Research Description by Daniele Quercia

During my PhD at University College London, I have been sponsored by Microsoft Research Cambridge to study new ways for co-located mobile users to share digital content (e.g., pictures, news). In doing so, I have published in top-tier conferences in the areas of:

- Ubiquitous Computing (Ubicomp: acceptance rate 19%)
- Data Mining (ICDM: 7%)
- Mobile Systems (MobiQuitous: 22%)
- Security (SECURECOMM: 19%)

I have also been interested in commercializing new technologies while being: (1) a student of MBA "technology electives" at London Business School (consistently ranked as one of the top business schools in the world); and (2) a member of the "UCL Centre for Security and Crime Science", which offers holistic security solutions to organisations from industry, commerce and government.

Filtering Mobile Content using Trust Models

Problem: Mobile users can create digital content (e.g., photos, videos) and, using Bluetooth, they can share it with other users in proximity. However, what happens if everybody is distributing content? In that case, to paraphrase Italo Calvino, mobile users would live in an unending rainfall of content.

Existing Research: One way of solving this problem is to select reputable sources. Mobile users may do so by running trust models on their portable devices. A trust model is a piece of software that keeps track of which devices are trusted (for sending quality content) and which are not. Portable devices cannot use existing trust models because they either are centralized or have been largely designed for peer-to-peer networks and not for mobile networks.

My Research: I have designed algorithms with which portable devices automatically keep track of reputable sources and learn from each other's recommendations. More specifically, I have designed and evaluated:

- **A new set of protocols with which portable devices collect and store recommendations.** To learn from each other's recommendations, portable devices need to collect and store recommendations first. But mobile networks suffer from content providers who tweak recommendations to their own advantage. That is why, at Ubicomp'08, I will present a new decentralized mechanism (dubbed MobiRate) with which portable devices store recommendations in (local) tamper-evident tables and check the integrity of those tables through a gossiping protocol[1]. Using real mobility and social network data, I have showed that MobiRate considerably reduces the impact of "tweaked" recommendations. I have also showed that MobiRate runs on mobile phones at little computational and communication costs.
- **A new distributed algorithm that builds trust on recommendations.** At ICDM'07, I presented a way with which devices set their trust from third-party recommendations[2]. Each device organizes ratings of content producers that it knows and trusts in a graph (called "Web of Trust"). It then uses a graph-based learning technique to form opinions about content producers with whom it has never interacted before. This algorithm shows high predictive accuracy on a large real-world dataset, and, in contrast to existing approaches, it is fully decentralized. Plus, its JAVA implementation runs reasonably fast on mobile phones.
- **A new algorithm that builds trust across (content) categories.** Say that device *A* received financial news from *B* and found them interesting. *B* now produces economic news. From its past experience, can *A* conclude that *B*'s economic news are also of good quality? *A* may well conclude so since "economics" and "finance" are (semantically) similar. In general, to decide whether two categories are similar, existing algorithms typically use an ontology (e.g., a taxonomy of content categories). However, in so doing, they require that the same ontology is shared by all users (which is hardly the case in reality) and that those users agree on that ontology for good (i.e., the ontology is not supposed to change over time). That is why I presented at *MobiQuitous'07* an algorithm (called TRULLO) that does away with these two problems[3]. The idea is that devices learn statistical properties from ratings and upon those properties they then find out which categories are similar. I have showed that TRULLO effectively determines category similarity using simulations based on eBay's data and that its implementation is reasonably fast on mobile phones.

Who cares: To understand the potential impact of this research, consider that these algorithms can be used:

- *By mobile users* to run applications that, for example, automatically find: *songs of emerging musicians* - who upload their latest tracks into publicly-available WiFi hotspots and add the date of their next gig as a note to the track for some free publicity; *prices of outlets* - that embed their latest offerings or discounts or seasonal menus within their clickable logos displayed on navigational maps; or impromptu *street performances* - advertised by disseminating electronic flyers.
- *By software developers* as ready-made tools when designing Web 2.0 mobile applications (whose market is forecasted at $22.4bn in 2013[4]).

---

[1] Quercia, Hailes, and Capra. *MobiRate: Making Mobile Raters Stick to their Word*. 10th ACM Ubicomp. 2008.
[2] Quercia, Hailes, and Capra. *Lightweight Distributed Trust Propagation*. 7th IEEE ICDM. 2007.
[3] Quercia, Hailes, and Capra. *TRULLO - local trust bootstrapping for ubiquitous devices.* 4th IEEE Mobiquitous. 2007.
[4] Share, Collaborate, Exploit - Defining Mobile Web 2.0. Juniper Research. 2008.