

RTEMS CPU Architecture Supplement

Edition 4.10.99.0, for RTEMS 4.10.99.0

23 November 2011

On-Line Applications Research Corporation

COPYRIGHT © 1988 - 2012.
On-Line Applications Research Corporation (OAR).

The authors have used their best efforts in preparing this material. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. No warranty of any kind, expressed or implied, with regard to the software or the material contained in this document is provided. No liability arising out of the application or use of any product described in this document is assumed. The authors reserve the right to revise this material and to make changes from time to time in the content hereof without obligation to notify anyone of such revision or changes.

The RTEMS Project is hosted at <http://www.rtems.com>. Any inquiries concerning RTEMS, its related support components, its documentation, or any custom services for RTEMS should be directed to the contacts listed on that site. A current list of RTEMS Support Providers is at <http://www.rtems.com/oarsupport>.

Table of Contents

Preface	1
1 Port Specific Information	3
1.1 CPU Model Dependent Features	3
1.1.1 CPU Model Name	4
1.1.2 Floating Point Unit	4
1.2 Calling Conventions	4
1.2.1 Calling Mechanism	4
1.2.2 Register Usage	4
1.2.3 Parameter Passing	5
1.2.4 User-Provided Routines	5
1.3 Memory Model	5
1.3.1 Flat Memory Model	5
1.4 Interrupt Processing	5
1.4.1 Vectoring of an Interrupt Handler	6
1.4.2 Interrupt Levels	6
1.4.3 Disabling of Interrupts by RTEMS	6
1.5 Default Fatal Error Processing	7
1.6 Board Support Packages	7
1.6.1 System Reset	7
2 ARM Specific Information	9
2.1 CPU Model Dependent Features	9
2.1.1 CPU Model Name	9
2.1.2 Count Leading Zeroes Instruction	9
2.1.3 Floating Point Unit	9
2.2 Calling Conventions	9
2.3 Memory Model	9
2.4 Interrupt Processing	9
2.4.1 Interrupt Levels	10
2.4.2 Interrupt Stack	10
2.5 Default Fatal Error Processing	10
3 Atmel AVR Specific Information	11
3.1 CPU Model Dependent Features	11
3.1.1 Count Leading Zeroes Instruction	11
3.2 Calling Conventions	11
3.2.1 Processor Background	11
3.2.2 Register Usage	11
3.2.3 Parameter Passing	11
3.3 Memory Model	12
3.4 Interrupt Processing	12

3.4.1	Vectoring of an Interrupt Handler	12
3.4.2	Disabling of Interrupts by RTEMS	12
3.4.3	Interrupt Stack	12
3.5	Default Fatal Error Processing	12
3.6	Board Support Packages	12
3.6.1	System Reset	12
4	Blackfin Specific Information	13
4.1	CPU Model Dependent Features	13
4.1.1	Count Leading Zeroes Instruction	13
4.2	Calling Conventions	13
4.2.1	Processor Background	13
4.2.2	Register Usage	13
4.2.3	Parameter Passing	14
4.3	Memory Model	14
4.4	Interrupt Processing	14
4.4.1	Vectoring of an Interrupt Handler	14
4.4.2	Disabling of Interrupts by RTEMS	14
4.4.3	Interrupt Stack	14
4.5	Default Fatal Error Processing	14
4.6	Board Support Packages	15
4.6.1	System Reset	15
5	Intel/AMD x86 Specific Information	17
5.1	CPU Model Dependent Features	17
5.1.1	bswap Instruction	17
5.2	Calling Conventions	17
5.2.1	Processor Background	17
5.2.2	Calling Mechanism	17
5.2.3	Register Usage	17
5.2.4	Parameter Passing	18
5.3	Memory Model	18
5.3.1	Flat Memory Model	18
5.4	Interrupt Processing	18
5.4.1	Vectoring of Interrupt Handler	19
5.4.2	Interrupt Stack Frame	19
5.4.3	Interrupt Levels	19
5.4.4	Interrupt Stack	19
5.5	Default Fatal Error Processing	19
5.6	Board Support Packages	20
5.6.1	System Reset	20
5.6.2	Processor Initialization	20

6	Lattice Mico32 Specific Information	23
6.1	CPU Model Dependent Features	23
6.2	Register Architecture	23
6.3	Calling Conventions	24
6.3.1	Calling Mechanism	24
6.3.2	Register Usage	24
6.3.3	Parameter Passing	24
6.4	Memory Model.....	24
6.5	Interrupt Processing	24
6.6	Default Fatal Error Processing	25
6.7	Board Support Packages	25
6.7.1	System Reset	25
7	M68xxx and Coldfire Specific Information ..	27
7.1	CPU Model Dependent Features	27
7.1.1	BFFFO Instruction	27
7.1.2	Vector Base Register	27
7.1.3	Separate Stacks	27
7.1.4	Pre-Indexing Address Mode	27
7.1.5	Extend Byte to Long Instruction	27
7.2	Calling Conventions	28
7.2.1	Calling Mechanism.....	28
7.2.2	Register Usage	28
7.2.3	Parameter Passing	28
7.3	Memory Model.....	28
7.4	Interrupt Processing	29
7.4.1	Vectoring of an Interrupt Handler	29
7.4.1.1	Models Without Separate Interrupt Stacks.....	29
7.4.1.2	Models With Separate Interrupt Stacks.....	29
7.4.2	CPU Models Without VBR and RAM at 0	30
7.4.3	Interrupt Levels.....	31
7.5	Default Fatal Error Processing	31
7.6	Board Support Packages	31
7.6.1	System Reset	31
7.6.2	Processor Initialization.....	32
8	MIPS Specific Information	33
8.1	CPU Model Dependent Features	33
8.1.1	Another Optional Feature	33
8.2	Calling Conventions	33
8.2.1	Processor Background.....	33
8.2.2	Calling Mechanism.....	33
8.2.3	Register Usage	33
8.2.4	Parameter Passing	33
8.3	Memory Model.....	33
8.3.1	Flat Memory Model.....	33
8.4	Interrupt Processing	34

8.4.1	Vectoring of an Interrupt Handler	34
8.4.2	Interrupt Levels	34
8.5	Default Fatal Error Processing	34
8.6	Board Support Packages	34
8.6.1	System Reset	34
8.6.2	Processor Initialization	34
9	PowerPC Specific Information	35
9.1	CPU Model Dependent Features	36
9.1.1	Alignment	36
9.1.2	Cache Alignment	36
9.1.3	Maximum Interrupts	36
9.1.4	Has Double Precision Floating Point	36
9.1.5	Critical Interrupts	36
9.1.6	Use Multiword Load/Store Instructions	36
9.1.7	Instruction Cache Size	36
9.1.8	Data Cache Size	36
9.1.9	Debug Model	37
9.1.9.1	Low Power Model	37
9.2	Calling Conventions	37
9.2.1	Programming Model	37
9.2.1.1	Non-Floating Point Registers	37
9.2.1.2	Floating Point Registers	38
9.2.1.3	Special Registers	38
9.2.2	Call and Return Mechanism	38
9.2.3	Calling Mechanism	39
9.2.4	Register Usage	39
9.2.5	Parameter Passing	39
9.3	Memory Model	39
9.3.1	Flat Memory Model	39
9.4	Interrupt Processing	40
9.4.1	Synchronous Versus Asynchronous Exceptions	40
9.4.2	Vectoring of Interrupt Handler	41
9.4.3	Interrupt Levels	41
9.5	Default Fatal Error Processing	42
9.6	Board Support Packages	42
9.6.1	System Reset	42
9.6.2	Processor Initialization	42

10	SuperH Specific Information	43
10.1	CPU Model Dependent Features	43
10.1.1	Another Optional Feature	43
10.2	Calling Conventions	43
10.2.1	Calling Mechanism	43
10.2.2	Register Usage	43
10.2.3	Parameter Passing	43
10.3	Memory Model	43
10.3.1	Flat Memory Model	44
10.4	Interrupt Processing	44
10.4.1	Vectoring of an Interrupt Handler	44
10.4.2	Interrupt Levels	44
10.5	Default Fatal Error Processing	44
10.6	Board Support Packages	44
10.6.1	System Reset	44
10.6.2	Processor Initialization	44
11	SPARC Specific Information	45
11.1	CPU Model Dependent Features	46
11.1.1	CPU Model Feature Flags	46
11.1.1.1	CPU Model Name	46
11.1.1.2	Floating Point Unit	46
11.1.1.3	Bitscan Instruction	47
11.1.1.4	Number of Register Windows	47
11.1.1.5	Low Power Mode	47
11.1.2	CPU Model Implementation Notes	47
11.2	Calling Conventions	48
11.2.1	Programming Model	48
11.2.1.1	Non-Floating Point Registers	48
11.2.1.2	Floating Point Registers	49
11.2.1.3	Special Registers	49
11.2.2	Register Windows	49
11.2.3	Call and Return Mechanism	51
11.2.4	Calling Mechanism	51
11.2.5	Register Usage	51
11.2.6	Parameter Passing	51
11.2.7	User-Provided Routines	51
11.3	Memory Model	51
11.3.1	Flat Memory Model	52
11.4	Interrupt Processing	52
11.4.1	Synchronous Versus Asynchronous Traps	53
11.4.2	Vectoring of Interrupt Handler	53
11.4.3	Traps and Register Windows	54
11.4.4	Interrupt Levels	54
11.4.5	Disabling of Interrupts by RTEMS	54
11.4.6	Interrupt Stack	55
11.5	Default Fatal Error Processing	55
11.5.1	Default Fatal Error Handler Operations	55

11.6	Board Support Packages	55
11.6.1	System Reset	55
11.6.2	Processor Initialization	56
12	SPARC-64 Specific Information	57
12.1	CPU Model Dependent Features	57
12.1.1	CPU Model Feature Flags	57
12.1.1.1	CPU Model Name	57
12.1.1.2	Floating Point Unit	58
12.1.1.3	Number of Register Windows	58
12.1.2	CPU Model Implementation Notes	58
12.1.2.1	sun4u Notes	58
12.1.3	sun4v Notes	58
12.2	Calling Conventions	58
12.2.1	Programming Model	58
12.2.1.1	Non-Floating Point Registers	58
12.2.1.2	Floating Point Registers	59
12.2.1.3	Special Registers	59
12.2.2	Register Windows	60
12.2.3	Call and Return Mechanism	61
12.2.4	Calling Mechanism	62
12.2.5	Register Usage	62
12.2.6	Parameter Passing	62
12.2.7	User-Provided Routines	62
12.3	Memory Model	62
12.3.1	Flat Memory Model	62
12.4	Interrupt Processing	63
12.4.1	Synchronous Versus Asynchronous Traps	63
12.4.2	Vectoring of Interrupt Handler	63
12.4.3	Traps and Register Windows	65
12.4.4	Interrupt Levels	65
12.4.5	Disabling of Interrupts by RTEMS	65
12.4.6	Interrupt Stack	65
12.5	Default Fatal Error Processing	65
12.5.1	Default Fatal Error Handler Operations	65
12.6	Board Support Packages	65
12.6.1	HelenOS and Open Firmware	66
	RTEMS CPU Architecture Supplement	67
	Command and Variable Index	69
	Concept Index	71

Preface

The Real Time Executive for Multiprocessor Systems (RTEMS) is designed to be portable across multiple processor architectures. However, the nature of real-time systems makes it essential that the application designer understand certain processor dependent implementation details. These processor dependencies include calling convention, board support package issues, interrupt processing, exact RTEMS memory requirements, performance data, header files, and the assembly language interface to the executive.

Each architecture represents a CPU family and usually there are a wide variety of CPU models within it. These models share a common Instruction Set Architecture (ISA) which often varies based upon some well-defined rules. There are often multiple implementations of the ISA and these may be from one or multiple vendors.

On top of variations in the ISA, there may also be variations which occur when a CPU core implementation is combined with a set of peripherals to form a system on chip. For example, there are many ARM CPU models from numerous semiconductor vendors and a wide variety of peripherals. But at the ISA level, they share a common compatibility.

RTEMS depends upon this core similarity across the CPU models and leverages that to minimize the source code that is specific to any particular CPU core implementation or CPU model.

This manual is separate and distinct from the RTEMS Porting Guide. That manual is a guide on porting RTEMS to a new architecture. This manual is focused on the more mundane CPU architecture specific issues that may impact application development. For example, if you need to write a subroutine in assembly language, it is critical to understand the calling conventions for the target architecture.

The first chapter in this manual describes these issues in general terms. In a sense, it is posing the questions one should be aware may need to be answered and understood when porting an RTEMS application to a new architecture. Each subsequent chapter gives the answers to those questions for a particular CPU architecture.

1 Port Specific Information

This chapter provides a general description of the type of architecture specific information which is in each of the architecture specific chapters that follow. The outline of this chapter is identical to that of the architecture specific chapters.

In each of the architecture specific chapters, this introductory section will provide an overview of the architecture

Architecture Documents

In each of the architecture specific chapters, this section will provide pointers on where to obtain documentation.

1.1 CPU Model Dependent Features

Microprocessors are generally classified into families with a variety of CPU models or implementations within that family. Within a processor family, there is a high level of binary compatibility. This family may be based on either an architectural specification or on maintaining compatibility with a popular processor. Recent microprocessor families such as the SPARC or PowerPC are based on an architectural specification which is independent of any particular CPU model or implementation. Older families such as the Motorola 68000 and the Intel x86 evolved as the manufacturer strived to produce higher performance processor models which maintained binary compatibility with older models.

RTEMS takes advantage of the similarity of the various models within a CPU family. Although the models do vary in significant ways, the high level of compatibility makes it possible to share the bulk of the CPU dependent executive code across the entire family. Each processor family supported by RTEMS has a list of features which vary between CPU models within a family. For example, the most common model dependent feature regardless of CPU family is the presence or absence of a floating point unit or coprocessor. When defining the list of features present on a particular CPU model, one simply notes that floating point hardware is or is not present and defines a single constant appropriately. Conditional compilation is utilized to include the appropriate source code for this CPU model's feature set. It is important to note that this means that RTEMS is thus compiled using the appropriate feature set and compilation flags optimal for this CPU model used. The alternative would be to generate a binary which would execute on all family members using only the features which were always present.

The set of CPU model feature macros are defined in the file `cpukit/score/cpu/CPU/rtems/score/cpu.h` based upon the GNU tools multilib variant that is appropriate for the particular CPU model defined on the compilation command line.

In each of the architecture specific chapters, this section presents the set of features which vary across various implementations of the architecture that may be of importance to RTEMS application developers.

The subsections will vary amongst the target architecture chapters as the specific features may vary. However, each port will include a few common features such as the CPU Model

Name and presence of a hardware Floating Point Unit. The common features are described here.

1.1.1 CPU Model Name

The macro `CPU_MODEL_NAME` is a string which designates the name of this CPU model. For example, for the MC68020 processor model from the m68k architecture, this macro is set to the string "mc68020".

1.1.2 Floating Point Unit

In most architectures, the presence of a floating point unit is an option. It does not matter whether the hardware floating point support is incorporated on-chip or is an external coprocessor as long as it appears an FPU per the ISA. However, if a hardware FPU is not present, it is possible that the floating point emulation library for this CPU is not reentrant and thus context switched by RTEMS.

RTEMS provides two feature macros to indicate the FPU configuration:

- `CPU_HARDWARE_FP` is set to `TRUE` to indicate that a hardware FPU is present.
- `CPU_SOFTWARE_FP` is set to `TRUE` to indicate that a hardware FPU is not present and that the FP software emulation will be context switched.

1.2 Calling Conventions

Each high-level language compiler generates subroutine entry and exit code based upon a set of rules known as the compiler's calling convention. These rules address the following issues:

- register preservation and usage
- parameter passing
- call and return mechanism

A compiler's calling convention is of importance when interfacing to subroutines written in another language either assembly or high-level. Even when the high-level language and target processor are the same, different compilers may use different calling conventions. As a result, calling conventions are both processor and compiler dependent.

1.2.1 Calling Mechanism

In each of the architecture specific chapters, this subsection will describe the instruction(s) used to perform a *normal* subroutine invocation. All RTEMS directives are invoked as *normal* C language functions so it is important to the user application to understand the call and return mechanism.

1.2.2 Register Usage

In each of the architecture specific chapters, this subsection will detail the set of registers which are **NOT** preserved across subroutine invocations. The registers which are not preserved are assumed to be available for use as scratch registers. Therefore, the contents of these registers should not be assumed upon return from any RTEMS directive.

In some architectures, there may be a set of registers made available automatically as a side-effect of the subroutine invocation mechanism.

1.2.3 Parameter Passing

In each of the architecture specific chapters, this subsection will describe the mechanism by which the parameters or arguments are passed by the caller to a subroutine. In some architectures, all parameters are passed on the stack while in others some are passed in registers.

1.2.4 User-Provided Routines

All user-provided routines invoked by RTEMS, such as user extensions, device drivers, and MPCIE routines, must also adhere to these calling conventions.

1.3 Memory Model

A processor may support any combination of memory models ranging from pure physical addressing to complex demand paged virtual memory systems. RTEMS supports a flat memory model which ranges contiguously over the processor's allowable address space. RTEMS does not support segmentation or virtual memory of any kind. The appropriate memory model for RTEMS provided by the targeted processor and related characteristics of that model are described in this chapter.

1.3.1 Flat Memory Model

Most RTEMS target processors can be initialized to support a flat address space. Although the size of addresses varies between architectures, on most RTEMS targets, an address is 32-bits wide which defines addresses ranging from 0x00000000 to 0xFFFFFFFF (4 gigabytes). Each address is represented by a 32-bit value and is byte addressable. The address may be used to reference a single byte, word (2-bytes), or long word (4 bytes). Memory accesses within this address space may be performed in little or big endian fashion.

On smaller CPU architectures supported by RTEMS, the address space may only be 20 or 24 bits wide.

If the CPU model has support for virtual memory or segmentation, it is the responsibility of the Board Support Package (BSP) to initialize the MMU hardware to perform address translations which correspond to flat memory model.

In each of the architecture specific chapters, this subsection will describe any architecture characteristics that differ from this general description.

1.4 Interrupt Processing

Different types of processors respond to the occurrence of an interrupt in its own unique fashion. In addition, each processor type provides a control mechanism to allow for the proper handling of an interrupt. The processor dependent response to the interrupt modifies the current execution state and results in a change in the execution stream. Most processors require that an interrupt handler utilize some special control mechanisms to return to the normal processing stream. Although RTEMS hides many of the processor dependent details

of interrupt processing, it is important to understand how the RTEMS interrupt manager is mapped onto the processor's unique architecture.

RTEMS supports a dedicated interrupt stack for all architectures. On architectures with hardware support for a dedicated interrupt stack, it will be initialized such that when an interrupt occurs, the processor automatically switches to this dedicated stack. On architectures without hardware support for a dedicated interrupt stack which is separate from those of the tasks, RTEMS will support switching to a dedicated stack for interrupt processing.

Without a dedicated interrupt stack, every task in the system **MUST** have enough stack space to accommodate the worst case stack usage of that particular task and the interrupt service routines **COMBINED**. By supporting a dedicated interrupt stack, RTEMS significantly lowers the stack requirements for each task.

A nested interrupt is processed similarly with the exception that since the CPU is already executing on the interrupt stack, there is no need to switch to the interrupt stack.

In some configurations, RTEMS allocates the interrupt stack from the Workspace Area. The amount of memory allocated for the interrupt stack is user configured and based upon the `confdefs.h` parameter `CONFIGURE_INTERRUPT_STACK_SIZE`. This parameter is described in detail in the Configuring a System chapter of the User's Guide. On configurations in which RTEMS allocates the interrupt stack, during the initialization process, RTEMS will also install its interrupt stack. In other configurations, the interrupt stack is allocated and installed by the Board Support Package (BSP).

In each of the architecture specific chapters, this section discusses the interrupt response and control mechanisms of the architecture as they pertain to RTEMS.

1.4.1 Vectoring of an Interrupt Handler

In each of the architecture specific chapters, this subsection will describe the architecture specific details of the interrupt vectoring process. In particular, it should include a description of the Interrupt Stack Frame (ISF).

1.4.2 Interrupt Levels

In each of the architecture specific chapters, this subsection will describe how the interrupt levels available on this particular architecture are mapped onto the 255 reserved in the task mode. The interrupt level value of zero (0) should always mean that interrupts are enabled.

Any use of an interrupt level that is not undefined on a particular architecture may result in behavior that is unpredictable.

1.4.3 Disabling of Interrupts by RTEMS

During the execution of directive calls, critical sections of code may be executed. When these sections are encountered, RTEMS disables all external interrupts before the execution of this section and restores them to the previous level upon completion of the section. RTEMS has been optimized to ensure that interrupts are disabled for the shortest number of instructions possible. Since the precise number of instructions and their execution time varies based upon target CPU family, CPU model, board memory speed, compiler version, and optimization level, it is not practical to provide the precise number for all possible RTEMS configurations.

Historically, the measurements were made by hand analyzing and counting the execution time of instruction sequences during interrupt disable critical sections. For reference purposes, on a 16 Mhz Motorola MC68020, the maximum interrupt disable period was typically approximately ten (10) to thirteen (13) microseconds. This architecture was memory bound and had a slow bit scan instruction. In contrast, during the same period a 14 Mhz SPARC would have a worst case disable time of approximately two (2) to three (3) microseconds because it had a single cycle bit scan instruction and used fewer cycles for memory accesses.

If you are interested in knowing the worst case execution time for a particular version of RTEMS, please contact OAR Corporation and we will be happy to product the results as a consulting service.

Non-maskable interrupts (NMI) cannot be disabled, and ISRs which execute at this level MUST NEVER issue RTEMS system calls. If a directive is invoked, unpredictable results may occur due to the inability of RTEMS to protect its critical sections. However, ISRs that make no system calls may safely execute as non-maskable interrupts.

1.5 Default Fatal Error Processing

Upon detection of a fatal error by either the application or RTEMS during initialization the `rtems_fatal_error_occurred` directive supplied by the Fatal Error Manager is invoked. The Fatal Error Manager will invoke the user-supplied fatal error handlers. If no user-supplied handlers are configured or all of them return without taking action to shutdown the processor or reset, a default fatal error handler is invoked.

Most of the action performed as part of processing the fatal error are described in detail in the Fatal Error Manager chapter in the User's Guide. However, the if no user provided extension or BSP specific fatal error handler takes action, the final default action is to invoke a CPU architecture specific function. Typically this function disables interrupts and halts the processor.

In each of the architecture specific chapters, this describes the precise operations of the default CPU specific fatal error handler.

1.6 Board Support Packages

An RTEMS Board Support Package (BSP) must be designed to support a particular processor model and target board combination.

In each of the architecture specific chapters, this section will present a discussion of architecture specific BSP issues. For more information on developing a BSP, refer to BSP and Device Driver Development Guide and the chapter titled Board Support Packages in the RTEMS Applications User's Guide.

1.6.1 System Reset

An RTEMS based application is initiated or re-initiated when the processor is reset or transfer is passed to it from a boot monitor or ROM monitor.

In each of the architecture specific chapters, this subsection describes the actions that the BSP must tak assuming the application gets control when the microprocessor is reset.

2 ARM Specific Information

This chapter discusses the [ARM architecture](#) dependencies in this port of RTEMS. The ARM family has a wide variety of implementations by a wide range of vendors. Consequently, there are many, many CPU models within it. Currently the ARMv5 (and compatible) architecture version as defined in the [ARMv5 Architecture Reference Manual](#) is supported by RTEMS.

Architecture Documents

For information on the ARM architecture refer to the [ARM Infocenter](#).

2.1 CPU Model Dependent Features

This section presents the set of features which vary across ARM implementations and are of importance to RTEMS. The set of CPU model feature macros are defined in the file `'cpukit/score/cpu/arm/rtems/score/arm.h'` based upon the particular CPU model flags specified on the compilation command line.

2.1.1 CPU Model Name

The macro `CPU_MODEL_NAME` is a string which designates the architectural level of this CPU model. See in `'cpukit/score/cpu/arm/rtems/score/arm.h'` for the values.

2.1.2 Count Leading Zeroes Instruction

The ARMv5 and later has the count leading zeroes `clz` instruction which could be used to speed up the find first bit operation. The use of this instruction should significantly speed up the scheduling associated with a thread blocking. This is currently not used.

2.1.3 Floating Point Unit

A floating point unit is currently not supported.

2.2 Calling Conventions

Please refer to the [Procedure Call Standard for the ARM Architecture](#).

2.3 Memory Model

A flat 32-bit memory model is supported. The board support package must take care about the MMU if necessary.

2.4 Interrupt Processing

The ARMv5 (and compatible) architecture has seven exception types:

- Reset
- Undefined
- Software Interrupt (SWI)
- Prefetch Abort

- Data Abort
- Interrupt (IRQ)
- Fast Interrupt (FIQ)

Of these types only the IRQ has explicit operating system support. It is intentional that the FIQ is not supported by the operating system. Without operating system support for the FIQ it is not necessary to disable them during critical sections of the system.

2.4.1 Interrupt Levels

The RTEMS interrupt level mapping scheme for the ARM is not a numeric level as on most RTEMS ports. It is a bit mapping that corresponds the enable bit positions in the Current Program Status Register (CPSR). There are only two levels: IRQ enabled and IRQ disabled.

2.4.2 Interrupt Stack

The board support package must initialize the interrupt stack. The memory for the stacks is usually reserved in the linker script.

2.5 Default Fatal Error Processing

The default fatal error handler for this architecture performs the following actions:

- disables operating system supported interrupts (IRQ),
- places the error code in `r0`, and
- executes an infinite loop to simulate a halt processor instruction.

3 Atmel AVR Specific Information

This chapter discusses the AVR architecture dependencies in this port of RTEMS.

Architecture Documents

For information on the AVR architecture, refer to the following documents available from Atmel.

TBD

- See other CPUs for documentation reference formatting examples.

3.1 CPU Model Dependent Features

CPUs of the AVR 53X only differ in the peripherals and thus in the device drivers. This port does not yet support the 56X dual core variants.

3.1.1 Count Leading Zeroes Instruction

The AVR CPU has the `XXX` instruction which could be used to speed up the find first bit operation. The use of this instruction should significantly speed up the scheduling associated with a thread blocking.

3.2 Calling Conventions

3.2.1 Processor Background

The AVR architecture supports a simple call and return mechanism. A subroutine is invoked via the call (`call`) instruction. This instruction saves the return address in the `RETS` register and transfers the execution to the given address.

It is the called functions responsibility to use the link instruction to reserve space on the stack for the local variables. Returning from a subroutine is done by using the `RTS` (`RTS`) instruction which loads the PC with the address stored in `RETS`.

It is important to note that the `call` instruction does not automatically save or restore any registers. It is the responsibility of the high-level language compiler to define the register preservation and usage convention.

3.2.2 Register Usage

A called function may clobber all registers, except `RETS`, `R4-R7`, `P3-P5`, `FP` and `SP`. It may also modify the first 12 bytes in the callers stack frame which is used as an argument area for the first three arguments (which are passed in `R0...R3` but may be placed on the stack by the called function).

3.2.3 Parameter Passing

RTEMS assumes that the AVR GCC calling convention is followed. The first three parameters are stored in registers `R0`, `R1`, and `R2`. All other parameters are put pushed on the stack. The result is returned through register `R0`.

3.3 Memory Model

The AVR family architecture support a single unified 4 GB byte address space using 32-bit addresses. It maps all resources like internal and external memory and IO registers into separate sections of this common address space.

The AVR architecture supports some form of memory protection via its Memory Management Unit. Since the AVR port runs in supervisor mode this memory protection mechanisms are not used.

3.4 Interrupt Processing

Discussed in this chapter are the AVR's interrupt response and control mechanisms as they pertain to RTEMS.

3.4.1 Vectoring of an Interrupt Handler

TBD

3.4.2 Disabling of Interrupts by RTEMS

During interrupt disable critical sections, RTEMS disables interrupts to level N (N) before the execution of this section and restores them to the previous level upon completion of the section. RTEMS uses the instructions CLI and STI to enable and disable Interrupts. Emulation, Reset, NMI and Exception Interrupts are never disabled.

3.4.3 Interrupt Stack

The AVR Architecture works with two different kind of stacks, User and Supervisor Stack. Since RTEMS and its Application run in supervisor mode, all interrupts will use the interrupted tasks stack for execution.

3.5 Default Fatal Error Processing

The default fatal error handler for the AVR performs the following actions:

- disables processor interrupts,
- places the error code in **r0**, and
- executes an infinite loop (`while(0)`); to simulate a halt processor instruction.

3.6 Board Support Packages

3.6.1 System Reset

TBD

4 Blackfin Specific Information

This chapter discusses the Blackfin architecture dependencies in this port of RTEMS.

Architecture Documents

For information on the Blackfin architecture, refer to the following documents available from Analog Devices.

TBD

- "*ADSP-BF533 Blackfin Processor Hardware Reference.*"
'http://www.analog.com/UploadedFiles/Associated_Docs/892485982bf533_hwr.pdf'

4.1 CPU Model Dependent Features

CPUs of the Blackfin 53X only differ in the peripherals and thus in the device drivers. This port does not yet support the 56X dual core variants.

4.1.1 Count Leading Zeroes Instruction

The Blackfin CPU has the BITTST instruction which could be used to speed up the find first bit operation. The use of this instruction should significantly speed up the scheduling associated with a thread blocking.

4.2 Calling Conventions

This section is heavily based on content taken from the Blackfin uCLinux documentation wiki which is edited by Analog Devices and Arcturus Networks. '<http://docs.blackfin.uclinux.org/>'

4.2.1 Processor Background

The Blackfin architecture supports a simple call and return mechanism. A subroutine is invoked via the call (`call`) instruction. This instruction saves the return address in the `RETS` register and transfers the execution to the given address.

It is the called functions responsibility to use the link instruction to reserve space on the stack for the local variables. Returning from a subroutine is done by using the `RTS` (`RTS`) instruction which loads the PC with the adress stored in `RETS`.

It is important to note that the `call` instruction does not automatically save or restore any registers. It is the responsibility of the high-level language compiler to define the register preservation and usage convention.

4.2.2 Register Usage

A called function may clobber all registers, except `RETS`, `R4-R7`, `P3-P5`, `FP` and `SP`. It may also modify the first 12 bytes in the callers stack frame which is used as an argument area for the first three arguments (which are passed in `R0...R3` but may be placed on the stack by the called function).

4.2.3 Parameter Passing

RTEMS assumes that the Blackfin GCC calling convention is followed. The first three parameters are stored in registers R0, R1, and R2. All other parameters are put pushed on the stack. The result is returned through register R0.

4.3 Memory Model

The Blackfin family architecture support a single unified 4 GB byte address space using 32-bit addresses. It maps all resources like internal and external memory and IO registers into separate sections of this common address space.

The Blackfin architecture supports some form of memory protection via its Memory Management Unit. Since the Blackfin port runs in supervisor mode this memory protection mechanisms are not used.

4.4 Interrupt Processing

Discussed in this chapter are the Blackfin's interrupt response and control mechanisms as they pertain to RTEMS. The Blackfin architecture support 16 kinds of interrupts broken down into Core and general-purpose interrupts.

4.4.1 Vectoring of an Interrupt Handler

RTEMS maps levels 0 -15 directly to Blackfins event vectors EVT0 - EVT15. Since EVT0 - EVT6 are core events and it is suggested to use EVT15 and EVT15 for Software interrupts, 7 Interrupts (EVT7-EVT13) are left for periferical use.

When installing an RTEMS interrupt handler RTEMS installs a generic Interrupt Handler which saves some context and enables nested interrupt servicing and then vectors to the users interrupt handler.

4.4.2 Disabling of Interrupts by RTEMS

During interrupt disable critical sections, RTEMS disables interrupts to level four (4) before the execution of this section and restores them to the previous level upon completion of the section. RTEMS uses the instructions CLI and STI to enable and disable Interrupts. Emulation, Reset, NMI and Exception Interrupts are never disabled.

4.4.3 Interrupt Stack

The Blackfin Architecture works with two different kind of stacks, User and Supervisor Stack. Since RTEMS and its Application run in supervisor mode, all interrupts will use the interrupted tasks stack for execution.

4.5 Default Fatal Error Processing

The default fatal error handler for the Blackfin performs the following actions:

- disables processor interrupts,
- places the error code in **r0**, and
- executes an infinite loop (`while(0);`) to simulate a halt processor instruction.

4.6 Board Support Packages

4.6.1 System Reset

TBD

5 Intel/AMD x86 Specific Information

This chapter discusses the Intel x86 architecture dependencies in this port of RTEMS. This family has multiple implementations from multiple vendors and suffers more from having evolved rather than being designed for growth.

For information on the i386 processor, refer to the following documents:

- *386 Programmer's Reference Manual, Intel, Order No. 230985-002.*
- *386 Microprocessor Hardware Reference Manual, Intel, Order No. 231732-003.*
- *80386 System Software Writer's Guide, Intel, Order No. 231499-001.*
- *80387 Programmer's Reference Manual, Intel, Order No. 231917-001.*

5.1 CPU Model Dependent Features

This section presents the set of features which vary across i386 implementations and are of importance to RTEMS. The set of CPU model feature macros are defined in the file `cpukit/score/cpu/i386/i386.h` based upon the particular CPU model specified on the compilation command line.

5.1.1 bswap Instruction

The macro `I386_HAS_BSWAP` is set to 1 to indicate that this CPU model has the `bswap` instruction which endian swaps a thirty-two bit quantity. This instruction appears to be present in all CPU models i486's and above.

5.2 Calling Conventions

5.2.1 Processor Background

The i386 architecture supports a simple yet effective call and return mechanism. A subroutine is invoked via the call (`call`) instruction. This instruction pushes the return address on the stack. The return from subroutine (`ret`) instruction pops the return address off the current stack and transfers control to that instruction. It is important to note that the i386 call and return mechanism does not automatically save or restore any registers. It is the responsibility of the high-level language compiler to define the register preservation and usage convention.

5.2.2 Calling Mechanism

All RTEMS directives are invoked using a call instruction and return to the user application via the `ret` instruction.

5.2.3 Register Usage

As discussed above, the call instruction does not automatically save any registers. RTEMS uses the registers EAX, ECX, and EDI as scratch registers. These registers are not preserved by RTEMS directives therefore, the contents of these registers should not be assumed upon return from any RTEMS directive.

5.2.4 Parameter Passing

RTEMS assumes that arguments are placed on the current stack before the directive is invoked via the call instruction. The first argument is assumed to be closest to the return address on the stack. This means that the first argument of the C calling sequence is pushed last. The following pseudo-code illustrates the typical sequence used to call a RTEMS directive with three (3) arguments:

```
push third argument
push second argument
push first argument
invoke directive
remove arguments from the stack
```

The arguments to RTEMS are typically pushed onto the stack using a push instruction. These arguments must be removed from the stack after control is returned to the caller. This removal is typically accomplished by adding the size of the argument list in bytes to the stack pointer.

5.3 Memory Model

5.3.1 Flat Memory Model

RTEMS supports the i386 protected mode, flat memory model with paging disabled. In this mode, the i386 automatically converts every address from a logical to a physical address each time it is used. The i386 uses information provided in the segment registers and the Global Descriptor Table to convert these addresses. RTEMS assumes the existence of the following segments:

- a single code segment at protection level (0) which contains all application and executive code.
- a single data segment at protection level zero (0) which contains all application and executive data.

The i386 segment registers and associated selectors must be initialized when the `initialize_executive` directive is invoked. RTEMS treats the segment registers as system registers and does not modify or context switch them.

This i386 memory model supports a flat 32-bit address space with addresses ranging from 0x00000000 to 0xFFFFFFFF (4 gigabytes). Each address is represented by a 32-bit value and is byte addressable. The address may be used to reference a single byte, half-word (2-bytes), or word (4 bytes).

5.4 Interrupt Processing

Although RTEMS hides many of the processor dependent details of interrupt processing, it is important to understand how the RTEMS interrupt manager is mapped onto the processor's unique architecture. Discussed in this chapter are the the processor's response and control mechanisms as they pertain to RTEMS.

5.4.1 Vectoring of Interrupt Handler

Although the i386 supports multiple privilege levels, RTEMS and all user software executes at privilege level 0. This decision was made by the RTEMS designers to enhance compatibility with processors which do not provide sophisticated protection facilities like those of the i386. This decision greatly simplifies the discussion of i386 processing, as one need only consider interrupts without privilege transitions.

Upon receipt of an interrupt the i386 automatically performs the following actions:

- pushes the EFLAGS register
- pushes the far address of the interrupted instruction
- vectors to the interrupt service routine (ISR).

A nested interrupt is processed similarly by the i386.

5.4.2 Interrupt Stack Frame

The structure of the Interrupt Stack Frame for the i386 which is placed on the interrupt stack by the processor in response to an interrupt is as follows:

Old EFLAGS	Register	ESP+8
UNUSED	Old CS	ESP+4
Old	EIP	ESP

5.4.3 Interrupt Levels

Although RTEMS supports 256 interrupt levels, the i386 only supports two – enabled and disabled. Interrupts are enabled when the interrupt-enable flag (IF) in the extended flags (EFLAGS) is set. Conversely, interrupt processing is inhibited when the IF is cleared. During a non-maskable interrupt, all other interrupts, including other non-maskable ones, are inhibited.

RTEMS interrupt levels 0 and 1 such that level zero (0) indicates that interrupts are fully enabled and level one that interrupts are disabled. All other RTEMS interrupt levels are undefined and their behavior is unpredictable.

5.4.4 Interrupt Stack

The i386 family does not support a dedicated hardware interrupt stack. On this processor, RTEMS allocates and manages a dedicated interrupt stack. As part of vectoring a non-nested interrupt service routine, RTEMS switches from the stack of the interrupted task to a dedicated interrupt stack. When a non-nested interrupt returns, RTEMS switches back to the stack of the interrupted stack. The current stack pointer is not altered by RTEMS on nested interrupt.

5.5 Default Fatal Error Processing

The default fatal error handler for this architecture disables processor interrupts, places the error code in EAX, and executes a HLT instruction to halt the processor.

5.6 Board Support Packages

5.6.1 System Reset

An RTEMS based application is initiated when the i386 processor is reset. When the i386 is reset,

- The EAX register is set to indicate the results of the processor's power-up self test. If the self-test was not executed, the contents of this register are undefined. Otherwise, a non-zero value indicates the processor is faulty and a zero value indicates a successful self-test.
- The DX register holds a component identifier and revision level. DH contains 3 to indicate an i386 component and DL contains a unique revision level indicator.
- Control register zero (CR0) is set such that the processor is in real mode with paging disabled. Other portions of CR0 are used to indicate the presence of a numeric coprocessor.
- All bits in the extended flags register (EFLAG) which are not permanently set are cleared. This inhibits all maskable interrupts.
- The Interrupt Descriptor Register (IDTR) is set to point at address zero.
- All segment registers are set to zero.
- The instruction pointer is set to 0x0000FFF0. The first instruction executed after a reset is actually at 0xFFFFFFF0 because the i386 asserts the upper twelve address until the first intersegment (FAR) JMP or CALL instruction. When a JMP or CALL is executed, the upper twelve address lines are lowered and the processor begins executing in the first megabyte of memory.

Typically, an intersegment JMP to the application's initialization code is placed at address 0xFFFFFFF0.

5.6.2 Processor Initialization

This initialization code is responsible for initializing all data structures required by the i386 in protected mode and for actually entering protected mode. The i386 must be placed in protected mode and the segment registers and associated selectors must be initialized before the `initialize_executive` directive is invoked.

The initialization code is responsible for initializing the Global Descriptor Table such that the i386 is in the thirty-two bit flat memory model with paging disabled. In this mode, the i386 automatically converts every address from a logical to a physical address each time it is used. For more information on the memory model used by RTEMS, please refer to the Memory Model chapter in this document.

Since the processor is in real mode upon reset, the processor must be switched to protected mode before RTEMS can execute. Before switching to protected mode, at least one descriptor table and two descriptors must be created. Descriptors are needed for a code segment and a data segment. (This will give you the flat memory model.) The stack can be placed in a normal read/write data segment, so no descriptor for the stack is needed. Before the GDT can be used, the base address and limit must be loaded into the GDTR register using an LGDT instruction.

If the hardware allows an NMI to be generated, you need to create the IDT and a gate for the NMI interrupt handler. Before the IDT can be used, the base address and limit for the idt must be loaded into the IDTR register using an LIDT instruction.

Protected mode is entered by setting the PE bit in the CR0 register. Either a LMSW or MOV CR0 instruction may be used to set this bit. Because the processor overlaps the interpretation of several instructions, it is necessary to discard the instructions from the read-ahead cache. A JMP instruction immediately after the LMSW changes the flow and empties the processor of instructions which have been pre-fetched and/or decoded. At this point, the processor is in protected mode and begins to perform protected mode application initialization.

If the application requires that the IDTR be some value besides zero, then it should set it to the required value at this point. All tasks share the same i386 IDTR value. Because interrupts are enabled automatically by RTEMS as part of the initialize_executive directive, the IDTR MUST be set properly before this directive is invoked to insure correct interrupt vectoring. If processor caching is to be utilized, then it should be enabled during the reset application initialization code. The reset code which is executed before the call to initialize_executive has the following requirements:

For more information regarding the i386 data structures and their contents, refer to Intel's 386 Programmer's Reference Manual.

6 Lattice Mico32 Specific Information

This chapter discusses the Lattice Mico32 architecture dependencies in this port of RTEMS. The Lattice Mico32 is a 32-bit Harvard, RISC architecture "soft" microprocessor, available for free with an open IP core licensing agreement. Although mainly targeted for Lattice FPGA devices the microprocessor can be implemented on other vendors' FPGAs, too.

Architecture Documents

For information on the Lattice Mico32 architecture, refer to the following documents available from Lattice Semiconductor '<http://www.latticesemi.com/>'.

- "*LatticeMico32 Processor Reference Manual*"
'http://www.latticesemi.com/dynamic/view_document.cfm?document_id=20890'

6.1 CPU Model Dependent Features

The Lattice Mico32 architecture allows for different configurations of the processor. This port is based on the assumption that the following options are implemented:

- hardware multiplier
- hardware divider
- hardware barrel shifter
- sign extension instructions
- instruction cache
- data cache
- debug

6.2 Register Architecture

This section gives a brief introduction to the register architecture of the Lattice Mico32 processor.

The Lattice Mico32 is a RISC architecture processor with a 32-register file of 32-bit registers.

Register Name	Function
r0	holds value zero
r1-r25	general purpose
r26/gp	general purpose / global pointer
r27/fp	general purpose / frame pointer
r28/sp	stack pointer
r29/ra	return address
r30/ea	exception address
r31/ba	breakpoint address

Note that on processor startup all register values are undefined including r0, thus r0 has to be initialized to zero.

6.3 Calling Conventions

6.3.1 Calling Mechanism

A call instruction places the return address to register r29 and a return from subroutine (ret) is actually a branch to r29/ra.

6.3.2 Register Usage

A subroutine may freely use registers r1 to r10 which are **not** preserved across subroutine invocations.

6.3.3 Parameter Passing

When calling a C function the first eight arguments are stored in registers r1 to r8. Registers r1 and r2 hold the return value.

6.4 Memory Model

The Lattice Mico32 processor supports a flat memory model with a 4 Gbyte address space with 32-bit addresses.

The following data types are supported:

Type	Bits	C Compiler Type
unsigned byte	8	unsigned char
signed byte	8	char
unsigned half-word	16	unsigned short
signed half-word	16	short
unsigned word	32	unsigned int / unsigned long
signed word	32	int / long

Data accesses need to be aligned, with unaligned accesses result are undefined.

6.5 Interrupt Processing

The Lattice Mico32 has 32 interrupt lines which are however served by only one exception vector. When an interrupt occurs following happens:

- address of next instruction placed in r30/ea
- IE field of IE CSR saved to EIE field and IE field cleared preventing further exceptions from occurring.
- branch to interrupt exception address EBA CSR + 0xC0

The interrupt exception handler determines from the state of the interrupt pending registers (IP CSR) and interrupt enable register (IE CSR) which interrupt to serve and jumps to the interrupt routine pointed to by the corresponding interrupt vector.

For now there is no dedicated interrupt stack so every task in the system **MUST** have enough stack space to accommodate the worst case stack usage of that particular task and the interrupt service routines **COMBINED**.

Nested interrupts are not supported.

6.6 Default Fatal Error Processing

Upon detection of a fatal error by either the application or RTEMS during initialization the `rtems_fatal_error_occurred` directive supplied by the Fatal Error Manager is invoked. The Fatal Error Manager will invoke the user-supplied fatal error handlers. If no user-supplied handlers are configured or all of them return without taking action to shutdown the processor or reset, a default fatal error handler is invoked.

Most of the action performed as part of processing the fatal error are described in detail in the Fatal Error Manager chapter in the User's Guide. However, the if no user provided extension or BSP specific fatal error handler takes action, the final default action is to invoke a CPU architecture specific function. Typically this function disables interrupts and halts the processor.

In each of the architecture specific chapters, this describes the precise operations of the default CPU specific fatal error handler.

6.7 Board Support Packages

An RTEMS Board Support Package (BSP) must be designed to support a particular processor model and target board combination.

In each of the architecture specific chapters, this section will present a discussion of architecture specific BSP issues. For more information on developing a BSP, refer to BSP and Device Driver Development Guide and the chapter titled Board Support Packages in the RTEMS Applications User's Guide.

6.7.1 System Reset

An RTEMS based application is initiated or re-initiated when the processor is reset.

7 M68xxx and Coldfire Specific Information

This chapter discusses the Freescale (formerly Motorola) MC68xxx and Coldfire architectural dependencies. The MC68xxx family has a wide variety of CPU models within it based upon different CPU core implementations. Ignoring the Coldfire parts, the part numbers for these models are generally divided into MC680xx and MC683xx. The MC680xx models are more general purpose processors with no integrated peripherals. The MC683xx models, on the other hand, are more specialized and have a variety of peripherals on chip including sophisticated timers and serial communications controllers.

Architecture Documents

For information on the MC68xxx and Coldfire architecture, refer to the following documents available from Freescale website (<http://www.freescale.com/>):

- *M68000 Family Reference, Motorola, FR68K/D.*
- *MC68020 User's Manual, Motorola, MC68020UM/AD.*
- *MC68881/MC68882 Floating-Point Coprocessor User's Manual, Motorola, MC68881UM/AD.*

7.1 CPU Model Dependent Features

This section presents the set of features which vary across m68k/Coldfire implementations that are of importance to RTEMS. The set of CPU model feature macros are defined in the file `cpukit/score/cpu/m68k/m68k.h` based upon the particular CPU model selected on the compilation command line.

7.1.1 BFFFO Instruction

The macro `M68K_HAS_BFFFO` is set to 1 to indicate that this CPU model has the `bfffo` instruction.

7.1.2 Vector Base Register

The macro `M68K_HAS_VBR` is set to 1 to indicate that this CPU model has a vector base register (`vbr`).

7.1.3 Separate Stacks

The macro `M68K_HAS_SEPARATE_STACKS` is set to 1 to indicate that this CPU model has separate interrupt, user, and supervisor mode stacks.

7.1.4 Pre-Indexing Address Mode

The macro `M68K_HAS_PREINDEXING` is set to 1 to indicate that this CPU model has the pre-indexing address mode.

7.1.5 Extend Byte to Long Instruction

The macro `M68K_HAS_EXTB_L` is set to 1 to indicate that this CPU model has the `extb.l` instruction. This instruction is supposed to be available in all models based on the `cpu32` core as well as `mc68020` and `up` models.

7.2 Calling Conventions

The MC68xxx architecture supports a simple yet effective call and return mechanism. A subroutine is invoked via the branch to subroutine (**bsr**) or the jump to subroutine (**jsr**) instructions. These instructions push the return address on the current stack. The return from subroutine (**rts**) instruction pops the return address off the current stack and transfers control to that instruction. It is important to note that the MC68xxx call and return mechanism does not automatically save or restore any registers. It is the responsibility of the high-level language compiler to define the register preservation and usage convention.

7.2.1 Calling Mechanism

All RTEMS directives are invoked using either a **bsr** or **jsr** instruction and return to the user application via the **rts** instruction.

7.2.2 Register Usage

As discussed above, the **bsr** and **jsr** instructions do not automatically save any registers. RTEMS uses the registers D0, D1, A0, and A1 as scratch registers. These registers are not preserved by RTEMS directives therefore, the contents of these registers should not be assumed upon return from any RTEMS directive.

7.2.3 Parameter Passing

RTEMS assumes that arguments are placed on the current stack before the directive is invoked via the **bsr** or **jsr** instruction. The first argument is assumed to be closest to the return address on the stack. This means that the first argument of the C calling sequence is pushed last. The following pseudo-code illustrates the typical sequence used to call a RTEMS directive with three (3) arguments:

```
push third argument
push second argument
push first argument
invoke directive
remove arguments from the stack
```

The arguments to RTEMS are typically pushed onto the stack using a move instruction with a pre-decremented stack pointer as the destination. These arguments must be removed from the stack after control is returned to the caller. This removal is typically accomplished by adding the size of the argument list in bytes to the current stack pointer.

7.3 Memory Model

The MC68xxx family supports a flat 32-bit address space with addresses ranging from 0x00000000 to 0xFFFFFFFF (4 gigabytes). Each address is represented by a 32-bit value and is byte addressable. The address may be used to reference a single byte, word (2-bytes), or long word (4 bytes). Memory accesses within this address space are performed in big endian fashion by the processors in this family.

Some of the MC68xxx family members such as the MC68020, MC68030, and MC68040 support virtual memory and segmentation. The MC68020 requires external hardware support such as the MC68851 Paged Memory Management Unit coprocessor which is typically

used to perform address translations for these systems. RTEMS does not support virtual memory or segmentation on any of the MC68xxx family members.

7.4 Interrupt Processing

Discussed in this section are the MC68xxx's interrupt response and control mechanisms as they pertain to RTEMS.

7.4.1 Vectoring of an Interrupt Handler

Depending on whether or not the particular CPU supports a separate interrupt stack, the MC68xxx family has two different interrupt handling models.

7.4.1.1 Models Without Separate Interrupt Stacks

Upon receipt of an interrupt the MC68xxx family members without separate interrupt stacks automatically perform the following actions:

- To Be Written

7.4.1.2 Models With Separate Interrupt Stacks

Upon receipt of an interrupt the MC68xxx family members with separate interrupt stacks automatically perform the following actions:

- saves the current status register (SR),
- clears the master/interrupt (M) bit of the SR to indicate the switch from master state to interrupt state,
- sets the privilege mode to supervisor,
- suppresses tracing,
- sets the interrupt mask level equal to the level of the interrupt being serviced,
- pushes an interrupt stack frame (ISF), which includes the program counter (PC), the status register (SR), and the format/exception vector offset (FVO) word, onto the supervisor and interrupt stacks,
- switches the current stack to the interrupt stack and vectors to an interrupt service routine (ISR). If the ISR was installed with the `interrupt_catch` directive, then the RTEMS interrupt handler will begin execution. The RTEMS interrupt handler saves all registers which are not preserved according to the calling conventions and invokes the application's ISR.

A nested interrupt is processed similarly by these CPU models with the exception that only a single ISF is placed on the interrupt stack and the current stack need not be switched.

The FVO word in the Interrupt Stack Frame is examined by RTEMS to determine when an outer most interrupt is being exited. Since the FVO is used by RTEMS for this purpose, the user application code **MUST NOT** modify this field.

The following shows the Interrupt Stack Frame for MC68xxx CPU models with separate interrupt stacks:

Status Register	0x0
Program Counter High	0x2
Program Counter Low	0x4
Format/Vector Offset	0x6

7.4.2 CPU Models Without VBR and RAM at 0

This is from a post by Zoltan Kocsi <zoltan@bendor.com.au> and is a nice trick in certain situations. In his words:

I think somebody on this list asked about the interrupt vector handling w/o VBR and RAM at 0. The usual trick is to initialise the vector table (except the first 2 two entries, of course) to point to the same location BUT you also add the vector number times 0x1000000 to them. That is, bits 31-24 contain the vector number and 23-0 the address of the common handler. Since the PC is 32 bit wide but the actual address bus is only 24, the top byte will be in the PC but will be ignored when jumping onto your routine.

Then your common interrupt routine gets this info by loading the PC into some register and based on that info, you can jump to a vector in a vector table pointed by a virtual VBR:

```
//
// Real vector table at 0
//

        .long   initial_sp
        .long   initial_pc
        .long   myhandler+0x02000000
        .long   myhandler+0x03000000
        .long   myhandler+0x04000000
        ...
        .long   myhandler+0xff000000

//
// This handler will jump to the interrupt routine   of which
// the address is stored at VBR[ vector_no ]
// The registers and stackframe will be intact, the interrupt
// routine will see exactly what it would see if it was called
// directly from the HW vector table at 0.
//

        .comm   VBR,4,2           // This defines the 'virtual' VBR
                                   // From C: extern void *VBR;

myhandler:                               // At entry, PC contains the full vector
        move.l  %d0,-(%sp)         // Save d0
        move.l  %a0,-(%sp)         // Save a0
        lea    0(%pc),%a0          // Get the value of the PC
        move.l  %a0,%d0            // Copy it to a data reg, d0 is VV??????
        swap   %d0                 // Now d0 is ???V??
```

```

and.w   #0xff00,%d0    // Now d0 is ???V00 (1)
lsl.w   #6,%d0        // Now d0.w contains the VBR table offset
move.l  VBR,%a0       // Get the address from VBR to a0
move.l  (%a0,%d0.w),%a0 // Fetch the vector
move.l  4(%sp),%d0    // Restore d0
move.l  %a0,4(%sp)    // Place target address to the stack
move.l  (%sp)+,%a0    // Restore a0, target address is on TOS
ret                                           // This will jump to the handler and
                                           // restore the stack

```

- (1) If 'myhandler' is guaranteed to be in the first 64K, e.g. just after the vector table then that insn is not needed.

There are probably shorter ways to do this, but it I believe is enough to illustrate the trick. Optimisation is left as an exercise to the reader :-)

7.4.3 Interrupt Levels

Eight levels (0-7) of interrupt priorities are supported by MC68xxx family members with level seven (7) being the highest priority. Level zero (0) indicates that interrupts are fully enabled. Interrupt requests for interrupts with priorities less than or equal to the current interrupt mask level are ignored.

Although RTEMS supports 256 interrupt levels, the MC68xxx family only supports eight. RTEMS interrupt levels 0 through 7 directly correspond to MC68xxx interrupt levels. All other RTEMS interrupt levels are undefined and their behavior is unpredictable.

7.5 Default Fatal Error Processing

The default fatal error handler for this architecture disables processor interrupts to level 7, places the error code in D0, and executes a `stop` instruction to simulate a halt processor instruction.

7.6 Board Support Packages

7.6.1 System Reset

An RTEMS based application is initiated or re-initiated when the MC68020 processor is reset. When the MC68020 is reset, the processor performs the following actions:

- The tracing bits of the status register are cleared to disable tracing.
- The supervisor interrupt state is entered by setting the supervisor (S) bit and clearing the master/interrupt (M) bit of the status register.
- The interrupt mask of the status register is set to level 7 to effectively disable all maskable interrupts.
- The vector base register (VBR) is set to zero.
- The cache control register (CACR) is set to zero to disable and freeze the processor cache.

- The interrupt stack pointer (ISP) is set to the value stored at vector 0 (bytes 0-3) of the exception vector table (EVT).
- The program counter (PC) is set to the value stored at vector 1 (bytes 4-7) of the EVT.
- The processor begins execution at the address stored in the PC.

7.6.2 Processor Initialization

The address of the application's initialization code should be stored in the first vector of the EVT which will allow the immediate vectoring to the application code. If the application requires that the VBR be some value besides zero, then it should be set to the required value at this point. All tasks share the same MC68020's VBR value. Because interrupts are enabled automatically by RTEMS as part of the context switch to the first task, the VBR MUST be set by either RTEMS or the BSP before this occurs ensure correct interrupt vectoring. If processor caching is to be utilized, then it should be enabled during the reset application initialization code.

In addition to the requirements described in the Board Support Packages chapter of the Applications User's Manual for the reset code which is executed before the call to initialize executive, the MC68020 version has the following specific requirements:

- Must leave the S bit of the status register set so that the MC68020 remains in the supervisor state.
- Must set the M bit of the status register to remove the MC68020 from the interrupt state.
- Must set the master stack pointer (MSP) such that a minimum stack size of `MINIMUM_STACK_SIZE` bytes is provided for the initialize executive directive.
- Must initialize the MC68020's vector table.

8 MIPS Specific Information

This chapter discusses the MIPS architecture dependencies in this port of RTEMS. The MIPS family has a wide variety of implementations by a wide range of vendors. Consequently, there are many, many CPU models within it.

Architecture Documents

IDT docs are online at <http://www.idt.com/products/risc/Welcome.html>

For information on the XXX architecture, refer to the following documents available from VENDOR (`'http://www.XXX.com/'`):

- *XXX Family Reference, VENDOR, PART NUMBER.*

8.1 CPU Model Dependent Features

This section presents the set of features which vary across MIPS implementations and are of importance to RTEMS. The set of CPU model feature macros are defined in the file `cpukit/score/cpu/mips/mips.h` based upon the particular CPU model specified on the compilation command line.

8.1.1 Another Optional Feature

The macro XXX

8.2 Calling Conventions

8.2.1 Processor Background

TBD

8.2.2 Calling Mechanism

TBD

8.2.3 Register Usage

TBD

8.2.4 Parameter Passing

TBD

8.3 Memory Model

8.3.1 Flat Memory Model

The MIPS family supports a flat 32-bit address space with addresses ranging from `0x00000000` to `0xFFFFFFFF` (4 gigabytes). Each address is represented by a 32-bit value and is byte addressable. The address may be used to reference a single byte, word (2-bytes), or long word (4 bytes). Memory accesses within this address space are performed in big endian fashion by the processors in this family.

Some of the MIPS family members such as the support virtual memory and segmentation. RTEMS does not support virtual memory or segmentation on any of these family members.

8.4 Interrupt Processing

Although RTEMS hides many of the processor dependent details of interrupt processing, it is important to understand how the RTEMS interrupt manager is mapped onto the processor's unique architecture. Discussed in this chapter are the MIPS's interrupt response and control mechanisms as they pertain to RTEMS.

8.4.1 Vectoring of an Interrupt Handler

Upon receipt of an interrupt the XXX family members with separate interrupt stacks automatically perform the following actions:

- TBD

A nested interrupt is processed similarly by these CPU models with the exception that only a single ISF is placed on the interrupt stack and the current stack need not be switched.

8.4.2 Interrupt Levels

TBD

8.5 Default Fatal Error Processing

The default fatal error handler for this target architecture disables processor interrupts, places the error code in **XXX**, and executes a XXX instruction to simulate a halt processor instruction.

8.6 Board Support Packages

8.6.1 System Reset

An RTEMS based application is initiated or re-initiated when the processor is reset. When the processor is reset, it performs the following actions:

- TBD

8.6.2 Processor Initialization

TBD

9 PowerPC Specific Information

This chapter discusses the PowerPC architecture dependencies in this port of RTEMS. The PowerPC family has a wide variety of implementations by a range of vendors. Consequently, there are many, many CPU models within it.

It is highly recommended that the PowerPC RTEMS application developer obtain and become familiar with the documentation for the processor being used as well as the specification for the revision of the PowerPC architecture which corresponds to that processor.

PowerPC Architecture Documents

For information on the PowerPC architecture, refer to the following documents available from Motorola and IBM:

- *PowerPC Microprocessor Family: The Programming Environment* (Motorola Document MPRPPCFPE-01).
- *IBM PPC403GB Embedded Controller User's Manual*.
- *PowerPC MPC500 Family RISC Central Processing Unit Reference Manual* (Motorola Document RCPUURM/AD).
- *PowerPC 601 RISC Microprocessor User's Manual* (Motorola Document MPR601UM/AD).
- *PowerPC 603 RISC Microprocessor User's Manual* (Motorola Document MPR603UM/AD).
- *PowerPC 603e RISC Microprocessor User's Manual* (Motorola Document MPR603EUM/AD).
- *PowerPC 604 RISC Microprocessor User's Manual* (Motorola Document MPR604UM/AD).
- *PowerPC MPC821 Portable Systems Microprocessor User's Manual* (Motorola Document MPC821UM/AD).
- *PowerQUICC MPC860 User's Manual* (Motorola Document MPC860UM/AD).

Motorola maintains an on-line electronic library for the PowerPC at the following URL:

<http://www.mot.com/powerpc/library/library.html>

This site has a wealth of information and examples. Many of the manuals are available from that site in electronic format.

PowerPC Processor Simulator Information

PSIM is a program which emulates the Instruction Set Architecture of the PowerPC microprocessor family. It is reely available in source code form under the terms of the GNU General Public License (version 2 or later). PSIM can be integrated with the GNU Debugger (gdb) to execute and debug PowerPC executables on non-PowerPC hosts. PSIM supports the addition of user provided device models which can be used to allow one to develop and debug embedded applications using the simulator.

The latest version of PSIM is included in GDB and enabled on pre-built binaries provided by the RTEMS Project.

9.1 CPU Model Dependent Features

This section presents the set of features which vary across PowerPC implementations and are of importance to RTEMS. The set of CPU model feature macros are defined in the file `cpukit/score/cpu/powerpc/powerpc.h` based upon the particular CPU model specified on the compilation command line.

9.1.1 Alignment

The macro `PPC_ALIGNMENT` is set to the PowerPC model's worst case alignment requirement for data types on a byte boundary. This value is used to derive the alignment restrictions for memory allocated from regions and partitions.

9.1.2 Cache Alignment

The macro `PPC_CACHE_ALIGNMENT` is set to the line size of the cache. It is used to align the entry point of critical routines so that as much code as possible can be retrieved with the initial read into cache. This is done for the interrupt handler as well as the context switch routines.

In addition, the "shortcut" data structure used by the PowerPC implementation to ease access to data elements frequently accessed by RTEMS routines implemented in assembly language is aligned using this value.

9.1.3 Maximum Interrupts

The macro `PPC_INTERRUPT_MAX` is set to the number of exception sources supported by this PowerPC model.

9.1.4 Has Double Precision Floating Point

The macro `PPC_HAS_DOUBLE` is set to 1 to indicate that the PowerPC model has support for double precision floating point numbers. This is important because the floating point registers need only be four bytes wide (not eight) if double precision is not supported.

9.1.5 Critical Interrupts

The macro `PPC_HAS_RFCI` is set to 1 to indicate that the PowerPC model has the Critical Interrupt capability as defined by the IBM 403 models.

9.1.6 Use Multiword Load/Store Instructions

The macro `PPC_USE_MULTIPLE` is set to 1 to indicate that multiword load and store instructions should be used to perform context switch operations. The relative efficiency of multiword load and store instructions versus an equivalent set of single word load and store instructions varies based upon the PowerPC model.

9.1.7 Instruction Cache Size

The macro `PPC_I_CACHE` is set to the size in bytes of the instruction cache.

9.1.8 Data Cache Size

The macro `PPC_D_CACHE` is set to the size in bytes of the data cache.

9.1.9 Debug Model

The macro `PPC_DEBUG_MODEL` is set to indicate the debug support features present in this CPU model. The following debug support feature sets are currently supported:

`PPC_DEBUG_MODEL_STANDARD`

indicates that the single-step trace enable (SE) and branch trace enable (BE) bits in the MSR are supported by this CPU model.

`PPC_DEBUG_MODEL_SINGLE_STEP_ONLY`

indicates that only the single-step trace enable (SE) bit in the MSR is supported by this CPU model.

`PPC_DEBUG_MODEL_IBM4xx`

indicates that the debug exception enable (DE) bit in the MSR is supported by this CPU model. At this time, this particular debug feature set has only been seen in the IBM 4xx series.

9.1.9.1 Low Power Model

The macro `PPC_LOW_POWER_MODE` is set to indicate the low power model supported by this CPU model. The following low power modes are currently supported.

`PPC_LOW_POWER_MODE_NONE`

indicates that this CPU model has no low power mode support.

`PPC_LOW_POWER_MODE_STANDARD`

indicates that this CPU model follows the low power model defined for the PPC603e.

9.2 Calling Conventions

RTEMS supports the Embedded Application Binary Interface (EABI) calling convention. Documentation for EABI is available by sending a message with a subject line of "EABI" to eabi@goth.sis.mot.com.

9.2.1 Programming Model

This section discusses the programming model for the PowerPC architecture.

9.2.1.1 Non-Floating Point Registers

The PowerPC architecture defines thirty-two non-floating point registers directly visible to the programmer. In thirty-two bit implementations, each register is thirty-two bits wide. In sixty-four bit implementations, each register is sixty-four bits wide.

These registers are referred to as `gpr0` to `gpr31`.

Some of the registers serve defined roles in the EABI programming model. The following table describes the role of each of these registers:

Register Name	Alternate Names	Description
r1	sp	stack pointer
r2	NA	global pointer to the Small Constant Area (SDA2)
r3 - r12	NA	parameter and result passing
r13	NA	global pointer to the Small Data Area (SDA2)

9.2.1.2 Floating Point Registers

The PowerPC architecture includes thirty-two, sixty-four bit floating point registers. All PowerPC floating point instructions interpret these registers as 32 double precision floating point registers, regardless of whether the processor has 64-bit or 32-bit implementation.

The floating point status and control register (fpSCR) records exceptions and the type of result generated by floating-point operations. Additionally, it controls the rounding mode of operations and allows the reporting of floating exceptions to be enabled or disabled.

9.2.1.3 Special Registers

The PowerPC architecture includes a number of special registers which are critical to the programming model:

Machine State Register

The MSR contains the processor mode, power management mode, endian mode, exception information, privilege level, floating point available and floating point exception mode, address translation information and the exception prefix.

Link Register

The LR contains the return address after a function call. This register must be saved before a subsequent subroutine call can be made. The use of this register is discussed further in the **Call and Return Mechanism** section below.

Count Register

The CTR contains the iteration variable for some loops. It may also be used for indirect function calls and jumps.

9.2.2 Call and Return Mechanism

The PowerPC architecture supports a simple yet effective call and return mechanism. A subroutine is invoked via the "branch and link" (**bl**) and "branch and link absolute" (**bla**) instructions. These instructions place the return address in the Link Register (LR). The callee returns to the caller by executing a "branch unconditional to the link register" (**blr**) instruction. Thus the callee returns to the caller via a jump to the return address which is stored in the LR.

The previous contents of the LR are not automatically saved by either the **bl** or **bla**. It is the responsibility of the callee to save the contents of the LR before invoking another subroutine. If the callee invokes another subroutine, it must restore the LR before executing the **blr** instruction to return to the caller.

It is important to note that the PowerPC subroutine call and return mechanism does not automatically save and restore any registers.

The LR may be accessed as special purpose register 8 (SPR8) using the "move from special register" (`mfspr`) and "move to special register" (`mtspr`) instructions.

9.2.3 Calling Mechanism

All RTEMS directives are invoked using the regular PowerPC EABI calling convention via the `bl` or `bla` instructions.

9.2.4 Register Usage

As discussed above, the call instruction does not automatically save any registers. It is the responsibility of the callee to save and restore any registers which must be preserved across subroutine calls. The callee is responsible for saving callee-preserved registers to the program stack and restoring them before returning to the caller.

9.2.5 Parameter Passing

RTEMS assumes that arguments are placed in the general purpose registers with the first argument in register 3 (`r3`), the second argument in general purpose register 4 (`r4`), and so forth until the seventh argument is in general purpose register 10 (`r10`). If there are more than seven arguments, then subsequent arguments are placed on the program stack. The following pseudo-code illustrates the typical sequence used to call a RTEMS directive with three (3) arguments:

```
load third argument into r5
load second argument into r4
load first argument into r3
invoke directive
```

9.3 Memory Model

9.3.1 Flat Memory Model

The PowerPC architecture supports a variety of memory models. RTEMS supports the PowerPC using a flat memory model with paging disabled. In this mode, the PowerPC automatically converts every address from a logical to a physical address each time it is used. The PowerPC uses information provided in the Block Address Translation (BAT) to convert these addresses.

Implementations of the PowerPC architecture may be thirty-two or sixty-four bit. The PowerPC architecture supports a flat thirty-two or sixty-four bit address space with addresses ranging from `0x00000000` to `0xFFFFFFFF` (4 gigabytes) in thirty-two bit implementations or to `0xFFFFFFFFFFFFFFFF` in sixty-four bit implementations. Each address is represented by either a thirty-two bit or sixty-four bit value and is byte addressable. The address may be used to reference a single byte, half-word (2-bytes), word (4 bytes), or in sixty-four bit implementations a doubleword (8 bytes). Memory accesses within the address space are performed in big or little endian fashion by the PowerPC based upon the current setting of the Little-endian mode enable bit (LE) in the Machine State Register (MSR). While the

processor is in big endian mode, memory accesses which are not properly aligned generate an "alignment exception" (vector offset 0x00600). In little endian mode, the PowerPC architecture does not require the processor to generate alignment exceptions.

The following table lists the alignment requirements for a variety of data accesses:

Data Type	Alignment Requirement
byte	1
half-word	2
word	4
doubleword	8

Doubleword load and store operations are only available in PowerPC CPU models which are sixty-four bit implementations.

RTEMS does not directly support any PowerPC Memory Management Units, therefore, virtual memory or segmentation systems involving the PowerPC are not supported.

9.4 Interrupt Processing

Although RTEMS hides many of the processor dependent details of interrupt processing, it is important to understand how the RTEMS interrupt manager is mapped onto the processor's unique architecture. Discussed in this chapter are the PowerPC's interrupt response and control mechanisms as they pertain to RTEMS.

RTEMS and associated documentation uses the terms interrupt and vector. In the PowerPC architecture, these terms correspond to exception and exception handler, respectively. The terms will be used interchangeably in this manual.

9.4.1 Synchronous Versus Asynchronous Exceptions

In the PowerPC architecture exceptions can be either precise or imprecise and either synchronous or asynchronous. Asynchronous exceptions occur when an external event interrupts the processor. Synchronous exceptions are caused by the actions of an instruction. During an exception SRR0 is used to calculate where instruction processing should resume. All instructions prior to the resume instruction will have completed execution. SRR1 is used to store the machine status.

There are two asynchronous nonmaskable, highest-priority exceptions system reset and machine check. There are two asynchronous maskable low-priority exceptions external interrupt and decremter. Nonmaskable exceptions are never delayed, therefore if two nonmaskable, asynchronous exceptions occur in immediate succession, the state information saved by the first exception may be overwritten when the subsequent exception occurs.

The PowerPC arcitecture defines one imprecise exception, the imprecise floating point enabled exception. All other synchronous exceptions are precise. The synchronization occurring during asynchronous precise exceptions conforms to the requirements for context synchronization.

9.4.2 Vectoring of Interrupt Handler

Upon determining that an exception can be taken the PowerPC automatically performs the following actions:

- an instruction address is loaded into SRR0
- bits 33-36 and 42-47 of SRR1 are loaded with information specific to the exception.
- bits 0-32, 37-41, and 48-63 of SRR1 are loaded with corresponding bits from the MSR.
- the MSR is set based upon the exception type.
- instruction fetch and execution resumes, using the new MSR value, at a location specific to the exception type.

If the interrupt handler was installed as an RTEMS interrupt handler, then upon receipt of the interrupt, the processor passes control to the RTEMS interrupt handler which performs the following actions:

- saves the state of the interrupted task on it's stack,
- saves all registers which are not normally preserved by the calling sequence so the user's interrupt service routine can be written in a high-level language.
- if this is the outermost (i.e. non-nested) interrupt, then the RTEMS interrupt handler switches from the current stack to the interrupt stack,
- enables exceptions,
- invokes the vectors to a user interrupt service routine (ISR).

Asynchronous interrupts are ignored while exceptions are disabled. Synchronous interrupts which occur while are disabled result in the CPU being forced into an error mode.

A nested interrupt is processed similarly with the exception that the current stack need not be switched to the interrupt stack.

9.4.3 Interrupt Levels

The PowerPC architecture supports only a single external asynchronous interrupt source. This interrupt source may be enabled and disabled via the External Interrupt Enable (EE) bit in the Machine State Register (MSR). Thus only two level (enabled and disabled) of external device interrupt priorities are directly supported by the PowerPC architecture.

Some PowerPC implementations include a Critical Interrupt capability which is often used to receive interrupts from high priority external devices.

The RTEMS interrupt level mapping scheme for the PowerPC is not a numeric level as on most RTEMS ports. It is a bit mapping in which the least three significant bits of the interrupt level are mapped directly to the enabling of specific interrupt sources as follows:

Critical Interrupt	Setting bit 0 (the least significant bit) of the interrupt level enables the Critical Interrupt source, if it is available on this CPU model.
Machine Check	Setting bit 1 of the interrupt level enables Machine Check exceptions.
External Interrupt	Setting bit 2 of the interrupt level enables External Interrupt exceptions.

All other bits in the RTEMS task interrupt level are ignored.

9.5 Default Fatal Error Processing

The default fatal error handler for this architecture performs the following actions:

- places the error code in r3, and
- executes a trap instruction which results in a Program Exception.

If the Program Exception returns, then the following actions are performed:

- disables all processor exceptions by loading a 0 into the MSR, and
- goes into an infinite loop to simulate a halt processor instruction.

9.6 Board Support Packages

9.6.1 System Reset

An RTEMS based application is initiated or re-initiated when the PowerPC processor is reset. The PowerPC architecture defines a Reset Exception, but leaves the details of the CPU state as implementation specific. Please refer to the User's Manual for the CPU model in question.

In general, at power-up the PowerPC begin execution at address 0xFFF00100 in supervisor mode with all exceptions disabled. For soft resets, the CPU will vector to either 0xFFF00100 or 0x00000100 depending upon the setting of the Exception Prefix bit in the MSR. If during a soft reset, a Machine Check Exception occurs, then the CPU may execute a hard reset.

9.6.2 Processor Initialization

If this PowerPC implementation supports on-chip caching and this is to be utilized, then it should be enabled during the reset application initialization code. On-chip caching has been observed to prevent some emulators from working properly, so it may be necessary to run with caching disabled to use these emulators.

In addition to the requirements described in the **Board Support Packages** chapter of the RTEMS C Applications User's Manual for the reset code which is executed before the call to `rtems_initialize_executive`, the PowerPC version has the following specific requirements:

- Must leave the PR bit of the Machine State Register (MSR) set to 0 so the PowerPC remains in the supervisor state.
- Must set stack pointer (sp or r1) such that a minimum stack size of `MINIMUM_STACK_SIZE` bytes is provided for the RTEMS initialization sequence.
- Must disable all external interrupts (i.e. clear the EI (EE) bit of the machine state register).
- Must enable traps so window overflow and underflow conditions can be properly handled.
- Must initialize the PowerPC's initial Exception Table with default handlers.

10 SuperH Specific Information

This chapter discusses the SuperH architecture dependencies in this port of RTEMS. The SuperH family has a wide variety of implementations by a wide range of vendors. Consequently, there are many, many CPU models within it.

Architecture Documents

For information on the SuperH architecture, refer to the following documents available from VENDOR (`'http://www.XXX.com/'`):

- *SuperH Family Reference, VENDOR, PART NUMBER.*

10.1 CPU Model Dependent Features

This chapter presents the set of features which vary across SuperH implementations and are of importance to RTEMS. The set of CPU model feature macros are defined in the file `cpukit/score/cpu/sh/sh.h` based upon the particular CPU model specified on the compilation command line.

10.1.1 Another Optional Feature

The macro XXX

10.2 Calling Conventions

10.2.1 Calling Mechanism

All RTEMS directives are invoked using a XXX instruction and return to the user application via the XXX instruction.

10.2.2 Register Usage

The SH1 has 16 general registers (r0..r15).

- r0..r3 used as general volatile registers
- r4..r7 used to pass up to 4 arguments to functions, arguments above 4 are passed via the stack)
- r8..13 caller saved registers (i.e. push them to the stack if you need them inside of a function)
- r14 frame pointer
- r15 stack pointer

10.2.3 Parameter Passing

XXX

10.3 Memory Model

10.3.1 Flat Memory Model

The SuperH family supports a flat 32-bit address space with addresses ranging from 0x00000000 to 0xFFFFFFFF (4 gigabytes). Each address is represented by a 32-bit value and is byte addressable. The address may be used to reference a single byte, word (2-bytes), or long word (4 bytes). Memory accesses within this address space are performed in big endian fashion by the processors in this family.

Some of the SuperH family members support virtual memory and segmentation. RTEMS does not support virtual memory or segmentation on any of the SuperH family members. It is the responsibility of the BSP to initialize the mapping for a flat memory model.

10.4 Interrupt Processing

Although RTEMS hides many of the processor dependent details of interrupt processing, it is important to understand how the RTEMS interrupt manager is mapped onto the processor's unique architecture. Discussed in this chapter are the MIPS's interrupt response and control mechanisms as they pertain to RTEMS.

10.4.1 Vectoring of an Interrupt Handler

Upon receipt of an interrupt the XXX family members with separate interrupt stacks automatically perform the following actions:

- TBD

A nested interrupt is processed similarly by these CPU models with the exception that only a single ISF is placed on the interrupt stack and the current stack need not be switched.

10.4.2 Interrupt Levels

TBD

10.5 Default Fatal Error Processing

The default fatal error handler for this architecture disables processor interrupts, places the error code in **XXX**, and executes a **XXX** instruction to simulate a halt processor instruction.

10.6 Board Support Packages

10.6.1 System Reset

An RTEMS based application is initiated or re-initiated when the processor is reset. When the processor is reset, it performs the following actions:

- TBD

10.6.2 Processor Initialization

TBD

11 SPARC Specific Information

The Real Time Executive for Multiprocessor Systems (RTEMS) is designed to be portable across multiple processor architectures. However, the nature of real-time systems makes it essential that the application designer understand certain processor dependent implementation details. These processor dependencies include calling convention, board support package issues, interrupt processing, exact RTEMS memory requirements, performance data, header files, and the assembly language interface to the executive.

This document discusses the SPARC architecture dependencies in this port of RTEMS. Currently, only implementations of SPARC Version 7 are supported by RTEMS.

It is highly recommended that the SPARC RTEMS application developer obtain and become familiar with the documentation for the processor being used as well as the specification for the revision of the SPARC architecture which corresponds to that processor.

SPARC Architecture Documents

For information on the SPARC architecture, refer to the following documents available from SPARC International, Inc. (<http://www.sparc.com>):

- SPARC Standard Version 7.
- SPARC Standard Version 8.
- SPARC Standard Version 9.

ERC32 Specific Information

The European Space Agency's ERC32 is a three chip computing core implementing a SPARC V7 processor and associated support circuitry for embedded space applications. The integer and floating-point units (90C601E & 90C602E) are based on the Cypress 7C601 and 7C602, with additional error-detection and recovery functions. The memory controller (MEC) implements system support functions such as address decoding, memory interface, DMA interface, UARTs, timers, interrupt control, write-protection, memory reconfiguration and error-detection. The core is designed to work at 25MHz, but using space qualified memories limits the system frequency to around 15 MHz, resulting in a performance of 10 MIPS and 2 MFLOPS.

Information on the ERC32 and a number of development support tools, such as the SPARC Instruction Simulator (SIS), are freely available on the Internet. The following documents and SIS are available via anonymous ftp or pointing your web browser at <ftp://ftp.estec.esa.nl/pub/ws/wsd/erc32>.

- ERC32 System Design Document
- MEC Device Specification

Additionally, the SPARC RISC User's Guide from Matra MHS documents the functionality of the integer and floating point units including the instruction set information. To obtain this document as well as ERC32 components and VHDL models contact:

Matra MHS SA
3 Avenue du Centre, BP 309,

78054 St-Quentin-en-Yvelines,
Cedex, France
VOICE: +31-1-30607087
FAX: +31-1-30640693

Amar Guennon (amar.guennon@matramhs.fr) is familiar with the ERC32.

11.1 CPU Model Dependent Features

Microprocessors are generally classified into families with a variety of CPU models or implementations within that family. Within a processor family, there is a high level of binary compatibility. This family may be based on either an architectural specification or on maintaining compatibility with a popular processor. Recent microprocessor families such as the SPARC or PowerPC are based on an architectural specification which is independent of any particular CPU model or implementation. Older families such as the M68xxx and the iX86 evolved as the manufacturer strived to produce higher performance processor models which maintained binary compatibility with older models.

RTEMS takes advantage of the similarity of the various models within a CPU family. Although the models do vary in significant ways, the high level of compatibility makes it possible to share the bulk of the CPU dependent executive code across the entire family.

11.1.1 CPU Model Feature Flags

Each processor family supported by RTEMS has a list of features which vary between CPU models within a family. For example, the most common model dependent feature regardless of CPU family is the presence or absence of a floating point unit or coprocessor. When defining the list of features present on a particular CPU model, one simply notes that floating point hardware is or is not present and defines a single constant appropriately. Conditional compilation is utilized to include the appropriate source code for this CPU model's feature set. It is important to note that this means that RTEMS is thus compiled using the appropriate feature set and compilation flags optimal for this CPU model used. The alternative would be to generate a binary which would execute on all family members using only the features which were always present.

This section presents the set of features which vary across SPARC implementations and are of importance to RTEMS. The set of CPU model feature macros are defined in the file `cpukit/score/cpu/sparc/sparc.h` based upon the particular CPU model defined on the compilation command line.

11.1.1.1 CPU Model Name

The macro `CPU_MODEL_NAME` is a string which designates the name of this CPU model. For example, for the European Space Agency's ERC32 SPARC model, this macro is set to the string "erc32".

11.1.1.2 Floating Point Unit

The macro `SPARC_HAS_FPU` is set to 1 to indicate that this CPU model has a hardware floating point unit and 0 otherwise.

11.1.1.3 Bitscan Instruction

The macro `SPARC_HAS_BITSCAN` is set to 1 to indicate that this CPU model has the bitscan instruction. For example, this instruction is supported by the Fujitsu SPARC lite family.

11.1.1.4 Number of Register Windows

The macro `SPARC_NUMBER_OF_REGISTER_WINDOWS` is set to indicate the number of register window sets implemented by this CPU model. The SPARC architecture allows for a maximum of thirty-two register window sets although most implementations only include eight.

11.1.1.5 Low Power Mode

The macro `SPARC_HAS_LOW_POWER_MODE` is set to one to indicate that this CPU model has a low power mode. If low power is enabled, then there must be CPU model specific implementation of the IDLE task in `cpukit/score/cpu/sparc/cpu.c`. The low power mode IDLE task should be of the form:

```
while ( TRUE ) {
    enter low power mode
}
```

The code required to enter low power mode is CPU model specific.

11.1.2 CPU Model Implementation Notes

The ERC32 is a custom SPARC V7 implementation based on the Cypress 601/602 chipset. This CPU has a number of on-board peripherals and was developed by the European Space Agency to target space applications. RTEMS currently provides support for the following peripherals:

- UART Channels A and B
- General Purpose Timer
- Real Time Clock
- Watchdog Timer (so it can be disabled)
- Control Register (so powerdown mode can be enabled)
- Memory Control Register
- Interrupt Control

The General Purpose Timer and Real Time Clock Timer provided with the ERC32 share the Timer Control Register. Because the Timer Control Register is write only, we must mirror it in software and insure that writes to one timer do not alter the current settings and status of the other timer. Routines are provided in `erc32.h` which promote the view that the two timers are completely independent. By exclusively using these routines to access the Timer Control Register, the application can view the system as having a General Purpose Timer Control Register and a Real Time Clock Timer Control Register rather than the single shared value.

The RTEMS Idle thread take advantage of the low power mode provided by the ERC32. Low power mode is entered during idle loops and is enabled at initialization time.

11.2 Calling Conventions

Each high-level language compiler generates subroutine entry and exit code based upon a set of rules known as the compiler's calling convention. These rules address the following issues:

- register preservation and usage
- parameter passing
- call and return mechanism

A compiler's calling convention is of importance when interfacing to subroutines written in another language either assembly or high-level. Even when the high-level language and target processor are the same, different compilers may use different calling conventions. As a result, calling conventions are both processor and compiler dependent.

11.2.1 Programming Model

This section discusses the programming model for the SPARC architecture.

11.2.1.1 Non-Floating Point Registers

The SPARC architecture defines thirty-two non-floating point registers directly visible to the programmer. These are divided into four sets:

- input registers
- local registers
- output registers
- global registers

Each register is referred to by either two or three names in the SPARC reference manuals. First, the registers are referred to as r0 through r31 or with the alternate notation r[0] through r[31]. Second, each register is a member of one of the four sets listed above. Finally, some registers have an architecturally defined role in the programming model which provides an alternate name. The following table describes the mapping between the 32 registers and the register sets:

Register Number	Register Names	Description
0 - 7	g0 - g7	Global Registers
8 - 15	o0 - o7	Output Registers
16 - 23	l0 - l7	Local Registers
24 - 31	i0 - i7	Input Registers

As mentioned above, some of the registers serve defined roles in the programming model. The following table describes the role of each of these registers:

Register Name	Alternate Names	Description
g0	NA	reads return 0; writes are ignored
o6	sp	stack pointer
i6	fp	frame pointer
i7	NA	return address

11.2.1.2 Floating Point Registers

The SPARC V7 architecture includes thirty-two, thirty-two bit registers. These registers may be viewed as follows:

- 32 single precision floating point or integer registers (f0, f1, ... f31)
- 16 double precision floating point registers (f0, f2, f4, ... f30)
- 8 extended precision floating point registers (f0, f4, f8, ... f28)

The floating point status register (fpsr) specifies the behavior of the floating point unit for rounding, contains its condition codes, version specification, and trap information.

A queue of the floating point instructions which have started execution but not yet completed is maintained. This queue is needed to support the multiple cycle nature of floating point operations and to aid floating point exception trap handlers. Once a floating point exception has been encountered, the queue is frozen until it is emptied by the trap handler. The floating point queue is loaded by launching instructions. It is emptied normally when the floating point completes all outstanding instructions and by floating point exception handlers with the store double floating point queue (stdfq) instruction.

11.2.1.3 Special Registers

The SPARC architecture includes two special registers which are critical to the programming model: the Processor State Register (psr) and the Window Invalid Mask (wim). The psr contains the condition codes, processor interrupt level, trap enable bit, supervisor mode and previous supervisor mode bits, version information, floating point unit and coprocessor enable bits, and the current window pointer (cwp). The cwp field of the psr and wim register are used to manage the register windows in the SPARC architecture. The register windows are discussed in more detail below.

11.2.2 Register Windows

The SPARC architecture includes the concept of register windows. An overly simplistic way to think of these windows is to imagine them as being an infinite supply of "fresh" register sets available for each subroutine to use. In reality, they are much more complicated.

The save instruction is used to obtain a new register window. This instruction decrements the current window pointer, thus providing a new set of registers for use. This register set includes eight fresh local registers for use exclusively by this subroutine. When done with a register set, the restore instruction increments the current window pointer and the previous register set is once again available.

The two primary issues complicating the use of register windows are that (1) the set of register windows is finite, and (2) some registers are shared between adjacent registers windows.

Because the set of register windows is finite, it is possible to execute enough save instructions without corresponding restore's to consume all of the register windows. This is easily accomplished in a high level language because each subroutine typically performs a save instruction upon entry. Thus having a subroutine call depth greater than the number of register windows will result in a window overflow condition. The window overflow condition generates a trap which must be handled in software. The window overflow trap handler is responsible for saving the contents of the oldest register window on the program stack.

Similarly, the subroutines will eventually complete and begin to perform restore's. If the restore results in the need for a register window which has previously been written to memory as part of an overflow, then a window underflow condition results. Just like the window overflow, the window underflow condition must be handled in software by a trap handler. The window underflow trap handler is responsible for reloading the contents of the register window requested by the restore instruction from the program stack.

The Window Invalid Mask (wim) and the Current Window Pointer (cwp) field in the psr are used in conjunction to manage the finite set of register windows and detect the window overflow and underflow conditions. The cwp contains the index of the register window currently in use. The save instruction decrements the cwp modulo the number of register windows. Similarly, the restore instruction increments the cwp modulo the number of register windows. Each bit in the wim represents whether a register window contains valid information. The value of 0 indicates the register window is valid and 1 indicates it is invalid. When a save instruction causes the cwp to point to a register window which is marked as invalid, a window overflow condition results. Conversely, the restore instruction may result in a window underflow condition.

Other than the assumption that a register window is always available for trap (i.e. interrupt) handlers, the SPARC architecture places no limits on the number of register windows simultaneously marked as invalid (i.e. number of bits set in the wim). However, RTEMS assumes that only one register window is marked invalid at a time (i.e. only one bit set in the wim). This makes the maximum possible number of register windows available to the user while still meeting the requirement that window overflow and underflow conditions can be detected.

The window overflow and window underflow trap handlers are a critical part of the run-time environment for a SPARC application. The SPARC architectural specification allows for the number of register windows to be any power of two less than or equal to 32. The most common choice for SPARC implementations appears to be 8 register windows. This results in the cwp ranging in value from 0 to 7 on most implementations.

The second complicating factor is the sharing of registers between adjacent register windows. While each register window has its own set of local registers, the input and output registers are shared between adjacent windows. The output registers for register window N are the same as the input registers for register window $((N - 1) \text{ modulo } RW)$ where RW is the number of register windows. An alternative way to think of this is to remember how parameters are passed to a subroutine on the SPARC. The caller loads values into what are its output registers. Then after the callee executes a save instruction, those parameters are available in its input registers. This is a very efficient way to pass parameters as no data is actually moved by the save or restore instructions.

11.2.3 Call and Return Mechanism

The SPARC architecture supports a simple yet effective call and return mechanism. A subroutine is invoked via the call (call) instruction. This instruction places the return address in the caller's output register 7 (o7). After the callee executes a save instruction, this value is available in input register 7 (i7) until the corresponding restore instruction is executed.

The callee returns to the caller via a jmp to the return address. There is a delay slot following this instruction which is commonly used to execute a restore instruction – if a register window was allocated by this subroutine.

It is important to note that the SPARC subroutine call and return mechanism does not automatically save and restore any registers. This is accomplished via the save and restore instructions which manage the set of registers windows.

11.2.4 Calling Mechanism

All RTEMS directives are invoked using the regular SPARC calling convention via the call instruction.

11.2.5 Register Usage

As discussed above, the call instruction does not automatically save any registers. The save and restore instructions are used to allocate and deallocate register windows. When a register window is allocated, the new set of local registers are available for the exclusive use of the subroutine which allocated this register set.

11.2.6 Parameter Passing

RTEMS assumes that arguments are placed in the caller's output registers with the first argument in output register 0 (o0), the second argument in output register 1 (o1), and so forth. Until the callee executes a save instruction, the parameters are still visible in the output registers. After the callee executes a save instruction, the parameters are visible in the corresponding input registers. The following pseudo-code illustrates the typical sequence used to call a RTEMS directive with three (3) arguments:

```
load third argument into o2
load second argument into o1
load first argument into o0
invoke directive
```

11.2.7 User-Provided Routines

All user-provided routines invoked by RTEMS, such as user extensions, device drivers, and MPCPI routines, must also adhere to these calling conventions.

11.3 Memory Model

A processor may support any combination of memory models ranging from pure physical addressing to complex demand paged virtual memory systems. RTEMS supports a flat memory model which ranges contiguously over the processor's allowable address space. RTEMS does not support segmentation or virtual memory of any kind. The appropriate

memory model for RTEMS provided by the targeted processor and related characteristics of that model are described in this chapter.

11.3.1 Flat Memory Model

The SPARC architecture supports a flat 32-bit address space with addresses ranging from 0x00000000 to 0xFFFFFFFF (4 gigabytes). Each address is represented by a 32-bit value and is byte addressable. The address may be used to reference a single byte, half-word (2-bytes), word (4 bytes), or doubleword (8 bytes). Memory accesses within this address space are performed in big endian fashion by the SPARC. Memory accesses which are not properly aligned generate a "memory address not aligned" trap (type number 7). The following table lists the alignment requirements for a variety of data accesses:

Data Type	Alignment Requirement
byte	1
half-word	2
word	4
doubleword	8

Doubleword load and store operations must use a pair of registers as their source or destination. This pair of registers must be an adjacent pair of registers with the first of the pair being even numbered. For example, a valid destination for a doubleword load might be input registers 0 and 1 (i0 and i1). The pair i1 and i2 would be invalid. [NOTE: Some assemblers for the SPARC do not generate an error if an odd numbered register is specified as the beginning register of the pair. In this case, the assembler assumes that what the programmer meant was to use the even-odd pair which ends at the specified register. This may or may not have been a correct assumption.]

RTEMS does not support any SPARC Memory Management Units, therefore, virtual memory or segmentation systems involving the SPARC are not supported.

11.4 Interrupt Processing

Different types of processors respond to the occurrence of an interrupt in its own unique fashion. In addition, each processor type provides a control mechanism to allow for the proper handling of an interrupt. The processor dependent response to the interrupt modifies the current execution state and results in a change in the execution stream. Most processors require that an interrupt handler utilize some special control mechanisms to return to the normal processing stream. Although RTEMS hides many of the processor dependent details of interrupt processing, it is important to understand how the RTEMS interrupt manager is mapped onto the processor's unique architecture. Discussed in this chapter are the SPARC's interrupt response and control mechanisms as they pertain to RTEMS.

RTEMS and associated documentation uses the terms interrupt and vector. In the SPARC architecture, these terms correspond to traps and trap type, respectively. The terms will be used interchangeably in this manual.

11.4.1 Synchronous Versus Asynchronous Traps

The SPARC architecture includes two classes of traps: synchronous and asynchronous. Asynchronous traps occur when an external event interrupts the processor. These traps are not associated with any instruction executed by the processor and logically occur between instructions. The instruction currently in the execute stage of the processor is allowed to complete although subsequent instructions are annulled. The return address reported by the processor for asynchronous traps is the pair of instructions following the current instruction.

Synchronous traps are caused by the actions of an instruction. The trap stimulus in this case either occurs internally to the processor or is from an external signal that was provoked by the instruction. These traps are taken immediately and the instruction that caused the trap is aborted before any state changes occur in the processor itself. The return address reported by the processor for synchronous traps is the instruction which caused the trap and the following instruction.

11.4.2 Vectoring of Interrupt Handler

Upon receipt of an interrupt the SPARC automatically performs the following actions:

- disables traps (sets the ET bit of the psr to 0),
- the S bit of the psr is copied into the Previous Supervisor Mode (PS) bit of the psr,
- the cwp is decremented by one (modulo the number of register windows) to activate a trap window,
- the PC and nPC are loaded into local register 1 and 2 (l0 and l1),
- the trap type (tt) field of the Trap Base Register (TBR) is set to the appropriate value, and
- if the trap is not a reset, then the PC is written with the contents of the TBR and the nPC is written with TBR + 4. If the trap is a reset, then the PC is set to zero and the nPC is set to 4.

Trap processing on the SPARC has two features which are noticeably different than interrupt processing on other architectures. First, the value of psr register in effect immediately before the trap occurred is not explicitly saved. Instead only reversible alterations are made to it. Second, the Processor Interrupt Level (pil) is not set to correspond to that of the interrupt being processed. When a trap occurs, ALL subsequent traps are disabled. In order to safely invoke a subroutine during trap handling, traps must be enabled to allow for the possibility of register window overflow and underflow traps.

If the interrupt handler was installed as an RTEMS interrupt handler, then upon receipt of the interrupt, the processor passes control to the RTEMS interrupt handler which performs the following actions:

- saves the state of the interrupted task on its stack,
- insures that a register window is available for subsequent traps,
- if this is the outermost (i.e. non-nested) interrupt, then the RTEMS interrupt handler switches from the current stack to the interrupt stack,
- enables traps,

- invokes the vectors to a user interrupt service routine (ISR).

Asynchronous interrupts are ignored while traps are disabled. Synchronous traps which occur while traps are disabled result in the CPU being forced into an error mode.

A nested interrupt is processed similarly with the exception that the current stack need not be switched to the interrupt stack.

11.4.3 Traps and Register Windows

One of the register windows must be reserved at all times for trap processing. This is critical to the proper operation of the trap mechanism in the SPARC architecture. It is the responsibility of the trap handler to insure that there is a register window available for a subsequent trap before re-enabling traps. It is likely that any high level language routines invoked by the trap handler (such as a user-provided RTEMS interrupt handler) will allocate a new register window. The save operation could result in a window overflow trap. This trap cannot be correctly processed unless (1) traps are enabled and (2) a register window is reserved for traps. Thus, the RTEMS interrupt handler insures that a register window is available for subsequent traps before enabling traps and invoking the user's interrupt handler.

11.4.4 Interrupt Levels

Sixteen levels (0-15) of interrupt priorities are supported by the SPARC architecture with level fifteen (15) being the highest priority. Level zero (0) indicates that interrupts are fully enabled. Interrupt requests for interrupts with priorities less than or equal to the current interrupt mask level are ignored.

Although RTEMS supports 256 interrupt levels, the SPARC only supports sixteen. RTEMS interrupt levels 0 through 15 directly correspond to SPARC processor interrupt levels. All other RTEMS interrupt levels are undefined and their behavior is unpredictable.

11.4.5 Disabling of Interrupts by RTEMS

During the execution of directive calls, critical sections of code may be executed. When these sections are encountered, RTEMS disables interrupts to level seven (15) before the execution of this section and restores them to the previous level upon completion of the section. RTEMS has been optimized to insure that interrupts are disabled for less than RTEMS_MAXIMUM_DISABLE_PERIOD microseconds on a RTEMS_MAXIMUM_DISABLE_PERIOD_MHZ Mhz ERC32 with zero wait states. These numbers will vary based the number of wait states and processor speed present on the target board. [NOTE: The maximum period with interrupts disabled is hand calculated. This calculation was last performed for Release RTEMS_RELEASE_FOR_MAXIMUM_DISABLE_PERIOD.]

[NOTE: It is thought that the length of time at which the processor interrupt level is elevated to fifteen by RTEMS is not anywhere near as long as the length of time ALL traps are disabled as part of the "flush all register windows" operation.]

Non-maskable interrupts (NMI) cannot be disabled, and ISRs which execute at this level MUST NEVER issue RTEMS system calls. If a directive is invoked, unpredictable results

may occur due to the inability of RTEMS to protect its critical sections. However, ISRs that make no system calls may safely execute as non-maskable interrupts.

11.4.6 Interrupt Stack

The SPARC architecture does not provide for a dedicated interrupt stack. Thus by default, trap handlers would execute on the stack of the RTEMS task which they interrupted. This artificially inflates the stack requirements for each task since EVERY task stack would have to include enough space to account for the worst case interrupt stack requirements in addition to its own worst case usage. RTEMS addresses this problem on the SPARC by providing a dedicated interrupt stack managed by software.

During system initialization, RTEMS allocates the interrupt stack from the Workspace Area. The amount of memory allocated for the interrupt stack is determined by the `interrupt_stack_size` field in the CPU Configuration Table. As part of processing a non-nested interrupt, RTEMS will switch to the interrupt stack before invoking the installed handler.

11.5 Default Fatal Error Processing

Upon detection of a fatal error by either the application or RTEMS the fatal error manager is invoked. The fatal error manager will invoke the user-supplied fatal error handlers. If no user-supplied handlers are configured, the RTEMS provided default fatal error handler is invoked. If the user-supplied fatal error handlers return to the executive the default fatal error handler is then invoked. This chapter describes the precise operations of the default fatal error handler.

11.5.1 Default Fatal Error Handler Operations

The default fatal error handler which is invoked by the `fatal_error_occurred` directive when there is no user handler configured or the user handler returns control to RTEMS. The default fatal error handler disables processor interrupts to level 15, places the error code in `g1`, and goes into an infinite loop to simulate a halt processor instruction.

11.6 Board Support Packages

An RTEMS Board Support Package (BSP) must be designed to support a particular processor and target board combination. This chapter presents a discussion of SPARC specific BSP issues. For more information on developing a BSP, refer to the chapter titled Board Support Packages in the RTEMS Applications User's Guide.

11.6.1 System Reset

An RTEMS based application is initiated or re-initiated when the SPARC processor is reset. When the SPARC is reset, the processor performs the following actions:

- the enable trap (ET) of the psr is set to 0 to disable traps,
- the supervisor bit (S) of the psr is set to 1 to enter supervisor mode, and
- the PC is set 0 and the nPC is set to 4.

The processor then begins to execute the code at location 0. It is important to note that all fields in the psr are not explicitly set by the above steps and all other registers retain

their value from the previous execution mode. This is true even of the Trap Base Register (TBR) whose contents reflect the last trap which occurred before the reset.

11.6.2 Processor Initialization

It is the responsibility of the application's initialization code to initialize the TBR and install trap handlers for at least the register window overflow and register window underflow conditions. Traps should be enabled before invoking any subroutines to allow for register window management. However, interrupts should be disabled by setting the Processor Interrupt Level (pil) field of the psr to 15. RTEMS installs its own Trap Table as part of initialization which is initialized with the contents of the Trap Table in place when the `rtems_initialize_executive` directive was invoked. Upon completion of executive initialization, interrupts are enabled.

If this SPARC implementation supports on-chip caching and this is to be utilized, then it should be enabled during the reset application initialization code.

In addition to the requirements described in the Board Support Packages chapter of the C Applications Users Manual for the reset code which is executed before the call to `rtems_initialize_executive`, the SPARC version has the following specific requirements:

- Must leave the S bit of the status register set so that the SPARC remains in the supervisor state.
- Must set stack pointer (sp) such that a minimum stack size of `MINIMUM_STACK_SIZE` bytes is provided for the `rtems_initialize_executive` directive.
- Must disable all external interrupts (i.e. set the pil to 15).
- Must enable traps so window overflow and underflow conditions can be properly handled.
- Must initialize the SPARC's initial trap table with at least trap handlers for register window overflow and register window underflow.

12 SPARC-64 Specific Information

This document discusses the SPARC Version 9 (aka SPARC-64, SPARC64 or SPARC V9) architecture dependencies in this port of RTEMS.

The SPARC V9 architecture leaves a lot of undefined implementation dependencies which are defined by the processor models. Consult the specific CPU model section in this document for additional documents covering the implementation dependent architectural features.

sun4u Specific Information

sun4u is the subset of the SPARC V9 implementations comprising the UltraSPARC I through UltraSPARC IV processors.

The following documents were used in developing the SPARC-64 sun4u port:

- UltraSPARC Users Manual (<http://www.sun.com/microelectronics/manuals/ultrasparc/802-7220-02.pdf>)
- UltraSPARC IIIi Processor (datasheets.chipdb.org/Sun/UltraSparc-IIIi.pdf)

sun4v Specific Information

sun4v is the subset of the SPARC V9 implementations comprising the UltraSPARC T1 or T2 processors.

The following documents were used in developing the SPARC-64 sun4v port:

- UltraSPARC Architecture 2005 Specification (<http://opensparc-t1.sunsource.net/specs/UA2005-current-draft-P-EXT.pdf>)
- UltraSPARC T1 supplement to UltraSPARC Architecture 2005 Specification (<http://opensparc-t1.sunsource.net/specs/UST1-UASuppl-current-draft-P-EXT.pdf>)

The defining feature that separates the sun4v architecture from its predecessor is the existence of a super-privileged hypervisor that is responsible for providing virtualized execution environments. The impact of the hypervisor on the real-time guarantees available with sun4v has not yet been determined.

12.1 CPU Model Dependent Features

12.1.1 CPU Model Feature Flags

This section presents the set of features which vary across SPARC-64 implementations and are of importance to RTEMS. The set of CPU model feature macros are defined in the file `cpukit/score/cpu/sparc64/sparc64.h` based upon the particular CPU model defined on the compilation command line.

12.1.1.1 CPU Model Name

The macro `CPU_MODEL_NAME` is a string which designates the name of this CPU model. For example, for the UltraSPARC T1 SPARC V9 model, this macro is set to the string `"sun4v"`.

12.1.1.2 Floating Point Unit

The macro `SPARC_HAS_FPU` is set to 1 to indicate that this CPU model has a hardware floating point unit and 0 otherwise.

12.1.1.3 Number of Register Windows

The macro `SPARC_NUMBER_OF_REGISTER_WINDOWS` is set to indicate the number of register window sets implemented by this CPU model. The SPARC architecture allows for a maximum of thirty-two register window sets although most implementations only include eight.

12.1.2 CPU Model Implementation Notes

This section describes the implementation dependencies of the CPU Models `sun4u` and `sun4v` of the SPARC V9 architecture.

12.1.2.1 sun4u Notes

XXX

12.1.3 sun4v Notes

XXX

12.2 Calling Conventions

Each high-level language compiler generates subroutine entry and exit code based upon a set of rules known as the compiler's calling convention. These rules address the following issues:

- register preservation and usage
- parameter passing
- call and return mechanism

A compiler's calling convention is of importance when interfacing to subroutines written in another language either assembly or high-level. Even when the high-level language and target processor are the same, different compilers may use different calling conventions. As a result, calling conventions are both processor and compiler dependent.

The following document also provides some conventions on the global register usage in SPARC V9: <http://developers.sun.com/solaris/articles/sparcv9abi.html>

12.2.1 Programming Model

This section discusses the programming model for the SPARC architecture.

12.2.1.1 Non-Floating Point Registers

The SPARC architecture defines thirty-two non-floating point registers directly visible to the programmer. These are divided into four sets:

- input registers
- local registers

- output registers
- global registers

Each register is referred to by either two or three names in the SPARC reference manuals. First, the registers are referred to as r0 through r31 or with the alternate notation r[0] through r[31]. Second, each register is a member of one of the four sets listed above. Finally, some registers have an architecturally defined role in the programming model which provides an alternate name. The following table describes the mapping between the 32 registers and the register sets:

Register Number	Register Names	Description
0 - 7	g0 - g7	Global Registers
8 - 15	o0 - o7	Output Registers
16 - 23	l0 - l7	Local Registers
24 - 31	i0 - i7	Input Registers

As mentioned above, some of the registers serve defined roles in the programming model. The following table describes the role of each of these registers:

Register Name	Alternate Names	Description
g0	NA	reads return 0; writes are ignored
o6	sp	stack pointer
i6	fp	frame pointer
i7	NA	return address

12.2.1.2 Floating Point Registers

The SPARC V9 architecture includes sixty-four, thirty-two bit registers. These registers may be viewed as follows:

- 32 32-bit single precision floating point or integer registers (f0, f1, ... f31)
- 32 64-bit double precision floating point registers (f0, f2, f4, ... f62)
- 16 128-bit extended precision floating point registers (f0, f4, f8, ... f60)

The floating point state register (fsr) specifies the behavior of the floating point unit for rounding, contains its condition codes, version specification, and trap information.

12.2.1.3 Special Registers

The SPARC architecture includes a number of special registers:

Ancillary State Registers (ASRs)

The ancillary state registers (ASRs) are optional state registers that may be privileged or nonprivileged. ASRs 16-31 are implementation-dependent. The SPARC V9 ASRs include: y, ccr, asi, tick, pc, fprs. The sun4u ASRs include: pcr, pic, dcr, gsr, softint set, softint clr, softint, and tick cmpr. The sun4v ASRs include: pcr, pic, gsr, softint set, softint clr, softint, tick cmpr, stick, and stick cmpr.

Processor State Register (pstate)

The privileged pstate register contains control fields for the processors current state. Its flag fields include the interrupt enable, privileged mode, and enable FPU.

Processor Interrupt Level (pil)

The PIL specifies the interrupt level above which interrupts will be accepted.

Trap Registers

The trap handling mechanism of the SPARC V9 includes a number of registers, including: trap program counter (tpc), trap next pc (tnpc), trap state (tstate), trap type (tt), trap base address (tba), and trap level (tl).

Alternate Globals

The AG bit of the pstate register provides access to an alternate set of global registers. On sun4v, the AG bit is replaced by the global level (gl) register, providing access to at least two and at most eight alternate sets of globals.

Register Window registers

A number of registers assist in register window management. These include the current window pointer (cwp), savable windows (cansave), restorable windows (canrestore), clean windows (cleanwin), other windows (otherwin), and window state (wstate).

12.2.2 Register Windows

The SPARC architecture includes the concept of register windows. An overly simplistic way to think of these windows is to imagine them as being an infinite supply of "fresh" register sets available for each subroutine to use. In reality, they are much more complicated.

The save instruction is used to obtain a new register window. This instruction increments the current window pointer, thus providing a new set of registers for use. This register set includes eight fresh local registers for use exclusively by this subroutine. When done with a register set, the restore instruction decrements the current window pointer and the previous register set is once again available.

The two primary issues complicating the use of register windows are that (1) the set of register windows is finite, and (2) some registers are shared between adjacent registers windows.

Because the set of register windows is finite, it is possible to execute enough save instructions without corresponding restore's to consume all of the register windows. This is easily accomplished in a high level language because each subroutine typically performs a save instruction upon entry. Thus having a subroutine call depth greater than the number of register windows will result in a window overflow condition. The window overflow condition generates a trap which must be handled in software. The window overflow trap handler is responsible for saving the contents of the oldest register window on the program stack.

Similarly, the subroutines will eventually complete and begin to perform restore's. If the restore results in the need for a register window which has previously been written to memory as part of an overflow, then a window underflow condition results. Just like the

window overflow, the window underflow condition must be handled in software by a trap handler. The window underflow trap handler is responsible for reloading the contents of the register window requested by the restore instruction from the program stack.

The `cansave`, `canrestore`, `otherwin`, and `cwp` are used in conjunction to manage the finite set of register windows and detect the window overflow and underflow conditions. The first three of these registers must satisfy the invariant $\text{cansave} + \text{canrestore} + \text{otherwin} = \text{nwindow} - 2$, where `nwindow` is the number of register windows. The `cwp` contains the index of the register window currently in use. RTEMS does not use the `cleanwin` and `otherwin` registers.

The `save` instruction increments the `cwp` modulo the number of register windows, and if `cansave` is 0 then it also generates a window overflow. Similarly, the `restore` instruction decrements the `cwp` modulo the number of register windows, and if `canrestore` is 0 then it also generates a window underflow.

Unlike with the SPARC model, the SPARC-64 port does not assume that a register window is available for a trap. The window overflow and underflow conditions are not detected without hardware generating the trap. (These conditions can be detected by reading the register window registers and doing some simple arithmetic.)

The window overflow and window underflow trap handlers are a critical part of the run-time environment for a SPARC application. The SPARC architectural specification allows for the number of register windows to be any power of two less than or equal to 32. The most common choice for SPARC implementations appears to be 8 register windows. This results in the `cwp` ranging in value from 0 to 7 on most implementations.

The second complicating factor is the sharing of registers between adjacent register windows. While each register window has its own set of local registers, the input and output registers are shared between adjacent windows. The output registers for register window `N` are the same as the input registers for register window $((N + 1) \text{ modulo } RW)$ where `RW` is the number of register windows. An alternative way to think of this is to remember how parameters are passed to a subroutine on the SPARC. The caller loads values into what are its output registers. Then after the callee executes a `save` instruction, those parameters are available in its input registers. This is a very efficient way to pass parameters as no data is actually moved by the `save` or `restore` instructions.

12.2.3 Call and Return Mechanism

The SPARC architecture supports a simple yet effective call and return mechanism. A subroutine is invoked via the `call` (`call`) instruction. This instruction places the return address in the caller's output register 7 (`o7`). After the callee executes a `save` instruction, this value is available in input register 7 (`i7`) until the corresponding `restore` instruction is executed.

The callee returns to the caller via a `jmp` to the return address. There is a delay slot following this instruction which is commonly used to execute a `restore` instruction – if a register window was allocated by this subroutine.

It is important to note that the SPARC subroutine call and return mechanism does not automatically save and restore any registers. This is accomplished via the `save` and `restore`

instructions which manage the set of registers windows. This allows for the compiler to generate leaf-optimized functions that utilize the callers output registers without using save and restore.

12.2.4 Calling Mechanism

All RTEMS directives are invoked using the regular SPARC calling convention via the call instruction.

12.2.5 Register Usage

As discussed above, the call instruction does not automatically save any registers. The save and restore instructions are used to allocate and deallocate register windows. When a register window is allocated, the new set of local registers are available for the exclusive use of the subroutine which allocated this register set.

12.2.6 Parameter Passing

RTEMS assumes that arguments are placed in the caller's output registers with the first argument in output register 0 (o0), the second argument in output register 1 (o1), and so forth. Until the callee executes a save instruction, the parameters are still visible in the output registers. After the callee executes a save instruction, the parameters are visible in the corresponding input registers. The following pseudo-code illustrates the typical sequence used to call a RTEMS directive with three (3) arguments:

```
load third argument into o2
load second argument into o1
load first argument into o0
invoke directive
```

12.2.7 User-Provided Routines

All user-provided routines invoked by RTEMS, such as user extensions, device drivers, and MPCPI routines, must also adhere to these calling conventions.

12.3 Memory Model

A processor may support any combination of memory models ranging from pure physical addressing to complex demand paged virtual memory systems. RTEMS supports a flat memory model which ranges contiguously over the processor's allowable address space. RTEMS does not support segmentation or virtual memory of any kind. The appropriate memory model for RTEMS provided by the targeted processor and related characteristics of that model are described in this chapter.

12.3.1 Flat Memory Model

The SPARC-64 architecture supports a flat 64-bit address space with addresses ranging from 0x0000000000000000 to 0xFFFFFFFFFFFFFFFF. Each address is represented by a 64-bit value (and an 8-bit address space identifier or ASI) and is byte addressable. The address may be used to reference a single byte, half-word (2-bytes), word (4 bytes), doubleword (8 bytes), or quad-word (16 bytes). Memory accesses within this address space are performed in big endian fashion by the SPARC. Memory accesses which are not properly aligned

generate a "memory address not aligned" trap (type number 0x34). The following table lists the alignment requirements for a variety of data accesses:

Data Type	Alignment Requirement
byte	1
half-word	2
word	4
doubleword	8
quadword	16

RTEMS currently does not support any SPARC Memory Management Units, therefore, virtual memory or segmentation systems involving the SPARC are not supported.

12.4 Interrupt Processing

RTEMS and associated documentation uses the terms interrupt and vector. In the SPARC architecture, these terms correspond to traps and trap type, respectively. The terms will be used interchangeably in this manual. Note that in the SPARC manuals, interrupts are a subset of the traps that are delivered to software interrupt handlers.

12.4.1 Synchronous Versus Asynchronous Traps

The SPARC architecture includes two classes of traps: synchronous (precise) and asynchronous (deferred). Asynchronous traps occur when an external event interrupts the processor. These traps are not associated with any instruction executed by the processor and logically occur between instructions. The instruction currently in the execute stage of the processor is allowed to complete although subsequent instructions are annulled. The return address reported by the processor for asynchronous traps is the pair of instructions following the current instruction.

Synchronous traps are caused by the actions of an instruction. The trap stimulus in this case either occurs internally to the processor or is from an external signal that was provoked by the instruction. These traps are taken immediately and the instruction that caused the trap is aborted before any state changes occur in the processor itself. The return address reported by the processor for synchronous traps is the instruction which caused the trap and the following instruction.

12.4.2 Vectoring of Interrupt Handler

Upon receipt of an interrupt the SPARC automatically performs the following actions:

- The trap level is set. This provides access to a fresh set of privileged trap-state registers used to save the current state, in effect, pushing a frame on the trap stack.

$$TL \leftarrow TL + 1$$
- Existing state is preserved
 - $TSTATE[TL].CCR \leftarrow CCR$
 - $TSTATE[TL].ASI \leftarrow ASI$
 - $TSTATE[TL].PSTATE \leftarrow PSTATE$
 - $TSTATE[TL].CWP \leftarrow CWP$

- $TPC[TL] \leftarrow PC$
- $TNPC[TL] \leftarrow nPC$
- The trap type is preserved. $TT[TL] \leftarrow$ the trap type
- The PSTATE register is updated to a predefined state
 - PSTATE.MM is unchanged
 - $PSTATE.RED \leftarrow 0$
 - $PSTATE.PEF \leftarrow 1$ if FPU is present, 0 otherwise
 - $PSTATE.AM \leftarrow 0$ (address masking is turned off)
 - $PSTATE.PRIV \leftarrow 1$ (the processor enters privileged mode)
 - $PSTATE.IE \leftarrow 0$ (interrupts are disabled)
 - $PSTATE.AG \leftarrow 1$ (global regs are replaced with alternate globals)
 - $PSTATE.CLE \leftarrow PSTATE.TLE$ (set endian mode for traps)
- For a register-window trap only, CWP is set to point to the register window that must be accessed by the trap-handler software, that is:
 - If $TT[TL] = 0x24$ (a clean window trap), then $CWP \leftarrow CWP + 1$.
 - If $(0x80 \leq TT[TL] \leq 0xBF)$ (window spill trap), then $CWP \leftarrow CWP + CANSERVE + 2$.
 - If $(0xC0 \leq TT[TL] \leq 0xFF)$ (window fill trap), then $CWP \leftarrow CWP + 1$.
 - For non-register-window traps, CWP is not changed.
- Control is transferred into the trap table:
 - $PC \leftarrow TBA\langle 63:15 \rangle (TL > 0) \parallel TT[TL] \parallel 0000$
 - $nPC \leftarrow TBA\langle 63:15 \rangle (TL > 0) \parallel TT[TL] \parallel 00100$
 - where $(TL > 0)$ is 0 if $TL = 0$, and 1 if $TL > 0$.

In order to safely invoke a subroutine during trap handling, traps must be enabled to allow for the possibility of register window overflow and underflow traps.

If the interrupt handler was installed as an RTEMS interrupt handler, then upon receipt of the interrupt, the processor passes control to the RTEMS interrupt handler which performs the following actions:

- saves the state of the interrupted task on its stack,
- switches the processor to trap level 0,
- if this is the outermost (i.e. non-nested) interrupt, then the RTEMS interrupt handler switches from the current stack to the interrupt stack,
- enables traps,
- invokes the vectors to a user interrupt service routine (ISR).

Asynchronous interrupts are ignored while traps are disabled. Synchronous traps which occur while traps are disabled may result in the CPU being forced into an error mode.

A nested interrupt is processed similarly with the exception that the current stack need not be switched to the interrupt stack.

12.4.3 Traps and Register Windows

XXX

12.4.4 Interrupt Levels

Sixteen levels (0-15) of interrupt priorities are supported by the SPARC architecture with level fifteen (15) being the highest priority. Level zero (0) indicates that interrupts are fully enabled. Interrupt requests for interrupts with priorities less than or equal to the current interrupt mask level are ignored.

Although RTEMS supports 256 interrupt levels, the SPARC only supports sixteen. RTEMS interrupt levels 0 through 15 directly correspond to SPARC processor interrupt levels. All other RTEMS interrupt levels are undefined and their behavior is unpredictable.

12.4.5 Disabling of Interrupts by RTEMS

XXX

12.4.6 Interrupt Stack

The SPARC architecture does not provide for a dedicated interrupt stack. Thus by default, trap handlers would execute on the stack of the RTEMS task which they interrupted. This artificially inflates the stack requirements for each task since EVERY task stack would have to include enough space to account for the worst case interrupt stack requirements in addition to it's own worst case usage. RTEMS addresses this problem on the SPARC by providing a dedicated interrupt stack managed by software.

During system initialization, RTEMS allocates the interrupt stack from the Workspace Area. The amount of memory allocated for the interrupt stack is determined by the `interrupt_stack_size` field in the CPU Configuration Table. As part of processing a non-nested interrupt, RTEMS will switch to the interrupt stack before invoking the installed handler.

12.5 Default Fatal Error Processing

Upon detection of a fatal error by either the application or RTEMS the fatal error manager is invoked. The fatal error manager will invoke the user-supplied fatal error handlers. If no user-supplied handlers are configured, the RTEMS provided default fatal error handler is invoked. If the user-supplied fatal error handlers return to the executive the default fatal error handler is then invoked. This chapter describes the precise operations of the default fatal error handler.

12.5.1 Default Fatal Error Handler Operations

The default fatal error handler which is invoked by the `fatal_error_occurred` directive when there is no user handler configured or the user handler returns control to RTEMS. The default fatal error handler disables processor interrupts to level 15, places the error code in `g1`, and goes into an infinite loop to simulate a halt processor instruction.

12.6 Board Support Packages

An RTEMS Board Support Package (BSP) must be designed to support a particular processor and target board combination. This chapter presents a discussion of SPARC specific

BSP issues. For more information on developing a BSP, refer to the chapter titled Board Support Packages in the RTEMS Applications User's Guide.

12.6.1 HelenOS and Open Firmware

The provided BSPs make use of some bootstrap and low-level hardware code of the HelenOS operating system. These files can be found in the `shared/helenos` directory of the `sparc64` bsp directory. Consult the sources for more detailed information.

The shared BSP code also uses the Open Firmware interface to re-use firmware code, primarily for console support and default trap handlers.

RTEMS CPU Architecture Supplement

Command and Variable Index

There are currently no Command and Variable Index entries.

Concept Index

There are currently no Concept Index entries.

