life.augmented

AN5928

Application note

How to build a short range wireless application with STM32WBA MCUs

Introduction

This application note guides designers through the steps required to build specific Bluetooth® Low Energy applications based on STM32WBA series microcontrollers.

It groups together the most important information and lists the aspects to be addressed.

To fully benefit from the information in this document and to develop an application, the user must be familiar with STM32 microcontrollers, Bluetooth® Low Energy technology, and master system services, such as low-power management and task sequencing.

For more information, such as deep dive concepts, instructions, and examples, visit the wiki page or refer to www.st.com.

1 List of acronyms and abbreviations

Table 1. Acronyms and abbreviations

Acronyms	Definitions
ACI	Application command interface
ATT	Attribute protocol
BLE	Bluetooth® Low Energy
BSP	Board support package
EXTI	Extended interrupts and event controller
FM	Flash manager
FUOTA	Firmware update over the air
FW	Firmware
GAP	Generic access profile
GATT	Generic attribute profile
HAL	Hardware abstraction layer
HCI	Host controller interface
HSE	High-speed external oscillator
HW	Hardware
IFS	Interframe space
IP	Semiconductor intellectual property core
ISR	Interrupt service routine
L2CAP	Logical link control and adaptation protocol
LL	Link layer
LPBAM	Low-power background autonomous mode
LSE	Low-speed external oscillator
NVIC	Nested vectored interrupt controller
NVM	Nonvolatile memory
OS	Operating system
PHY	Physical layer
PLL	Phase-locked loops
QOS	Quality of service
RTC	Real-time clock
RTOS	Real time operating system
SNVMA	Simple NVM arbiter
SoC	System on a chip
SW	software
TZ	TrustZone®
WFI	Wait for interrupt (Arm assembly code)

2 General information

The STM32WBA device is a multi-protocol wireless and ultra-low-power device. It embeds an Arm® CPU Cortex®-M33 core alongside a 2.4 GHz radio compliant with Bluetooth® Low Energy (BLE).

It is also composed of a multitude of smart and high-performance peripherals, a large set of advanced and low-power analog features, and several peripherals tuned for low-power modes.

In addition to these, the STM32WBA also provides a dedicated security framework.

For more precise explanation relative to protocols , for example, BLE, etc. or system concepts, for example, low-power, refer to our wiki website.

3 STM32WBA overview

3.1 STM32WBA series key features

This section lists the key features of the STM32WBA series.

Figure 1. STM32WBA key features

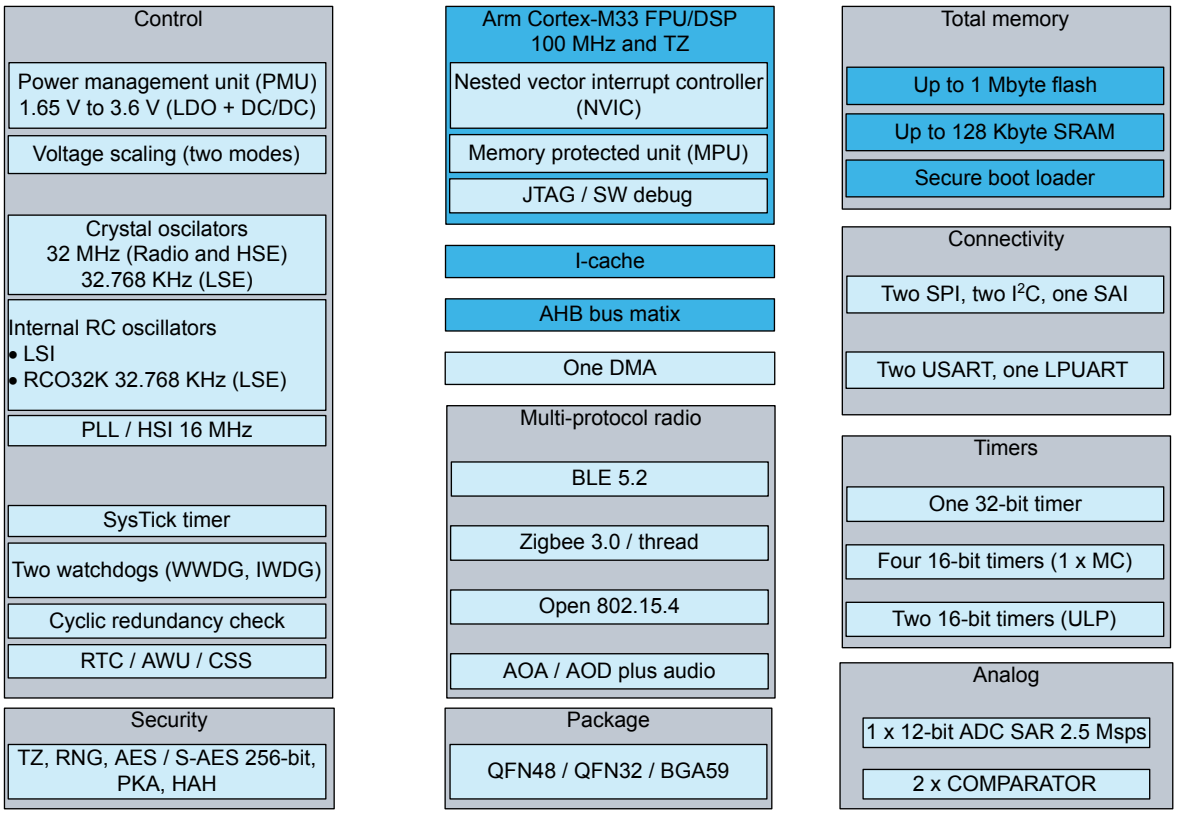

For more information, including GPIOs, NVIC groups, memory mapping, etc., visit the wiki page or refer to the user manual on www.st.com.

3.2 SW architecture

The software architecture is explained from two points of view:

- the application
- the project

3.2.1 Application view

From the application point of view, the user can rely on multiple modules and SW blocks to create an application. The framework is referred to as the Application Firmware Framework.

The architecture is designed over a 3-level organization:

- Application firmware
- BLE stack
- Platform resources

Figure 2. **Application Firmware Framework**

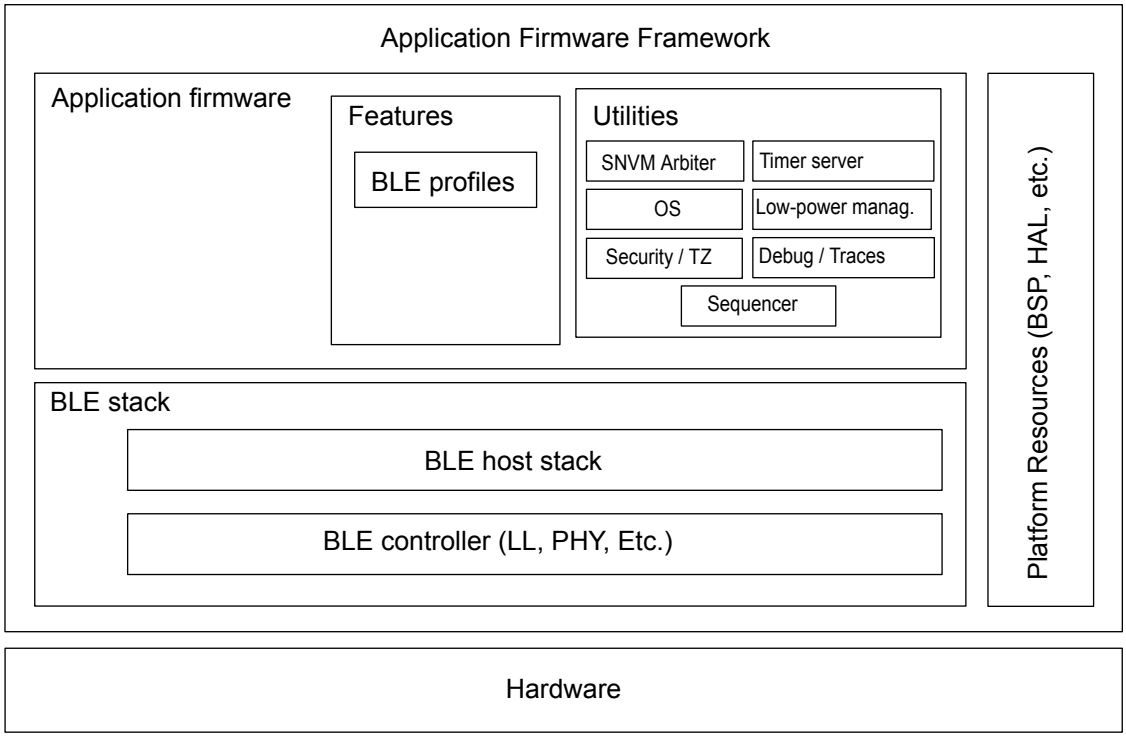

3.2.1.1 *Application firmware*

The user application is based on:

- IPs either from STMicroelectronics or from a third-party supplier.
- Features developed by STMicroelectronics for short-range protocols, such as Bluetooth® Low Energy profiles, etc.
- Utilities required by the application and the protocol stacks.
- Generic utilities that provide easy access to basic features such as low-power management, sequencer, etc.

3.2.1.2 *Bluetooth® Low Energy stack*

The Bluetooth® Low Energy stack is the main interface between the application and the hardware for BLE purposes. It is composed of two layers:

- BLE stack: It manages all networks and transport protocols that allow the application to communicate with other devices. Provided by STMicroelectronics, it can also be substituted by the host stack of the user.
- BLE controller: The BLE controller manages the hardware part, the RF state, and guarantees that the RF protocol is correctly followed. It includes the physical layer, the link layer, and the host controller interface.

3.2.1.3 *Platform resources*

The platform resources include all the HALs, BSPs, and drivers that ease the platform hardware use.

3.2.2 Project view

The project organization differs slightly from the application architecture, as it is repository based. Nevertheless, the above-mentioned concepts are found without difficulties.

This organization is divided into four parts as shown below.

Figure 3. **Project view organization**

3.2.2.1 *Application*

Application is the core part of the project. It regroups the main information and actions that define the user application.

This section is divided into four subdirectories:

- Core: This directory is the main entry point of the application. It contains all the setup and entry codes such as hardware initialization routines, IRQ's setup, scheduler configuration, tasks registration, etc.
  The files contained below this directory are `main.c`, `app_entry.c`, etc.
- Startup: Includes the startup file of the project.
- STM32_WPAN: This part is composed of two subparties:
  - App: Main applicative source files, for the example of HR, this repository contains the main application file `app_ble.c` and the BLE services implementation.
  - Target: Dedicated to interface and integration of the libraries and modules present in middleware.
- System: System-related modules, interfaces, and utilities (PKA, RNG, advanced memory manager, USART, RTDebug, etc.). The user also has access to several configuration files to customize the system experience.

3.2.2.2 *Drivers*

The drivers are divided into four component sets:

- Hardware abstraction layer (HAL):
  This layer provides the hardware abstraction drivers and the hardware interfacing methods to interact with the upper layers (application, libraries, and stacks). The HAL is composed of:
  - HAL drivers:
    A set of portable abstraction APIs based on high level services built around standalone processes. The HAL drivers are functionality oriented, for example to the timer peripheral, for which it is possible to split APIs into several categories following the functions offered by the IPs (basic timer, capture, PWM, etc.).
  - Low layer drivers:
    A set of basic drivers with direct hardware access with no standalone processes. This layer is called either by the applications or by the HAL drivers.
  - HAL core drivers:
    A set of internal drivers providing low level operations to the “complex” peripherals like the USB and Ethernet. They all come with a dedicated source and header files. Some drivers may have an additional file containing extended features, identified by the file extension "_ex".
- BSP drivers:
  This layer contains the high-level board services for the different functionalities offered by the hardware (LED, audio, pushbuttons, etc.) and the drivers for the external components mounted on the used boards (audio codec, IO expander, etc.).
- CMSIS drivers:
  Cortex® microcontroller software interface standard (CMSIS) drivers that provide a single standard across all Cortex-Mx processor series vendors. It enables code reuse and code sharing across software projects.
- Basic peripheral usage examples:
  This layer encloses the examples built over the STM32 peripheral using the HAL APIs and the low layer drivers.

3.2.2.3 *Middleware*

Libraries and protocol-based components (BLE stack, for example). This directory contains the BLE services management, system commands, etc.

Horizontal interactions between the components of this layer are directly performed by calling the features APIs while the vertical interaction with the low-level drivers are performed through specific callbacks and static macros implemented in the library interface.

3.2.2.4 *Utilities*

Miscellaneous software utilities that provide additional system and media resources services like sequencer tools box, time server, low-power manager, along with several trace utilities and standard library services like, memory, string, timer, and math services.

3.3 SW concepts/features

This chapter covers the main software components available on the STM32WBA. Each module concept is described briefly to enhance user appreciation.

For further information, visit the online wiki pages, for example system modules. Each module has its own page, and regroups the interface, the current ways of use and different examples. This information is kept up to date with the latest software release.

3.3.1 Sequencer

The integration cost of a real time operating system can be disproportional in the case of a simple application as a result of:

- RTOS knowledge required
- Increase in the complexity of the application
- Impact on sizing of RAM/ROM

This is the reason why the sequencer module has been developed, as an alternative to a real time OS.

3.3.1.1 *Concepts*

The sequencer utility is designed as a simple alternative for simple application cases. However, it does not cover all the services provided by an operating system. For instance, this software solution does not provide preemption mechanisms, so this must be considered in the application design. Instead of tasks designed with a risk of freezing the system, it is recommended to use reentrant functions based on state machines.

Important concepts to consider:

- Task creation: initialize the task and render it callable by the internal scheduler of the sequencer.
- Task enable: enable the task, through a task or an interrupt, so the task can be executed by the scheduler.
- Task pause/resume: pause/resume the task execution from the scheduler point of view, independent whether the task is enabled or not.
- Idle task: if the scheduler has no task to execute, it calls an optional hook function to manage entry in idle mode.
- Task execution: calls the function associated to the task, and the scheduler is locked until the function returns.
- Sequencer: embed a task scheduler that sequences the tasks execution and allows the task to stop until an event reception.

Additionally, the sequencer provides the following features:

- Up to 32 tasks registered
- Request a task to be executed
- Task pause and resume
- Wait for a specific event (not blocking)
- Priority on tasks
- Allow management of an IDLE task

3.3.2 System clock manager

The system clock manager is a feature for system clock management between multiple users. It determines which frequency is best suited to the system needs. This module also handles the clock management during the low-power modes.

Note: The clock is cut during several low-power modes.

3.3.2.1 *Concepts*

The module is based on a request behavior. Any user, up to 32 users, can request for clock frequency modification. Among the requests, the system clock manager determines the request that fulfills the lower possible configuration.

The best system clock evaluation is realized at each new request.

Figure 4. System clock evaluation

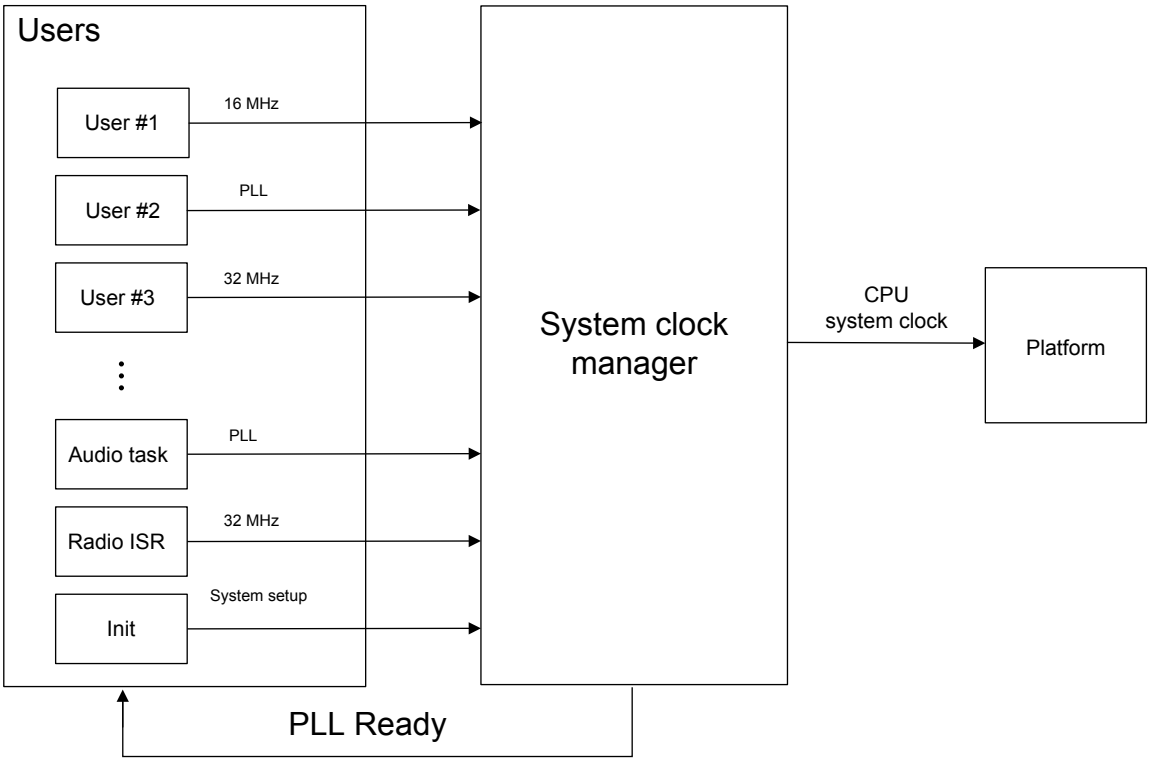

The decrease clock speed request is handled immediately.

However, the increase clock speed request requires more time to setup. This is not blocking as an interrupt mechanism is used to switch to the requested speed. This permits the firmware to keep running at the actual clock speed until the system is ready.

Different clock configurations are possible to request. The supported ones are the following:

- HSE32
- HSE16
- PLL (with a dedicated interface for configuration)

The module also handles the configuration of the flash memory and SRAM latencies, the CPU wait states, the configuration of the VOS and the AHB5 divider configuration, depending on the determined clock speed. For the other impacted modules or systems, the user must adequate and adapt the configurations to the newly set CPU system clock frequency.

3.3.2.2 *Going further*

For more information on the interface, using the interface, and for examples, refer to the wiki system page.

3.3.3 Low-power management

The low-power manager provides a simple interface to receive the input from up to 32 different users and computes the lowest possible power mode the system can use. It also provides hooks to the application before entering or on exiting low-power mode.

3.3.3.1 *Concepts*

From an architecture point of view, the low-power manager has a two section organization. One for the BLE LL and one for the SoC. Both impact each other on their low-power management.

Figure 5. Low-power management

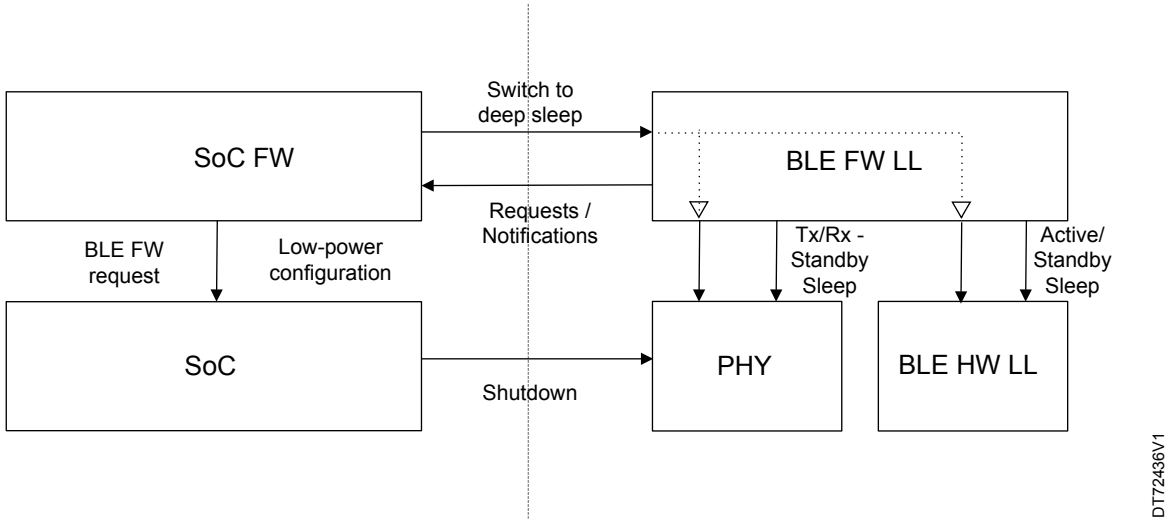

There are numerous different low-power modes available. They all differ according to the section they impact and the intended objective. The different low-power modes of the two sections are as follows:

- BLE IP low-power modes:
  - Hardware LL:
    - Active/Standby: These modes are requested when the radio is operating. Standby is used during IFS.
    - Sleep: This mode is used when there is no radio activity. This is fully handled by the BLE firmware LL. It is possible to access BLE hardware LL registers in this mode. For this, the bus clock needs to be active. To access RXTX SRAM or sequence SRAM, the baseband clock must remain active. This is managed by the LL.
    - Deep-sleep: This mode is used when there is no radio activity and no need to read/write any register in the BLE hardware LL (except the sleep timer and the register to set the low-power mode of the BLE hardware LL). This mode is entered when the BLE hardware LL has remained in idle mode long enough to ensure that the overhead to exit from deep-sleep mode is not significant. This requires profiling to define the threshold value.
  - PHY:
    - Tx/Rx/Standby: These modes are requested when the radio is operating.
    - Sleep: Standby is used during IFS.
    - Shutdown: This mode is applied only when the SoC enters standby (no retention). All radio configurations are lost (calibration, etc.).

- SoC low-power modes:
  - Run mode: The CPU and the system clock are running. It must be used only when it is required to execute code.
  - Sleep mode: Only the CPU is in low-power mode (stops fetching code). The system clock is running. This must be used when a peripheral (that does not have a kernel clock) requires the system clock to operate while the CPU is stopped.
  - Stop0/Stop1 modes: The CPU, the system clock, and the high-speed oscillators are all stopped. LPBAM peripheral can still operate. The 2.4 GHz radio can still be active when in Stop0 Range1. For this, the high-speed system clock and peripheral kernel clocks may be kept running. When a kernel clock is enabled, Stop0 is selected by the SoC hardware. This should be used when standby is not possible.
  - Standby retention: This is the lowest low-power mode targeted in the application. The CPU, the system clock, the high-speed oscillators, and most SoC registers are all lost. It is required to save/restore all contexts and hardware register configurations (that are not retained in the SRAM) to select which content is not lost.
    The time required to save/restore all contexts and hardware registers depends on how many peripherals are used in the application (for example, most of the time GPIO, UART, CPU, DMA, Clocks are used and need to be restored).
    Due to the extra CPU execution time required to exit/enter this mode (mostly due to startup and save/restore operation), it should be used only when the latency to wake-up does not matter and when the “idle” time is long enough so that the power saving in this mode compensates the extra power consumption due to extra CPU processing to enter/exit this mode.
    The radio and BLE hardware LL registers, except for the sleep timer, are lost and need to be restored by the BLE LL firewall. (The BLE LL is in deep-sleep). This should be used, when possible, between radio events to save additional power.
  - Standby mode: SoC registers, radio sleep timer, and BLE hardware LL SRAMs are lost. Only RTC, TAMP, and WKUP are retained. This should be used only when it is decided to restart everything from wake-up. For example, when there is no more radio activity and the application may stop for several days to wake up on a GPIO event or an alarm from the RTC.

Application idle state

The application is either in active state when the CPU executes code or in idle state when the CPU is stopped.

This idle state may be split into three parts:

- Pre-idle task: It holds the last code to be executed before stopping the CPU. It runs outside the critical section so the user may implement a long execution time code without impacting any interrupt latency.
- Idle task: It holds the code to enter low-power mode. A critical section is required to avoid interrupt loss due to race conditions (when WFI is used to stop the CPU). When moving out of low-power mode, the mandatory system setup shall be executed here before any interrupt run. This is required as this is the last code executed before jumping in the interrupt handler.
- Post-idle task: It holds the first code to be executed when moving out of the idle task. It must be executed when the pre-idle task has been executed and only in that case. It runs outside critical sections so the user can implement code with a long execution time without impacting any interrupt latency. At this point, any system setup required must be executed to run the application in Active, but it is not mandatory when followed by an interrupt handler.

3.3.3.2 *Going further*

For more information on the interface, using the interface, and for examples, refer to the wiki system page.

3.3.4 Memory management

An enhanced version of dynamic allocation is available with the STM32WBA series: the advanced memory manager.

It eases dynamic allocations in a multitask/shared environment without adding much complexity.

The new features provided by AMM are the following:

- Support several virtual memories, up to 255
- Minimum size warranty per virtual memory (QOS)
- Callback in case of memory allocation failure

3.3.4.1 *Concepts*

The advanced memory manager comes on top of a basic memory manager, such as Alloc and Free.

Figure 6. Advanced memory manager

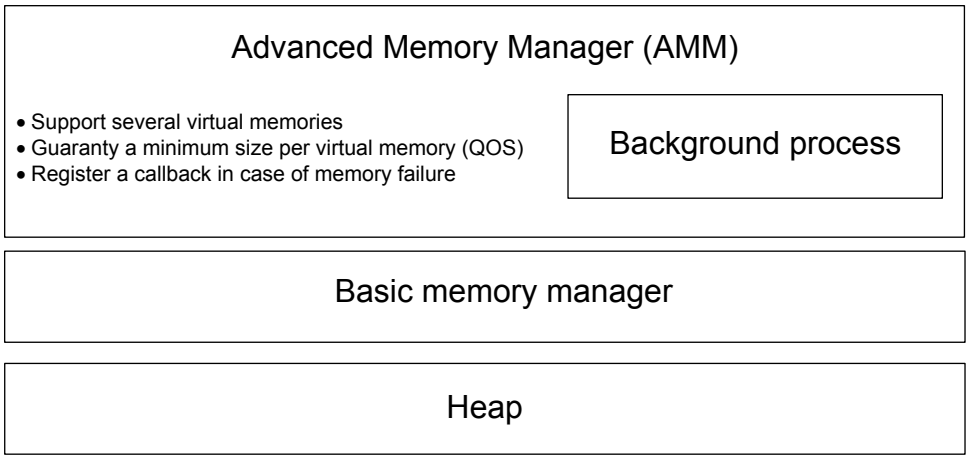

Basic memory manager and heap

The basic memory manager can be anything that is capable to allocate/free memory. It must support the capability to merge into one single memory two continuous memories that have been freed (coalescence). It is provided by the user before using the AMM through a registering function.

The heap can either be a dedicated memory provided by the application, or the legacy default heap provided by most applications at link time. The size of the heap shall be at least equal to the sum of all the virtual memories combined size.

Virtual memories

In a concurrent execution application, dynamic allocation can encounter some issues with heap sharing. For instance, a task may request a major part of the available heap leading to allocation failure for all other tasks requesting too much memory.

To avoid that kind of issue, the advanced memory manager introduces the concept of virtual memories and shared pool.

Figure 7. Advanced memory manager (AMM)

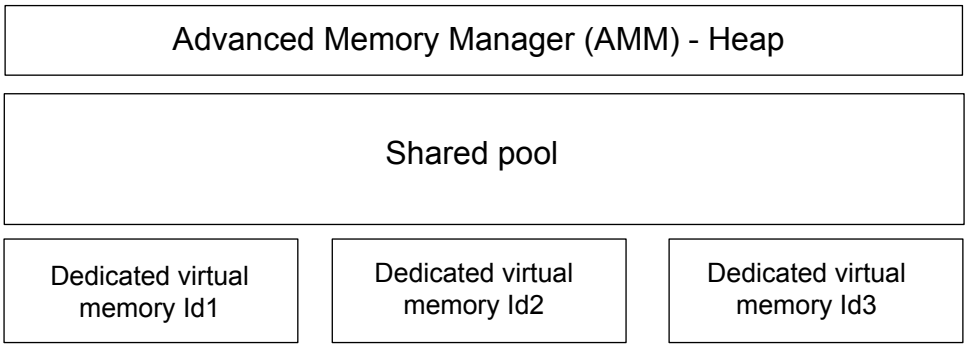

Virtual memory:

- Dedicated user memory pool, for example, specific VM IDs with up to 255 different IDs
- Ensures that the users have access to the memory amount needed for a minimal execution

Shared pool:

- Mutual memory pool is used by any virtual ID

- Provides memory for an optimal execution
- Size corresponding to the BMM pool size minus the space required for the virtual memories

Retry callback

During program execution, users may encounter some memory allocation failure if there is not enough memory available. Instead of setting up a polling mechanism, in an allocation request, the AMM offers the possibility to register a callback. This callback informs the requester (in an asynchronous way) that some space has been freed, either in the shared pool or in its dedicated virtual pool, and a new allocation request can be submitted.

In some cases, the memory can be freed from a different context than the one it has been allocated from. The callback should not implement any active code and should be used only to set up a new allocation request from the main context.

3.3.4.2 *Going further*

For more information on the interface, using the interface, and for examples, refer to the wiki system page.

3.3.5 Flash memory management

The flash memory management proposes a simple interface to the upper layers to execute operations in flash memory. It manages synchronization between flash memory operations and the BLE LL activity. Thus, users do not have to spend additional effort to synchronize RF timing and flash memory operations.

3.3.5.1 *Concepts*

The flash memory management relies on different concepts:

1. Asynchronous flash memory operations
   All flash memory operations are requested via the flash memory manager interface and executed afterward by the module. The requester is, later on, notified on the status of the flash memory operation. However, in case of a write operation, the user buffer is held as long as the flash memory operation is not over. If the buffer is updated during the write operation, the user must restart a brand new flash memory operation.
2. Four layer organization
   Flash memory management is based upon a four layers distribution composed by:
   - Flash memory manager: Main user interface for flash memory operation (flash memory write, flash erase, etc.).
   - RF timing synchro: Module that realizes synchronization between BLE LL and flash memory operations by activating or deactivating the dedicated flash memory control status.
   - Flash memory driver: Low-level driver abstraction layer. Controlled by the flash memory control statuses.
   - HAL flash memory: Low level driver that interacts directly with the flash memory hardware.

Figure 8. Four layer organization

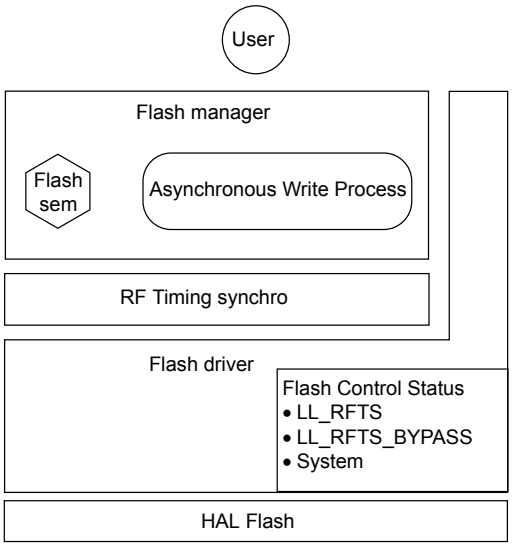

3. Flash memory access
Flash memory access is protected at two different levels:

– Flash memory semaphore:
A semaphore is required to share the flash memory interface between several SW modules. The owner of the semaphore is the only one that can request flash memory operations. The semaphore is attributed to the first requester and released once its operation is over.

– Flash memory control statuses
Independently from the flash memory semaphore, the flash memory driver provides flags, the flash memory control status, to prevent flash memory operation depending on the system activity. These flags/statuses are checked before any flash memory operation by the flash memory driver.

Simple routine

The classic flash memory operation routine is represented as follows:

Figure 9. Classic flash memory operation

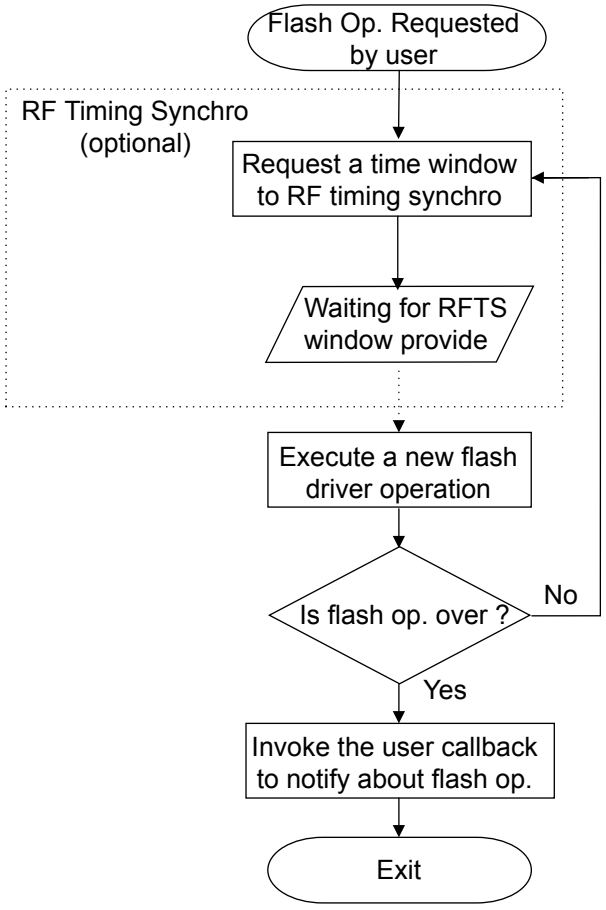

According to the RF state, the flash memory manager adapts its behavior. During RF activity, it requests the help of the RF timing synchro to synchronize flash memory operation and radio activity. Otherwise, the flash memory manager simply executes the flash memory operation until all work is done.

3.3.5.2 *Simple NVM arbiter*

The simple NVM arbiter is a different interface than flash memory operations. It relies on the flash memory manager but adds the possibility to create and manage NVMs, up to 32 NVMs.

NVMs are identified by a unique ID and are composed of multiple banks (at least two, with one for restore and one ready for write). Banks have a sector boundary.

Figure 10. **Simple NVM arbiter**

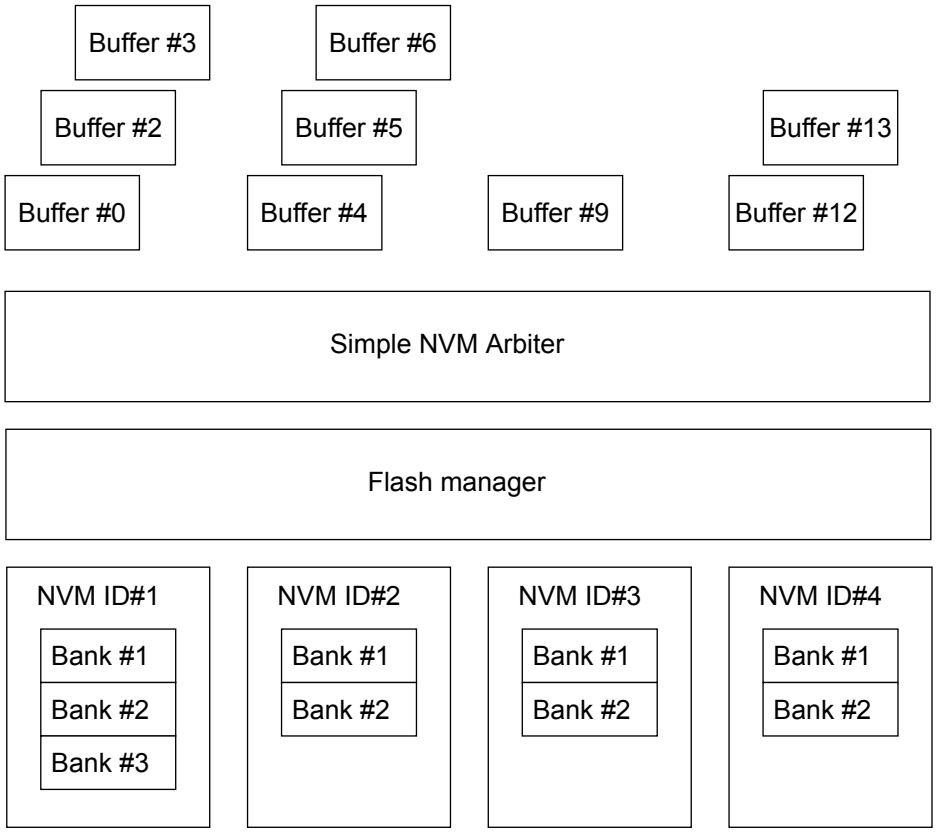

The user can register up to four buffers in one NVM. During the write operation, all the registered buffers are written in flash memory, that is, a whole bank update. During restore operation, the user restores only one identified buffer.

The classic use of the simple NVM arbiter is shown below:

Figure 11. **Simple NVM arbiter schema**

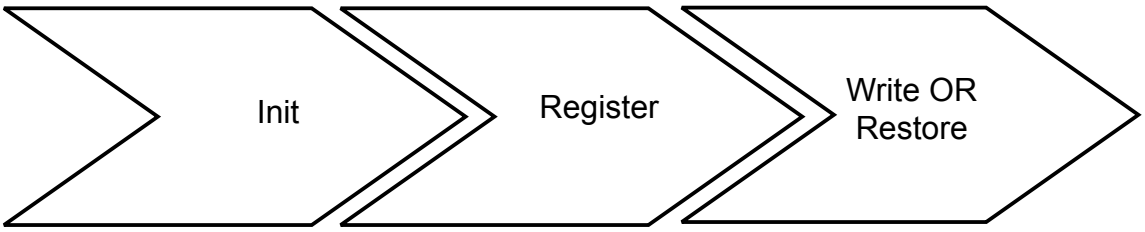

1. Initialize the SNVMA.
2. Register the buffer(s).
3. Execute, as many times as required, the write or restore operation.

3.3.5.3 *Going further*

For more information on the interface, using the interface, and for examples, refer to the wiki system page.

3.3.6 Trace management: advanced trace

An application often needs to manage a system of trace for debug purposes or to communicate with another system (console, monitoring, etc.). The advanced trace is a simple module designed to provides all the trace services.

This module has been designed with three major constraints:

- Low impact on the real time execution
- No specific associated processing, the trace evacuations are managed in the background
- Low footprint

3.3.6.1 *Concepts*

The advanced trace is split into two layers:

- Advanced trace services:
  API that regroups all the services available for traces.
- Trace interface:
  Interface between the hardware and the advance trace services layer.

Advanced trace services

The advanced trace offers interesting features to users for trace debugging or logging:

- Trace FIFO management:
  The advanced trace module manages a circular FIFO (the size is customizable) to store all the application data before pushing them on the media. The module guarantees the data integrity and parallel access to the trace services.
- Unchunk mode:
  The unchunk mode is an optional mode that can be enabled inside the utilities' configuration header. The goal of this mode is to guarantee that a frame is not split inside two transfers.
- Timestamp management:
  The timestamp management allows the user to add a timestamp buffer ahead of the frame to output. However, the following steps are necessary:
  - Create a callback function that returns a timestamp buffer and its buffer size.
  - Use the adequate functions that allow timestamp management. This mode can be activated within the utilities' configuration header.
- Verbose level and region management:
  The verbose level and the region are two different levels used to determine whether a frame can be sent or not. If one of these two conditions fails, the frame is discarded.
  Verbose level:
  - The user can define the current verbose level. The default setup is 0, meaning only value frames with a verbose level equal to zero are sent through the advanced trace. If the application verbose level is lower or equal to the current verbose level, the frame is sent or otherwise discarded.

  Region:
  - The user can define the region mask. The default setup is 0, meaning only the application frames with a region equal to zero are sent through the advanced trace. If the application region value is equal to zero or if the value is aligned with the mask region, the frame is sent or otherwise discarded.
- Overrun management:
  This feature aims to indicate when a frame has been discarded due to a full FIFO. The switch enabled mode is not enough to make it functional, the user must provide a callback function that returns an overrun frame and the size of this frame.
  This feature can be activated within the utilities' configuration header.

Trace interface

The trace interface handles the hardware interface. The only requirements for the hardware are to provide characteristics such as:

- Ability to send data (8-bit data)
- Ability able to receive data

This hardware layer has been built to match with the UART, but there are other hardware IPs compatible with these requirements: USB, I²C, SPI, etc. The choice is managed by the application designer according to the available resources and the application requirements.

3.3.6.2 *Going further*

For more information on the interface, using the interface, and for examples, refer to the wiki system page.

3.3.7 Real time SW debug

Debug on GPIO allows the user to get real time debugging traces of the application. Debug on GPIO is present at any relevant places. It concerns:

- Start/end of interrupt.
- Start/end of each background process.
- Identification of specific procedures and machine states.

3.3.7.1 Concepts

The real time SW debug has a two stages in the pipeline architecture:

- Local layer: The first stage determines which local signal is linked to which global signal. The local signal association table is responsible for matching the local and global signals.
- Global layer: The second stage determines which global signal is linked to which GPIO pin. The global signal association table is responsible for the match between the global signals and the GPIO pin.

Figure 12. Real time SW debug

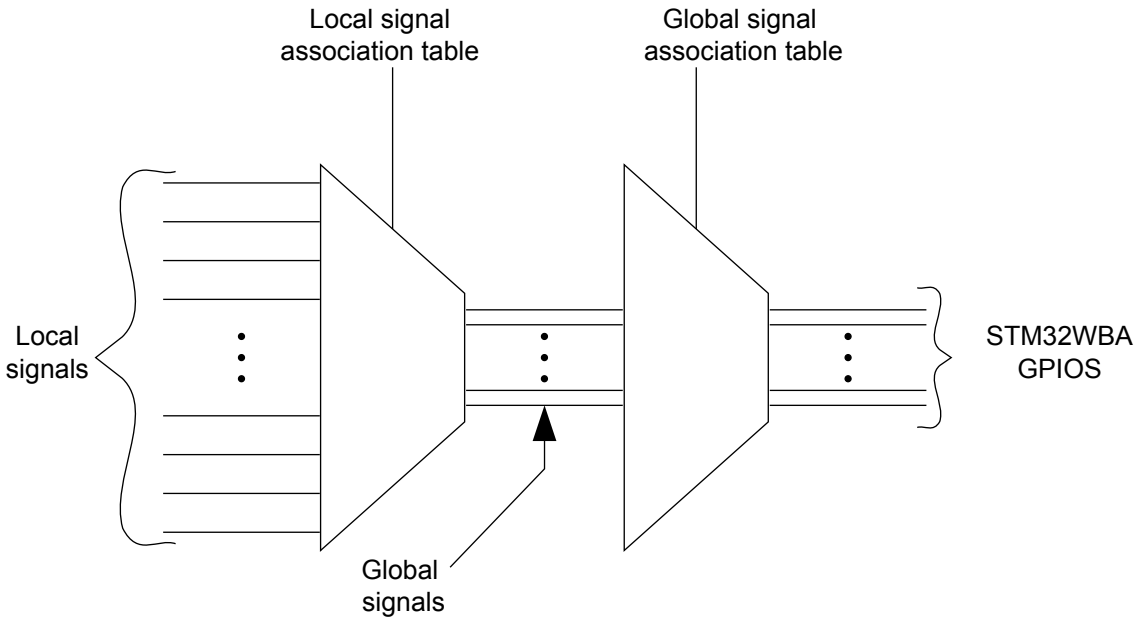

Signal selection

Debug signals are divided into different categories regarding the concerned layer:

- System SoC signals (interrupts, system services, etc.)
- Link layer signals
- MAC signals
- Host stack signals (BLE, Open Thread, Zigbee, Matter)
- Application signals

As some of these layers are delivered in library format, the debug signal selection cannot be done in the SW component itself.

- The desired signal is selected in a general configuration file.
- In every SW component that supports RT SW debug, all the **local** debugging signals are used. Regarding the **global** signal selection, the **global** signal is used or not.

Note: The local signals in different modules can have the same ID/number. Therefore, it is necessary to associate the local signals (possibly all used) to the global ones (effectively used).

GPIO configuration

GPIO associated to debug signal should be entirely configurable.

For optimization purposes, there is no dedicated callback for getting the associated GPIO at runtime (and no registering). The goal is to reduce code size and processing time (not modify real time on complicated use cases).

Note: It is necessary to associate the SW signal to the desired GPIO at compilation time.

3.3.7.2 *Going further*

For more information on the interface, using the interface, and for examples, refer to the wiki system page.

3.3.8 FUOTA

The Firmware Update Over the Air (FUOTA) allows the user to update the application during runtime without any wire connection.

3.3.8.1 *Concepts*

The FUOTA architecture is based on two components:

- OTA application: This component is responsible for downloading the raw data of the new application version. It can be based on any short-range wireless protocol handled by the STM32WBA.
- Application install manager application: This component is responsible for updating the firmware. It substitutes the older version with the new one at the restart. When there is no firmware update to perform, the boot manager application simply acts as a normal boot manager. This application part cannot be updated by OTA.

Memory mapping

The FUOTA memory mapping is organized along four fixed and presized regions:

- Application install manager region: This region starts at the boot address of the CPU. It contains the application install manager binary with its own vector table.
- NVM/User data region: This region handles the NVM and user data. These parts can be split into two if necessary. They can be positioned anywhere in the memory mapping.
- Active slot region: The active slot region is composed of the user app binaries, for example, user application, OTA application, and wireless protocol stack (including its vector table). Its size is fixed and is the same as the download slot. The active slot region address is fixed, but can be placed anywhere.
- Download slot region: This region has the same amount of space as the active slot region. It is designed to host the raw data of the new application version. The download slot region address is fixed, but can be placed anywhere.

Figure 13. FUOTA overview

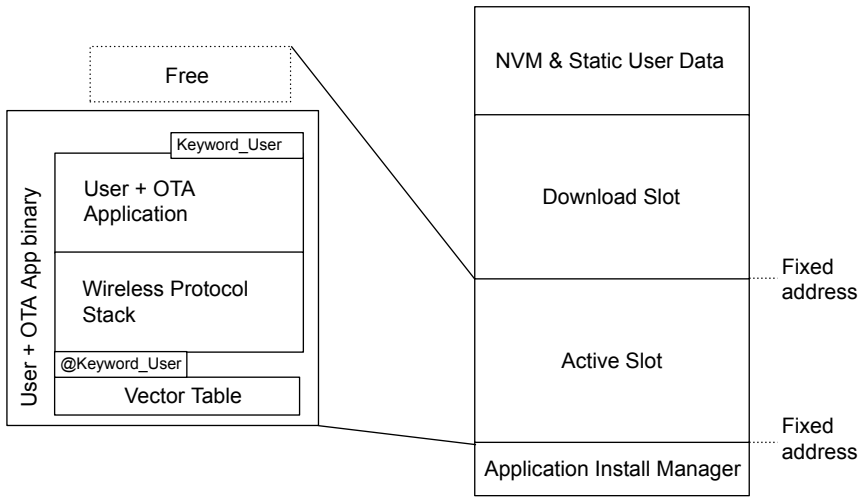

Active slot

As previously stated, the user application part is composed of different components:

- User + OTA application: This is a single binary implementing both the user and OTA applications. The OTA application is responsible for loading the update firmware in the download slot. This binary (including both the user and OTA application) can be updated with OTA.
- Wireless protocol stack: This is the protocol stack of the application.
- Keyword user: The keyword User is a milestone that indicates the end of the active slot binaries. It ensures that the firmware has been fully loaded. The value of the keyword user is a number, but can be a CRC. It is up to the user to implement a postscript for the integrity computation. The keyword user address is placed inside a variable located right after the active slot vector table end. This variable has a fixed address.

3.3.8.2 *Going further*

For more information on the interface, using the interface, and for examples, refer to the FUOTA wiki page

4 System initialization

4.1 General concepts

All applications have a four layer architecture:

- `main.c` First entry point, all hardware configuration that is common to any application
- `app_entry.c` All SW configuration and implementation that is common to any application
- `app_ble.c/app_xxx.c` Application dedicated files
- BLE LL/YYY LL: Link layer interface

The system is initialized using the steps below.

Figure 14. General architecture

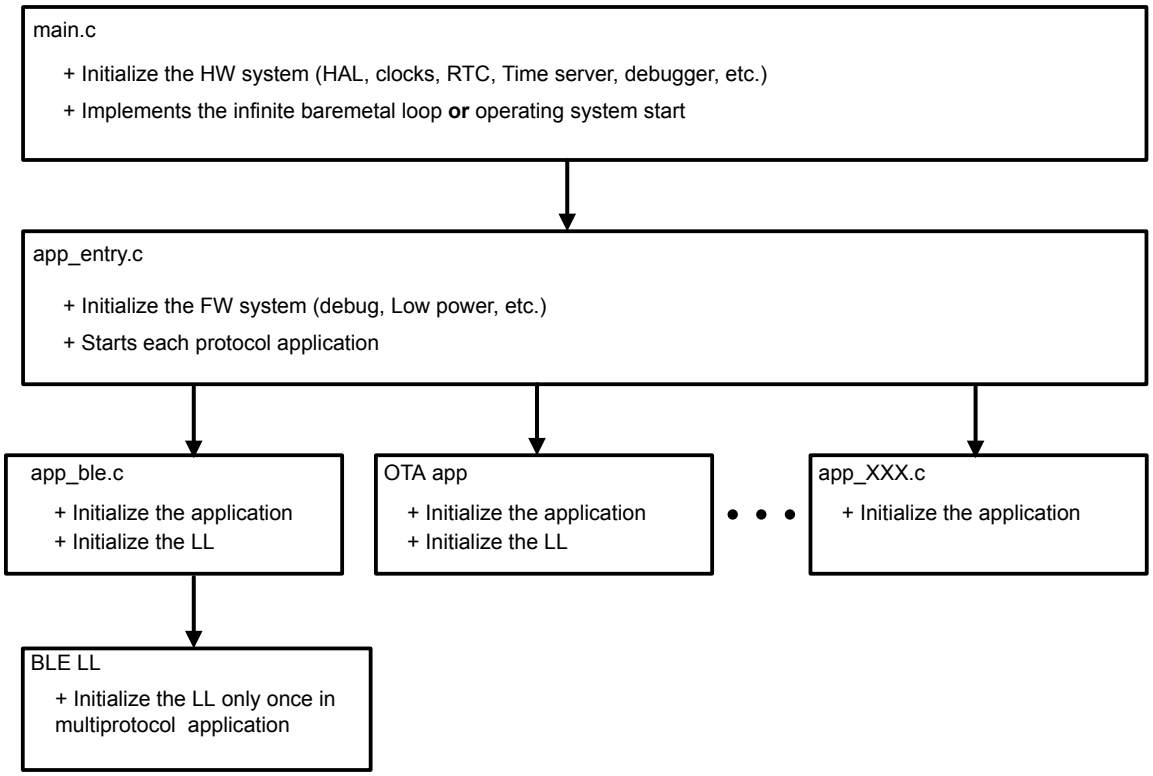

For detailed information and examples, refer to the wiki page on the system initialization,

5 Design of a short-range wireless application

5.1 Bluetooth® Low Energy

This chapter aims to highlight a generic Bluetooth® Low Energy application architecture on a STM32WBA.

5.1.1 Overview

A BLE application is composed of different components and modules, but it mostly depends on BLE profiles. These profiles define the capacities and the purpose of the BLE application.

A BLE profile is a collection of one or more BLE services. It also defines the behavior of its services mainly in the context of GAP management.

A BLE profile is composed of three components:

- A GATT management module
- An application service for each selected BLE service
- A GAP management module

The figure below deals with the design of a generic BLE application:

Figure 15. Generic BLE application design

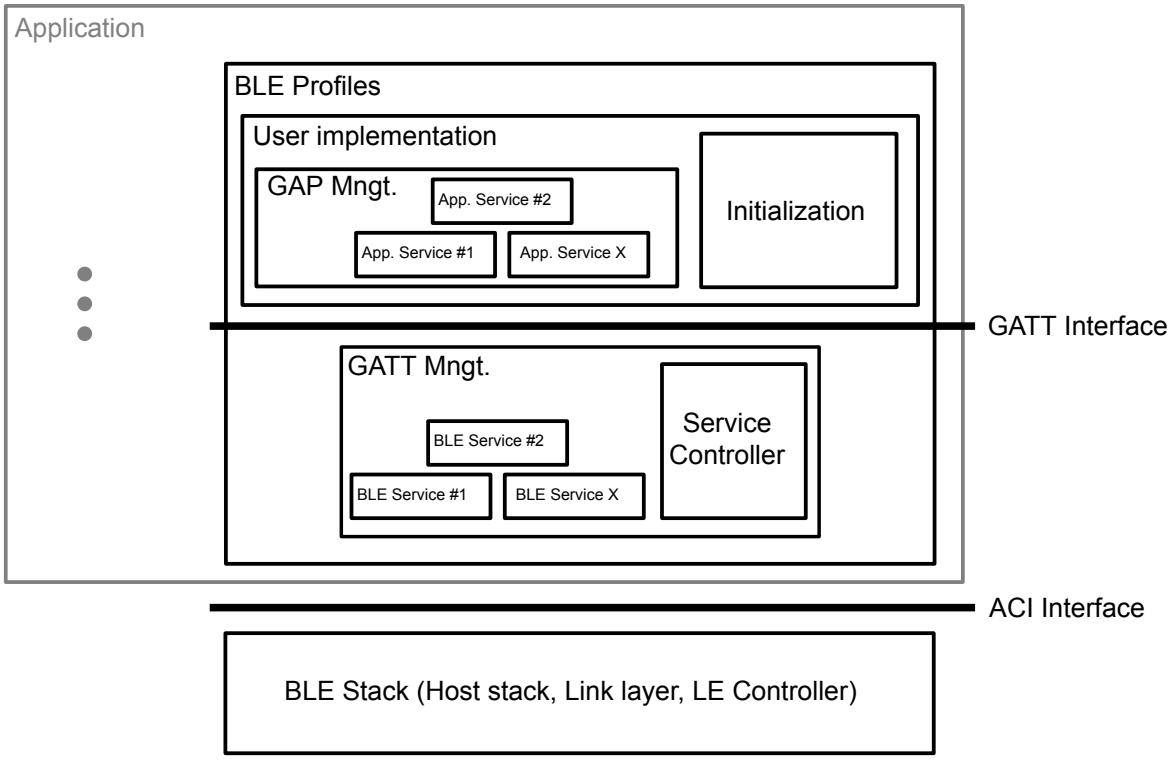

5.1.1.1 *Layer descriptions*

The BLE firmware is built on a layering architecture. For BLE configuration, the user can select one of the two available interfaces for the application.

ACI interface

The ACI interface is a specific STMicroelectronics interface provided to access all the available features in the BLE host stack. This interface follows two rules:

- Each command is blocking and does not return until the response is received.
- It reports to the user all asynchronous events received from the BLE host stack.

GATT interface

The GATT interface is also a specific STMicroelectronics interface provided for use with the ACI interface. Its purpose is to ease management of the BLE services.

When this interface is used, the ACI GATT interface should no longer be used. There are two exceptions:

- GATT initialization
- Application already implements its BLE service outside the GATT management module

GATT management

The GATT management has two components:

- Service controller:
  - Manages the BLE service initializations, and the events received by the BLE host stack
  - Initializes all enabled BLE services provided in the GATT management module
  - Provides an interface so that the application may initialize its own BLE service implemented outside the GATT management module
  - Manages all events reported by the BLE host stack
  - Forwards the events that are not GATT events to the application
  - Forwards GATT events to each registered BLE service until one of these services reports that the GATT event has been processed

Note: If no registered BLE services is concerned with the GATT event, it is forwarded to the application.

- BLE services:
  - Can either be specified by the BT SIG or custom to STMicroelectronics
  - Each BLE service is implemented in a dedicated file named with the three or four letters, as specified by the BT SIG
  - Each BLE service provides the same features:
    - Creates the service in the BLE stack
    - Implements all mechanisms (that do not require user input) for GATT event reception
    - Provides a simplified interface for value update/receive from the server
    - Adds its characteristics

User implementation

This is the application part of BLE profile creation. The user should implement three components to build a BLE profile:

- Initialization:
  Initializes all the SW modules of the application, for example, GAP initialization, Service controller initialization, etc.
- GAP management:
  Manages everything that is not related to a GATT event
  Has the BLE connection management role
  Receives notification from the application services to be compliant with a BLE profile
- Application services:
  Application Service is the application management of a BLE Service. There is one application service file per BLE Service.
  For an application service, all GATT events (and only those events) are handled in this file.
  Each application service provides at least:
  - An initialization function
  - An implementation of the corresponding BLE service notification function
  - For some profiles, a notification to the GAP management

5.1.2 Configuration

The BLE stack is a two-layer block. One is for the host and the other is for the controller.

Figure 16. **BLE stack**

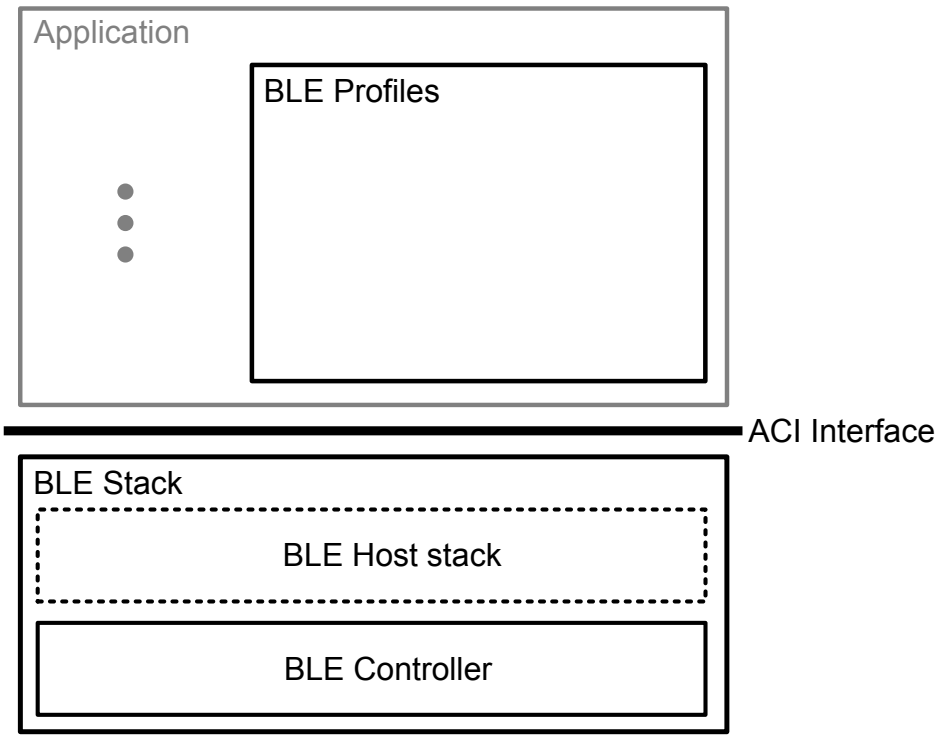

Both parts features are provided as two distinct libraries. These two libraries are mandatory for any BLE project, one relying on the other.

5.1.2.1 *BLE host stack*

The BLE host stack library is configurable in different ways and is provided in four different variants:

- Two variants containing the BLE host stack and the STMicroelectronics LL controller:
  - Full stack (`stm32wba_ble_stack_full.a`)
  - Basic stack (`stm32wba_ble_stack_basic.a`)
- Two variants without the BLE host stack, with the STMicroelectronics LL controller only:
  - Link layer only stack (`stm32wba_ble_stack_llo.a`)
  - Link layer only basic stack (`stm32wba_ble_stack_llobasic.a`)

Depending on the application purpose(s), the user can select one or another configuration. Each variant has its own features and details. The differences between the configurations are listed below.

- Full stack (`stm32wba_ble_stack_full.a`):
  Contains all the legacy stack supported features plus the extended advertising, caching GATT, ACI HCI flow control, isochronous support for audio, L2CAP connection-oriented channels and host enabled with GATT, GAP, and L2CAP features.

Note: This BLE stack variant can be configured to run in controller only mode (host stack is bypassed).

- Basic stack (`stm32wba_ble_stack_basic.a`):
  Contains only basic supported features or BLE legacy features without extended advertising, neither caching GATT, ACI HCI flow control, isochronous support, nor L2CAP connection-oriented channels. However, the host is here enabled and supports all the basic GATT, GAP, and L2CAP features.
- Link layer only stack (`stm32wba_ble_stack_llo.a`):
  Contains all the features supported by the full stack but does not include the host part (GATT, GAP and L2CAP features).
- Link layer only basic stack (`stm32wba_ble_stack_llobasic.a`):
  Contains all the features supported by the basic stack but does not include the host part (GATT, GAP, and L2CAP features).

5.1.2.2 *Controller stack*

The controller stack library is only composed of the link layer features. This library is provided in a basic and a full version, for example, LinkLayer_BLE_Full_lib.a and LinkLayer_BLE_Basic_lib.a.

5.1.2.3 *Footprints*

The library footprints are listed in Table 2. Library footprints

Table 2. Library footprints

Tiny Blue Fish memory footprint Version: STEP3.1 (v1.0.0)		Flash memory footprint (KB)				RAM memory footprint (KB)			
BLE application	Library used	Total size	Application	BLE stack	SNPS LL	Total size	Application	BLE stack	SNPS LL
BLE_Heartrate	basic	163.4	43	42.5	77.9	36.6	22.3	0.9	13.4
BLE_Beacon		168.2	51.7	39.5	77	40.3	25.9	0.9	13.4
BLE_p2pServer_Ext	full	285	58.8	51.5	174.6	56.6	32.5	1.1	23
BLE_TransparentMode		333.6	39.6	81.5	212.4	82.7	58.6	1.1	23

Note: The information in the table above is provided as an example and may not represent the current state of the libraries' footprints. Always refer to the wiki or to the release page for the most up to date information.

5.1.3 Going further

For more information on the BLE initialization procedure, refer to to the wiki page.

6 Requirements/guidelines

This chapter contains detailed requirements and best practices for the optimal design of a short-range wireless application.

6.1 System

6.1.1 Interrupts

Some restrictions should be followed concerning interruption management. There are some events that cannot be left over, in terms of scheduling. For instance, BLE radio interrupts, for example, LL_High_ISR, must not be delayed for more than 40 µs when in use.

Therefore, the user must organize the interrupt levels as follows:

Table 3. Interrupt levels

Interrupt name	Priority(1)
LL_High_ISR (when the radio is active)	0
RCC (HSERDY and PLL1RDY)	1
User system and peripheral interrupts	2 - 4
LL_High_ISR (when radio is inactive)	5
User system and peripheral interrupts	6 – 12
Systick	13
LL_Low_ISR	14
GPDMA_CHx	15

1. *The lower the number, the higher the priority*

Additionally, the LL_Low_ISR runs in an SW interrupt handler. Any unused NVIC ISR can be used to host the LL_Low_ISR.

6.1.2 Performances

There are a few rules to follow for nominal clock configuration:

- Default firmware run speed: 16 MHz - HSE/2 Range2
- Radio event period, run speed: at least 32MHz–HSE Range1
- PLL required run speed: 100 MHz

6.1.3 Processes

The design of each process must give the opportunity for each of them to be executed within a correct time frame window.

Although this is not a strict requirement and some specific processes may not comply with these timings, the recommendations are:

- The sum of each process timing execution should not exceed 300 ms.
- The time execution of each process should not exceed 30 ms.

These timing requirements are especially relevant in a bare metal implementation, as there is no way to stop a running process to execute another one.

When using an operating system, these timing requirements still apply to improve the dynamic between processes that have the same priority.

Revision history

Table 4. **Document revision history**

Date	Version	Changes
08-Mar-2023	1	Initial release.

Contents

List of figures

List of tables