MaxPace: Speed-Constrained Location Queries

Per Hallgren, Martín Ochoa, and Andrei Sabelfeld

With the increasing proliferation of mobile devices, location-based services enjoy increasing popularity. At the same time, this raises concerns regarding location privacy, as seen in many publicized cases when user location is illegitimately tracked both by malicious users and by invasive service providers. This paper is focused on privacy for the location proximity problem, with the goal of revealing the proximity of a user without disclosing any other data about the user's location. A key challenge is attacks by multiple requests, when a malicious user requests proximity to a victim from multiple locations in order to position the user by trilateration.

To mitigate these concerns we develop MaxPace, a general policy framework to restrict proximity queries based on the speed of the requester. MaxPace boosts the privacy guarantees, which is demonstrated by comparative bounds on how the knowledge about the users' location changes over time.

The effect of speed constraints is illustrated in Figure 1. Each query corresponds to a disk. Large disk overlaps with speed-constrained queries means that the attacker learns little information from each query compared to the unrestricted attacker, and thus needs to issue more requests to learn the victims location.

MaxPace applies to both a centralized setting, where the server can enforce the policy on the actual locations, and a decentralized setting, dispensing with the trust to the service provider. In the centralized setting, we are encouraged by the recent changes in the policies of the popular centralized location-based services Facebook, Swarm, and Tinder to incorporate forms of speed-based constraints. Our study is intended to provide rigorous analysis and understanding of guarantees provided by this type of constraints.

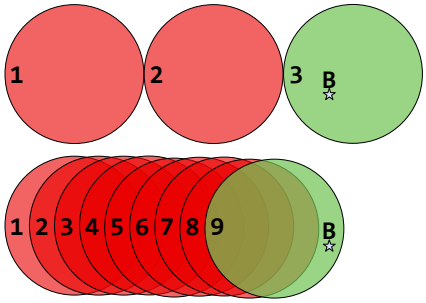

Fig. 1. Two different proximity disclosure protocols

In the decentralized setting, we develop a secure multi-party computation protocol. This offers a contribution beyond the state of the art. A dominant assumption in the most recent literature on securing location proximity is the assumption of a single run. In contrast, our approach does not impose such an assumption, allowing us to reason about multi-run attacks.