

embOS

Real-Time Operating System

CPU & Compiler specifics
for ARM using GCC

Document: UM01052
Software Version: 5.20.0.0
Revision: 0
Date: May 19, 2025



A product of SEGGER Microcontroller GmbH

www.segger.com

Disclaimer

The information written in this document is assumed to be accurate without guarantee. The information in this manual is subject to change for functional or performance improvements without notice. SEGGER Microcontroller GmbH (SEGGER) assumes no responsibility for any errors or omissions in this document. SEGGER disclaims any warranties or conditions, express, implied or statutory for the fitness of the product for a particular purpose. It is your sole responsibility to evaluate the fitness of the product for any specific use.

Copyright notice

You may not extract portions of this manual or modify the PDF file in any way without the prior written permission of SEGGER. The software described in this document is furnished under a license and may only be used or copied in accordance with the terms of such a license.

© 2010-2025 SEGGER Microcontroller GmbH, Monheim am Rhein / Germany

Trademarks

Names mentioned in this manual may be trademarks of their respective companies.

Brand and product names are trademarks or registered trademarks of their respective holders.

Contact address

SEGGER Microcontroller GmbH

Ecolab-Allee 5
D-40789 Monheim am Rhein

Germany

Tel. +49 2173-99312-0
Fax. +49 2173-99312-28
E-mail: support@segger.com*
Internet: www.segger.com

*By sending us an email your (personal) data will automatically be processed. For further information please refer to our privacy policy which is available at <https://www.segger.com/legal/privacy-policy/>.

Manual versions

This manual describes the current software version. If you find an error in the manual or a problem in the software, please inform us and we will try to assist you as soon as possible. Contact us for further information on topics or functions that are not yet documented.

Print date: May 19, 2025

Software	Revision	Date	By	Description
5.20.0.0	0	250519	MC	New software version.
5.18.0.0	0	221104	MM	New software version.
5.16.1.1	0	220707	MM	Chapter "Libraries" and "VFP and NEON support" updated.
5.16.1.0	0	220210	MM	New software version. Chapter "CPU and compiler specifics" updated.
5.14.0.0	0	210719	MM	New software version. Chapter "VFP and NEON support" updated.
5.10.2.0	0	200911	MM	New software version.
5.02	0	180625	TS	New software version.
4.40	0	180322	MC	New software version.
4.32	0	170125	TS	New Cache API description added.
4.22	0	160620	MC	New software version.
4.10b	0	150618	TS	New software version.
4.04a	1	150217	TS	Chapter "Naming conventions for prebuilt libraries" updated.
4.04a	0	150123	TS	New software version
3.90a	0	140430	AW	New software version.
3.90	0	140306	AW	New software version.
3.88h	0	131223	MC	Initial version.

About this document

Assumptions

This document assumes that you already have a solid knowledge of the following:

- The software tools used for building your application (assembler, linker, C compiler).
- The C programming language.
- The target processor.
- DOS command line.

If you feel that your knowledge of C is not sufficient, we recommend *The C Programming Language* by Kernighan and Richie (ISBN 0--13--1103628), which describes the standard in C programming and, in newer editions, also covers the ANSI C standard.

How to use this manual

This manual explains all the functions and macros that the product offers. It assumes you have a working knowledge of the C language. Knowledge of assembly programming is not required.

Typographic conventions for syntax

This manual uses the following typographic conventions:

Style	Used for
Body	Body text.
Keyword	Text that you enter at the command prompt or that appears on the display (that is system functions, file- or pathnames).
Parameter	Parameters in API functions.
Sample	Sample code in program examples.
Sample comment	Comments in program examples.
Reference	Reference to chapters, sections, tables and figures or other documents.
GUIElement	Buttons, dialog boxes, menu names, menu commands.
Emphasis	Very important sections.

Table of contents

1	Using embOS	9
1.1	Installation	10
1.2	First Steps	11
1.3	The example application OS_StartLEDBlink.c	12
1.4	Stepping through the sample application	13
2	Build your own application	16
2.1	Introduction	17
2.2	Required files for an embOS	17
2.3	Change library mode	17
2.4	Select another CPU	17
3	Libraries	18
3.1	Naming conventions for prebuilt libraries	19
4	CPU and compiler specifics	20
4.1	Standard system libraries	21
4.2	Interrupt and thread safety	21
4.3	Thread-Local Storage TLS	23
4.3.1	API functions	23
4.3.1.1	OS_TLS_Set()	24
4.3.1.2	OS_TLS_SetTaskContextExtension()	25
5	Stacks	27
5.1	Task stack	28
5.2	System stack	28
5.3	Interrupt stack	28
5.4	Stack specifics	30
6	Interrupts	31
6.1	What happens when an interrupt occurs?	32
6.2	Defining interrupt handlers in C	32
6.3	Interrupt handling without vectored interrupt controller	32
6.4	Interrupt handling with vectored interrupt controller	34
6.4.1	API functions	34
6.4.1.1	OS_ARM_InstallISRHandler()	35
6.4.1.2	OS_ARM_EnableISR()	36
6.4.1.3	OS_ARM_DisableISR()	37

6.4.1.4	OS_ARM_ISRSetPrio()	38
6.4.1.5	OS_ARM_ClearPendingFlag()	39
6.4.1.6	OS_ARM_IsPending()	40
6.4.1.7	OS_ARM_AssignISRSource()	41
6.4.1.8	OS_ARM_EnableISRSource()	42
6.4.1.9	OS_ARM_DisableISRSource()	43
6.4.1.10	OS_ARM_SetISRCfg()	44
6.4.1.11	OS_ARM_SetVBAR()	45
6.5	Interrupt-stack switching	46
6.6	Fast Interrupt (FIQ)	47
7	MMU/MPU and cache support	48
7.1	Introduction	49
7.2	MMU handling for ARMv5/ARMv7-A CPUs	50
7.2.1	API functions	50
7.2.1.1	OS_ARM_MMU_InitTT()	51
7.2.1.2	OS_ARM_MMU_AddTTEntries()	52
7.2.1.3	OS_ARM_MMU_Enable()	54
7.2.1.4	OS_ARM_MMU_GetVirtualAddr()	55
7.2.1.5	OS_ARM_MMU_v2p()	56
7.3	MPU handling for ARMv7-R CPUs	57
7.3.1	API functions	57
7.3.1.1	OS_ARM_MPU_AddEntry()	58
7.3.1.2	OS_ARM_MPU_Enable()	60
7.3.1.3	OS_ARM_MPU_GetMinRegionSize()	61
7.3.1.4	OS_ARM_MPU_GetNumRegions()	62
7.3.1.5	OS_ARM_MPU_Init()	63
7.4	Cache handling for ARMv5/ARMv7 CPUs	64
7.4.1	API functions	64
7.4.1.1	OS_ARM_ICACHE_Enable()	65
7.4.1.2	OS_ARM_ICACHE_Invalidate()	66
7.4.1.3	OS_ARM_DCACHE_Enable()	67
7.4.1.4	OS_ARM_DCACHE_Invalidate()	68
7.4.1.5	OS_ARM_DCACHE_Clean()	69
7.4.1.6	OS_ARM_DCACHE_CleanRange()	70
7.4.1.7	OS_ARM_DCACHE_InvalidateRange()	71
7.4.1.8	OS_ARM_CACHE_Sync()	72
7.4.1.9	OS_ARM_AddL2Cache()	73
7.4.1.10	OS_ARM_CACHE_GetLineSize()	74
7.5	MMU and cache handling program sample	75
7.6	MPU and cache handling program sample	76
8	VFP and NEON support	77
8.1	Introduction	78
8.2	Using embOS libraries with VFP/NEON support	78
8.3	Using the VFP/NEON unit in interrupt service routines	78
9	Technical data	79
9.1	Resource Usage	80

Chapter 1

Using embOS

1.1 Installation

This chapter describes how to get started with embOS. You should follow these steps to become familiar with embOS.

embOS is shipped as a zip-file in electronic form. To install it, you should extract the zip-file to any folder of your choice while preserving its directory structure (i.e. keep all files in their respective sub directories). Ensure the files are not read-only after extraction. Assuming that you are using an IDE to develop your application, no further installation steps are required.

Note

The projects at `/Start/BoardSupport/<DeviceManufacturer>/<Board>` assume a relative location for the `/Start/Lib` and `/Start/Inc` folders. If you copy a BSP folder to another location, you will need to adjust the include paths of the project accordingly.

At `/Start/BoardSupport/<DeviceManufacturer>/<Board>` you should find several example start projects, which you may adapt to write your application. To do so, follow the instructions of section *First Steps* on page 11.

In order to become familiar with embOS, consider using the example projects (even if you will not use the IDE for application development).

If you do not or do not want to work with an IDE, you may copy either all library files or only the library that is used with your project into your work directory. embOS does not rely on an IDE, but may be used without an IDE just as well, e.g. using batch files or a make utility.

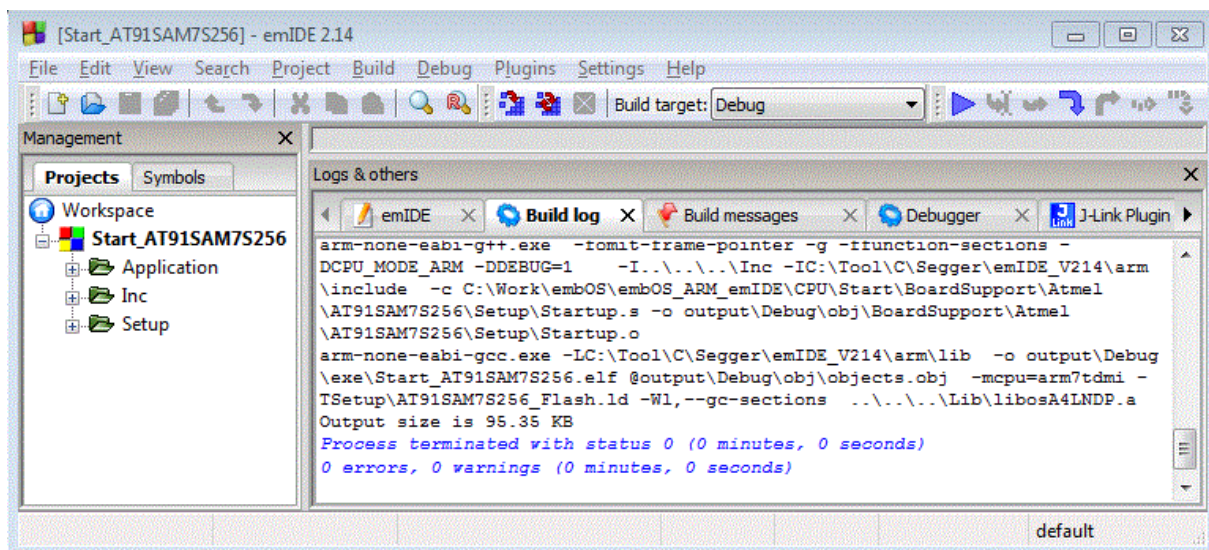
1.2 First Steps

After installation of embOS, you can create your first multitasking application. You received several ready-to-go sample workspaces and projects as well as all required embOS files inside the subfolder `Start`. The subfolder `Start/BoardSupport` contains the workspaces and projects, sorted into manufacturer- and board-specific subfolders. It is a good idea to use one of the projects as a starting point for any application development.

To get your new application running, you should:

- Create a directory for your development.
- Copy the whole `Start` folder from your embOS shipment into the directory.
- Clear the read-only attribute of all files in the copied `Start` folder.
- Open one sample workspace/project in
`Start/BoardSupport/<DeviceManufacturer>/<Board>` with your IDE (for example, by double clicking it).
- Build the project. It should be built without any error or warning messages.

After building the project of your choice, the screen should look like this:



For additional information, you should open the `ReadMe.txt` file that is part of every BSP. It describes the different configurations of the project and, if required, gives additional information about specific hardware settings of the supported evaluation board(s).

1.3 The example application OS_StartLEDBlink.c

The following is a printout of the example application OS_StartLEDBlink.c. It is a good starting point for your application (the actual file shipped with your port of embOS may differ slightly).

What happens is easy to see:

After initialization of embOS, two tasks are created and started. The two tasks get activated and execute until they run into a delay, thereby suspending themselves for the specified time, and eventually continue execution.

```

/*****
*                               SEGGER Microcontroller GmbH                               *
*                               The Embedded Experts                                   *
*****/

----- END-OF-HEADER -----
File      : OS_StartLEDBlink.c
Purpose   : embOS sample program running two simple tasks, each toggling
            an LED of the target hardware (as configured in BSP.c).
*/

#include "RTOS.h"
#include "BSP.h"

static OS_STACKPTR int StackHP[128], StackLP[128]; // Task stacks
static OS_TASK      TCBHP, TCBLP;                 // Task control blocks

static void HPTask(void) {
    while (1) {
        BSP_ToggleLED(0);
        OS_TASK_Delay(50);
    }
}

static void LPTask(void) {
    while (1) {
        BSP_ToggleLED(1);
        OS_TASK_Delay(200);
    }
}

/*****
*
*      main()
*
*****/
int main(void) {
    OS_Init(); // Initialize embOS
    OS_InitHW(); // Initialize required hardware
    BSP_Init(); // Initialize LED ports
    OS_TASK_CREATE(&TCBHP, "HP Task", 100, HPTask, StackHP);
    OS_TASK_CREATE(&TCBLP, "LP Task", 50, LPTask, StackLP);
    OS_Start(); // Start embOS
    return 0;
}

/***** End of file *****/

```

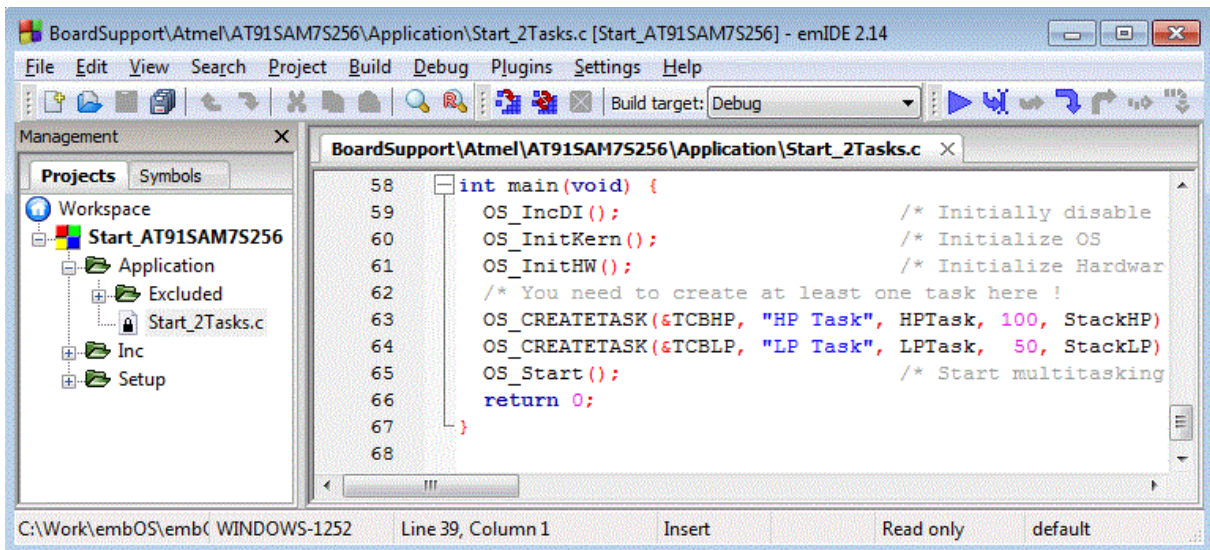
1.4 Stepping through the sample application

When starting the debugger, you will see the `main()` function (see example screenshot below). The `main()` function appears as long as project option `Run to main` is selected, which it is enabled by default. Now you can step through the program.

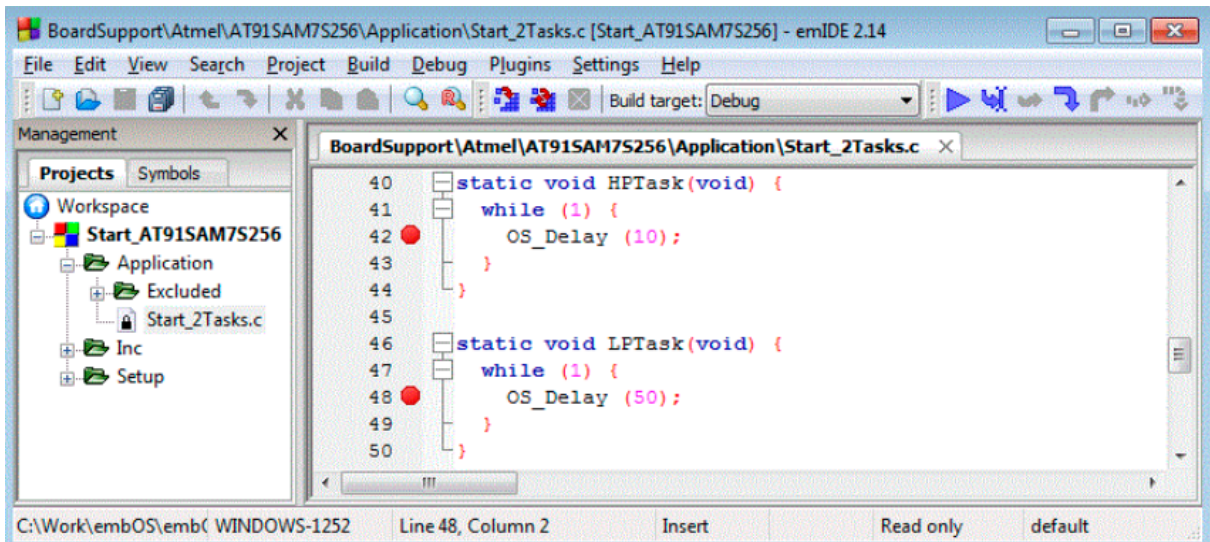
`OS_Init()` is part of the embOS library and written in assembler; you can therefore only step into it in disassembly mode. It initializes the relevant OS variables.

`OS_InitHW()` is part of `RTOSInit.c` and therefore part of your application. Its primary purpose is to initialize the hardware required to generate the system tick interrupt for embOS. Step through it to see what is done.

`OS_Start()` should be the last line in `main()`, because it starts multitasking and does not return.

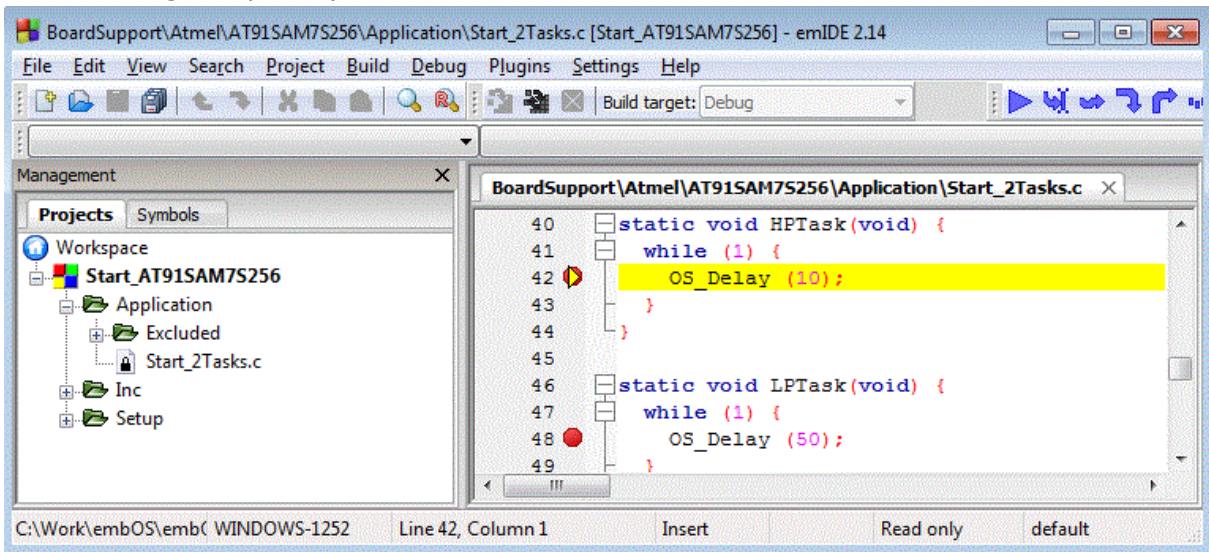


Before you step into `OS_Start()`, you should set two breakpoints in the two tasks as shown below.

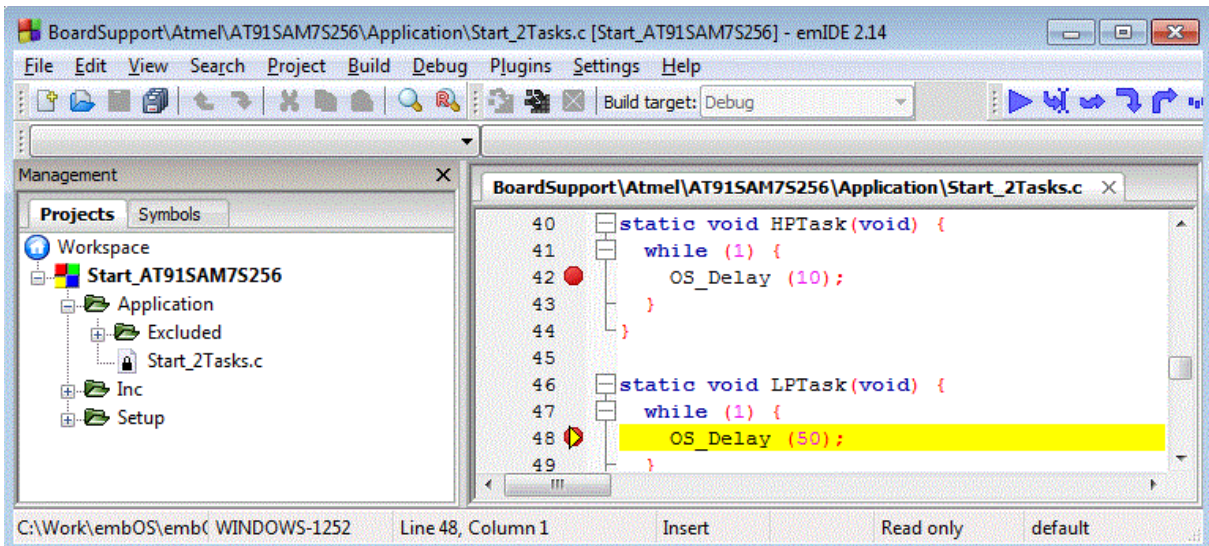


As `OS_Start()` is part of the embOS library, you can step through it in disassembly mode only.

Click **GO**, step over `OS_Start()`, or step into `OS_Start()` in disassembly mode until you reach the highest priority task.

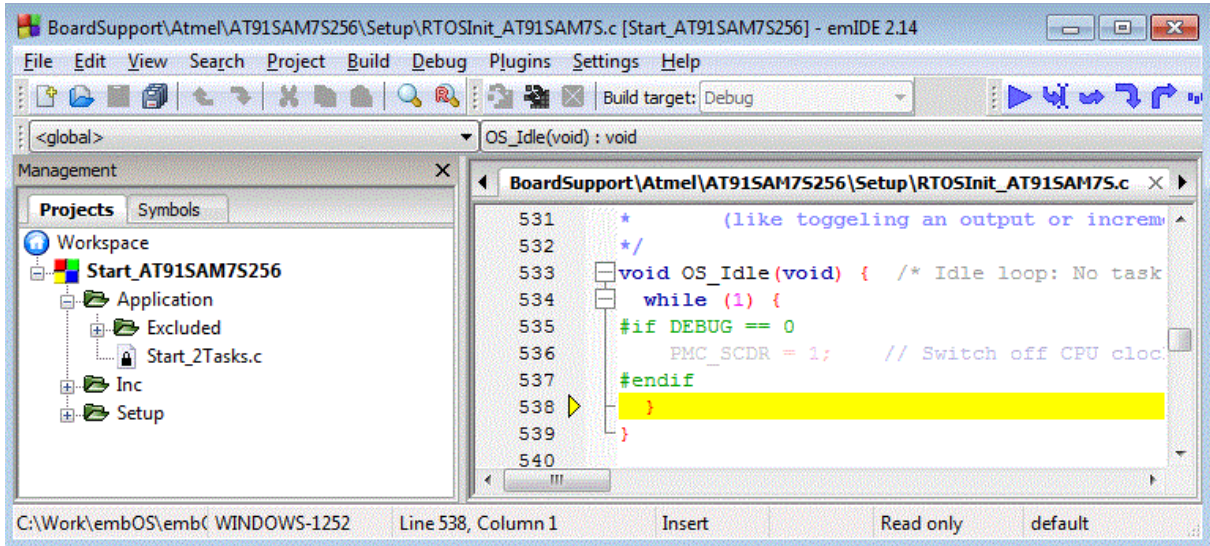


If you continue stepping, you will arrive at the task that has lower priority:



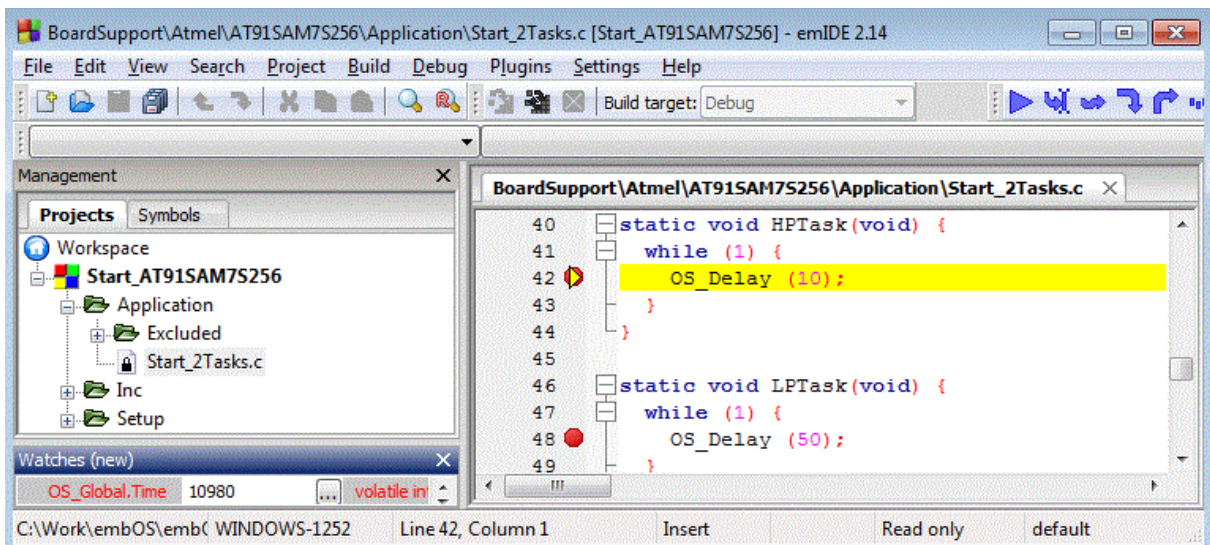
Continue to step through the program, there is no other task ready for execution. embOS will therefore start the idle-loop, which is an endless loop always executed if there is nothing else to do (no task is ready, no interrupt routine or timer executing).

You will arrive there when you step into the `OS_Task_Delay()` function in disassembly mode. `OS_Idle()` is part of `RTOSInit.c`. You may also set a breakpoint there before stepping over the delay in `LPTask()`.



If you set a breakpoint in one or both of our tasks, you will see that they continue execution after the given delay.

As can be seen by the value of embOS timer variable `OS_Global.Time`, shown in the Watch window, `HPTask()` continues operation after expiration of the delay.



Chapter 2

Build your own application

2.1 Introduction

This chapter provides all information to set up your own embOS project. To build your own application, you should always start with one of the supplied sample workspaces and projects. Therefore, select an embOS workspace as described in chapter *First Steps* on page 11 and modify the project to fit your needs. Using an embOS start project as starting point has the advantage that all necessary files are included and all settings for the project are already done.

2.2 Required files for an embOS

To build an application using embOS, the following files from your embOS distribution are required and have to be included in your project:

- **RTOS.h** from the directory `.\Start\Inc`. This header file declares all embOS API functions and data types and has to be included in any source file using embOS functions.
- **RTOSInit*.c** from one target specific `.\Start\BoardSupport\<Manufacturer>\<MCU>` subfolder. It contains hardware-dependent initialization code for embOS. It initializes the system timer interrupt but can also initialize or set up the interrupt controller, clocks and PLLs, the memory protection unit and its translation table, caches and so on.
- **OS_Error.c** from one target specific subfolder `.\Start\BoardSupport\<Manufacturer>\<MCU>`. The error handler is used only if a debug library is used in your project.
- One **embOS library** from the subfolder `.\Start\Lib`.
- Additional CPU and compiler specific files may be required according to CPU.

When you decide to write your own startup code or use a low level `init()` function, ensure that non-initialized variables are initialized with zero, according to C standard. This is required for some embOS internal variables. Your `main()` function has to initialize embOS by calling `OS_Init()` and `OS_InitHW()` prior to any other embOS functions that are called.

2.3 Change library mode

For your application you might want to choose another library. For debugging and program development you should always use an embOS debug library. For your final application you may wish to use an embOS release library or a stack check library.

Therefore you have to select or replace the embOS library in your project or target:

- If your selected library is already available in your project, just select the appropriate project configuration.
- To add a library, you may add the library to the existing Lib group. Exclude all other libraries from your build, delete unused libraries or remove them from the configuration.
- Check and set the appropriate `OS_LIBMODE_*` define as preprocessor option and/or modify the `OS_Config.h` file accordingly.

2.4 Select another CPU

embOS contains CPU-specific code for various CPUs. Manufacturer- and CPU-specific sample start workspaces and projects are located in the subfolders of the `.\Start\BoardSupport` directory. To select a CPU which is already supported, just select the appropriate workspace from a CPU-specific folder.

If your CPU is currently not supported, examine all `RTOSInit.c` files in the CPU-specific subfolders and select one which almost fits your CPU. You may have to modify `OS_InitHW()`, the interrupt service routines for the embOS system tick timer and the low level initialization.

Chapter 3

Libraries

3.1 Naming conventions for prebuilt libraries

embOS is shipped with different pre-built libraries with different combinations of features. Note that not all combinations are available (e.g., there is no VFP support for Armv4).

The libraries are named as follows:

```
libos<CpuMode><Architecture><VFP_NEON><Endianness><Interwork><Libmode>.a
```

Parameter	Meaning	Values
CpuMode	CPU mode	A : ARM T : Thumb / Thumb2
Architecture	ARM architecture	4 : ARMv4 5 : ARMv5 7A : ARMv7-A 7R : ARMv7-R
VFP_NEON	Floating point / NEON support	: No hardware VFP support V16 : VFP-D16 with softfp floating-point ABI V16H: VFP-D16 with hardware floating-point ABI V32 : VFP-D32 with softfp floating-point ABI V32H: VFP-D32 with hardware floating-point ABI
Endianness	Byte order	B : Big endian L : Little endian
Interwork	Interwork support	N : Non interwork (ARMv4, ARMv5) I : Interwork
Libmode	embOS library mode	XR : Extreme Release R : Release S : Stack check SP : Stack check + profiling D : Debug DP : Debug + profiling + stack check DT : Debug + profiling + stack check + trace

Example

libosT7AV32LIDP.a is the library for an ARMv7-A core, thumb2 mode, support for VFP/NEON D32, little endian mode, interwork, with debug and profiling support.

Note

When updating from an earlier embOS version you might need to update to an embOS library with VFP/NEON support. For example, if you use libosT7ALIDP.a for an ARMv7A CPU with VFP/NEON unit you will get a linker error message like "undefined symbol OS_Init_VFPD32 referenced by symbol main".

This check avoids that the project and the used embOS library use different VFP/NEON settings. If your project settings allow the compiler to generate VFP/NEON instructions, an embOS library with VFP/NEON support like libosT7AV32LIDP.a must be used.

Chapter 4

CPU and compiler specifics

4.1 Standard system libraries

embOS for and GCC may be used with the Red Hat newlib C libraries for most of all projects without any modification.

Since heap management with newlib depends on a working implementation of an `__sbrk()` function, that implementation is provided with embOS inside the source module `OS_Syscalls.c`, which itself is included in the "Setup" subdirectory of every embOS start project. Using that source file requires the symbols `__heap_start__` and `__heap_end__` to be appropriately defined in the respective project's linker file.

Alternatively, if modification of the linker file is not feasible (e.g. when it is auto-generated by a project generator), the heap symbols may also be defined as additional linker flags. For example, if the linker file defines the symbols `LOWEST_HEAP_ADDRESS` and `HIGHEST_HEAP_ADDRESS`, the linker flags would need to be set as shown below:

```
--defsym=__heap_start__=LOWEST_HEAP_ADDRESS  
--defsym=__heap_end__=HIGHEST_HEAP_ADDRESS
```

Heap management and file operation functions of standard system libraries are not reentrant and require a special initialization or additional modules when used with embOS when those non-thread-safe functions are used from different tasks (refer to *Interrupt and thread safety* on page 21).

4.2 Interrupt and thread safety

Using embOS with specific calls to standard library functions (e.g. heap management functions) may require thread-safe system libraries if these functions are called from several tasks or interrupts. Newlib provides functions, which can be overwritten to implement a locking mechanism making the system library functions thread-safe.

The Setup directory in each embOS BSP contains the file `OS_ThreadSafe.c` which overwrites these functions. By default they disable and restore embOS interrupts to ensure thread safety in tasks, embOS interrupts, `OS_Idle()` and software timers. Zero latency interrupts are not disabled and therefore unprotected. If you need to call e.g. `malloc()` also from within a zero latency interrupt additional handling needs to be added. If you don't call such functions from within embOS interrupts, `OS_Idle()` or software timers, you can instead use thread safety for tasks only. This reduces the interrupt latency because a mutex is used instead of disabling embOS interrupts.

You can choose the safety variant with the macro `OS_INTERRUPT_SAFE`.

- When defined to 1 thread safety is guaranteed in tasks, embOS interrupts, `OS_Idle()` and software timers.
- When defined to 0 thread safety is guaranteed only in tasks. In this case you must not call e.g. heap functions from within an ISR, `OS_Idle()` or embOS software timers.

Alternatively, embOS delivers its own thread-safe functions for heap management. These are described in the embOS generic manual.

4.2.1 `__malloc_lock()`, lock the heap against mutual access

`__malloc_lock()` is the locking function which is called by the system library whenever the heap management has to be locked against mutual access. The implementation delivered with embOS claims a mutex or disables interrupts to achieve this.

4.2.2 `__malloc_unlock()`

`__malloc_unlock()` is the counterpart to `__malloc_lock()`. It is called by the system library whenever the heap management locking can be released. The implementation delivered with embOS releases the mutex or restores the interrupt state.

None of these functions has to be called directly by the application. They are called from the system library functions when required. The functions are delivered in source form to allow replacement of the dummy functions in the system library.

4.3 Thread-Local Storage TLS

Newlib supports usage of thread-local storage. Several library objects and functions need local variables which have to be unique to a thread. Thread-local storage will be required when these functions are called from multiple threads.

embOS for GCC is prepared to support thread-local storage, but does not use it per default. This has the advantage of no additional overhead as long as thread-local storage is not needed by the application. The embOS implementation of thread-local storage allows activation of TLS separately for each task.

Only tasks that are accessing TLS variables, for instance by calling functions from the system library, need to initialize their TLS by calling an initialization function when the task is started. For each task that uses TLS the memory for the thread-local storage is allocated on the heap. Therefore, thread-safe heap management should be used together with TLS. For information on thread safety, please refer to *Interrupt and thread safety* on page 21.

Library objects that need thread-local storage when used in multiple tasks are for example:

- error functions - `errno`, `strerror`.
- locale functions - `localeconv`, `setlocale`.
- time functions - `asctime`, `localtime`, `gmtime`, `mktime`.
- multibyte functions - `mbrlen`, `mbrtowc`, `mbsrtowc`, `mbtowc`, `wcrtomb`, `wcsrtomb`, `wctomb`.
- rand functions - `rand`, `srand`.
- etc functions - `atexit`, `strtok`.
- C++ exception engine.

4.3.1 API functions

Routine	Description	main	Priv Task	Unpriv Task	ISR	SW Timer
<code>OS_TLS_Set()</code>	Initializes the thread-local storage for the current task.		•			
<code>OS_TLS_SetTaskContextExtension()</code>	Initializes the thread-local storage and sets the TLS task context extension for the current task.		•			

4.3.1.1 OS_TLS_Set()

Description

Initializes the thread-local storage for the current task.

Prototype

```
void OS_TLS_Set(struct _reent* pReentStruct);
```

Parameters

Parameter	Description
<code>pReentStruct</code>	Pointer to the thread local storage. It is the address of the variable of type <code>struct _reent</code> which holds the thread-local data.

Additional information

`OS_TLS_Set()` shall be the first function called from a task when TLS should be used in this task. This function has to be used only in combination with `OS_TASK_AddContextExtension()` or `OS_TASK_SetContextExtension()` and `OS_TLS_ContextExtension` as argument to these functions. When `OS_TLS_SetTaskContextExtension()` is used, `OS_TLS_Set()` will be called automatically.

Please ensure sufficient task stack if the `_reent` structure variable is placed on the task stack. For details on the `_reent` structure, `_impure_ptr`, and library functions which require precautions on reentrancy, refer to the GNU documentation.

Example

```
static void Task(void) {
    struct _reent TaskReentStruct;

    OS_TLS_Set(&TaskReentStruct);
    OS_TASK_SetContextExtension(&OS_TLS_ContextExtension);
    while (1) {
    }
}
```


4.3.1.2 OS_TLS_SetTaskContextExtension()

Description

Initializes the thread-local storage and sets the TLS task context extension for the current task.

Prototype

```
void OS_TLS_SetTaskContextExtension(struct _reent* pReentStruct);
```

Parameters

Parameter	Description
<code>pReentStruct</code>	Pointer to the thread local storage. It is the address of the variable of type <code>struct _reent</code> which holds the thread-local data.

Additional information

`OS_TLS_SetTaskContextExtension()` shall be the first function called from a task when TLS should be used in this task.

If the task already contains a task context extension, `OS_TLS_SetTaskContextExtension()` cannot be used. Instead, `OS_TASK_AddContextExtension()` needs to be called with `OS_TLS_ContextExtension` as argument. Furthermore, `OS_TLS_Set()` needs to be called to initialize TLS for this task.

Please ensure sufficient task stack if the `_reent` structure variable is placed on the task stack. For details on the `_reent` structure, `_impure_ptr`, and library functions which require precautions on reentrancy, refer to the GNU documentation.

Example

The following printout demonstrates the usage of task specific TLS in an application.

```
#include "RTOS.h"

static OS_STACKPTR int StackHP[512], StackLP[512]; // Task stacks
static OS_TASK      TCBHP, TCBLP;                 // Task control blocks

static void HPTask(void) {
    struct _reent TaskReentStruct;

    OS_TLS_SetTaskContextExtension(&TaskReentStruct);
    while (1) {
        errno = 42; // errno specific to HPTask
        OS_TASK_Delay(50);
    }
}

static void LPTask(void) {
    struct _reent TaskReentStruct;

    OS_TLS_SetTaskContextExtension(&TaskReentStruct);
    while (1) {
        errno = 1; // errno specific to LPTask
        OS_TASK_Delay(200);
    }
}

int main(void) {
    errno = 0; // errno not specific to any task
    OS_Init(); // Initialize embOS
    OS_Inithw(); // Initialize required hardware
    OS_TASK_CREATE(&TCBHP, "HP Task", 100, HPTask, StackHP);
    OS_TASK_CREATE(&TCBLP, "LP Task", 50, LPTask, StackLP);
}
```

```
OS_Start();    // Start embOS  
return 0;  
}
```

Chapter 5

Stacks

5.1 Task stack

Each task uses its individual stack. The stack pointer is initialized and set every time a task is activated by the scheduler. The task stack needs to be able to accommodate the stack content of any (sub-)function plus the basic stack size.

The basic stack size is the size of memory required to store the context of the task on the stack. The minimum basic task stack size is 72 bytes for CPUs without a VFP/NEON unit and up to 332 bytes for CPUs with a VFP/NEON unit. We recommend at least 256 bytes stack as a start for CPUs without a VFP/NEON unit and 512 bytes for CPUs with a VFP/NEON unit.

Note

Stacks for ARM devices need to be 8-byte aligned. embOS ensures that task stacks are properly aligned. If an unaligned stack was aligned, the first few bytes up to the aligned address will not be used. Thus, the application should ensure that task stacks are properly aligned. This can be achieved by defining an array using a 64-bit data type like `OS_U64`.

5.2 System stack

The embOS scheduler executes in supervisor (SVC) mode. However, embOS doesn't use the dedicated SVC stack symbol in order to initialize the SVC stack pointer. After `OS_Start()` was called embOS uses the stack symbol of the system stack which was used in the `main()` routine. This avoids the need to allocate dedicated SVC stack space.

The minimum system stack size required by embOS is about 160 bytes (stack check & profiling build, no VFP/NEON unit). Since the system stack is also used by the application before the start of multitasking (the call to `OS_Start()`), and because software timers and C-level interrupt handlers also use the system stack, the actual stack requirements depend on the application.

The size of the system stack can be changed by modifying the stack size definition in your linker file or within the project settings. We recommend a minimum stack size of 512 bytes for the system stack for CPUs without a VFP/NEON unit and 1024 bytes for CPUs with a VFP/NEON unit.

The symbol `__stack_end__` is required by embOS and thus needs to be defined in the linker file.

In order to perform overflow checks on the system stack and to provide stack usage information, embOS requires in addition the symbol `__stack_start__` to be defined in the linker file.

Alternatively, if modification of the linker file is not feasible (e.g. when it is auto-generated by a project generator), the stack symbols may also be defined as additional linker flags. For example, if the linker file defines the symbols `LOWEST_STACK_ADDRESS` and `HIGHEST_STACK_ADDRESS`, the linker flags would need to be set as shown below:

```
--defsym=__stack_start__=LOWEST_STACK_ADDRESS
--defsym=__stack_end__=HIGHEST_STACK_ADDRESS
```

5.3 Interrupt stack

If a normal hardware exception occurs, the ARM core switches to IRQ mode, which has a separate stack pointer. After saving the scratch registers as well as `LR_irq` and `SPSR_irq` (and `FPSCR`, if VFP/NEON unit is present) onto the IRQ stack embOS switches to SVC mode. Only the previously mentioned registers are saved onto the IRQ stack. Thus, every interrupt requires 32 bytes on the IRQ stack. The maximum IRQ stack size required by the application

can be calculated as “Maximum interrupt nesting level * 32 bytes”. For the interrupt routine itself, the system stack is used, because they’re executed in SVC mode.

The size of the IRQ stack can be changed by modifying the stack size definition in your linker file or within the project settings.

The symbol `__stack_irq_end__` is required by embOS and thus needs to be defined in the linker file.

In order to perform overflow checks on the interrupt stack and to provide stack usage information, embOS requires in addition the symbol `__stack_irq_start__` to be defined in the linker file.

Alternatively, if modification of the linker file is not feasible (e.g. when it is auto-generated by a project generator), the stack symbols may also be defined as additional linker flags. For example, if the linker file defines the symbols `LOWEST_IRQ_STACK_ADDRESS` and `HIGHEST_IRQ_STACK_ADDRESS`, the linker flags would need to be set as shown below:

```
--defsym=__stack_irq_start__=LOWEST_IRQ_STACK_ADDRESS  
--defsym=__stack_irq_end__=HIGHEST_IRQ_STACK_ADDRESS
```

5.4 Stack specifics

There are two stacks which have to be declared in the linker script file or project settings:

- The system stack.
- The IRQ stack.

The system stack is used by the startup, the `main()` routine, embOS internal functions, and C-level interrupt handlers.

The IRQ stack is used when an interrupt exception is triggered. The exception handler saves some registers and then performs a mode switch which then uses the system stack for further execution.

When the CPU starts, it runs in supervisor mode. Then the startup code initializes the various stack pointer registers for each mode with their assigned stack and finally jumps into the `main()` routine. embOS expects the `main()` routine to use the system stack, no matter in which CPU mode.

When `OS_Init()` is called, embOS initializes the supervisor stack pointer to point to the system stack. After embOS is started with `OS_Start()`, the embOS scheduler runs in supervisor mode using the system stack while each task is running in system mode using its own dedicated stack.

Note

Stacks for ARM devices need to be 8-byte aligned.

Chapter 6

Interrupts

6.1 What happens when an interrupt occurs?

- The CPU-core receives an interrupt request.
- As soon as the interrupts are enabled, the interrupt is executed.
- The CPU switches to the IRQ mode which uses the IRQ stack.
- The CPU saves PC and flags in registers `LR_irq` and `SPSR_irq`.
- The CPU jumps to offset `0x18` in the exception vector table which contains an instruction to branch to the embOS low-level `IRQ_Handler()`.
- `embOS IRQ_Handler()`: Saves scratch registers as well as `LR_irq`, `SPSR_irq` and `FPSCR` on the IRQ stack.
- `embOS IRQ_Handler()`: Switches to supervisor mode and system stack.
- `embOS IRQ_Handler()`: Saves scratch VFP registers on the system stack.
- `embOS IRQ_Handler()`: Executes `OS_irq_handler()` (defined in `RTOSInit*.c`).
- `embOS OS_irq_handler()`: Checks for interrupt source and executes the according ISR handler. The implementation of this functions depends on the implemented interrupt controller.
- `embOS IRQ_Handler()`: Restores scratch VFP registers from the system stack.
- `embOS IRQ_Handler()`: Switches to IRQ mode and IRQ stack.
- `embOS IRQ_Handler()`: Restores scratch registers as well as `LR_irq`, `SPSR_irq` and `FPSCR` from the IRQ stack.
- `embOS IRQ_Handler()`: Returns from interrupt.

Note

FPSCR and VFP registers are only preserved by embOS libraries with VFP support.

6.2 Defining interrupt handlers in C

Interrupt handlers called from the default C interrupt handler `OS_irq_handler()` located in `RTOSInit*.c` are just normal functions which do not take parameters and do not return any value. `OS_irq_handler()` first calls `OS_INT_Enter()` or `OS_INT_EnterNestable()` to inform embOS that interrupt code is running. Then this handler examines the source of interrupt and calls the related interrupt handler function. Finally, `OS_irq_handler()` calls `OS_INT_Leave()` or `OS_INT_LeaveNestable()` and returns to the primary low level interrupt handler `IRQ_Handler()`.

Depending on the interrupting source, it may be required to reset the interrupt pending condition of the related peripherals.

Example

Simple interrupt routine:

```
void Timer_IRQHandler(void) {
    static unsigned long Time = 0;

    //
    // Handle timer IRQ
    //
    Time++;
}
```

6.3 Interrupt handling without vectored interrupt controller

When using an ARM CPU without implementation of a vectored interrupt controller, the application is responsible to examine which interrupting source triggered the IRQ.

The reaction to an interrupt is as follows:

- `IRQ_Handler()` calls `OS_irq_handler()`.
- `OS_irq_handler()` informs embOS that interrupt code is running by calling `OS_INT_Enter()`.
- `OS_irq_handler()` determines the interrupt sources and handles all pending IRQs.
- `OS_irq_handler()` informs embOS that interrupt handling ended by calling `OS_INT_Leave()`.
- `IRQ_Handler()` returns to `IRQ_Handler()`.

Example

Simple interrupt routine:

```
void OS_irq_handler(void) {
    OS_INT_Enter();
    if (Timer_IsPending()) { // Interrupt pending?
        Timer_IRQHandler(); // Handle interrupt
    }
    if (UART_IsPending()) { // Interrupt pending ?
        UART_IRQHandler(); // Handle interrupt
    }
    OS_INT_Leave();
}
```

During interrupt processing, you should not re-enable interrupts, as this would lead in recursion.

6.4 Interrupt handling with vectored interrupt controller

For ARM derivatives with built in vectored interrupt controller, embOS uses a different interrupt handling procedure and delivers additional functions to install and setup interrupt handler functions. You should not program the interrupt controller for IRQ handling directly. You should use the functions delivered with embOS.

The reaction to an interrupt with vectored interrupt controller is as follows:

- `IRQ_Handler()` calls `OS_irq_handler()`.
- `OS_irq_handler()` examines the interrupting source by reading the interrupt vector from the interrupt controller.
- `OS_irq_handler()` informs embOS that interrupt code is running by calling `OS_INT_Enter()`.
- `OS_irq_handler()` calls the interrupt handler function which is addressed by the interrupt vector.
- `OS_irq_handler()` resets the interrupt controller to re-enable acceptance of new interrupts.
- `OS_irq_handler()` informs embOS that interrupt handling ended by calling `OS_INT_Leave()`.
- `OS_irq_handler()` returns to `IRQ_Handler()`.

Note

Different ARM CPUs may have different versions of vectored interrupt controller hardware, and usage of embOS supplied functions varies depending on the type of interrupt controller. Refer to the samples delivered with embOS which are used in the CPU specific RTOSInit module.

6.4.1 API functions

To handle interrupts with vectored interrupt controller, embOS offers the following functions:

Function	Description
<code>OS_ARM_InstallISRHandler()</code>	Installs an interrupt handler
<code>OS_ARM_EnableISR()</code>	Enables an interrupt
<code>OS_ARM_DisableISR()</code>	Disables an interrupt
<code>OS_ARM_ISRSetPrio()</code>	Sets the priority of an interrupt
<code>OS_ARM_ClearPendingFlag()</code>	Clears an interrupt pending flag
<code>OS_ARM_IsPending()</code>	Checks if an interrupt is pending
<code>OS_ARM_AssignISRSource()</code>	Assigns a hardware interrupt channel to an interrupt vector
<code>OS_ARM_EnableISRSource()</code>	Enables an interrupt channel of a VIC type interrupt controller
<code>OS_ARM_DisableISRSource()</code>	Disables an interrupt channel of a VIC type interrupt controller
<code>OS_ARM_SetISRCfg()</code>	Sets the interrupt configuration.
<code>OS_ARM_SetVBAR()</code>	Writes the vector table address register.

6.4.1.1 OS_ARM_InstallISRHandler()

Description

`OS_ARM_InstallISRHandler()` is used to install a specific interrupt vector when ARM CPUs with vectored interrupt controller are used.

Prototype

```
OS_ISR_HANDLER* OS_ARM_InstallISRHandler(int                ISRIndex,  
                                           OS_ISR_HANDLER* pISRHandler);
```

Parameters

Parameter	Description
<i>ISRIndex</i>	Index of the interrupt source, usually the interrupt vector number.
<i>pISRHandler</i>	Address of the interrupt handler function.

Return Value

`OS_ISR_HANDLER*`: The address of the interrupt handler that was previously installed with the addressed interrupt source.

Additional Information

This function just installs the interrupt vector but does not modify the priority and does not automatically enable the interrupt.

6.4.1.2 OS_ARM_EnableISR()

Description

`OS_ARM_EnableISR()` is used to enable interrupt acceptance of a specific interrupt source in a vectored interrupt controller.

Prototype

```
void OS_ARM_EnableISR(int ISRIndex);
```

Parameters

Parameter	Description
<code>ISRIndex</code>	Index of the interrupt source which should be enabled.

Additional Information

This function just enables the interrupt inside the interrupt controller. It does not enable the interrupt of any peripherals. This has to be done elsewhere.

Note

For ARM CPUs with VIC type interrupt controller, this function just enables the interrupt vector itself. To enable the hardware assigned to that vector, you have to call `OS_ARM_EnableISRSource()` also.

6.4.1.3 OS_ARM_DisableISR()

Description

`OS_ARM_DisableISR()` is used to disable interrupt acceptance of a specific interrupt source in a vectored interrupt controller which is not of the VIC type.

Prototype

```
void OS_ARM_DisableISR(int ISRIndex);
```

Parameters

Parameter	Description
<code>ISRIndex</code>	Index of the interrupt source which should be disabled.

Additional Information

This function just disables the interrupt controller. It does not disable the interrupt of any peripherals. This has to be done elsewhere.

Note

When using an ARM CPU with built in interrupt controller of VIC type, use `OS_ARM_DisableISRSource()` to disable a specific interrupt.

6.4.1.4 OS_ARM_ISRSetPrio()

Description

OS_ARM_ISRSetPrio() is used to set or modify the priority of a specific interrupt source by programming the interrupt controller.

Prototype

```
int OS_ARM_ISRSetPrio(int ISRIndex,  
                     int Prio);
```

Parameters

Parameter	Description
ISRIndex	Index of the interrupt source which should be modified.
Prio	The priority which should be set for the specific interrupt.

Return Value

Previous priority which was assigned before the call of OS_ARM_ISRSetPrio().

Additional Information

This function sets the priority of an interrupt channel by programming the interrupt controller. Refer to CPU-specific manuals about allowed priority levels.

6.4.1.5 OS_ARM_ClearPendingFlag()

Description

OS_ARM_ClearPendingFlag() is used to clear an interrupt pending flag

Prototype

```
void OS_ARM_ClearPendingFlag(int ISRIndex);
```

Parameters

Parameter	Description
ISRIndex	Index of the interrupt source which should be cleared

Additional Information

This function just clears the interrupt pending flag inside the interrupt controller. It does not clear the interrupt pending flag in any peripheral.

6.4.1.6 OS_ARM_IsPending()

Description

OS_ARM_IsPending() is used to check if an interrupt is pending

Prototype

```
unsigned int OS_ARM_IsPending(int ISRIndex);
```

Parameters

Parameter	Description
ISRIndex	Index of the interrupt source which should be checked

Return value

= 0 Interrupt is not pending.
= 1 Interrupt is pending.

Additional Information

This function just checks the interrupt pending flag inside the interrupt controller. It does not check the interrupt pending flag in any peripheral.

6.4.1.7 OS_ARM_AssignISRSource()

Description

`OS_ARM_AssignISRSource()` is used to assign a hardware interrupt channel to an interrupt vector in an interrupt controller of VIC type.

Prototype

```
void OS_ARM_AssignISRSource(int ISRIndex,  
                           int Source);
```

Parameters

Parameter	Description
<code>ISRIndex</code>	Index of the interrupt source which should be modified.
<code>Source</code>	The source channel number which should be assigned to the specified interrupt vector.

Additional Information

This function assigns a hardware interrupt line to an interrupt vector of VIC type only. It cannot be used for other types of vectored interrupt controllers. The hardware interrupt channel number of specific peripherals depends on specific CPU derivatives and has to be taken from the hardware manual of the CPU.

Example

```
/* Install UART interrupt handler */  
OS_ARM_InstallISRHandler(UART_ID, &COM_ISR);           // UART interrupt vector  
OS_ARM_ISRSetPrio(UART_ID, UART_PRIO);                // UART interrupt priority  
OS_ARM_EnableISR(UART_ID);                             // Enable UART interrupt  
/* Install UART interrupt handler with VIC type interrupt controller*/  
OS_ARM_InstallISRHandler(UART_INT_INDEX, &COM_ISR);    // UART interrupt vector  
OS_ARM_AssignISRSource(UART_INT_INDEX, UART_INT_SOURCE);  
OS_ARM_EnableISR(UART_INT_INDEX);                     // Enable UART interrupt vector  
OS_ARM_EnableISRSource(UART_INT_SOURCE);               // Enable UART interrupt source
```

6.4.1.8 OS_ARM_EnableISRSource()

Description

`OS_ARM_EnableISRSource()` is used to enable an interrupt input channel of an interrupt controller of VIC type.

Prototype

```
void OS_ARM_EnableISRSource(int SourceIndex);
```

Parameters

Parameter	Description
<code>SourceIndex</code>	Index of the interrupt channel which should be enabled.

Additional Information

This function enables a hardware interrupt input of a VIC-type interrupt controller. It cannot be used for other types of vectored interrupt controllers. The hardware interrupt channel number of specific peripherals depends on specific CPU derivatives and has to be taken from the hardware manual of the CPU.

Example

```
/* Install UART interrupt handler */
OS_ARM_InstallISRHandler(UART_ID, &COM_ISR);           // UART interrupt vector
OS_ARM_ISRSetPrio(UART_ID, UART_PRIO);               // UART interrupt priority
OS_ARM_EnableISR(UART_ID);                           // Enable UART interrupt
/* Install UART interrupt handler with VIC type interrupt controller*/
OS_ARM_InstallISRHandler(UART_INT_INDEX, &COM_ISR);   // UART interrupt vector
OS_ARM_AssignISRSource(UART_INT_INDEX, UART_INT_SOURCE);
OS_ARM_EnableISR(UART_INT_INDEX);                   // Enable UART interrupt vector
OS_ARM_EnableISRSource(UART_INT_SOURCE);             // Enable UART interrupt source
```

6.4.1.9 OS_ARM_DisableISRSource()

Description

OS_ARM_DisableISRSource() is used to disable an interrupt input channel of an interrupt controller of VIC type.

Prototype

```
void OS_ARM_DisableISRSource(int SourceIndex);
```

Parameters

Parameter	Description
SourceIndex	Index of the interrupt channel which should be disabled.

Additional Information

This function disables a hardware interrupt input of a VIC-type interrupt controller. It cannot be used for other types of vectored interrupt controllers. The hardware interrupt channel number of specific peripherals depends on specific CPU derivatives and has to be taken from the hardware manual of the CPU.

6.4.1.10 OS_ARM_SetISRCfg()

Description

OS_ARM_SetISRCfg() sets the interrupt configuration.

Prototype

```
void OS_ARM_SetISRCfg(int    ISRIndex,  
                     OS_U32 Cfg);
```

Parameters

Parameter	Description
ISRIndex	Index of the interrupt source.
Cfg	0: Corresponding interrupt is level-sensitive. 1: Corresponding interrupt is edge-triggered.

6.4.1.11 OS_ARM_SetVBAR()

Description

OS_ARM_SetVBAR() writes the vector table address register.

Prototype

```
void OS_ARM_SetVBAR(OS_U32 Addr);
```

Parameters

Parameter	Description
Addr	Address of the vector table.

6.5 Interrupt-stack switching

Because ARM core based controllers have a separate stack pointer for interrupts, there is no need for explicit stack-switching in an interrupt routine. The routines `OS_INT_EnterIntStack()` and `OS_INT_LeaveIntStack()` are supplied for source compatibility to other processors only and have no functionality.

The ARM interrupt stack is used for the low-level interrupt handler `IRQ_Handler()` in the embOS library only.

6.6 Fast Interrupt (FIQ)

The FIQ interrupt cannot be used with embOS functions, it is reserved for high speed user functions.

Note the following:

- FIQ is never disabled by embOS.
- Never call any embOS function from an FIQ handler.
- Do not assign any embOS interrupt handler to FIQ.

Note

When you decide to use FIQ, ensure the FIQ stack is initialized during startup and that an interrupt vector for FIQ handling is included in your application.

Chapter 7

MMU/MPU and cache support

7.1 Introduction

This chapter describes the MMU/MPU and cache support for ARM CPUs. With the ARM core the MMU is part of the Virtual Memory System Architecture (VMSA) and the MPU is part of the Protected Memory System Architecture (PMSA). embOS comes with functions to support the MMU/MPU and cache of ARMv4, ARMv5 and ARMv7-A/ARMv7-R CPUs.

7.2 MMU handling for ARMv5/ARMv7-A CPUs

The MMU allows virtual-to-physical address mapping with sections of one MByte and cache control. The MMU requires a translation table which can be located in any data area, RAM or ROM, but has to be aligned at a 16Kbyte boundary. A translation table in RAM has to be set up during run time. embOS delivers API functions to set up this table.

7.2.1 API functions

Function	Description
<code>OS_ARM_MMU_InitTT()</code>	Initialize the MMU translation table.
<code>OS_ARM_MMU_AddTTEntries()</code>	Add address entries to the table.
<code>OS_ARM_MMU_Enable()</code>	Enable the MMU.
<code>OS_ARM_MMU_GetVirtualAddr()</code>	Translates a physical address into a virtual address
<code>OS_ARM_MMU_v2p()</code>	Translates a virtual address into a physical address.

7.2.1.1 OS_ARM_MMU_InitTT()

Description

OS_ARM_MMU_InitTT() is used to initialize an MMU translation table which is located in RAM. The table is filled with zeroes, thus all entries are marked as OS_ARM_MMU_UNMAPPED initially.

Prototype

```
void OS_ARM_MMU_InitTT(unsigned int* pTranslationTable);
```

Parameters

Parameter	Description
pTranslationTable	Points to the base address of the translation table.

7.2.1.2 OS_ARM_MMU_AddTTEntries()

Description

OS_ARM_MMU_AddTTEntries() is used to add entries to the MMU address translation table. The start address of the virtual address, physical address, area size and cache modes are passed as parameter.

Prototype

```
void OS_ARM_MMU_AddTTEntries(unsigned int* pTranslationTable,
                             unsigned int  CacheMode,
                             unsigned int  VIndex,
                             unsigned int  PIndex,
                             unsigned int  NumEntries);
```

Parameters

Parameter	Description
pTranslationTable	Points to the base address of the translation table.
CacheMode	<p>Specifies the cache operating mode and memory access permissions which should be used for the selected area. May be one of the following modes:</p> <p>ARMv4/ARMv5:</p> <ul style="list-style-type: none"> OS_ARM_MMU_UNMAPPED: The associated MVA is unmapped, and attempting to access it generates a translation fault OS_ARM_CACHEMODE_NC_NB: non-cacheable, non-bufferable OS_ARM_CACHEMODE_C_NB: cacheable, non-bufferable OS_ARM_CACHEMODE_NC_B: non-cacheable, bufferable OS_ARM_CACHEMODE_C_B: cacheable, bufferable <p>ARMv7-A:</p> <ul style="list-style-type: none"> OS_ARM_MMU_UNMAPPED: The associated MVA is unmapped, and attempting to access it generates a translation fault OS_ARM_CACHEMODE_STRONGLY_ORDERED: Strongly ordered OS_ARM_CACHEMODE_SHAREABLE_DEVICE: Shareable Device OS_ARM_CACHEMODE_WRITE_THROUGH: Outer and Inner Write-Through, no Write-Allocate OS_ARM_CACHEMODE_WRITE_BACK_NO_ALLOC: Outer and Inner Write-Back, no Write-Allocate OS_ARM_CACHEMODE_NON_CACHEABLE: Outer and Inner Non-cacheable OS_ARM_CACHEMODE_WRITE_BACK_ALLOC: Outer and Inner Write-Back, Write-Allocate OS_ARM_CACHEMODE_NON_SHAREABLE_DEVICE: Non-shareable Device OS_ARM_MMU_NOACCESS: All accesses generate Permission faults OS_ARM_MMU_READWRITE: Full access OS_ARM_MMU_READONLY: Read-only access OS_ARM_MMU_EXECUTE_NEVER: Determines whether the memory region is executable

Parameter	Description
<code>VIndex</code>	Virtual address index, which is the start address of the virtual memory address range with MBytes resolution. <code>VIndex</code> = (virtual address >> 20)
<code>PIndex</code>	Physical address index, which is the start address of the physical memory area range with MBytes resolution. <code>PIndex</code> = (physical address >> 20)
<code>NumEntries</code>	Specifies the size of the memory area in MBytes.

Additional information

The function adds entries for every section of one MegaByte size into the translation table for the specified memory area.

The macros for normal memory, i.e. `OS_ARM_CACHEMODE_WRITE_THROUGH`, `OS_ARM_CACHEMODE_WRITE_BACK_NO_ALLOC`, `OS_ARM_CACHEMODE_NON_CACHEABLE` and `OS_ARM_CACHEMODE_WRITE_BACK_ALLOC`, can be OR-red with `OS_ARM_MMU_SHAREABLE` to mark normal memory as shareable.

`OS_ARM_MMU_NOACCESS`, `OS_ARM_MMU_READWRITE`, `OS_ARM_MMU_READONLY` and `OS_ARM_MMU_EXECUTE_NEVER` can be used in combination with the cache attribute defines. If no memory access permissions are set full memory access is allowed per default.

`OS_ARM_MMU_InitTT()` sets all entries to `OS_ARM_MMU_UNMAPPED`. The MMU table does not need to define such entries.

7.2.1.3 OS_ARM_MMU_Enable()

Description

OS_ARM_MMU_Enable() is used to enable the MMU which will then perform the address mapping.

Prototype

```
void OS_ARM_MMU_Enable(unsigned int* pTranslationTable);
```

Parameters

Parameter	Description
<code>pTranslationTable</code>	Points to the base address of the translation table.

Additional information

As soon as the function was called, the address translation is active. The MMU table has to be setup before calling OS_ARM_MMU_Enable().

OS_ARM_MMU_Enable() also enables the branch prediction unit of Cortex-A CPUs.

7.2.1.4 OS_ARM_MMU_GetVirtualAddr()

Description

OS_ARM_MMU_GetVirtualAddr() is used to translate a physical address into a virtual address with specified cache mode.

Prototype

```
void* OS_ARM_MMU_GetVirtualAddr(unsigned long PAddr,
                                unsigned int  NumEntries);
```

Parameters

Parameter	Description
PAddr	The physical address as unsigned long.
CacheMode	<p>The cache mode of the requested virtual address May be one of the defined cache modes:</p> <p>ARMv4/ARMv5:</p> <ul style="list-style-type: none"> OS_ARM_CACHEMODE_NC_NB OS_ARM_CACHEMODE_C_NB OS_ARM_CACHEMODE_NC_B OS_ARM_CACHEMODE_C_B OS_ARM_CACHEMODE_ANY <p>ARMv7-A:</p> <ul style="list-style-type: none"> OS_ARM_CACHEMODE_STRONGLY_ORDERED OS_ARM_CACHEMODE_SHAREABLE_DEVICE OS_ARM_CACHEMODE_WRITE_THROUGH OS_ARM_CACHEMODE_WRITE_BACK_NO_ALLOC OS_ARM_CACHEMODE_NON_CACHEABLE OS_ARM_CACHEMODE_WRITE_BACK_ALLOC OS_ARM_CACHEMODE_NON_SHAREABLE_DEVICE OS_ARM_CACHEMODE_ANY

Return value

void* to the first virtual address found. A value of 0xFFFFFFFF indicates that no entry was found.

Additional information

The function may be useful to examine an address of memory mapped to a virtual address with specific cache mode. For the CPU it may be necessary to write into a specific memory in uncached mode. This can be done by setting up the MMU table with different virtual address for the same physical memory with different cache modes. For efficiency reasons, the CPU should access the memory fully cached for normal operation. When a peripheral or DMA accesses the same memory for reading, for example an LCD controller accesses the display buffer, or an Ethernet MAC access a transferbuffer, the CPU has to write the data uncached into this memory, or has to clean the cache after writing. The function OS_ARM_MMU_GetVirtualAddress() can be used to find the address for uncached access. The MMU table has to be setup before the function is called.

7.2.1.5 OS_ARM_MMU_v2p()

Description

OS_ARM_MMU_v2p() is used to translate a virtual address into a physical address.

Prototype

```
unsigned long OS_ARM_MMU_v2p(void* pVAddr);
```

Parameters

Parameter	Description
<code>pVAddr</code>	Pointer which represents the virtual address.

Return value

The physical address which is mapped to the virtual address passed as parameter.

Additional information

The function can be used to examine the physical addresss of memory. The CPU normally operates with virtual addresses which may differ from the physical address of the memory. When a peripheral or DMA has to be programmed to access the same memory, the peripheral has to be programmed to access the physical memory. The function OS_ARM_MMU_v2p() can be used to find the physical address of a memory area. The MMU table has to be setup before the function is called.

7.3 MPU handling for ARMv7-R CPUs

The ARMv7-R MPU is used to set cache and access settings for memory regions.

7.3.1 API functions

Function	Description
<code>OS_ARM_MPU_AddEntry()</code>	Sets an ARMv7-R PMSA MPU memory region.
<code>OS_ARM_MPU_Enable()</code>	Enables the ARMv7-R PMSA MPU.
<code>OS_ARM_MPU_GetMinRegionSize()</code>	Returns the ARMv7-R PMSA minimum memory region size.
<code>OS_ARM_MPU_GetNumRegions()</code>	Returns the number of available memory regions.
<code>OS_ARM_MPU_Init()</code>	Initializes the ARMv7-R PMSA MPU.

7.3.1.1 OS_ARM_MPU_AddEntry()

Description

OS_ARM_MPU_AddEntry() sets an ARMv7-R PMSA MPU memory region.

Prototype

```
void OS_ARM_MPU_AddEntry(OS_U32 Region,
                        void* BaseAddr,
                        OS_U32 Size,
                        OS_U32 Permissions,
                        OS_U32 Attributes);
```

Parameters

Parameter	Description
Region	Region index
BaseAddr	Memory region address
Size	Memory region size in bytes
Permissions	OS_ARM_MPU_NOACCESS: No read or write access OS_ARM_MPU_READWRITE: Read and write access OS_ARM_MPU_READONLY: Read access only OS_ARM_MPU_EXECUTE_NEVER: No code execution allowed
Attributes	OS_ARM_CACHEMODE_STRONGLY_ORDERED: Strongly ordered OS_ARM_CACHEMODE_SHAREABLE_DEVICE: Shareable Device OS_ARM_CACHEMODE_WRITE_THROUGH: Outer and Inner Write-Through, no Write-Allocate OS_ARM_CACHEMODE_WRITE_BACK_NO_ALLOC: Outer and Inner Write-Back, no Write-Allocate OS_ARM_CACHEMODE_NON_CACHEABLE: Outer and Inner Non-cacheable OS_ARM_CACHEMODE_WRITE_BACK_ALLOC: Outer and Inner Write-Back, Write-Allocate OS_ARM_CACHEMODE_NON_SHAREABLE_DEVICE: Non-shareable Device

Additional information

The region index starts at zero for the first region. The number of available regions can be read with OS_ARM_MPU_GetNumRegions(). The regions size must be aligned to the PMSA regions size which can be read with OS_ARM_MPU_GetMinRegionSize().

The macros for normal memory, i.e. OS_ARM_CACHEMODE_WRITE_THROUGH, OS_ARM_CACHEMODE_WRITE_BACK_NO_ALLOC, OS_ARM_CACHEMODE_NON_CACHEABLE and OS_ARM_CACHEMODE_WRITE_BACK_ALLOC, can be OR-red with OS_ARM_MPU_SHAREABLE to mark normal memory as shareable.

Example

```
int __low_level_init(void) {
    //
    // Sets access and cache settings for 256 Bytes at address 0x00
    //
    OS_ARM_MPU_AddEntry(0u, (void*)0x00000000, 256,
```

```
OS_ARM_MPU_READWRITE,  
OS_ARM_CACHEMODE_WRITE_BACK_ALLOC);  
  
return 1;  
}
```

7.3.1.2 OS_ARM_MPU_Enable()

Description

OS_ARM_MPU_Enable() enables the ARMv7-R PMSA MPU.

Prototype

```
void OS_ARM_MPU_Enable(void);
```

Additional information

OS_ARM_MPU_Enable() has to be called after the the MPU was initialized and configured.

7.3.1.3 OS_ARM_MPU_GetMinRegionSize()

Description

OS_ARM_MPU_GetMinRegionSize() returns the ARMv7-R PMSA minimum memory region size.

Prototype

```
OS_U32 OS_ARM_MPU_GetMinRegionSize(void);
```

Return value

Minimum memory region size which can be used with this PMSA implementation.

7.3.1.4 OS_ARM_MPU_GetNumRegions()

Description

`OS_ARM_MPU_GetNumRegions()` returns the number of available memory regions.

Prototype

```
OS_U32 OS_ARM_MPU_GetNumRegions(void);
```

Return value

Number of available memory regions.

7.3.1.5 OS_ARM_MPU_Init()

Description

OS_ARM_MPU_Init() initializes the ARMv7-R PMSA MPU.

Prototype

```
void OS_ARM_MPU_Init(void);
```

7.4 Cache handling for ARMv5/ARMv7 CPUs

ARM CPUs with MMU/MPU and cache have separate data and instruction caches. embOS delivers the following functions to setup and handle the MMU and caches.

7.4.1 API functions

Function	Description
<code>OS_ARM_ICACHE_Enable()</code>	Enable the instruction cache.
<code>OS_ARM_ICACHE_Invalidate()</code>	Invalidates the complete instruction cache.
<code>OS_ARM_DCACHE_Enable()</code>	Enable the data cache.
<code>OS_ARM_DCACHE_Invalidate()</code>	Invalidates the complete data cache.
<code>OS_ARM_DCACHE_Clean()</code>	Clean data cache.
<code>OS_ARM_DCACHE_CleanRange()</code>	Clean data cache range.
<code>OS_ARM_DCACHE_InvalidateRange()</code>	Invalidate the data cache.
<code>OS_ARM_CACHE_Sync()</code>	Syncs data and instruction cache.
<code>OS_ARM_AddL2Cache()</code>	Sets 2nd level cache API table.
<code>OS_ARM_CACHE_GetLineSize()</code>	Returns cache line size of the specified cache level.

7.4.1.1 OS_ARM_ICACHE_Enable()

Description

OS_ARM_ICACHE_Enable() is used to enable the instruction cache of the CPU.

Prototype

```
void OS_ARM_ICACHE_Enable(void);
```

Additional information

As soon as the function was called, the instruction cache is active. It is CPU implementation defined whether the instruction cache works without MMU. Normally, the MMU should be setup before activating instruction cache.

7.4.1.2 OS_ARM_ICACHE_Invalidate()

Description

`OS_ARM_ICACHE_Invalidate()` invalidates the complete instruction cache. Invalidating means, mark all entries in the specified area as invalid. Invalidation forces re-reading the code from memory into the cache, when the specified area is accessed again.

Prototype

```
void OS_ARM_ICACHE_Invalidate(void);
```

7.4.1.3 OS_ARM_DCACHE_Enable()

Description

OS_ARM_DCACHE_Enable() is used to enable the data cache of the CPU.

Prototype

```
void OS_ARM_DCACHE_Enable(void);
```

Additional information

The function must not be called before the MMU translation table was set up correctly and the MMU was enabled. As soon as the function was called, the data cache is active, according to the cache mode settings which are defined in the MMU translation table. It is CPU implementation defined whether the data cache is a write through, a write back, or a write through/write back cache. Most modern CPUs will have implemented a write through/write back cache.

7.4.1.4 OS_ARM_DCACHE_Invalidate()

Description

`OS_ARM_DCACHE_Invalidate()` invalidates the complete data cache. Invalidating means, mark all entries in the specified area as invalid. Invalidation forces re-reading the data from memory into the cache, when the specified area is accessed again.

Prototype

```
void OS_ARM_DCACHE_Invalidate(void);
```

7.4.1.5 OS_ARM_DCACHE_Clean()

Description

`OS_ARM_DCACHE_Clean()` is used to clean the data cache memory without invalidating the instruction cache.

Prototype

```
void OS_ARM_DCACHE_Clean(void);
```

Additional information

Cleaning the data cache is needed, when data should be transferred by a DMA or other BUS master that does not use the data cache. When the CPU writes data into a cacheable area, the data might not be written into the memory immediately. When then a DMA cycle is started to transfer the data from memory to any other location or peripheral, the wrong data will be written.

The cache is cleaned line by line. Cleaning one cache line takes approximately 10 CPU cycles. The total time to invalidate a range may be calculated as:

$$t = (\text{NumBytes} / \text{Cache line size}) * (10 [\text{CPU clock cycles}] + \text{Memory write time}).$$

The real time depends on the content of the cache. If data in the cache is marked as dirty, the cache line has to be written to memory. The memory write time depends on the memory BUS clock and memory speed. If data has to be written to memory, the required cycles for this memory operation has to be added to the 10 CPU clock cycles for every cache line to be cleaned.

7.4.1.6 OS_ARM_DCACHE_CleanRange()

Description

`OS_ARM_DCACHE_CleanRange()` is used to clean a range in the data cache memory to ensure that the data is written from the data cache into the memory.

Prototype

```
void OS_ARM_DCACHE_CleanRange(void* p,  
                               unsigned int NumEntries);
```

Parameters

Parameter	Description
<code>p</code>	Points to the base address of the memory area that should be updated.
<code>NumBytes</code>	Number of bytes which have to be written from cache to memory.

Additional information

Cleaning the data cache is needed, when data should be transferred by a DMA or other BUS master that does not use the data cache. When the CPU writes data into a cacheable area, the data might not be written into the memory immediately. When then a DMA cycle is started to transfer the data from memory to any other location or peripheral, the wrong data will be written.

Before starting a DMA transfer, a call of `OS_ARM_DCACHE_CleanRange()` ensures, that the data is transferred from the data cache into the memory and the write buffers are drained.

The cache is cleaned line by line. Cleaning one cache line takes approximately 10 CPU cycles. The total time to invalidate a range may be calculated as:

$$t = (\text{NumBytes} / \text{Cache line size}) * (10 [\text{CPU clock cycles}] + \text{Memory write time}).$$

The real time depends on the content of the cache. If data in the cache is marked as dirty, the cache line has to be written to memory. The memory write time depends on the memory BUS clock and memory speed. If data has to be written to memory, the required cycles for this memory operation has to be added to the 10 CPU clock cycles for every cache line to be cleaned.

Note

Unfortunately, only complete cache lines can be cleaned. Therefore, it is required, that the base address of the memory area has to be located at a cache line size byte boundary and the number of bytes to be cleaned has to be a multiple of the cache line size. The debug version of embOS will call `OS_Error()` with error code `OS_ERR_NON_ALIGNED_INVALIDATE`, if one of these restrictions is violated.

7.4.1.7 OS_ARM_DCACHE_InvalidateRange()

Description

`OS_ARM_DCACHE_InvalidateRange()` is used to invalidate a memory area in the data cache. Invalidating means, mark all entries in the specified area as invalid. Invalidating forces re-reading the data from memory into the cache, when the specified area is accessed again.

Prototype

```
void OS_ARM_DCACHE_InvalidateRange(void* p,
                                   unsigned int NumBytes);
```

Parameters

Parameter	Description
<code>p</code>	Points to the base address of the memory area that should be updated.
<code>NumBytes</code>	Number of bytes which have to be written from cache to memory.

Additional information

This function is needed, when a DMA or other BUS master is used to transfer data into the main memory and the CPU has to process the data after the transfer.

To ensure, that the CPU processes the updated data from the memory, the cache has to be invalidated. Otherwise the CPU might read invalid data from the cache instead of the memory.

Special care has to be taken, before the data cache is invalidated. Invalidating a data area marks all entries in the data cache as invalid. If the cache contained data which was not written into the memory before, the data gets lost. The cache is invalidated line by line. Invalidating one cache line takes approximately 10 CPU cycles. The total time to invalidate a range may be calculated as: $t = (\text{NumBytes} / \text{Cache line size}) * 10$ [CPU clock cycles]. Notes

Note

Unfortunately, only complete cache lines can be invalidated. Therefore, it is required, that the base address of the memory area has to be located at a cache line size byte boundary and the number of bytes to be invalidated has to be a multiple of the cache line size. The debug version of embOS will call `OS_Error()` with error code `OS_ERR_NON_ALIGNED_INVALIDATE`, if one of these restrictions is violated.

7.4.1.8 OS_ARM_CACHE_Sync()

Description

`OS_ARM_CACHE_Sync()` cleans the data cache and invalidates the instruction cache to ensure cache coherency.

Prototype

```
void OS_ARM_CACHE_Sync(void);
```

Additional information

This function is for example needed, when code is copied into RAM and code is then executed from RAM.

7.4.1.9 OS_ARM_AddL2Cache()

Description

OS_ARM_AddL2Cache() is used to add.

Prototype

```
void OS_ARM_v7_AddL2Cache(const OS_ARM_L2CACHE_API* pCacheAPI,  
                          void* pParam);
```

Parameters

Parameter	Description
pCacheAPI	Pointer to 2nd level Cache API table.
pParam	Additional parameter (e.g. base address or cache registers).

Additional information

This function is needed to enable the L2 cache. Nothing else is necessary to do since the actual L2 cache routines are automatically called by the L1 cache routines. For example OS_ARM_DCACHE_InvalidateRange() calls also internally the according L2 cache routine.

Example

```
#define L2CACHE_BASE_ADDR 0x3FFF000u  
  
//  
// Set API functions and base address for L2 Cache  
//  
OS_ARM_AddL2Cache(&OS_L2CACHE_L2C310, (void*)L2CACHE_BASE_ADDR);
```

7.4.1.10 OS_ARM_CACHE_GetLineSize()

Description

`OS_ARM_CACHE_GetLineSize()` returns the cache line size of the specified cache level.

Prototype

```
OS_U32 OS_ARM_CACHE_GetLineSize(OS_U32 CIndex);
```

Parameters

Parameter	Description
<code>CIndex</code>	Index of the cache level of which the cache line size shall be returned.

Additional information

The returned cache line size can be used to calculate the alignment and number of bytes passed to the `OS_ARM_DCACHE_InvalidateRange()` and `OS_ARM_DCACHE_CleanRange()` functions.

7.5 MMU and cache handling program sample

The MMU and cache handling has to be set up before the data segments are initialized. Otherwise a virtual address mapping would not work. The startup code must call a `__low_level_init()` function before sections are initialized.

It is a good idea to initialize memory access, the MMU table and the cache control during `__low_level_init()`. The following sample is an excerpt from one `__low_level_init()` function which is part of an `RTOSInit.c` file:

```

/*****
 *
 *  MMU and cache configuration
 *
 *  The MMU translation table has to be aligned to 16KB boundary
 *  and has to be located in uninitialized data area
 */
#pragma data_alignment=16384
__no_init static unsigned int _TranslationTable [0x1000]; // OS_INTERWORK int
int __low_level_init(void) {
    //
    // Init MMU and caches
    //
    OS_ARM_MMU_InitTT (&_TranslationTable[0]);
    //
    // Internal SRAM, the first MB remapped to 0,
    // cacheable, bufferable, region not executable
    //
    OS_ARM_MMU_AddTTEntries ( &_TranslationTable[0],
                             OS_ARM_CACHEMODE_C_B | OS_ARM_MMU_EXECUTE_NEVER,
                             0x000, 0x200, 0x001);

    //
    // Internal SRAM, original address, NON cachable, NON bufferable
    //
    OS_ARM_MMU_AddTTEntries ( &_TranslationTable[0],
                             OS_ARM_CACHEMODE_NC_NB,
                             0x200, 0x200, 0x001);

    OS_ARM_MMU_Enable (&_TranslationTable[0]);
    OS_ARM_ICACHE_Enable();
    OS_ARM_DCACHE_Enable();
    return 1;
}

```

Other samples are included in the CPU specific `RTOSInit*.c` files delivered with embOS.

7.6 MPU and cache handling program sample

```
int __low_level_init(void) {  
    //  
    // Enable MPU, Caches and branch prediction unit  
    //  
    OS_ARM_DCACHE_Enable();  
    OS_ARM_ICACHE_Enable();  
    OS_ARM_MPU_Init();  
    //  
    // Add MPU regions to set cache settings for different memory sections  
    //  
    OS_ARM_MPU_AddEntry(0u, (void*)0x00000000, 0x00140000,  
        OS_ARM_MPU_READONLY, OS_ARM_CACHEMODE_WRITE_BACK_ALLOC); // FLASH  
    OS_ARM_MPU_AddEntry(1u, (void*)0x08000000, 0x00030000,  
        OS_ARM_MPU_READWRITE, OS_ARM_CACHEMODE_WRITE_BACK_ALLOC); // RAM  
    OS_ARM_MPU_Enable();  
  
    return 1;  
}
```

Chapter 8

VFP and NEON support

8.1 Introduction

Some ARM MCUs come with integrated Arm VFP and NEON units.

When activating the VFP or NEON support in the project options, the compiler and linker will add efficient code which uses the VFP/NEON register bank and VFP/NEON instructions where possible in the application.

With embOS, the VFP/NEON registers are automatically saved and restored when preemptive or cooperative task switches are performed. embOS also automatically saves and restores VFP/NEON registers for all embOS interrupt routines.

The VFP register bank consists of either 16 or 32 double-precision registers. The VFP register bank is also shared between the VFP and NEON units. If a NEON unit is implemented the VFP register bank consists of 32 64-bit double-precision registers. For a VFP unit with 32 double-precision registers or NEON unit all 32 double-precision registers are preserved, while for a VFP unit with 16 double-precision registers only the 16 double-precision registers need to be preserved.

embOS comes with libraries which preserve 16 double-precision registers D0-D15, 32 double-precision registers D0-D31 or none. Please have a look in the chapter *Libraries* on page 18 for more details.

Note

embOS ARM until V5.16.1.0 used task context extensions and ISR macros to preserve VFP/NEON registers. These API functions and macros are kept for compatibility but have no functionality anymore.

8.2 Using embOS libraries with VFP/NEON support

When VFP/NEON support is selected as project option, one of the embOS libraries with VFP/NEON support have to be used in the project. The embOS libraries with VFP/NEON support require that the VFP/NEON unit is switched on during startup and remains switched on during program execution. When the VFP/NEON unit is not switched on, the embOS scheduler will fail. Using your own startup code, ensure that the VFP/NEON unit is switched on during startup.

The debug version of embOS checks in `OS_Init()` whether the VFP/NEON unit is switched on. If not, embOS calls `OS_Error()` with the error code `OS_ERR_HW_NOT_AVAILABLE`.

8.3 Using the VFP/NEON unit in interrupt service routines

Using the VFP/NEON unit in embOS interrupt service routines does not require any additional functions to save and restore the VFP/NEON registers. embOS automatically saves and restores these registers.

VFP/NEON registers are not automatically saved and restored in zero latency interrupts. If the VFP/NEON unit is used in zero latency interrupts, it is the user's responsibility to preserve these registers.

Chapter 9

Technical data

9.1 Resource Usage

The memory requirements of embOS (RAM and ROM) differs depending on the used features, CPU, compiler, and library model. The following values are measured using embOS library mode `OS_LIBMODE_XR`.

Module	Memory type	Memory requirements
embOS kernel	ROM	~1700 bytes
embOS kernel	RAM	~136 bytes
Task control block	RAM	36 bytes
Software timer	RAM	20 bytes
Task event	RAM	0 bytes
Event object	RAM	12 bytes
Mutex	RAM	16 bytes
Semaphore	RAM	8 bytes
RWLock	RAM	28 bytes
Mailbox	RAM	24 bytes
Queue	RAM	32 bytes
Watchdog	RAM	12 bytes
Fixed Block Size Memory Pool	RAM	32 bytes