AN14865

Channel Sounding Fundamentals for the KW47 and MCX W72

Rev. 1.0 — 10 December 2025 Application note

Document information

Information	Content
Keywords	AN14865, KW47, MCX W72, Channel Sounding Fundamentals
Abstract	This document provides an overview of the fundamentals for Channel Sounding technology and how it can be used for custom solutions and applications.

1 Introduction

Bluetooth Low Energy (LE) technology is now used in many devices. These devices include smartphones, tablets, smartwatches, headphones, Internet of Things (IoT) devices, and edge devices, such as smart appliances and smart lights.

Over the years, the possible applications of this technology have evolved toward more complex applications like device positioning, proximity detection, asset tracking, and indoor navigation.

The latest Bluetooth Core Specification version 6.0, from the Bluetooth LE Special Interest Group (SIG), includes Channel Sounding (CS). This feature makes Bluetooth LE a tool for secure fine ranging between two devices.

This document provides an overview of the fundamentals for CS technology and how it can be used for custom solutions and applications.

2 Bluetooth CS for distance estimation

This section focuses on the CS standard description. The CS consists of a combination of software, Link Layer (LL), and Physical Layer (PHY) procedures that produce a result for the application layer. Each of these steps is explained separately.

2.1 Architecture

Bluetooth CS takes place in a one-to-one topology between two devices. The one that starts the communication is known as initiator and the other device that answers this request is known as Reflector.

For two devices to be able to participate in a CS procedure, both must have a Bluetooth LE controller, which supports the CS feature. The CS feature of the Bluetooth LE stack impacts the PHY, LL, Host Controller Interface (HCI), and Generic Access Profile (GAP).

To support the CS feature, a new Bluetooth LE profile is introduced called the Ranging Profile (RAP). The ranging client has a direct connection to the distance measurement application. It has information from both the devices to perform a calculation. The ranging server assists the ranging client by sending its local device capabilities, receiving the ranging client configurations, and sending the ranging result at the end of the CS procedures.

The device acting as the ranging client assumes the initiator role and the ranging server assumes the reflector role. However, these roles are configurable and interchangeable depending on the application. It is also possible that a single host device supports both client and server implementations and switches according to the application needs.

Bluetooth CS requires that the application layer calculates the distances using the data that the controller provides. This data is the result of signal exchanges and low-level measurements made by both devices during a CS procedure. The data is then passed to the application layer in HCI events.

2.2 CS specification overview

There are several procedures that can be performed at the LL level before an application is able to compute a distance estimation. The Bluetooth LE standard defines some of these procedures while others are open to the implementation of application developer. Figure 1 shows the main procures that takes place before a CS measurement.

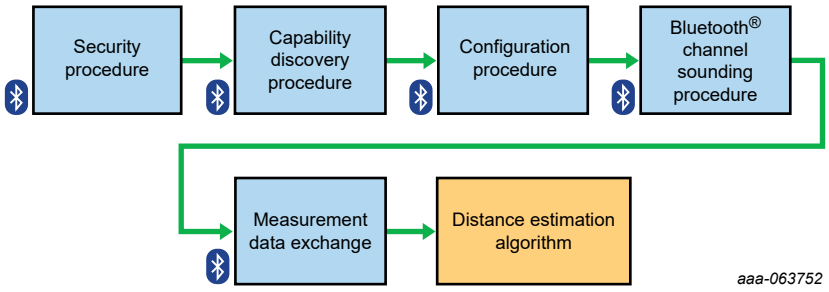

Figure 1. Bluetooth CS flowchart

2.2.1 Security procedure

Bluetooth CS has its own security capabilities. These security capabilities are different from the capabilities that are associated with the standard Bluetooth LE connection between the devices. This security procedure allows the two devices to exchange the parameters securely that are later used in the Bluetooth CS security functions.

The Bluetooth CS security procedure starts with the Bluetooth LE central device generating three random numbers. The following are the three random numbers:

- Initialization vector (CS_IV_C)
- Instantiation nonce (CS_IN_C)
- Personalization vector (CS_PV_C)

Then, the central device sends these numbers to the peripheral using a LL_CS_SEC_REQ PDU. The peripheral generates its own set of numbers with the same rules (CS_IV_P, CS_IN_P, CS_PV_P), and sends them back to the central in a LL_CS_SEC_RSP PDU. Each of the LL of the device concatenates these two sets of security parameters. It results in both devices having the same security parameters CS_IV, CS_IN, and CS_PV.

2.2.2 Capabilities exchange procedure

Two devices forming a Bluetooth LE connection can have a different set of capabilities. Therefore, a mutually supported configuration must be established before a CS procedure. These capabilities include:

- PHY support
- Round-Trip Time (RTT) accuracy
- Supported CS modes
- Attack detection support
- Maximum number of antenna paths

Any of the two devices can initiate the exchange of capabilities. It consists of interchanging its device details through specific PDUs called LL_CS_CAPABILITIES_REQ and LL_CS_CAPABILITIES_RSP. A device can store capabilities received by a device and select not to perform this exchange in the future.

2.2.3 Configuration procedure

During this procedure, the devices select the specific configuration that is used according to the information exchanged during the capabilities exchange procedure.

A host can maintain various configuration parameter sets with a unique identifier. These configuration parameter sets can be used in specific scenarios according to the needs or capabilities of the peer device.

The device that transmits the LL_CS_CONFIG_REQ Protocol Data Unit (PDU), selects its CS role. The CS roles are initiator or reflector. After the other device responds with the LL_CS_CONFIG_RSP PDU, it takes the other role.

2.2.4 CS start

After the devices finish the security procedure, exchange capabilities, and agree on a configuration, the CS procedure can be initiated. The initiation can be achieved via the following three PDUs:

- LL_CS_REQ
- LL_CS_RSP
- LL_CS_IND

The LL_CS_REQ and LL_CS_RSP PDUs include proposed timing and structural parameters for each device. The proposed timing and structural parameters control how the time is divided during the CS procedure. When the central role sends the LL_CS_IND after it receives LL_CS_REQ or LL_CS_RSP, it indicates that the CS procedure must now start. It uses the parameters values that both devices accept based on the proposals in the other PDUs.

2.2.5 CS procedure

The Bluetooth CS procedure can be defined as a set of time and frequency slots. In this procedure, two devices exchange a combination of RF signals to estimate the physical characteristics of the transmission channel. These exchanges are always bidirectional.

A CS procedure is divided into one or more CS events. Each CS event is also divided into CS subevents. These subevents are also partitioned into CS steps. It is within the CS steps that packets and tones are transmitted and received.

Figure 2 shows the graphical view of the CS hierarchical time divisions.

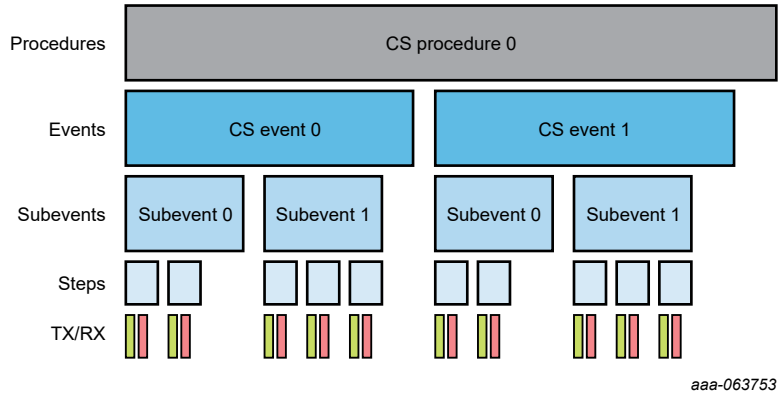

Figure 2. Bluetooth CS procedure divisions

Several parameters control the timing, duration, and scheduling of procedures, events, subevents, and steps. These parameters are configured during the initialization procedures. Table 1 shows some of the main parameters.

Table 1. CS procedure parameters

Parameter	Range of values	Description
Number of CS procedure repetitions	0 to 65535	The number of CS procedures that are executed before CS is terminated. A value of 0 indicates that CS procedures

Table 1. CS procedure parameters*...continued*

Parameter	Range of values	Description
		must be executed until a CS procedure repeat termination is invoked.
Number of subevents per event	1 to 16	The number of subevents anchored off the same event.
Subevent interval	0 µs or 625 µs to 40959.375 ms	Time interval between the beginning of the CS subevent and the next CS subevent within the same CS event. A value of 0 means that events are not divided into subevents.
Number of steps per subevent	2 to 160	Configured by application. A maximum of 256 steps per procedure.

There can be programmable gaps between CS events to allow for other Bluetooth LE data transmission or simultaneous connections.

All procedure, event, subevent, and step start times are directly or indirectly connected to a selected connection event. This event is part of the Bluetooth LE connection that runs the CS initiation procedures.

Table 2 contains all the timing parameters that control a CS procedure.

Table 2. CS timing configurations

Parameter	Description	Time duration (µs)
T_FCS	Time for frequency hopping between every CS step.	15 to 150
T_FM	Time for frequency measurement in mode 0.	80
T_SY	Time for synchronization sequence.	Depends on the PHY used and CS_Sync packet length.
T_IP1	Interlude period 1. Transition time from TX to RX and the vice versa when transmitting packets.	10 to 145
T_IP2	Interlude period 2. Transition time from TX to RX and vice versa when transmitting tones.	10 to 145
T_GD	Transition time between a packet and a tone.	10
T_RD	Ramp-down time for transmission. Used to remove energy from the RF channel.	5
T_PM	Single-antenna phase measurement period.	10, 20, and 40
T_SW	Time period reserved for antenna switching.	1, 2, 4, and 10

2.2.6 CS steps

The CS procedure consists of a set number of steps in which the radio activity takes place. The CS steps consist of a combination of modulated and unmodulated RF signals. The modulated sections convey

synchronization information and extract the RTT information. However, the unmodulated sections are used to measure the phase shift.

The four different modes for CS steps, each with a different usage, are explained in the following sections.

2.2.6.1 Mode-0

CS Mode-0 is related to calibration and time synchronization between devices. It measures the frequency offset between the initiator and the Reflector Device (RD) at a given frequency. The support for Mode-0 is mandatory for all devices that include a CS feature.

The initiator transmits a CS_SYNC packet on a selected channel and frequency. Then the reflector responds with a CS_SYNC packet and a tone on the same frequency as received from the initiator. The initiator is able to tune its receiver and set its gain with the preamble from the CS_SYNC packet. When receiving the tone from the reflector, the initiator calculates a value called the Fractional Frequency Offset (FFO). The FFO is used to compensate for the difference between the devices. It is also used to improve the accuracy of the results within the subsequent CS steps in a CS event.

2.2.6.2 Mode-1

CS Mode-1 measures the RTT of CS_SYNC packets exchanged between the initiator device and RD. The support for Mode-1 is mandatory for all devices using the CS feature. To measure the RTT, the initiator records a timestamp when it sends the CS_SYNC packet. This timestamp is known as the Time of Departure (ToD). Then it records a second timestamp when it receives the CS_SYNC packet back from the reflector. This timestamp is known as the Time of Arrival (ToA).

All the timing parameters have been pre-agreed during the configuration procedure. The initiator includes the following times in its RTT calculation:

- The time the receiver needs to prepare and send its packet (T_IP1)
- The packet duration (T_SY)
- The ramp-down time (T_RD)

Figure 3 shows the transmission of a CS_SYNC packet in a mode-1 transmission.

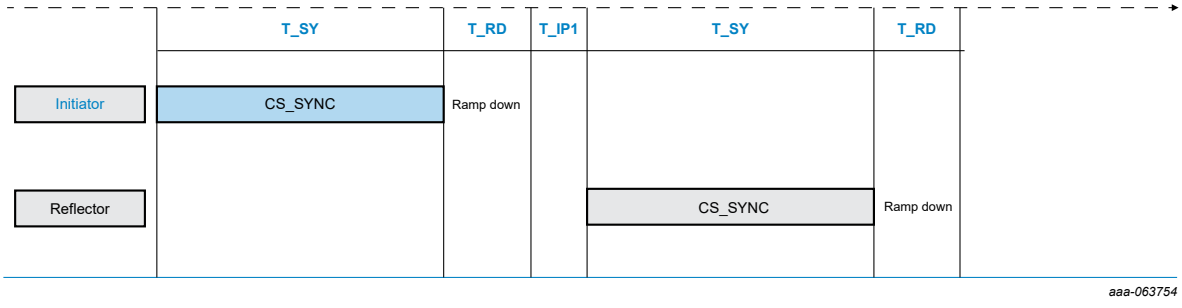

Figure 3. Mode-1 transmission of CS_SYNC packets

For more details on the different methods for timestamping packets, see Section 3.

2.2.6.3 Mode-2

CS Mode-2 is used for Phase-Based Ranging (PBR). It measures the phase shift between RF tones exchanged between the initiator device and RD. For more details on this method, see Section 4. Support for Mode-2 is mandatory for any device supporting the CS feature.

The Mode-2 step starts when the initiator sends an unmodulated RF carrier signal, known as CS tone. It goes through the selected channel and each available antenna path. After the ramp-down and interlude period, the reflector replies with a CS tone in the same frequency through each of its available antenna paths (N_AP).

The initiator device measures the phase of the CS tone received from the reflector during the T_PM period. It measures the phase once for each antenna path. It reserves T_SW time for antenna switching and uses the compensation values from the Mode-0 steps. These measurements are passed to the application layer in the form of an array of in-phase, and quadrature components, known as IQ data or IQ samples.

For added security, after each time slot allocated for CS tone transmissions, an additional time period is available, which is known as CS tone extension. The transmission during this time slot is randomized. When needed, the device transmits a CS tone using the same antenna path used in the last transmission. This results in a total transmission time of (T_SW + T_PM)*(N_AP + 1), as it can be observed in Figure 4.

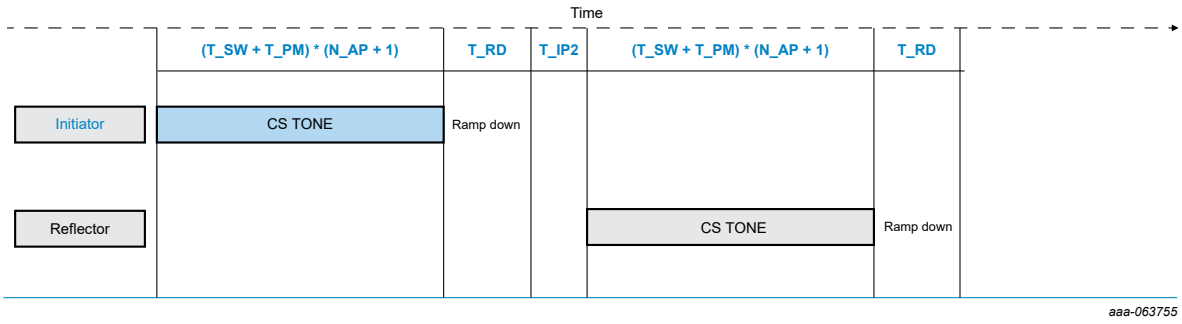

Figure 4. Mode-2 transmission of CS tones

2.2.6.4 Mode-3

CS Mode-3 measures the RTT of CS_SYNC packets. It also measures the phase shift of CS tones exchanged between the initiator device and RD. This mode is a combination of the modes 1 and 2 that allows for RTT and PBR calculations. The support for Mode-3 is optional in the CS feature implementation.

As the implementation of Mode-3 is optional, an initiator device can discover during the capabilities exchange phase that its peer RD does not support it. In this case, the initiator device can use a combination of Mode-1 and Mode-2 measurements, in mode sequence. For more details, see Section 2.2.7.

Figure 5 shows the timing diagram of a Mode-3 transmission.

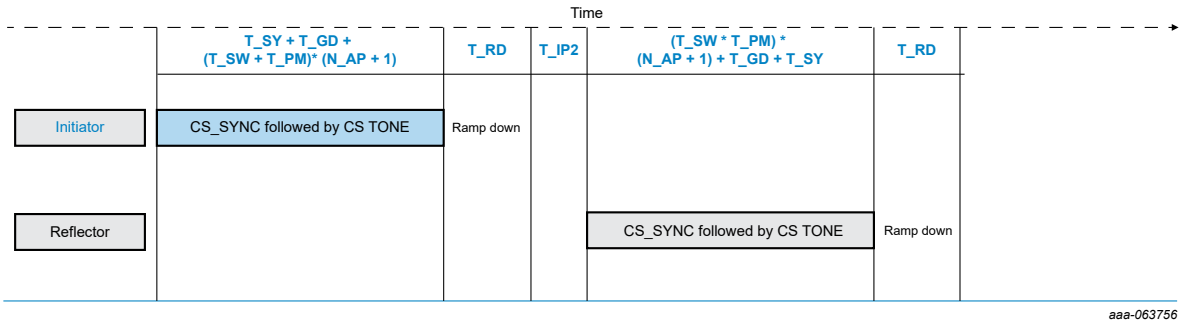

Figure 5. Mode-3 transmission of CS_SYNC packets and CS tones

2.2.7 Mode sequencing

A Bluetooth CS procedure always involves the execution of a sequence of at least two different modes. Section 2.2.6.1, Section 2.2.6.2, Section 2.2.6.3, and Section 2.2.6.4 have only focused on individual steps and how the information must be transmitted between devices. The distance estimation requires multiple exchanges

of information for accuracy or because the method requires it. PBR relies on phase difference calculation, which can only be done with at least two steps of CS tones exchange. It involves more than one signal and one frequency.

In general, more CS steps give more data. It lets the application layer create more accurate distance measurements. However, it takes more time to execute.

The sequencing of these steps, the mode variation, and the patterns of their execution are covered in this section.

2.2.7.1 Mode combinations

Any CS event and subevent always start with one or more Mode-0 steps for frequency offset measurement and time synchronization. It is followed by a sequence of non-mode-0 steps in the same event. It is possible to use a combination of two non-mode-0 modes in the same event. The primary mode is known as Main_Mode. If there is a secondary mode, it is known as Sub_Mode.

Table 3 is extracted from the Bluetooth Core Specification. It lists all the permitted non-mode-0 mode combinations.

Table 3. Permitted non-mode-0 mode combinations

Main_Mode	Sub_Mode
Mode-1	None
Mode-2	None
Mode-3	None
Mode-2	Mode-1
Mode-2	Mode-3
Mode-3	Mode-2

2.2.7.2 Mode sequence configuration and Sub_Mode insertion

The step mode sequence is defined during the CS configuration and start procedures. Some of the key parameters that decide mode sequencing behavior are shown in Table 4.

Table 4. Mode sequencing parameters

Parameter	Range	Purpose
Mode_0_Steps	1, 2, or 3	Defines the number of consecutive Mode-0 steps that can be executed at the start of every CS subevent.
Main_Mode_Type	1, 2, or 3	Indicates which mode is the primary mode.
Sub_Mode_Type	1, 2, 3, or None	Indicates which mode is the secondary mode.
Min_Main_Mode_Steps	-	Determines the minimum number of main mode steps that must be executed before a submode step.
Max_Main_Mode_Steps	>= Min_Main_Mode_Steps	Determines the maximum number of main mode steps that must be executed before a submode step.

Table 4. Mode sequencing parameters*...continued*

Parameter	Range	Purpose
Main_Mode_Repetition	-	Specifies the number of main mode steps to be repeated in the next subevent.

The number of main mode steps to be executed consecutively before a submode step is selected randomly in a range between Min_Main_Mode_Steps and Max_Main_Mode_Steps, inclusive.

The step mode sequences are not tied to subevent boundaries and can span across more than one subevent. However, as described in Section 2.2.7.1, the subevents must always start with one or more Mode-0 steps.

Figure 6 shows an example of a CS procedure with mode sequencing.

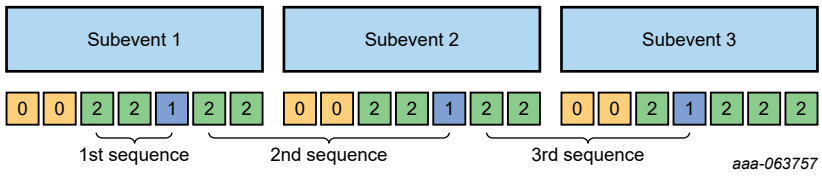

Figure 6. Mode sequencing example

In this example, the following steps are set as follows:

- Main_Mode_Type to 2
- Sub_Mode_Type to 1
- Min_Main_Mode_Steps to 2
- Max_Main_Mode_Steps to 4
- Mode_0_Steps to 2

The first subevent starts with the mandatory mode-0 steps, then two main mode steps, followed by a submode step at the end. The first subevent still has some time remaining. The next sequence is started with a randomly selected four main mode steps. Two of these steps are included in the current subevent, while the next two are executed in the next subevent after the required mode-0 steps, followed by a submode step at the end. This pattern is followed until the number of subevents specified for the CS procedure have been completed.

The final parameter available for mode sequencing is called Main_Mode_Repetition. This parameter specifies various main mode steps to be repeated in the next subevent. When this feature is applied, the steps repeated in the next subevent use the same channel index as the ones in the previous subevent. It ensures that the transmission has the same frequency. The purpose of repeating the main mode steps is to address the possible frequency drift and Doppler shift effects. It makes possible to correlate the properties of these shifts to allow the tracking of moving devices. Steps that are executed due to main mode repetition are not counted toward the mode sequencing process described above.

2.2.8 CS channel map

For typical Bluetooth LE applications, the 2.4 GHz ISM frequency band is divided into 40 channels, each 2 MHz wide. However, this is not the case when using CS. For CS purposes, there are 72 channels, each being 1 MHz wide. The channel map for CS procedures is defined in such a way that the standard Bluetooth LE primary advertising channels are avoided.

The channel width of 1 MHz, rather than 2 MHz is selected specifically for PBR purposes. Using CS tones at adjacent channels using 2 MHz channel width creates a distance ambiguity at 75 meters due to the phase shifting behavior. When using 1 MHz-wide channels, this ambiguity does not arise until 150 meters.

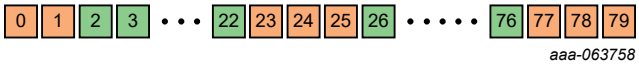

Figure 7. Allowed CS channel indexes

Figure 7 shows a diagram of allowed channels for CS use, with its corresponding indexes. It starts with channel 0, which has a center frequency of 2402 MHz and is not allowed to use due to overlapping with Bluetooth LE Advertising channel 37. Channels from 23 to 25 and from 77 to 79 are also not allowed for CS use, but the rest of the channels are available.

2.2.8.1 Channel Selection Algorithms (CSA)

During a CS procedure, the available channels are not swept linearly, but randomly. The standard defines a set of CSA, which is collectively known as CSA #3 and individually known as CSA #3a, CSA #3b, and CSA #3c. Each of these algorithms is used for different purposes.

The CSA #3a is used exclusively for Mode-0 steps channel selection. This algorithm creates a channel index list with randomly shuffled available channel numbers. When all entries on the shuffled channel list are used, it is regenerated, creating a new randomized list of channels.

The CSA #3b and CSA #3c are used for non-Mode-0 steps. However, only one is used per procedure. The CSA #3b uses the same shuffling method as CSA #3a. It allows for a single channel to be used more than once in a single CS procedure before regenerating. The application layer controls this behavior with a parameter. The CSA #3c uses a whole different method for channel list creation. It groups the subsets of the input channel list to form patterns, named “hat” and “X”.

For more information about each of the CSA, see the Bluetooth Core Specifications.

2.2.9 Antenna switching

Devices supporting CS can include multiple antennas for use during CS tone exchanges in Mode-2 and Mode-3 steps. The maximum number of antennas that a device can have is four. A given pair of antenna configurations, one for each device, provides the number of antenna paths between the two devices.

The Bluetooth Core Specification defines a total of eight antenna permutations, listed in Table 5.

Table 5. Antenna Configurations

Antenna Configuration Index (ACI)	Number of Antennas in Device A	Number of Antennas in Device B	Number of Antenna Paths (N_AP)
0	1	1	1
1	2	1	2
2	3	1	3
3	4	1	4
4	1	2	2
5	1	3	3
6	1	4	4
7	2	2	4

During PBR, a tone is exchanged in each available antenna path. The sequence of paths used is randomized using a Deterministic Random Bit Generator (DRBG) at every CS step.

In theory, a larger number of antenna paths provide better results for distance estimation algorithms based on phase measurements, by lessening the impact of multi-path propagation.

2.3 Host applications and distance measurement algorithms

The Bluetooth stack does not generate distance measurements directly. It is the responsibility of the application-layer code to gather low-level data from the controller like phase correction terms, IQ samples, time-of-flight timestamps. It uses the data as an input for a custom implementation of a distance measurement algorithm.

As the Bluetooth Core Specification does not cover the implementation of distance measurement algorithms, it is up to each vendor to develop the best possible solution for distance estimation algorithms. Superior algorithms produce the superior results.

The fundamentals for each type of distance estimation are covered in the following sections.

3 RTT estimation fundamentals

The RTT is the method for distance estimation. It is previously known as Time of Flight (ToF). It involves the exchange of CS_SYNC packets between two devices. It also involves the precise measurement of the time difference between transmission and reception of these packets, which accounts for radio latencies and processing time overheads.

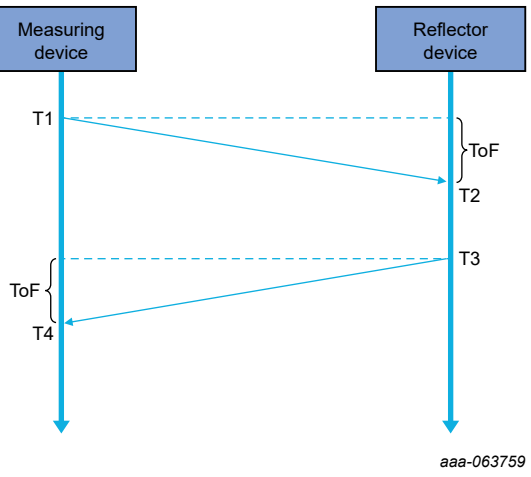

Figure 8. RTT Timestamping Sequence

As it can be observed in Figure 8, the RTT measurement requires two devices that can exchange packets and precisely timestamp the moments. The measuring device (MD) timestamps the initial transmission (T1) when it sends the packet. The RD timestamps the time of reception (T2) when it receives the packet. After processing, the RD timestamp the moment (T3) when it starts the transmission back to the MD. The MD timestamp the moment (T4) when it receives the packet back from the RD.

With these timestamps, a ToF value can be generated according to the following equation:

$$ToF = \frac{\{T4\text{-}T1\}\text{-}\{T3\text{-}T2\}}{2} \quad (1)$$

Once the ToF value is calculated, the distance can be estimated using the equation:

$$Distance_m = ToF * c \quad (2)$$

A precise timestamping mechanism is required on both ends of the RTT measurement to achieve a reasonable level of accuracy. As radio waves travel at the speed of light (c), a time difference of only 3.3 nanoseconds already represents a deviation of 1 m in the distance estimation. This imposes requirements on the stability and relative accuracy of the timestamping mechanism in both MD and RD.

The Bluetooth Core Specification defines several anchor points for timestamping incoming CS_SYNC packets. It includes the Access Address field and optional Random Sequence or Sounding Sequence appended to the normal CS_SYNC packet structure. These three methods offer different degrees of accuracy, security, and latency for application developers. The last two optional methods offer the best accuracy at the cost of extra implementation complexity.

For more details about the definition of the timestamping mechanisms, see the Blueetooth Core Specifications.

4 PBR estimation fundamentals

The PBR is a technique that uses the phase difference in the tones exchanged between the MD and RD. This method is also known as Round-Trip Phase (RTP) and previously known as Phase Distance Estimation (PDE).

In its simplest form, the MD transmits a continuous wave to the RD. The RD notes its phase and sends it back at the same frequency to the MD. The MD then compares the phase of the received signal with the phase of the original signal. It obtains a phase difference used in multiple ways to estimate the distance between the devices.

The simplest PBR uses one tone exchange and the physical properties of a radio wave. It multiplies the signal wavelength by the number of times the signal wraps as it goes from the MD to the RD and back. With this information, the distance can be calculated using the following equation:

$$Distance_m = \frac{\lambda}{2} * \left(\frac{\varphi}{2\pi} + n \right) \quad (3)$$

Where:

- φ is the phase difference measured by MD
- n is the number of wraps
- λ is the wavelength of the carrier frequency, also defined as:

$$\lambda = \frac{c}{f} \quad (4)$$

If the distance to be measured is larger than the λ, the phase wraps. It creates a distance ambiguity and must be accounted for in the distance estimation.

A sweep of multiple frequencies can be used to reduce the mentioned distance ambiguity. If the phase difference is obtained at two different frequencies, the following equation can be used to determine the distance:

$$Distance_m = \frac{c}{4\pi} * \left(\frac{\varphi_2 - \varphi_1}{f_2 - f_1} \right) \quad (5)$$

4.1 Slope-based PDE

If a frequency sweep is performed at a fixed distance and a constant spacing between frequencies Δf, the phases obtained at each frequency in the sweep produce a slope.

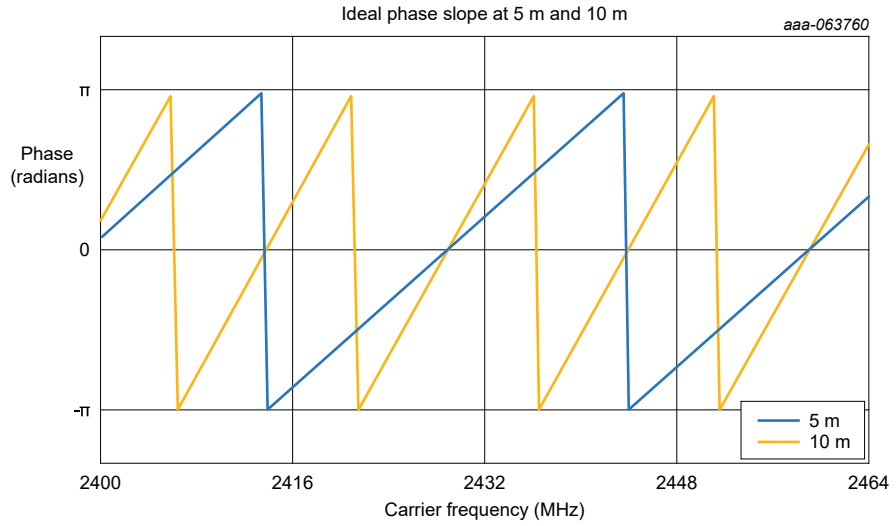

Figure 9. Phase slope at 5 m and 10 m

The relationship between the phase slope and the distance between the devices is linear. The smaller the slope, the shorter the distance. This can be observed in Figure 9.

Using this slope, the distance can be obtained using the following equation:

$$d = \frac{c}{4\pi} * slope \tag{6}$$

If Equation (5) is applied multiple times with multiple pairs of phase differences and frequencies, the following equation is created:

$$d_n = \frac{c}{4\pi} * \left(\frac{\varphi_n - \varphi_{n-1}}{f_n - f_{n-1}} \right) \tag{7}$$

The final slope-based distance estimate can be obtained as follows:

$$d_{slope} = \frac{1}{n-1} \sum_{k=1}^{n-1} d_k \tag{8}$$

4.2 Complex Distance Estimation (CDE)

In CDE, the phase measurements taken at each frequency are converted into a complex signal in the frequency domain. An Inverse Fast Fourier Transform (IFFT) changes the signal into the time domain and the distribution of the propagation delays in the signal can be obtained.

The maximum propagation delay that can be measured without ambiguity is calculated using the frequency spacing, as shown in the following equation:

$$t_{max} = \frac{1}{2 * \Delta f} \tag{9}$$

Therefore, the maximum propagation delay determines the maximum measurable distance that this method can obtain, perform the following:

$$d_{max} = t_{max} * c \tag{10}$$

The distance estimate is obtained using the following equation:

$$d_{cde} = d_{max} * \frac{n_{peak} - 1}{M} \tag{11}$$

Where:

- M is the number of bins used in the IFFT
- n_{peak} is the bin with the highest peak
- d_{max} is the largest unambiguous distance

Using the defined channel frequency spacing in the CS specification of 1 MHz, the maximum measurable distance can be obtained as follows:

$$d_{max} = t_{max} * c = \frac{c}{2*\Delta f} = \frac{299792458\ m/s}{2*1MHz} = 149.896\ m \quad (12)$$

Extract the wrapped phase from a frequency sweep with 1 MHz spacing. Then apply a 128-bin IFFT. This gives an output like the one shown in Figure 10.

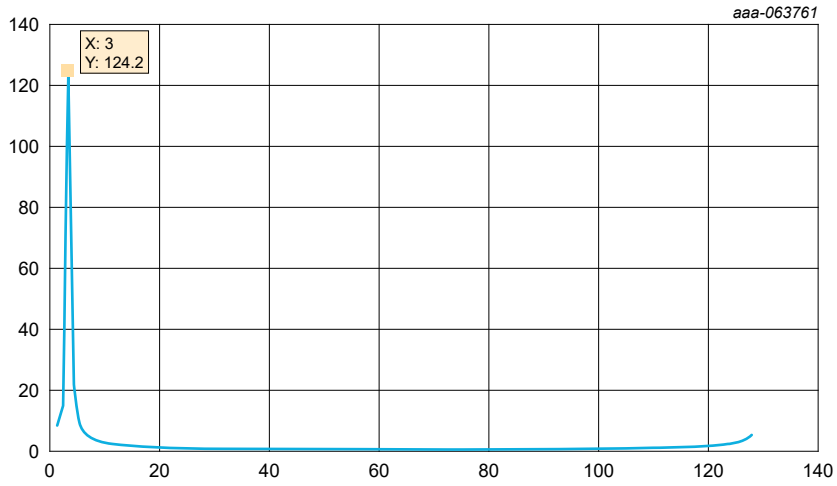

Figure 10. 128-bin IFFT for a phase slope dataset

It is observed that the bin holding the biggest peak is number 3. The estimated distance can be obtained with Equation (9), as shown:

$$d_{cde} = 149.896 * \frac{3-1}{128} = 2.34\ m \quad (13)$$

5 Using NXP KW47/MCX W72 for CS

The KW47/MCX W72 product family is a low-power, highly secure, single-chip wireless MCU. It integrates a high-performance Bluetooth LE radio supporting the new Bluetooth Core Specification version 6.0 that makes it a Bluetooth LE CS compatible device.

The KW47/MCX W72 SoC contains a DSP-V co-processor known as Localization Compute Engine (LCE). It reduces algorithm latency by up to 45 % compared to a Cortex M33 core.

The KW47-LOC/MCXW72-LOC board is a highly configurable, cost-effective evaluation, and development platform for NXP KW47/MCX W72 MCU dedicated for Bluetooth LE CS technology. Several board features are designed directly for the Bluetooth LE CS specification, integrating hardware features tailored for this use case, including:

- A dedicated RF switch with dual input/output paths.
- Two onboard wideband monopole antennas optimized for spatial diversity.
- SMA connectors to connect external antennas or run cable-based measurements for high-precision validation.

These features allow the KW47-LOC/MCXW72-LOC to perform both radiated and conducted CS tests, making it suitable for advanced localization scenarios and RF performance benchmarking.

The KW47 Software Development Kit (SDK) offers RADE, an NXP proprietary high accuracy localization algorithm. It is designed to remove noise and reflection from incoming RF signals to provide the best results for ranging applications.

Table 6 shows the CS-specific capabilities that the KW47/MCX W72 MCU supports.

Table 6. KW47/MCX W72 supported CS capabilities

Capability	Symbol	Supported
CS step modes	-	Mode-0, Mode-1, Mode-2, and Mode-3 steps
Number of antenna paths	N_AP	1, 2, 3, and 4
Antenna configurations	-	1x1, 1x2, 1x3, 1x4, 2x1, 3x1, 4x1, 2x2
Time allocated to swap channel	T_FCS	50, 80, 150 µs
Interlude time between RTT packets	T_IP1	40, 80, 145 µs
Interlude time between CS tones	T_IP2	40, 80, 145 µs
Phase measuring time	T_PM	20, 40 µs
Time for power amplifier ramp-down	T_RD	5 µs
Guard-time between tone and packet transmission	T_GD	10 µs
Antenna switching time	T_SW	2, 4, 10 µs
Frequency measurement time	T_FM	80 µs
CS_SYNC packet duration	T_SY	44 (1 Mbit PHY and 26 (2 Mbit PHY)
CS_SYNC packet payload	RTT_TYPE	32, 64, 96, 128-bit random sequence
Channel selection algorithms	CSA	#3a, #3b, #3c
Tone quality indicator	TQI	Low and high
Normalized attack detection metric	NADM	Supported

For more details, see the *KW47 Reference Manual* (document KW47RM), KW47 Product Family Data Sheet, *MCXW72 Reference Manual* (document MCXW72xRM), and MCXW72 Product Family Data Sheet.

6 Acronyms

Table 7 lists the acronyms used in this document.

Table 7. Acronyms

Acronym	Description
CDE	Complex Distance Estimation
CS	Channel Sounding
CSA	Channel Selection Algorithms
DRBG	Deterministic Random Bit Generator
DSP	Digital Signal Processor
FFO	Fractional Frequency Offset
GAP	Generic Access Profile
HCI	Host Controller Interface

Table 7. Acronyms...*continued*

Acronym	Description
IFFT	Inverse Fast Fourier Transform
IoT	Internet of Things
LCE	Localization Compute Engine
LE	Low Energy
LL	Link Layer
MCU	Microcontroller Unit
MD	Measuring Device
PBR	Phase-Based Ranging
PDE	Phase Distance Estimation
PDU	Protocol Data Unit
PHY	Physical Layer
RAP	Ranging Profile
RD	Reflector Device
RF	Radio Frequency
RTP	Round-Trip Phase
RTT	Round-Trip Time
SDK	Software Development Kit
SIG	Special Interest Group
SMA	SubMiniature version A
SoC	System-on-Chip
ToA	Time of Arrival
ToD	Time of Departure
ToF	Time of Flight

7 References

Table 8 lists the references used to supplement this document.

Table 8. Related documentation/resources

Document	Link/how to access
KW47 Reference Manual (document KW47RM)	KW47RM
KW47 Product Family Data Sheet	Contact local FAE or sales representative
MCXW72 Reference Manual (document MCXW72xRM)	MCXW72xRM
MCXW72x Product Family Data Sheet	Contact local FAE or sales representative
Bluetooth Core Specifications	Bluetooth Core Specifications

8 Revision history

Table 9 summarizes the revisions to this document.

Table 9. Revision history

Document ID	Release date	Description
AN14865 v.1.0	12 December 2025	Initial public release

Legal information

Definitions

Draft — A draft status on a document indicates that the content is still under internal review and subject to formal approval, which may result in modifications or additions. NXP Semiconductors does not give any representations or warranties as to the accuracy or completeness of information included in a draft version of a document and shall have no liability for the consequences of use of such information.

Disclaimers

Limited warranty and liability — Information in this document is believed to be accurate and reliable. However, NXP Semiconductors does not give any representations or warranties, expressed or implied, as to the accuracy or completeness of such information and shall have no liability for the consequences of use of such information. NXP Semiconductors takes no responsibility for the content in this document if provided by an information source outside of NXP Semiconductors.

In no event shall NXP Semiconductors be liable for any indirect, incidental, punitive, special or consequential damages (including - without limitation - lost profits, lost savings, business interruption, costs related to the removal or replacement of any products or rework charges) whether or not such damages are based on tort (including negligence), warranty, breach of contract or any other legal theory.

Notwithstanding any damages that customer might incur for any reason whatsoever, NXP Semiconductors' aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms and conditions of commercial sale of NXP Semiconductors.

Right to make changes — NXP Semiconductors reserves the right to make changes to information published in this document, including without limitation specifications and product descriptions, at any time and without notice. This document supersedes and replaces all information supplied prior to the publication hereof.

Suitability for use — NXP Semiconductors products are not designed, authorized or warranted to be suitable for use in life support, life-critical or safety-critical systems or equipment, nor in applications where failure or malfunction of an NXP Semiconductors product can reasonably be expected to result in personal injury, death or severe property or environmental damage. NXP Semiconductors and its suppliers accept no liability for inclusion and/or use of NXP Semiconductors products in such equipment or applications and therefore such inclusion and/or use is at the customer's own risk.

Applications — Applications that are described herein for any of these products are for illustrative purposes only. NXP Semiconductors makes no representation or warranty that such applications will be suitable for the specified use without further testing or modification.

Customers are responsible for the design and operation of their applications and products using NXP Semiconductors products, and NXP Semiconductors accepts no liability for any assistance with applications or customer product design. It is customer's sole responsibility to determine whether the NXP Semiconductors product is suitable and fit for the customer's applications and products planned, as well as for the planned application and use of customer's third party customer(s). Customers should provide appropriate design and operating safeguards to minimize the risks associated with their applications and products.

NXP Semiconductors does not accept any liability related to any default, damage, costs or problem which is based on any weakness or default in the customer's applications or products, or the application or use by customer's third party customer(s). Customer is responsible for doing all necessary testing for the customer's applications and products using NXP Semiconductors products in order to avoid a default of the applications and the products or of the application or use by customer's third party customer(s). NXP does not accept any liability in this respect.

Terms and conditions of commercial sale — NXP Semiconductors products are sold subject to the general terms and conditions of commercial sale, as published at https://www.nxp.com/profile/terms, unless otherwise agreed in a valid written individual agreement. In case an individual agreement is concluded only the terms and conditions of the respective agreement shall apply. NXP Semiconductors hereby expressly objects to applying the customer's general terms and conditions with regard to the purchase of NXP Semiconductors products by customer.

Export control — This document as well as the item(s) described herein may be subject to export control regulations. Export might require a prior authorization from competent authorities.

Suitability for use in non-automotive qualified products — Unless this document expressly states that this specific NXP Semiconductors product is automotive qualified, the product is not suitable for automotive use. It is neither qualified nor tested in accordance with automotive testing or application requirements. NXP Semiconductors accepts no liability for inclusion and/or use of non-automotive qualified products in automotive equipment or applications.

In the event that customer uses the product for design-in and use in automotive applications to automotive specifications and standards, customer (a) shall use the product without NXP Semiconductors' warranty of the product for such automotive applications, use and specifications, and (b) whenever customer uses the product for automotive applications beyond NXP Semiconductors' specifications such use shall be solely at customer's own risk, and (c) customer fully indemnifies NXP Semiconductors for any liability, damages or failed product claims resulting from customer design and use of the product for automotive applications beyond NXP Semiconductors' standard warranty and NXP Semiconductors' product specifications.

HTML publications — An HTML version, if available, of this document is provided as a courtesy. Definitive information is contained in the applicable document in PDF format. If there is a discrepancy between the HTML document and the PDF document, the PDF document has priority.

Translations — A non-English (translated) version of a document, including the legal information in that document, is for reference only. The English version shall prevail in case of any discrepancy between the translated and English versions.

Security — Customer understands that all NXP products may be subject to unidentified vulnerabilities or may support established security standards or specifications with known limitations. Customer is responsible for the design and operation of its applications and products throughout their lifecycles to reduce the effect of these vulnerabilities on customer's applications and products. Customer's responsibility also extends to other open and/or proprietary technologies supported by NXP products for use in customer's applications. NXP accepts no liability for any vulnerability. Customer should regularly check security updates from NXP and follow up appropriately.

Customer shall select products with security features that best meet rules, regulations, and standards of the intended application and make the ultimate design decisions regarding its products and is solely responsible for compliance with all legal, regulatory, and security related requirements concerning its products, regardless of any information or support that may be provided by NXP.

NXP has a Product Security Incident Response Team (PSIRT) (reachable at PSIRT@nxp.com) that manages the investigation, reporting, and solution release to security vulnerabilities of NXP products.

NXP B.V. — NXP B.V. is not an operating company and it does not distribute or sell products.

Trademarks

Notice: All referenced brands, product names, service names, and trademarks are the property of their respective owners.

NXP — wordmark and logo are trademarks of NXP B.V.

Contents