



Adding Multitasking Capability to Applications Using Multiprocessing Services

For Multiprocessing Services 2.1



October 21, 1999
Technical Publications
© 1999 Apple Computer, Inc.

 Apple Computer, Inc.
© 1999 Apple Computer, Inc.
All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, mechanical, electronic, photocopying, recording, or otherwise, without prior written permission of Apple Computer, Inc., except to make a backup copy of any documentation provided on CD-ROM.

The Apple logo is a trademark of Apple Computer, Inc. Use of the “keyboard” Apple logo (Option-Shift-K) for commercial purposes without the prior written consent of Apple may constitute trademark infringement and unfair competition in violation of federal and state laws.

No licenses, express or implied, are granted with respect to any of the technology described in this book. Apple retains all intellectual property rights associated with the technology described in this book. This book is intended to assist application developers to develop applications only for Apple-labeled or Apple-licensed computers.

Every effort has been made to ensure that the information in this manual is accurate. Apple is not responsible for typographical errors.

Apple Computer, Inc.
1 Infinite Loop
Cupertino, CA 95014
408-996-1010

Apple, the Apple logo, Mac, MacinTalk, and Macintosh are trademarks of Apple Computer, Inc., registered in the United States and other countries.

Adobe, Acrobat, and PostScript are trademarks of Adobe Systems Incorporated or its subsidiaries and may be registered in certain jurisdictions.

Helvetica and Palatino are registered trademarks of Linotype-Hell AG and/or its subsidiaries.

ITC Zapf Dingbats is a registered trademark of International Typeface Corporation.

PowerPC is a trademark of International Business Machines Corporation, used under license therefrom.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company, Ltd.

Simultaneously published in the United States and Canada.

Even though Apple has reviewed this manual, APPLE MAKES NO WARRANTY OR REPRESENTATION, EITHER EXPRESS OR IMPLIED, WITH RESPECT TO THIS MANUAL, ITS QUALITY, ACCURACY, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. AS A RESULT, THIS MANUAL IS SOLD “AS IS,” AND YOU, THE PURCHASER, ARE ASSUMING THE ENTIRE RISK AS TO ITS QUALITY AND ACCURACY.

IN NO EVENT WILL APPLE BE LIABLE FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES RESULTING FROM ANY DEFECT OR INACCURACY IN THIS MANUAL, even if advised of the possibility of such damages.

THE WARRANTY AND REMEDIES SET FORTH ABOVE ARE EXCLUSIVE AND IN LIEU OF ALL OTHERS, ORAL OR WRITTEN, EXPRESS OR IMPLIED. No Apple dealer, agent, or employee is authorized to make any modification, extension, or addition to this warranty.

Some states do not allow the exclusion or limitation of implied warranties or liability for incidental or consequential damages, so the above limitation or exclusion may not apply to you. This warranty gives you specific legal rights, and you may also have other rights which vary from state to state.

Contents

Figures, Tables, and Listings 7

Chapter 1 Introduction 9

Chapter 2 About Multitasking on the Mac OS 11

Multitasking Basics 13
Multitasking and Multiprocessing 14
Tasks and Address Spaces 15
Task Scheduling 16
Shared Resources and Task Synchronization 18
 Semaphores 19
 Message Queues 19
 Event Groups 20
 Kernel Notifications 20
 Critical Regions 21
Tasking Architectures 22
 Multiple Independent Tasks 22
 Parallel Tasks With Parallel I/O Buffers 23
 Parallel Tasks With a Single Set of I/O Buffers 24
 Sequential Tasks 25

Chapter 3 Using Multiprocessing Services 27

Multiprocessing Services in Mac OS 9 and Mac OS X 30
Compatibility with Older System Software 30
Criteria for Creating Tasks 31
Checking for the Availability of Multiprocessing Services 32
Determining the Number of Processors 32
Creating Tasks 33
Terminating Tasks 38

Synchronizing and Notifying Tasks	40
Handling Periodic Actions	43
Notifying Tasks at Interrupt Time	44
Using Critical Regions	45
Allocating Memory in Tasks	46
Using Task-Specific Storage	46
Using Timers	47
Making Remote Procedure Calls	48
Handling Exceptions and Debugging	49

Chapter 4 Multiprocessing Services Reference 51

Gestalt Constants	55
Functions	56
Determining Processor Availability	57
Creating and Scheduling Tasks	59
Creating and Handling Message Queues	67
Creating and Handling Semaphores	73
Handling Event Groups	76
Handling Kernel Notifications	80
Handling Critical Regions	83
Timer Services Functions	87
Time Utility Functions	93
Accessing Per-Task Storage Variables	95
Memory Allocation Functions	99
Exception Handling Functions	103
Debugger Support Functions	109
Remote Calling Function	111
Application-Defined Function	112
Data Types	113
Constants	123
Timer Duration Constants	124
Timer Option Masks	125
Memory Allocation Alignment Constants	126
Memory Allocation Option Constants	127
Task Creation Options	128
Task State Constants	129

Task Run State Constants	130
Task Information Structure Version Constant	131
Task Exception Disposal Constants	131
Remote Call Context Option Constants	132
Result Codes	133

Appendix A	Preemptive Task–Safe Mac OS System Software Functions	135
------------	---	-----

Appendix B	Calculating the Intertask Signaling Time	145
------------	--	-----

Appendix C	Changes From Previous Versions of Multiprocessing Services	153
------------	--	-----

Appendix D	Document Version History	159
------------	--------------------------	-----

Glossary	163
----------	-----

Index	165
-------	-----

Figures, Tables, and Listings

Chapter 2	About Multitasking on the Mac OS	11
<hr/>		
	Figure 2-1	Tasks within processes 16
	Figure 2-2	The Mac OS task and other preemptive tasks 17
	Figure 2-3	Parallel tasks with parallel I/O buffers 24
	Figure 2-4	Parallel tasks with a single set of I/O buffers 25
	Figure 2-5	Sequential tasks 26
Chapter 3	Using Multiprocessing Services	27
<hr/>		
	Listing 3-1	Creating tasks 33
	Listing 3-2	A sample task 36
	Listing 3-3	Terminating tasks 38
	Listing 3-4	Assigning work to tasks 41
	Listing 3-5	Using a semaphore to perform periodic actions 43
	Listing 3-6	Performing actions periodically and on demand 44
Appendix A	Preemptive Task–Safe Mac OS System Software Functions	135
<hr/>		
	Table A-1	Preemptive task–safe Device Manager functions 136
	Table A-2	Preemptive task–safe HFS Plus parameter block functions 136
	Table A-3	Preemptive task–safe HFS Plus file system reference functions 138
	Table A-4	Preemptive task–safe HFS parameter block functions 139
	Table A-5	Preemptive task–safe high-level HFS functions 142
	Table A-6	Preemptive task–safe file system specification functions 143
	Table A-7	Preemptive task–safe functions shared by the Device Manager and File Manager 143
	Table A-8	Miscellaneous preemptive task–safe functions 144
Appendix B	Calculating the Intertask Signaling Time	145
<hr/>		
	Listing B-1	Calculating the intertask signaling time 145

Appendix C Changes From Previous Versions of Multiprocessing Services 153

Table C-1	New functions introduced with version 2.1	153
Table C-2	Functions introduced with version 2.0	153
Table C-3	Older functions supported in version 2.0	155
Table C-4	Unofficial functions still supported in version 2.0	156
Table C-5	Debugging functions unsupported in version 2.0	157

Appendix D Document Version History 159

Table D-1	Multiprocessing Services documentation revision history	159
------------------	---	-----

Introduction

Multiprocessing Services is a technology that allows your application to create tasks that run independently on one or more microprocessors. For example, you can have your application perform graphical calculations while writing data to a hard drive. Unlike the cooperative model (such as used by the Thread Manager or the Mac OS Process Manager), Multiprocessing Services automatically divides processor time among available tasks so no particular task can “hog” the system. On computers with multiple microprocessors, you can actually perform multiple tasks simultaneously. This feature allows you to divide up time-intensive calculations among several microprocessors.

This document is a complete guide to Multiprocessing Services 2.1 . This technology is available with Mac OS 9.0 and later, although some functions may work with earlier system software versions.

You should read this document if you want to add multitasking capability to Mac OS applications. This document assumes you are familiar with programming Macintosh computers. For more information about how the Mac OS handles applications in memory, see the documents *Inside Macintosh: Processes* and *Mac OS Runtime Architectures*.

This document covers Multiprocessing Services in the following chapters:

- “About Multitasking on the Mac OS” (page 13) describes the basics of multitasking and multiprocessing, as well as information about how Multiprocessing Services implements these capabilities on the Mac OS.
- “Using Multiprocessing Services” (page 29) contains programming examples and other detailed information about adding Multiprocessing Services to your application.
- “Multiprocessing Services Reference” (page 55) contains a complete programming reference, documenting the functions, data types, and constants available with Multiprocessing Services.

Introduction

- Appendix A, “Preemptive Task-Safe Mac OS System Software Functions,” lists Mac OS system software functions that you can call from preemptive tasks.
- Appendix B, “Calculating the Intertask Signaling Time,” contains sample code you can use to determine the amount of time it takes to notify a task.
- Appendix C, “Changes From Previous Versions of Multiprocessing Services,” describes changes and additions to the Multiprocessing Services API between version 1.4 and 2.1.
- Appendix D, “Document Version History,” describes changes to this document.

For additional information about Multiprocessing Services, you should check the Apple Developer Web site:

<<http://www.apple.com/developer/>>

About Multitasking on the Mac OS

Contents

Multitasking Basics	13
Multitasking and Multiprocessing	14
Tasks and Address Spaces	15
Task Scheduling	16
Shared Resources and Task Synchronization	18
Semaphores	19
Message Queues	19
Event Groups	20
Kernel Notifications	20
Critical Regions	21
Tasking Architectures	22
Multiple Independent Tasks	22
Parallel Tasks With Parallel I/O Buffers	23
Parallel Tasks With a Single Set of I/O Buffers	24
Sequential Tasks	25

This chapter describes the basic concepts underlying multitasking and how Multiprocessing Services uses them on Macintosh computers.

You should read this chapter if you are not familiar with multitasking or multiprocessing concepts. Note that this chapter covers mostly concepts rather than implementation or programming details. For information about actually using the Multiprocessing Services API in your application, see “Using Multiprocessing Services” (page 29)

The topics covered in this chapter include the following:

- “Multitasking Basics” (page 13)
- “Multitasking and Multiprocessing” (page 14)
- “Tasks and Address Spaces” (page 15)
- “Task Scheduling” (page 16)
- “Shared Resources and Task Synchronization” (page 18)
- “Tasking Architectures” (page 22)

Multitasking Basics

Multitasking is essentially the ability to do many things concurrently. For example, you may be working on a project, eating lunch, and talking to a colleague at the same time. Not everything may be happening simultaneously, but you are jumping back and forth, devoting your attention to each task as necessary.

In programming, a **task** is simply an independent execution path. On a computer, the system software can handle multiple tasks, which may be applications or even smaller units of execution. For example, the system may execute multiple applications, and each application may have independently executing tasks within it. Each such task has its own stack and register set.

Multitasking may be either cooperative or preemptive. **Cooperative multitasking** requires that each task voluntarily give up control so that other tasks can execute. An example of cooperative multitasking is an unsupervised group of children wanting to look at a book. Each child can theoretically get a chance to look at the book. However, if a child is greedy, he or she may spend

About Multitasking on the Mac OS

an inordinate amount of time looking at the book or refuse to give it up altogether. In such cases, the other children are deprived.

Preemptive multitasking allows an external authority to delegate execution time to the available tasks. Preemptive multitasking would be the case where a teacher (or other supervisor) was in charge of letting the children look at the book. He or she would assign the book to each child in turn, making sure that each one got a chance to look at it. The teacher could vary the amount of time each child got to look at the book depending on various circumstances (for example, some children may read more slowly and therefore need more time).

The Mac OS 8 operating system implements cooperative multitasking between applications. The Process Manager can keep track of the actions of several applications. However, each application must voluntarily yield its processor time in order for another application to gain it. An application does so by calling `WaitNextEvent`, which cedes control of the processor until an event occurs that requires the application's attention.

Multiprocessing Services allows you to create preemptive tasks within an application (or process). For example, you can create tasks within an application to process information in the background (such as manipulating an image or retrieving data over the network) while still allowing the user to interact with the application using the user interface.

Note

The definition of *task* in this document is analogous to the use of the term *thread* in some other operating systems such as UNIX[®]. In older documentation, Apple has sometimes referred to separate units of execution as *threads*. For example, the Thread Manager allows you to create cooperatively scheduled threads within a task. You should not confuse these threads with the preemptively scheduled tasks created by Multiprocessing Services. ♦

Multitasking and Multiprocessing

Multitasking and multiprocessing are related concepts, but it is important to understand the distinctions between them. Multitasking is the ability to handle several different tasks at once. **Multiprocessing** is the ability of a computer to use more than one processor simultaneously. Typically, if multiple processors

About Multitasking on the Mac OS

are available in a multitasking environment, then tasks can be divided among the processors. In such cases, tasks can run simultaneously. For example, if you have a large image that you need to manipulate with a filter, you can break up the image into sections, and then assign a task to process each section. If the user's computer contains multiple processors, several tasks can be executed simultaneously, reducing the overall execution time. If only one processor exists, then the tasks are preempted in turn to give each access to the processor.

Note

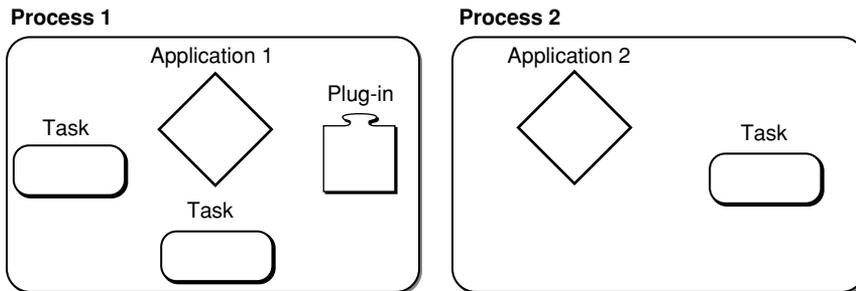
Because multitasking allows an operating system to attend to several different operations at the same time, these operations may *appear* to occur simultaneously on even a single-processor system, due to the speed of the processor.



Multiple processor support is transparent in Multiprocessing Services. If multiple processors exist, then Multiprocessing Services divides the tasks among all the available processors to maximize efficiency (a technique often called **symmetric multiprocessing**). If only one processor exists, then Multiprocessing Services simply schedules the available tasks with the processor to make sure that each task receives attention.

Tasks and Address Spaces

On the Mac OS, all applications are assigned a **process** or *application context* at runtime. The process contains all the resources required by the application, such as allocated memory, stack space, plug-ins, nonglobal data, and so on. Tasks created with Multiprocessing Services are automatically associated with the creating application's process, as shown in Figure 2-1.

Figure 2-1 Tasks within processes

All resources within a process occupy the same address space, so tasks created by the same application are free to share memory. For example, if you want to divide an image filtering operation among multiple identical tasks, you can allocate space for the entire image in memory, and then assign each task the address and length of the portion it should process.

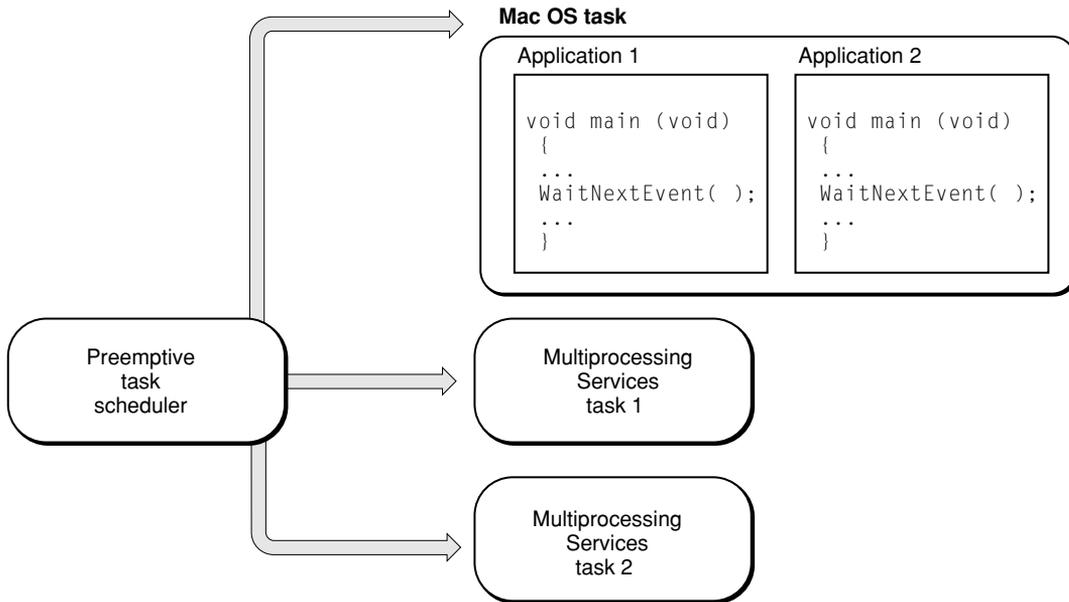
IMPORTANT

Although all processes share the same address space in Mac OS 9.0, you should not assume that this will remain the case; your application or task should not attempt to access data or code residing in another process. ▲

Task Scheduling

Multitasking environments require one or more **task schedulers**, which control how processor time (for one or more processors) is divided among available tasks. Each task you create is added to a task queue. The task scheduler then assigns processor time to each task in turn.

As seen by the Mac OS 9.0 task scheduler, all cooperatively multitasked programs (that is, all the applications that are scheduled by the Process Manager) occupy a single preemptive task called the **Mac OS task**, as shown in Figure 2-2.

Figure 2-2 The Mac OS task and other preemptive tasks

For example, if your cooperatively scheduled application creates a task, the task is preemptively scheduled. The application task (containing the main event loop) is not preemptively scheduled, but it resides within the Mac OS task, which *is* preemptively scheduled. Within the Mac OS task, the application must cooperatively share processor time with any other applications that are currently running.

A task executes until it completes, is blocked, or is preempted. A task is **blocked** when it is waiting for some event or data. For example, a task may need output from a task that has not yet completed, or it may need certain resources that are currently in use by another task.

A blocked task is removed from the task queue until it becomes eligible for execution. The task becomes eligible if either the event that it was waiting for occurs or the waiting period allowed for the event expires.

If the task does not complete or block within a certain amount of time (determined by the scheduler), the task scheduler preempts the task, placing it

About Multitasking on the Mac OS

at the end of the task queue, and gives processor time to the next task in the queue.

Note that if the main application task is blocked while waiting for a Multiprocessing Services event, the blocking application does not get back to its event loop until the event finally occurs. This delay may be unacceptable for long-running tasks. Therefore, in general your application should poll for Multiprocessing Services events from its event loop, rather than block while waiting for them. The task notification mechanisms described in “Shared Resources and Task Synchronization” (page 18) are ideal for this purpose.

Note

In the future, application tasks will run as individual preemptive tasks, rather than within the Mac OS task. However, calls to nonreentrant Mac OS system software functions will cause the task to be blocked for the duration of the call, in a manner similar to a remote procedure call. See “Making Remote Procedure Calls” (page 48) for more information. ♦

Shared Resources and Task Synchronization

Although each created task may execute separately, it may need to share information or otherwise communicate with other tasks. For example, Task 1 may write information to memory that will be read by Task 2. In order for such operations to occur successfully, some synchronization method must exist to make sure that Task 2 does not attempt to read the memory until Task 1 has completed writing the data and until Task 2 knows that valid data actually exists in memory. The latter scenario can be an issue when using multiple processors, because the PowerPC™ architecture allows for writes to memory to be deferred. In addition, if multiple tasks are waiting for another task to complete, synchronization is necessary to ensure that only one task can respond at a time.

Multitasking environments offer several ways for tasks to coordinate and synchronize with each other. The sections that follow describe four **notification mechanisms** (or signaling mechanisms), which allow tasks to pass information between them, and one task sharing method.

About Multitasking on the Mac OS

Note that the time required to perform the work in a given request should be much more than the amount of time it takes to communicate the request and its results. Otherwise, delegating work to tasks may actually reduce overall performance. Typically the work performed should be greater than the intertask signalling time (20-50 microseconds).

Note that you should avoid creating your own synchronization or sharing methods, because they may work on some Mac OS implementations but not on others.

Semaphores

A **semaphore** is a single variable that can be incremented or decremented between zero and some specified maximum value. The value of the semaphore can communicate state information. A mail box flag is an example of a semaphore. You raise the flag to indicate that a letter is waiting in the mailbox. When the mailman picks up the letter, he lowers the flag again. You can use semaphores to keep track of how many occurrences of a particular thing are available for use.

Binary semaphores, which have a maximum value of one, are especially efficient mechanisms for indicating to some other task that something is ready. When a task or application has finished preparing data at some previously agreed to location, it raises the value of a binary semaphore that the target task waits on. The target task lowers the value of the semaphore, performs any necessary processing, and raises the value of a different binary semaphore to indicate that it is through with the data.

Semaphores are quicker and less memory intensive than other notification mechanisms, but due to their size they can convey only limited information.

Message Queues

A **message queue** is a collection of data (messages) that must be processed by tasks in a first-in, first-out order. Several tasks can wait on a single queue, but only one will obtain any particular message. Messages are useful for telling a task what work to do and where to look for information relevant to the request being made. They are also useful for indicating that a given request has been processed and, if necessary, what the results are.

Typically a task has two message queues, one for input and one for output. You can think of message queues as In boxes and Out boxes. For example, your In

About Multitasking on the Mac OS

box at work may contain a number of papers (messages) indicating work to do. After completing a job, you would place another message in the Out box. Note that if you have more than one person assigned to an In box/Out box pair, each person can work independently, allowing data to be processed in parallel.

In Multiprocessing Services, a message is 96-bits of data that can convey any desired information.

Message queues incur more overhead than the other two notification mechanisms. If you must synchronize frequently, you should try to use semaphores instead of message queues whenever possible.

Event Groups

An **event group** is essentially a group of binary semaphores. You can use event groups to indicate a number of simple events. For example, a task running on a server may need to be aware of multiple message queues. Instead of trying to poll each one in turn, the server task can wait on an event group. Whenever a message is posted on a queue, the poster can also set the bit corresponding to that queue in the event group. Doing so notifies the task, and it then knows which queue to access to extract the message. In Multiprocessing Services, an event group consists of thirty-two 1-bit flags, each of which may be set independently. When a task receives an event group, it receives all 32-bits at once (that is, it cannot poll individual bits), and all the bits in the event group are subsequently cleared.

Kernel Notifications

A **kernel notification** is a set of simple notification mechanisms (semaphores, message queues, and event groups) which can be notified with only one signal. For example, a kernel notification might contain both a semaphore and a message queue. When you signal the kernel notification, the semaphore is signaled and a message is sent to the specified queue. A kernel notification can contain one of each type of simple notification mechanism.

You use kernel notifications to hide complexity from the signaler. For example, say a server has three queues to process, ranked according to priority (queue 1 holds the most important tasks, queue 2 holds lesser priority tasks, and queue 3 holds low priority tasks). The server can wait on an event group, which indicates when a task is posted to a queue.

About Multitasking on the Mac OS

If you do not use a kernel notification, then when a client has a message to deliver, it must both send the message to the appropriate queue and signal the event group with the correct priority value. Doing so requires the client to keep track of which queue to send to as well as the proper event group bit to set.

A simpler method is to set up three kernel notifications, one for each priority. To signal a high priority message, the client simply loads the message and signals the high priority kernel notification.

Critical Regions

In addition to notification mechanisms, you can also specify **critical regions** in a multitasking environment. A critical region is a section of code that can be accessed by only one task at a time. For example, say part of a task's job is to search a data tree and modify it. If multiple tasks were allowed to search and try to modify the tree at the same time, the tree would quickly become corrupted. An easy way to avoid the problem is to form a critical region around the tree searching and modification code. When a task tries to enter the critical region it can do so only if no other task is currently in it, thus preserving the integrity of the tree.

Critical regions differ from semaphores in that critical regions can handle recursive entries and code that has multiple entry points. For example, if three functions `func1`, `func2`, and `func3` access some common resource (such as the tree described above), but never call each other, then you can use a semaphore to synchronize access. However, suppose `func3` calls `func1` internally. In that case, `func3` would obtain the semaphore, but when it calls `func1`, it will deadlock. Synchronizing using a critical region instead allows the same task to enter multiple times, so `func1` can enter the region again when called from `func3`.

Because critical regions introduce forced linearity into task execution, improper use can create bottlenecks that reduce performance. For example, if the tree search described above constituted the bulk of a task's work, then the tasks would spend most of their time just trying to get permission to enter the critical region, at great cost to overall performance. A better approach in this case might be to use different critical regions to protect subtrees. You can then have one task search one part of the tree while others were simultaneously working on other parts of the tree.

Tasking Architectures

Determining how to divide work into tasks depends greatly on the type of work you need to do and how the individual tasks rely on each other.

For a computer running multiple processors, you should optimize your multitasking application to keep them as busy as possible. You can do so by creating a number of tasks and letting the task scheduler assign them to available processors, or you can query for the number of available processors and then create enough tasks to keep them all busy.

A simple method is to determine the number of processors available and create as many tasks as there are processors. The application can then split the work into that many pieces and have each processor work on a piece. The application can then poll the tasks from its event loop until all the work is completed.

IMPORTANT

Even if only one processor exists, you should create preemptive tasks to handle faceless computations (filtering, spellchecking, background updating, and so on). Doing so gives the task scheduler more flexibility in assigning processor time, and it will also scale transparently if multiple processors are available. The application should do all the work only if Multiprocessing Services is not available. ▲

The sections that follow describe several common tasking architectures you can use to divide work among multiple processors. You might want to combine these approaches to solve specific problems or come up with your own if none described here are appropriate.

Multiple Independent Tasks

In many cases, you can break down applications into different sections that do not necessarily depend on each other but would ideally run concurrently. For example, your application may have one code section to render images on the screen, another to do graphical computations in the background, and a third to download data from a server. Each such section is a natural candidate for preemptive tasking. Even if only one processor is available, it is generally

About Multitasking on the Mac OS

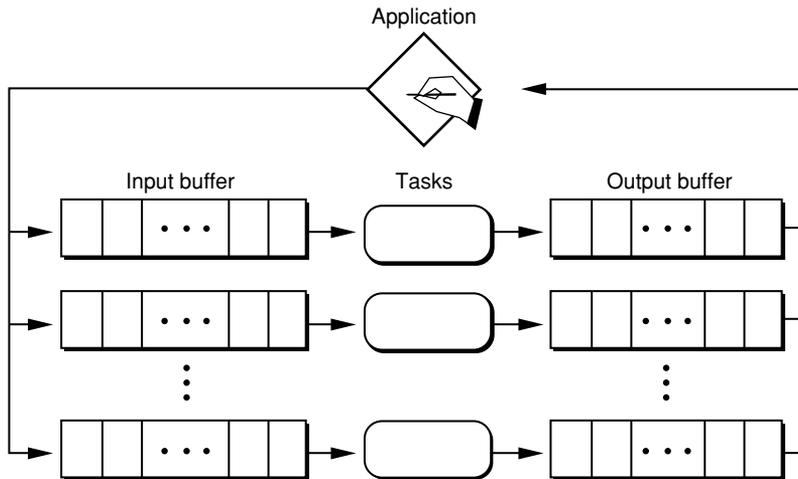
advantageous to have such independent sections run as preemptive tasks. The application can notify the tasks (using any of the three notification mechanisms) and then poll for results within its event loop.

Parallel Tasks With Parallel I/O Buffers

If you can divide the computational work of your application into several similar portions, each of which takes about the same amount of time to complete, you can create as many tasks as the number of processors and divide the work evenly among the tasks (“divide and conquer”). An example would be a filtering task on a large image. You could divide the image into as many equal sections as there are processors and have each do a fraction of the total work.

This method for using Multiprocessing Services involves creating two buffers per task: one for receiving work requests and one for posting results. You can create these buffers using either message queues or semaphores.

As shown in Figure 2-3, the application splits the data evenly among the tasks and posts a work request, which defines the work a task is expected to perform, to each task’s input buffer. Each task asks for a work request from its input buffer, and blocks if none is available. When a request arrives, the task performs the required work and then posts a message to its output buffer indicating that it has finished and providing the application with the results.

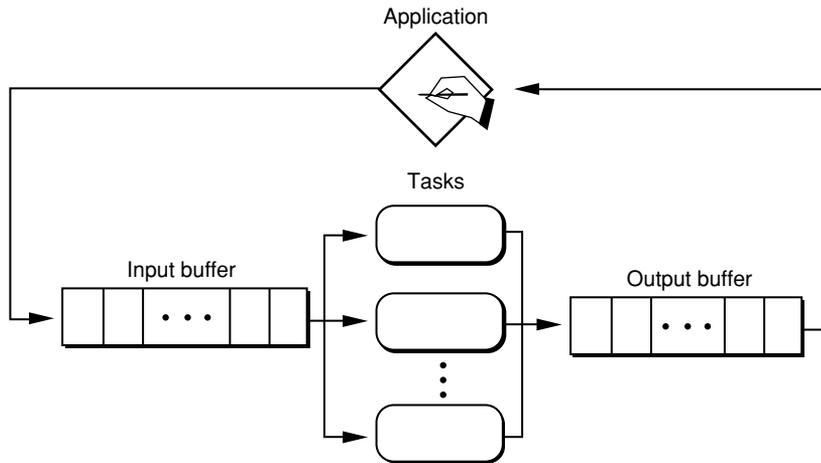
Figure 2-3 Parallel tasks with parallel I/O buffers

Parallel Tasks With a Single Set of I/O Buffers

If you can divide the computational work of your application into portions that can be performed by identical tasks but you can't predict how long each computation will take, you can use a single input buffer for all your tasks. The application places each work request in the input buffer, and each free task asks for a work request. When a task finishes processing the request, it posts the result to a single output buffer shared by all the tasks and asks for a new request from the input buffer. This method is analogous to handling a queue of customers waiting in a bank line. There is no way to predict which task will process which request, and there is no way to predict the order in which results will be placed in the output buffer. For this reason, you might want to have the task include the original work request with the result so the application can determine which result is which.

As in the "divide and conquer" architecture, the application can check events, control data flow, and perform some of the calculations while the tasks are running.

Figure 2-4 illustrates this "bank line" tasking architecture.

Figure 2-4 Parallel tasks with a single set of I/O buffers

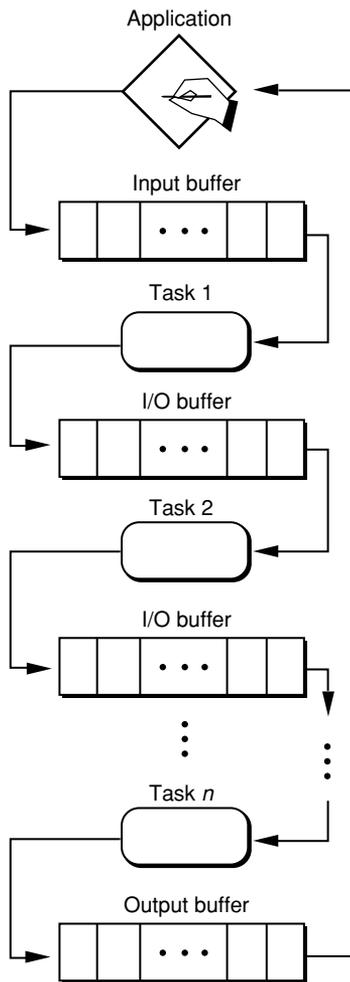
Sequential Tasks

For some applications, you may want to create several different tasks, each of which performs a different function. Some of these tasks might be able to operate in parallel, in which case you can adapt the “divide and conquer” model to perform the work. However, sometimes the output of one task is needed as the input of another. For example a sound effects application may need to process several different effects in sequence (reverb, limiting, flanging, and so on). Each task can process a particular effect, and the output of one task would feed into the input of the next. In this case, a sequential or “pipeline” architecture is appropriate.

Note that a sequential task architecture benefits from multiple processors only if several of the tasks can operate at the same time; that is, new data must be entering the pipeline while other tasks are processing data further down the line. It is harder with this architecture than with parallel task architectures to ensure that all processors are being used at all times, so you should add parallel tasks to individual stages of the pipeline whenever possible.

As shown in Figure 2-5, the sequential task architecture requires a single input buffer and a single output buffer for the pipeline, and intermediate buffers between sequential tasks.

Figure 2-5 Sequential tasks



Using Multiprocessing Services

Contents

Multiprocessing Services in Mac OS 9 and Mac OS X	30
Compatibility with Older System Software	30
Criteria for Creating Tasks	31
Checking for the Availability of Multiprocessing Services	32
Determining the Number of Processors	32
Creating Tasks	33
Terminating Tasks	38
Synchronizing and Notifying Tasks	40
Handling Periodic Actions	43
Notifying Tasks at Interrupt Time	44
Using Critical Regions	45
Allocating Memory in Tasks	46
Using Task-Specific Storage	46
Using Timers	47
Making Remote Procedure Calls	48
Handling Exceptions and Debugging	49

Using Multiprocessing Services

This chapter describes how to incorporate Multiprocessing Services into your Mac OS application. You should read this chapter if you are interested in adding preemptive tasks to your application.

The topics covered in this chapter include:

- “Multiprocessing Services in Mac OS 9 and Mac OS X” (page 30)
- “Compatibility with Older System Software” (page 30)
- “Criteria for Creating Tasks” (page 31)
- “Checking for the Availability of Multiprocessing Services” (page 32)
- “Determining the Number of Processors” (page 32)
- “Creating Tasks” (page 33)
- “Terminating Tasks” (page 38)
- “Synchronizing and Notifying Tasks” (page 40)
- “Allocating Memory in Tasks” (page 46)
- “Using Timers” (page 47)
- “Making Remote Procedure Calls” (page 48)
- “Handling Exceptions and Debugging” (page 49)

Note

This document describes version 2.1 of Multiprocessing Services. For a list of functions changed or added between versions 1.4 and 2.1, see Appendix C. Applications built using older versions of Multiprocessing Services can execute without modification under version 2.1. ◆

IMPORTANT

Preemptive tasks cannot execute 68K code. If you must call 68K code in a task, you must do so through a remote procedure call as described in “Making Remote Procedure Calls” (page 48). ▲

Multiprocessing Services in Mac OS 9 and Mac OS X

Multiprocessing Services 2.1 runs on system software Mac OS 9 and later. All PowerPC Macintosh computers are supported..

Multiprocessing Services 2.1 is packaged as part of the Mac OS System file, so you cannot install version 2.1 on older versions of system software.

Multiprocessing Services allows your application to create preemptive tasks within your application's process (or execution context). However, the individual applications are still cooperatively scheduled by the Process Manager. In Mac OS X, both applications and tasks created by applications will be preemptively scheduled. Multiprocessing Services is Carbon-compliant, so applications built following the Carbon specification can run transparently on both Mac OS 8 and Mac OS X systems.

Compatibility with Older System Software

Multiprocessing Services 2.0 was introduced with Mac OS 8.6. Unlike earlier versions of Multiprocessing Services, you can create and execute preemptive tasks with virtual memory turned on. Multiprocessing Services 2.0 runs on all Power Macintosh computers except for 6100/7100/8100 and 5200/6200 series computers.

Multiprocessing Services 1.0 functions can run on System 7.5.2 and later if the Multiprocessing Services 1.x shared library is available. Pre-2.0 versions of the library were installed as part of system software for Mac OS 8 through Mac OS 8.5 but must be explicitly installed for earlier releases. The 1.x versions of the Multiprocessing Services library can run on all PowerPC Macintosh computers.

For a listing of functions introduced with versions 1.0, 2.0, and 2.1 of Multiprocessing Services, see Appendix C.

Criteria for Creating Tasks

Although you can in theory designate almost any type of code as a task, in practice you should use the following guidelines to make best use of the available processors and to avoid unnecessary bottlenecks.

- Tasks should generally perform faceless processing, such as calculation-intensive work or I/O operations.
- The work performed by a task should be substantially more than the time required to process the request and result notifications. If it takes much longer to notify a task and retrieve results than to execute the task itself, an application's performance will be dramatically worse with multiprocessing. Assuming a typical intertask signaling time is 20-50 microseconds, your tasks should take at least 200-500 microseconds to execute. If you want to explicitly calculate the intertask signaling time, you can use the sample code provided in Appendix B.
- If your task needs to allocate memory, you must allocate the memory prior to signaling the task, or use the function `MPAllocateAligned` (page 99).
- Tasks should not call 68K code. The 68K emulator runs only cooperatively within the Mac OS task, and not within any preemptive task. If you must call 68K code, you can do so using a remote procedure call. See "Making Remote Procedure Calls" (page 48) for more information.
- You can call only preemptive task-safe Mac OS system software functions directly from a task. See "Gestalt Constants" (page 55) and Appendix A to determine which functions are preemptive task-safe. Other (unsafe) system software functions must be called indirectly through remote procedure calls. See "Making Remote Procedure Calls" (page 48) for more information.
- Tasks should not access low-memory global data. A task may be executing at any time, including when applications that did not create them are running.
- Tasks should not call into unknown code. If you allow third parties to specify a callback function, you should never call that function from a task, since you cannot control what the callback will do. Calling back into nonreentrant code could easily corrupt data or cause a system crash.
- Avoid global variables. The main cause of nonreentrancy is the manipulation of global data. Tasks that manipulate global variables, global states, or

Using Multiprocessing Services

buffers pointed to by global variables must use synchronization techniques to prevent other tasks from attempting to do so at the same time. Read-only global data are allowed.

- Do not call any Multiprocessing Services functions at interrupt time unless you are signaling a notification mechanism. See “Notifying Tasks at Interrupt Time” (page 44) for more information.

Checking for the Availability of Multiprocessing Services

You should always determine the availability of Multiprocessing Services before attempting to call any of its functions. If you are programming in C, you should do so by calling the Boolean macro `MPLibraryIsLoaded`. A return value of true indicates that the Multiprocessing Services library is present and available for use. Listing 3-1 (page 33) in “Creating Tasks” shows an example of using the `MPLibraryIsLoaded` macro.

Note

For historical reasons, the Multiprocessing Services shared library may be prepared by the Code Fragment Manager and yet still be unusable; checking for resolved imported symbols is not enough to ensure that the functions are available. Therefore, you must always check for the presence of Multiprocessing Services by using the `MPLibraryIsLoaded` macro. ♦

You probably want your application to run even if Multiprocessing Services is not available, so you should specify a weak link to the Multiprocessing Services shared library. Doing so allows your application to run even if the shared library is not present.

Determining the Number of Processors

You may want to determine the number of processors available on the host computer before creating any tasks. Typically, you would create one task per processor; even if only one processor is present, it is generally more efficient to

assign faceless work to a task and have the cooperatively scheduled main application handle only user interaction.

Multiprocessing Services uses two functions to determine the number of processors. The function `MPPProcessors` (page 57) returns the number of physical processors available on the host computer. The function `MPPProcessorsScheduled` (page 57) returns the number of active processors available (that is, the number that are currently available to execute tasks). The number of active processors may vary over time (due to changing priorities, power consumption issues, and so on).

Creating Tasks

After determining how many processors are available, you can go ahead and create tasks for your application.

Each task must be a function that takes one 32-bit parameter and returns a 32-bit result of type `OSStatus` when it finishes. The 32-bit input parameter can be any information that the task needs to perform its function. Some examples of input are

- a message queue ID that indicates which queue the task should go to for information
- a pointer to a structure containing data to process
- a pointer to a C++ object
- a pointer to a task-specific block of memory through which the application can communicate information for the life of the task

You create a task by calling the function `MPCreateTask` (page 59). The code in Listing 3-1 shows how you can create a number of identical tasks. Identical tasks can be useful when you want to divide up a large calculation (such as a image filtering operation) among several processors to improve performance.

Listing 3-1 Creating tasks

```
#define kMPStackSize 0 // use default stack size
#define kMPTaskOptions 0 // use no options
```

Using Multiprocessing Services

```

typedef struct {
    long firstThing;
    long totalThings;
    } sWorkParams, *sWorkParamsPtr;

typedef struct {
    MPTaskID taskID;
    MPQueueID requestQueue;
    MPQueueID resultQueue;
    sWorkParams params;
    } sTaskData, *sTaskDataPtr;

sTaskDataPtr myTaskData;
UInt32 numProcessors;
MPQueueID notificationQueue;

void CreateMPTasks( void ) {

    OSErr theErr;
    UInt32 i;

    theErr = noErr;

    /* Assume single processor mode */
    numProcessors = 1;

    /* Initialize remaining globals */
    myTaskData = NULL;
    notificationQueue = NULL;

    /* If the library is present then create the tasks */
    if( MPLibraryIsLoaded() ) {
        numProcessors = MPProcessorsScheduled();
        myTaskData = (sTaskDataPtr)NewPtrClear
            ( numProcessors * sizeof( sTaskData ) );
        theErr = MemError();
        if( theErr == noErr )
            theErr = MPCreateQueue( &notificationQueue );
        for( i = 0; i < numProcessors && theErr == noErr; i++ ) {
            if( theErr == noErr )
                theErr = MPCreateQueue(&myTaskData[i].requestQueue );
        }
    }
}

```

Using Multiprocessing Services

```

        if( theErr == noErr )
            theErr = MPPCreateQueue(&myTaskData[i].resultQueue );
        if( theErr == noErr )
            theErr = MPPCreateTask( MyTask, &myTaskData[i],
                                   kMPStackSize, notificationQueue,
                                   NULL, NULL, kMPTaskOptions,
                                   &myTaskData[i].taskID );
    }
}

/* If something went wrong, just go back to single processor mode */
if( theErr != noErr ) {
    StopMPTasks();
    numProcessors = 1;
}
}

```

The `sTaskData` structure defines a number of values to be used with the task, such as the task ID, the IDs of the message queues used with the task, and a pointer to parameters to pass to the task. A pointer to a structure of this type is passed in the function `MPPCreateTask` (page 59).

The variable `notificationQueueID` holds the ID of the notification queue to associate with the tasks. When a task terminates, it sends a message to this queue. After sending a termination request, the application typically polls this queue to determine when the task has actually terminated.

The `CreateMPTasks` function creates as many identical tasks as there are available processors (as stored in `numProcessors`). If for some reason the tasks cannot be created (for example, if Multiprocessing Services is not available), the variable `numProcessors` is set to 1 and the application should do the work of the tasks itself without making any Multiprocessing Services calls.

Before creating the tasks, `CreateMPTasks` calls the function `MPPCreateQueue` (page 67) to create a notification queue to be used by all the tasks. It then calls the Memory Manager function `NewPtrClear` to allocate memory for all the `myTaskData` structures required (in the case of this example, one per task).

Next, `CreateMPTasks` iterates over the number of requested tasks. For each iteration, it does the following:

Using Multiprocessing Services

- Makes two calls to `MPCreateQueue` (page 67) to create a request queue and a result queue for each task. The IDs for these queues are stored in the task's `myTaskData` structure.
- Fills out the `myTaskData` structure for that task as necessary.
- Calls the function `MPCreateTask` (page 59). When calling this function, you must specify the following values:
 - the entry point of the task and its input parameters
 - the size of the stack to associate with the task
 - the notification queue to associate with the task (by passing the ID of the queue obtained in the function `MPCreateQueue` (page 67).)

Each task is assigned its own unique ID, which is passed back in the `taskID` field of the `myTaskData` task structure.

Although not a requirement, you can assign a relative weight to each task by calling the function `MPSetTaskWeight` (page 65). The task weight is a value that indicates the amount of processor attention to give this task relative to all other eligible tasks. If, as in this example, you create a number of identical tasks, each would probably be given equal weight.

The sample task in Listing 3-2 calls one of two different functions depending on the request is placed on its queue.

Listing 3-2 A sample task

```
#define kMyRequestOne           1
#define kMyRequestTwo          2

#define kMyResultException     -1

OSStatus MyTask( void *parameter ) {

    OSErr theErr;
    sTaskDataPtr p;
    Boolean finished;
    UInt32 message;

    theErr = noErr;
```

Using Multiprocessing Services

```

    /* Get a pointer to this task's unique data */
    p = (sTaskDataPtr)parameter ;

    /* Process each request handed to the task and return a result */
    finished = false;
    while( !finished ) {
        theErr = MPWaitOnQueue( p->requestQueue, (void **)&message,
                               NULL, NULL, kDurationForever );
        if( theErr == noErr ) {
            /* Pick a function to call and pass in the parameters. */
            /* The parameters should be set up prior to sending the */
            /* message just received. Note that we could also just */
            /* pass in a pointer to the desired function instead of */
            /* using a selector. */
            /*
            switch( message ) {
                case kMyRequestOne:
                    theErr = fMyTaskFunctionOne( &p->params );
                    break;
                case kMyRequestTwo:
                    theErr = fMyTaskFunctionTwo( &p->params );
                    break;
                default:
                    finished = true;
                    theErr = kMyResultException;
            }
            MPNotifyQueue( p->resultQueue, (void *)theErr, NULL, NULL );
        }
        else
            finished = true;
    }

    /* Task is finished now */
    return( theErr );
}

```

This task takes one parameter, a pointer to its task data structure. This structure contains all the information that is needed for the life of the task, such as the request and result queues created for it, and any input necessary when processing a task request. The input parameters are passed along to the requested function.

Using Multiprocessing Services

After some initialization, the task sets the `finished` flag to `false` and then spends the rest of its time in a `while` loop processing message requests. The task calls the function `MPWaitOnQueue` (page 70), which waits indefinitely until a message appears on its request queue. In this case, the message indicates which function the task is to call. When a message is received, `MyTask` checks the request message to determine which function is desired and calls through to that function. Upon return, it posts a message on the result queue by calling `MPNotifyQueue` (page 69) and then calls `MPWaitOnQueue` again to wait for the next message.

Note that if you are creating tasks on-the-fly, you may want to have your task dispose of its task record (pointed to by `p`) upon completion of the task. For more information about allocating and disposing of memory in tasks, see “Allocating Memory in Tasks” (page 46).

Terminating Tasks

When you want to terminate a task, you should call the function `MPTerminateTask` (page 62). Doing so deletes the task, but you are still responsible for disposing of any memory you may have allocated for the task. In addition, because the tasks run asynchronously, the task may not actually terminate until sometime after the `MPTerminateTask` function returns. Therefore, you should not assume that the task has terminated until you have received a termination message from the notification queue you specified in the function `MPCreateTask` (page 59). Listing 3-3 shows how you might terminate the tasks created in Listing 3-1 (page 33).

Listing 3-3 Terminating tasks

```
void StopMPTasks(void)
{
    UInt32 i;

    if (myTaskData != NULL)
    {
        for (i = 0; i < numProcessors; i++)
        {
            if (myTaskData[i].TaskID != NULL)
```

Using Multiprocessing Services

```

    {
        MPTerminateTask(myTaskData[i].TaskID, noErr);
        MPWaitOnQueue(notificationQueue, NULL, NULL, NULL,
            kDurationForever);
    }

    if (myTaskData[i].fRequestQueue != NULL)
        MPDeleteQueue(myTaskData[i].RequestQueue);
    if (myTaskData[i].fResultQueue != NULL)
        MPDeleteQueue (myTaskData[i].ResultQueue);
}

if (notificationQueue != NULL)
{
    MPDeleteQueue (notificationQueue);
    notificationQueue = NULL;
}

DisposePtr((Ptr)myTaskData);
myTaskData = NULL;
}
}

```

The `StopMPTasks` function iterates through all the task data structures that were created in `CreateMPTasks` and checks for those with valid task IDs. It then calls the function `MPTerminateTask` (page 62) for each valid task ID.

After making the termination call, `StopMPTasks` then waits for a message to appear on the notification queue indicating that the task has in fact been terminated. It does so by waiting continuously on the notification queue until the termination message arrives. It then clears the task ID and disposes of the queues allocated for the task.

Note

If you call `MPWaitOnQueue` (page 70) from a cooperative task, you should specify only the `kDurationImmediate` wait time. You must use a `while` loop that continuously calls `MPWaitOnQueue` (page 70) until the termination message appears in the notification queue. Doing so also allows you to process events that may occur between calls. For additional information, see “Synchronizing and Notifying Tasks” (page 40). ♦

After terminating all the existing tasks, `StopMPTasks` then deletes the notification queue and disposes of the task data structures.

Synchronizing and Notifying Tasks

As described in “About Multitasking on the Mac OS” (page 13) tasks often need to coordinate with the main application or with other tasks to avoid data corruption or synchronization problems. To coordinate tasks, Multiprocessing Services provides three simple notification mechanisms (semaphores, event groups, and message queues), one complex one (kernel notifications), and critical regions.

Of the three simple notification mechanisms, message queues are the easiest to use, but they are the slowest. Typically a task has two message queues associated with it. It takes messages off an input queue, processes the information accordingly, and, when done, posts a message to an output queue.

IMPORTANT

You should never use only one instance of a notification mechanism to convey both input and output information, because doing so can easily cause confusion. For example, after posting a request, an application will at some point start waiting for results. If it waits on the same mechanism where the request was posted, the request itself may appear to be the result. The application may then clear the request in the mistaken belief that it was a result and no actual work gets done. ▲

Before notifying a task, your application should make sure that everything the task needs is in memory. That is, you should have created any necessary queues

Using Multiprocessing Services

and allocated space for any data the task may require. For each task, your application establishes the parameters of the work that it wants the task to perform and then it must signal the task through either a queue or a semaphore to begin performing that work. The specific work that the task is to perform can be completely defined within a message, or possibly within a block of memory reserved for that task. You can also pass in a pointer to the function that the task should call to perform the work. Doing so allows one task to perform many different types of chores.

Listing 3-4 shows a function that divides up a large amount of data among multiple tasks, placing requests on each task's request queue and waiting for the results.

Listing 3-4 Assigning work to tasks

```

OSErr NotifyTasks( UInt32 realFirstThing, UInt32 realTotalThings ) {

    UInt32 i;
    OSErr theErr;
    UInt32 thingsPerTask;
    UInt32 message;
    sWorkParams appData;

    theErr = noErr;

    thingsPerTask = realTotalThings / numProcessors;

    /* Start each task working on a unique piece of the total data */
    for( i = 0; i < numProcessors; i++ ) {
        myTaskData[i].params.firstThing =
            realFirstThing + thingsPerTask * i;
        myTaskData[i].params.totalThings = thingsPerTask;
        message = kMyRequestOne;
        MPNotifyQueue( myTaskData[i].requestQueue, (void *)message,
            NULL, NULL );
    }

    /* Now wait for the tasks to finish */
    for( i = 0; i < numProcessors; i++ )
        MPWaitOnQueue( myTaskData[i].resultQueue, (void **)&message,

```

Using Multiprocessing Services

```

        NULL, NULL, kDurationForever );

    return( theErr );
}

```

For each task, it calls `MPNotifyQueue` (page 69) to place the pointer to the task's portion of the data on the task's request queue. It then calls `MPWaitOnQueue` (page 70) to wait for confirmation that the task has completed.

Note

A message queue message is passed to the queue as three 32-bit parameters. Because the message in Listing 3-4 is only 32-bits long, the remaining two parameters are set to `NULL`. ♦

If you want to use semaphores or event groups instead of message queues, you would call the following functions to set up, notify, and wait on them, in a manner similar to that shown in Listing 3-4:

- `MPCreateSemaphore` (page 73), `MPSignalSemaphore` (page 75), `MPWaitOnSemaphore` (page 75), and `MPDeleteSemaphore` (page 74) for semaphores
- `MPCreateEvent` (page 77), `MPSetEvent` (page 78), `MPWaitForEvent` (page 79), and `MPDeleteEvent` (page 77) for event groups

However, if you use the simpler notification mechanisms, you have to find another way to pass the function pointer to the task. One possibility is to assign the pointer to a field in the task's task data structure.

To use kernel notifications, you should call the following functions:

- `MPCreateNotification` (page 80), `MPCauseNotification` (page 83), `MPModifyNotification` (page 82), and `MPDeleteNotification` (page 81)

There is no function for waiting on a kernel notification, as the task would wait on the appropriate subcomponents of the kernel notification (for example, a semaphore and a message queue).

Note that the example in Listing 3-4 will wait forever (`kDurationForever`) for a message to appear on its result queue. While this method is fine if called from a preemptive task, it can cause problems if called from a cooperative task. If the task takes a significant amount of time to execute, the calling task "hangs" for that time, since it can't call `WaitNextEvent` to give other applications processor time. If you want to wait on a task from a cooperative task, your application

Using Multiprocessing Services

should post the message and then return to its event loop. From within the event loop it can then poll the result queue using `kDurationImmediate` waits until a message appears.

If you specify `kDurationImmediate` for the waiting time for either `MPWaitOnQueue` (page 70), `MPWaitOnSemaphore` (page 75), `MPWaitForEvent` (page 79), or `MPEnterCriticalRegion` (page 85), the function always returns immediately. If the return value is `KMPTimeoutErr`, then the task generated no new results since the last time the application checked. That is, no message was available, the semaphore was zero, or the critical region was being executed by another processor. If the value is `noErr`, a result was present and obtained by the call.

Handling Periodic Actions

You can use notification mechanisms to do more than simply signal tasks. For example, Listing 3-5 shows a task that uses a semaphore to do periodic actions.

Listing 3-5 Using a semaphore to perform periodic actions

```
void MyTask(void) {
    MPSemaphoreID delay;

    MPCreateSemaphore(1, 0, &delay); // a binary semaphore
    while(true)
    {
        DoIt(); // do something interesting
        (void) MPWaitOnSemaphore(delay, 10 * kDurationMillisecond);
    }
}
```

This example uses a semaphore solely to create a delay. After each call to the `DoIt` function, `MyTask` waits for a notification that never arrives and times out after 10ms.

You can combine the delaying and notification aspects of a semaphore to add more flexibility as shown in Listing 3-6.

Listing 3-6 Performing actions periodically and on demand

```

main(void) {

    MPSemaphoreID delay;
    ...
    MPCreateSemaphore(2, 0, &delay);
    MPCreateTask(...);

    while(true) {
        // Event loop.

        if ( /* something important happened */ )
        {
            MPSignalSemaphore(delay);
        }
    }
}

void MyTask(void) {

    while(true) {
        DoIt(); // Do interesting things.
        (void) MPWaitOnSemaphore(work, 100 * kDurationMillisecond);
    }
}

```

In this example, the `MyTask` task runs essentially as before, except that the main application creates the semaphore. If no signal is sent to the semaphore, the `DoIt` function in `MyTask` executes every 100ms. However, in this example the application can signal the semaphore, which unblocks the task and allows the `DoIt` function to execute. That is, the `DoIt` function executes whenever the application signals the semaphore, or every 100ms otherwise.

Notifying Tasks at Interrupt Time

If you want to send a notification to a task from a 68K-style interrupt handler, you can do so using the functions `MPSignalSemaphore` (page 75), `MPSetEvent` (page 78), or `MPNotifyQueue` (page 69). The `MPSignalSemaphore` and `MPSetEvent`

Using Multiprocessing Services

functions are always interrupt-safe, while the `MPNotifyQueue` (page 69) function becomes interrupt-safe if you reserve notifications on the message queue. See the `MPSetQueueReserve` (page 72) function description for more information about reserving notifications.

▲ **WARNING**

Aside from these three notification functions, only `MPCurrentTaskID` (page 63) and `MPYield` (page 65) are interrupt-safe; attempting to call other Multiprocessing Services functions at interrupt time, or at a deferred-task time, may cause a system crash. ▲

Using Critical Regions

If your tasks need access to code that is nonreentrant, (that is, only one task can be executing the code at any particular time), you must designate that code as being a critical region. You do so by calling the function `MPCreateCriticalRegion` (page 84). Doing so returns a critical region ID that you use to identify the region when you want to enter or exit it later. To enter a critical region, the task must call `MPEnterCriticalRegion` (page 85) and specify the ID of the region to enter. This function acts much like the functions that wait on message queues and semaphores; if the critical region is not currently available, the task can wait for a specified time for it to become available (after which it will time out).

After the task has completed using the critical region, you must call `MPExitCriticalRegion` (page 86). Doing so “frees” the critical region so that another task that is waiting on it can enter. Note that a task can call `MPEnterCriticalRegion` multiple times during execution (as in a recursive call) as long as it balances each such call with `MPExitCriticalRegion` when it leaves the critical region.

Note that the area of code designated as a critical region is not “tagged” as such in any way. You must make sure that your code is synchronized to properly isolate the critical region. For example, if you have a critical region that will be shared by two different tasks, you must create the critical region outside the tasks that will require it and pass the critical region ID to the tasks. This method ensures that, even if multiple instances of a task were created and running, only one could access a particular critical region at a time.

If a task contains more than one critical region, each critical region must have its own unique ID; otherwise, a task entering a critical region may block another task from entering a different critical region.

Allocating Memory in Tasks

If you need to allocate a block of memory for a task, you must call the function `MPAllocateAligned` (page 99). Doing so returns a pointer to allocated memory with the alignment you specify. You should always use the Multiprocessing Services memory allocation functions if your task needs to allocate, deallocate, or otherwise manipulate memory. For example, if your task deallocates its task data structure after it has finished processing, it must call `MPFree` (page 100). Note however, that since the memory is being deallocated by a preemptive task, you must have initially allocated the task record by calling `MPAllocateAligned`, even if this allocation didn't occur in a preemptive task.

Using Task-Specific Storage

Task-specific storage is useful for storing small pieces of data, such as pointers to task-specific information. For example, if you create several identical tasks, each of which requires some unique data, you can store that data as task-specific storage. Task-specific storage locations are cross-referenced by an index value and the task ID, so the same code can easily refer to “per-instance” variables. Each such storage location holds a 32-bit value.

Task-specific storage is automatically allocated when a task is created; the amount is fixed and cannot change for the life of the task. To access the task-specific storage, you call the function `MPAllocateTaskStorageIndex` (page 96). Doing so returns an index number which references a storage location in each available task in the process. Subsequent calls to `MPAllocateTaskStorageIndex` return new task index values to access more of the task-specific storage. Note that, aside from the fact that each index value is unique, you should not assume anything about the actual values of the index. For example, you cannot assume that successive calls to `MPAllocateTaskStorageIndex` will monotonically increase the index value.

Since the amount of task-specific storage is fixed, you may use up the available storage (and corresponding index values) if you make many `MPAllocateTaskStorageIndex` calls. In such cases, further calls to `MPAllocateTaskStorageIndex` return an error indicating insufficient resources.

Using Multiprocessing Services

You call `MPSetTaskStorageValue` (page 97) and `MPGetTaskStorageValue` (page 98) to set and retrieve the storage data. After you are finished using the storage locations, you must call `MPDeallocateTaskStorageIndex` (page 96) to free the index.

Using Timers

On occasion you may want to use timers in your preemptive tasks. For example, say you want a task to send a message to a given queue every 20 milliseconds. To do so, you can set a timer to block your task for 20ms after sending the notification by calling the function `MPDelayUntil` (page 87).

Note

Note that in some cases you may want to use notification mechanisms to accomplish periodic actions, as described in “Handling Periodic Actions” (page 43). ♦

In addition, you can create timers that will signal a specified notification mechanism after the timer expires. For example, say you have a task that is prompting the user to enter a name and password. Once you bring up the input dialog box, you may have another task (or the application) create a timer object to expire after five minutes. If the user has not entered a password during those five minutes, the timer expires and sends a message to the task, signaling that it should terminate.

You create a timer using the function `MPCreateTimer` (page 88) and arm it by calling the function `MPArmTimer` (page 91). To specify the notification mechanisms to signal when the timer expires, you call the function `MPSetTimerNotify` (page 90). Note that you can signal one notification mechanism of each type if desired. For example, the timer can send a message to a queue and also set a bit in an event group when it expires.

The timers in Multiprocessing Services use time units of type `AbsoluteTime`, which increases monotonically since the host computer was started up. You can obtain the time since startup by calling the function `UpTime` (page 93). Multiprocessing Services also provides the functions `DurationToAbsolute` (page 94) and `AbsoluteToDuration` (page 95) which let you convert time between units of `AbsoluteTime` and units of type `Duration`. Note that you should not make any assumptions about what the `AbsoluteTime` units are based upon.

Making Remote Procedure Calls

At times a preemptive task may need to call a system software function, and doing so may cause problems. For example, many calls to Mac OS system software manipulate global variables, so data could easily be corrupted if more than one task attempts to make similar calls. To work around this problem, Multiprocessing Services allows you to make **remote procedure calls** if you need to call system software from a preemptive task. A remote procedure call also allows your task to call 68K code.

IMPORTANT

A subset of system software calls, as described in Appendix A, are preemptive task-safe in Mac OS 9. However, with the exception of the functions in Multiprocessing Services, you should assume that all other system calls are unsafe. Even if some system software function appears to work today when called from a preemptive task, unless explicitly stated otherwise there is no guarantee that subsequent versions of the same routine will continue to work in future versions of system software. In the Mac OS 8 and Mac OS 9 implementations of Multiprocessing Services, the only exceptions to this rule are the atomic memory operations (such as `AddAtomic`) exported in the `InterfaceLib` shared library. Even these functions may switch to 68K mode if the operands to them are not properly aligned. ▲

To make a remote procedure call, you must designate an application-defined function that will make the actual calls to system software. You then pass a pointer to this function as well as any required parameters in the `MPRemoteCall` (page 111) function.

Note

Since your application-defined function must be written in PowerPC code, you do not need to build a universal procedure pointer to pass to the `MPRemoteCall` function. ◆

When you call the function `MPRemoteCall` from a task, that task is blocked, and the application-defined function you designated then executes as a

cooperatively scheduled task, which can make system software calls with no danger.

Note that when you call `MPRemoteCall`, you can designate which context (or process) you want your application-defined function to execute in. If you specify that the function should execute in the same context that owns the task, the function has access to data available to the main application (just as if the application had called the function). However, the function cannot execute until the owning context becomes active (and then not until the application calls `WaitNextEvent`). Alternatively, you can designate that the function execute in any available context. Doing so minimizes possible lag time, but the function cannot access any resources specific to the task's context.

IMPORTANT

In the future, individual application processes may not always share the same address space, so in general you should never attempt to access code or data in another process. ▲

After your application-defined function returns, the task is unblocked and execution proceeds normally.

Handling Exceptions and Debugging

Multiprocessing Services provides a number of functions you can use to handle exceptions and to aid in debugging.

By default, if you do not register an exception handler, and no debuggers are registered, a task terminates when it takes an exception. If debuggers or exception handlers exist, then the task is suspended when an exception occurs and a message is sent to the appropriate debugger or handler.

If desired, you can install an exception handler for a task by calling the function `MPSetExceptionHandler` (page 104). When an exception occurs, a message is sent to a designated queue, which your exception handler can wait upon.

In addition, you can register one or more debuggers with Multiprocessing Services by calling the function `MPRegisterDebugger` (page 109). When calling `MPRegisterDebugger`, you must specify the queue to which you want the exception message to be sent as well as a debugger level. The debugger level is simply an integer that indicates where to place this debugger in the hierarchy of

Using Multiprocessing Services

registered debuggers. In addition, When an exception occurs, the order of notification for handlers is as follows:

- The debugger with the highest debugger level (for example, a debugger registered at level 3 will have precedence over one registered as level 2).
- The debugger with the next highest level (and so on, for all the registered debuggers).
- The task's exception handler.
- The task's termination function.

At each level, the handler can choose to do either of the following:

- Set or retrieve the task's register or stack information using `MPSetTaskState` (page 108) or `MPExtractTaskState` (page 107).
- Call `MPDisposeTaskException` (page 106), which allows you to do any of the following:
 - Resume the task.
 - Resume the task and enable single-stepping or branch-stepping.
 - Propagate the exception to the next lower level. For example, instead of handling the exception itself, a debugger can pass the exception message to the next debugger (or exception handler) in the hierarchy.

If you want to throw an exception to a task, you can use the `MPTThrowException` (page 105) function.

Multiprocessing Services Reference

Contents

Gestalt Constants	55
Functions	56
Determining Processor Availability	57
MPProcessors	57
MPProcessorsScheduled	57
MPGetNextCpuID	58
Creating and Scheduling Tasks	59
MPCreateTask	59
MPTerminateTask	62
MPExit	63
MPCurrentTaskID	63
MPGetNextTaskID	64
MPYield	65
MPSetTaskWeight	65
MPTaskIsPreemptive	66
Creating and Handling Message Queues	67
MPCreateQueue	67
MPDeleteQueue	68
MPNotifyQueue	69
MPWaitOnQueue	70
MPSetQueueReserve	72
Creating and Handling Semaphores	73
MPCreateSemaphore	73
MPDeleteSemaphore	74
MPSignalSemaphore	75
MPWaitOnSemaphore	75
Handling Event Groups	76

MPCreateEvent	77
MPDeleteEvent	77
MPSetEvent	78
MPWaitForEvent	79
Handling Kernel Notifications	80
MPCreateNotification	80
MPDeleteNotification	81
MPModifyNotification	82
MPCauseNotification	83
Handling Critical Regions	83
MPCreateCriticalRegion	84
MPDeleteCriticalRegion	84
MPEnterCriticalRegion	85
MPExitCriticalRegion	86
Timer Services Functions	87
MPDelayUntil	87
MPCreateTimer	88
MPDeleteTimer	89
MPSetTimerNotify	90
MPArmTimer	91
MPCancelTimer	92
Time Utility Functions	93
UpTime	93
DurationToAbsolute	94
AbsoluteToDuration	95
Accessing Per-Task Storage Variables	95
MPAllocateTaskStorageIndex	96
MPDeallocateTaskStorageIndex	96
MPSetTaskStorageValue	97
MPGetTaskStorageValue	98
Memory Allocation Functions	99
MPAllocateAligned	99
MPFree	100
MPGetAllocatedBlockSize	101
MPBlockCopy	101
MPBlockClear	102
MPDataToCode	103
Exception Handling Functions	103

MPSetExceptionHandler	104
MPThrowException	105
MPDisposeTaskException	106
MPExtractTaskState	107
MPSetTaskState	108
Debugger Support Functions	109
MPRegisterDebugger	109
MPUnregisterDebugger	110
Remote Calling Function	111
MPRemoteCall	111
Application-Defined Function	112
MyRemoteProcedure	113
Data Types	113
MPProcessID	114
MPCpuID	115
MPCoherenceID	115
MPTaskID	115
TaskProc	116
MPTaskOptions	116
MPTaskWeight	117
MPTaskInfo	117
MPQueueID	118
MPSemaphoreID	119
MPEventID	119
MPCriticalRegionID	120
MPNotificationID	120
MPTimerID	120
MPOpaqueID	121
TaskStorageIndex	121
TaskStorageValue	121
MPSemaphoreCount	122
MPEventFlags	122
MPExceptionKind	122
MPDebuggerLevel	123
MPRemoteProcedure	123
Constants	123
Timer Duration Constants	124
Timer Option Masks	125

Memory Allocation Alignment Constants	126
Memory Allocation Option Constants	127
Task Creation Options	128
Task State Constants	129
Task Run State Constants	130
Task Information Structure Version Constant	131
Task Exception Disposal Constants	131
Remote Call Context Option Constants	132
Result Codes	133

Multiprocessing Services Reference

This chapter describes the Multiprocessing Services application programming interface (API) available with Mac OS 9.0. This chapter contains the following sections:

- “Gestalt Constants” (page 55)
- “Functions” (page 56)
- “Application-Defined Function” (page 112)
- “Data Types” (page 113)
- “Constants” (page 123)
- “Result Codes” (page 133)

Note

This document describes version 2.1 of Multiprocessing Services. For a list of functions changed or added between versions 1.4 and 2.1, see Appendix C. ♦

Gestalt Constants

Before calling any Mac OS system software functions from a preemptive task, you should call the `Gestalt` function with the `gestaltMPCallableAPIsAttr` selector set to determine which preemptively safe system calls are allowed.

```
enum {
    gestaltMPCallableAPIsAttr = 'mpsc'),
    gestaltMPFileManager      = 0,
    gestaltMPDeviceManager    = 1
};
```

Constant Descriptions

`gestaltMPCallableAPIsAttr`

The Gestalt selector passed to determine the availability of preemptive system software functions. The `Gestalt` function produces a 32-bit value that you should test to determine which what type of preemptive calls are allowed.

Multiprocessing Services Reference

`gestaltMPFileManager`

If this bit is set, you can call the preemptively safe File Manager functions listed in Appendix A.

`gestaltMPDeviceManager`

If this bit is set, you can call the preemptively safe Device Manager functions listed in Appendix A.

Note that for functions that are shared between managers (for example, `PBCloseSync`), you should check the bit that is appropriate for the manager you want to call.

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

Functions

This section describes Multiprocessing Services functions in the following categories:

- “Determining Processor Availability” (page 57)
- “Creating and Scheduling Tasks” (page 59)
- “Creating and Handling Message Queues” (page 67)
- “Creating and Handling Semaphores” (page 73)
- “Handling Event Groups” (page 76)
- “Handling Critical Regions” (page 83)
- “Timer Services Functions” (page 87)
- “Time Utility Functions” (page 93)
- “Accessing Per-Task Storage Variables” (page 95)
- “Memory Allocation Functions” (page 99)
- “Exception Handling Functions” (page 103)
- “Debugger Support Functions” (page 109)
- “Remote Calling Function” (page 111)

Determining Processor Availability

Multiprocessing Services provides the following functions for determining the number of processors available on the host computer:

- `MPProcessors` (page 57) returns the number of processors on the host computer.
- `MPProcessorsScheduled` (page 57) returns the number of active processors available on the host computer.
- `MPGetNextCpuID` (page 58) obtains the next CPU ID in the list of physical processors.

MPProcessors

Returns the number of processors on the host computer.

```
ItemCount MPProcessors (void);
```

function result A value of type `ItemCount` that indicates the number of physical processors on the host computer.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPProcessorsScheduled` (page 57).

MPProcessorsScheduled

Returns the number of active processors available on the host computer.

```
ItemCount MPProcessorsScheduled (void);
```

Multiprocessing Services Reference

function result A value of type `ItemCount` that indicates the number of active processors available on the host computer.

DISCUSSION

The number of active processors is defined as the number of processors scheduled to run tasks. This number varies while the system is running. Advanced power management facilities may stop or start scheduling processors in the system to control power consumption or to maintain a proper operating temperature.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPProcessors` (page 57).

MPGetNextCpuID

Obtains the next CPU ID in the list of physical processors of the specified memory coherence group.

```
OSStatus MPGetNextCpuID (
    MPCoherenceID owningCoherenceID,
    MPCpuID * cpuID);
```

`owningCoherenceID`

The ID of the memory coherence group whose physical processor IDs you want to obtain. Pass `kInvalidID`, as only one coherence group, internal RAM, is currently defined.

`cpuID`

A pointer to a variable of type `MPCpuID` (page 115). On return, `cpuID` points to the ID of the next physical processor.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DESCRIPTION

By iterating on this function (after calling `MPPProcessors` (page 57), for example), you can obtain the IDs of all the processors available on the host computer. Generally you would only use this function in diagnostic programs.

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

Creating and Scheduling Tasks

Multiprocessing Services provides the following functions for creating and manipulating preemptive tasks:

- `MPCreateTask` (page 59) creates a preemptive task.
- `MPTerminateTask` (page 62) terminates an existing task.
- `MPExit` (page 63) allows a task to terminate itself.
- `MPCurrentTaskID` (page 63) obtains the task ID of the current preemptive task.
- `MPGetNextTaskID` (page 64) obtains the next task ID in the list of available tasks.
- `MPYield` (page 65) allows a task to yield the processor to another task.
- `MPSetTaskWeight` (page 65) assigns a relative weight to a task, indicating how much processor time it should receive compared to other available tasks.
- `MPTaskIsPreemptive` (page 66) determines whether a task is preemptively scheduled.

MPCreateTask

Creates a preemptive task.

```
OSStatus MPCreateTask (
    TaskProc entryPoint,
    void *parameter,
    ByteCount stackSize,
```

Multiprocessing Services Reference

```

    MPQueueID notifyQueue,
    void *terminationParameter1,
    void *terminationParameter2,
    MPTaskOptions options,
    MPTaskID *task);

```

<code>entryPoint</code>	A pointer of type <code>TaskProc</code> (page 116) that references the task function. The task function should take a single 32-bit parameter and return a value of type <code>OSStatus</code> .
<code>parameter</code>	The parameter to pass to the task function.
<code>stackSize</code>	A value of type <code>ByteCount</code> that specifies the size of the stack assigned to the task. Note that you should be careful not to exceed the bounds of the stack, since stack overflows may not be detected. Specifying zero for the size will result in a default stack size of 4KB.
<code>notifyQueue</code>	A value of type <code>MPQueueID</code> (page 118) that specifies the ID of the message queue to which the system will send a message when the task terminates. You specify the first 64-bits of the message in the parameters <code>terminationParameter1</code> and <code>terminationParameter2</code> respectively. The last 32-bits contain the result code of the task function.
<code>terminationParameter1</code>	A 32-bit value that is sent to the message queue specified by the parameter <code>notifyQueue</code> when the task terminates.
<code>terminationParameter2</code>	A 32-bit value that is sent to the message queue specified by the parameter <code>notifyQueue</code> when the task terminates.
<code>options</code>	A value of type <code>MPTaskOptions</code> (page 116) that specifies optional attributes of the preemptive task. See “Task Creation Options” (page 128) for a list of possible values.
<code>task</code>	A pointer to a variable of type <code>MPTaskID</code> (page 115). On return, the variable contains the ID of the newly created task.
<i>function result</i>	A result code. See “Result Codes” (page 133) for a list of possible values. If <code>MPCreateTask</code> could not create the task because some critical resource was not available, the function returns <code>kMPInsufficientResourcesErr</code> . Usually this is due to lack of

memory to allocate the internal data structures associated with the task or the stack. The function also returns `kMPIInsufficientResourcesErr` if any reserved option bits are set.

DISCUSSION

Tasks are created in the unblocked state, ready for execution. A task can terminate in the following ways:

- By returning from its entry point
- By calling `MPExit` (page 63)
- When specified as the target of an `MPTerminateTask` (page 62) call
- If a hardware-detected exception or programming exception occurs and no exception handler is installed
- If the application calls `ExitToShell`

Task resources (its stack, active timers, internal structures related to the task, and so on) are reclaimed by the system when the task terminates. The task's address space is inherited from the process address space. All existing tasks are terminated when the owning process terminates.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPTerminateTask` (page 62).

The function `MPSetTaskWeight` (page 65).

MPTerminateTask

Terminates an existing task.

```
OSStatus MPTerminateTask (
    MPTaskID task,
    OSStatus terminationStatus);
```

task A value of type `MPTaskID` (page 115) that specifies the ID of the task you wish to terminate.

terminationStatus A 32-bit value of type `OSStatus` indicating termination status. This value is sent to the termination status message queue you specified in `MPCreateTask` (page 59) in place of the task function's result code.

function result A result code. See "Result Codes" (page 133) for a list of possible values. If the task to be terminated is already in the process of termination, `MPTerminateTask` returns `kMPInsufficientResourcesErr`. You do not need to take any additional action if this occurs.

DISCUSSION

You should not assume that the task has completed termination when this call returns; the proper way to synchronize with task termination is to wait on the termination queue (specified in `MPCreateTask` (page 59)) until a message appears. Because task termination is a multistage activity, it is possible for a preemptive task to attempt to terminate a task that is already undergoing termination. In such cases, `MPTerminateTask` returns `kMPInsufficientResourcesErr`.

Note that Multiprocessing Services resources (event groups, queues, semaphores, and critical regions) owned by a preemptive task are not released when that task terminates. If a task has a critical region locked when it terminates, the critical region remains in the locked state. Multiprocessing Services resources no longer needed should be explicitly deleted by the task that handles the termination message. All Multiprocessing Services resources created by tasks are released when their owning process (that is, the host application) terminates.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

MPExit

Allows a task to terminate itself

```
void MPExit (OSStatus status);
```

status An application-defined value of type `OSStatus` that indicates termination status. This value is sent to the termination message queue in place of the task's result code.

DISCUSSION

When called from within a preemptive task, the task terminates, and the value indicated by the parameter `status` is sent to the termination message queue you specified in `MPCreateTask` (page 59). Note that you cannot call `MPExit` from outside a preemptive task.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

MPCurrentTaskID

Obtains the task ID of the currently-executing preemptive task

```
MPTaskID MPCurrentTaskID (void);
```

function result A value of type `MPTaskID` (page 115) that specifies the task ID of the current preemptive task.

DISCUSSION

Returns the ID of the current preemptive task. If called from a cooperative task, this function returns an ID which is different than the ID of any preemptive task. Nonpreemptive processes may or may not have different task IDs for each application; future implementations of this API may behave differently in this regard.

Note that you can call this function from an interrupt handler.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

MPGetNextTaskID

Obtains the next task ID in the list of available tasks.

```
OSStatus MPGetNextTaskID (
    MPProcessID owningProcessID,
    MPTaskID * taskID);
```

owningProcessID

A value of type `MPProcessID` (page 114) that indicates the ID of the process (typically the application) that owns the tasks. This ID is the same as the process ID handled by the Code Fragment Manager.

taskID

A pointer to variable of type `MPTaskID` (page 115). On return, `taskID` points to ID of the next task in the list of tasks.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DESCRIPTION

By iterating on this function, you can obtain the IDs of all the tasks in a given process. These tasks may be running, ready, or blocked. Generally you would only use this function in diagnostic programs.

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

MPYield

Allows a task to yield the processor to another task.

```
void MPYield (void);
```

DISCUSSION

This function indicates to the scheduler that another task can run. Other than possibly yielding the processor to another task or application, the call has no effect. Note that since tasks are preemptively scheduled, an implicit yield may occur at any point, whether or not this function is called.

In most cases you should not need to call this function. The most common use of `MPYield` is to release the processor when a task is in a loop in which further progress is dependent on other tasks, and the task cannot be blocked by waiting on a Multiprocessing Services resource. You should avoid such busy waiting whenever possible.

Note that you can call this function from an interrupt handler.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

MPSetTaskWeight

Assigns a relative weight to a task, indicating how much processor time it should receive compared to other available tasks.

```
OSStatus MPSetTaskWeight (
    MPTaskID task,
    MPTaskWeight weight);
```

Multiprocessing Services Reference

<i>task</i>	A value of type <code>MPTaskID</code> (page 115) that specifies the ID of the task to which you want to assign a weighting.
<i>weight</i>	A value of type <code>MPTaskWeight</code> (page 117) indicating the relative weight to assign. This value can range from 1 to 10,000, with the default value being 100.
<i>function result</i>	A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

The approximate processor share is defined as

$$\text{weight of the task} / \text{total weight of available tasks}$$

For a set of ready tasks, the amount of CPU time dedicated to the tasks will be determined by the dynamically computed share. Note that the processor share devoted to tasks may deviate from the suggested weighting if critical tasks require attention. For example, a real-time task (such as a QuickTime movie) may require more than its relative weight of processor time, and the scheduler will adjust proportions accordingly.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPTaskIsPreemptive

Determines whether a task is preemptively scheduled.

```
Boolean MPTaskIsPreemptive (MPTaskID taskID);
```

<i>taskID</i>	A value of type <code>MPTaskID</code> (page 115) that specifies the task you want to check. Pass <code>kInvalidID</code> if you want to specify the current task.
---------------	---

<i>function result</i>	A value of type <code>Boolean</code> . If true, the task is preemptively scheduled. If false, the task is cooperatively scheduled.
------------------------	--

DISCUSSION

If you have code that may be called from either cooperative or preemptive tasks, that code can call `MPTaskIsPreemptive` if its actions should differ depending on its execution environment.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Creating and Handling Message Queues

Multiprocessing Services provides the following functions for creating and handling message queues:

- `MPCreateQueue` (page 67) creates a message queue.
- `MPDeleteQueue` (page 68) deletes a message queue.
- `MPNotifyQueue` (page 69) sends a message to the specified message queue.
- `MPWaitOnQueue` (page 70) obtains a message from a specified message queue.
- `MPSetQueueReserve` (page 72) reserves space for messages on a specified message queue.

MPCreateQueue

Creates a message queue.

```
OSStatus MPCreateQueue (MPQueueID *queue);
```

queue A pointer to a variable of type `MPQueueID` (page 118). On return, the variable contains the ID of the newly created message queue.

function result A result code. See “Result Codes” (page 133) for a list of possible values. If a queue could not be created, `MPCreateQueue` returns `KMPInsufficientResourcesErr`.

DISCUSSION

This call creates a message queue, which can be used to notify (that is, send) and wait for (that is, receive) messages consisting of three 32-bit words in a preemptively safe manner.

Message queues are created from dynamically allocated internal resources. Other tasks may be competing for these resources so it is possible this function may not be able to create a queue.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPDeleteQueue` (page 68).

The function `MPSetQueueReserve` (page 72).

MPDeleteQueue

Deletes a message queue.

```
OSStatus MPDeleteQueue (MPQueueID queue);
```

queue A value of type `MPQueueID` (page 118) that specifies the ID of the message queue you want to delete.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

After calling `MPDeleteQueue`, the specified queue ID becomes invalid, and all internal resources associated with the queue (including queued messages) are reclaimed. Any tasks waiting on the queue are unblocked and their respective `MPWaitOnQueue` (page 70) calls will return with the result code `kMPDeletedErr`.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPCreateQueue` (page 67).

MPNotifyQueue

Sends a message to the specified message queue.

```
OSStatus MPNotifyQueue (
    MPQueueID queue,
    void *param1,
    void *param2,
    void *param3);
```

<code>queue</code>	A value of type <code>MPQueueID</code> (page 118) that specifies the queue ID of the message queue you want to notify.
<code>param1</code>	The first 32-bits of the message to send.
<code>param2</code>	The second 32-bits of the message to send.
<code>param3</code>	The third 32-bits of the message to send.
<i>function result</i>	A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

This function sends a message to the specified queue, which consist of the three parameters, `param1`, `param2`, and `param3`. The system does not interpret the three 32-bit words which comprise the text of the message. If tasks are waiting on the specified queue, the first waiting task is unblocked and the task’s `MPWaitOnQueue` (page 70) function completes.

Depending on the queue mode, the system either allocates messages dynamically or assigns them to memory reserved for the queue. In either case,

Multiprocessing Services Reference

if no more memory is available for messages `MPNotifyQueue` returns `kMPInsufficientResourcesErr`.

You can call this function from an interrupt handler if messages are reserved on the queue. For more information about queueing modes and reserving messages, see `MPSetQueueReserve` (page 72).

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPWaitOnQueue` (page 70).

MPWaitOnQueue

Obtains a message from a specified message queue.

```
OSStatus MPWaitOnQueue (
    MPQueueID queue,
    void **param1,
    void **param2,
    void **param3,
    Duration timeout);
```

<code>queue</code>	A value of type <code>MPQueueID</code> (page 118) that specifies the ID of the message queue from which to receive the notification.
<code>param1</code>	A pointer to a 32-bit variable. On return, this variable contains the first 32-bits of the notification message. Pass <code>NULL</code> if you do not need this portion of the message.
<code>param2</code>	A pointer to a 32-bit variable. On return, the variable contains the second 32-bits of the notification message. Pass <code>NULL</code> if you do not need this portion of the message.
<code>param3</code>	A pointer to a 32-bit variable. On return, the variable contains the third 32-bits of the notification message. Pass <code>NULL</code> if you do not need this portion of the message.

Multiprocessing Services Reference

- timeout* A value of type `Duration` specifying how long to wait for a notification before timing out. See “Timer Duration Constants” (page 124) for a list of constants you can use to specify the wait interval.
- function result* A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

This function receives a message from the specified message queue. If no messages are currently available, the *timeout* specifies how long the function should wait for one. Tasks waiting on the queue are handled in a first in, first out manner; that is, the first task to wait on the queue receives the message from the `MPNotifyQueue` (page 69) call.

After calling this function, when a message appears, it is removed from the queue and the three fields, `param1`, `param2`, and `param3` are set to the values specified by the message text. Note these parameters are pointers to variables to be set with the message text.

If you call this function from a cooperative task, you should specify only `kDurationImmediate` for the *timeout* length; other waits will cause the task to block.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPNotifyQueue` (page 69).

MPSetQueueReserve

Reserves space for messages on a specified message queue.

```
OSStatus MPSetQueueReserve (
    MPQueueID queue,
    ItemCount count);
```

queue A value of type `MPQueueID` (page 118) that specifies the ID of the queue whose messages you want to reserve.

count A value of type `ItemCount` that specifies the number of messages to reserve.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

`MPNotifyQueue` (page 69) allocates spaces for messages dynamically; that is, memory to hold the message is allocated for the queue at the time of the call. In most cases this method is both speed and storage efficient. However, it is possible that, due to lack of memory resources, space for the message may not be available at the time of the call; in such cases, `MPNotifyQueue` (page 69) will return `kInsufficientResourcesErr`.

If you must have guaranteed message delivery, or if you need to call `MPNotifyQueue` (page 69) from an interrupt handler, you should reserve space on the specified queue by calling `MPSetQueueReserve`. Because such allocated space is reserved for duration of the queue’s existence, you should avoid straining internal system resources by reserving messages only when absolutely necessary. Note that if you have reserved messages on a queue, additional space *cannot* be added dynamically if the number of messages exceeds the number reserved for that queue.

The number of reserved messages is set to `count`, lowering or increasing the current number of reserved messages as required. If `count` is set to zero, no messages are reserved for the queue, and space for messages is allocated dynamically.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Creating and Handling Semaphores

Multiprocessing Services provides the following functions for creating and handling semaphores:

- `MPCreateSemaphore` (page 73) creates a semaphore.
- `MPDeleteSemaphore` (page 74) removes a semaphore.
- `MPSignalSemaphore` (page 75) signals a semaphore.
- `MPWaitOnSemaphore` (page 75) waits on a semaphore.

MPCreateSemaphore

Creates a semaphore.

```
OSStatus MPCreateSemaphore (
    MPSemaphoreCount maximumValue,
    MPSemaphoreCount initialValue,
    MPSemaphoreID *semaphore);
```

`maximumValue` A value of type `MPSemaphoreCount` (page 122) that specifies the maximum allowed value of the semaphore.

`initialValue` A value of type `MPSemaphoreCount` (page 122) that specifies the initial value of the semaphore.

`semaphore` A pointer to a variable of type `MPSemaphoreID` (page 119). On return, the variable contains the ID of the newly-created semaphore.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

If you want to create a binary semaphore, you can call the macro `MPCreateBinarySemaphore (MPSemaphoreID *semaphore)` instead, which simply calls `MPCreateSemaphore` with both `maximumValue` and `initialValue` set to 1.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPDeleteSemaphore` (page 74).

MPDeleteSemaphore

Removes a semaphore.

```
OSStatus MPDeleteSemaphore (MPSemaphoreID semaphore);
```

`semaphore` A value of type `MPSemaphoreID` (page 119) that specifies the ID of the semaphore you want to remove.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

Calling this function unblocks all tasks waiting on the semaphore and the tasks’ respective `MPWaitOnSemaphore` (page 75) calls will return with the result code `kMPDeletedErr`.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPCreateSemaphore` (page 73).

MPSignalSemaphore

Signals a semaphore.

```
OSStatus MPSignalSemaphore (MPSemaphoreID semaphore);
```

semaphore A value of type `MPSemaphoreID` (page 119) that specifies the ID of the semaphore you want to signal.

function result A result code. See “Result Codes” (page 133) for a list of possible values. If the value of the semaphore was already at the maximum, `MPSignalSemaphore` returns `kInsufficientResourcesErr`.

DISCUSSION

If tasks are waiting on the semaphore, the oldest (first queued) task is unblocked so that the corresponding `MPWaitOnSemaphore` (page 75) call for that task completes. Otherwise, if the value of the semaphore is not already equal to its maximum value, it is incremented by one.

Note that you can call this function from an interrupt handler.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPWaitOnSemaphore` (page 75).

MPWaitOnSemaphore

Waits on a semaphore

```
OSStatus MPWaitOnSemaphore (
    MPSemaphoreID semaphore,
    Duration timeout);
```

Multiprocessing Services Reference

<code>semaphore</code>	A value of type <code>MPSemaphoreID</code> (page 119) that specifies the ID of the semaphore you want to wait on.
<code>timeout</code>	A value of type <code>Duration</code> that specifies the maximum time the function should wait before timing out. See “Timer Duration Constants” (page 124) for a list of constants you can use to specify the wait interval.
<i>function result</i>	A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

If the value of the semaphore is greater than zero, the value is decremented and the function returns with `noErr`. Otherwise, the task is blocked awaiting a signal until the specified timeout is exceeded.

If you call this function from a cooperative task, you should specify only `kDurationImmediate` for the timeout length; other waits will cause the task to block.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPSignalSemaphore` (page 75).

Handling Event Groups

Multiprocessing Services provides the following functions for creating and handling event groups:

- `MPCreateEvent` (page 77) creates an event group.
- `MPDeleteEvent` (page 77) removes an event group.
- `MPSetEvent` (page 78) merges event flags into a specified event group.
- `MPWaitForEvent` (page 79) retrieves event flags from a specified event group.

MPCreateEvent

Creates an event group.

```
OSStatus MPMCreateEvent (MPEventID *event);
```

event A pointer to a variable of type `MPEventID` (page 119). On return, the variable contains the ID of the newly created event group.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

Event groups are created from dynamically allocated internal resources. Other tasks may be competing for these resources so it is possible that this function will not be able to create an event group.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPDeleteEvent` (page 77).

MPDeleteEvent

Removes an event group.

```
OSStatus MPDeleteEvent (MPEventID event);
```

event A value of type `MPEventID` (page 119) that specifies the ID of the event group you want to remove.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

After deletion, the event ID becomes invalid, and all internal resources associated with the event group are reclaimed. Calling this function unblocks all tasks waiting on the event group and their respective `MPWaitForEvent` (page 79) calls will return with the result code `kMPDeletedErr`.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPCreateEvent` (page 77).

MPSetEvent

Merges event flags into a specified event group.

```
OSStatus MPSetEvent (
    MPEventID event,
    MPEventFlags flags);
```

event A value of type `MPEventID` (page 119) that specifies the ID of the event group you want to set.

flags A 32-bit value of type `MPEventFlags` (page 122) that contains the flags you want to merge into the event group.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

The flags are logically ORed with the current flags in the event group. This procedure is an atomic operation to ensure that multiple updates do not get lost. If tasks are waiting on this event group, the first waiting task is unblocked.

Note that you can call this function from an interrupt handler.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPWaitForEvent` (page 79).

MPWaitForEvent

Retrieves event flags from a specified event group.

```
OSStatus MPWaitForEvent (
    MPEventID event,
    MPEventFlags *flags,
    Duration timeout);
```

<code>event</code>	A value of type <code>MPEventID</code> (page 119) that specifies the event group whose flags you want to retrieve.
<code>flags</code>	A pointer to a 32-bit variable of type <code>MPEventFlags</code> (page 122). On return, the variable contains the flags of the specified event group. Pass <code>NULL</code> if you do not need any flag information.
<code>timeout</code>	A value of type <code>Duration</code> that specifies the maximum time to wait for events before timing out. See “Timer Duration Constants” (page 124) for a list of constants you can use to specify the wait interval.
<i>function result</i>	A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

This function obtains event flags from the specified event group. The timeout specifies how long to wait for events if none are present when the call is made. If any flags are set when this function is called, all the flags in the event group are moved to the `flag` field and the event group is cleared. This obtaining and clearing action is an atomic operation to ensure that no updates are lost. If

Multiprocessing Services Reference

multiple tasks are waiting on an event group, only one can obtain any particular set of flags.

If you call this function from a cooperative task, you should specify only `kDurationImmediate` for the timeout length; other waits will cause the task to block.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPSetEvent` (page 78).

Handling Kernel Notifications

Multiprocessing Services provides the following functions for manipulating kernel notifications:

- `MPCreateNotification` (page 80) creates a kernel notification.
- `MPDeleteNotification` (page 81) removes a kernel notification.
- `MPSModifyNotification` (page 82) adds a simple notification to a kernel notification.
- `MPCauseNotification` (page 83) signals a kernel notification.

MPCreateNotification

Creates a kernel notification

```
OSStatus MPCreateNotification (MPNotificationID * notificationID);
```

```
notificationID
```

A pointer to a variable of type `MPNotificationID` (page 120). On return, `notificationID` points to the newly created kernel notification.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DESCRIPTION

After creating the kernel notification object, you can add simple notifications by calling the function `MPModifyNotification` (page 82).

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

SEE ALSO

The function `MPDeleteNotification` (page 81).

MPDeleteNotification

Removes a kernel notification.

```
OSStatus MPDeleteNotification (MPNotificationID notificationID);
```

`notificationID`

The ID of the notification you want to remove.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

SEE ALSO

The function `MPCreateNotification` (page 80).

MPMModifyNotification

Adds a simple notification to a kernel notification.

```
OSStatus MPMModifyNotification (
    MPNotificationID notificationID,
    MPOpaqueID anID,
    void * notifyParam1,
    void * notifyParam2,
    void * notifyParam3);
```

notificationID

The ID of the kernel notification you want to add to..

anID

The ID of the simple notification (semaphore, message group, or event group) you want to add to the kernel notification.

notifyParam1

If *anID* specifies an event group, this parameter should contain the flags to set in the event group when `MPCauseNotification` (page 83) is called. If *anID* specifies a message queue, this parameter should contain the first 32-bits of the message to be sent to the message queue when `MPCauseNotification` is called.

notifyParam2

If *anID* specifies a message queue, this parameter should contain the second 32-bits of the message to be sent to the message queue when `MPCauseNotification` is called. Pass `NULL` if you don't need this parameter.

notifyParam3

If *anID* specifies a message queue, this parameter should contain the third 32-bits of the message sent to the message queue when `MPCauseNotification` is called. Pass `NULL` if you don't need this parameter.

function result

A result code. See "Result Codes" (page 133) for a list of possible values.

DESCRIPTION

You specify the parameters for the simple notifications just as if you were calling the `MPSetTimerNotify` (page 90) function.

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

MPCauseNotification

Signals a kernel notification.

```
OSStatus MPCauseNotification (MPNotificationID notificationID);
```

`notificationID`

A value of type `MPNotificationID` (page 120) that indicates the ID of the kernel notification you want to signal.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DESCRIPTION

You call this function to signal a kernel notification much as you would signal any simple notification (for example, `MPNotifyQueue` (page 69)).

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

Handling Critical Regions

Multiprocessing Services provides the following functions for creating and handling critical regions:

- `MPCreateCriticalRegion` (page 84) creates a critical region object.
- `MPDeleteCriticalRegion` (page 84) removes the specified critical region object.
- `MPEnterCriticalRegion` (page 85) attempts to enter a critical region.
- `MPExitCriticalRegion` (page 86) exits a critical region.

MPCreateCriticalRegion

Creates a critical region object.

```
OSStatus MPMCreateCriticalRegion (MPCriticalRegionID *criticalRegion);
```

```
criticalRegion
```

A pointer to a variable of type `MPCriticalRegionID` (page 120). On return, the variable contains the ID of the newly created critical region object.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPDeleteCriticalRegion` (page 84).

MPDeleteCriticalRegion

Removes the specified critical region object.

```
OSStatus MPDeleteCriticalRegion (MPCriticalRegionID criticalRegion);
```

```
criticalRegion
```

A value of type `MPCriticalRegionID` (page 120) that specifies the critical region object you want to remove.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

Calling this function unblocks all tasks waiting to enter the critical region and their respective `MPEnterCriticalRegion` (page 85) calls will return with the result code `kMPDeletedErr`.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPCreateCriticalRegion` (page 84).

MPEnterCriticalRegion

Attempts to enter a critical region.

```
OSStatus MPEnterCriticalRegion (
    MPCriticalRegionID criticalRegion,
    Duration timeout);
```

criticalRegion

A value of type `MPCriticalRegionID` (page 120) that specifies the ID of the critical region you want to enter.

timeout

A value of type `Duration` that specifies the maximum time to wait for entry before timing out. See “Timer Duration Constants” (page 124) for a list of constants you can use to specify the wait interval.

function result

A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

If another task currently occupies the critical region, the current task is blocked until the critical region is released or until the designated timeout expires. Otherwise the task enters the critical region and `MPEnterCriticalRegion` increments the region’s use count.

Multiprocessing Services Reference

Once a task enters a critical region it can make further calls to `MPEnterCriticalRegion` without blocking (its use count increments for each call). However, each call to `MPEnterCriticalRegion` must be balanced by a call to `MPExitCriticalRegion` (page 86); otherwise the region is not released for use by other tasks.

Note that you can enter a critical region from a cooperative task. Each cooperative task is treated as unique and different from any preemptive task. If you call this function from a cooperative task, you should specify only `kDurationImmediate` for the timeout length; other waits will cause the task to block.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPExitCriticalRegion` (page 86).

MPExitCriticalRegion

Exits a critical region.

```
OSStatus MPExitCriticalRegion (MPCriticalRegionID criticalRegion);
```

`criticalRegion`

A value of type `MPCriticalRegionID` that specifies the ID of the critical region you want to exit.

function result A result code. See “Result Codes” (page 133) for a list of possible values. If the task does not own the critical region specified by `criticalRegion`, `MPExitCriticalRegion` returns `kMPInsufficientResourcesErr`.

DISCUSSION

This function decrements the use count of the critical region object. When the use count reaches zero, ownership of the critical region object is released (which allows another task to use the critical region).

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPEnterCriticalRegion` (page 85).

Timer Services Functions

Multiprocessing Services provides the following timer services functions:

- `MPDelayUntil` (page 87) blocks the calling task until a specified time.
- `MPCreateTimer` (page 88) creates a timer.
- `MPDeleteTimer` (page 89) removes a timer.
- `MPSetTimerNotify` (page 90) sets the notification information associated with a timer.
- `MPArmTimer` (page 91) arms the timer to expire at a given time.
- `MPCancelTimer` (page 92) cancels an armed timer.

MPDelayUntil

Blocks the calling task until a specified time.

```
OSStatus MPDelayUntil (AbsoluteTime *expirationTime);
```

```
expirationTime
```

A value of type `AbsoluteTime` specifying when to unblock the task.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

You cannot call this function from a cooperative task.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPCreateTimer

Creates a timer.

```
OSStatus MPCreateTimer (MPTimerID *timerID);
```

timerID A pointer to a variable of type `MPTimerID` (page 120). On return, the variable contains the ID of the newly created timer.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

You can use a timer to notify an event, queue, or semaphore after a specified amount of time has elapsed.

Timer objects are created from dynamically-allocated internal resources. Other tasks may be competing for these resources so it is possible this function may not be able to create one.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPDeleteTimer` (page 89).

The function `MPArmTimer` (page 91).

The function `MPSetTimerNotify` (page 90).

MPDeleteTimer

Removes a timer.

```
OSStatus MPDeleteTimer (MPTimerID timerID);
```

`timerID` A value of type `MPTimerID` (page 120) that specifies the ID of the timer you want to remove.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

After deletion, the timer ID becomes invalid, and all internal resources associated with the timer are reclaimed.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPCreateTimer` (page 88).

MPSetTimerNotify

Sets the notification information associated with a timer.

```
OSStatus MPSetTimerNotify (
    MPTimerID timerID,
    MPOpaqueID anID,
    void *notifyParam1,
    void *notifyParam2,
    void *notifyParam3);
```

<code>timerID</code>	A value of type <code>MPTimerID</code> (page 120) that specifies the ID of the timer whose notification information you want to set.
<code>anID</code>	A value specifying the ID of the notification mechanism to associate with the timer. This value should be the ID of an event group, a message queue, or a semaphore.
<code>notifyParam1</code>	If <code>anID</code> specifies an event group, this parameter should contain the flags to set in the event group when the timer expires. If <code>anID</code> specifies a message queue, this parameter should contain the first 32-bits of the message to be sent to the message queue when the timer expires.
<code>notifyParam2</code>	If <code>anID</code> specifies a message queue, this parameter should contain the second 32-bits of the message to be sent to the message queue when the timer expires. Pass <code>NULL</code> if you don't need this parameter.
<code>notifyParam3</code>	If <code>anID</code> specifies a message queue, this parameter should contain the third 32-bits of the message sent to the message queue when the timer expires. Pass <code>NULL</code> if you don't need this parameter.
<i>function result</i>	A result code. See "Result Codes" (page 133) for a list of possible values.

DISCUSSION

When the timer expires, Multiprocessing Services checks the notification ID, and if it is valid, notifies the related notification mechanisms (that is, event groups, queues, or semaphores) you had specified in your `MPSetTimerNotify` (page 90) calls.

Multiprocessing Services Reference

You can specify multiple notification mechanisms by calling this function several times. For example, you can call `MPSetTimerNotify` to specify a message queue and then call it again to specify a semaphore. When the timer expires, a message is sent to the message queue and the appropriate semaphore is signaled. You cannot, however, specify more than one notification per notification mechanism (for example, if you call `MPSetTimerNotify` twice, specifying different messages or message queues in each call, the second call will overwrite the first). Note that if a call to `MPSetTimerNotify` returns an error, any previous calls specifying the same timer are still valid; previously set notifications will still be notified when the timer expires.

You can set the notification information at any time. If the timer is armed, it will modify the notification parameters dynamically. If the timer is disarmed, it will modify the notification parameters to be used for the next `MPArmTimer` (page 91) call.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPArmTimer

Arms the timer to expire at a given time.

```
OSStatus MPArmTimer (
    MPTimerID timerID,
    AbsoluteTime *expirationTime,
    OptionBits options);
```

`timerID` A value of type `MPTimerID` (page 120) that specifies the ID of the timer you want to arm.

`expirationTime` A pointer to a value of type `AbsoluteTime` that specifies when you want the timer to expire. Note that if you arm the timer with a time that has already passed, the timer expires immediately.

`options` A value of type `OptionBits` specifying any optional action. See “Timer Option Masks” (page 125) for a list of possible values.

function result A result code. See “Result Codes” (page 133) for a list of possible values. If the timer has already expired, the reset does not take place and the function returns `kMPInsufficientResourcesErr`.

DISCUSSION

The expiration time is an absolute time, which you can generate by calling the `UpTime` (page 93) function. When the timer expires, a notification is sent to the notification mechanism specified in the last `MPSetTimerNotify` (page 90) call. If the specified notification ID has become invalid, no action is taken when the timer expires. The timer itself is deleted when it expires unless you specified the `kMPPreserveTimerID` option in the options parameter.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPCancelTimer` (page 92).

The function `MPSetTimerNotify` (page 90).

MPCancelTimer

Cancels an armed timer.

```
OSStatus MPCancelTimer (
    MPTimerID timerID,
    AbsoluteTime *timeRemaining);
```

`timerID` A value of type `MPTimerID` that specifies the ID of the armed timer you want to cancel.

`timeRemaining` A pointer to a variable of type `AbsoluteTime`. On return, the variable contains the time remaining before the timer would have expired.

function result A result code. See “Result Codes” (page 133) for a list of possible values. If the timer has already expired, this function returns `kMPIInsufficientResourcesErr`.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPArmTimer` (page 91).

Time Utility Functions

Multiprocessing Services provides the following time utility functions:

- `UpTime` (page 93) obtains the time elapsed since system startup.
- `DurationToAbsolute` (page 94) converts time from units of type `Duration` to units of type `AbsoluteTime`.
- `AbsoluteToDuration` (page 95) converts time from units of type `AbsoluteTime` to units of type `Duration`.

UpTime

Obtains the time elapsed since machine startup.

```
AbsoluteTime UpTime (void);
```

function result A value of type `AbsoluteTime` that indicates the time elapsed since the host computer was started up. Absolute time is a 64-bit monotonically increasing value. You should not make any assumptions about what units absolute time is based upon.

DISCUSSION

You can call this function from the task level, software interrupt level, or hardware interrupt level.

This function is identical to the `UpTime` function found in the Driver Services Library.

VERSION NOTES

Available in the Multiprocessing Services library with version 2.0.

DurationToAbsolute

Converts time from units of type `Duration` to units of type `AbsoluteTime`.

```
AbsoluteTime DurationToAbsolute (Duration duration);
```

`duration` A value of type `Duration` that you want to convert to type `AbsoluteTime`.

function result A value of type `AbsoluteTime` which is the converted time.

DISCUSSION

This function is identical to the `DurationToAbsolute` function found in the Driver Services Library.

VERSION NOTES

Available in the Multiprocessing Services library with version 2.0.

SEE ALSO

The function `AbsoluteToDuration` (page 95).

AbsoluteToDuration

Converts time from units of type `AbsoluteTime` to units of type `Duration`.

```
Duration AbsoluteToDuration (AbsoluteTime time);
```

`time` A value of type `AbsoluteTime` that you want to convert to type `Duration`.

function result A value of type `Duration` which is the converted time.

DISCUSSION

This function is identical to the `AbsoluteToDuration` function found in the Driver Services Library.

VERSION NOTES

Available in the Multiprocessing Services library with version 2.0.

SEE ALSO

The function `DurationToAbsolute` (page 94).

Accessing Per-Task Storage Variables

Multiprocessing Services provides the following functions for creating and manipulating per-task storage variables:

- `MPAllocateTaskStorageIndex` (page 96) returns an index number to access per-task storage.
- `MPDeallocateTaskStorageIndex` (page 96) frees an index number used to access per-task storage.
- `MPSetTaskStorageValue` (page 97) sets the storage value for a given index number.
- `MPGetTaskStorageValue` (page 98) gets the storage value stored at a specified index number.

MPAllocateTaskStorageIndex

Returns an index number to access per-task storage.

```
OSStatus MPAllocateTaskStorageIndex (TaskStorageIndex *index);
```

index A pointer to a variable of type `TaskStorageIndex` (page 121). On return, the variable contains an index number you can use to store task data.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

A call to the function `MPAllocateTaskStorageIndex` returns an index number that is common across all tasks in the current process. You can use this index number in calls to `MPSetTaskStorageValue` (page 97) and `MPGetTaskStorageValue` (page 98) to set a different value for each task using the same index.

You can think of the task storage area as a two dimensional array cross-referenced by the task storage index number and the task ID. Note that since the amount of per-task storage is determined when the task is created, the number of possible index values associated with a task is limited.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPDeallocateTaskStorageIndex` (page 96).

MPDeallocateTaskStorageIndex

Frees an index number used to access per-task storage

```
OSStatus MPDeallocateTaskStorageIndex (TaskStorageIndex index);
```

Multiprocessing Services Reference

- index* A value of type `TaskStorageIndex` (page 121) that specifies the index number you want to deallocate.
- function result* A result code. See “Result Codes” (page 133) for a list of possible values.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPAllocateTaskStorageIndex` (page 96).

MPSetTaskStorageValue

Sets the storage value for a given index number.

```
OSStatus MPSetTaskStorageValue (
    TaskStorageIndex index,
    TaskStorageValue value);
```

- index* A 32-bit value of type `TaskStorageIndex` (page 121) specifying the index number whose storage value you want to set.
- value* A value of type `TaskStorageValue` (page 121) that specifies the value you want to set.
- function result* A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

Typically you use `MPSetTaskStorageValue` to store pointers to task-specific structures or data.

Calling this function from within a task effectively assigns a value in a two-dimensional array cross-referenced by task storage index value and the task ID.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPGetTaskStorageValue` (page 98).

MPGetTaskStorageValue

Gets the storage value stored at a specified index number.

```
TaskStorageValue MPGetTaskStorageValue (TaskStorageIndex index);
```

index A 32-bit value of type `TaskStorageIndex` (page 121) specifying the index number of the storage value you want to obtain.

function result A value of type `TaskStorageValue` (page 121) that is the value stored at the specified index number.

DISCUSSION

Calling this function from within a task effectively reads a value in a two-dimensional array cross-referenced by task storage index value and the task ID.

Note that since this function does not return any status information, it may not be immediately obvious whether the returned storage value is valid.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPSetTaskStorageValue` (page 97).

Memory Allocation Functions

Multiprocessing Services provides the following functions for allocating and manipulating blocks of memory from within preemptive tasks:

- `MPAllocateAligned` (page 99) allocates a nonrelocatable memory block.
- `MPFree` (page 100) frees memory allocated by `MPAllocateAligned` (page 99).
- `MPGetAllocatedBlockSize` (page 101) returns the size of a memory block.
- `MPBlockCopy` (page 101) copies a block of memory.
- `MPBlockClear` (page 102) clears a block of memory.
- `MPDataToCode` (page 103) designates the specified block of memory as executable code.

MPAllocateAligned

Allocates a nonrelocatable memory block.

```
LogicalAddress MPAllocateAligned (
    ByteCount size,
    UInt8 alignment,
    OptionBits options);
```

size	A value of type <code>ByteCount</code> specifying the size, in bytes, of the memory block to allocate.
alignment	An integer specifying the desired alignment of the allocated memory block. See “Memory Allocation Alignment Constants” (page 126) for a list of possible values to pass. Note that there will be a minimum alignment regardless of the requested alignment. If the requested memory block is 4 bytes or smaller, the block will be at least 4-byte aligned. If the requested block is greater than 4 bytes, the block will be at least 8-byte aligned.
options	A value of type <code>OptionBits</code> that specifies any optional information to use with this call. See “Memory Allocation Option Constants” (page 127) for a list of possible values to pass.

Multiprocessing Services Reference

function result A pointer to the allocated memory. If the function cannot allocate the requested memory or the requested alignment, the returned address is `NULL`.

DISCUSSION

The memory referenced by the returned address is guaranteed to be accessible by the application's cooperative task and any preemptive tasks that it creates, but not by other applications or their preemptive tasks. Any existing nonglobal heap blocks are freed when the application terminates. As with all shared memory, you must explicitly synchronize access to allocated heap blocks using a notification mechanism.

You can replicate the effect of the older `MPAllocate` function by calling `MPAllocateAligned` with 32-byte alignment and no options.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPFree` (page 100).

MPFree

Frees memory allocated by `MPAllocateAligned` (page 99).

```
void MPFree (LogicalAddress object);
```

object A pointer of type `LogicalAddress` to the memory you want to release.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

SEE ALSO

The function `MPAllocateAligned` (page 99).

MPGetAllocatedBlockSize

Returns the size of a memory block.

```
ByteCount MPGetAllocatedBlockSize (LogicalAddress object);
```

object A value of type `LogicalAddress` that specifies the address of the memory block whose size you want to determine.

function result A value of type `ByteCount` that is the size of the allocated memory block, in bytes.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPBlockCopy

Copies a block of memory.

```
void MPBlockCopy (
    LogicalAddress source,
    LogicalAddress destination,
    ByteCount size);
```

source A pointer of type `LogicalAddress` that specifies the starting address of the memory block you want to copy.

destination A pointer of type `LogicalAddress` that specifies the location to which you want to copy the memory block.

size A value of type `ByteCount` that specifies the number of bytes to copy.

DISCUSSION

As with all shared memory, your application must synchronize access to the memory blocks to avoid data corruption. `MPBlockCopy` ensures the copying stays within the bounds of the area specified by `size`, but the calling task can be preempted during the copying process.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

MPBlockClear

Clears a block of memory.

```
void MPBlockClear (
    LogicalAddress address,
    ByteCount size);
```

`address` A pointer of type `LogicalAddress` that specifies the starting address of the memory block you want to clear.

`size` A value of type `ByteCount` that specifies the number of bytes you want to clear.

DISCUSSION

As with all shared memory, your application must synchronize access to the memory blocks to avoid data corruption. `MPBlockClear` ensures the clearing stays within the bounds of the area specified by `size`, but the calling task can be preempted during the copying process.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPDataToCode

Designates the specified block of memory as executable code.

```
void MPDataToCode (  
    LogicalAddress address,  
    ByteCount size);
```

address A pointer of type `LogicalAddress` that specifies the starting address of the memory block you want to designate as code.

size A value of type `ByteCount` that specifies the size of the memory block.

DISCUSSION

Since PowerPC processors need to differentiate between code and data in memory, you should call this function to tag any executable code that your tasks may generate.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Exception Handling Functions

Multiprocessing Services provides the following functions for handling exceptions and examining the state of a suspended task:

- `MPSetExceptionHandler` (page 104) sets an exception handler for a task.
- `MPTthrowException` (page 105) throws an exception to a specified task.
- `MPDisposeTaskException` (page 106) removes a task exception.
- `MPExtractTaskState` (page 107) extracts state information from a suspended task.
- `MPSetTaskState` (page 108) sets state information for a suspended task.

MPSetExceptionHandler

Sets an exception handler for a task.

```
OSStatus MPSetExceptionHandler (
    MPTaskID task,
    MPQueueID exceptionQ);
```

task A value of type `MPTaskID` (page 115) that specifies the task to associate with the exception handler.

exceptionQ A value of type `MPQueueID` (page 118) that specifies the message queue to which an exception message will be sent.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

DISCUSSION

When an exception handler is set and an exception occurs, the task is suspended and a message is sent to the message queue specified by `exceptionQ`. The message contains the following information:

- The first 32-bits contain the ID of the task in which the exception occurred.
- The second 32-bits contain the type of exception that occurred. See the header file `MachineExceptions.h` for a listing of exception types.
- The last 32-bits are set to `NULL` (reserved for future use).

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPThrowException

Throws an exception to a specified task.

```
OSStatus MPThrowException (
    MPTaskID task,
    MPExceptionKind kind);
```

task A value of type `MPTaskID` (page 115) that specifies the task to which the exception should be thrown.

kind A value of type `ExceptionKind` that specifies the type of exception to give to the task.

function result A result code. See “Result Codes” (page 133) for a list of possible values. If the task is already suspended or if the task is not defined to take thrown exceptions, the function returns `kMPIInsufficientResourcesErr`.

DISCUSSION

The exception is treated in the same manner as any other exception taken by a task. However, since it is asynchronous, it may not be presented immediately.

By convention, you should set the exception kind to `kMPTaskStoppedErr` if you want to suspend a task. In general, you should do so only if you are debugging and wish to examine the state of the task. Otherwise you should block the task using one of the traditional notification mechanisms (such as a message queue).

An exception can be thrown at any time, whether that task is running, eligible to be run (that is, ready), or blocked. The task is suspended and an exception message may be generated the next time the task is about to run. Note that this may never occur—for example, if the task is deadlocked or the resource it is waiting on is never released. If the task is currently blocked when this function is executed, `kMPTaskBlockedErr` is returned. If the task was suspended immediately at the conclusion of this function call the return value is `kMPTaskStoppedErr`.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPDisposeTaskException

Removes a task exception.

```
OSStatus MPDisposeTaskException (
    MPTaskID task,
    OptionBits action);
```

task A value of type `MPTaskID` (page 115) that specifies the task whose exception you want to remove.

action A value of type `OptionBits` that specifies actions to perform on the task. For example, you can enable single-stepping when the task resumes, or you can pass the exception on to another handler. See “Task Exception Disposal Constants” (page 131) for a listing of possible values.

function result A result code. See “Result Codes” (page 133) for a list of possible values. If the specified action is invalid or unsupported, or if the specified task is not suspended, this function returns `kMPInsufficientResourcesErr`.

DISCUSSION

This function removes the task exception and allows the task to resume operation. If desired, you can enable single-stepping or branch-stepping, or propagate the exception instead.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPExtractTaskState

Extracts state information from a suspended task.

```
OSStatus MPExtractTaskState (
    MPTaskID task,
    MPTaskStateKind kind,
    void *info);
```

task A value of type `MPTaskID` that specifies the task whose state information you want to obtain.

kind A value of type `MPTaskStateKind` that specifies the kind of state information you want to obtain. See “Task State Constants” (page 129) for a listing of possible values.

info A pointer to a data structure to hold the state information. On return, the data structure holds the desired state information. The format of the data structure varies depending on the state information you want to retrieve. See the header file `MachineExceptions.h` for the formats of the various state information structures.

function result A result code. See “Result Codes” (page 133) for a list of possible values. If you attempt to extract state information for a running task, this function returns `kMPIInsufficientResourcesErr`.

DISCUSSION

You can use this function to obtain register contents or exception information about a particular task.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPSetTaskState` (page 108).

MPSetTaskState

Sets state information for a suspended task.

```
OSStatus MPSetTaskState (
    MPTaskID task,
    MPTaskStateKind kind,
    void *info);
```

task A value of type `MPTaskID` (page 115) that specifies the task whose state information you want to set.

kind A value of type `MPTaskStateKind` that specifies the kind of state information you want to set. See “Task State Constants” (page 129) for a listing of possible values. Note that some state information is read-only and cannot be changed using this function.

info A pointer to a data structure holding the state information you want to set. The format of the data structure varies depending on the state information you want to set. See the header file `MachineExceptions.h` for the formats of the various state information structures.

function result A result code. See “Result Codes” (page 133) for a list of possible values. If you specify `kMPTaskState32BitMemoryException` for the state information, this function returns `kMPInsufficientResourcesErr`, since the exception state information is read-only. Attempting to set state information for a running task will also return `kMPInsufficientResourcesErr`.

DISCUSSION

You can use this function to set register contents or exception information for a particular task. However, some state information, such as the exception information (as specified by `kMPTaskState32BitMemoryException`) as well as the MSR, `ExceptKind`, `DSISR`, and `DAR` machine registers (specified under `kMPTaskStateMachine`) are read-only. Attempting to set the read-only machine registers will do nothing, while attempting to set the exception information will return an error.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPExtractTaskState` (page 107).

Debugger Support Functions

Multiprocessing Services supplies two functions for use with third-party debuggers:

- `MPRegisterDebugger` (page 109) registers a debugger.
- `MPUnregisterDebugger` (page 110) unregisters a debugger.

MPRegisterDebugger

Registers a debugger.

```
OSStatus MPRegisterDebugger (
    MPQueueID queue,
    MPDebuggerLevel level);
```

queue A value of type `MPQueueID` (page 118) that specifies the ID of the queue to which you want exception messages and other information to be sent.

level A value of type `MPDebuggerLevel` (page 123) that specifies the level of this debugger with respect to other debuggers. Exceptions and informational messages are sent first to the debugger with the highest level. If more than one debugger attempts to register at a particular level, only the first debugger is registered. Other attempts return an error.

function result A result code. See “Result Codes” (page 133) for a list of possible values. If the number of registered debuggers exceeds the system limit, the function returns `kMPInsufficientResourcesErr`.

DISCUSSION

Exception messages are sent when tasks are suspended. When a task exception occurs, Multiprocessing Services notifies registered debuggers and other handlers in the following order:

- the registered debugger queue of the highest level
- lower level registered debugger queues according to level
- the local exception handler
- task termination

The notification moves to the next lower level when a debugger (or eventually the task's local exception handler) calls the `MPDisposeTaskException` (page 106) function with the indication that the exception be propagated. The exception messages sent to the debugger's queue are in the same format as those described in `MPSetExceptionHandler` (page 104).

The system may implement a limited number of debugger slots. It is possible no debugger support is provided.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPUnregisterDebugger` (page 110).

MPUnregisterDebugger

Unregisters a debugger.

```
OSStatus MPUnregisterDebugger (MPQueueID queue);
```

queue A value of type `MPQueueID` (page 118) that specifies the ID of the queue whose debugger you want to unregister.

function result A result code. See “Result Codes” (page 133) for a list of possible values.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

SEE ALSO

The function `MPRegisterDebugger` (page 109).

Remote Calling Function

If you want to call a nonreentrant function (such as a system software function) from a preemptive task, you must call it remotely using the `MPRemoteCall` (page 111) function.

MPRemoteCall

Calls a nonreentrant function and blocks the current task.

```
void *MPRemoteCall (
    MPRemoteProcedure remoteProc,
    void *parameter,
    MPRemoteContext context);
```

<code>remoteProc</code>	A pointer of type <code>MPRemoteProcedure</code> (page 123) that references the application-defined function you want to call. See <code>MyRemoteProcedure</code> (page 113) for more information about the form of this function.
<code>parameter</code>	A pointer to a parameter to pass to the application-defined function. For example, this value could point to a data structure or a memory location.
<code>context</code>	A value of type <code>MPRemoteContext</code> that specifies which contexts (that is processes) are allowed to execute the function. See “Remote Call Context Option Constants” (page 132) for a list of possible values.

DISCUSSION

You use this function primarily to indirectly execute Mac OS system software functions. The task making the remote call is blocked until the call completes. The amount of time taken to schedule the remote procedure depends on the choice of the designated operating context. Specifying `kMPAnyRemoteContext` offers the lowest latency, but the called procedure may not have access to process-specific resources such as some low-memory values. Specifying `kMPOwningProcessRemoteContext` has higher latency because the remote procedure is deferred until the owning process becomes active. However, the remote procedure is guaranteed to execute within the owning process.

Note that with the exception of functions in Multiprocessing Services, you cannot safely call any system software functions directly from a preemptive task. Even if some system software function appears to work today when called from a preemptive task, unless explicitly stated otherwise there is no guarantee that subsequent versions of the same function will continue to work in future versions of system software. In Mac OS 8 implementations of Multiprocessing Services, the only exceptions to this rule are the atomic memory operations (such as `AddAtomic`) exported in the `InterfaceLib` shared library. Even these functions may switch to 68K mode if the operands to them are not aligned. If you need to access system software functions from a preemptive task, you must do so using the `MPRemoteCall` function.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Application-Defined Function

Multiprocessing Services allows you to create an application-defined function such as `MyRemoteProcedure` (page 113) to call nonreentrant functions from your preemptive task.

MyRemoteProcedure

When calling `MPRemoteCall` (page 111), you must designate an application-defined function to handle any calls to nonreentrant functions (such as Mac OS system software calls). For example, this is how you would declare the application-defined function if you were to name the function

`MyRemoteProcedure`:

```
void* MyRemoteProcedure (void *parameter);
```

`parameter` A pointer to any information you want to pass to `MyRemoteProcedure`. For example, `parameter` might point to a parameter list that `MyRemoteProcedure` could then pass to a Mac OS system software function.

DISCUSSION

Note that your application-defined function must be PowerPC native code, since Multiprocessing Services tasks cannot call 68K code.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Data Types

This section describes the data types used with Multiprocessing Services:

- `MPPProcessID` (page 114)
- `MPCpuID` (page 115)
- `MPCoherenceID` (page 115)
- `MPTaskID` (page 115)
- `TaskProc` (page 116)
- `MPTaskOptions` (page 116)
- `MPTaskWeight` (page 117)

Multiprocessing Services Reference

- MPTaskInfo (page 117)
- MPQueueID (page 118)
- MPSemaphoreID (page 119)
- MPEventID (page 119)
- MPCriticalRegionID (page 120)
- MPNotificationID (page 120)
- MPTimerID (page 120)
- TaskStorageIndex (page 121)
- TaskStorageValue (page 121)
- MPSemaphoreCount (page 122)
- MPEventFlags (page 122)
- MPExceptionKind (page 122)
- MPDebuggerLevel (page 123)
- MPRemoteProcedure (page 123)

MPPProcessID

Multiprocessing Services manipulates processes by passing a process ID, which is defined by the `MPPProcessID` type:

```
typedef struct OpaqueMPPProcessID* MPPProcessID;
```

Note that this process ID is identical to the process ID (or context ID) handled by the Code Fragment Manager.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPCpuID

Multiprocessing Services manipulates CPUs (processors) by passing a CPU ID, which is defined by the `MPCpuID` type:

```
typedef struct OpaqueMPCpuID *MPCpuID;
```

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

MPCoherenceID

Multiprocessing Services handles memory coherence groups by passing a coherence ID, which is defined by the `MPCoherenceID` type:

```
typedef struct OpaqueMPCoherenceID *MPCoherenceID;
```

A coherence group is the set of processors and other bus controllers that have cache-coherent access to memory. Mac OS 9 defines only one coherence group, which is all the processors that can access internal memory (RAM). Other coherence groups are possible; for example, a PCI card with its own memory and processors can comprise a coherence group.

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

MPTaskID

Multiprocessing Services manipulates tasks by passing a task ID, which is defined by the `MPTaskID` type:

```
typedef struct OpaqueMPTaskID* MPTaskID;
```

You obtain a task ID by calling the function `MPCreateTask` (page 59).

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

TaskProc

When calling the `MPCreateTask` (page 59) function, you specify the task to create by passing a pointer of type `TaskProc` which references the code to be executed as the task:

```
typedef OSStatus (*TaskProc) (void *parameter);
```

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

MPTaskOptions

When calling the `MPCreateTask` (page 59) function, you can specify any optional attributes for the task created by passing a value of type `MPTaskOptions`, which has the following type definition:

```
typedef OptionBits MPTaskOptions;
```

See “Task Creation Options” (page 128) for a list of possible values.

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

MPTaskWeight

The Multiprocessing Services function `MPSetTaskWeight` (page 65) handles the weight of a task using the `MPTaskWeight` type, which has the following definition.

```
typedef UInt32 MPTaskWeight;
```

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPTaskInfo

If you specify the `kMPTaskStateTaskInfo` constant when calling the function `MPExtractTaskState` (page 107), Multiprocessing Services returns state information in a data structure of type `MPTaskInfo`, which has the following definition:

```
struct MPTaskInfo {
    PBVersion      version; /* Version of the data structure requested*/
    OSType         name;
    OSType         queueName;
    UInt16         runState;
    UInt16         lastCPU;
    UInt32         weight;
    MPProcessID    processID;
    AbsoluteTime   cpuTime;
    AbsoluteTime   schedTime;
    AbsoluteTime   creationTime;
    ItemCount      codePageFaults;
    ItemCount      dataPageFaults;
    ItemCount      preemptions;
    MPCpuID        cpuID;
};
typedef struct MPTaskInfo MPTaskInfo;
```

Multiprocessing Services Reference

Field Descriptions

version	The version of this data structure.
name	The name of the task.
queueName	A four-byte code indicating the status of the queue waiting on the task. This field can contain any of the following values: 'SEMA' - waiting on a semaphore 'MSGQ' - waiting on a message queue 'CRGN' - waiting on a critical region 'EVNT' - waiting on an event group 'DEBUG' - task is suspended 'DLYQ' - waiting on a synchronous delay 'RDYQ' - on the ready queue
runState	The current state of the task (running, ready, or blocked).
lastCPU	The address of the last processor that ran this task.
weight	The weighting assigned to this task.
processID	The ID of the process that owns this task.
cpuTime	The accumulated CPU time used by the task.
schedTime	The time when the task was last scheduled.
creationTime	The time when the task was created.
codePageFaults	The number of page faults that occurred during code execution.
dataPageFaults	The number of page faults that occurred during data access.
preemptions	The number of times this task was preempted.
cpuID	The ID of the last processor that ran this task.

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

MPQueueID

Multiprocessing Services manipulates message queues by passing a queue ID, which is defined by the `MPQueueID` type:

Multiprocessing Services Reference

```
typedef struct OpaqueMPQueueID* MPQueueID;
```

You obtain a queue ID by calling the function `MPCreateQueue` (page 67).

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

MPSemaphoreID

Multiprocessing Services manipulates semaphores by passing a semaphore ID, which is defined by the `MPSemaphoreID` type:

```
typedef struct OpaqueMPSemaphoreID* MPSemaphoreID;
```

You obtain a semaphore ID by calling the function `MPCreateSemaphore` (page 73).

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

MPEventID

Multiprocessing Services manipulates event groups by passing an event group ID, which is defined by the `MPEventID` type:

```
typedef struct OpaqueMPEventID* MPEventID;
```

You obtain an event group ID by calling the function `MPCreateEvent` (page 77).

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPCriticalRegionID

Multiprocessing Services manipulates critical region objects by passing a critical region ID, which is defined by the `MPCriticalRegionID` type:

```
typedef struct OpaqueMPCriticalRegionID*  MPCriticalRegionID;
```

You obtain a critical region ID by calling the function `MPCreateCriticalRegion` (page 84).

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

MPNotificationID

Multiprocessing Services manipulates kernel notifications by passing a notification ID, which is defined by the `MPNotificationID` type.

```
typedef struct OpaqueMPNotificationID *MPNotificationID;
```

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

MPTimerID

Multiprocessing Services manipulates timers by passing a timer ID, which is defined by the `MPTimerID` type:

```
typedef struct OpaqueMPTimerID*  MPTimerID;
```

You obtain a timer ID by calling the function `MPCreateTimer` (page 88).

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPOpaqueID

The Multiprocessing Services function handles a generic notification ID (that is, an ID that could be a queue ID, event ID, or semaphore ID) by passing an opaque ID, which is defined by the `MPOpaqueID` type:

```
typedef struct OpaqueMPOpaqueID* MPOpaqueID;
```

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

TaskStorageIndex

The Multiprocessing Services functions described in “Accessing Per-Task Storage Variables” (page 95) manipulate task storage index values using the `TaskStorageIndex` type, which has the following definition:

```
typedef UInt32 TaskStorageIndex;
```

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

TaskStorageValue

The Multiprocessing Services functions described in “Accessing Per-Task Storage Variables” (page 95) manipulate task storage values using the `TaskStorageValue` type, which has the following definition:

```
typedef UInt32 TaskStorageValue;
```

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPSemaphoreCount

The Multiprocessing Services function `MPCreateSemaphore` (page 73) handles the minimum and maximum semaphore values using the `MPSemaphoreCount` type, which has the following definition.

```
typedef ItemCount MPSemaphoreCount;
```

VERSION NOTES

Introduced with Multiprocessing Services 1.0.

MPEventFlags

The Multiprocessing Services functions described in “Handling Event Groups” (page 76) handle event information using the `MPEventFlags` type, which has the following definition:

```
typedef UInt32 MPEventFlags;
```

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPExceptionKind

The Multiprocessing Services function `MPTThrowException` (page 105) indicates the kind of exception thrown using the `MPExceptionKind` type, which has the following definition:

```
typedef UInt32 MPExceptionKind;
```

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPDebuggerLevel

The Multiprocessing Services function `MPRegisterDebugger` (page 109) indicates the debugger level using the `MPDebuggerLevel` type, which has the following definition:

```
typedef UInt32 MPDebuggerLevel;
```

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

MPRemoteProcedure

When calling the `MPRemoteCall` (page 111) function the remote procedure call (that is, the application-defined function) you want to designate must have the following type definition:

```
typedef void* (*MPRemoteProcedure)(void *parameter);
```

See `MyRemoteProcedure` (page 113) for more information about how to implement the application-defined function.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Constants

This section describes the constants available with Multiprocessing Services.

Multiprocessing Services Reference

- “Timer Duration Constants” (page 124)
- “Timer Option Masks” (page 125)
- “Memory Allocation Alignment Constants” (page 126)
- “Memory Allocation Option Constants” (page 127)
- “Task Creation Options” (page 128)
- “Task State Constants” (page 129)
- “Task Run State Constants” (page 130)
- “Task Information Structure Version Constant” (page 131)
- “Task Exception Disposal Constants” (page 131)
- “Remote Call Context Option Constants” (page 132)

Timer Duration Constants

Several Multiprocessing Services functions take a parameter of type `Duration`, which specifies the maximum time a task should wait for an event to occur. Multiprocessing Services recognizes four constants which you can use when specifying a duration. Note that you can use these constants in conjunction with other values to indicate specific wait intervals. For example, to wait 1 second, you can pass `kDurationMillisecond * 1000`.

```
enum {
    kDurationImmediate      = 0L,
    kDurationForever       = 0x7FFFFFFF,
    kDurationMillisecond    = 1,
    kDurationMicrosecond   = -1
};
```

Constant Descriptions

`kDurationImmediate`

The task times out immediately, whether or not the event has occurred. If the event occurred, the return status is `noErr`. If the event did not occur, the return status is `kMPTimeoutErr` (assuming no other errors occurred).

Multiprocessing Services Reference

<code>kDurationForever</code>	The task waits forever. The blocking call waits until either the event occurs, or until the object being waited upon (such as a message queue) is deleted.
<code>kDurationMillisecond</code>	The task waits one millisecond before timing out.
<code>kDurationMicrosecond</code>	The task waits one microsecond before timing out.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Timer Option Masks

When calling `MPArmTimer` (page 91), you can pass values of type `OptionBits` to specify any optional actions.

```
enum {
    kMPPreserveTimerIDMask      = 1L << 0,
    kMPTimeIsDeltaMask          = 1L << 1,
    kMPTimeIsDurationMask       = 1L << 2
};
```

Constant Descriptions

<code>kMPPreserveTimerID</code>	Specifying this mask prevents the timer from being deleted when it expires.
<code>kMPTimeIsDeltaMask</code>	Specifying this mask indicates that the specified time should be added to the previous expiration time to form the new expiration time. You can use this mask to compensate for timing drift caused by the finite amount of time required to arm the timer, receive the notification, and so on.
<code>kMPTimeIsDurationMask</code>	Specifying this mask indicates that the specified expiration time is of type <code>Duration</code> . You can use this mask to avoid

having to call time conversion routines when specifying an expiration time.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Memory Allocation Alignment Constants

When calling the `MPAllocateAligned` (page 99) function, you must specify the alignment of the desired memory block by passing a constant of type `UInt8` in the `alignment` parameter.

```
enum {
    /* Values for the alignment parameter to MPAllocateAligned.*/
    kMPAllocateDefaultAligned    = 0,
    kMPAllocate8ByteAligned      = 3,
    kMPAllocate16ByteAligned     = 4,
    kMPAllocate32ByteAligned     = 5,
    kMPAllocate1024ByteAligned   = 10,
    kMPAllocate4096ByteAligned   = 12,
    kMPAllocateMaxAlignment      = 16,
    kMPAllocateAltivecAligned    = kMPAllocate16ByteAligned,
    kMPAllocateVMXAligned        = kMPAllocateAltivecAligned,
    kMPAllocateVMPPageAligned    = 254,
    kMPAllocateInterlockAligned  = 255
};
```

Constant Descriptions

`kMPAllocateDefaultAligned`
Use the default alignment.

`kMPAllocate8ByteAligned`
Use 8-byte alignment.

`kMPAllocate16ByteAligned`
Use 16-byte alignment.

`kMPAllocate32ByteAligned`
Use 32-byte alignment.

Multiprocessing Services Reference

`kMPAllocate1024ByteAligned`

Use 1024-byte alignment.

`kMPAllocate4096ByteAligned`

Use 4096-byte alignment.

`kMPAllocateMaxAlignment`

Use the maximum alignment (65536 byte).

`kMPAllocateAltiVecAligned`

Use AltiVec alignment.

`kMPAllocateVMXAligned`

Use VMX (now called AltiVec) alignment.

`kMPAllocateVMPageAligned`

Use virtual memory page alignment. This alignment is set at runtime.

`kMPAllocateInterlockAligned`

Use interlock alignment, which is the alignment needed to allow the use of CPU interlock instructions (that is, `lwarx` and `stwcx`.) on the returned memory address. This alignment is set at runtime. In most cases you would never need to use this alignment.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Memory Allocation Option Constants

When calling the `MPAllocateAligned` (page 99) function, you can specify optional actions by passing a constant of type `OptionBits` in the `options` parameter.

```
enum {
    /* Values for the options parameter to MPAllocateAligned.*/
    kMPAllocateClearMask    = 0x0001,
    kMPAllocateGloballyMask = 0x0002,
    kMPAllocateResidentMask = 0x0004,
    kMPAllocateNoGrowthMask = 0x0010
};
```

Constant Descriptions`kMPAllocateClearMask`

Zero out the allocated memory block.

`kMPAllocateGloballyMask`Allocate memory from in memory space that is visible to all processes. Note that such globally-allocated space is not automatically reclaimed when the allocating process terminates. By default, `MPAllocateAligned` (page 99) allocates memory from process-specific (that is, not global) memory.`kMPAllocateResidentMask`

Allocate memory from resident memory only (that is, the allocated memory is not pageable).

`kMPAllocateNoGrowthMask`

Do not attempt to grow the pool of available memory. Specifying this option is useful, as attempting to grow memory may cause your task to block until such memory becomes available.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Task Creation Options

When creating a task using `MPCreateTask` (page 59), you can specify the following constants in the `options` field:

```
enum {
    kMPCreateTaskSuspendedMask          = 1L << 0,
    kMPCreateTaskTakesAllExceptionsMask = 1L << 1,
    kMPCreateTaskValidOptionsMask = kMPCreateTaskSuspendedMask |
                                    kMPCreateTaskTakesAllExceptionsMask
};
```

Constant Descriptions`kMPCreateTaskSuspendedMask`

The task is initially suspended.

Multiprocessing Services Reference

`kMPCreateTaskTakesAllExceptionsMask`

The task will take all exceptions, including those normally handled by the system, such as page faults.

`kMPCreateTaskValidOptionsMask`

Include all valid options for this task.

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

Task State Constants

When calling the `MPExtractTaskState` (page 107) or `MPSetTaskState` (page 108) functions you must specify what states you want to obtain or set by passing a constant of type `TaskStateKind` in the `kind` parameter.

```
enum {
    kMPTaskStateRegisters          = 0,
    kMPTaskStateFPU                = 1,
    kMPTaskStateVectors            = 2,
    kMPTaskStateMachine            = 3,
    kMPTaskState32BitMemoryException = 4,
    kMPTaskStateTaskInfo           = 5
};
typedef UInt32 MPTaskStateKind;
```

Constant Descriptions

`kMPTaskStateRegisters`

The task's general-purpose (GP) registers. The `RegisterInformationPowerPC` structure in `MachineExceptions.h` defines the format of this information.

`kMPTaskStateFPU`

The task's floating point registers. The `FPUInformationPowerPC` structure in `MachineExceptions.h` defines the format of this information.

`kMPTaskStateVectors`

The task's vector registers. The `VectorInformationPowerPC` structure in `MachineExceptions.h` defines the format of this information.

Multiprocessing Services Reference

`kMPTaskStateMachine`

The task's machine registers. The `MachineInformationPowerPC` structure in `MachineExceptions.h` defines the format of this information. Note that the MSR, `ExceptKind`, DSISR, and DAR registers are read-only.

`kMPTaskState32BitMemoryException`

The task's exception information for older 32-bit memory exceptions (that is, memory exceptions on 32-bit CPUs). The `MemoryExceptionInformation` structure in `MachineExceptions.h` defines the format of this information. This exception information is read-only.

`kMPTaskStateTaskInfo`

Static and dynamic information about the task, as described by the data structure `MPTaskInfo` (page 117). This task information is read-only.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Task Run State Constants

The task run state constants indicate the state of the task when returned as part of the `MPTaskInfo` (page 117) data structure.

```
enum {
    kMPTaskBlocked = 0,
    kMPTaskReady   = 1,
    kMPTaskRunning = 2
};
```

Constant Descriptions

<code>kMPTaskBlocked</code>	The task is blocked..
<code>kMPTaskReady</code>	The task is ready for execution.
<code>kMPTaskRunning</code>	The task is currently running.

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

Task Information Structure Version Constant

When returning a structure of type `MPTaskInfo` (page 117), Multiprocessing Services always indicates the version of the structure in the first field.

```
enum {
    kMPTaskInfoVersion = 2
};
```

Constant Descriptions

`kMPTaskInfoVersion` The current version of the task information structure.

VERSION NOTES

Introduced with Multiprocessing Services 2.1.

Task Exception Disposal Constants

When calling the `MPDisposeTaskException` (page 106) function, you must pass a constant of type `OptionBits` in the `action` parameter specifying what actions to take.

```
enum {
    /* Option bits for MPDisposeTaskException.*/
    kMPTaskResumeMask      = 0x0000,
    kMPTaskPropagateMask  = 1 << kMPTaskPropagate,
    kMPTaskResumeStepMask = 1 << kMPTaskResumeStep,
    kMPTaskResumeBranchMask = 1 << kMPTaskResumeBranch
};
```

Constant Descriptions

`kMPTaskResumeMask` Resume the task.

`kMPTaskPropagateMask` Propagate the exception to the next debugger level.

Multiprocessing Services Reference

kMPTaskResumeStepMask

Resume the task and enable single stepping.

kMPTaskResumeBranchMask

Resume the task and enable branch stepping.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Remote Call Context Option Constants

When making a remote call to an application-defined function using `MPRemoteCall` (page 111), you must pass a constant of type `MPRemoteContext` in the `context` parameter specifying which contexts are allowed to execute the function.

```
enum {
    kMPAnyRemoteContext          = 0,
    kMPOwningProcessRemoteContext = 1
};
typedef UInt8 MPRemoteContext;
```

Constant Descriptions

`kMPAnyRemoteContext`

Any cooperative context can execute the function. Note that the called function may not have access to any of the owning context's process-specific low-memory values.

`kMPOwningProcessRemoteContext`

Only the context that owns the task can execute the function.

VERSION NOTES

Introduced with Multiprocessing Services 2.0.

Result Codes

Many Multiprocessing Services functions return result codes. In addition, Multiprocessing Services functions may also return File Manager, Code Fragment Manager, and Process Manager result codes, which are described in *Inside Macintosh*.

Multiprocessing Services Reference

The most common result codes returned by Multiprocessing Services are listed below

noErr	No error
paramErr	Invalid parameter in function call
memFullErr	Out of memory
kMPDeletedErr	The desired notification the function was waiting upon was deleted.
kMPInsufficientResourcesErr	Could not complete task due to unavailable Multiprocessing Services resources.
	Note that many functions return this value as a general error when the desired action could not be performed.
kMPInvalidID	Invalid ID value (for example, an invalid message queue ID was passed to <code>MPNotifyQueue</code> (page 69)).
kMPTaskBlockedErr	The desired task is blocked.
kMPTaskStoppedErr	The desired task is stopped.
kMPTimeoutErr	The designated timeout interval passed before the function could take action.

Preemptive Task–Safe Mac OS System Software Functions

This appendix lists Mac OS 9 system software calls that you can call directly from a Multiprocessing Services task without using a remote procedure call. Other system software functions are not preemptive task–safe, so you must call them through a callback as described in “Making Remote Procedure Calls” (page 48).

These preemptive task–safe functions cause the calling task to block until the call is completed. However, unlike a remote procedure call, the called function does not have to wait until the calling process receives processor time; it executes asynchronously as if it was an interrupt call.

IMPORTANT

Before calling any system software function from a task, you should call `Gestalt` with the appropriate selector to determine whether it is safe to do so. ▲

This appendix contains the following tables of functions::

- Device Manager functions: Table A-1 (page 136)
- File Manager:
 - Table A-2 (page 136), “Preemptive task–safe HFS Plus parameter block functions”
 - Table A-3 (page 138), “Preemptive task–safe HFS Plus file system reference functions”
 - Table A-4 (page 139), “Preemptive task–safe HFS parameter block functions”
 - Table A-5 (page 142), “Preemptive task–safe high-level HFS functions”
 - Table A-6 (page 143), “Preemptive task–safe file system specification functions”
- Device Manager and File Manager shared functions: Table A-7 (page 143)

Preemptive Task–Safe Mac OS System Software Functions

- Miscellaneous functions: Table A-8 (page 144)

Table A-1 Preemptive task–safe Device Manager functions

Name	Comments
PBOpenSync	<p>Not supported in Carbon, so you should not use this function for File Manager open requests. For file open requests, you should call PBHOpenDFSync, PBHOpenDF, PBOpenForkSync, or FSOpenFork instead.</p> <p>May be handled by the File Manager in Mac OS 9.</p>
PBControlSync	
PBStatusSync	
PBKillIOSync	

Table A-2 Preemptive task–safe HFS Plus parameter block functions

Name	Comments
PBMakeFSRefSync	
PBMakeFSRefUnicodeSync	
PBCompareFSRefsSync	
PBCreateFileUnicodeSync	
PBCreateDirectoryUnicodeSync	
PBDeleteObjectSync	
PBMoveObjectSync	
PBExchangeObjectsSync	
PBRenameUnicodeSync	

Preemptive Task–Safe Mac OS System Software Functions

Table A-2 Preemptive task–safe HFS Plus parameter block functions

Name	Comments
PBGetCatalogInfoSync	
PBSetCatalogInfoSync	
PBOpenIteratorSync	
PBCloseIteratorSync	
PBGetCatalogInfoBulkSync	
PBCatalogSearchSync	
PBCreateForkSync	
PBDeleteForkSync	
PBIterateForksSync	
PBOpenForkSync	
PBReadForkSync	
PBWriteForkSync	
PBGetForkPositionSync	
PBSetForkPositionSync	
PBGetForkSizeSync	
PBSetForkSizeSync	
PBAllocateForkSync	
PBFlushForkSync	
PBCloseForkSync	
PBGetForkCBInfoSync	
PBGetVolumeInfoSync	
PBSetVolumeInfoSync	

Table A-3 Preemptive task–safe HFS Plus file system reference functions

Name	Comments
FSpMakeFSRef	
FSMakeFSRefUnicode	
FSCompareFSRefs	
FSCreateFileUnicode	
FSCreateDirectoryUnicode	
FSDeleteObject	
FSMoveObject	
FSExchangeObjects	
FSRenameUnicode	
FSGetCatalogInfo	
FSSetCatalogInfo	
FSOpenIterator	
FSCloseIterator	
FSGetCatalogInfoBulk	
FSCatalogSearch	
FSCreateFork	
FSDeleteFork	
FSIterateForks	
FSOpenFork	
FSReadFork	
FSWriteFork	
FSGetForkPosition	
FSSetForkPosition	
FSGetForkSize	
FSSetForkSize	
FSAllocateFork	

Preemptive Task–Safe Mac OS System Software Functions

Table A-3 Preemptive task–safe HFS Plus file system reference functions

Name	Comments
FSFlushFork	
FSCloseFork	
FSGetForkCBInfo	
FSGetVolumeInfo	
FSSetVolumeInfo	
FSGetDataForkName	
FSGetResourceForkName	

Table A-4 Preemptive task–safe HFS parameter block functions

Name	Comments
PBXGetVolInfoSync	
PBFlushVolSync	
PBAllocateSync	
PBGetEOFSync	
PBSetEOFSync	
PBGetFPosSync	
PBSetFPosSync	
PBFlushFileSync	
PBCatSearchSync	
PBHSetVolSync	
PBHGetVolSync	
PBCatMoveSync	
PBDirCreateSync	
PBGetFCBInfoSync	
PBGetCatInfoSync	

Preemptive Task–Safe Mac OS System Software Functions

Table A-4 Preemptive task–safe HFS parameter block functions

Name	Comments
PBSetCatInfoSync	
PBAllocContigSync	
PBLockRangeSync	
PBUnlockRangeSync	
PBSetVInfoSync	
PBHVInfoSync	
PBHOOpenSync	
PBHOOpenRFSync	
PBHOOpenDFSsync	
PBHCreateSync	
PBHDeleteSync	
PBHRenameSync	
PBHRstFlockSync	
PBHSetFlockSync	
PBHGetFInfoSync	
PBHSetFInfoSync	
PBMakeFSSpecSync	
PBHGetVolParmsSync	
PBHGetLogInInfoSync	
PBHGetDirAccessSync	
PBHSetDirAccessSync	
PBHMapIDSync	
PBHMapNameSync	
PBHCopyFileSync	
PBHMoveRenameSync	
PBHOOpenDenySync	
PBHOOpenRFDenySync	
PBGetXCatInfoSync	

Preemptive Task–Safe Mac OS System Software Functions

Table A-4 Preemptive task–safe HFS parameter block functions

Name	Comments
PBExchangeFilesSync	
PBCreateFileIDRefSync	
PBResolveFileIDRefSync	
PBDeleteFileIDRefSync	
PBGetForeignPrivsSync	
PBSetForeignPrivsSync	
PBDTAddIconSync	
PBDTGetIconSync	
PBDTGetIconInfoSync	
PBDTAddAPPLSync	
PBDTRemoveAPPLSync	
PBDTGetAPPLSync	
PBDTSetCommentSync	
PBDTRemoveCommentSync	
PBDTGetCommentSync	
PBDTFlushSync	
PBDTResetSync	
PBDTGetInfoSync	
PBDTDeleteSync	
PBShareSync	
PBUnshareSync	
PBGetUGEntrySync	

Table A-5 Preemptive task–safe high-level HFS functions

Name	Comments
FSClose	
FSRead	
FSWrite	
FlushVol	
Allocate	
GetEOF	
SetEOF	
GetFPos	
SetFPos	
HGetVol	
HSetVol	
HOpen	
HOpenDF	
HOpenRF	
AllocContig	
HCreate	
DirCreate	
HDelete	
HGetFInfo	
HSetFInfo	
HSetFLock	
HRstFLock	
HRename	
CatMove	

Preemptive Task–Safe Mac OS System Software Functions

Table A-6 Preemptive task–safe file system specification functions

Name	Comments
FSMakeFSSpec	
FSpOpenDF	
FSpOpenRF	
FSpCreate	
FSpDirCreate	
FSpDelete	
FSpGetFInfo	
FSpSetFInfo	
FSpSetFLock	
FSpRstFLock	
FSpRename	
FSpCatMove	
FSpExchangeFiles	

Table A-7 Preemptive task–safe functions shared by the Device Manager and File Manager

Name	Comments
PBCloseSync	
PBReadSync	
PBWriteSync	

Preemptive Task–Safe Mac OS System Software Functions

Table A-8 Miscellaneous preemptive task–safe functions

Name	Comments
DTInstall	Deferred Task Manager function
WakeUpProcess	Process Manager function

Calculating the Intertask Signaling Time

When using Multiprocessing Services tasks, the amount of time used by the task should be much greater than the time taken to pass notifications to the task. This intertask signaling time is generally between 20 and 50 microseconds. If you want to explicitly calculate the signaling time, you can use the code in Listing B-1 to do so.

Listing B-1 Calculating the intertask signaling time

```
#include <Multiprocessing.h>

#include <Types.h>
#include <stdio.h>
#include <stdlib.h>
#include <sioux.h>
#include <math64.h>
#include <DriverServices.h>

enum {
    aQueue = 0,
    aSemaphore,
    anEvent
} reflectOP;

MPOpaqueID waiterID, postID;

static OSStatus Reflector ( )
{
    void*p1,*p2,*p3;
    MPEventFlags events;
```

APPENDIX B

Calculating the Intertask Signaling Time

```
while (true)
{
    switch (reflectOP)
    {
        case aQueue:
            MPWaitOnQueue((MPQueueID) waiterID, &p1, &p2, &p3, kDurationForever);
            break;

        case aSemaphore:
            MPWaitOnSemaphore((MPSemaphoreID) waiterID, kDurationForever);
            break;

        case anEvent:
            MPWaitForEvent((MPEventID) waiterID, &events, kDurationForever);
            break;

        default:
            return -123;
    }

    switch (reflectOP)
    {
        case aQueue:
            MPNotifyQueue((MPQueueID) postID, &p1, &p2, &p3);
            break;

        case aSemaphore:
            MPSignalSemaphore((MPSemaphoreID) postID);
            break;

        case anEvent:
            MPSetEvent((MPEventID) postID, 0x01010101);
            break;

        default:
            return -123;
    }
}

return -123;
}
```

APPENDIX B

Calculating the Intertask Signaling Time

```
static float HowLong(
    AbsoluteTime endTime,
    AbsoluteTime bgnTime
)
{
    AbsoluteTime absTime;
    Nanoseconds nanosec;

    absTime = SubAbsoluteFromAbsolute(endTime, bgnTime);
    nanosec = AbsoluteToNanoseconds(absTime);
    return (float) UnsignedWideToUInt64( nanosec ) / 1000.0;
}

void main ( void )
{
    OSStatus      err;
    MPTaskID      task;
    UInt32        i, count;
    void          *p1,*p2,*p3;
    MPEventFlags events;
    AbsoluteTime  nowTime, bgnTime;
    float         uSec;
    char          buff[10];

    /* Set the console window defaults */
    /* (this is a Metrowerks CodeWarrior thing). */
    SIOUXSettings.autocloseonquit = true;
    SIOUXSettings.asktosaveonclose = false;
    SIOUXSettings.showstatusline = false;
    SIOUXSettings.columns = 100;
    SIOUXSettings.rows = 20;
    SIOUXSettings.fontsize = 10;
    // SIOUXSettings.fontid = monaco;
    SIOUXSettings.standalone = true;

    // DebugStr ( "\pStarting" );
```

Calculating the Intertask Signaling Time

```

/* Can't get very far without this one. */
if (!MPLibraryIsLoaded())
{
    printf("The MP library did not load.\n");
    return;
}

/* Find the overhead up UpTime. Perform a bunch of calls to average out */
/* cache effects. */
printf("\n");

bgnTime = UpTime();
for (i=0; i<16; i++)
{
    nowTime = UpTime();
}

uSec = HowLong(nowTime, bgnTime);
uSec /= 16.0;
printf(" UpTime overhead: %.3f usec \n", uSec);

/* Time intertask communication. */
printf("\n Queues\n");

reflectOP = aQueue;
MPCreateQueue((MPQueueID*) &waiterID);
MPCreateQueue((MPQueueID*) &postID);

bgnTime = UpTime();
err = MPCreateTask( Reflector,
                    NULL,
                    0,
                    NULL,
                    0,
                    0,
                    kNilOptions,
                    &task );
nowTime = UpTime();
uSec = HowLong(nowTime, bgnTime);

```

APPENDIX B

Calculating the Intertask Signaling Time

```
printf(" MPCreateTask overhead: %.3f usec (may vary significantly) \n", uSec);
if (err != noErr)
{
    printf(" Task not created!\n");
    return;
}

count = 100000;
bgnTime = UpTime();
for (i=0; i<count; i++)
{
    MPNotifyQueue((MPQueueID) waiterID, 0, 0, 0);
    while (true)
    {
        err = MPWaitOnQueue((MPQueueID) postID, &p1, &p2, &p3, kDurationImmediate);
        if (err != kMPTimeoutErr) break;
    }
}
nowTime = UpTime();
uSec = HowLong(nowTime, bgnTime);
uSec /= ((float) count / 2.0); // Two trips.
printf(" Intertask signalling using queues overhead: %.3f usec \n", uSec);

/* Time intertask communication.
*/
MPTerminateTask(task, 123);
printf("\n Semaphores\n");

reflectOP = aSemaphore;

MPCreateSemaphore(1, 0, (MPSemaphoreID*) &waiterID);
MPCreateSemaphore(1, 0, (MPSemaphoreID*) &postID);

bgnTime = UpTime();
err = MPCreateTask( Reflector,
                    NULL,
                    0,
                    NULL,
                    0,
                    0,
                    kNilOptions,
```

APPENDIX B

Calculating the Intertask Signaling Time

```
        &task );
nowTime = UpTime();
uSec = HowLong(nowTime, bgnTime);
printf(" MPCreateTask overhead: %.3f usec (may vary significantly) \n", uSec);
if (err != noErr)
{
    printf(" Task not created!\n");
    return;
}

count = 100000;
bgnTime = UpTime();
for (i=0; i<count; i++)
{
    MPSignalSemaphore((MPSemaphoreID) waiterID);
    while (true)
    {
        err = MPWaitOnSemaphore((MPSemaphoreID) postID, kDurationImmediate);
        if (err != kMPTimeoutErr) break;
    }
}
nowTime = UpTime();
uSec = HowLong(nowTime, bgnTime);
uSec /= ((float) count / 2.0); // Two trips.
printf(" Intertask signalling using semaphores overhead: %.3f usec \n", uSec);

/* Time intertask communication. */
MPTerminateTask(task, 123);
printf("\n Event Groups\n");

reflectOP = anEvent;
MPCreateEvent((MPEventID*) &waiterID);
MPCreateEvent((MPEventID*) &postID);

bgnTime = UpTime();
err = MPCreateTask( Reflector,
                    NULL,
                    0,
                    NULL,
                    0,
```

APPENDIX B

Calculating the Intertask Signaling Time

```
        0,
        kNilOptions,
        &task );
nowTime = UpTime();
uSec = HowLong(nowTime, bgnTime);
printf(" MPCreateTask overhead: %.3f usec (may vary significantly) \n", uSec);
if (err != noErr)
{
    printf(" Task not created!\n");
    return;
}

count = 100000;
bgnTime = UpTime();
for (i=0; i<count; i++)
{
    MPSetEvent((MPEventID) waiterID, 0x01);
    while (true)
    {
        err = MPWaitForEvent((MPEventID) postID, &events, kDurationImmediate);
        if (err != kMPTimeoutErr) break;
    }
}
nowTime = UpTime();
uSec = HowLong(nowTime, bgnTime);
uSec /= ((float) count / 2.0); // Two trips.
printf(" Intertask signalling using events overhead: %.3f usec \n", uSec);

gets(buff);
}
```


Changes From Previous Versions of Multiprocessing Services

Multiprocessing Services 2.1 supports all the functions available with version 2.0. For compatibility between version 2.0 and older versions, see Table C-3, Table C-4, and Table C-5.

Table C-1 lists Multiprocessing Services functions introduced with version 2.1:

Table C-1 New functions introduced with version 2.1

Name	Comments
MPGetNextCpuID (page 58)	
MPGetNextTaskID (page 64)	
MPCreateNotification (page 80)	For manipulating kernel notifications. See “Kernel Notifications” (page 20) for more information about this notification mechanism.
MPDeleteNotification (page 81)	
MPModifyNotification (page 82)	
MPCauseNotification (page 83)	

Table C-2 lists Multiprocessing Services functions that were introduced in version 2.0.

Table C-2 Functions introduced with version 2.0

Name	Comments
MPProcessorsScheduled (page 57)	
MPSetTaskWeight (page 65)	
MPTaskIsPreemptive (page 66)	

Changes From Previous Versions of Multiprocessing Services

Table C-2 Functions introduced with version 2.0 (continued)

Name	Comments
MPAllocateTaskStorageIndex	(page 96)
MPDeallocateTaskStorageIndex	(page 96)
MPSetTaskStorageValue	(page 97)
MPGetTaskStorageValue	(page 98)
MPSetQueueReserve	(page 72)
MPCreateEvent	(page 77)
MPDeleteEvent	(page 77)
MPSetEvent	(page 78)
MPWaitForEvent	(page 79)
UpTime	(page 93)
DurationToAbsolute	(page 94)
AbsoluteToDuration	(page 95)
MPDelayUntil	(page 87)
MPCreateTimer	(page 88)
MPDeleteTimer	(page 89)
MPSetTimerNotify	(page 90)
MPArmTimer	(page 91)
MPCancelTimer	(page 92)
MPSetExceptionHandler	(page 104)
MPThrowException	(page 105)
MPDisposeTaskException	(page 106)
MPExtractTaskState	(page 107)
MPSetTaskState	(page 108)
MPRegisterDebugger	(page 109)
MPRegisterDebugger	(page 109)

Changes From Previous Versions of Multiprocessing Services

Table C-2 Functions introduced with version 2.0 (continued)

Name	Comments
MPAllocateAligned (page 99)	Preferred over MPAllocate.
MPGetAllocatedBlockSize (page 101)	
MPBlockClear (page 102)	
MPDataToCode (page 103)	
MPRemoteCall (page 111)	Preferred over _MPRPC

Table C-3 lists the functions that were introduced in version 1.0 that are still supported in version 2.0.

Table C-3 Older functions supported in version 2.0

Name	Comments
MPProcessors (page 57)	
MPCreateTask (page 59)	
MPTerminateTask (page 62)	
MPCurrentTaskID (page 63)	
MPYield (page 65)	
MPExit (page 63)	
MPCreateQueue (page 67)	
MPDeleteQueue (page 68)	
MPNotifyQueue (page 69)	
MPWaitOnQueue (page 70)	
MPCreateSemaphore (page 73)	
MPCreateBinarySemaphore	In C, a macro that calls MPCreateSemaphore (page 73).
MPDeleteSemaphore (page 74)	

Changes From Previous Versions of Multiprocessing Services

Table C-3 Older functions supported in version 2.0 (continued)

Name	Comments
MPSignalSemaphore (page 75)	
MPWaitOnSemaphore (page 75)	
MPCreateCriticalRegion (page 84)	
MPDeleteCriticalRegion (page 84)	
MPEnterCriticalRegion (page 85)	
MPExitCriticalRegion (page 86)	
MPAllocate	Deprecated. Use MPAllocateAligned instead.
MPFree (page 100)	
MPBlockCopy (page 101)	
MPLibraryIsLoaded	In C, a macro that checks to see if the MPProcessors symbol is resolved.

Table C-4 shows unofficial functions included in earlier header files that remain supported in version 2.0. Note, however, that future versions may not support these functions.

Table C-4 Unofficial functions still supported in version 2.0

Name	Comments
_MPRPC	Deprecated. Use MPRemoteCall (page 111) instead.
_MPAllocateSys	Deprecated. Use MPAllocateAligned (page 99) instead.
_MPTaskIsToolboxSafe	
_MPLibraryVersion	
_MPLibraryIsCompatible	

APPENDIX C

Changes From Previous Versions of Multiprocessing Services

Table C-5 shows functions used for debugging that are no longer supported in version 2.0. You can access these functions for older builds if you `#define MPIncludeDefunctServices` to be nonzero.

Table C-5 Debugging functions unsupported in version 2.0

Name	Comments
<code>_MPInitializePrintf</code>	
<code>_MPPrintf</code>	
<code>_MPDebugStr</code>	
<code>_MPStatusPString</code>	
<code>_MPStatusCString</code>	

Document Version History

This document has had the following releases:

Table D-1 Multiprocessing Services documentation revision history

Version	Notes
October 21, 1999	<p>Revised for Multiprocessing Services 2.1. The following changes were made from the 2.0 version:</p> <p>Added “Gestalt Constants” (page 55), which you can use to determine the availability of preemptively safe Mac OS system software functions, and added a list of these functions in Appendix A. Changed text in “Criteria for Creating Tasks” (page 31) and “Making Remote Procedure Calls” (page 48) to reflect these additions.</p> <p>Added new functions: <code>MPGetNextCpuID</code> (page 58) and <code>MPGetNextTaskID</code> (page 64).</p> <p>Added new kernel notification functions: <code>MPCreateNotification</code> (page 80), <code>MPCauseNotification</code> (page 80), <code>MPModifyNotification</code> (page 82), and <code>MPCauseNotification</code> (page 83). Added “Kernel Notifications” (page 20) section to Chapter 2.</p> <p>Added new options for <code>MPCreateTask</code> (page 59). See “Task Creation Options” (page 128) for details.</p> <p>Added new task information structure, <code>MPTaskInfo</code> (page 117), for use with <code>MPExtractTaskState</code> (page 107).</p> <p>Added new task run state constants to be used with the new <code>MPTaskInfo</code> (page 117) structure. See “Task Run State Constants” (page 130) for details.</p> <p>Added glossary entries for coherence groups and kernel notifications.</p>

Table D-1 Multiprocessing Services documentation revision history (continued)

Version	Notes
April 30, 1999	<p data-bbox="368 345 1265 402">Initial public release. The following changes were made from the previous (seed draft) version:</p> <p data-bbox="368 425 1265 539">Added Chapter 1, “Introduction,” Chapter 2, “About Multitasking on the Mac OS,” and Chapter 3, “Using Multiprocessing Services,” which include introductory information, conceptual information, programming discussions, and sample code.</p> <p data-bbox="368 562 1265 619">Added versioning information to functions, data types, and constants in Chapter 4, “Multiprocessing Services Reference.”</p> <p data-bbox="368 642 1265 725">Added discussion and parameter information to <code>MPTaskIsPreemptive</code> (page 66) indicating how and why you can specify <code>kInvalidID</code> to determine the preemptiveness of the current task. Also added Version Notes section.</p> <p data-bbox="368 747 1265 887">Correction in <code>MPWaitOnQueue</code> (page 70), <code>MPWaitOnSemaphore</code> (page 75), <code>MPWaitForEvent</code> (page 79), and <code>MPEnterCriticalRegion</code> (page 85): When calling from a cooperative task, you should specify only <code>kDurationImmediate</code> waits; others are allowable, but they will cause the task to block.</p> <p data-bbox="368 910 1265 1049">Added information stating that setting event bits in <code>MPSetEvent</code> (page 78) and obtaining and clearing and event group in <code>MPWaitForEvent</code> (page 79) are atomic operations. For example, bits cannot be set between when a task obtains an event group and when the event group is cleared, so no data can be lost.</p> <p data-bbox="368 1072 1265 1128">Changed wording in <code>MPDelayUntil</code> (page 87) to clarify that you must indicate a specific time to unblock the task, not a duration.</p> <p data-bbox="368 1151 1265 1234">Changed wording for <code>MPAllocateTaskStorageIndex</code> (page 96) and <code>MPDeallocateTaskStorageIndex</code> (page 96) to indicate that these functions do not actually allocate or deallocate memory.</p> <p data-bbox="368 1257 1265 1379">Added disclaimer to <code>MPTThrowException</code> (page 105) indicating that you should throw an exception to a task to stop it only if you are debugging and plan to examine the state of the task. Otherwise, you should block the task using a traditional notification method (such as a message queue).</p> <p data-bbox="368 1402 1265 1458">Modified discussion of <code>MPSetExceptionHandler</code> (page 104) to indicate the format of the message sent to the exception handler.</p> <p data-bbox="368 1481 1265 1525">Discussion of informative messages in <code>MPRegisterDebugger</code> (page 109) removed to reflect status as of version 2.0.</p>

Table D-1 Multiprocessing Services documentation revision history (continued)

Version	Notes
	<p>Added information to <code>MPExtractTaskState</code> (page 107) and <code>MPSetTaskState</code> (page 108) indicating that attempting to set or read state information for a nonsuspended task returns the error <code>kMPIInsufficientResourcesErr</code>.</p> <p>Added information to <code>MPSetTaskState</code> (page 108) and “Task State Constants” (page 129): the exception state information and some machine registers (MRS, <code>ExceptKind</code>, DSISR, and DAR) are read-only. Attempting to set the exception state information will return an error. Attempts to change the MRS, <code>ExceptKind</code>, DSISR, and DAR registers will simply have no effect.</p> <p>Discussion in <code>MPRemoteCall</code> (page 111) modified to reflect this clarification: If you specify that the function should execute in the same context that owns the task, the function has access to data available to the main application (just as if the application had called the function). However, the function cannot execute until the owning context becomes active (and then not until the application calls <code>WaitNextEvent</code>).</p> <p>Atomic memory operations mentioned in <code>MPRemoteCall</code> (page 111) and “Making Remote Procedure Calls” (page 48) are now located in <code>InterfaceLib</code>, not the Driver Services Library.</p> <p>Data types <code>MPAddressSpaceID</code> and <code>MPCpuID</code> removed to reflect status as of version 2.0.</p> <p>Clarified that you can use the constants in “Timer Duration Constants” (page 124) to specify any number of waiting times by adding multipliers.</p> <p>Correction in “Memory Allocation Alignment Constants” (page 126): CPU interlock instruction <code>swarx</code> should be <code>stwcx</code>.</p> <p>Specified the <code>MachineExceptions.h</code> structures that correspond to the state information constants in “Task State Constants” (page 129)</p> <p>Removed nonbitmask values (<code>kMPTaskPropagate</code>, <code>kMPTaskResumeStep</code>, and <code>kMPTaskResumeBranch</code>) and descriptions from “Task Exception Disposal Constants” (page 131).</p>
Feb. 26, 1999	First seed draft release

Glossary

atomic operation An action that executes as one indivisible sequence. For example, obtaining and clearing an event group occurs as an atomic operation. It is not possible for another task to set an event bit *after* the event group is obtained but *before* it is cleared.

cooperative multitasking A multitasking model that requires a task to voluntarily suspend its execution in order to give processor time to other tasks.

blocked The state where a task is not executing because it is waiting for data or a resource to become available. Blocked tasks are not assigned processor time by the scheduler.

coherence group The set of processors and other bus controllers that have cache-coherent access to memory. Mac OS 9 defines only one coherence group, which is all the processors that can access internal memory. Other coherence groups are possible; for example, a PCI card with its own memory and processors can comprise a coherence group.

context The execution environment for an application. Each running application has its own context which is scheduled (cooperatively or preemptively) by system software. For example, in Mac OS 8, contexts are cooperatively-scheduled by the Process Manager. Also called a *process* or an *execution context*.

critical region A section of executable code that can only be accessed by one task at a time.

event group A set of flags which you can use to signal events for a task.

kernel notification A set of simple notification mechanisms (semaphores, message groups, and event groups) that can be signaled together. A kernel notification can contain one of each type of simple notification mechanism.

latency A significant time delay between when your code requests an action and when the action actually occurs. For example, there is typically a latency between when you request that a preemptive task be terminated and when it actually terminates.

Mac OS task The preemptive task that contains all the cooperatively scheduled programs handled by the Process Manager.

message A 96-bit block of data that passes information to a task. Messages are passed to message queues which then can be read by a task. *See also* message queue.

message queue A storage location for messages. Tasks can place or retrieve messages from a queue. *See also* message.

multiprocessing The ability of a computer and operating system to use more than one microprocessor at a time.

Multiprocessing Services Apple Computer's preemptive multitasking technology for Macintosh computers.

multitasking The ability to handle several tasks at a time.

notification mechanism An indicator that passes information to a task. Some notification mechanisms include messages, semaphores, and event groups.

nonreentrant function A function that cannot be called simultaneously or sequentially by different preemptive tasks. Typically functions that manipulate global data are nonreentrant.

preemptive multitasking The ability of an operating system to divide processor time among many tasks, allowing them to execute in a simultaneous or near-simultaneous manner.

preemptive scheduling A multitasking model that allows a scheduler to determine the amount of processor time to assign to tasks. The scheduler uses well-defined rules to determine when to stop execution of one task and resume another.

process *See context.*

remote procedure call A call by a preemptive task that must execute in a nonpreemptive environment. Typically you use remote procedure calls to call nonreentrant functions.

scheduler An authority that assigns processor time to individual tasks. Also called a *task scheduler*.

semaphore A counter that can be incremented or decremented to signal an event to a task.

task An individual unit of program execution. Each task has its own stack and register set.

symmetric multiprocessing A multiprocessing technique where tasks are divided equally among all available processors. This method is more efficient than *asymmetric multiprocessing* where one master processor assigns work to a number of slave processors.

thread As defined for the Mac OS operating system, a task that is cooperatively-scheduled by the Thread Manager. Compare **task**.

Index

Numerals

68K code in preemptive tasks 29

A

absolute time, defined 47
AbsoluteToDuration function 95
address spaces 16, 49
allocating memory in tasks 46

B

"bank line" tasking architecture 24
binary semaphores 19
blocked task, defined 17

C

calling 68K code 48
calling nonreentrant functions 48
changes from previous versions of
 Multiprocessing Services 153
checking for the availability of Multiprocessing
 Services 32
coherence group, defined 115
critical regions
 creating 45
 defined 21

D

debuggers 49
determining the number of processors 32
"divide and conquer" tasking architecture 23
DurationToAbsolute function 94

E

event groups, defined 20
exception handlers 49

H

hardware requirements 30

I

interrupt handlers 44
intertask signaling time, calculating 145

K

kDurationForever constant 125
kDurationImmediate constant 124
kDurationMicrosecond constant 125
kDurationMillisecond constant 125
kernel notifications, defined 20
kMPAllocate1024ByteAligned constant 127
kMPAllocate16ByteAligned constant 126
kMPAllocate32ByteAligned constant 126
kMPAllocate4096ByteAligned constant 127

kMPAllocate8ByteAligned **constant** 126
 kMPAllocateAltivecAligned **constant** 127
 kMPAllocateClearMask **constant** 128
 kMPAllocateDefaultAligned **constant** 126
 kMPAllocateGloballyMask **constant** 128
 kMPAllocateInterlockAligned **constant** 127
 kMPAllocateMaxAlignment **constant** 127
 kMPAllocateNoGrowthMask **constant** 128
 kMPAllocateResidentMask **constant** 128
 kMPAllocateVMPageAligned **constant** 127
 kMPAllocateVMXAligned **constant** 127
 kMPAnyRemoteContext **constant** 132
 kMPCreateTaskSuspendedMask **constant** 128
 kMPCreateTaskTakesAllExceptionsMask
 constant 129
 kMPCreateTaskValidOptionsMask **constant** 129
 kMPDeletedErr **result code** 134
 kMPInsufficientResourcesErr **result code** 134
 kMPInvalidID **result code** 134
 kMPOwningProcessRemoteContext **constant** 132
 kMPPreserveTimerID **constant** 125
 kMPTaskBlocked **constant** 130
 kMPTaskBlockedErr **result code** 134
 kMPTaskInfoVersion **constant** 131
 kMPTaskPropagateMask **constant** 131
 kMPTaskReady **constant** 130
 kMPTaskResumeBranchMask **constant** 132
 kMPTaskResumeMask **constant** 131
 kMPTaskResumeStepMask **constant** 132
 kMPTaskRunning **constant** 130
 kMPTaskState32BitMemoryException
 constant 130
 kMPTaskStateFPU **constant** 129
 kMPTaskStateMachine **constant** 130
 kMPTaskStateRegisters **constant** 129
 kMPTaskStateTaskInfo **constant** 130
 kMPTaskStateVectors **constant** 129
 kMPTaskStoppedErr **result code** 134
 kMPTimeIsDeltaMask **constant** 125
 kMPTimeIsDurationMask **constant** 125
 kMPTimeoutErr **result code** 134

M

Mac OS task, **defined** 16
 message queues
 defined 19
 used when creating tasks 35
 microprocessors *See* processors
 MPAllocateAligned **function** 99
 MPAllocateTaskStorageIndex **function** 96
 MPArmTimer **function** 91
 MPBlockClear **function** 102
 MPBlockCopy **function** 101
 MPCancelTimer **function** 92
 MPCauseNotification **function** 83
 MPCreateCriticalRegion **function** 84
 MPCreateEvent **function** 77
 MPCreateNotification **function** 80
 MPCreateQueue **function** 67
 MPCreateSemaphore **function** 73
 MPCreateTask **function** 59
 MPCreateTimer **function** 88
 MPCriticalRegionID **type** 120
 MPCurrentTaskID **function** 63
 MPDataToCode **function** 103
 MPDeallocateTaskStorageIndex **function** 96
 MPDebuggerLevel **type** 123
 MPDelayUntil **function** 87
 MPDeleteCriticalRegion **function** 84
 MPDeleteEvent **function** 77
 MPDeleteNotification **function** 81
 MPDeleteQueue **function** 68
 MPDeleteSemaphore **function** 74
 MPDeleteTimer **function** 89
 MPDisposeTaskException **function** 106
 MPEnterCriticalRegion **function** 85
 MPEventFlags **type** 122
 MPEventID **type** 119
 MPExceptionKind **type** 122
 MPExitCriticalRegion **function** 86
 MPExit **function** 63
 MPExtractTaskState **function** 107
 MPFree **function** 100
 MPGetAllocatedBlockSize **function** 101
 MPGetNextCpuID **function** 58
 MPGetNextTaskID **function** 64

MPGetTaskStorageValue function 98
 MPMotifyNotification function 82
 MPNotifyQueue function 69
 MPOpaqueID type 121
 MPProcessID type 114
 MPProcessors function 57
 MPProcessorsScheduled function 57
 MPQueueID type 118
 MPRegisterDebugger function 109
 MPRemoteCall function 111
 MPSemaphoreCount type 122
 MPSemaphoreID type 119
 MPSetEvent function 78
 MPSetExceptionHandler function 104
 MPSetQueueReserve function 72
 MPSetTaskState function 108
 MPSetTaskStorageValue function 97
 MPSetTaskWeight function 65
 MPSetTimerNotify function 90
 MPSignalSemaphore function 75
 MPTaskID type 115
 MPTaskInfo type 117
 MPTaskIsPreemptive function 66
 MPTaskOptions type 116
 MPTaskWeight type 117
 MPTerminateTask function 62
 MPThrowException function 105
 MPTimerID type 120
 MPUnregisterDebugger function 110
 MPWaitForEvent function 79
 MPWaitOnQueue function 70
 MPWaitOnSemaphore function 75
 MPYield function 65
 multiprocessing, defined 14
 Multiprocessing Services on Mac OS 8 versus
 Mac OS X 30
 multitasking
 cooperative versus preemptive 13
 defined 13

N

new functions added with version 2.0 153
 notification methods 18
 notification queue 35

P

"pipeline" tasking architecture 25
 preemptive task–safe system software functions
 31, 48, 135
 preemptive tasks *See* tasks
 processors
 determining the number of 32
 versus number of tasks 22

Q

queues *See* message queues

R

recursion in critical regions 21
 remote procedure calls 48

S

schedulers 16
 semaphores
 defined 19
 used for periodic actions 43
 shared resources 18
 synchronization of tasks 18, 40
 system requirements 30
 system software, preemptive task–safe functions
 in 31, 48, 135

T

tasking architectures

- multiple independent tasks 22
- parallel tasks with a single set of I/O buffers 24
- parallel tasks with parallel I/O buffers 23
- sequential tasks 25

tasks

- blocked 17
- compared to threads 14
- creating 33
- defined 13
- intertask signaling time 145
- notifying at interrupt time 44
- setting weight of 36
- synchronization of 18, 40
- terminating 38
- versus number of processors 22

task schedulers 16task-specific storage 46TaskStorageIndex type 121TaskStorageValue type 121task weight 36terminating tasks 38termination queue 35Thread Manager, compared to Multiprocessing Services 14timers 47

U

UpTime function 93

V

version 2.0, new functions added for 153virtual memory 30

W

waiting on a queue 38, 42

This Apple manual was written, edited, and composed on a desktop publishing system using Apple Macintosh computers and FrameMaker software. Line art was created using Adobe™ Illustrator and Adobe Photoshop.

Text type is Palatino® and display type is Helvetica®. Bullets are ITC Zapf Dingbats®. Some elements, such as program listings, are set in Adobe Letter Gothic.

WRITER

Jun Suzuki

ILLUSTRATOR

Ruth Anderson

DEVELOPMENTAL EDITOR

Laurel Rezeau

PRODUCTION EDITOR

Gerri Gray

Special thanks to René Vega, Alan Lillich, Jim Murphy, Roger Pantos, and George Warner.