



Prototyping a Dynamic Indoor Positioning Solution

November 2018

Introduction

Over the past decade, Real Time Location Systems (RTLS) have become more common and consumer-accessible through smartphones and mobile applications. Outdoor RTLS like GPS have demonstrated widespread impact and utility; they can help drivers avoid traffic on the road, help emergency responders pinpoint distress signals, and provide valuable location data to countless use cases. Although GPS is able to provide highly accurate location data over a wide geographic area, satellite signals are unable to penetrate through solid objects like buildings. Indoor Positioning Systems (IPS), or indoor RTLS, are being built to overcome the physical limitations of GPS. However, these newer applications have been adopted more slowly due to technical hurdles and infrastructure costs.

As with outdoor RTLS, there are countless use cases for indoor RTLS. Applications range from medical asset tracking in healthcare to supply chain management. For example, locating lost medical devices can help reduce replacement costs and improve hospital logistics. There are also multiple approaches to building indoor tracking solutions that leverage technologies like Bluetooth, WiFi, RFID, Ultra Wide Band (UWB), and Ultrasound. While there are many opportunities for IPS product development in this space, the technologies required to build useful IPS products are not nearly as well understood or as ubiquitous as outdoor technologies like GPS.

To better understand and characterize the issues associated with indoor positioning, the **Leverege R&D team** is building a dynamic, high-resolution indoor tracking system that fuses data from multiple sources. In this paper, we present one of our Bluetooth RTLS prototypes and walk through some of the core research, design, and engineering considerations that go into building a robust indoor RTLS system.

Working with Bluetooth RTLS Hardware

Our R&D team used Bluetooth Low Energy (BLE) tracking tags and programmable location “hub” devices that served as general proximity tag “sniffers.” The devices report tag data as a Received Signal Strength Indicator (RSSI) and an associated Media Access Control (MAC) address, which paired serve as a unique device ID.

There is no “right” hardware for all Bluetooth RTLS solutions since the spaces in which the RTLS will operate place unique constraints and requirements on the hardware. [Although the maximum range of BLE signal exceeds 20m under ideal conditions](#), indoor RTLS applications are rarely “ideal” scenarios. They’re dynamic, asymmetrical, and often unique. Signal obstructions like pillars, furniture, and walls can reduce effective signal range significantly. Hub placement and hub density in an indoor space are two important considerations when designing the system. Our team mapped out the physical space in which the prototype would be tested early on.

We chose Active RadBeacon Dots as our prototypical asset tracking tags:



*A RadBeacon Dot
Image Credit: Leverage*

We chose RadBeacon Dots for several reasons:

- **Compatibility with iBeacon, Eddystone, and Altbeacon Protocols** – Depending on which BLE protocol is used (iBeacon, Eddystone, or AltBeacon), additional data such as manufacturing ID is also accessible. Although this data isn’t necessary for this prototype

IPS, being protocol agnostic means that the tags can be reused in other indoor tracking applications.

- **Configurable Transmission (Tx) Power** – While maximizing Tx power on the tags ensures that the location hubs are able to receive data, this can result in noisy or inaccurate data. Because RSSI is easily attenuated by multiple factors, this can create ambiguity in location determination. By lowering the Tx power, hubs that are farther away are effectively filtered out by being “out of range.” Having a configurable Tx power provides a way to experimentally determine the optimal tag settings for a particular environment.

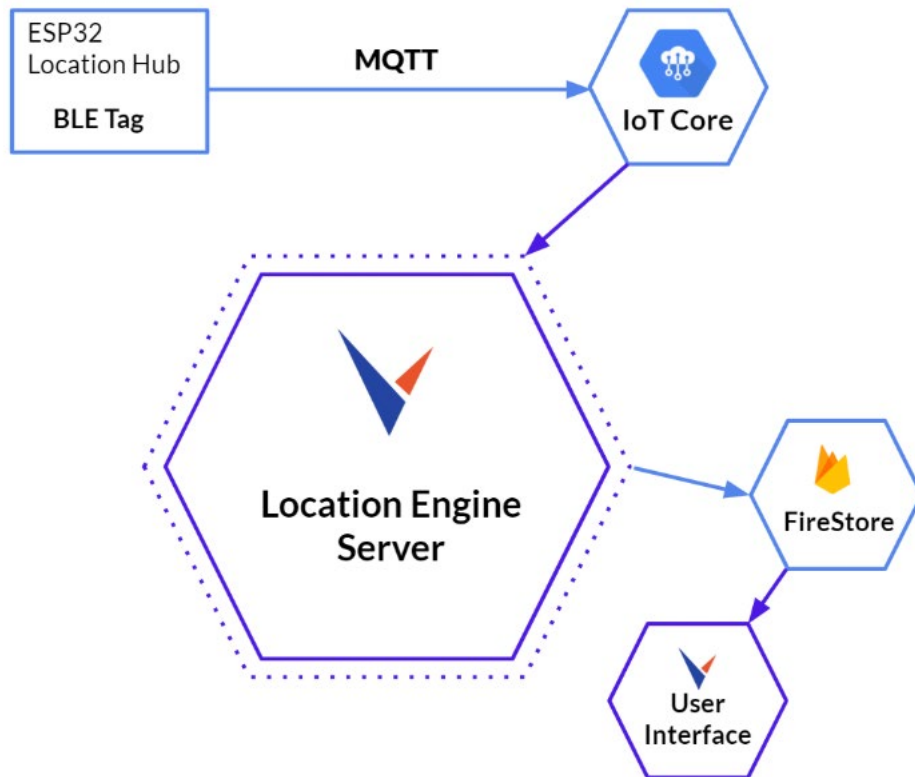
For the hubs, while there were many programmable microcontrollers with built-in WiFi and Bluetooth chips, the ESP32 development board has extensive documentation and support via its open source community. It is also significantly more cost effective than other WiFi + Bluetooth microcontrollers. Moreover, commercial demos of Bluetooth RTLS solutions can be upwards of [\\$3500 per 100 square meters](#), so we wanted to assess the true deployment cost of this type of solution through prototyping.



*ESP32 Microcontroller
Image Credit: Leverage*

System Design and Engineering

Our team collaborated with [Pure Engineering](#) to create custom firmware for the ESP32 location hubs. We configured the boards to scan the vicinity for any Bluetooth devices for four seconds before reporting the paired RSSI and MAC address of every device found in [proto3 format](#) to Google Cloud IoT Core via the MQTT protocol. Leverage's Location Engine server processed the data before being pushed to Firestore and displayed in real-time on a user interface that we built and optimized for indoor positioning solutions.



System diagram of Indoor Tracking System Prototype
Image Credit: Leverage

Developing the Leverage Location Engine


Indoor RTLS solutions require not only tuning hardware to the nuances of each physical space but also creating dynamic and reliable positioning algorithms. Our R&D team developed Leverage's Location Engine through a series of prototypes. Each prototype was more accurate, precise, and robust than the previous version. Each brought us one step closer to the mature algorithms required in a business-critical Location Engine.

As an early proof of concept, we first applied a simple filter to calculate which location hub was closest to a given tag based on nearest RSSI. Because of how noisy the raw BLE data was, we delved into increasingly sophisticated algorithms to improve RTLS accuracy. After testing geolocation techniques, such as multilateration and Kalman filtering, our team discovered that a particle filter delivered the most accurate and stable location results in dynamic indoor environments.

Building the Engine Through Prototyping

Prototype I: Nearest RSSI Algorithm

At Leverage, rapid prototyping is our culture. Every product that crosses the finish line carries a whole chain of past iterations. The first Location Engine prototype consisted of a simple filter that identified the location hub which reported the strongest RSSI value. Since signal attenuates with increasing distance, the strongest signal corresponds to the nearest location hub under ideal conditions. Although this method is highly reproducible, its accuracy depends on location hub density.



Experimental Layout for Nearest RSSI. Small blue circles and large purple circles represent ESP32 location hubs and RadBeacon Dot tags respectively. Image credit: Leverage

This approach comes with two main drawbacks that make it infeasible at scale:


1. It requires a large investment in infrastructure: providing access to WiFi backhaul and power, the physical set-up of hubs, and the mapping/registration of the hubs to a user interface. Hardware and labor costs are high and would increase rapidly at scale.
2. Based on the number of devices transmitting 2.4GHz signals (WiFi and Bluetooth

frequency band), there can be erratic signal behavior due to interference from other location hubs and connected devices like smartphones and wireless headsets. Human bodies also attenuate the 2.4GHz signal significantly.

Prototype II: Multilateration Algorithms

Our team explored numerous lateration methods. Trilateration algorithms compute position by measuring distances to three known points. Multilateration algorithms have the same working principle but rely on more than three known points in their calculations. To perform trilateration and multilateration on the BLE location data, we first converted RSSI values into distances.

We arranged nine hubs in a rough 3x3 grid spanning ~100 square meters in our furnished office environment, mimicking the hub density of commercial Bluetooth RTLS products. Our trilateration algorithm consistently achieved 3-meter location accuracy.



Experimental System Layout with a 3x3 grid. Small blue circles and large purple circles correspond to ESP32 location hubs and RadBeacon Dot tags respectively. Image credit: Leverage

Prototype III: The Particle Filter

Developing robust and dynamic IoT solutions is all about learning from one prototype and leveraging that insight toward building the next, more powerful prototype. We knew we would need to get creative to increase RTLS resolution.

Our team applied a particle filter to the BLE data to predict an asset's location stochastically for greater accuracy. Rather than using a singular point on a map, or a series of X, Y coordinates, this filter created a small cloud of candidate locations that were then overlaid onto the floor plan UI.

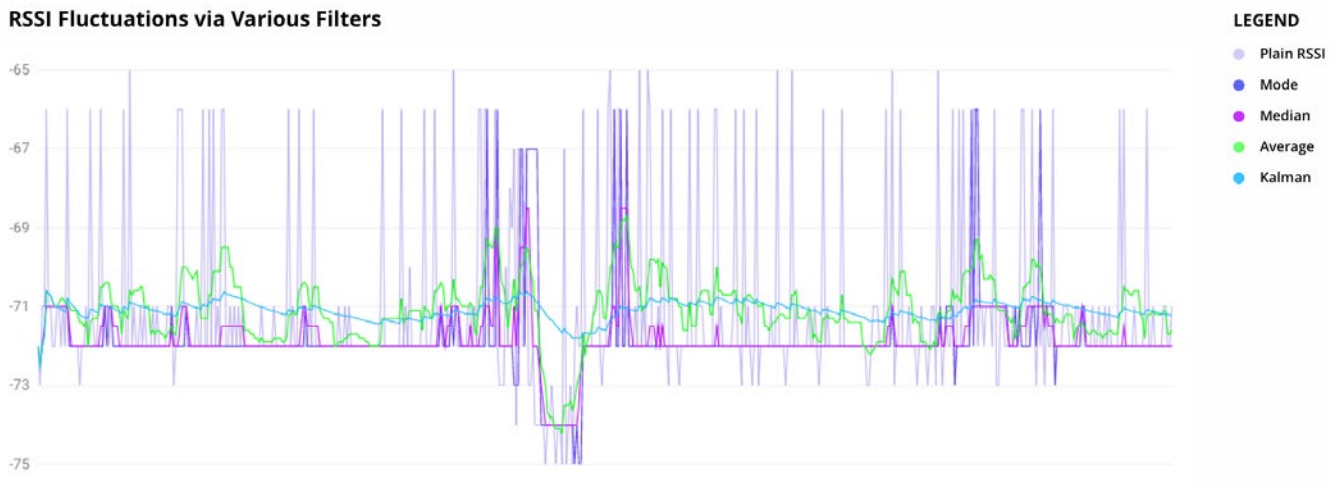
The particle filter predictions were very similar to those produced by the trilateration algorithm but slightly more accurate. Filtered particles clustered around the predicted trilateration location. With a particle filter, our budding location engine began boasting <3 meter accuracy.

Unlike with trilateration or multilateration, increasing the number of location hub inputs does not affect the complexity of the calculation algorithm. Furthermore, introducing new sources of data from other indoor tracking technologies, like WiFi or RFID, simply means applying the corresponding probability distribution function to the existing particle filter. Our particle filter approach has proven effective for the location engine component of an accurate, reliable indoor RTLS.

Enhancing Location Accuracy

One major shortcoming of Bluetooth data is the large fluctuations in RSSI values that are reported when a tag is stationary. We observed RSSI fluctuations ranging between 6-15 dBm in all of the location hubs used for testing. These fluctuations were large enough to skew the results and cause distant hubs to appear much closer to an asset than they actually were. As shown below, our team tested several filters to help reduce this noise. The Kalman filter and moving average filters consistently produced the best results.

RSSI Fluctuations via Various Filters



There are two likely causes of these RSSI fluctuations:

1. **Signal attenuation or interference by other devices and moving objects** – Although we still observed RSSI fluctuations when the office was emptier during the evenings, those fluctuations were considerably more stable when there were fewer moving bodies and personal devices causing signal interference.
2. **Bluetooth chip on the ESP32** – Aside from price, the other major differences between the ESP Bluetooth module and Nordic Semiconductors' nRF modules lie in power consumption and processor speed. Our R&D team plans to conduct more rigorous testing of nRF modules to determine the best chip for indoor RTLS.

For a prototypical Bluetooth indoor RTLS, the multilateration approach produces satisfactory results. However, the particle filter approach is more scalable and practical for consumer or

industry-grade RTLS. Because a particle filter is able to incorporate multi-source data from any number of technologies, it can be adapted to numerous use cases and we plan to layer in other sources of data in the future.

The Future of Indoor Positioning Solutions

Bluetooth RTLS can determine approximate locations down to several meter accuracy, but many business use cases require much more precise positioning. Some rooms are only a few meters squared after all. We have determined that successful indoor positioning systems often require balancing several location technologies—Bluetooth, RFID, Ultrasonic, etc.—and processing their data through a sophisticated particle filter-based location engine to create a robust, high-resolution RTLS.

Indoor RTLS pose significant challenges for solutions providers given the need for embedded engineering expertise that ranges across multiple hardware and software systems—from Bluetooth to RFID, from Kalman filtering to machine learning. And while some solutions providers may advertise expertise in a few of these domains, at its core, indoor positioning is all about striking a harmony between numerous interlocking hardware and software systems. We plan to continue to allocate significant R&D resources to prototype and build a multi-source location engine that provides businesses with the most reliable, highest-resolution indoor positioning solutions possible.