CS 598 WSI - Lecture 12: RF Sensing and Privacy

Privacy Risks

What does a network know about you?

- While you're connected:
  - Device type
  - Rough location
  - Time of connection
  - Traffic patterns
- Over time: can build a profile with this data and answer higher-level questions.
  - When do you take your lunch break?
  - When do you come to work?
  - Etc.

Privacy Leakage

How do we discover networks?

- Passive discovery:
  - Listen for periodically-transmitted beacons from AP.
  - This is slow! You have to listen on all channels to discover the APs.
- Active discovery:
  - Your device sends out messages asking if your known networks are available.
  - Quicker since you don't have to wait for the AP to send you a message.
  - This leaks info - now surrounding APs know which networks you have saved
    - Maybe restaurants you visit, your office network, etc.
    - May give information about your habits
  - This makes it easier to fingerprint you!

Fingerprinting: can I uniquely identify a user based on information they share?

- Once you have a fingerprint, it's possible to identify users across networks.
  - For example, if different AP operators share data, they can track your habits and behavior on a larger scale.
  - Additionally, they can figure out your friends, etc. based on how much time you spend together.
- Gets more dangerous based on the granularity of information that it leaked.
  - Location level: which buildings do you go to?
  - Room level: when are you working, and when are you taking a break?
  - Step level: how are you feeling? How is your mental and physical health?

Mitigation & Limitations

MAC address randomization:

- Tracking often utilizes the MAC address, since it is typically unique for a device
- We can make things more difficult by using a random MAC.
- Two ways to do this:
  - Send each probe from a different MAC address
    - Tough implementation
  - Probe for all networks under the same MAC address, but switch MAC addresses periodically
    - Simpler implementation; this method is used in practice.
- There are limitations to this approach, especially the second variant
  - MAC may be different, but set of networks is the same
  - If you have timing information, you can link the MAC addresses together.
  - Timing info can be used for fingerprinting, since different devices may have different intervals and patterns.

Privacy & Sensing

Passive sensing: tracking people without snooping on their devices

- For example, using reflections to track people through walls
- Is someone home? Is a room empty? Etc.
- More granular information possible too - heart rate, breathing patterns, etc.

Location privacy: AP wants to sense the user's location. Can we avoid this?

- AP uses RSSI:
  - User can vary transmission power.
    - Data rate may be reduced due to rate adaptation
    - Inconsistent data rate may cause issues for higher-layer protocols
- AP uses ToF:
  - If AP uses user-provided timestamps, user can falsify its timestamps
  - If AP only uses its own timestamps, user can increase delay before ACKs
    - WiFi dictates 10 microseconds between reception and ACK, so if user sends ACK, the AP can estimate ToF on its own
    - Hard to change, since the delay is baked into the hardware
- AP uses angle-based tracking w/ multiple antennas:
  - Very difficult to defend against.
  - Angle-based attacks also can break other types of privacy tools.
    - e.g. different MAC, same exact location is probably the same user
  - User can break channel estimation, and therefore also angle estimation
    - Sadly, estimating the wrong channel breaks communication!
    - AP measures channel h_i from received signal y_i and known preamble x with the equation $h_i = \frac{y_i}{x}$.

        - User can trick the AP by sending x' instead of known preamble x.
    - Other approaches exist, but lead to antenna war
        - User can fool AP if they have a greater number of antennas
            - Impractical since usually APs have more antennas than clients.
        - General idea: prevent AP from estimating channel precisely.
            - User has n antennas, AP has m antennas, $n > m$.
            - User can tweak each of their n antenna's transmitted x.
            - At the AP, this creates an equation with n unknowns.
            - However, the AP only has m antennas, so it doesn't have enough equations to solve the system.
        - Attackers may be tricky and not reveal their exact number of antennas!

RF-Protect

Goal: protect against passive eavesdroppers.

- Your activities can be tracked passively due to reflections from your body.
- Nothing you can practically do to prevent your body from reflecting.

Jamming: may be impractical or illegal.

- Method 1: Block radio signals from entering your house
    - Inconvenient: will also block cell signals, etc.
- Method 2: Transmit high-power signal in the spectrum you want to jam
    - Could be illegal: you may be jamming licensed frequencies, or you may exceed legal power levels
    - Inconvenient: may also interfere with our devices

Idea behind RF-Protect: Introduce fake people into the environment.

- This may cause an attacker to make erroneous inferences.
    - e.g. a fake human at home may confuse an attacker who is waiting for the home to be vacant.
- Adding fake people into the environment makes an attacker's inferences noisy.
    - Important: you can't make the real human disappear!
- If you want, you can allow for legitimate tracking by disclosing the fake reflection.

What are our requirements?

- Can't use a static reflector, since people move, etc. all the time.
- We need to mimic human motion.
    - Can't just be random motion, or a static object - these are easy to identify and filter out.

Mimicking human motion:

- Spoofing distance: add small amounts of frequency shift
    - Causes the FMCW radar to misinterpret the distance

- Spoofing angle: switch antennas when reflecting the signal
    - Each antenna covers a sector of the angular space.
    - With appropriate antenna placement, the transition between sectors can appear smooth to the radar.
    - Switching between antennas can then create a varying angle measurement.

How to make convincing human-like trajectories?
- Random trajectories or static patterns are easy to filter out.
- Instead, RF-Protect uses a GAN.

Limitations and discussion:
- The RF-Protect device would need to be placed along all potential attack surfaces.
    - You need info about the attacking radar.
        - Where is the attack coming from?
        - What frequency bands are they operating on?
            - Although RF-protect works on multiple bands due to being a reflector, there is a limit to how wide the band can be.
    - This system requires some deployment effort.
    - However, let's say you have one device in particular you don't want spying on you.
    - You could create a reflector for this device that would direct its signals to a surface with RF-Protect installed.
- Vulnerable to side-channel and contextual information.
    - If you know how many people are in the house, you can tell that RF-Protect is running
    - If you know the room layout and are aware of where people can and cannot go, you may be able to identify a path generated by RF-Protect

Future work:
- Making the system distributed, i.e. having small antennas everywhere
- Addressing the limitations discussed earlier