is available. Gathering a large-scale trace of mobility data has substantial logistical and cost issues. Similarly, continuously monitoring people's locations raises numerous privacy concerns.

Instead, we propose studying a simpler problem: analyzing people's encounters, as opposed to analyzing their mobility patterns. A person's encounters describe *when* and *who* they meet without specifying *where* the encounter occurred. While losing geographical information, studying encounters helps us to understand how information propagates in the context of new, mobile Internet systems. For example, encounters drive the workloads of the previously mentioned Bluetooth dating application [4].

3 Analyzing and Modeling Encounters

This poster presents several findings from our analysis of a large-scale trace of encounters collected by the "Reality Mining" project at the MIT Media Lab [5]. This trace instrumented 100 people with Bluetooth-enabled cell-phones. More than 20,000 distinct Bluetooth devices were contacted over one-and-a-half years.

Our analysis is centered on answering three questions related to the temporal and social aspects of human encounters: (1) do people behave consistently over time? (2) is everyone equally popular? and (3) do people mostly meet friends (i.e., people encountered over and over again) or strangers (i.e., people encountered on fewer than ten days) ?

Temporal Analysis: People's encounters are driven by diurnal and weekly cycles. Most people's encounters occur on afternoons during week-days with a peak at 4:00pm. There are 65% more encounters on afternoons (2-5pm) compared with mornings (9am-12pm) and 168% more encounters occur during week-days than during week-end days.

Once we account for time-of-day and week-end/week-day effects, the average person's inter-encounter times are exponentially distributed. While the average person encounters other people at a rate that varies according to diurnal and weekly patterns, this rate remains constant during the same hour on the same day of the week. This implies that the new Internet systems' average load per user is predictable.

Social Analysis: Global popularity, or the number of people encountered by an individual, follows a long-tailed distribution: while a few people are very popular, most people are quite unpopular. This implies that the new Internet systems' workloads will be highly unbalanced and these systems must accommodate for "hotspots" created by "popular" users.

Finally, we find that while people spend most of their encounter-time with "friends", most of the people they encounter are "strangers". This implies that caching is likely to be effective for these new Internet systems. For example, because people spend most of their time in the proximity of their friends, a local cache on a Bluetooth dating user's cell-phone will have a high hit rate.

Modeling Encounters: We have constructed a preliminary model that incorporates the patterns we observed. Based on our model, we have analyzed the spread of computer worms in the context of new, mobile Internet systems [8]. In future work, we plan to use our model to characterize these new systems' workloads.

[1] J.-Y. L. Boudec and M. Vojnović. Perfect simulation and stationarity of a class of mobility models. In *Infocom*, 2005.

[2] Filter UK Ltd. Bluecasting: The proximity marketing system, 2006. `http://www.bluecasting.com`.

[3] InfoWorld. More mobile Internet users than wired in Japan, July, 2006. `http://www.infoworld.com/article/06/07/05/HNjapannetusers_1.html`.

[4] Kangourouge. Proxidating, the first ever Bluetooth dating software for mobile phones, 2006. `http://www.proxidating.com`.

[5] MIT Media Lab. Reality Mining. `http://reality.media.mit.edu/`.

[6] C. News.com. Mobile browsing becomes mainstream, April, 2006. `http://news.com.com/Mobile+browsing+becoming+mainstream/2100-1039_3-6062365.html`.

[7] A. Seth, D. Kroeker, M. Zaharia, S. Guo, and S. Keshav. Low-cost communication for rural Internet kiosks using mechanical backhaul. In *MobiCom*, 2006.

[8] J. Su, K. K. W. Chan, A. G. Miklas, K. Po, A. Akhavan, S. Saroiu, E. de Lara, and A. Goel. A preliminary investigation of worm infections in a Bluetooth environment. In *Proceedings of WORM*, 2006.

[9] J. Yoon, B. Noble, M. Liu, and M. Kim. Building realistic mobility models from coarse-grain traces. In *Mobisys*, 2006.