Atlas

How Sigfox computes geolocation information based on the Network Location service

Technical Paper

External

Version history

Version	Description of the change	Date
1.0	Document publication	September 1st 2019

Authors

Sigfox

Diffusion list

Name/Team	Company - Function
All	**External**

Table of contents

1. Glossary

Term	Description
Atlas Native	Geolocation service offer providing the geographic coordinates (latitude, longitude and a radius) of an activated device. Atlas Native is an offer solely based on Sigfox Network topology, without any additional positioning information from the device.
Activated device	A device is considered as activated if it has an active token, thus communicating and sending messages.
Estimated position	The estimated position of a device is the position computed through the Network location service. It is an estimation of the location of the device whereas the true position is a location computed with a precise positioning system, e.g. GPS, Bluetooth or WiFi.
Geolocated message	A Sigfox message for which we know both the metadata and the true position.
Metadata	Metadata is defined as the data providing information about one or more aspects of the data; it is used to summarize basic information about data which can make monitoring and processing with specific data easier. For a Sigfox message, it includes all the data generated during the message reception, i.e. base stations receiving the message, RSSI and so on.
Network location	Technical service used by Atlas Native, computing and delivering the geographic coordinates (latitude, longitude) and radius from a device through a callback or an API interface for each message, using message's metadata and network topology.
Positioning	Mechanism computing the position (latitude, longitude) of a device.
Radius	A position is always given with a confidence area. The radius is used to determine the precision circle, centered at the estimated position. The radius is provided for a fixed probability, i.e. 80% of probability that a device is within the x-meter radius.

RSSI	Received Signal Strength Indication: measurement of the radio signal reception power in dBm for a given base station.
RSSI vector	For a Sigfox message it refers to the list of RSSI of all the base stations that received the same message.
True error	Distance between the true position and the estimated position.
True position	Exact position of the device when sending the message. This position shall be provided by a precise geolocation source. What is called a precise geolocation source is a source with an accuracy significantly more precise than the network geolocation accuracy, i.e. GPS location or Wi-Fi location.

2. Introduction

Launched in January 2017, Network Location is a service providing the approximate position of a device in the format of geographic coordinates (longitude and latitude) and radius.

This proprietary algorithm leverages the RSSI (Received Signal Strength Indicator) to calculate the most probable location of a device. This RSSI is measured for every message that is received over the Sigfox 0G network.

To benefit from the Network Location service, a device only needs to be Sigfox Ready and capable of sending messages over the Sigfox 0G network. There is no requirement for GPS or other hardware/software components.

The Sigfox Atlas Suite is a series of offers, which are all based on Network Location. While the Atlas Native offer is solely based on the Network Location service, Atlas WiFi offer combines WiFi location with Network Location to deliver enhanced positioning for a device.

Before diving into the document, here is a general overview of the Network based geolocation, widely used in wireless communication systems like 3G, 4G, Wi-Fi and so on.

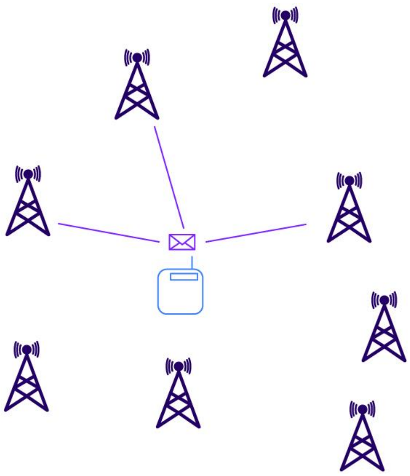

- A device sends a message
- This message is received by several Base Stations (3 in this example)
- The signal strength is measured by each of the Base Stations
- The position of the device is estimated by using the strength of the message signal as recorded by each of the Base Stations

Figure 1: Network Based Geolocation principle

Starting with the introduction of the positioning algorithm and the radius computation, the document progresses with the explanation of machine learning techniques and input data which is used to improve geolocation accuracy.

To complete this document, a Q&A section provides some further insights about the service.

3. Sigfox positioning calculation

In this section, the positioning algorithm is explained, describing the method for calculating the position (latitude, longitude) for a device.

3.1. Algorithm

Everything begins with a device sending a message through the Sigfox network. As reflected on the example below, three base stations are receiving the same device's message. The Received Signal Strength is measured for each base station.

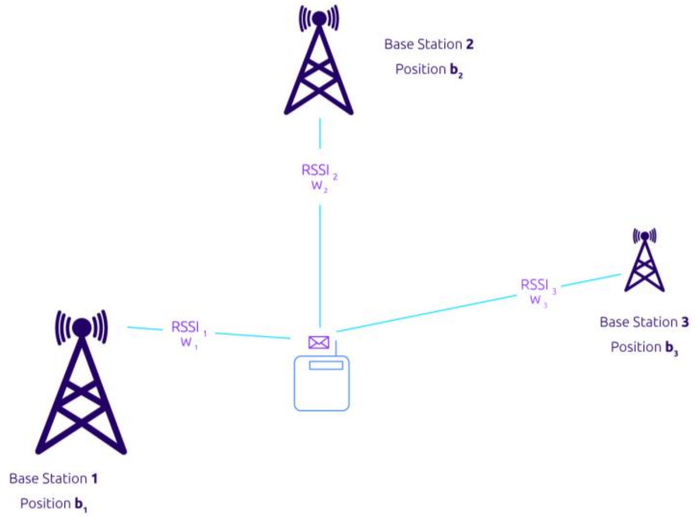

Figure 2: Example of a device message received by three base stations. Each base station has a position b and is receiving the device's message with a RSSI and computed weight w.

The information received through a message from the device will be used as input into the positioning algorithm. The metadata included in a message when received, e.g. the RSSI vector, is used to calculate the most probable point, $\hat{z}$.
$\hat{z}$ is a vector with z as the position, i.e. latitude, longitude. We can write $\hat{z}$ as:

$$\hat{z} = \sum_{i=1}^{i=N_{BS}} w_i b_i$$

Where:

- i is one of the base stations receiving the message
- b_i is the position of the ith base station receiving the message
- w_i is the weight associated to the metadata of the ith base station which received the message
- N_{BS} is the number of base stations which received the message

The value of the base station weight is given by:

$$w_i = 10^{RSSI_{dB}(i)/10}$$

Where $RSSI_{dB}(i)$ is the RSSI (in dBm) of the i^{th} base station.

3.2. Exceptions

When only 1 base station receives the message from a device, the estimated position for the device is identical to the location of the base station:

$$\hat{z} = b$$

Where *b* is the base station position.

In these situations where only 1 base station is used to calculate the position, some random noise is added to the output to obfuscate the exact location of that base station:

$$\hat{z} = b + n$$

Where n is a noise composed by two components:

- A bias that depends only on the base station coordinates
  Amplitude: 0.001 degree

- A random gaussian noise that varies for each message
  Standard deviation of 0.001 degree, saturated at 0.05 degree

In cases where the difference in signal strength between the two strongest RSSIs is equal to or larger than 20 dB, only the message received by the base station with the strongest signal will be used in the algorithm.

4. Sigfox radius calculation

The Network Location algorithm provides the predicted position of a device together with the radius. The radius defines an area around the predicted position, where the device should be located within a certain level of confidence. The typical radius

We can see in the figure below that the estimated radius, i.e. the typical radius, is larger than the true error to ensure that the true position of the device falls within the indicated area.

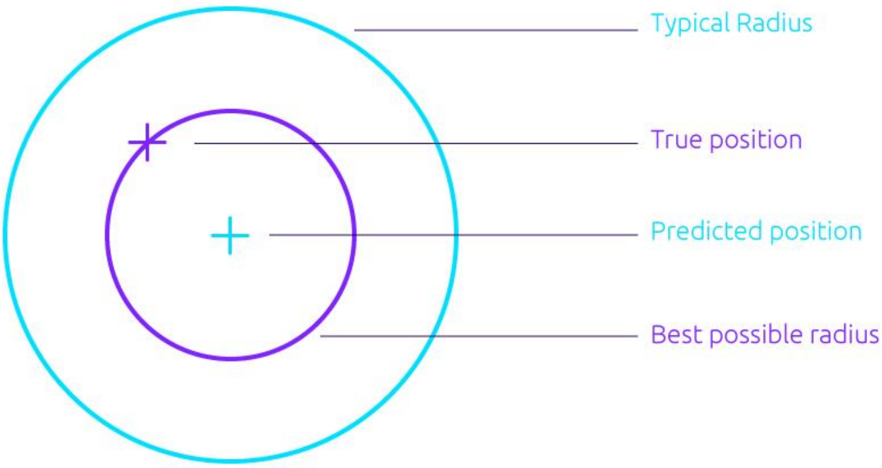

When computing a radius for the estimated location, the target level of confidence is set to 80% - 90%.

Figure 3: Illustration of the true position of the device sending a message (GPS or WI-FI location), the predicted position using Network Location, the overall estimated radius, i.e. typical radius, and the best possible radius.

4.1. Algorithm

The radius estimation algorithm aims to compute the typical radius the closest possible of the best possible radius. This algorithm has two primary contributors:

- Message metadata
  The estimation for the radius considers the signal strength of the message received (RSSI).

- Sigfox 0G network topology
  The estimation for the radius takes into account the Sigfox 0G network characteristics. These characteristics are independent of the message metadata and include base stations density, grey and black-lists.

Message metadata are contributing the most to compute the device location, trusting mostly the highest RSSIs, i.e. from -110dBm, since it’s highly likely reflecting a close position of a device to a base station. Sigfox 0G network characteristics-based radius estimation is used to refine the overall estimated radius.

4.1.1. Metadata contribution

Geolocated messages are used to model the radius as a function of the RSSI. After a series of statistical analysis, the metadata contribution has been incorporated into the radius calculation model.

4.1.2. Network contribution

To complete the estimation of the radius a second estimation method has been developed based on the network characteristics and independent of the RSSI.

The estimation is performed following these steps:

1. Predict the geolocation using Network Location algorithm for geolocated messages

2. Compute the true error between the Network Location position and the true position (GPS or WI-FI). The true error is the target for the machine learning model (step 4)

3. The objective is to identify dimensions to be included in the algorithm which will reduce the true error.

   In Sigfox's implementation, two dimensions are currently used:

   - The distance between the predicted location and the closest 3 base stations

   - The network density in the vicinity of the predicted location. Currently, the density is measured by number of base stations within a fixed perimeter.

   Both dimensions are not used to estimate the true position of a device but only to estimate the error induced, i.e. the radius, during the computation of the estimated location.

4. Train a machine learning model using the features and target computed on geolocated messages to predict a quantile of the error

   The machine learning model is directly tied to the required level of confidence, i.e. currently a 90%-target. A model has been selected which provides a quantile of the true error for easy optimization.

5. Discretize the world on a grid

   A fixed grid size with WGS84 coordinates (latitude and longitude) has been implemented. The grid path is set to 0.02 degree, which corresponds approximately to 2.2 km.

6. Compute the dimensions (step 3) for each grid point

7. Determine the network characteristics-based radius estimation as the predicted radius at each point of the grid using the trained model

Once the radius has been estimated, a business rule is applied to provide an upper and lower threshold. In the current implementation, the minimum and maximum values are set respectively to 50 m and 30 km.

The network contribution is updated once every 2 months, or on demand, i.e. through Sigfox support, Sigfox operators or product management.
The automation of the generation and update of the network model is scheduled for 2019: every other week is targeted for an update.

4.1.3. Additional business rules

Besides the two dimensions previously mentioned, the radius estimation, once computed, is limited with an upper threshold.
This upper bound is directly tied to the number of base stations which received the device message, and whether or not they are included on a grey list. Grey and black lists will be introduced at a later stage.

The following rules are applied:

- The radius is set to a maximum default value, i.e. 30 km, if the location for a device is predicted using only grey listed base stations

- The radius is multiplied by a factor 2 for “orphan” messages.
  A message is considered as “orphan”: when received by only 1 base station.

4.2. Positioning performance

A precision better than 10 km in 80% of the cases is the targeted performance for Network Location: it means that for 80% of the messages, Network Location manages to estimate a position with an error, i.e. the distance between the estimated position and the true location, lower than 10 km.

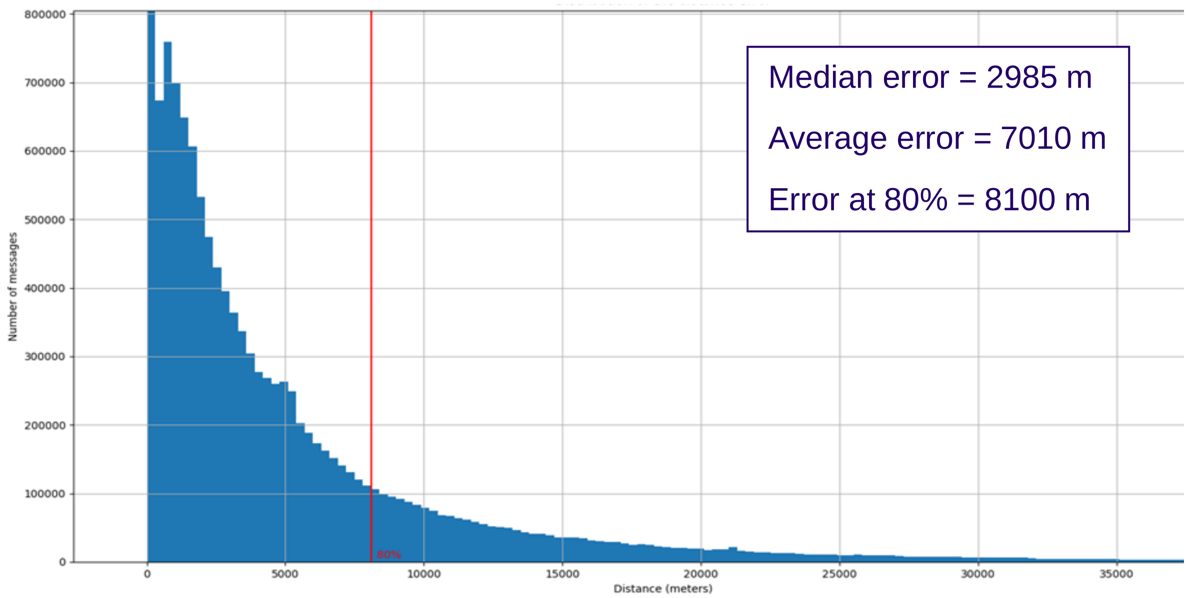

Figure 4: Distribution of the distance error. The error at 80% is lower than the 10 km target.

The figure above shows the performance of Network Location service measured in the field. Geolocated messages are used to compare the estimated location and the GPS location. The distance error distribution is computed with 13 million of geolocated messages over 2018. (c.f. 6.1 for geographical distribution)

4.3. Radius performance

To assess the performances of the radius estimation, the first KPI of interest is the cumulative distribution of the ratio error / radius (figure 5).

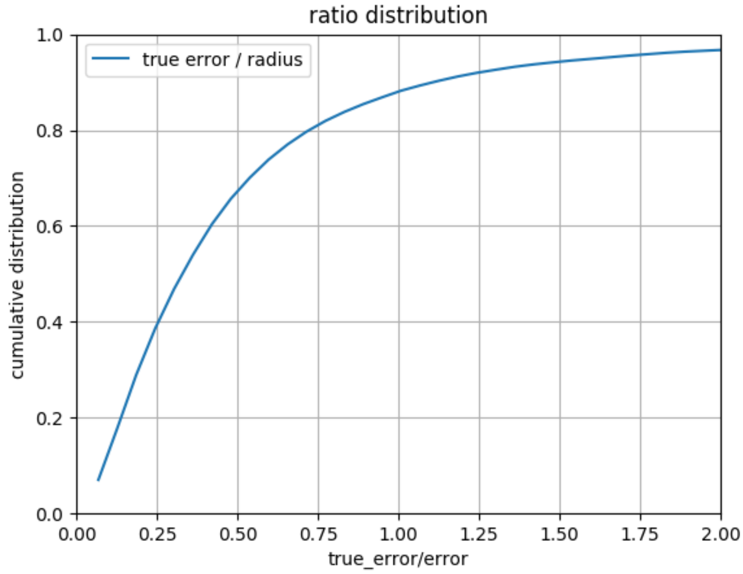

Figure 5: Cumulative probability distribution of the ratio between the true error and the estimated radius using production parameters.

This result has been computed using 225 345 geolocated messages during a period of 5 days.

In particular, when this cumulative probability is equal to 1, it corresponds to the rate of messages for which the radius is larger or equal to the true error, i.e. includes the true error, as drawn in Figure 3.
This value represents the radius confidence level, showing that our radius is larger than the error for 90% of messages.

Figure 6: Cumulative probability distribution of the true error and the estimated radius using the production's parameters. This result has been computed using 225 345 geolocated messages during a period of 5 days.

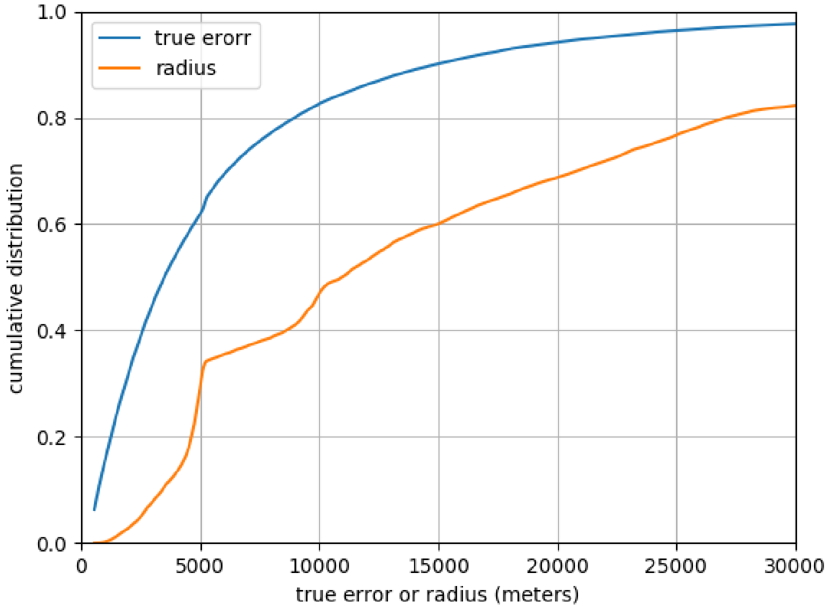

In the figure right above, the distribution of the radius is compared to the one of the true error. It highlights the over-estimation on the radius when setting a confidence level to 90%. Note that this discrepancy between the true error and the estimated radius can be reduced if the constraint associated with the radius confidence level is relaxed.

5. Base station position

The base station location (latitude, longitude) represents crucial input for the geolocation calculation. Obviously, if the base station location is not accurate then device positioning calculation will render incorrect results.
That's why, it is critical that the accurate location is declared during the commissioning of the base station. Furthermore, the location information must be updated whenever the base station is moved.

5.1. Base station eligibility

To determine the most probable location for a device, the network location algorithm uses the signal strength of a message received by a base station with the known location for that same base station. However, in some cases, either the signal strength or the base station location may be dubious. There may be several reasons for this such as base station mobility, base station lifecycle status, urban interference, or geological constraints.

To ensure the most accurate positioning, those base stations must be excluded from the algorithm. This is achieved by:

- Manually specifying whether a base station should be included or excluded from the location algorithm

- Automatically generating listings with base stations which are not reliable enough.

5.1.1. Contribution setting

By default, to ensure accurate positioning calculation, only base stations that have been completely commissioned should be considered by the network location algorithm. However, in some circumstances, this standard guideline may need to be overruled. To accommodate these exceptions, a specific attribute has been implemented: "Base station contribution setting". This attribute provides 3 possible configurations for a base station:

- **Default** – include the base station in the network location algorithm depending on its life-cycle status

- **Enabled** – include the base station in the location algorithm regardless of its life-cycle status

- **Disabled** - exclude the base station in the location algorithm regardless of its life-cycle status

This base station attribute can only be modified by Sigfox Support and any modification may take 48 to 72 hours to take effect.

Note: base station grey/black listings supersede the Contribution settings attribute, i.e. even if a base station is set to “Enabled”, if that same base station is also included on the black listing it **will not** be considered in the location algorithm.

5.1.2. Grey listed and black listed base stations

To further improve positioning accuracy, an automated mechanism has been implemented which evaluates base station fitness-for-use through machine learning. If a base station does not meet minimal acceptance criteria, it will be added onto the exclusion list. Depending on the amount of deviation, a base station may be added onto the “grey-list” or onto the “black-list”.

To determine the contribution of a base station to the positioning accuracy, geolocated devices (through GPS, WiFi, etc) are used for input. Geolocated devices provide the true location of the device which is compared to its calculated position. The “geolocation error” represents the difference between the 2 locations. The base station will be added to the grey list or the black list if:

- the geolocation error *exceeds* a certain threshold and/or

- the positioning accuracy *improves* if the base station is *removed* from the location algorithm

A base station will be reinstated if the above situations no longer apply.

- If a base station is added onto the “black-list”, it will never be considered by the network location algorithm (regardless of its Contribution setting).
- If a base station is added onto the “grey-list”, it will only be included in the location algorithm if there are no other base stations available.

5.2. Additional rules and specific cases

In addition to the grey list and the black list, a couple of sanity checks have been introduced:

- When a message from a device is received by multiple base stations, the top 5 base stations with the strongest signal strength will be used by the location algorithm

- When a message from a device is received by more than 3 base stations, any remote base stations are excluded from the location algorithm. For example, base station A and B are 1 kilometre apart and base station C is 100 kilometres apart then base station C will be discarded.

- When there is a concern with the declared location for a base station, then no positioning will be provided for the device. For example, this would be the case if a message is received by 2 base stations which are supposedly more than 220 kilometres apart.

Note: considering a remote base station as an outlier and ignoring it for the position prediction does not mean the base station is malfunctioning. It means that if the base station would be included in the positioning calculation, it would decrease the accuracy prediction for this device.

Lastly, in some specific cases, Sigfox network location algorithm will not provide any position:

- If more than 100 base stations have received the message.
  This rule has been implemented to prevent a location prediction when a device is sending messages from an airplane.

6. Geolocated messages

Aiming better performances for Network location requires a significant amount of data to refine the position of a device through sequences of machine learning.

Feeding the algorithm with geolocated data means that we can learn for the geolocated positions, devices behaviour, associated RSSI signatures and network information.
This way, when another device requests a location through the Network location service, we are able to provide a more refine location that we used to, reducing the radius based on the RSSI signatures of devices with precise positions.

As of today, geolocated messages come from two main sources: GPS and Wi-Fi devices.

6.1. GPS devices

A growing number of Sigfox device makers create tracking devices including an embedded GPS module. These devices include their current latitude and longitude in the payload before sending it through the Sigfox protocol. The accuracy of a GPS system is a few tens of meters, which is useful for many applications, and it works at almost any location.

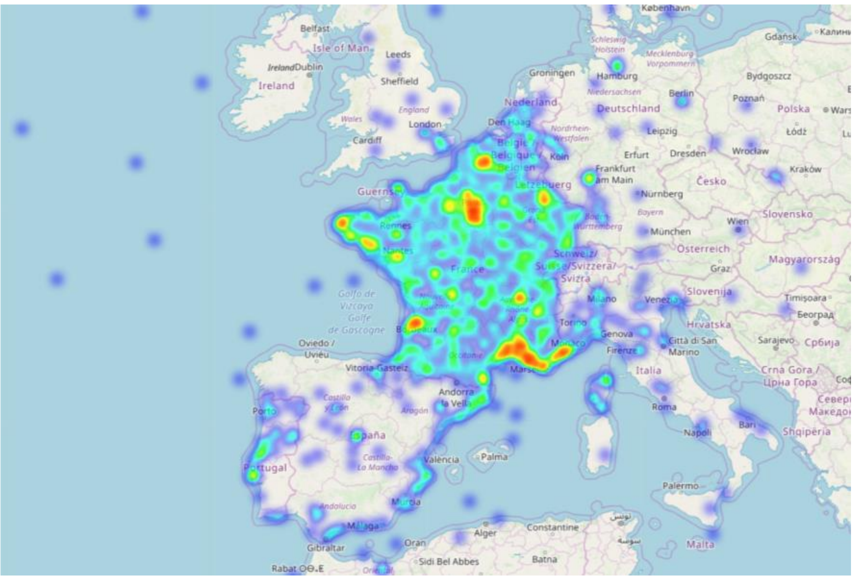

Figure 7: An example of GPS devices heatmap

Figure 6 shows the geographical distribution of the GPS messages used for the tests and learning phases.

6.2. Wi-Fi devices

Wi-Fi location positioning is a Sigfox service relying on devices with a Wi-Fi module. These devices collecting MAC addresses from Wi-Fi network equipment through the Wi-Fi protocol and send them in the message payload. Once a message reaches the Sigfox Cloud, we are asking from internal or external databases to match the MAC addresses to coordinates, in order to retrieve a position.

Wi-Fi Location works only in areas where Wi-Fi infrastructure is available, with an accuracy of 100 m for 80% of the messages.

6.3. Collecting geolocated data

Both GPS and Wi-Fi geolocated messages are used to:

- Test the performance of geolocation algorithms
- Create models with supervised learning techniques, e.g. POI Location
- Verify the consistency of base station locations

There are two ways of collecting geolocated messages:

- Decode the message payloads for GPS devices and extract latitude, longitude and other relevant data.
  It implies that device makers provide payloads layouts to allow Sigfox to decode them. This solution isn't scalable and add a complexity and maintenance to the equation.

- The device makers send the latitudes, longitudes, and metadata of their messages through a web API back to Sigfox. This is the recommended approach.
  Device makers oversee the decoding of their payloads but can set up a following callback or any other mechanism to send the metadata of their messages and their locations through the API.

7. Q&A

Since Network Location service was deployed (January 2017), many questions were raised by customers. In this section, we will answer some of most frequently asked questions.

If your question is not answered below, please go to the Sigfox support page.

Questions summary

Why isn't my message geolocated?

Here are the main reasons why a message is not geolocated:

- Most of the cases, none of the base stations receiving the message are used for geolocation (c.f. 5. Base station position). The possible causes are:
  - Site status is not allowed (5.1)
  - Base station contribution setting is disabled (5.1.1)
  - Base station is blacklisted (5.1.2)
- Special reception conditions, e.g. if the message is received by a large number of base stations, i.e. > 200, then it will not be geolocated (high altitude devices)
- Data infrastructure failure: A public cloud solution is used to host geolocation algorithms; in some rare cases a failure can happen. In that case a replay mechanism in implemented, making an absence of response extremely rare (The public cloud shall be down for a continuous and long period of time, which never happened)
- Type of received message: ACK messages are not geolocated

What precision can I expect?

The target precision is 10 km in 80% of the cases i.e.80 % of the messages will be geolocated within a distance less than 10 km from the exact location. This is a SLO given for static devices.

The global measured precision is given in 4.2.

To go more into details, the precision will depend mostly on 2 parameters:

- The base station density in the area of the device
- RSSI level

The amount of geolocated data in the device area can also slightly improve the precision.

Why is the distance between Network Location and the true point so large?

Several reasons can induce large distances:

- **Long range**
  Sigfox signal can be received by base station relatively far from the device. In fact, some partners recorded a distance higher than 1000 km, that's why in some cases it is difficult to precisely estimate the device position.

- **Mobility**
  In those conditions, geolocation algorithm performances are significantly degraded.

- **Reception conditions**
  Low RSSI and low number of received base stations impact the location computation.

- **Device is on the sea/ocean or close to the coast**
  Signal propagation over the sea or the ocean is very good, leading to base station reception quite far from the device. In practice, geolocation performances on a waterfront are significantly degraded.

- **Devices at altitude**
  In that case, a lot of base stations can receive the message, some being in line of sight conditions, leading to a misinterpretation of the metadata collected.

- **Base stations at altitude**

These base stations will capture signals emitted by devices far from them. That will increase the average error, that’s why this kind of base stations are often in grey lists, but when no geolocated data is available, these base stations are used.

Why the radius is so large?

In some cases, the radius seems very large compared to the true precision of the algorithm, the main reasons of this are the following:

- Base station density
  Even with a high number of base stations receiving the message, it can happen that the estimated position is in an area where the base station density is low, for example when the highest RSSI is received by a base station in a low dense area.

- Radius network contribution is not up to date
  It happens when some base stations are recently installed and are not considered for the network contribution calculation.

- Highest RSSI base station is not used
  It happens when a message is received by several base stations, one of them having a high RSSI (ex: -70 dBm). In that case, the customer expects that the radius is very low. This is true except when this base station is not used for the geolocation.

Why is the true point outside the precision circle?

In that case, the precision circle, in green hereafter: center = estimated location, radius = estimated radius, will not contain the true location of the device.

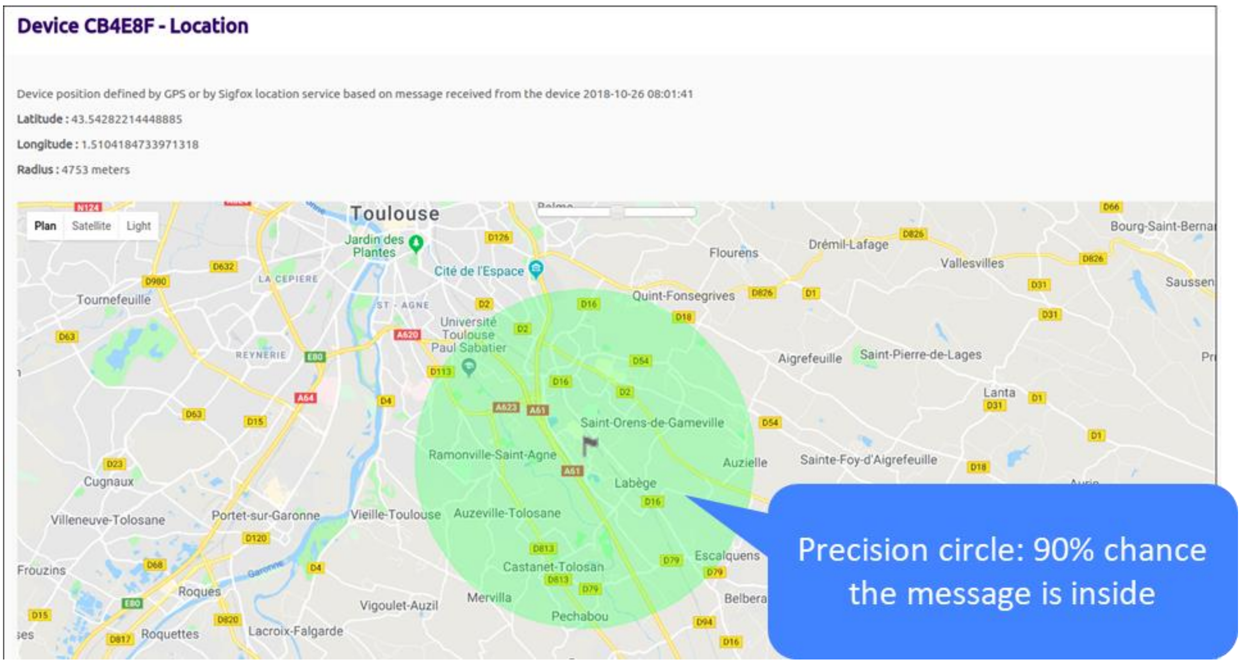

In practice, 90% of the messages will be inside the precision circle. It means that 10% of the messages will be outside the precision circle.

It can be explained with conditions being not favorable, e.g. high altitude, speed > 5 km/h.
One other possible cause is when a message is received by several base stations, one of them having a high RSSI (ex: -70 dBm), others low RSSI. In that case, the customer expects that the position and the radius are accurate. In practice, these cases will be accurate, except when the high RSSI base station is not used for the geolocation.

Why is Network Location positioning moving over the time?

Signal reception can slightly vary across the time. If a base station receives a message at a relatively low level, e.g. -140 dBm, it is possible that few seconds later, the same base station will not receive the message anymore.
In that case, the Network Location position can vary significantly. Let's take the following example:

- A first message is sent by a static device
- The message is received by 2 base stations:
  BS_A at -140 dBm, BS_B at -137 dBm
- 1 minute later, the same device sends another message without moving
- This 2nd message is received only by BS_A at -141 dBm
- The 2 base stations are 20 km apart

Then the position of the 1st and the 2nd message will be several kilometres apart.

Why is the radius larger than 10 km?

Sometimes, there could be a confusion between the radius and the true error. As the true error target is 10 km in 80% of the cases (c.f. 4.2), some can expect that the radius is lower than 10 km. In practice, the radius is not equal to the true error and can exceed 10 km with some conditions (c.f. section 4.1.3)

Why do we see geolocated points in the ocean or in the Mediterranean Sea?

It is possible when a message is received by two base stations located in a coastal area with a strip of sea (or lake) in the middle. Today, there is no mechanism to force the estimated position to be located on-land. Note also that devices can send messages from a boat.

Why is geolocation accuracy poor on coastal areas?

Radio propagation is very good on the water. Therefore, a message sent from a coastal area can be received by base stations very far from the device (often tens of kilometers) with a high RSSI, while some other base stations closer to the device will receive the same message at a lower RSSI.
It is quite difficult to overcome this issue, but this effect can be reduced using learning techniques.