A Needle in the Haystack - Delay Based User Identification in Cellular Networks[*]

Marco V. Barbera, Simone Bronzini, Alessandro Mei, and Vasile C. Perta

Sapienza University, Rome, Italy
{barbera,bronzini,perta,mei}@di.uniroma1.it

Abstract. In this work, we discuss a technique for identifying users in cellular networks that exploits the effect that RRC state machine transitions have on the measured round-trip time of mobile devices. Our preliminary experiments performed in a controlled environment, show that it is possible to leverage popular real-time messaging apps, such as Facebook, WhatsApp and Viber, to trigger an observable delay pattern on a user's device, and use it to identify the device.

Keywords: Cellular Networks, Security, Privacy

1 Introduction

With respect to broadband fixed networks, cellular networks are very constrained in terms of both energy and radio resources available to each mobile device. To balance between efficiency and user experience, mobile devices (referred to as "user equipment" by the standard) are assigned radio resources depending on the volume of data they send or receive from the network. This process is regulated by means of transitions in a Radio Resource Control (RRC) state machine that is associated to each device. RRC states are typically `CELL_IDLE`, `CELL_FACH`, and `CELL_DCH`, corresponding to no, low, or full radio resources respectively. Promotions from lower to higher resource states are not immediate. Rather, they introduce an observable extra delay (*i.e.,* 1 or 2 seconds) to packets sent to a mobile device that has not recently used network resources (*e.g.,* is in the `CELL_IDLE` state). Because of this, round-trip times are sufficient to remotely characterise the RRC state machine used by devices in a target cellular network, as recently shown by Qian *et al.* [3]. In this work we verify whether round-trip time variations due to RRC state machine transitions, in conjunction with network activity triggered by mobile push notifications, may, in principle, allow to remotely identify the IP address of users of popular mobile messaging apps. This could represent a potential threat to mobile users, as it would permit an adversary to perform focused attacks on a specific set of devices, such as the stealth-spam-attack discussed by Peng *et al.* [2]. More in general, this is another example of attack exploiting the unique characteristics of mobile networks and devices [1,5,4].

[*] This work has been partially supported by a Google Faculty Research Grant 2013.

2 Identifying mobile devices from RTT variations

In our model, an adversary produces some network traffic on the target user device (*e.g.,* using a sequence of instant messages), which triggers RRC state transitions on the device and induce some observable pattern on the round-trip times towards it. At the same time, the adversary looks for similar delay patterns towards all, or a subset of, the devices of the mobile network operator. This results in a set of candidate IPs `IP_range` that is reduced in size by iteratively applying the same procedure multiple times. To produce traffic on the user's device, we propose the use of real-time messaging apps such as Google Talk, or near real-time apps such as Facebook Messenger, Viber, and Whatsapp. In fact, to improve the detection accuracy, messages to the user's device should be delivered within a short time after they have been sent, assuming the user is online. Note that it is not necessary for the adversary to be socially very close to the target user. For instance, both WhatsApp and Viber allow messages to be sent to any user, given her mobile phone number. To measure round-trip times, the adversary has to be able to directly reach the target user device from a vantage point. This is possible if the cellular network assigns public, reachable IP addresses to the devices, or, if device-to-device probing is allowed between devices with a private IP address. This has been recently estimated to be the case for around 50% of the cellular networks [4]. The initial `IP_range` can be set to the whole set of IPs of the cellular network carrier, if no extra information on the target user is known. If the user's coarse-grained location is known, to speed up the process, the set can instead be restricted by mapping IP addresses of mobile devices to a given geographical area using the method proposed by Qian *et al.* [4].

3 Evaluation

To test the effectiveness of our detection methodology, we used as a target device a Samsung Galaxy S Plus attached to a popular Italian cellular network. During th test, the device was left idle, under stable network conditions and signal strength. The device was assigned an IP in a /19 subnet, which we used as the initial `IP_range` set. The adversary, a Linux host attached to our university's network, continuously probed the round-trip times (RTTs) towards the `IP_range` set by means of low-rate ping packets sent every 10 seconds. The probing rate has to be chosen in such a way that ping packets alone cannot keep the devices' state-machines at a high-power state (*i.e.,* `CELL_DCH`). If a too-high rate is used, then all the devices `IP_range` would roughly show the same low-delay pattern, thus making detection impossible. To trigger periods of network traffic on the target's user device pattern we alternated a 3 minutes interval where no data is sent to the device, with a 2 minutes interval where the RRC machine state of the device is kept on `CELL_DCH` by means of, *e.g.,*, Facebook, WhatsApp, or Viber instant messages. To restrict `IP_range` to the set of possible IPs of the target device, the adversary looks for devices whose RTT suddenly drops by at least α milliseconds after the message sequence has been sent. These corresponded to the devices that, during the probing period, switched from a low-power state to a higher-power state (*e.g.,* `CELL_IDLE` → `CELL_DCH`). This is exemplified in Figure 1, where it can also be observed how network

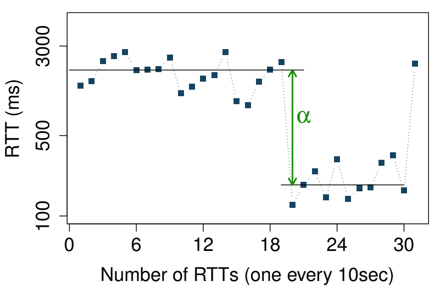

Fig. 1. Sudden drop in the RTT when the target device passes from `CELL_IDLE` to `CELL_DCH`.

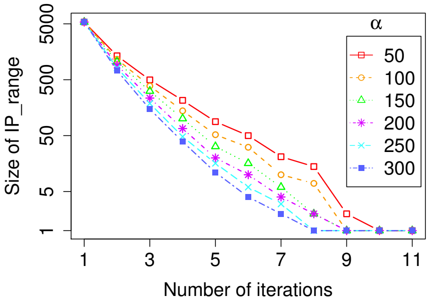

Fig. 2. Convergence of the algorithm with different values of α (in ms). Each iteration lasts 5 minutes.

delays are at least an order of magnitude lower than delays produced by RRC state transitions. In Figure 2 we show how the `IP_range` set shrinks by iteratively applying our detection methodology with different values of α. Interestingly, with just one iteration, the `IP_range` candidate set shrinks by around the 90% almost independently on the α parameter used. This is probably because many devices in the initial `IP_range` set are in `CELL_IDLE` state. Discarding these devices is very easy, as they always yield very high RTTs. After some iterations, the percentage of devices that is discarded each time slowly decreases, while the parameter α has a higher impact. In this case, increasing the threshold α helps taking into account only the strong delay variations given by actual RRC machine state transitions. Overall, we were able to correctly guess the IP of the target user's device in a few iterations (around 8 in the example). Detecting the user's device IP address with this level of accuracy may take more time when starting with a larger `IP_range` set. However, depending on the scenario, the number of iterations could be reduced by terminating the search when the `IP_range` set becomes smaller than a certain threshold. For instance, to save bandwidth resources in a spam attack [2], an adversary could be satisfied with just a small enough set of possible IPs of the target user's device.

References

1. Lee, P.P., Bu, T., Woo, T.: On the detection of signaling DoS attacks on 3G wireless networks. In: INFOCOM. IEEE (2007)
2. Peng, C., Li, C.y., Tu, G.h., Lu, S., Zhang, L.: Mobile data charging: new attacks and countermeasures. In: CCS. ACM (2012)
3. Qian, F., Wang, Z., Gerber, A., Mao, Z.M., Sen, S., Spatscheck, O.: Characterizing radio resource allocation for 3G networks. In: IMC. ACM (2010)
4. Qian, Z., Wang, Z., Xu, Q., Mao, Z.M., Zhang, M., Wang, Y.M.: You can run, but you cant hide: Exposing network location for targeted DoS attacks in cellular networks. In: NDSS (2012)
5. Traynor, P., Lin, M., Ongtang, M., Rao, V., Jaeger, T., McDaniel, P., La Porta, T.: On cellular botnets: measuring the impact of malicious devices on a cellular network core. In: CCS. ACM (2009)