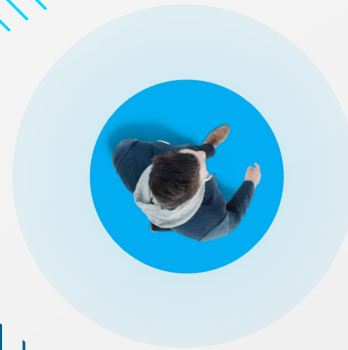SILICON LABS

Whitepaper

Bluetooth Angle Estimation for Real-Time Locationing

By Sauli Lehtimäki
Senior Software Engineer, Silicon Labs

Bluetooth Angle Estimation for Real-Time Locationing

By Sauli Lehtimäki, Senior Software Engineer, Silicon Labs

Bluetooth Angle of Arrival (AoA) and Angle of Departure (AoD) are new technologies that establish a standardized framework for indoor locationing. With these technologies, the fundamental problem of locationing comes down to solving the arrival and departure angles of radio frequency signals. In this paper, we explain the basics of these technologies and give some theory for estimating direction of arrival. Currently the Bluetooth AoA/AoD specifications are in a mature state but not yet public. Therefore, this paper will only cover the general concepts without going into the details of the specification.

Locationing technologies have many useful applications, one example being GPS, which is widely used all over the world. Unfortunately, GPS does not work very well indoors, so there is a real need for better indoor positioning technologies. Our goal is to track the locations (or angles) of individual objects with an external tracking system, or for a device to track its own location in an indoor environment. This kind of locationing system can be used to track assets in a warehouse or people in a shopping mall, or people can use locationing for their own wayfinding.

Bluetooth Angle of Arrival and Angle of Departure

Let's consider a device with a multiple antenna linear array for a receiver and a device with one antenna for a transmitter. Also, assume that the radio wave travels as a planar wave front rather than spherically, which we can safely assume when looking from a distance. If the transmitter, which is sending a sine wave through the air, lies on the normal line perpendicular to the array line, every antenna (channel) of the array will see the incoming signal in the same phase. If the transmitter does not lie on the normal line, then the receiving antennas will see phase differences between the channels. This phase difference information can be used to calculate the angle of arrival.

In practice, the receiver will need to have multiple ADC channels or use an RF switch to take samples from each individual channel. The samples are called "IQ-samples" since a sample pair of "In-phase" and "Quadrature-phase" readings are taken from the same input signal. These samples have a 90 degree phase difference in the sampling. When this pair is considered to be a complex value, each complex value contains both phase and amplitude information and can be an input for the arrival angle estimation algorithm.

Radio waves travel at the speed of light, which is 300,000km/s. When using frequencies around 2.4GHz, the corresponding wavelengths are about 0.125m. The maximum distance between two adjacent antennas for most estimation algorithms is a half wavelength. Many algorithms require this, otherwise we get effects similar to aliasing. There is no theoretical minimum distance limitation, but in practice, the minimum size is limited by the mechanical dimensions of the array plus, for example, mutual coupling between the antenna elements.

> The maximum distance between two adjacent antennas for most estimation algorithms is a half wavelength.

For Angle of Departure, the fundamental idea of measuring phase differences is the same, but device roles are swapped. In AoD, the device being tracked uses only one antenna, and the transmitter devices use multiple antennas. The transmitter device sequentially switches the transmitting antenna, and the receiving side knows the antenna array configuration and switching sequence.

When considering this from an application point of view, we can see a clear difference between these two techniques. In AoD, the receiving device is able to calculate its own position in space using angles from multiple beacons and their positions (by triangulation). In AoA, the receiving device tracks arrival angles for individual objects. Still, it is good to note that different combinations of these can be performed; so, these techniques do not limit what can be done at the application level. Both in Bluetooth AoA and AoD, the related control data is sent over a traditional data channel. Typically, these techniques can achieve a couple of degrees angular accuracy and around 0.5 m locationing accuracy, but these figures are highly dependent on the implementation of the locationing system.

Challenges

One of the biggest and perhaps most obvious challenges in this subject is answering the question: "How are angle estimates calculated based on the sample data?" It is not enough to calculate angle estimates in an ideal environment; we must also calculate them in environments with very heavy multi-path wheresignals are highly correlated or coherent. By coherent signal, we mean a signal that is delayed, and a scaled version of some other signal. This can be the case when radio waves are reflected from walls, for example.

Other challenges include signal polarization. In most cases, we cannot control the polarization of the mobile device, so the system has to take this into account. Also signal noise, clock jitter, and signal propagation delays add their own variables to the problem. Depending on the system scale, the RAM and especially CPU requirements can be demanding for an embedded system. Many of the well performing angle estimation algorithms require a significant amount of processing power from the CPU.

Angle of Arrival Theory

Angle estimation methods and antenna arrays are essential for the locationing system to work properly. The history of direction finding theory goes back more than 100 years when the first attempts to solve this problem were made using directional antennas and, obviously, purely analog systems. In the years following, test methods moved to the digital world, but the basic principles are still quite similar. These direction-finding methods are already used in many applications, such as medical equipment, security, and military devices.

In this section, we will discuss the basics of some typical antenna arrays and estimation algorithms. By direction finding, we refer to the general problem of estimating arrival and departure angles.

Antenna Arrays

Antenna arrays for direction finding are usually divided into categories. The most common ones discussed here are Uniform Linear Array (ULA), Uniform Rectangular Array (URA), and Uniform Circular Array (UCA). The linear array is a one-dimensional array, meaning that all the antennas in the array lie on a single line, whereas the rectangular and circular arrays are two-dimensional arrays, meaning that the antennas are spread in two dimensions (on a plane). By using a one-dimensional antenna array, one can reliably measure only the azimuth angle, assuming the tracked device moves consistently on the same plane. Furthermore, with two-dimensional arrays, one can reliably measure both azimuth and elevation angles in the 3D half-space. If the array is extended to a full 3D array (antennas spread on all three Cartesian coordinates), then we will be able to measure the full 3D space.

Designing an antenna array for direction finding is not a straightforward task. When antennas are placed in an array, they start affecting each other's responses; this is called mutual coupling. We also have to keep in mind that, in most cases, we cannot control the polarization of the transmitting end. This creates an additional challenge for the designer. In IoT applications, the devices are often expected to be small and even work in very high frequency bands. Estimation algorithms often require some certain properties from the array. For example the estimation algorithm called ESPRIT works on the mathematical assumption that the array is divided into two identical subarrays [3].

Angle Estimation Algorithms

Let's look at the mathematical/algorithmic problem of estimating the angle of arrival based on the input IQ-data. The problem definition itself is simple: estimate the arrival angle of an emitted (narrowband) signal arriving at the receiving array. While the problem statement sounds very trivial, a robust solution (that works in real life) for this problem is not easy and can require much processing power from the hardware.

Next, we will present two different approaches for solving this problem. The first one is basic and called a classical beamformer. The second is a more advanced technique called Multiple Signal Classification (MUSIC). We will not go through proofs of any theorems or reasons why these methods work, but rather give a high-level view of how the algorithms work. Deeper studies about these estimation algorithms can be found from [1] and [2].

Classical Beamformer

Let's begin with a mathematical model of a uniform linear array. We are given a data vector of IQ-samples for each antenna. Let this vector be called x. Now, there is a phase shift seen by each antenna (which can be 0) plus some noise, n, in the measurements, so x can be written as a function of time t:

$$x(t) = a(\theta)s(t) + n(t), \qquad (1)$$

where s is the signal sent over the air, and a is the steering vector of the antenna array:

$$a(\theta) = [1, e^{j2\pi d sin(\theta)/\lambda}, ..., e^{j2\pi(m-1)d sin(\theta)/\lambda}], \qquad (2)$$

where d is the distance between adjacent antennas; λ is the wavelength of the signal; m is the number of elements in the antenna array, and θ stands for the angle of arrival.

Steering vector (2) describes how signals on each antenna are phase shifted because of the varying distances to the transmitter. By using (1), we can calculate an approximation of the so-called sample covariance matrix, R_{xx}, by calculating

$$R_{xx} \approx \frac{1}{N} \sum_{t=1}^{N} x(t)x^H(t), \qquad (3)$$

where H stands for the Hermitian transpose of a matrix.

The sample covariance matrix (3) will be used as an input for the estimator algorithm as we will see.

The idea of the classical beamformer is to maximize the output power as a function of the angle, similar to how a mechanical radar works. If we attempt to maximize the power, we end up with the next formula:

$$P(\theta) = \frac{a^H(\theta)R_{xx}a(\theta)}{a^H(\theta)a(\theta)}$$

(4)

Now, to find the arrival angle, we need to loop through the arrival angle θ and find the peak maximum power, P. The angle, ortheta,producing the maximum power corresponds to the angle of arrival. While this approach is quite simple, its accuracy is not generally very good. Therefore, let's introduce another method, which is a bit better in terms of accuracy. See, for example, [4] for an algorithm accuracy comparison.

MUSIC (Multiple Signal Classification)

One type of estimation algorithm is the so-called subspace estimator, and one popular algorithm of that category is called MUSIC (Multiple Signal Classification). The idea of this algorithm is to perform eigendecomposition on the covariance matrix R_{xx}:

$$R_{xx} = VAV^{-1},$$

(5)

where A is a diagonal matrix containing eigenvalues and V containing the corresponding eigenvectors of R_{xx}.

Assume we are trying to estimate the angle of arrival for one transmitter with an n antenna linear array. It can be shown that the eigenvectors of R_{xx} either belong to so-called noise subspace or signal subspace. If the eigenvalues are sorted in ascending order, the corresponding $n-1$ eigenvectors span the noise subspace, which is orthogonal to the signal subspace. Based on the orthogonality information, we can calculate the pseudo spectrum P:

$$P(\theta) = \frac{1}{a^H(\theta)VV^Ha(\theta)}$$

(6)

As in a classical beamformer, we loop through the desired values of θ and find the maximum peak value of P, which corresponds the angle of arrival (argument θ) we wish to measure.

In an ideal case, MUSIC has excellentresolution in a good SNR environment and is very accurate. On the other hand, its performance is not very strong when the input signals are highly correlated, especially in an indoor environment. Multipath effects distort the pseudo-spectrum causing it to have maximums at the wrong locations. More information about the conventional beamformer and MUSIC estimators can be found from [3].

Spatial Smoothing

Spatial smoothing is a method for solving problems caused by multipathing (when coherent signals are present). It can be proven that the signal covariance matrix can be "decorrelated" by calculating an averaged covariance matrix using subarrays of the original covariance matrix. For a two-dimensional array, this can be written as the following

$$\overline{R} = \frac{1}{M_sN_s}\sum_{m=1}^{M_s}\sum_{n=1}^{N_s}R_{mn},$$

(7)

where M_s and N_s are the number of subarrays in x- and y-directions respectively and R_{mn} stands for the (m,n):the subarray covariance matrix. An example proof of this formula and more information can be found from [2].

The resulting covariance matrix can now be used as a "decorrelated" version of the covariance matrix and fed to the MUSIC algorithm to produce correct results. The downside of spatial smoothing is it reduces the size of the covariance matrix, which further reduces the accuracy of the estimate.

Other Locationing Technologies

In this section, we briefly present two other locationing technologies for comparison. These two methods use different kinds of algorithms/methods for locationing than those presented in this paper.

With Received Signal Strength Indicator (RSSI), the basic idea is to measure the signal strength of the received signal to get a distance approximation between RX and TX. This information can be used to trilaterate the position of a receiver device based on multiple distance measurements from different transmitter points. This technology requires only one antenna per device, but is not usually very accurate in an indoor environment.

With Time of Arrival / Time of Flight (ToA/ToF), we measure the travel time of a signal between RX and TX and use that to calculate the distance between the ends. This distance is then used to trilaterate the position of the receiver. In ToA, all devices are time-synchronized. This technology also requires only one antenna per device, but, on the other hand, it requires very high clock accuracy to get reasonable positioning accuracies. There is also a variant of this technology called TDoA, where only the receiver devices need to be time-synchronized, and the estimation algorithms use the time difference for calculating position estimates.

Conclusion

Bluetooth Angle of Arrival and Angle of Departure are new emerging technologies that can be used to track assets as andforindoor positioning and wayfinding. These are phase-based direction finding systems that require an antenna array, RF switches (or a multi-channel ADC), and processing power to run the estimation algorithms. Designing a proper antenna array, as well as an angle estimation algorithm are essential for a RTLS system. Strong performing estimator algorithms are often not computationally cheap. Other positioning technologies include (but are not limited to) RSSI based methods and ToA based methods, but only phase-based AoA/AoD currently have a standardized framework in Bluetooth.

Sources:

[1] H. Krim, M. Viberg, "Two Decades of Array Signal Processing", IEEE Signal Processing Magazine, July 1996, pp. 67-94

[2] Y.-M. Chen, "On Spatial Smoothing for Two-Dimensional Direction-of-Arrival Estimation of Coherent Signals", IEEE Transactions on Signal Processing, Vol. 45, No. 7, July 1997

[3] Z. Chen, G. Gokeda, Y. Yu, "Introduction to Direction-of-Arrival Estimation", Artech House, 2010

[4] N. A. Baig, M. B. Malik, "Comparison of Direction of Arrival (DOA) Esimation Techniques for Closely Spaced Targets", International Journal of Future Computer and Communication, Vol. 2, No. 6, December 2013