RTEMS Intel i960 Applications Supplement

Edition 1, for RTEMS 4.5.1-pre3

30 October 2001

On-Line Applications Research Corporation

On-Line Applications Research Corporation
TEXinfo 1999-04-25.15

COPYRIGHT © 1988 - 2000.
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.

Any inquiries concerning RTEMS, its related support components, or its documentation should be directed to either:

On-Line Applications Research Corporation
4910-L Corporate Drive
Huntsville, AL 35805
VOICE: (256) 722-9985
FAX: (256) 722-0985
EMAIL: rtems@OARcorp.com

Table of Contents

Preface 1

1 CPU Model Dependent Features 3

1.1 Introduction 3
1.2 CPU Model Name 3
1.3 Floating Point Unit 3

2 Calling Conventions 5

2.1 Introduction 5
2.2 Processor Background 5
2.3 Calling Mechanism 5
2.4 Register Usage 6
2.5 Parameter Passing 6
2.6 User-Provided Routines 6
2.7 Leaf Procedures 6

3 Memory Model 7

3.1 Introduction 7
3.2 Flat Memory Model 7

4 Interrupt Processing 9

4.1 Introduction 9
4.2 Vectoring of Interrupt Handler 9
4.3 Interrupt Record 9
4.4 Interrupt Levels 10
4.5 Disabling of Interrupts by RTEMS 10
4.6 Register Cache Flushing 10
4.7 Interrupt Stack 10

5 Default Fatal Error Processing 13

5.1 Introduction 13
5.2 Default Fatal Error Handler Operations 13

6 Board Support Packages 15

6.1 Introduction 15
6.2 System Reset 15
6.3 Processor Initialization 15

7 Processor Dependent Information Table 17

7.1 Introduction 17
7.2 CPU Dependent Information Table 17

8 Memory Requirements 19

8.1 Introduction 19
8.2 Data Space Requirements 19
8.3 Minimum and Maximum Code Space Requirements 19
8.4 RTEMS Code Space Worksheet 19
8.5 RTEMS RAM Workspace Worksheet 21

9 Timing Specification 23

9.1 Introduction 23
9.2 Philosophy 23
9.2.1 Determinancy 23
9.2.2 Interrupt Latency 24
9.2.3 Context Switch Time 25
9.2.4 Directive Times 25
9.3 Methodology 26
9.3.1 Software Platform 26
9.3.2 Hardware Platform 26
9.3.3 What is measured? 26
9.3.4 What is not measured? 27
9.3.5 Terminology 27

10 CVME961 Timing Data 29

10.1 Introduction 29
10.2 Hardware Platform 29
10.3 Interrupt Latency 29
10.4 Context Switch 29
10.5 Directive Times 30
10.6 Task Manager 31
10.7 Interrupt Manager 32
10.8 Clock Manager 32
10.9 Timer Manager 32
10.10 Semaphore Manager 33
10.11 Message Manager 33
10.12 Event Manager 34
10.13 Signal Manager 34
10.14 Partition Manager 34
10.15 Region Manager 35
10.16 Dual-Ported Memory Manager 35
10.17 I/O Manager 35
10.18 Rate Monotonic Manager 35

Command and Variable Index 37

Concept Index 39

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.

For information on the i960CA and the i960 processor family in general, refer to the following documents:

- *80960CA User's Manual, Intel, Order No. 270710.*
- *32-Bit Embedded Controller Handbook, Intel, Order No. 270647.*
- *Glenford J. Meyers and David L. Budde. The 80960 Microprocessor Architecture. Wiley. New York. 1988.*

It is highly recommended that the i960CA RTEMS application developer obtain and become familiar with Intel's i960CA User's Manual.

1 CPU Model Dependent Features

1.1 Introduction

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 PA-RISC are based on an architectural specification which is independent or 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. 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 chapter presents the set of features which vary across i960 implementations and are of importance to RTEMS. The set of CPU model feature macros are defined in the file c/src/exec/score/cpu/i960/i960.h based upon the particular CPU model defined on the compilation command line.

1.2 CPU Model Name

The macro CPU_MODEL_NAME is a string which designates the name of this CPU model. For example, for the Intel i960CA, this macro is set to the string "i960ca".

1.3 Floating Point Unit

The macro I960_HAS_FPU is set to 1 to indicate that this CPU model has a hardware floating point unit and 0 otherwise.

2 Calling Conventions

2.1 Introduction

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.

2.2 Processor Background

All members of the i960 architecture family support two methods for performing procedure calls: a RISC-style branch-and-link and an integrated call and return mechanism.

On a branch-and-link, the processor branches to the invoked procedure and saves the return address in a register, G14. Typically, the invoked procedure will not invoke another procedure and is referred to as a leaf procedure. Many high-level language compilers for the i960 family recognize leaf procedures and automatically optimize them to utilize the branch-and-link mechanism. Branch-and-link procedures are invoked using the bal and balx instructions and return control via the bx instruction. By convention, G14 is zero when not in a leaf procedure. It is the responsibility of the leaf procedure to clear G14 before returning.

The integrated call and return mechanism also branches to the invoked procedure and saves the return address as did the branch and link mechanism. However, the important difference is that the call, callx, and calls instructions save the local register set (R0 through R15) before transferring control to the invoked procedure. The ret instruction automatically restores the previous local register set. The i960CA provides a register cache which can be configured to retain the last five to sixteen recent register caches. When the register cache is full, the oldest cached register set is written to the stack.

2.3 Calling Mechanism

All RTEMS directives are invoked using either a call or callx instruction and return to the user via the ret instruction.

2.4 Register Usage

As discussed above, the call and callx instructions automatically save the current contents of the local register set (R0 through R15). The contents of the local registers will be restored as part of returning to the application. The contents of global registers G0 through G7 are not preserved by RTEMS directives.

2.5 Parameter Passing

RTEMS uses the standard i960 family C parameter passing mechanism in which G0 contains the first parameter, G1 the second, and so on for the remaining parameters. No RTEMS directive requires more than six parameters.

2.6 User-Provided Routines

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

2.7 Leaf Procedures

RTEMS utilizes leaf procedures internally to improve performance. This improves execution speed as well as reducing stack usage and the number of register sets which must be cached.

3 Memory Model

3.1 Introduction

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.

3.2 Flat Memory Model

The i960CA supports a flat 32-bit address space with addresses ranging from 0x00000000 to 0xFFFFFFFF (4 gigabytes). Although the i960CA reserves portions of this address space, application code and data may be placed in any non-reserved areas. 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), double-word (8 bytes), triple-word (12 bytes) or quad-word (16 bytes). The i960CA does not support virtual memory or segmentation.

The i960CA allows the memory space to be partitioned into sixteen regions which may be configured individually as big or little endian. RTEMS assumes that the memory regions in which its code, data, and the RTEMS Workspace reside are configured as little endian.

4 Interrupt Processing

4.1 Introduction

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 the proper handling of an interrupt. The processor dependent response to the interrupt which modifies the execution state and results in the modification of the execution stream. This modification usually requires that an interrupt handler utilize the provided 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 the processor's response and control mechanisms as they pertain to RTEMS.

4.2 Vectoring of Interrupt Handler

Upon receipt of an interrupt the i960CA automatically performs the following actions:

- saves the local register set,
- sets the Frame Pointer (FP) to point to the interrupt stack,
- increments the FP by sixteen (16) to make room for the Interrupt Record,
- saves the current values of the arithmetic-controls (AC) register, the process-controls (PC) register, and the interrupt vector number are saved in the Interrupt Record,
- the CPU sets the Instruction Pointer (IP) to the address of the first instruction in the interrupt handler,
- the return-status field of the Previous Frame Pointer (PFP or R0) register is set to interrupt return,
- sets the PC state bit to interrupted,
- sets the current interrupt disable level in the PC to the level of the current interrupt, and
- disables tracing.

A nested interrupt is processed similarly by the i960CA with the exception that the Frame Pointer (FP) already points to the interrupt stack. This means that the FP is NOT overwritten before space for the Interrupt Record is allocated.

The state flag bit of the saved PC register in the Interrupt Record is examined by RTEMS to determine when an outer most interrupt is being exited. Therefore, the user application code MUST NOT modify this bit.

4.3 Interrupt Record

The structure of the Interrupt Record for the i960CA which is placed on the interrupt stack by the processor in response to an interrupt is as follows:

Saved Process Controls	NFP-16
Saved Arithmetic Controls	NFP-12
UNUSED	NFP-8
UNUSED	NFP-4

4.4 Interrupt Levels

Thirty-two levels (0-31) of interrupt priorities are supported by the i960CA microprocessor with level thirty-one (31) 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 i960CA only supports thirty-two. RTEMS interrupt levels 0 through 31 directly correspond to i960CA interrupt levels. All other RTEMS interrupt levels are undefined and their behavior is unpredictable.

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 thirty-one (31) 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 2.5 microseconds on a 33 Mhz i960CA with zero wait states. These numbers will vary based the number of wait states and processor speed present on the target board. [NOTE: This calculation was most recently performed for Release 3.2.1.]

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.

4.6 Register Cache Flushing

The i960CA version of the RTEMS interrupt manager is optimized to insure that the flushreg instruction is only executed when a context switch is necessary. The flushreg instruction flushes the i960CA register set cache and takes (14 + 23 * number of sets flushed) cycles to execute. As the i960CA supports caching of from five to sixteen register sets, this instruction takes from 129 to 382 cycles (3.90 to 11.57 microseconds at 33 Mhz) to execute given no wait state memory. RTEMS flushes the register set cache only at the conclusion of the outermost ISR when a context switch is necessary. The register set cache will not be flushed as part of processing a nested interrupt or when a context switch is not necessary. This optimization is essential to providing high-performance interrupt management on the i960CA.

4.7 Interrupt Stack

On the i960CA, 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. During the initialization process, RTEMS will install its interrupt stack.

5 Default Fatal Error Processing

5.1 Introduction

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 is 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.

5.2 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 31, places the error code in G0, and executes a branch to self instruction to simulate a halt processor instruction.

6 Board Support Packages

6.1 Introduction

An RTEMS Board Support Package (BSP) must be designed to support a particular processor and target board combination. This chapter presents a discussion of i960CA 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.

6.2 System Reset

An RTEMS based application is initiated when the i960CA processor is reset. When the i960CA is reset, the processor reads an Initial Memory Image (IMI) to establish its state. The IMI consists of the Initialization Boot Record (IBR) and the Process Control Block (PRCB) from an Initial Memory Image (IMI) at location 0xFFFFFF00. The IBR contains the initial bus configuration data, the address of the first instruction to execute after reset, the address of the PRCB, and the checksum used by the processor's self-test.

6.3 Processor Initialization

The PRCB contains the base addresses for system data structures, and initial configuration information for the core and integrated peripherals. In particular, the PRCB contains the initial contents of the Arithmetic Control (AC) Register as well as the base addresses of the Interrupt Vector Table, System Procedure Entry Table, Fault Entry Table, and the Control Table. In addition, the PRCB is used to configure the depth of the instruction and register caches and the actions when certain types of faults are encountered.

The Process Controls (PC) Register is initialized to 0xC01F2002 which sets the i960CA's interrupt level to 0x1F (31 decimal). In addition, the Interrupt Mask (IMSK) Register (alternately referred to as Special Function Register 1 or sf1) is set to 0x00000000 to mask all external and DMA interrupt sources. Thus, all interrupts are disabled when the first instruction is executed.

For more information regarding the i960CA's data structures and their contents, refer to Intel's i960CA User's Manual.

7 Processor Dependent Information Table

7.1 Introduction

Any highly processor dependent information required to describe a processor to RTEMS is provided in the CPU Dependent Information Table. This table is not required for all processors supported by RTEMS. This chapter describes the contents, if any, for a particular processor type.

7.2 CPU Dependent Information Table

The i960CA version of the RTEMS CPU Dependent Information Table contains the information required to interface a Board Support Package and RTEMS on the i960CA. This information is provided to allow RTEMS to interoperate effectively with the BSP. The C structure definition is given here:

```
typedef struct {
  void        (*pretasking_hook)( void );
  void        (*predriver_hook)( void );
  void        (*postdriver_hook)( void );
  void        (*idle_task)( void );
  boolean       do_zero_of_workspace;
  unsigned32    idle_task_stack_size;
  unsigned32    interrupt_stack_size;
  unsigned32    extra_mpci_receive_server_stack;
  void       (*stack_free_hook)( void* );
  /* end of fields required on all CPUs */

  i960_PRCB    *Prcb;

} rtems_cpu_table;
```

The contents of the i960 Processor Control Block are discussed in the User's Manual for the particular i960 model being used. Structure definitions for the i960CA and i960HA PRCB and Control Table are provided by including the file `rtems.h`.

`pretasking_hook` is the address of the user provided routine which is invoked once RTEMS APIs are initialized. This routine will be invoked before any system tasks are created. Interrupts are disabled. This field may be NULL to indicate that the hook is not utilized.

`predriver_hook` is the address of the user provided routine that is invoked immediately before the the device drivers and MPCI are initialized. RTEMS initialization is complete but interrupts and tasking are disabled. This field may be NULL to indicate that the hook is not utilized.

`postdriver_hook` is the address of the user provided routine that is invoked immediately after the the device drivers and MPCI are initialized. RTEMS

initialization is complete but interrupts and tasking are disabled. This field may be NULL to indicate that the hook is not utilized.

`idle_task` is the address of the optional user provided routine which is used as the system's IDLE task. If this field is not NULL, then the RTEMS default IDLE task is not used. This field may be NULL to indicate that the default IDLE is to be used.

`do_zero_of_workspace`
indicates whether RTEMS should zero the Workspace as part of its initialization. If set to TRUE, the Workspace is zeroed. Otherwise, it is not.

`idle_task_stack_size`
is the size of the RTEMS idle task stack in bytes. If this number is less than MINIMUM_STACK_SIZE, then the idle task's stack will be MINIMUM_STACK_SIZE in byte.

`interrupt_stack_size`
is the size of the RTEMS allocated interrupt stack in bytes. This value must be at least as large as MINIMUM_STACK_SIZE.

`extra_mpci_receive_server_stack`
is the extra stack space allocated for the RTEMS MPCI receive server task in bytes. The MPCI receive server may invoke nearly all directives and may require extra stack space on some targets.

`stack_allocate_hook`
is the address of the optional user provided routine which allocates memory for task stacks. If this hook is not NULL, then a stack_free_hook must be provided as well.

`stack_free_hook` is the address of the optional user provided routine which frees memory for task stacks. If this hook is not NULL, then a stack_allocate_hook must be provided as well.

`Prcb` is the base address of the Processor Control Block. It is primarily used by RTEMS to install interrupt handlers.

8 Memory Requirements

8.1 Introduction

Memory is typically a limited resource in real-time embedded systems, therefore, RTEMS can be configured to utilize the minimum amount of memory while meeting all of the applications requirements. Worksheets are provided which allow the RTEMS application developer to determine the amount of RTEMS code and RAM workspace which is required by the particular configuration. Also provided are the minimum code space, maximum code space, and the constant data space required by RTEMS.

8.2 Data Space Requirements

RTEMS requires a small amount of memory for its private variables. This data area must be in RAM and is separate from the RTEMS RAM Workspace. The following illustrates the data space required for all configurations of RTEMS:

- Data Space: 128

8.3 Minimum and Maximum Code Space Requirements

A maximum configuration of RTEMS includes the core and all managers, including the multiprocessing manager. Conversely, a minimum configuration of RTEMS includes only the core and the following managers: initialization, task, interrupt and fatal error. The following illustrates the code space required by these configurations of RTEMS:

- Minimum Configuration: xx,129
- Maximum Configuration: xx,130

8.4 RTEMS Code Space Worksheet

The RTEMS Code Space Worksheet is a tool provided to aid the RTEMS application designer to accurately calculate the memory required by the RTEMS run-time environment. RTEMS allows the custom configuration of the executive by optionally excluding managers which are not required by a particular application. This worksheet provides the included and excluded size of each manager in tabular form allowing for the quick calculation of any custom configuration of RTEMS. The RTEMS Code Space Worksheet is below:

RTEMS Code Space Worksheet

Component	Included	Not Included	Size
Core	x,131	NA	
Initialization	x,132	NA	
Task	x,133	NA	
Interrupt	x,134	NA	
Clock	x,135	NA	
Timer	x,136	148	
Semaphore	x,137	149	
Message	x,138	150	
Event	x,139	151	
Signal	x,140	152	
Partition	x,141	153	
Region	x,142	154	
Dual Ported Memory	x,143	155	
I/O	x,144	156	
Fatal Error	x,145	NA	
Rate Monotonic	x,146	157	
Multiprocessing	x,147	158	
Total Code Space Requirements			

8.5 RTEMS RAM Workspace Worksheet

The RTEMS RAM Workspace Worksheet is a tool provided to aid the RTEMS application designer to accurately calculate the minimum memory block to be reserved for RTEMS use. This worksheet provides equations for calculating the amount of memory required based upon the number of objects configured, whether for single or multiple processor versions of the executive. This information is presented in tabular form, along with the fixed system requirements, allowing for quick calculation of any application defined configuration of RTEMS. The RTEMS RAM Workspace Worksheet is provided below:

RTEMS RAM Workspace Worksheet

Description	Equation	Bytes Required
maximum_tasks	* 159 =	
maximum_timers	* 160 =	
maximum_semaphores	* 161 =	
maximum_message_queues	* 162 =	
maximum_regions	* 163 =	
maximum_partitions	* 164 =	
maximum_ports	* 165 =	
maximum_periods	* 166 =	
maximum_extensions	* 167 =	
Floating Point Tasks	* 168 =	
Task Stacks	=	
Total Single Processor Requirements		
Description	**Equation**	**Bytes Required**
maximum_nodes	* 169 =	
maximum_global_objects	* 170 =	
maximum_proxies	* 171 =	
Total Multiprocessing Requirements		
Fixed System Requirements	x,172	
Total Single Processor Requirements		
Total Multiprocessing Requirements		
Minimum Bytes for RTEMS Workspace		

9 Timing Specification

9.1 Introduction

This chapter provides information pertaining to the measurement of the performance of RTEMS, the methods of gathering the timing data, and the usefulness of the data. Also discussed are other time critical aspects of RTEMS that affect an applications design and ultimate throughput. These aspects include determinancy, interrupt latency and context switch times.

9.2 Philosophy

Benchmarks are commonly used to evaluate the performance of software and hardware. Benchmarks can be an effective tool when comparing systems. Unfortunately, benchmarks can also be manipulated to justify virtually any claim. Benchmarks of real-time executives are difficult to evaluate for a variety of reasons. Executives vary in the robustness of features and options provided. Even when executives compare favorably in functionality, it is quite likely that different methodologies were used to obtain the timing data. Another problem is that some executives provide times for only a small subset of directives, This is typically justified by claiming that these are the only time-critical directives. The performance of some executives is also very sensitive to the number of objects in the system. To obtain any measure of usefulness, the performance information provided for an executive should address each of these issues.

When evaluating the performance of a real-time executive, one typically considers the following areas: determinancy, directive times, worst case interrupt latency, and context switch time. Unfortunately, these areas do not have standard measurement methodologies. This allows vendors to manipulate the results such that their product is favorably represented. We have attempted to provide useful and meaningful timing information for RTEMS. To insure the usefulness of our data, the methodology and definitions used to obtain and describe the data are also documented.

9.2.1 Determinancy

The correctness of data in a real-time system must always be judged by its timeliness. In many real-time systems, obtaining the correct answer does not necessarily solve the problem. For example, in a nuclear reactor it is not enough to determine that the core is overheating. This situation must be detected and acknowledged early enough that corrective action can be taken and a meltdown avoided.

Consequently, a system designer must be able to predict the worst-case behavior of the application running under the selected executive. In this light, it is important that a real-time system perform consistently regardless of the number of tasks, semaphores, or other resources allocated. An important design goal of a real-time executive is that all internal algorithms be fixed-cost. Unfortunately, this goal is difficult to completely meet without sacrificing the robustness of the executive's feature set.

Many executives use the term deterministic to mean that the execution times of their services can be predicted. However, they often provide formulas to modify execution times based upon the number of objects in the system. This usage is in sharp contrast to the notion of deterministic meaning fixed cost.

Almost all RTEMS directives execute in a fixed amount of time regardless of the number of objects present in the system. The primary exception occurs when a task blocks while acquiring a resource and specifies a non-zero timeout interval.

Other exceptions are message queue broadcast, obtaining a variable length memory block, object name to ID translation, and deleting a resource upon which tasks are waiting. In addition, the time required to service a clock tick interrupt is based upon the number of timeouts and other "events" which must be processed at that tick. This second group is composed primarily of capabilities which are inherently non-deterministic but are infrequently used in time critical situations. The major exception is that of servicing a clock tick. However, most applications have a very small number of timeouts which expire at exactly the same millisecond (usually none, but occasionally two or three).

9.2.2 Interrupt Latency

Interrupt latency is the delay between the CPU's receipt of an interrupt request and the execution of the first application-specific instruction in an interrupt service routine. Interrupts are a critical component of most real-time applications and it is critical that they be acted upon as quickly as possible.

Knowledge of the worst case interrupt latency of an executive aids the application designer in determining the maximum period of time between the generation of an interrupt and an interrupt handler responding to that interrupt. The interrupt latency of an system is the greater of the executive's and the applications's interrupt latency. If the application disables interrupts longer than the executive, then the application's interrupt latency is the system's worst case interrupt disable period.

The worst case interrupt latency for a real-time executive is based upon the following components:

- the longest period of time interrupts are disabled by the executive,
- the overhead required by the executive at the beginning of each ISR,
- the time required for the CPU to vector the interrupt, and
- for some microprocessors, the length of the longest instruction.

The first component is irrelevant if an interrupt occurs when interrupts are enabled, although it must be included in a worst case analysis. The third and fourth components are particular to a CPU implementation and are not dependent on the executive. The fourth component is ignored by this document because most applications use only a subset of a microprocessor's instruction set. Because of this the longest instruction actually executed is application dependent. The worst case interrupt latency of an executive is typically defined as the sum of components (1) and (2). The second component includes the time necessry for RTEMS to save registers and vector to the user-defined handler. RTEMS includes the

third component, the time required for the CPU to vector the interrupt, because it is a required part of any interrupt.

Many executives report the maximum interrupt disable period as their interrupt latency and ignore the other components. This results in very low worst-case interrupt latency times which are not indicative of actual application performance. The definition used by RTEMS results in a higher interrupt latency being reported, but accurately reflects the longest delay between the CPU's receipt of an interrupt request and the execution of the first application-specific instruction in an interrupt service routine.

The actual interrupt latency times are reported in the Timing Data chapter of this supplement.

9.2.3 Context Switch Time

An RTEMS context switch is defined as the act of taking the CPU from the currently executing task and giving it to another task. This process involves the following components:

- Saving the hardware state of the current task.
- Optionally, invoking the TASK_SWITCH user extension.
- Restoring the hardware state of the new task.

RTEMS defines the hardware state of a task to include the CPU's data registers, address registers, and, optionally, floating point registers.

Context switch time is often touted as a performance measure of real-time executives. However, a context switch is performed as part of a directive's actions and should be viewed as such when designing an application. For example, if a task is unable to acquire a semaphore and blocks, a context switch is required to transfer control from the blocking task to a new task. From the application's perspective, the context switch is a direct result of not acquiring the semaphore. In this light, the context switch time is no more relevant than the performance of any other of the executive's subroutines which are not directly accessible by the application.

In spite of the inappropriateness of using the context switch time as a performance metric, RTEMS context switch times for floating point and non-floating points tasks are provided for comparison purposes. Of the executives which actually support floating point operations, many do not report context switch times for floating point context switch time. This results in a reported context switch time which is meaningless for an application with floating point tasks.

The actual context switch times are reported in the Timing Data chapter of this supplement.

9.2.4 Directive Times

Directives are the application's interface to the executive, and as such their execution times are critical in determining the performance of the application. For example, an application using a semaphore to protect a critical data structure should be aware of the time required to acquire and release a semaphore. In addition, the application designer can uti-

lize the directive execution times to evaluate the performance of different synchronization and communication mechanisms.

The actual directive execution times are reported in the Timing Data chapter of this supplement.

9.3 Methodology

9.3.1 Software Platform

The RTEMS timing suite is written in C. The overhead of passing arguments to RTEMS by C is not timed. The times reported represent the amount of time from entering to exiting RTEMS.

The tests are based upon one of two execution models: (1) single invocation times, and (2) average times of repeated invocations. Single invocation times are provided for directives which cannot easily be invoked multiple times in the same scenario. For example, the times reported for entering and exiting an interrupt service routine are single invocation times. The second model is used for directives which can easily be invoked multiple times in the same scenario. For example, the times reported for semaphore obtain and semaphore release are averages of multiple invocations. At least 100 invocations are used to obtain the average.

9.3.2 Hardware Platform

Since RTEMS supports a variety of processors, the hardware platform used to gather the benchmark times must also vary. Therefore, for each processor supported the hardware platform must be defined. Each definition will include a brief description of the target hardware platform including the clock speed, memory wait states encountered, and any other pertinent information. This definition may be found in the processor dependent timing data chapter within this supplement.

9.3.3 What is measured?

An effort was made to provide execution times for a large portion of RTEMS. Times were provided for most directives regardless of whether or not they are typically used in time critical code. For example, execution times are provided for all object create and delete directives, even though these are typically part of application initialization.

The times include all RTEMS actions necessary in a particular scenario. For example, all times for blocking directives include the context switch necessary to transfer control to a new task. Under no circumstances is it necessary to add context switch time to the reported times.

The following list describes the objects created by the timing suite:

- All tasks are non-floating point.
- All tasks are created as local objects.

- No timeouts are used on blocking directives.
- All tasks wait for objects in FIFO order.

In addition, no user extensions are configured.

9.3.4 What is not measured?

The times presented in this document are not intended to represent best or worst case times, nor are all directives included. For example, no times are provided for the initialize executive and fatal_error_occurred directives. Other than the exceptions detailed in the Determinancy section, all directives will execute in the fixed length of time given.

Other than entering and exiting an interrupt service routine, all directives were executed from tasks and not from interrupt service routines. Directives invoked from ISRs, when allowable, will execute in slightly less time than when invoked from a task because rescheduling is delayed until the interrupt exits.

9.3.5 Terminology

The following is a list of phrases which are used to distinguish individual execution paths of the directives taken during the RTEMS performance analysis:

another task The directive was performed on a task other than the calling task.

available A task attempted to obtain a resource and immediately acquired it.

blocked task The task operated upon by the directive was blocked waiting for a resource.

caller blocks The requested resoure was not immediately available and the calling task chose to wait.

calling task The task invoking the directive.

messages flushed One or more messages was flushed from the message queue.

no messages flushed No messages were flushed from the message queue.

not available A task attempted to obtain a resource and could not immediately acquire it.

no reschedule The directive did not require a rescheduling operation.

NO_WAIT A resource was not available and the calling task chose to return immediately via the NO_WAIT option with an error.

obtain current The current value of something was requested by the calling task.

preempts caller The release of a resource caused a task of higher priority than the calling to be readied and it became the executing task.

ready task The task operated upon by the directive was in the ready state.

reschedule The actions of the directive necessitated a rescheduling operation.

returns to caller The directive succeeded and immediately returned to the calling task.

returns to interrupted task
: The instructions executed immediately following this interrupt will be in the interrupted task.

returns to nested interrupt
: The instructions executed immediately following this interrupt will be in a previously interrupted ISR.

returns to preempting task
: The instructions executed immediately following this interrupt or signal handler will be in a task other than the interrupted task.

signal to self
: The signal set was sent to the calling task and signal processing was enabled.

suspended task
: The task operated upon by the directive was in the suspended state.

task readied
: The release of a resource caused a task of lower or equal priority to be readied and the calling task remained the executing task.

yield
: The act of attempting to voluntarily release the CPU.

10 CVME961 Timing Data

NOTE: The CVME961 board used by the RTEMS Project to obtain i960CA times is currently broken. The information in this chapter was obtained using Release 3.2.1.

10.1 Introduction

The timing data for the i960CA version of RTEMS is provided along with the target dependent aspects concerning the gathering of the timing data. The hardware platform used to gather the times is described to give the reader a better understanding of each directive time provided. Also, provided is a description of the interrupt latency and the context switch times as they pertain to the i960CA version of RTEMS.

10.2 Hardware Platform

All times reported except for the maximum period interrupts are disabled by RTEMS were measured using a Cyclone Microsystems CVME961 board. The CVME961 is a 33 Mhz board with dynamic RAM which has two wait state dynamic memory (four CPU cycles) for read accesses and one wait state (two CPU cycles) for write accesses. The Z8536 on a SQUALL SQSIO4 mezzanine board was used to measure elapsed time with one-half microsecond resolution. All sources of hardware interrupts are disabled, although the interrupt level of the i960CA allows all interrupts.

The maximum interrupt disable period was measured by summing the number of CPU cycles required by each assembly language instruction executed while interrupts were disabled. Zero wait state memory was assumed. The total CPU cycles executed with interrupts disabled, including the instructions to disable and enable interrupts, was divided by 33 to simulate a i960CA executing at 33 Mhz with zero wait states.

10.3 Interrupt Latency

The maximum period with interrupts disabled within RTEMS is less than 2.5 microseconds including the instructions which disable and re-enable interrupts. The time required for the i960CA to generate an interrupt using the sysctl instruction, vectoring to an interrupt handler, and for the RTEMS entry overhead before invoking the user's interrupt handler are a total of 37 microseconds. These combine to yield a worst case interrupt latency of less than 2.5 + 37 microseconds. [NOTE: The maximum period with interrupts disabled within RTEMS was last calculated for Release 3.2.1.]

It should be noted again that the maximum period with interrupts disabled within RTEMS is hand-timed. The interrupt vector and entry overhead time was generated on the Cyclone CVME961 benchmark platform using the sysctl instruction as the interrupt source.

10.4 Context Switch

The RTEMS processor context switch time is 1 microseconds on the Cyclone CVME961 benchmark platform. This time represents the raw context switch time with no user exten-

sions configured. Additional execution time is required when a TSWITCH user extension is configured. The use of the TSWITCH extension is application dependent. Thus, its execution time is not considered part of the base context switch time.

The CVME961 has no hardware floating point capability and floating point tasks are not supported.

The following table summarizes the context switch times for the CVME961 benchmark platform:

No Floating Point Contexts	1
Floating Point Contexts	
restore first FP task	2
save initialized, restore initialized	3
save idle, restore initialized	4
save idle, restore idle	5

10.5 Directive Times

This sections is divided into a number of subsections, each of which contains a table listing the execution times of that manager's directives.

10.6 Task Manager

TASK_CREATE	6
TASK_IDENT	7
TASK_START	8
TASK_RESTART	
calling task	9
suspended task – returns to caller	9
blocked task – returns to caller	10
ready task – returns to caller	11
suspended task – preempts caller	12
blocked task – preempts caller	13
ready task – preempts caller	14
TASK_DELETE	
calling task	15
suspended task	16
blocked task	17
ready task	18
TASK_SUSPEND	
calling task	19
returns to caller	20
TASK_RESUME	
task readied – returns to caller	21
task readied – preempts caller	22
TASK_SET_PRIORITY	
obtain current priority	23
returns to caller	24
preempts caller	25
TASK_MODE	
obtain current mode	26
no reschedule	27
reschedule – returns to caller	28
reschedule – preempts caller	29
TASK_GET_NOTE	30
TASK_SET_NOTE	31
TASK_WAKE_AFTER	
yield – returns to caller	32
yield – preempts caller	33
TASK_WAKE_WHEN	34

10.7 Interrupt Manager

It should be noted that the interrupt entry times include vectoring the interrupt handler.

Interrupt Entry Overhead	
returns to nested interrupt	35
returns to interrupted task	36
returns to preempting task	37
Interrupt Exit Overhead	
returns to nested interrupt	38
returns to interrupted task	39
returns to preempting task	40

10.8 Clock Manager

CLOCK_SET	41
CLOCK_GET	42
CLOCK_TICK	43

10.9 Timer Manager

TIMER_CREATE	44
TIMER_IDENT	45
TIMER_DELETE	
inactive	46
active	47
TIMER_FIRE_AFTER	
inactive	48
active	49
TIMER_FIRE_WHEN	
inactive	50
active	51
TIMER_RESET	
inactive	52
active	53
TIMER_CANCEL	
inactive	54
active	55

10.10 Semaphore Manager

SEMAPHORE_CREATE	56
SEMAPHORE_IDENT	57
SEMAPHORE_DELETE	58
SEMAPHORE_OBTAIN	
available	59
not available – NO_WAIT	60
not available – caller blocks	61
SEMAPHORE_RELEASE	
no waiting tasks	62
task readied – returns to caller	63
task readied – preempts caller	64

10.11 Message Manager

MESSAGE_QUEUE_CREATE	65
MESSAGE_QUEUE_IDENT	66
MESSAGE_QUEUE_DELETE	67
MESSAGE_QUEUE_SEND	
no waiting tasks	68
task readied – returns to caller	69
task readied – preempts caller	70
MESSAGE_QUEUE_URGENT	
no waiting tasks	71
task readied – returns to caller	72
task readied – preempts caller	73
MESSAGE_QUEUE_BROADCAST	
no waiting tasks	74
task readied – returns to caller	75
task readied – preempts caller	76
MESSAGE_QUEUE_RECEIVE	
available	77
not available – NO_WAIT	78
not available – caller blocks	79
MESSAGE_QUEUE_FLUSH	
no messages flushed	80
messages flushed	81

10.12 Event Manager

EVENT_SEND	
no task readied	82
task readied – returns to caller	83
task readied – preempts caller	84
EVENT_RECEIVE	
obtain current events	85
available	86
not available – NO_WAIT	87
not available – caller blocks	88

10.13 Signal Manager

SIGNAL_CATCH	89
SIGNAL_SEND	
returns to caller	90
signal to self	91
EXIT ASR OVERHEAD	
returns to calling task	92
returns to preempting task	93

10.14 Partition Manager

PARTITION_CREATE	94
PARTITION_IDENT	95
PARTITION_DELETE	96
PARTITION_GET_BUFFER	
available	97
not available	98
PARTITION_RETURN_BUFFER	98

10.15 Region Manager

REGION_CREATE	100
REGION_IDENT	101
REGION_DELETE	102
REGION_GET_SEGMENT	
available	103
not available – NO_WAIT	104
not available – caller blocks	105
REGION_RETURN_SEGMENT	
no waiting tasks	106
task readied – returns to caller	107
task readied – preempts caller	108

10.16 Dual-Ported Memory Manager

PORT_CREATE	109
PORT_IDENT	110
PORT_DELETE	111
PORT_INTERNAL_TO_EXTERNAL	112
PORT_EXTERNAL_TO_INTERNAL	113

10.17 I/O Manager

IO_INITIALIZE	114
IO_OPEN	115
IO_CLOSE	116
IO_READ	117
IO_WRITE	118
IO_CONTROL	119

10.18 Rate Monotonic Manager

RATE_MONOTONIC_CREATE	120
RATE_MONOTONIC_IDENT	121
RATE_MONOTONIC_CANCEL	122
RATE_MONOTONIC_DELETE	
active	123
inactive	124
RATE_MONOTONIC_PERIOD	
initiate period – returns to caller	125
conclude period – caller blocks	126
obtain status	127

Command and Variable Index

There are currently no Command and Variable Index entries.

Concept Index

There are currently no Concept Index entries.