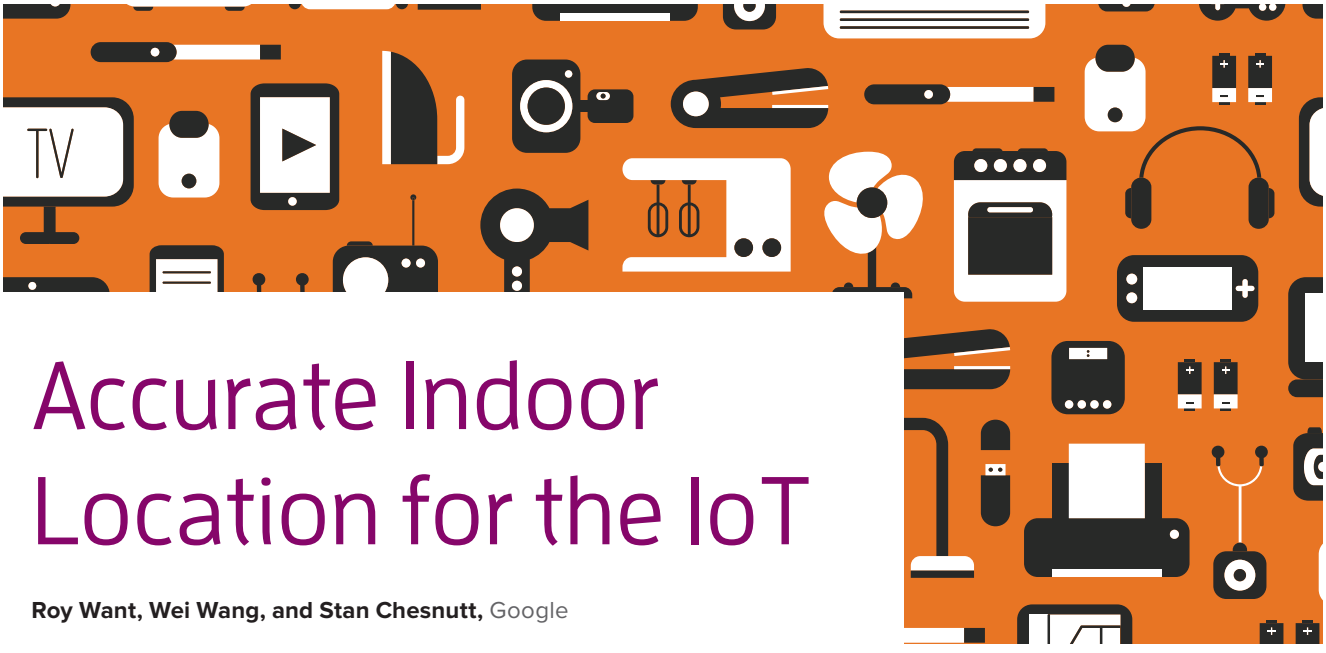


Accurate Indoor Location for the IoT

Roy Want, Wei Wang, and Stan Chesnutt, Google

The combination of IoT and accurate location-finding technologies will enable a wide range of novel context-aware services to automate tasks and support everyday work practice.

A defining characteristic of the Internet of Things (IoT) is that devices are both networked and embedded in the environment. Unlike traditional networked server technology that has very little connection with the environment, IoT devices are designed and positioned so that they can respond to and influence their surroundings. For example, if a user accesses a web server from an office in California, it makes very little difference whether the service is located in the next building or situated in Europe. However, if a user accesses an IoT device that is responsible for measuring the temperature of a room, the sensor and device need to be located in that specific room.

As a result, IoT design is a unique departure from infrastructural networking that tries to abstract away the network topology; instead, the topology and position of IoT nodes are assets in performing the networking task.

IoT devices often contain sensors and actuators. Consumer products for IoT sensors include light sensors, microphones, cameras, security switches, motion detectors, thermostats, barometers, moisture sensors, and when extended to industry, a wide variety of exotic sensors. Actuator outputs include lamps, valves (water sprinklers), speakers, and a variety of displays—again, industry has many additional needs for bespoke actuators. The interrelationship of the position and function of each of these devices allows us to design systems that are often called "smart homes" or "smart buildings."

An additional consideration of spatially arranged infrastructure is that mobile wireless components can move through it. In a smart home, people will be present and the infrastructure is intended to make their lives easier and more comfortable. Devices such as lights, HVAC, and security systems can be turned on or off by the presence and detection of people entering or leaving an area. This is not only convenient, but also addresses human forgetfulness and has cost-saving implications because it reduces energy consumption.

In the modern smart home, voice-driven assistants have become popular (such as Amazon Echo and Google Assistant) and allow devices in smart homes to be activated by casual voice commands. Typically, human speech is often tersely phrased with the assumption that the context is understood. This is challenging for a computer, which doesn't have the shared context to respond

0018-9162/18/$33.00 © 2018 IEEE

EDITOR **ROY WANT**
Google; roywant@google.com



to these commands. One of the biggest qualifiers of context is a user's position. For example, an ambiguous command such as "turn on this lamp" can be immediately qualified if the system knows the relative position of the user and all the nearby lamps.

Although it might be hard to determine the location of a person, it is now possible to accurately calculate the position of a mobile device using Wi-Fi radios (see www.ieee802.org/11) with positioning support. Since the early Wi-Fi (802.11b) products appeared in 1999, Received Signal Strength Indication (RSSI) has been used to determine the approximate position of a mobile device relative to surrounding access points installed at known positions.[1] More recently, Bluetooth low-energy beacons have been used in a similar way, but at higher density due to their small size and low cost. However, due to multipath fading, the RSSI fall-off from the antenna does not follow a uniform inverse square law and is distorted by the building materials that it travels through. As a result, the position accuracy for Wi-Fi can often have large errors of 10 to 20 meters. Improvements have been made by measuring the RSSI fingerprints for all detectable access points at each square meter of a building, but this requires considerable investment in surveys, which are labor intensive and have to be regularly updated.

A more accurate approach was needed that could localize a device to about 1 square meter (the area occupied by a person), but only required the knowledge of Wi-Fi access point positions.

THE IMPORTANCE OF A STANDARD

In December 2016, the IEEE 802.11 working group ratified amendment 802.11-REVmc[2] for the Wi-Fi standard, which included a new ranging capability based on time of flight (ToF) of Wi-Fi packets. As these packets are transmitted as RF signals that travel at the speed of light *c*, the distance can be calculated by knowing the ToF of a packet from a mobile to an access point and multiplying by the speed *c*. Knowing the distance or range to three or more access points (and using a process called multilateration), the mobile's position can be calculated (see Figure 1).

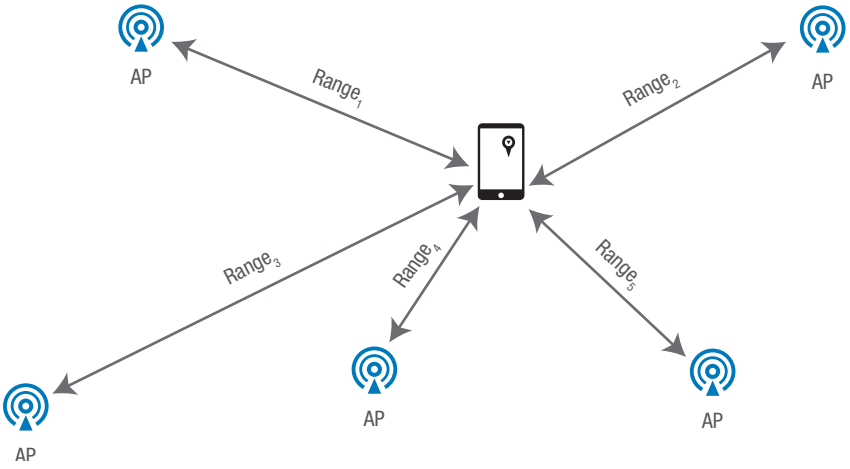


Figure 1. Multilateration, the process of converting ranges to a position. AP: access point.

In recent times, there have been other radio ranging standards (for example, ultrawideband [UWB], which was standardized in IEEE 802.15.4a[3]), but none have had much impact on the mobile industry. Wi-Fi is different because of its wide-scale use in mobile phones, laptops, and tablets throughout the world. And most importantly, Wi-Fi access points' infrastructure has grown as a grassroots effort and is deployed ubiquitously, including in many public spaces.

The spread of Wi-Fi has been enhanced because of three value propositions:

- it has a practical range of about 25m, and as a result the density and related cost of access points to cover a work area are not great;
- its power consumption is low enough that it can be embedded into small mobile devices with acceptable impact on battery life (for smartphones and tablets); and
- the data rate of Wi-Fi products has remained competitive and has typically had a x10 advantage over cellular data, which often has a data limit and results in additional cost if exceeded.

Wi-Fi has an additional characteristic in that it has already been improved several times to meet growing user demands. There is a natural upgrade cycle for Wi-Fi as radio technology improves and as users crave higher data rates to support video and faster downloads. These upgrades also serve as a vector for new features (such as ranging) that can be introduced in each upgrade cycle.

In short, Wi-Fi 802.11-REVmc is ideally suited as a standard to disseminate a new positioning technology with the potential for a wide reach and the opportunity for mobile application developers to build on its novel capabilities.

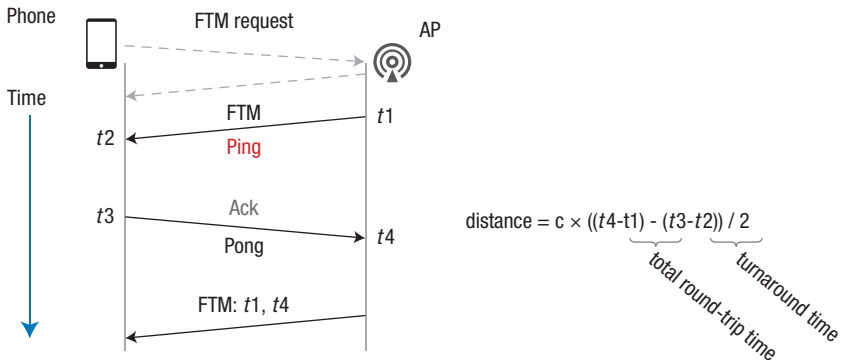


Figure 2. The 802.11mc Fine Timing Measurement (FTM) Protocol.

IEEE 802.11-REVMC AND WI-FI ROUND-TRIP TIME

The core of the 802.11-REVmc protocol is built around a new packet type called a Fine Timing Measurement (FTM) frame. The recipient responds to an FTM with an acknowledgement packet, and the transmission and reception times of each are recorded. It's basically a ping-pong protocol that calculates ToF by measuring the time it takes for a packet to be sent from an access point to a mobile and then back again. This is called the round-trip time (RTT), so the protocol is sometimes referred to as Wi-Fi RTT.

For a mobile device to measure the ToF, the precise times of the packet leaving the access point ($t1$) and arriving at the mobile receiver ($t2$) must be known. Further, the precise times the mobile device transmits the response ($t3$) and the time it is received back at the access point ($t4$) must also be recorded. From the mobile's perspective, it measures only $t2$ and $t3$ as they occur locally at its radio. However, $t1$ and $t4$ are measured at the access point, so for the mobile to calculate the RTT, the access point needs to send $t1$ and $t4$ as an additional message to the mobile (see Figure 2), which is piggybacked on the next FTM packet.

The phone can now calculate the RTT based on the total time $t4 - t1$, subtracting $t3 - t2$ to account for the turnaround time of the acknowledgement packet.

$$RTT = (t4 - t1) - (t3 - t2)$$

The desired range is determined by multiplying by the speed of light c and dividing by 2, as we are only interested in the distance from mobile to access point and not the round-trip distance.

$$range = c \times RTT / 2$$

Some general assumptions are made that the clocks are running at about the same rate on the two devices, and with a modern crystal oscillator with an accuracy of 40 parts per million, this is reasonable. Note that the clocks of each device don't need to be synchronized as the $(t4 - t1) - (t3 - t2)$ terms provide time differences for each separate clock at the beginning and ending of each exchange, and therefore do not require an absolute time reference. Given the typical clock stability and an 80 MHz bandwidth in the 5 GHz Wi-Fi band, the expected accuracy of this protocol is expected to be ~1m, which meets the design goals and considerably outperforms approaches based on RSSI.

INCREASED LOCATION ACCURACY ENABLES NOVEL APPLICATIONS

Important questions for the design of any positioning system are "What is the position system going to be used for?" and "What target accuracy is needed?" If a system has an accuracy of 10m and is improved to 9m or even 8m, does it make the system more useful? Perhaps not that much. However, if you were to improve it to 1m, it would cross a significant human threshold because a person typically occupies about 1 square meter of standing space. It opens up the possibility to place a person next to another person, product, or thing.

For Wi-Fi RTT, 1 to 2m is just less than the width of most corridors or aisles found in a building. This resolution actually enables in-building navigation, for example, directing a person to a product in a store or an office worker to an unfamiliar conference room. Errors two to three times this size would cause the indoor navigation system to place you in the wrong aisle, resulting in a confusing user experience. In the same way that GPS is accurate enough (~5m) to place a car on the correct road, in-building navigation is enabled by a 1 to 2m Wi-Fi location system. In combination with indoor maps and a routing system calculating the shortest path, Wi-Fi RTT can show your progress along a route and provide corrections by rerouting when you deviate from the path.

In the IoT world, ~1 meter accuracy can automatically determine the location of devices such as lights and thermostats, which might not have their exact position recorded. Positioning to ~1m can automate the documentation for an installation process, reducing contractor time and removing transcription errors that inevitably occur in hectic construction environments. Further, it can automatically keep records up to date as devices are replaced and during remodeling.

In daily work practices, users carrying mobile devices with Wi-Fi RTT can be located in a smart home (using an application with permission from the owner) and automate the operation of heating, lighting, door locks, and security systems. The modern smartphone is a multipurpose device, and owners typically carry the device with them most of the time. As home

monitoring and control capabilities grow via smartphone apps, we expect that users will make use of these apps in their everyday lives.

There are further advantages for accurate location in peer-to-peer interactions. When two devices are next to each other, there is the opportunity to share information. For example, if a user wishes to share a photo, it might be done with a simple flick gesture on a screen. A phone that can detect the proximity of nearby devices can limit the scope of gesture-based sharing to the closest device. You can imagine many other forms of digital peer-to-peer transfers that are made possible by automatically and accurately detecting close proximity, which includes ticketing and payments.

LOCATION DATA AND PRIVACY

We have already seen that accurate location can bring many advantages to daily work practice and social interactions. However, there are important privacy implications to consider, and great care needs to be taken when implementing and enabling these services for users.

Mobile devices that can repeatedly link a user's identity to the device's location create a location "trail," which could result from location data being stored persistently on the device or uploaded to a server. With nefarious access to this information over time, it's straightforward to determine the habitual whereabouts of the user, which could be used to tip off thieves or others with malicious intent.

Protecting the user's privacy starts with an agreement with the user. First, the software must request explicit permission to establish the user's location and to store that information. Second, persistent location information must be carefully protected after being uploaded to a central server (for example, protection against reuse in unexpected ways and against observation by a rogue employee). This need for protection also applies to raw ranging data, not just position, as it can also be converted to user location by a nefarious party.

Furthermore, an IoT device might have capabilities that can threaten user privacy when coupled with location data. For example, a security camera that uploads video data could also add identity and location information.

The European Union's General Protection Data Regulation (GDPR) establishes clear standards for data retention, storage, and use. These guidelines should be treated as a starting point for data stewardship.

THE FUTURE OF LOCATION TECHNOLOGIES

Going forward, we expect 802.11-REVmc to be adopted by a wide variety of client devices and access points, which will result in many novel location-based applications that can make use of ~1m indoor accuracy. The designers of the first Wi-Fi RTT protocol had many additional features that they wanted to include, but they had to be deferred due to the timing of the first release.

A more extensive set of location protocol features are currently being undertaken by the follow-up task group 802.11az (Next Generation Positioning), whose primary goals are to enable greater scalability and provide for secure operation.

The 802.11az task group is developing a protocol that takes advantage of the new high-efficiency (HE) data standard being developed in parallel in Task Group 802.11ax with four times the data throughput of existing 802.11ac-based systems. In addition to higher-density coding, the radio supports multi-user uplink MIMO (multiple input, multiple output), allowing more spatial parallelism in data streams than before and resulting in more efficient use of the communication channel. Location techniques can take advantage of this mechanism by ranging between an access point and several stations simultaneously, thus supporting a larger number of potential clients.

An additional technique that supports even larger populations of clients is called passive ranging. This mechanism is characterized by mobile devices that only need to listen to the transmissions made by other stations and never need to transmit themselves. The result is that the channel capacity is not affected by the number of users trying to measure their position—furthermore, privacy is guaranteed as a mobile never needs to announce its presence. In that sense it's similar to the operation of GPS, although it differs in the geometric calculation needed to determine a position.

An IoT device might have capabilities that can threaten user privacy when coupled with location data.

In this technique, special types of access points called anchor stations range to nearby access points. If a mobile device listens to the protocol exchanges between anchor stations and access points, it can use the reported timestamps to figure out the time difference of arrival of two packets at its receiver, as if they were sent simultaneously from two known positions. This information allows a mobile to place its position on a hyperbolic curve. When repeated for multiple pairs of anchor stations and access points and at known positions, the mobile can calculate its position by determining the intersection point of these hyperbolic curves.

Authentication and security are

always important—even more so when wireless devices are involved, because it's often hard to determine the entity that a mobile device is communicating with. For example, when you are in an unfamiliar location such as a large airport, are you aware of the access point your smartphone is associated with? 802.11az sets out to address this problem and provides the following mechanisms.

- A mobile device can ensure it's ranging to a trusted infrastructure. In other words, a mobile device should be able to reject potentially false range or positioning information from a nearby rogue access point.
- If a mobile device A is ranging to nearby access points, it should not be possible for another nearby nefarious mobile device B to overhear and decode the ranging packets. Therefore, B should not be able to calculate the position of A.
- Infrastructure that determines the location of a mobile device (assuming a user has given it permission to use that information) should be able to validate that the mobile is not falsifying its position.

Lastly, 802.11az provides for operation in the 60 GHz band, which promises greater range accuracy to a few centimeters with the increased bandwidth. In practice, this means position accuracy could have 10 times the resolution of 802.11-REVmc, which will open up new categories of location-based applications.

Location technology is coming of age at a time when IoT technologies are proliferating. The combination of IoT and accurate location-finding technologies will enable a wide range of novel context-aware services to automate tasks and support everyday work practice. Over the next few years, we expect fine-grained location-based services to become more prevalent and rapidly become a standard feature used by smart infrastructure.

REFERENCES

1. P. Bahl and V.N. Padmanabhan, "RADAR: An In-Building RF-based User Location and Tracking System," *Proc. 19th Ann. Joint Conf. IEEE Computer and Communications Societies* (INFOCOM 2000), 2000, doi:10.1109/INFCOM.2000.832252.
2. *P802.11-REVmc/D8.0, Aug 2016—IEEE Approved Draft Standard for Information Technology— Telecommunications and Information Exchange Between Systems—Local and Metropolitan Area Networks— Specific Requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, revision of IEEE Std. 802.11-2012, IEEE, 2016; https://ieeexplore.ieee.org/document/7558098.
3. E. Karapistoli et al., "An Overview of the IEEE 802.15.4a Standard," *IEEE Communications Magazine*, vol. 48, no. 1, 2010, pp. 47–53.

ROY WANT is a research scientist at Google. Contact him at roywant@google.com.

WEI WANG is a software engineer at Google. Contact him at weiwa@google.com.

STAN CHESNUTT is a software engineer at Google. Contact him at email chesnutt@google.com.



IEEE MultiMedia serves the community of scholars, developers, practitioners, and students who are interested in multiple media types and work in fields such as image and video processing, audio analysis, text retrieval, and data fusion.

Read It Today!

www.computer.org/multimedia

myCS Read your subscriptions through the myCS publications portal at **http://mycs.computer.org**