Technical Documentation

0559-APN-M2M-ASTRO

Astronode S
Low Energy Guidelines

Application Note

For feedback or questions: *support@astrocast.com*

Document history

Issue/rev.	Date	Modifications
1/0	26.01.2022	First Release
0/1	05.08.2021	Initial preliminary version

Reference documents

Document name	Document number
Astronode S - Datasheet	0532
Astronode Patch Antenna Datasheet & Integration guide	0534
Astronode S, Application Layer Messages	*docs.astrocast.com*

In this document

This is an information

This is a warning

Table of Content

1 Introduction

This document will help you understand how to integrate the Astronode S in your application from a system and firmware point of view to achieve the power efficiency you need.

2 Astronode S power phases

Before, during, and after sending and receiving data via the satellite constellation, the Astronode S goes through different power modes which are explained hereafter.

2.1 Power on

When power is applied to the Astronode S it starts up for up to 400ms, and then enters sleep mode or deep sleep mode (if configured). It is able to receive commands after 400ms.

2.2 Sleep mode

Most of time the module will be in sleep mode, minimizing power consumption between activities such as transmission cycles or application interface communication. The user may configure the Astronode S to activate deep sleep mode (Sleep mode is set by default). See the CFG_WR command's deep sleep enable bit in the serial communication guide.

2.2.1 Sleep mode

In sleep mode the module is always available via UART for asset interface communication such as queuing new messages.

The module consumes a current between 5µA and 10µA (See Chapter 5.4 for details)

All system data is maintained in this mode, including the message buffer.

2.2.2 Deep sleep mode

If the asset has to interact with the module e.g. for queuing a new message, the WAKEUP pin must be used in this mode.

In this mode the module consumes a current below 500nA (See Chapter 5.4 for details)

Only the RTC, ephemeris information and time of last satellite contact are maintained in this mode. The message queue cannot be stored in deep sleep. As soon as one or multiple messages are queued, the Astronode S remains in sleep mode to maintain the messages, even when configured for deep sleep mode:

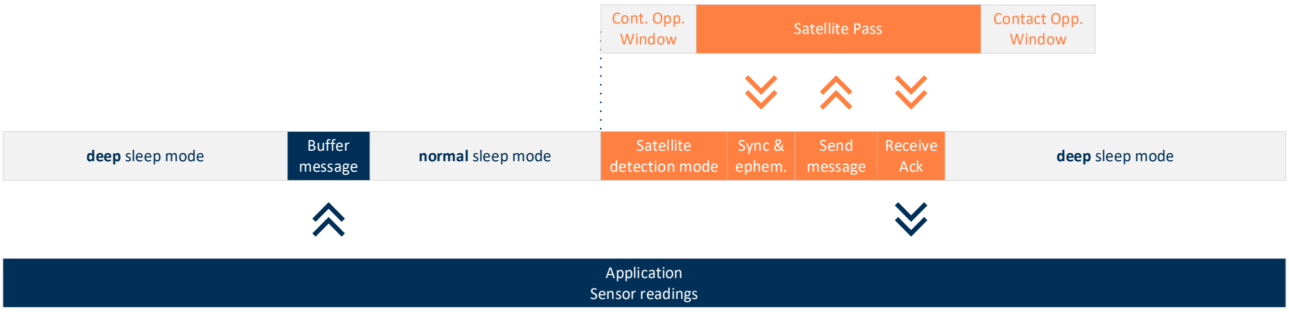

The message queue must be empty before the module will enter deep sleep, otherwise it will use sleep mode. This includes acknowledged messages that are waiting to be cleared by the asset (SAK_RR/SAK_CR).

When it wakes up from deep sleep, the module is effectively starting up as if it were a hardware reset. This requires a 100ms delay between the wakeup pin going high and the first UART communication on the asset interface. The module may sleep when the WAKEUP pin returns to a low state.

2.2.3 Enabling deep sleep mode

The CFG_WR has a bit enabling deep sleep. It is important to note the following points when setting this bit:

- If the wakeup pin is not connected, the only way to wake a device is to reset it
- In sleep mode, the previous RAM based configuration is maintained across wakeups.
- In deep sleep, the Non-Volatile Memory (NVM) based configuration is used on the next wakeup.
- CFG_SR is required to save the RAM configuration to NVM.

This has a consequence that may not be obvious:

- If deep sleep is enabled by writing to RAM, but not saved to NVM, and the module enters deep sleep, on wakeup it will read the configuration from NVM, and deep sleep could be disabled (depending on the previous NVM configuration).

2.3 Satellite detection mode

The Astronode S' satellite detection mode is designed to minimize power consumption while waiting for the next contact opportunity. The module is kept in sleep mode most of time and switches the RF part on for a short receive window every 17.9secs. This period is a system parameter and may be changed by Astrocast without prior notice.

The current consumption for the detection mode is shown Chapter 2.3.1. When searching for the satellite, the peak consumption of 45mA lasts for 50msec per receive window. After successful satellite contact, the receive window length can come down to 25msec.

The search window has about 150msec of power on and power off time without radio activity, see last Figure in Chapter 2.3.1.

As long as the message queue is empty, the module never starts in satellite detection mode and remains in sleep mode.

Once a message is queued, satellite detection mode will start as follows:

Use case	Satellite detection mode start time
Ephemeris is • disabled • not yet known • not valid anymore	Immediate
Ephemeris enabled, known and valid	• On next contact opportunity window • Immediate if within contact opportunity window

For details of ephemeris operation and validity, see Chapter 4.1

2.3.1 Low noise environment

Satellite detection mode over a ~60 sec period, with a detection on the right. Measured in a low noise environment comparable to an application installed outdoors (no other RF equipment close by):

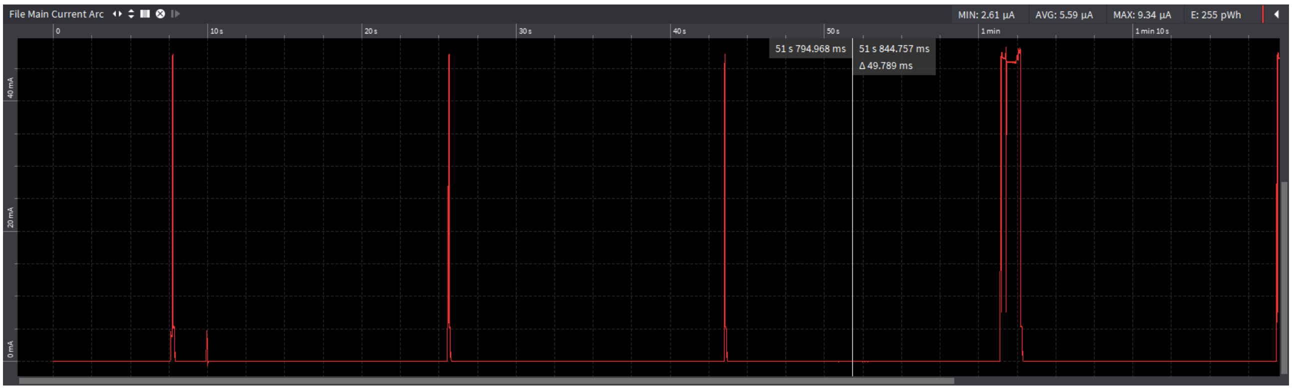

Here is the same visualization, but with a single period selected:

A receive window, zoomed in. This operation repeats every 17.9sec when searching. The ~45mA peak lasts for 50ms here:

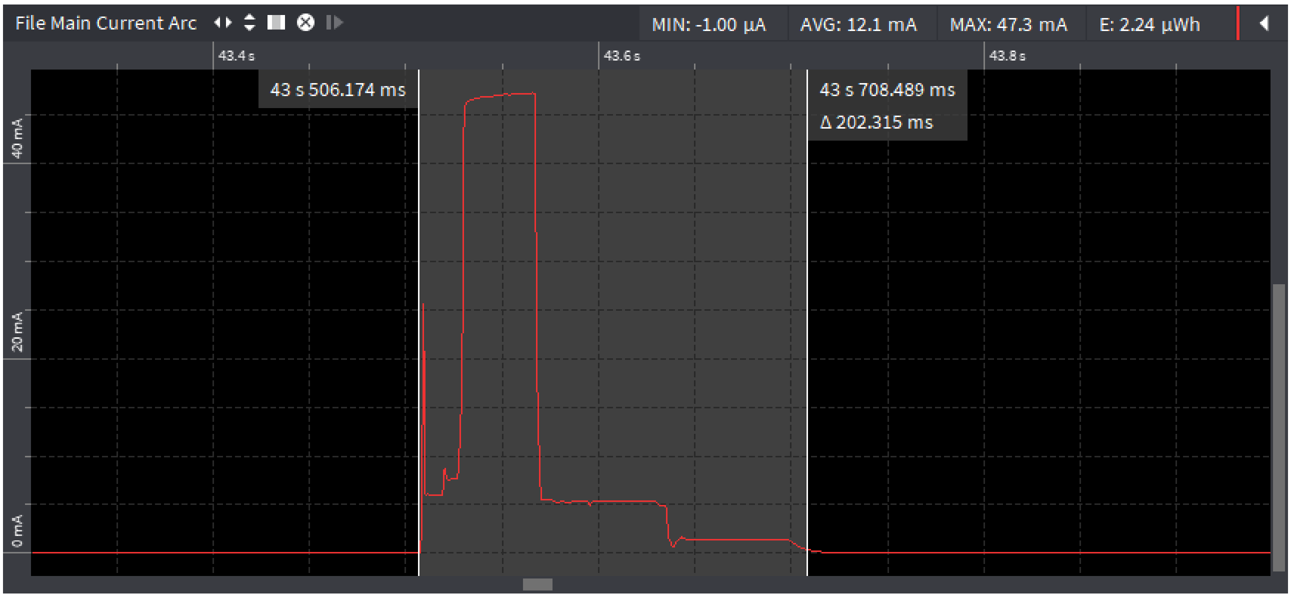

2.3.2 Noisy environment

A receive window, zoomed in, in a noisy environment. The noise present on the receiver triggers two additional failed attempts to demodulate the signal:

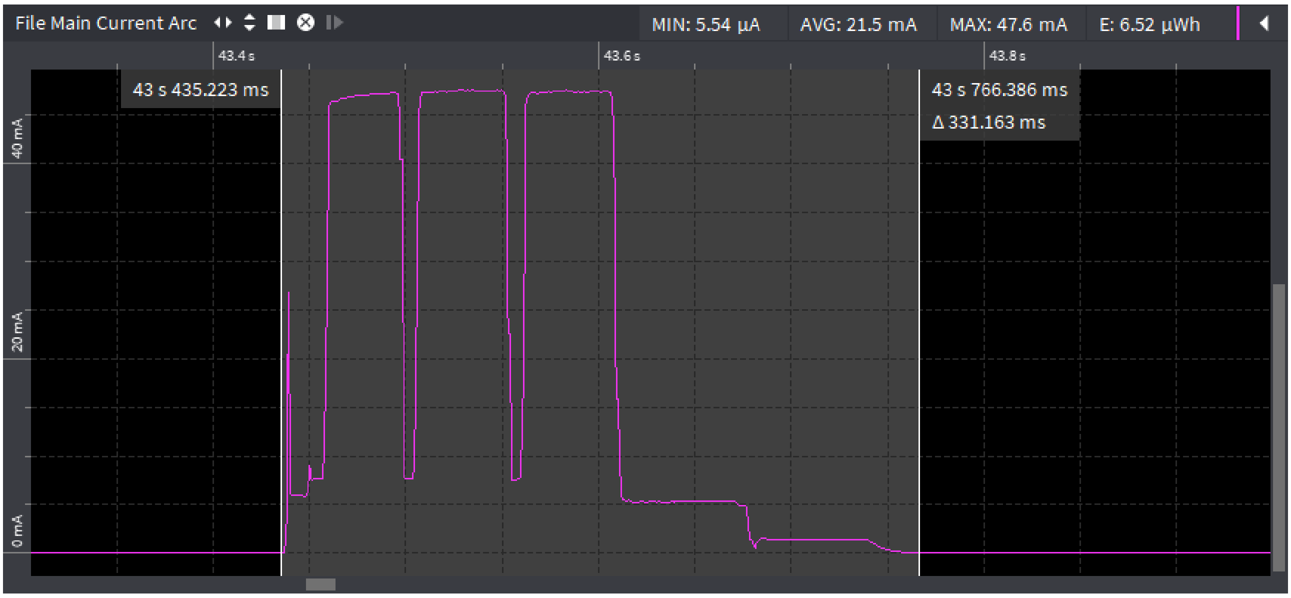

The same visualization, zoomed out, shows the average current for the 17.9s cycle:

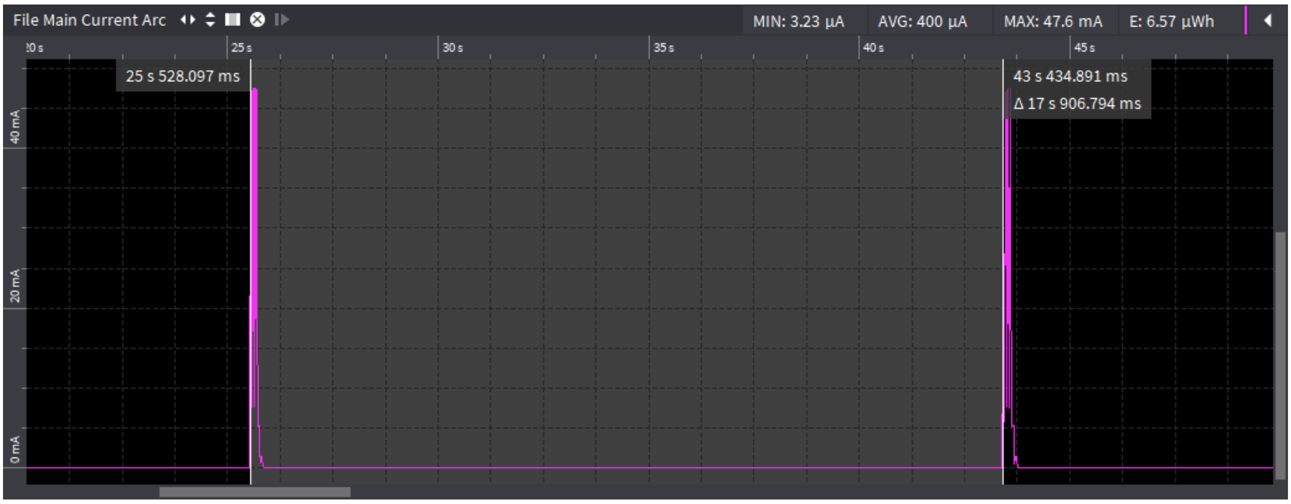

A receive window, zoomed in, in a noisy environment. The noise present on the receiver triggers two additional attempts to demodulate, with one false detection consuming even longer than in the above case:

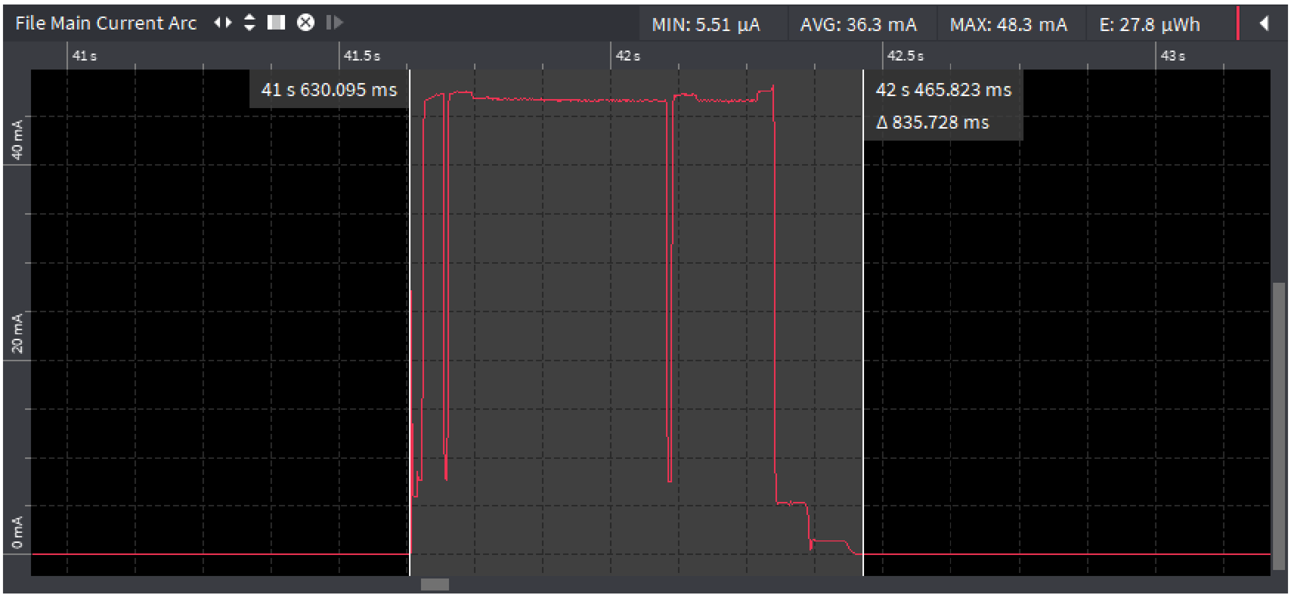

The same visualization zoomed out shows the average current for the 17.9s cycle:

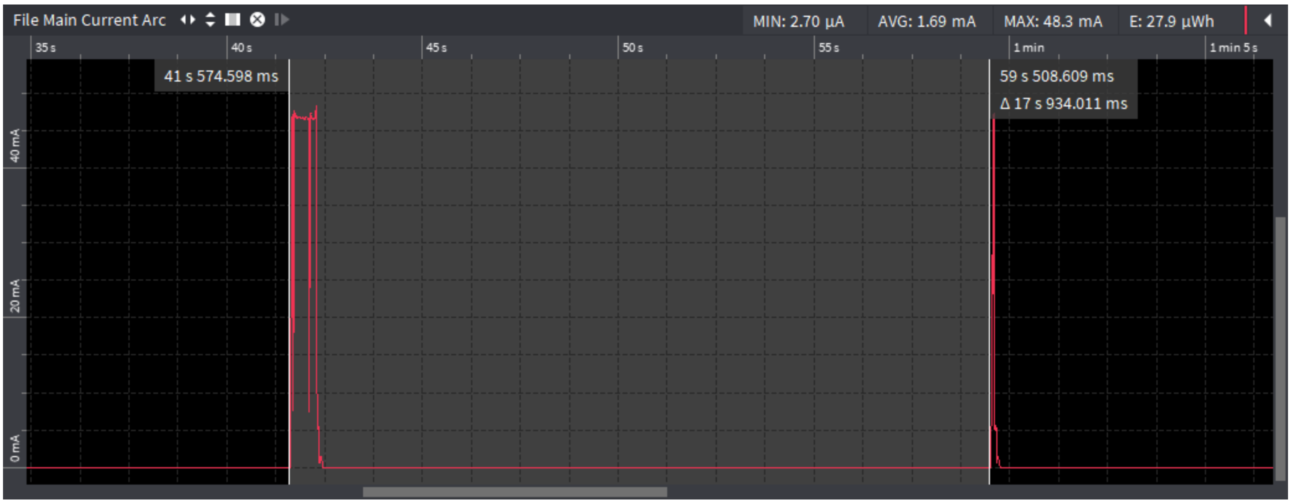

2.4 Sending data

The satellite detection phase attempts to find a satellite above the minimum elevation angle (See Chapter 5.2). Once detected, the transmission and acknowledgment events happen.

2.4.1 Message transmission mechanism

The following happens in the module when sending a queued message:

1. The queued message is split into multiple fragments (See Chapter 5.3 for details)
2. On satellite contact, the module receives signaling for synchronization between the module and the satellite
3. The module then sends a burst of fragments, containing up to 6 individual fragments (The number of fragments contained in a burst is a system parameter and may be changed by Astrocast without prior notice). Transmission of this burst lasts for 1.34 seconds
4. After transmission of the burst, the module listens for the fragment acknowledgements
5. When sending a message with more than 6 fragments, a second phase of signaling and burst transmission takes place after the first one, see Chapter 2.4.6.

This mechanism allows to send a message over multiple passes with limited connectivity. Only un-acknowledged fragments are sent again on a subsequent pass. This may lead to a longer latency but enables connectivity of applications which cannot be installed in optimal conditions.

The peak transmission current is 90mA but averages 80mA for each 1.34 sec burst transmission.

2.4.2 Acknowledgement

Each individual Fragment is acknowledged, but the acknowledgements from one transmission burst are sent grouped after each burst. The acknowledgement confirms the successful reception of the data by the satellite.

Once every individual fragment from a message is acknowledged by the satellite, the module prepares the message acknowledgement for the application, which is signaled via the event pin or has to be polled by the application, depending on the settings (See Acknowledgement Event Discussion for details).

2.4.3 Sending a 20 byte message

The satellite is detected, signaling is received, and the module transmits 2 fragments. It then waits for the acknowledgement.

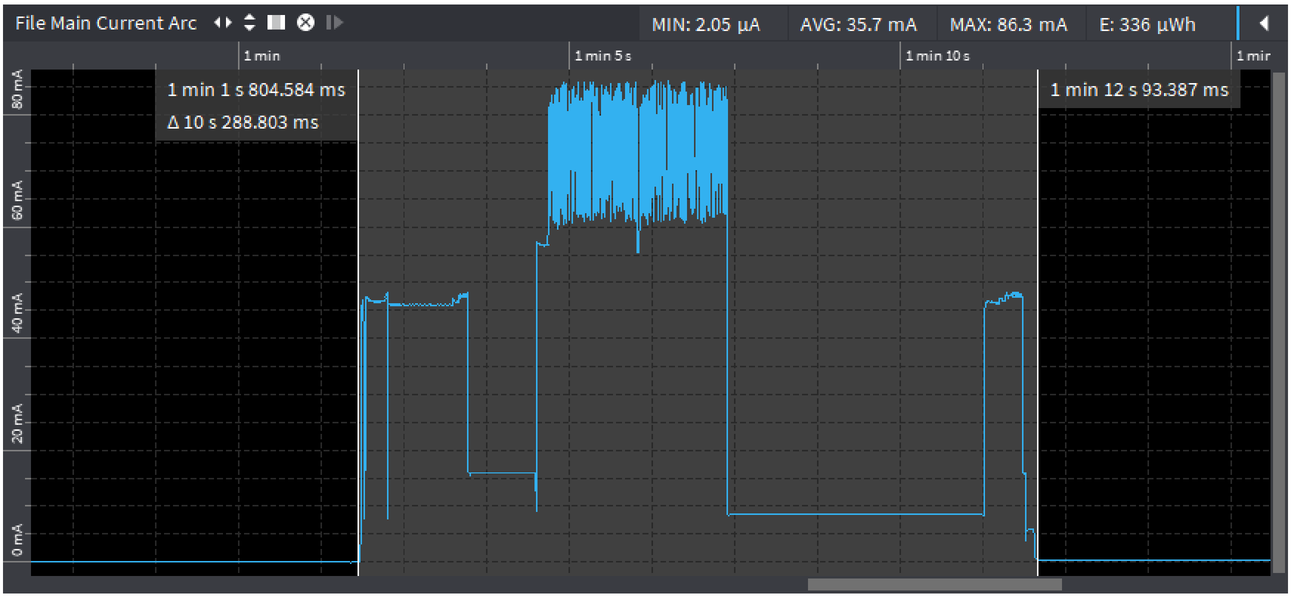

2.4.4 Sending a 30 byte message

The satellite is detected, signaling is received, and the module transmits 3 fragments. It then waits for the acknowledgement.

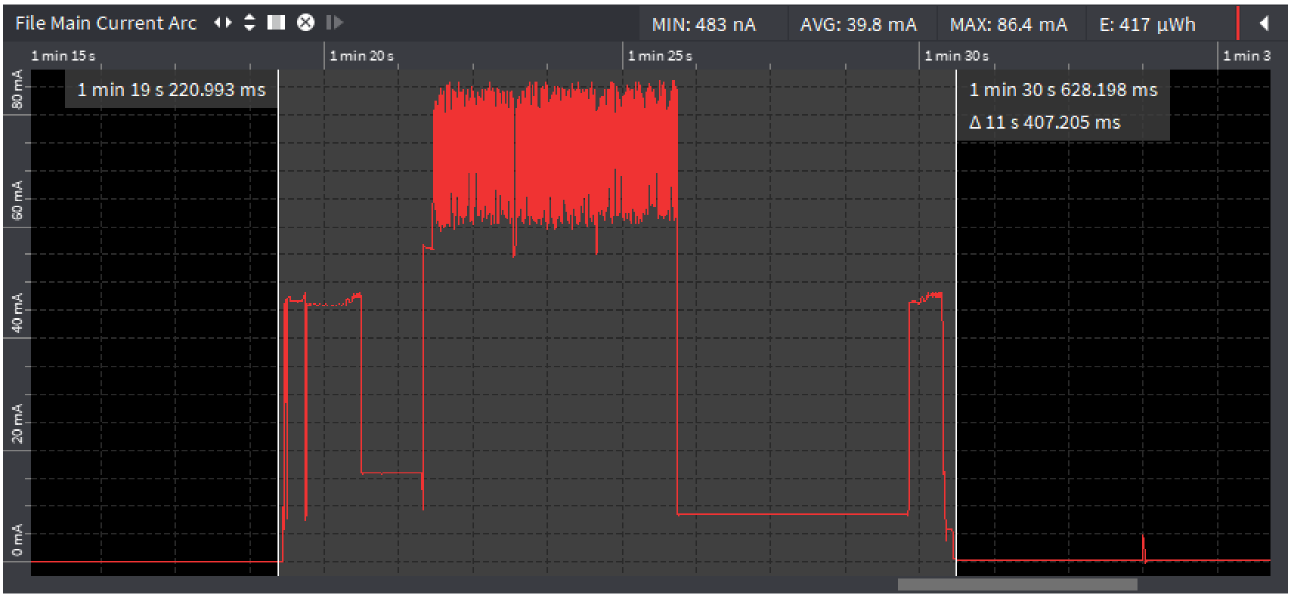

2.4.5 Sending an 80 byte message

The satellite is detected, signaling is received, and the module transmits 6 fragments. It then waits for the acknowledgement.

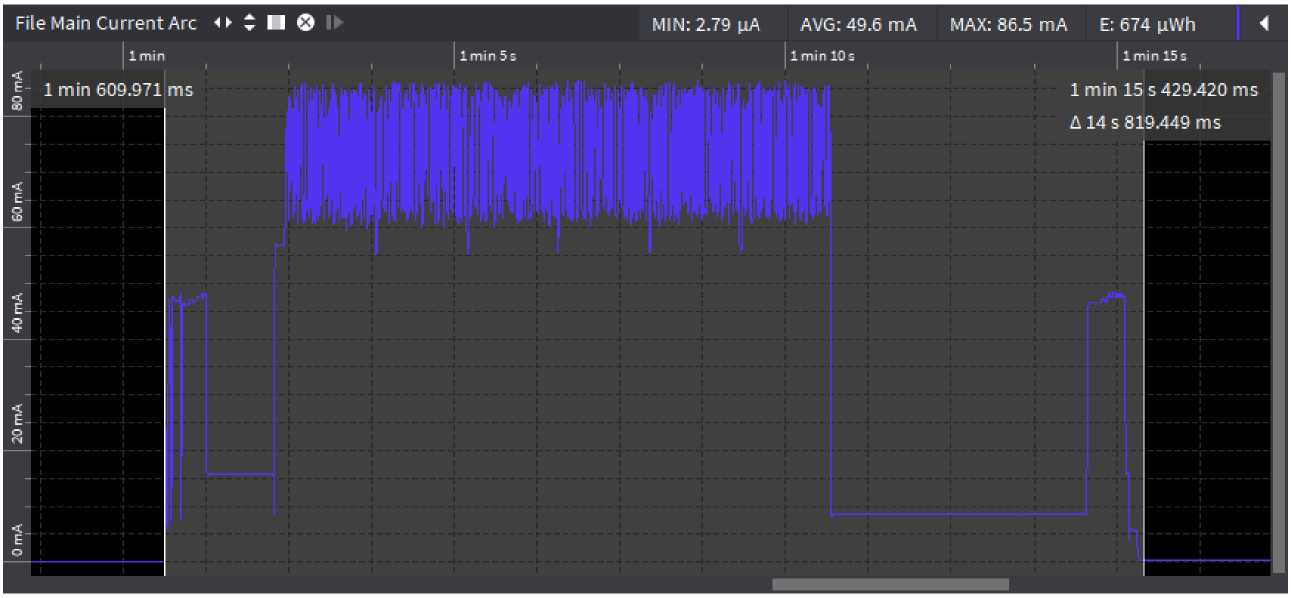

2.4.6 Sending a 160 byte message

The satellite is detected, signaling is received, and the module transmits 6 fragments. It then waits for the acknowledgement. It then repeats the process to send another 4 fragments and receives the acknowledgement for the latter.

2.5 After Satellite Contact

As long as there is data to send, the module will keep cycling through the satellite detection and transmission phases. As soon as the last fragment has been sent, the module is going in sleep mode. At some point, depending on the pass elevation, the satellite will move below the minimum elevation angle and the satellite detection will fail.

If the queue is not empty at this point, the module will continue to search for the satellite every 1.3 sec. This will last for 4 minutes counted from the first contact with the satellite. 4 minutes being the maximum connection time during a pass. After this period, the module will revert to a slower 17.9 sec search within the contact opportunity window. At the end of the window, the module is going to sleep mode.

The above timings are subject to change.

3 Typical energy profiles

The following profiles represent typical use cases considering typical values and certain assumptions on environmental factors. If you estimate your application's battery lifetime based on these profiles, consider adding some margin if some of the underlying assumptions do not match the environment your application will work in.

It is recommended you do your own measurements to refine the profiles provided in this document for your application and typical environment.

You should also factor in the elements described in Chapter 5 when doing your battery lifetime estimations.

3.1 Daily transmission with 2021 constellation

3.1.1 Assumptions

Parameter	Assumption
Maximum latency	24h Using Ephemeris data
Constellation	Satellites distributed on 1 orbit
Temperature	Constant 25°C
Payload size	Different scenarios
Message transmission	One message queued a time, complete message transmit in one satellite pass (buffer empty after pass)
Message queue time	At half the transmission period - 12h after last transmission

3.1.2 Sleep cycle

Once all messages have been sent, the module goes into deep sleep mode. However, after queuing another message 12h later, the module remains in sleep mode for 9h before starting into satellite detection mode:

Sleep	Deep sleep mode	Sleep mode
Average Current	500nA	5uA
Duration	12h	9h
Energy per day	**0.0198 mWh**	**0.149 mWh**

3.1.3 Satellite detection cycle

In this scenario the module is in satellite detection mode for 3h. There is a significant difference between the noise free and noisy environment.

Satellite detection mode	Noise free environment	Noisy environment
Average Current	145uA	400uA – 1690uA
Duration	3h	3h
Energy per day	**1.44mWh**	**3.96 – 16.7 mWh**

3.1.4 Transmission cycle

Transmit	20 bytes	30 bytes	80 bytes	160 bytes
Average Current	36mA	40mA	50mA	46mA
Duration	10.5s	11.5 sec	15s	28 sec
Energy per day	**0.347 mWh**	**0.421 mWh**	**0.688 mWh**	**1.18 mWh**

3.1.5 Total

Sending a 20byte message, 2021 constellation	Energy per day
Deep Sleep (12h)	0.0198mWh
Normal Sleep (9h)	0.149mWh
Satellite detection (3h)	1.44mWh
Transmission	0.346mWh
Retransmission rate (see Chapter 5.6)	1.5
Transmission total	0.519mWh
Total daily	**2.13mWh**

3.2 Daily transmission with full constellation

3.2.1 Assumptions

Parameter	Assumption
Maximum latency	15min
Constellation	Satellites distributed on 10 orbits
Temperature	Constant 25°C
Payload size	Different scenarios
Message transmission	One message queued a time, complete message transmit in one satellite pass (buffer empty after pass)
Message queue time	Period of 24h

3.2.2 Sleep cycle

Sleep	Deep sleep mode
Average Current	500nA
Duration	24h
Energy per day	**0.0396 mWh**

3.2.3 Satellite detection cycle

In this scenario the module is in satellite detection mode for 15 minutes after queueing the message (worst case). There is a significant difference between the noise free and noisy environment.

Satellite detection mode	Noise free environment	Noisy environment
Average Current	145uA	400uA – 1690uA
Duration	15min	15min
Energy per day	**0.12mWh**	**0.33 – 1.39 mWh**

3.2.4 Transmission cycle

Same as in Chapter 3.1.4

3.2.5 Total

Sending a 20byte message, full constellation	Energy per day
Deep Sleep (24h)	0.149mWh
Satellite detection (15min)	0.12mWh
Transmission	0.346mWh
Retransmission rate (see Chapter 5.6)	1.5
Transmission total	0.519mWh
Total daily	**0.79mWh**

4 Implementation strategies

4.1 Ephemeris

Astrocast satellites will be dispatched on different Sun-Synchronous Orbits (SSO) and equatorial orbits. The initial 2021 constellation started with one SSO, as represented in the Figures. See Chapter 5.1 for more information on the constellation development.

The ephemeris data sent by the satellite to the Astronode S contains information about the active orbits and their spacing. Based on this data, the module will calculate the start of the next contact opportunity window. The module will then wake up for satellite detection mode during the width of the window.

For illustration purposes we consider the orbits to be static, and the earth spinning between them. Once a deployed application crosses an orbit, a connection with a satellite can be established as every orbit has enough satellites deployed to assure the connection.

With the initial constellation it happens twice a day, with every new orbit deployed there will be two more crossings a day. But the crossings are not necessarily spaced equally, depending on the latitude the application is situated on, as illustrated in the following subsections.

A	*Location of Astronode S based application, crossing the orbit*
SSO	*Sun-Synchronous Orbit with multiple equidistant satellites*
Sat	*Satellites travelling on the orbit*
LS	Latitude South
LN	Latitude North
E	Equator
T	Time duration for the application to travel to the next orbit crossing

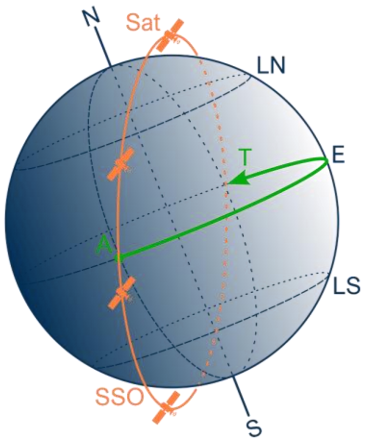

4.1.1 Equator

An application deployed on the equator will cross the orbit every 12 hours, time to travel to the first and second crossing of the day is the same:

4.1.2 Northern Hemisphere

With increasing latitude, an application on the northern hemisphere will take less than 12 hours to travel to the first crossing of the day, and from that point more than 12 hours to the second crossing:

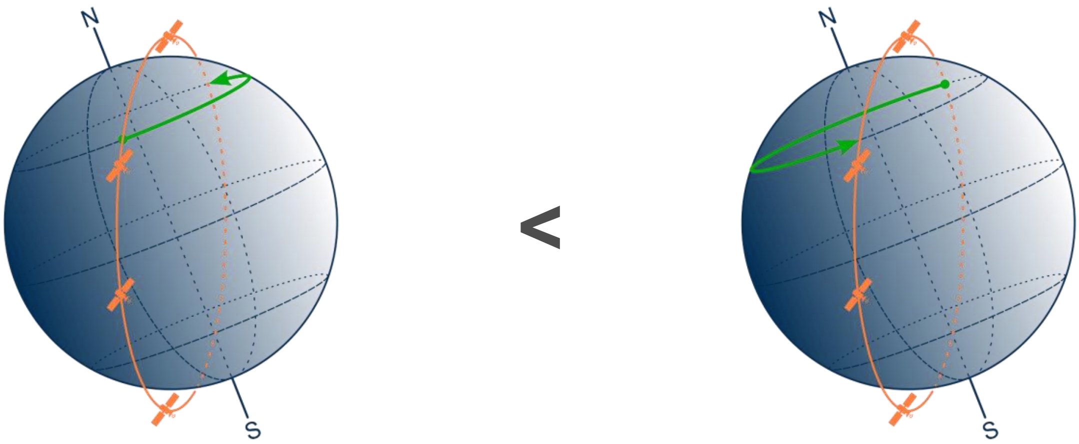

4.1.3 Southern Hemisphere

With increasing latitude, an application on the southern hemisphere will take more than 12 hours to travel to the first crossing of the day, and from that point less than 12 hours to the second crossing:

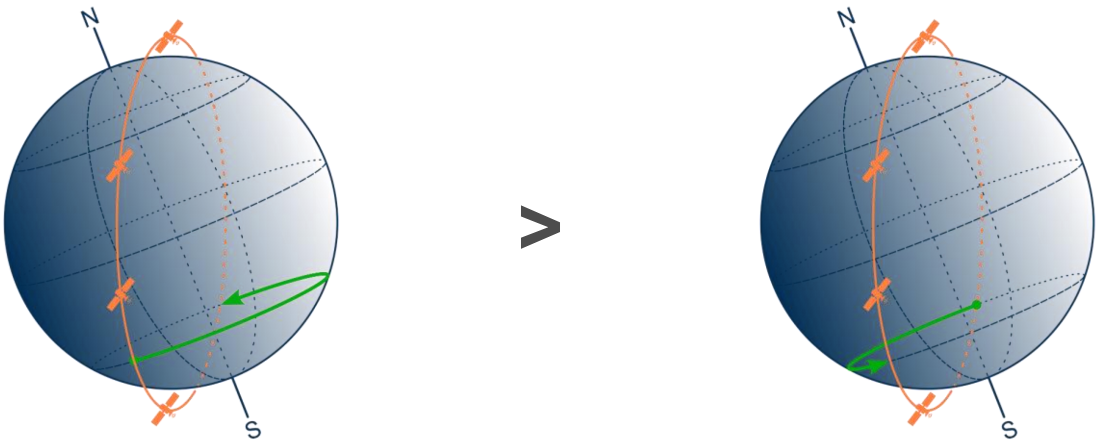

4.1.4 Contact opportunity window

Considering the system does not know if the application is installed

- on a northern or southern latitude
- before or after the first orbit crossing

the contact opportunity window compatible with the two locations in the example - and all locations inbetween - is the time span as marked in yellow.

The Astronode S calculates the start of the next contact opportunity to assure the application will establish a connection with a satellite on the next orbit crossing within an acceptable range of latitude.

The contact opportunity window width has been set to 4h 30min which allows operation between 55°S and 55°N. For applications closer to the poles see Chapter 4.1.7.

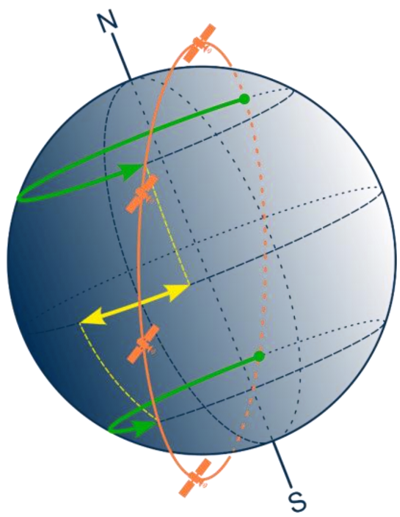

4.1.5 Using ephemeris data

When the ephemeris data use is enabled (On by default, can be disabled by the application), the Astronode S sets the timer for wake-up into satellite detection mode for the beginning of the next contact opportunity window.

With the initial Astrocast satellite constellation offering a 24h latency, it is most interesting for minimum power consumption to let the Astronode S handle the next transmission wake-up time on its own, based on the ephemeris data it receives.

If for any reason there is no contact to a satellite over the next 24 hours after queuing a message, the Astronode S will switch to continuous satellite detection mode.

4.1.6 Transition to lower latency

With the growing constellation, and as soon as the time span between two subsequent orbit crossings is falling below the contact opportunity window width, the application will constantly have a contact opportunity. At this point the use of ephemeris data will not make any difference and the Astronode S will start satellite detection mode as soon as a message is queued. The Next Contact Opportunity Read Command (NCO_R) will simply return zero at every call:

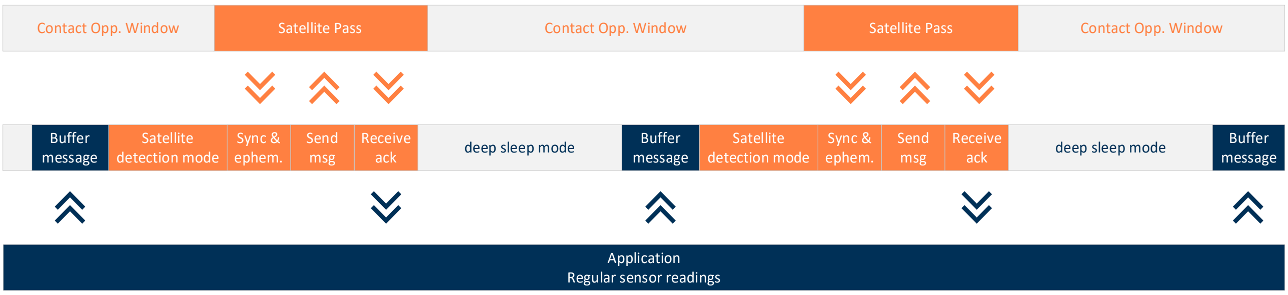

4.1.7 Use of ephemeris data not recommended

If the application is moving by more than 1000km between two transmissions, it is not recommended to work with ephemeris data. The application may start the contact opportunity window outside an orbit crossing, considering its new position, and therefore miss the pass.

If the application is installed fix above 55° latitude north or south (or moving in this region), the ephemeris data use should be deactivated. In these regions it is more efficient to have the module start in continuous satellite detection mode after queuing a message.

4.1.8 Mains powered applications

For applications with power available through mains it is recommended to disable the ephemeris data use. The Astronode S will start continuous satellite detection mode as soon as a message is queued.

4.2 Avoiding sleep mode

Queue the messages as close as possible to the next contact opportunity window, to keep the Astronode S most of time in deep sleep mode. Use the Ephemeris Read Command (EPH_R) to get the seconds left until the next contact opportunity window and plan the message buffering accordingly:

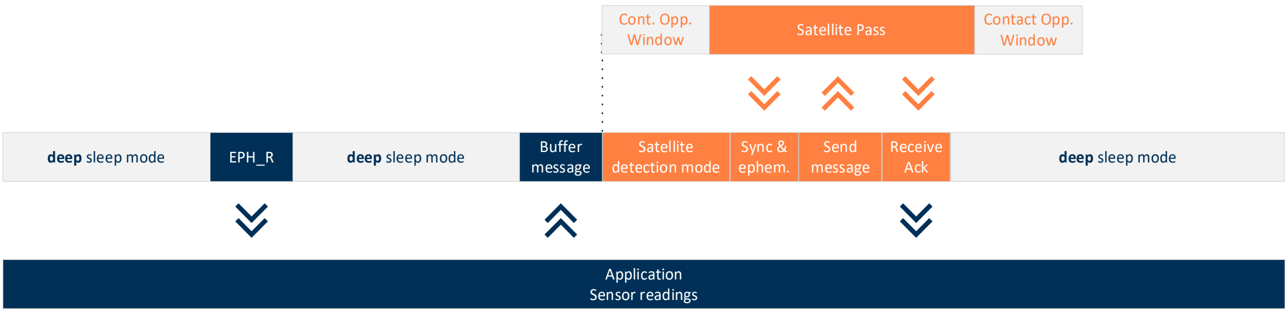

There will be a mismatch between the automatic time tag applied to the message and the time the message's data was gathered.

4.3 Daily transmission

The application can queue it's data at any time in the Astronode S' message buffer, which will be automatically sent on the next contact opportunity with ephemeris enabled. Therefore, a latency of 24h is ensured but there is no way to have a specific and periodic timing for data transmission:

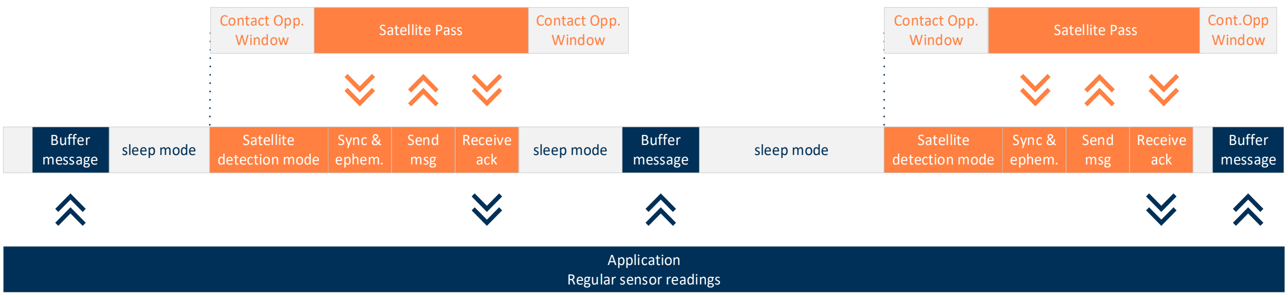

Carefully set the frequency at which the application is queuing data to the Astronode S' message buffer. Consider the buffer size limit of 8 messages, and factor in the elements described in Chapter 5 to evaluate how many messages may be sent per satellite pass.

When there is more data than can be handled in the daily transmissions, your application may have to prioritize messages before queuing them. Use the Dequeue (PLD_D) or Free (PLD_F) commands to (partially) empty the buffer for new messages.

4.4 Weekly or less frequent transmission

The ephemeris data from satellites have a validity period of 30 days. If the application is programmed to queue messages on a regular basis and more frequent than every 30 days, the Astronode S will wake up automatically for the next contact opportunity window.

Otherwise after more than 30 days in sleep mode, the module starts continuous satellite detection mode on the next message queued. In this use case, the Astronode S may also be switched off completely to economize the sleep current between the rare transmissions.

5 Influencing factors on power consumption

5.1 Satellite constellation

In beginning of 2021 Astrocast launched the first 5 commercial satellites. Since then, the network keeps growing until the full constellation of 80 satellites is reached. The satellites will be dispatched on 10 different orbits.

With every new orbit deployed, time to travel between two orbit crossings for a given location decreases drastically, with immediate effect on lower latency (End-to-end message transmission time). The following Figure does illustrate this evolution:

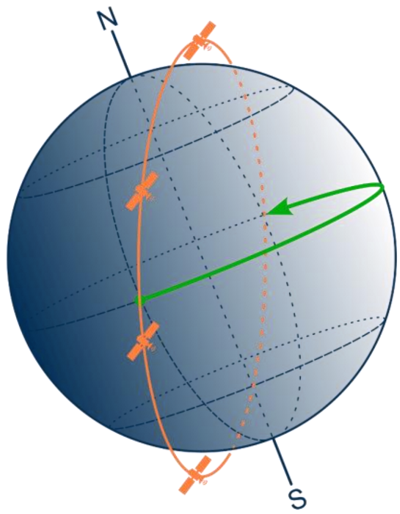

Initial constellation, one orbit

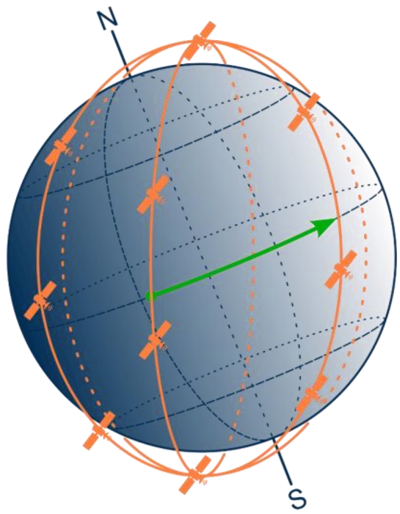

Constellation with two more orbits deployed

With the full constellation, latency will be down to 15min, till then time spent in satellite detection mode is more important and must be factored into battery lifetime calculations. Use ephemeris data if possible, see Chapter 4.1.

Please contact Astrocast for more information about the constellation agenda.

5.2 Vision angle, Elevation, Antenna performance

A satellite pass duration is about 4 minutes. The first and last part of the pass trajectory are at an elevation too low to establish the RF link. Thus, the time with opportunity for the module to send data to the satellite is limited and needs to be preserved as much as possible. The factors additionally limiting this time and therefore increasing power consumption over time are explained hereafter.

5.2.1 Explanation of symbols

Different use cases are described in this chapter with a simplified graphic representation. Find the different elements explained hereafter.

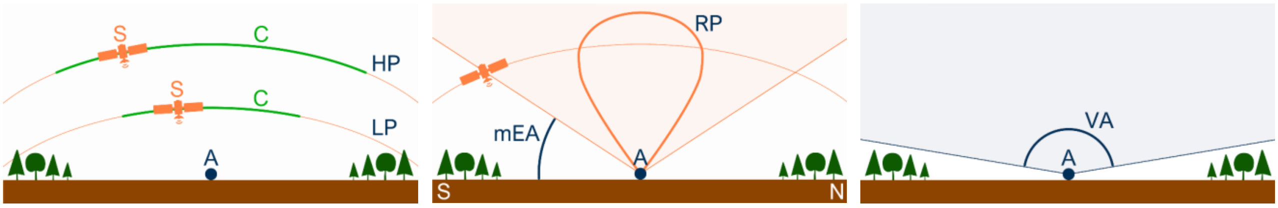

A	*Location of antenna with Astronode S based application*
S	*Satellite on its orbit*
C	*RF link possible during the green section of the pass*
LP	*Low elevation pass (passes below mEA are not considered)*
HP	*High elevation pass*
mEA	*Minimum elevation angle necessary to establish an RF link between satellite and application. Minimum angle given by the network, considering the use of an antenna with sufficient gain*
VA	*Vision angle (Antenna's unobstructed view to sky)*
RP	*Antenna Radiation Pattern in dB*

5.2.2 Ideal setup

With a free view to sky the RF link between satellite and antenna is only limited by the minimum elevation angle. The use of an antenna with sufficient gain (See *0534 Antenna Datasheet & Integration guide* for more details) provides a viable RF link for both high and low elevation passes.

- Free view to sky
- Antenna with good performance
- RF link only limited by minimum elevation angle

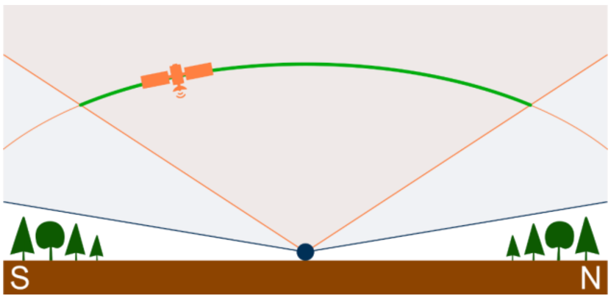

- The lower the pass, the shorter the time with RF link available
- With a low pass, less messages may be sent than with a high pass, or one message may need two consecutive low passes to be sent entirely

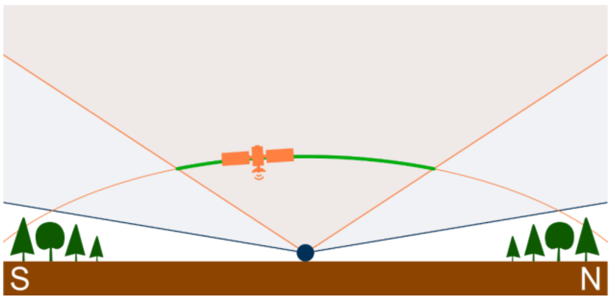

5.2.3 View to sky partly obstructed (North-South)

If the view south or north of the application is partly obstructed, the time to connect to the satellite is limited in each pass compared to the ideal setup 5.2.2. The use of an antenna with sufficient gain provides a viable RF link mainly in high elevation passes, but with a limited amount of data to transmit.

- Antenna with good performance
- RF link not available in obstructed part of satellite trajectory

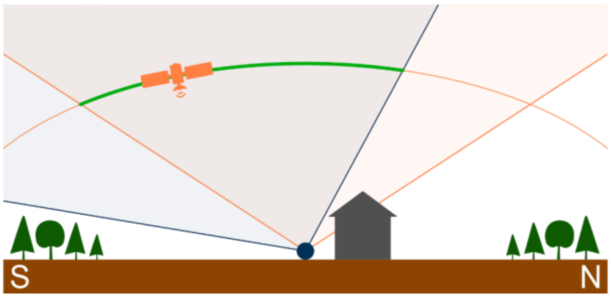

- The same situation with a low elevation pass may have an RF link time too short for a successful connection
- A message may need two consecutive low passes to be sent entirely

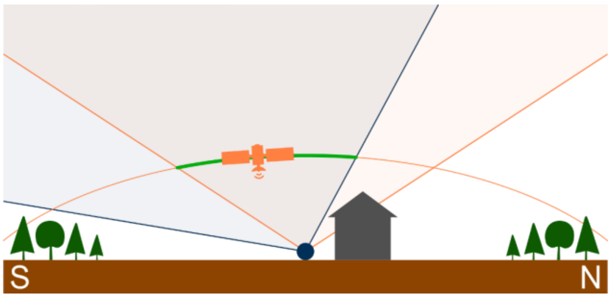

- With increasing obstruction, time with RF link available is getting too short for a successful connection

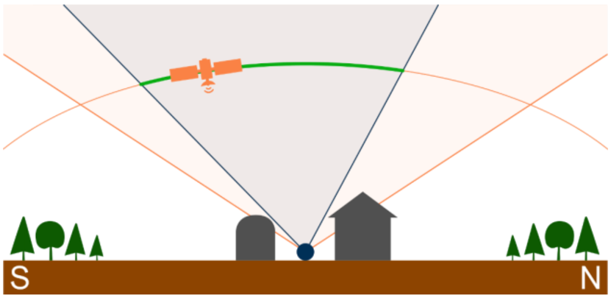

5.2.4 View to sky partly obstructed (West-Est)

If, for example, the view east of the application is partly obstructed, a satellite travelling on a polar orbit east of the application may be missed entirely. On the other hand, a similar pass of a satellite travelling west of the application provides a viable RF link as in the ideal setup 5.2.2.

- Antenna with good performance
- RF link is entirely available when pass is west of the application (in this example)
- No connection possible when the pass is east of the application (in this example)

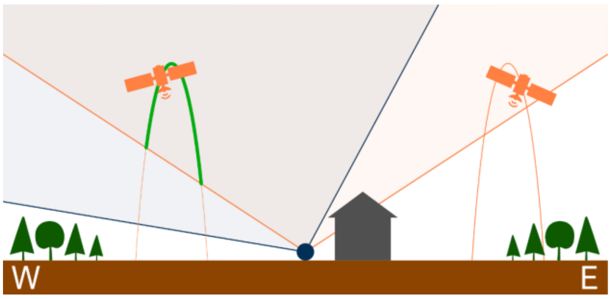

5.2.5 Antenna performance

The choice of antenna to be implemented in the application, the size of its ground plane (if required), its physical integration within the application and its enclosure's form and material can all have a huge impact on the performance of the antenna.

Carefully select and integrate the antenna to have an application with minimum power consumption and a good RF link quality to the satellite. As a reference refer to the *0534 Astronode Patch Antenna Datasheet and Integration Guide*.

Most applications require an antenna radiation pattern which features sufficient gain in all directions above the minimum elevation angle. This allows to have a maximum RF link time available per pass.

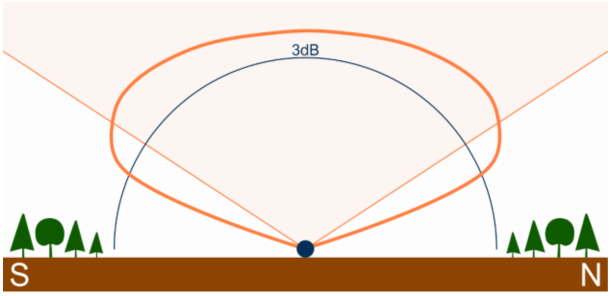

For some applications it may be sufficient to have good antenna gain but only for higher elevation angles. This will potentially result in higher latency as not all satellite passes can establish a connection.

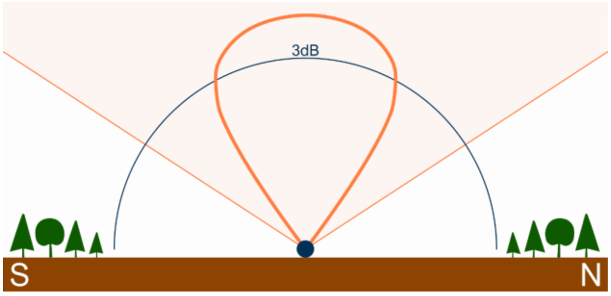

Some antennas may have an omnidirectional radiation pattern, but with insufficient gain to establish a connection. Carefully verify or test the gain versus elevation when considering the use of antennas other than those provided by Astrocast.

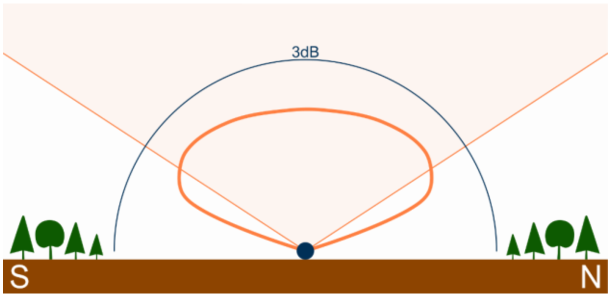

5.3 Payload size

Standard payload for one message is up to 160bytes. In addition, geolocation data can be enabled and added to the message, see GEO_W command. In this case the module will send an additional 8bytes on top of the payload.

If, for example, your application data counts 20 bytes and geolocation is enabled, the total payload will be 28 bytes and the module will send 3 fragments.

Optimizing payload size as explained in the following subsections does work best in ideal conditions. As soon as the application is facing important retransmission rates, the fragment and burst count is not controlled entirely by the application anymore.

5.3.1 Optimizing number of fragments

If the application allows it, the size of the message may be adjusted to the next smaller fragment number. Or data may be added to a message to use all the free space of an opened fragment. See table hereafter for details:

Payload [bytes]	Number of fragments	Number of bursts[1]
1 – 3	1	1
4 – 21	2	1
22 - 39	3	1
40 – 57	4	1
58 – 75	5	1
76 – 93	6	1
94 – 111	7	2
112 – 129	8	2
130 – 147	9	2
148 – 160	10	2
162 – 170	11	2

1. Considering bursts with maximum 6 fragments. The number of fragments contained in a burst is a system parameter and may be changed by Astrocast without prior notice.

Here an example of the two optimization approaches:

- If your message consists of 43 bytes, by optimizing the payload down to 39 bytes you would economize energy by sending 3 fragments instead of 4.
- If your message consists of 43 bytes, sending it would cost you the same amount of energy than sending 57 bytes of payload as in both cases the system will send 4 fragments.

5.3.2 Optimizing number of transmission bursts

Generally speaking, it is more efficient to send messages less frequent but with more payload.

Messages with a payload above 93 bytes or 6 fragments will be sent in two transmission bursts. Compared to smaller messages there is additional energy necessary for the second signaling and acknowledge phase. To optimize your application's energy consumption, you may choose one of the following options:

- Stay below 94 bytes of payload to avoid a second transmission burst
- Use a maximum of the 160 bytes payload available per message to amortize the second transmission burst

5.4 Temperature

Power consumption of the module is affected by temperature only in both normal and deep sleep mode. Consider the increased sleep mode currents especially if the application will operate in higher temperatures. However, temperature does not have any impact on the active modes of the Astronode S.

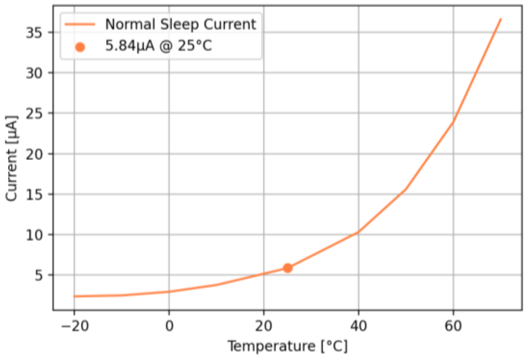

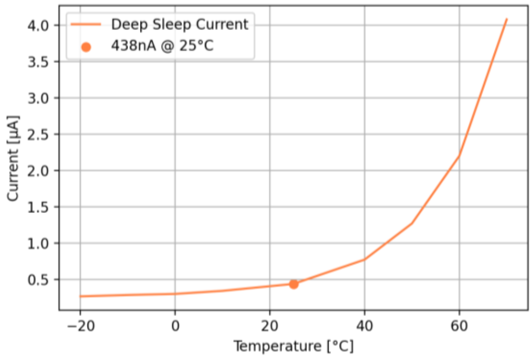

When powering the application with a battery, carefully check the maximum continuous current the battery can deliver at different temperatures, and the related voltage drop. To work as expected the Astronode S' power supply must stay at any time within the limits specified in the *0532 Astronode S Datasheet*.

5.5 Noisy environment

If one or multiple RF sources are close by the installed application, the present RF signals may lead to increased power consumption during the satellite detection mode. This can be the case in an office or urban environment. The search period of 17.9secs remains unchanged, but each single receive window may be longer, increasing average power consumption in the satellite detection mode. See examples in Chapter 2.3.2.

5.6 Retransmission rate

5.6.1 Network load

Characteristic of the Astrocast data transmission protocol is a high efficiency. Based on slotted ALOHA, it features several proprietary mechanisms allowing a high number of Astronode S communicating with a satellite during a pass without massively increasing collisions. Availability of the system is further assured by active load regulation on the satellite level:

- Even with a high number of Astronode S communicating at the same time with one satellite, the average retransmission rate will be lower than 2.2.
- In most cases, where only a few Astronode S are installed within the same spot, retransmission rate due to network load can be expected to be 1 (no collisions due to network load)

5.6.2 Environmental impact

Independent of the network load, a certain number of un-acknowledged fragments is expected in the system. This depends on several factors, such as the RF environment, satellite pass elevation, and RF link quality (See Chapter 5.2.5).

- Testing in urban conditions at Astrocast headquarters with optimal antenna integration and installation suggests a 1.2 times retransmission rate
- A conservative approach is to calculate with a 1.5 times retransmission rate as done in the Chapter 3 of this document