

VISUALDSP⁺™ 3.x

Kernel (VDK) User's Guide

Revision 4.1, April 2003

Part Number
82-000349-07

Analog Devices, Inc.
Digital Signal Processor Division
One Technology Way
Norwood, Mass. 02062-9106



Copyright Information

© 2003 Analog Devices, Inc., ALL RIGHTS RESERVED. This document may not be reproduced in any form without prior, express written consent from Analog Devices, Inc.

Printed in the USA.

Disclaimer

Analog Devices, Inc. reserves the right to change this product without prior notice. Information furnished by Analog Devices is believed to be accurate and reliable. However, no responsibility is assumed by Analog Devices for its use; nor for any infringement of patents or other rights of third parties which may result from its use. No license is granted by implication or otherwise under the patent rights of Analog Devices, Inc.

Trademark and Service Mark Notice

The Analog Devices logo, VisualDSP, the VisualDSP logo, Blackfin, the Blackfin logo, SHARC, the SHARC logo, TigerSHARC, and the TigerSHARC logo are registered trademarks of Analog Devices, Inc.

VisualDSP++, the VisualDSP++ logo, CROSSCORE, the CROSSCORE logo, and EZ-KIT Lite are trademarks of Analog Devices, Inc.

All other brand and product names are trademarks or service marks of their respective owners.

CONTENTS

PREFACE

Purpose of this Manual	xvii
Intended Audience	xviii
Manual Contents	xviii
What's New in this Manual	xix
Technical or Customer Support	xx
Supported Processors	xxi
Product Information	xxi
MyAnalog.com	xxi
DSP Product Information	xxii
Related Documents	xxii
Online Documentation	xxiii
From VisualDSP++	xxiv
From Windows	xxiv
From the Web	xxv
Printed Manuals	xxv
VisualDSP++ Documentation Set	xxv
Hardware Manuals	xxv
Data Sheets	xxvi

CONTENTS

Contacting DSP Publications	xxvi
Notation Conventions	xxvi

INTRODUCTION TO VDK

Motivation	1-1
Rapid Application Development	1-2
Debugged Control Structures	1-2
Code Reuse	1-3
Hardware Abstraction	1-3
Partitioning an Application	1-4
Scheduling	1-5
Priorities	1-5
Preemption	1-7
Protected Regions	1-7
Disabling Scheduling	1-7
Disabling Interrupts	1-8
Thread and Hardware Interaction	1-8
Thread Domain with Software Scheduling	1-9
Interrupt Domain with Hardware Scheduling	1-10
Device Drivers	1-10

CONFIGURATION AND DEBUGGING OF VDK PROJECTS

Configuring VDK Projects	2-1
Linker Description File	2-2

Thread Safe Libraries	2-2
Header Files for the VDK API	2-2
Debugging VDK Projects	2-3
Instrumented Build Information	2-3
VDK State History Window	2-3
Target Load Graph Window	2-4
VDK Status Window	2-4
General Tips	2-5

USING VDK

Threads	3-1
Thread Types	3-2
Thread Parameters	3-2
Stack Size	3-2
Priority	3-3
Required Thread Functionality	3-3
Run Function	3-3
Error Function	3-3
Create Function	3-4
Init Function/Constructor	3-5
Destructor	3-5
Writing Threads in Different Languages	3-6
C++ Threads	3-6
C and Assembly Threads	3-7
Global Variables	3-7

CONTENTS

Error Handling Facilities	3-8
Scheduling	3-8
Ready Queue	3-9
Scheduling Methodologies	3-10
Cooperative Scheduling	3-10
Round-robin Scheduling	3-10
Preemptive Scheduling	3-11
Disabling Scheduling	3-11
Entering the Scheduler From API Calls	3-13
Entering the Scheduler From Interrupts	3-13
Idle Thread	3-13
Signals	3-15
Semaphores	3-15
Behavior of Semaphores	3-16
Thread's Interaction With Semaphores	3-16
Pending on a Semaphore	3-17
Posting a Semaphore	3-18
Periodic Semaphores	3-20
Messages	3-21
Behavior of Messages	3-21
Thread's Interaction With Messages	3-22
Pending on a Message	3-23
Posting a Message	3-24
Events and Event Bits	3-25

Behavior of Events	3-26
Global State of Event Bits	3-26
Event Calculation	3-27
Effect of Unscheduled Regions on Event Calculation	3-28
Thread's Interaction With Events	3-29
Pending on an Event	3-29
Setting or Clearing of Event Bits	3-30
Loading New Event Data into an Event	3-31
Device Flags	3-32
Behavior of Device Flags	3-32
Thread's Interaction With Device Flags	3-32
Interrupt Service Routines	3-33
Enabling and Disabling Interrupts	3-33
Interrupt Architecture	3-34
Vector Table	3-34
Global Data	3-35
Communication with the Thread Domain	3-36
Timer ISR	3-37
Reschedule ISR	3-37
I/O Interface	3-38
I/O Templates	3-38
Device Drivers	3-38
Execution	3-39
Parallel Scheduling Domains	3-40

CONTENTS

Using Device Drivers	3-42
Dispatch Function	3-43
Device Flags	3-51
Pending on a Device Flag	3-51
Posting a Device Flag	3-52
General Notes	3-53
Variables	3-54
Notes on Critical/Unscheduled Regions	3-54
Memory Pools	3-55
Memory Pool Functionality	3-55

VDK DATA TYPES

Data Type Summary	4-1
Bitfield	4-3
DeviceDescriptor	4-4
DeviceFlagID	4-5
DispatchID	4-6
DispatchUnion	4-7
DSP_Family	4-9
DSP_Product	4-10
EventBitID	4-12
EventID	4-13
EventData	4-14
HistoryEnum	4-15
IMASKStruct	4-17

IOID	4-18
IOTemplateID	4-19
MessageID	4-20
MsgChannel	4-21
PoolID	4-23
Priority	4-24
SemaphoreID	4-25
SystemError	4-26
ThreadCreationBlock	4-29
ThreadID	4-31
ThreadStatus	4-32
ThreadType	4-33
Ticks	4-34
VersionStruct	4-35

VDK API REFERENCE

Calling Library Functions	5-2
Linking Library Functions	5-2
Working With VDK Library Header	5-3
Passing Function Parameters	5-3
Library Naming Conventions	5-3
API Summary	5-4
API Functions	5-10
AllocateThreadSlot()	5-11
AllocateThreadSlotEx()	5-13

CONTENTS

ClearEventBit()	5-15
ClearInterruptMaskBits()	5-17
ClearThreadError()	5-18
CloseDevice()	5-19
CreateDeviceFlag()	5-21
CreateMessage()	5-22
CreatePool()	5-24
CreatePoolEx()	5-26
CreateSemaphore()	5-28
CreateThread()	5-30
CreateThreadEx()	5-32
DestroyDeviceFlag()	5-34
DestroyMessage()	5-35
DestroyPool()	5-37
DestroySemaphore()	5-39
DestroyThread()	5-41
DeviceIOCtl()	5-43
DispatchThreadError()	5-45
FreeBlock()	5-46
FreeDestroyedThreads()	5-48
FreeThreadSlot()	5-49
GetEventBitValue()	5-51
GetEventData()	5-52
GetEventValue()	5-53

GetInterruptMask()	5-54
GetLastThreadError()	5-55
GetLastThreadErrorValue()	5-56
GetMessagePayload()	5-57
GetMessageReceiveInfo()	5-59
GetNumAllocatedBlocks()	5-61
GetNumFreeBlocks()	5-62
GetPriority()	5-63
GetSemaphoreValue()	5-64
GetThreadHandle()	5-65
GetThreadID()	5-66
GetThreadSlotValue()	5-67
GetThreadStackUsage()	5-68
GetThreadStatus()	5-69
GetThreadType()	5-70
GetUptime()	5-71
GetVersion()	5-72
LoadEvent()	5-73
LocateAndFreeBlock()	5-75
LogHistoryEvent()	5-76
MakePeriodic()	5-77
MallocBlock()	5-79
MessageAvailable()	5-81
OpenDevice()	5-83

CONTENTS

PendDeviceFlag()	5-85
PendEvent()	5-87
PendMessage()	5-89
PendSemaphore()	5-92
PopCriticalRegion()	5-94
PopNestedCriticalRegions()	5-96
PopNestedUnscheduledRegions()	5-98
PopUnscheduledRegion()	5-99
PostDeviceFlag()	5-101
PostMessage()	5-102
PostSemaphore()	5-104
PushCriticalRegion()	5-106
PushUnscheduledRegion()	5-107
RemovePeriodic()	5-108
ResetPriority()	5-110
SetEventBit()	5-111
SetInterruptMaskBits()	5-113
SetMessagePayload()	5-114
SetPriority()	5-116
SetThreadError()	5-118
SetThreadSlotValue()	5-119
Sleep()	5-120
SyncRead()	5-122
SyncWrite()	5-124

Yield()	5-126
Assembly Macros	5-127
VDK_ISR_ACTIVATE_DEVICE_()	5-129
VDK_ISR_CLEAR_EVENTBIT_()	5-130
VDK_ISR_LOG_HISTORY_EVENT_()	5-131
VDK_ISR_POST_SEMAPHORE_()	5-132
VDK_ISR_SET_EVENTBIT_()	5-133

PROCESSOR-SPECIFIC NOTES

VDK for Blackfin Processors (AD6532, ADSP-BF531, ADSP-BF532, ADSP-BF533, ADSP-BF535, and DM102)	A-1
User and Supervisor Modes	A-1
Thread, Kernel, and Interrupt Execution Levels	A-2
Critical and Unscheduled Regions	A-3
Exceptions	A-3
ISR APIs	A-3
Interrupts	A-4
Timer	A-5
ADSP-BF531, ADSP-BF532, ADSP-BF533, DM102 Processor Memory	A-6
ADSP-BF535 and AD6532 Processor Memory	A-6
Interrupt Nesting	A-7
Thread Stack Usage by Interrupts	A-7
Interrupt Latency	A-8

CONTENTS

VDK for ADSP-219x DSPs	
(ADSP-2191, ADSP-2192-12, ADSP-2195, and ADSP-2196)	A-8
Thread, Kernel, and Interrupt Execution Levels	A-8
Critical and Unscheduled Regions	A-9
Interrupts on ADSP-2192 DSPs	A-9
Interrupts on ADSP-2191 DSPs	A-10
Timer	A-10
Memory	A-11
Interrupt Nesting	A-11
Interrupt Latency	A-12
VDK for ADSP-TS101 TigerSHARC Processors	A-13
Thread, Kernel, and Interrupt Execution Levels	A-13
Critical and Unscheduled Regions	A-14
Interrupts on ADSP-TS101 Processors	A-14
Timer	A-14
ADSP-TS101 Processor Memory	A-15
Interrupt Nesting	A-15
Interrupt Latency	A-16
VDK for SHARC DSPs (ADSP-21060, ADSP-21061, ADSP-21062, ADSP-21065L, ADSP-21160, and ADSP-21161)	A-16
Thread, Kernel and Interrupt Execution Levels	A-16
Critical and Unscheduled Regions	A-17
Interrupts on ADSP-2106x DSPs	A-17
Interrupts on ADSP-2116x DSPs	A-18
Timer	A-19

Memory	A-19
Register Usage	A-20
40-bit Register Usage	A-21
Interrupt Nesting	A-22
Interrupt Latency	A-23

MIGRATING DEVICE DRIVERS

Step 1: Saving Existing Sources	B-1
Step 2: Revising Properties	B-2
Step 3: Revising Sources	B-3
Step 4: Creating Boot Objects	B-4

INDEX

CONTENTS

PREFACE

Thank you for purchasing Analog Devices (ADI) development software for digital signal processor (DSP) applications.

Purpose of this Manual

The *VisualDSP++ Kernel (VDK) User's Guide* contains information about VisualDSP++™ Kernel, a Real Time Operating System kernel integrated with the rest of the VisualDSP++ 3.x development tools. The VDK incorporates state-of-art scheduling and resource allocation techniques tailored specially for the memory and timing constraints of DSP programming. The kernel facilitates development of fast performance structured applications using frameworks of template files.

The kernel is specially designed for effective operations on Analog Devices DSP architectures: ADSP-219x, ADSP-BF53x Blackfin®, ADSP-21xxx SHARC®, and ADSP-TSxxx TigerSHARC® processors.

The majority of the information in this manual is generic. Information applicable to only a particular target processor, or to a particular processor family, is provided in Appendix A, “[Processor-Specific Notes](#)”.

This manual is designed so that you can quickly learn about the kernel internal structure and operation.

Intended Audience

The primary audience for this manual is programmers who are familiar with Analog Devices DSPs. This manual assumes the audience has a working knowledge of the appropriate processor architecture and instruction set. Programmers who are unfamiliar with Analog Devices DSPs can use this manual but should supplement it with other texts, such as *Hardware Reference* and *Instruction Set Reference* manuals, that describe your target architecture.

Manual Contents

The manual consists of:

- Chapter 1, “[Introduction to VDK](#)”

Concentrates on concepts, motivation, and general architectural principles of the VDK software.

- Chapter 2, “[Configuration and Debugging of VDK Projects](#)”

Describes the Integrated Development and Debugging Environment (IDDE) support for configuring and debugging a VDK enabled project. For specific procedures on how to create, modify, and manage the kernel’s components, refer to the VisualDSP++ online Help.

- Chapter 3, “[Using VDK](#)”

Describes the kernel’s internal structure and components.

- Chapter 4, “[VDK Data Types](#)”

Describes built-in data types supported in the current release of the VDK.

- Chapter 5, “[VDK API Reference](#)”
Describes library functions and macros included in the current release of the VDK.
- Appendix A, “[Processor-Specific Notes](#)”
Provides processor-specific information for ADSP-BF53x Blackfin, ADSP-219x, ADSP-TSxxx TigerSHARC, and ADSP-21xxx SHARC processor architectures.
- Appendix B, “[Migrating Device Drivers](#)”
Describes how to convert the device driver components created with VisualDSP++ 2.0 for use in projects built with VisualDSP++ 3.x.

What’s New in this Manual

This revision of the *VisualDSP++ 3.x Kernel (VDK) User’s Guide* documents VDK support for the new Blackfin processors (AD6532, ADSP-BF531, ADSP-BF533, and DM102) in addition to the existing processors, ADSP-BF532 and ADSP-BF535. Note that the older part numbers, “ADSP-21532” and “ADSP-21535”, are deprecated and replaced with “ADSP-BF532” and “ADSP-BF535”, respectively.

The Blackfin processors are embedded processors that sport a Media Instruction Set Computing (MISC) architecture. This architecture is the natural merging of RISC, media functions, and digital signal processing characteristics towards delivering signal processing performance in a microprocessor-like environment.

The manual documents VisualDSP++ Kernel version 3.0.00.

Technical or Customer Support

You can reach DSP Tools Support in the following ways.

- Visit the DSP Development Tools website at
www.analog.com/technology/dsp/developmentTools/index.html
- Email questions to
dsptools.support@analog.com
- Phone questions to **1-800-ANALOGD**
- Contact your ADI local sales office or authorized distributor
- Send questions by mail to

Analog Devices, Inc.
DSP Division
One Technology Way
P.O. Box 9106
Norwood, MA 02062-9106
USA

Supported Processors

VisualDSP++ 3.x Kernel currently supports the following Analog Devices DSPs.

- AD6532, ADSP-BF531, ADSP-BF532 (formerly ADSP-21532), ADSP-BF533, ADSP-BF535 (formerly ADSP-21535), and DM102
- ADSP-2191, ADSP-2192-12, ADSP-2195, and ADSP-2196
- ADSP-TS101
- ADSP-21060, ADSP-21061, ADSP-21062, ADSP-21065L, ADSP-21160, and ADSP-21161N

Product Information

You can obtain product information from the Analog Devices website, from the product CD-ROM, or from the printed publications (manuals).

Analog Devices is online at www.analog.com. Our website provides information about a broad range of products—analog integrated circuits, amplifiers, converters, and digital signal processors.

MyAnalog.com

MyAnalog.com is a free feature of the Analog Devices website that allows customization of a webpage to display only the latest information on products you are interested in. You can also choose to receive weekly email notification containing updates to the webpages that meet your interests. MyAnalog.com provides access to books, application notes, data sheets, code examples, and more.

Product Information

Registration:

Visit www.myanalog.com to sign up. Click **Register** to use MyAnalog.com. Registration takes about five minutes and serves as means for you to select the information you want to receive.

If you are already a registered user, just log on. Your user name is your email address.

DSP Product Information

For information on digital signal processors, visit our website at www.analog.com/dsp, which provides access to technical publications, data sheets, application notes, product overviews, and product announcements.

You may also obtain additional information about Analog Devices and its products in any of the following ways.

- Email questions or requests for information to dsp.support@analog.com
- Fax questions or requests for information to **1-781-461-3010** (North America) or **+49 (0) 89 76903-157** (Europe)
- Access the Digital Signal Processing Division's FTP website at [ftp.analog.com](ftp://ftp.analog.com) or **ftp 137.71.23.21** or <ftp://ftp.analog.com>

Related Documents

For information on product related development software, see the following publications for the appropriate processor family.

VisualDSP++ 3.x Getting Started Guide

VisualDSP++ 3.x User's Guide

VisualDSP++ 3.x C/C++ Compiler and Library Manual

VisualDSP++ 3.x Assembler and Preprocessor Manual

VisualDSP++ 3.x Linker and Utilities Manual

VisualDSP++ 3.x Kernel (VDK) User's Guide

Quick Installation Reference Card

For hardware information, refer to your DSP *Hardware Reference*, *Programming Reference*, and data sheet.

All documentation is available online. Most documentation is available in printed form.

Online Documentation

Online documentation comprises Microsoft HTML Help (.CHM), Adobe Portable Documentation Format (.PDF), and HTML (.HTM and .HTML) files. A description of each file type is as follows.

File	Description
.CHM	VisualDSP++ online Help system files and VisualDSP++ manuals are provided in Microsoft HTML Help format. Installing VisualDSP++ automatically copies these files to the VisualDSP\Help folder. Online Help is ideal for searching the entire tools manual set. Invoke Help from the VisualDSP++ Help menu or via the Windows Start button.
.PDF	Manuals and data sheets in Portable Documentation Format are located in the installation CD's Docs folder. Viewing and printing a .PDF file requires a PDF reader, such as Adobe Acrobat Reader (4.0 or higher). Running setup.exe on the installation CD provides easy access to these documents. You can also copy .PDF files from the installation CD onto another disk.
.HTM or .HTML	Dinkum Abridged C++ library and FlexLM network license manager software documentation is located on the installation CD in the Docs\Reference folder. Viewing or printing these files requires a browser, such as Internet Explorer 4.0 (or higher). You can copy these files from the installation CD onto another disk.

Product Information

Access the online documentation from the VisualDSP++ environment, Windows Explorer, or Analog Devices website.

From VisualDSP++

VisualDSP++ provides access to online Help. It does not provide access to .PDF files or the supplemental reference documentation (Dinkum Abridged C++ library and FlexLM network licence). Access Help by:

- Choosing **Contents**, **Search**, or **Index** from the VisualDSP++ **Help** menu
- Invoking context-sensitive Help on a user interface item (toolbar button, menu command, or window)

From Windows

In addition to shortcuts you may construct, Windows provides many ways to open VisualDSP++ online Help or the supplementary documentation.

Help system files (.CHM) are located in the `VisualDSP\Help` folder. Manuals and data sheets in PDF format are located in the `Docs` folder of the installation CD. The installation CD also contains the Dinkum Abridged C++ library and FlexLM network license manager software documentation in the `\Reference` folder.

Using Windows Explorer

- Double-click any file that is part of the VisualDSP++ documentation set.
- Double-click `vdsp-help.chm`, the master Help system, to access all the other .CHM files.

From the Web

To download the tools manuals, point your browser at www.analog.com/technology/dsp/developmentTools/gen_purpose.html.

Select a DSP family and book title. Download archive (.ZIP) files, one for each manual. Use any archive management software, such as WinZip, to decompress downloaded files.

Printed Manuals

For general questions regarding literature ordering, call the Literature Center at 1-800-ANALOGD (1-800-262-5643) and follow the prompts.

VisualDSP++ Documentation Set

Printed copies of VisualDSP++ manuals may be purchased through Analog Devices Customer Service at 1-781-329-4700; ask for a Customer Service representative. The manuals can be purchased only as a kit. For additional information, call 1-603-883-2430.

If you do not have an account with Analog Devices, you will be referred to Analog Devices distributors. To get information on our distributors, log onto www.analog.com/salesdir/continent.asp.

Hardware Manuals

Printed copies of hardware reference and instruction set reference manuals can be ordered through the Literature Center or downloaded from the Analog Devices website. The phone number is 1-800-ANALOGD (1-800-262-5643). The manuals can be ordered by a title or by product number located on the back cover of each manual.

Notation Conventions

Data Sheets

All data sheets can be downloaded from the Analog Devices website. As a general rule, printed copies of data sheets with a letter suffix (L, M, N, S) can be obtained from the Literature Center at **1-800-ANALOGD** (**1-800-262-5643**) or downloaded from the website. Data sheets without the suffix can be downloaded from the website only—no hard copies are available. You can ask for the data sheet by part name or by product number.

If you want to have a data sheet faxed to you, the phone number for that service is **1-800-446-6212**. Follow the prompts and a list of data sheet code numbers will be faxed to you. Call the Literature Center first to find out if requested data sheets are available.

Contacting DSP Publications

Please send your comments and recommendations on how to improve our manuals and online Help. You can contact us by:

- Emailing dsp.techpubs@analog.com
- Filling in and returning the attached Reader's Comments Card found in our manuals

Notation Conventions

The following table identifies and describes text conventions used in this manual.



Additional conventions, which apply only to specific chapters, may appear throughout this document.

Example	Description
Close command (File menu) or OK	Text in bold style indicates the location of an item within the VisualDSP++ environment's menu system and user interface items.
{this that}	Alternative required items in syntax descriptions appear within curly brackets separated by vertical bars; read the example as <i>this</i> or <i>that</i> .
[this that]	Optional items in syntax descriptions appear within brackets and separated by vertical bars; read the example as an optional <i>this</i> or <i>that</i> .
[this,...]	Optional item lists in syntax descriptions appear within brackets delimited by commas and terminated with an ellipsis; read the example as an optional comma-separated list of <i>this</i> .
.SECTION	Commands, directives, keywords, code examples, and feature names are in text with letter gothic font.
<i>filename</i>	Non-keyword placeholders appear in text with italic style format.
	A note providing information of special interest or identifying a related topic. In the online version of this book, the word Note appears instead of this symbol.
	A caution providing information about critical design or programming issues that influence operation of a product. In the online version of this book, the word Caution appears instead of this symbol.

1 INTRODUCTION TO VDK

This chapter concentrates on concepts, motivation, and general architectural principles of the operating system kernel. It also provides information on how to partition a VDK application into independent, reusable functional units that are easy to maintain and debug.

The following sections provide information about the operating system kernel concepts.

- [“Motivation” on page 1-1](#)
- [“Partitioning an Application” on page 1-4](#)
- [“Scheduling” on page 1-5](#)
- [“Protected Regions” on page 1-7](#)
- [“Thread and Hardware Interaction” on page 1-8](#)

Motivation

All applications require control code as support for the algorithms that are often thought of as the “real” program. The algorithms require data to be moved to and/or from peripherals, and many algorithms consist of more than one functional block. For some systems, this control code may be as simple as a “superloop” blindly processing data that arrives at a constant rate. However, as processors become more powerful, considerably more sophisticated control may be needed to realize the processor’s potential, to allow the DSP to absorb the required functionality of previously sup-

Motivation

ported chips, and to allow a single DSP to do the work of many. The following sections provide an overview of some of the benefits of using a kernel on a DSP.

Rapid Application Development

The tight integration between the VisualDSP++ environment and the VDK allows rapid development of applications compared to creating all of the control code required by hand. The use of automatic code generation and file templates, as well as a standard programming interface to device drivers, allows you to concentrate on the algorithms and the desired control flow rather than on the implementation details. VDK supports the use of C, C++, and assembly language. You are encouraged to develop code that is highly readable and maintainable, yet to retain the option of hand optimizing if necessary.

Debugged Control Structures

Debugging a traditional DSP application can be laborious because development tools (compiler, assembler, and linker among others) are not aware of the architecture of the target application and the flow of control that results. Debugging complex applications is much easier when instantaneous snapshots of the system state and statistical run-time data are clearly presented by the tools. To help offset the difficulties in debugging software, VisualDSP++ includes three versions of the VDK libraries containing full instrumentation (including error checking), only error checking, and neither instrumentation nor error checking.

In the instrumented mode, the kernel maintains statistical information and logging of all significant events into a history buffer. When the execution is paused, the debugger can traverse this buffer and present a graphical trace of the program's execution including context switches, pending and posting of signals, changes in a thread's status, and more. Statistics are presented for each thread in a tabular view and show the total

amount of time the thread has executed, the number of times it has been run, the signal it is currently blocked on, and other data. For more information, see [“Debugging VDK Projects” on page 2-3](#) and the online Help.

Code Reuse

Many programmers begin a new project by writing the infrastructure portions that transfers data to, from, and between algorithms. This necessary control logic usually is created from scratch by each design team and infrequently reused on subsequent projects. The VDK provides much of this functionality in a standard, portable and reusable library. Furthermore, the kernel and its tight integration with the VisualDSP++ environment are designed to promote good coding practice and organization by partitioning large applications into maintainable and comprehensible blocks. By isolating the functionality of subsystems, the kernel helps to prevent the morass all too commonly found in systems programming.

The kernel is designed specifically to take advantage of commonality in user applications and to encourage code reuse. Each thread of execution is created from a user defined template, either at boot time or dynamically by another thread. Multiple threads can be created from the same template, but the state associated with each created instance of the thread remains unique. Each thread template represents a complete encapsulation of an algorithm that is unaware of other threads in the system unless it has a direct dependency.

Hardware Abstraction

In addition to a structured model for algorithms, the VDK provides a hardware abstraction layer. Presented programming interfaces allow you to write most of the application in a platform independent, high level language (C or C++). The VDK API is identical for all Analog Devices processors, allowing code to be easily ported to a different DSP core.

Partitioning an Application

When porting an application to a new platform, programmers must address the two areas necessarily specific to a particular processor: interrupt service routines and device drivers. The VDK architecture identifies a crisp boundary around these subsystems and supports the traditionally difficult development with a clear programming framework and code generation. Both interrupts and device drivers are declared with a graphical user interface in the IDDE, which generates well commented code that can be compiled without further effort.

Partitioning an Application

A VDK *thread* is an encapsulation of an algorithm and its associated data. When beginning a new project, use this notion of a thread to leverage the kernel architecture and to reduce the complexity of your system. Since many algorithms may be thought of as being composed of “subalgorithm” building blocks, an application can be partitioned into smaller functional units that can be individually coded and tested. These building blocks then become reusable components in more robust and scalable systems.

You define the behavior of VDK threads by creating *thread types*. Types are templates that define the behavior and data associated with all threads of that type. Like data types in C or C++, thread types are not used directly until an instance of the type is created. Many threads of the same thread type can be created, but for each thread type, only one copy of the code is linked into the executable code. Each thread has its own private set of variables defined for the thread type, its own stack, and its own C run-time context.

When partitioning an application into threads, identify portions of your design in which a similar algorithm is applied to multiple sets of data. These are, in general, good candidates for thread types. When data is present in the system in sequential blocks, only one instance of the thread type is required. If the same operation is performed on separate sets of

data simultaneously, multiple threads of the same type can coexist and be scheduled for prioritized execution (based on when the results are needed).

Scheduling

The VDK is a *preemptive multitasking* kernel. Each thread begins execution at its entry point. Then, it either runs to completion or performs its primary function repeatedly in an infinite loop. It is the role of the scheduler to preempt execution of a thread and to resume its execution when appropriate. Each thread is given a *priority* to assist the scheduler in determining precedence of threads (see [Figure 1-1 on page 1-6](#)).

The scheduler gives processor time to the thread with the highest priority that is in the *ready* state (see [Figure 3-2 on page 3-14](#)). A thread is in the ready state when it is not waiting for any system resources it has requested. A reference to each ready thread is stored in a structure that is internal to the kernel and known as the *ready queue*. For more information, see [“Scheduling” on page 3-8](#).

Priorities

Each thread is assigned a dynamically modifiable priority based on the default for its thread type declared in VisualDSP++ environment’s **Project** window. An application is limited to either fourteen or thirty priority levels, depending on the processor’s architecture. However, the number of threads at each priority is limited, in practice, only by system memory. Priority level one is the highest priority, and priority fourteen (or thirty) is the lowest. The system maintains an idle thread that is set to a priority lower than that of the lowest user thread.

Assigning priorities is one of the most difficult tasks of designing a real time preemptive system. Although there has been research in the area of rigorous algorithms for assigning priorities based on deadlines (e.g., rate

Scheduling




Figure 1-1. VDK State Diagram

monotonic scheduling), most systems are designed by considering the interrupts and signals that are triggering the execution, while balancing the deadlines imposed by the system's input and output streams. For more information, see ["Thread Parameters" on page 3-2](#).

Preemption

A running thread continues execution unless it requests a system resource using a kernel API. When a thread requests a signal (semaphore, event, device flag, or message) and the signal is available, the thread resumes execution. If the signal is not available, the thread is removed from the ready queue—the thread is *blocked* (see [Figure 3-2 on page 3-14](#)). The kernel does not perform a context switch as long as the running thread maintains the highest priority in the ready queue, even if the thread frees a resource and enables other threads to move to the ready queue at the same or lower priority. A thread can also be interrupted. When an interrupt occurs, the kernel yields to the hardware interrupt controller. When the ISR completes, the highest priority thread resumes execution. For more information, see [“Preemptive Scheduling” on page 3-11](#).

Protected Regions

Frequently, system resources must be accessed atomically. The kernel provides two levels of protection for code that needs to execute sequentially—*unscheduled regions* and *critical regions*.

Critical and unscheduled regions can be intertwined. You can enter critical regions from within unscheduled regions, or enter unscheduled regions from within critical regions. For example, if you are in an unscheduled region and call a function that pushes and pops a critical region, the system is still in an unscheduled region when the function returns.

Disabling Scheduling

The VDK scheduler can be disabled by entering an unscheduled region. The ability to disable scheduling is necessary when you need to free multiple system resources without being switched out, or access global variables that are modified by other threads without preventing interrupts from being serviced. While in an unscheduled region, interrupts are still

Thread and Hardware Interaction

enabled and ISRs execute. However, the kernel does not perform a thread context switch even if a higher priority thread becomes ready. Unscheduled regions are implemented using a stack style interface. This enables you to begin and end an unscheduled region within a function without concern for whether or not the calling code is already in an unscheduled region.

Disabling Interrupts

On occasions, disabling the scheduler does not provide enough protection to keep a block of thread code reentrant. A critical region disables both scheduling and interrupts. Critical regions are necessary when a thread is modifying global variables that may also be modified by an Interrupt Service Routine. Similar to unscheduled regions, critical regions are implemented as a stack. Developers can enter and exit critical regions in a function without being concerned about the critical region state of the calling code. Care should be taken to keep critical regions as short as possible as they may increase interrupt latency.

Thread and Hardware Interaction

Threads should have minimal knowledge of hardware; rather, they should use *device drivers* for hardware control. A thread can control and interact with a device in a portable and hardware abstracted manner through a standard set of APIs.

The VDK Interrupt Service Routine framework encourages you to remove specific knowledge of hardware from the algorithms encapsulated in threads (see [Figure 1-2](#)). Interrupts relay information to threads through signals to device drivers or directly to threads. Using signals to connect hardware to the algorithms allows the kernel to schedule threads based on asynchronous events.




Figure 1-2. Device Drivers Entry Points

The VDK run-time environment can be thought of as a bridge between two domains, the *thread domain* and the *interrupt domain*. The interrupt domain services the hardware with minimal knowledge of the algorithms, and the thread domain is abstracted from the details of the hardware. Device drivers and signals bridge the two domains. For more information, see [“Threads” on page 3-1](#).

Thread Domain with Software Scheduling

The thread domain runs under a C/C++ run-time model. The prioritized execution is maintained by a software scheduler with full context switching. Threads should have little or no direct knowledge of the hardware; rather, threads should request resources and then wait for them to become available. Threads are granted processor time based on their priority and

Thread and Hardware Interaction

requested resources. Threads should minimize time spent in critical and unscheduled regions to avoid short-circuiting the scheduler and interrupt controller.

Interrupt Domain with Hardware Scheduling

The interrupt domain runs outside the C/C++ run-time model. The prioritized execution is maintained by the hardware interrupt controller. ISRs should be as small as possible. They should only do as much work as is necessary to acknowledge asynchronous external events and to allow peripherals to continue operations in parallel with the processor. ISRs should only signal that more processing can occur and leave the processing to threads. For more information, see [“Interrupt Service Routines” on page 3-33](#).

Device Drivers

Interrupt Service Routines can communicate with threads directly using signals. Alternatively, an interVisualDSP++ 3.x Kernel (VDK) User's Guidecomplex device-specific functionality that is abstracted from the algorithm. A device driver is a single function with multiple entry conditions and domains of execution. For more information, see [“Device Drivers” on page 3-38](#).

2 CONFIGURATION AND DEBUGGING OF VDK PROJECTS

This chapter contains information about the VisualDSP++ Integrated Development and Debugging Environment (IDDE) support for VDK enabled projects. You can access VDK components and services through the set of menus, commands, and windows in the IDDE.

If you are new to VisualDSP++ application development software, we recommend that you start with the *VisualDSP++ 3.x Getting Started Guide* for your target processor family.

The IDDE support for the VDK can be broken into two areas:

- [“Configuring VDK Projects” on page 2-1](#)
- [“Debugging VDK Projects” on page 2-3](#)

Configuring VDK Projects

VisualDSP++ is extended to manage all of the VDK components. You start developing a VDK based application by creating a set of source files. The IDDE automatically generates a source code framework for each user requested kernel object. Use the interface to supply the required information for these objects.

For specific procedures on how to set up VDK system parameters or how to create, modify, or delete a VDK component, refer to the VisualDSP++ online Help. Following the online procedures ensures your VDK projects

Configuring VDK Projects

build consistently and accurately with minimal project management. The process reduces development time and allows you to concentrate on algorithm development.

Linker Description File

When a new project makes use of the kernel, a reference to a VDK-specific default Linker Description File (.LDF) is added to the project. This file will be copied to your project directory to allow you to modify it to suit your individual hardware configurations.

Thread Safe Libraries

Just as user threads must be reentrant, special “thread safe” versions of the standard C and C++ libraries are included for use with the VDK. The default .LDF included in VDK projects links with these libraries. If you modify your Linker Description File, ensure that the file links with the thread safe libraries. Your project’s LDF resides in the **Linker Files** folder and is accessible via the **Project** tab of the **Project** window in VisualDSP++.

Header Files for the VDK API

When a VDK project is created in the development environment, one of the automatically generated files in the project directory is `vdk.h`. This header file contains enumerations for every user defined object in the development environment and all VDK API declarations. Your source files must include `vdk.h` to access any kernel services.

Debugging VDK Projects

Debugging embedded software is a difficult task. To help offset the initial difficulties present in debugging VDK enabled projects, the kernel offers special instrumented builds.

Instrumented Build Information

When building a VDK project, you have an option to include instrumentation in your executable by choosing **Full Instrumentation** as the instrumentation level in the **Kernel** tab of the **Project** window. An instrumented build differs from a release or non-instrumented build because the build includes extra code for thread statistic logging. In addition, an instrumented build creates a circular buffer of important system events. The extra logging introduces slight overhead in thread switches and certain API calls but helps you to trace system activities.

VDK State History Window

The VDK logs user defined events and certain system state changes in a circular buffer. An event is logged in the history buffer with a call to [LogHistoryEvent\(\)](#). The call to [LogHistoryEvent\(\)](#) logs four data values: the `ThreadID` of the calling thread, the tick when the call happened, the enumeration, and a value that is specific to the enumeration. Enumerations less than zero are reserved for use by the VDK. For more information about the history enumeration type, see [HistoryEnum](#) on page 4-15.

Using the history log, the IDDE displays a graph of running threads and system state changes in the **State History** window. Note that the data displayed in this window is only updated at halt. The **State History** window, the **Thread Status** and **Thread Event** legends are described in detail in the online Help.

Target Load Graph Window

Instrumented VDK builds allow you to analyze the average load of the processor over a period of time. The calculated load is displayed in the **Target Load** graph window. Although the calculation is not exact, the graph helps you to estimate the utilization level of the processor. Note that the information is updated at halt.

The **Target Load** graph shows the percent of time the target spent in the idle thread. A load of 0% means the VDK spent all of its time in the idle thread. A load of 100% means the target did not spend any time in the idle thread. Load data is processed using a moving window average. The load percentage is calculated for every clock tick, and all the ticks are averaged. The following formula is used to calculate the percentage of utilization for every clock tick.

$$\text{Load} = 1 - (\# \text{ of times idle thread ran this tick}) / (\# \text{ of threads run this tick})$$

For more information about the **Target Load** graph, refer to the online Help.

VDK Status Window

Besides history and processor load information, an instrumented build collects statistics for relevant VDK components, such as when a thread was created, last run, the number of times run, etc. This data is displayed in the **Status** window and is updated at halt.

For more information about the VDK Status window, refer to the online Help.

General Tips

Even with the data collection features built into the VDK, debugging thread code is a difficult task. Due to the fact that multiple threads in a system are interacting asynchronously with device drivers, interrupts, and the idle thread, it can become difficult to track down the source of an error.

Unfortunately, one of the oldest and easiest debugging methods—inserting breakpoints—can have uncommon side effects in VDK projects. Since multiple threads (either multiple instantiations of the same thread type or different threads of different thread types) can execute the same function with completely different contexts, the utilization of non-thread-aware breakpoints is diminished. One possible workaround is to insert some ‘thread-specific’ breakpoints:

```
if (VDK_GetThreadID() == <thread_with_bug>)
{
    <some statement>;    /* Insert breakpoint */
}
```


3 USING VDK

This chapter describes how the VDK implements the general concepts described in Chapter 2, [“Introduction to VDK”](#). For information about the kernel library, see Chapter 5, [“VDK API Reference”](#).

The following sections provide information about the operating system kernel components and operations.

- [“Threads” on page 3-1](#)
- [“Scheduling” on page 3-8](#)
- [“Signals” on page 3-15](#)
- [“Interrupt Service Routines” on page 3-33](#)
- [“I/O Interface” on page 3-38](#)
- [“Memory Pools” on page 3-55](#)

Threads

When designing an application, you partition it into threads, where each thread is responsible for a piece of the work. Each thread operates independently of the others. A thread performs its duty as if it has its own processor but can communicate with other threads.

Thread Types

You do not directly define threads; instead, you define thread types. A thread is an instance of a thread type and is similar to any other user defined type.

You can create multiple instantiations of the same thread type. Each instantiation of the thread type has its own stack, state, priority, and other local variables. Each thread is individually identified by its **ThreadID**, a handle that can be used to reference that thread in kernel API calls. A thread can gain access to its **ThreadID** by calling **GetThreadID()**. A **ThreadID** is valid for the life of the thread—once a thread is destroyed, the **ThreadID** becomes invalid.

Old **ThreadID**s are eventually reused, but there is significant time between a thread's destruction and the **ThreadID** reuse: other threads have to recognize that the original thread is destroyed.

Thread Parameters

When a thread is created, the system allocates space in the heap to store a data structure that holds the thread-specific parameters. The data structure contains internal information required by the kernel and the thread type specifications provided by the user.

Stack Size

Each thread has its own stack. The full C/C++ run-time model, as specified in the appropriate *VisualDSP++ 3.x C/C++ Compiler and Library Manual*, is maintained on a per thread basis. It is your responsibility to assure that each thread has enough room on its stack for all function calls' return addresses and passed parameters appropriate to the particular run-time model, user code structure, use of libraries, etc. Stack overflows do not generate an exception, so an undersized stack has the potential to cause difficulties when reproducing bugs in your system.

Priority

Each thread type specifies a default priority. Threads may change their own (or another thread's) priority dynamically using the [SetPriority\(\)](#) or [ResetPriority\(\)](#) functions. Priorities are predefined by the kernel as an enumeration of type [Priority](#) with a value of `kPriority1` being the highest priority (or the first to be scheduled) in the system. The priority enumeration is set up such that `kPriority1 > kPriority2 > ....` The number of priorities is limited to the processor's word size minus two.

Required Thread Functionality

Each thread type requires five particular functions to be declared and implemented. Default null implementations of all five functions are provided in the templates generated by the VisualDSP++ development environment. The thread's run function is the entry point for the thread. For many thread types, the thread's run and error functions are the only ones in the template you need to modify. The other functions allocate and free up system resources at appropriate times during the creation and destruction of a thread.

Run Function

The run function—called `Run()` in C++ and `RunFunction()` in C/assembly implemented threads—is the entry point for a fully constructed thread; `Run()` is roughly equivalent to `main()` in a C program. When a thread's run function returns, the thread is moved to the queue of threads waiting to free their resources. If the run function never returns, the thread remains running until destroyed.

Error Function

The thread's error function is called by the kernel when an error occurs in an API call made by the thread. The error function passes a description of the error in the form of an enumeration (see [SystemError](#) on page 4-26 for more details). It also can pass an additional piece of information whose

exact definition depends on the error enumeration. A thread's default error-handling behavior destroys the thread. See “[Error Handling Facilities](#)” on page 3-8 for more information about error handling in the VDK.

Create Function

The create function is similar to the C++ constructor. The function provides an abstraction used by the kernel APIs [CreateThread\(\)](#) and [CreateThreadEx\(\)](#) to enable dynamic thread creation. The create function is the first function called in the process of constructing a thread; it is also responsible for calling the thread's init function/constructor. Similar to the constructor, the create function executes in the context of the thread that is spawning a new thread by calling [CreateThread\(\)](#) or [CreateThreadEx\(\)](#). The thread being constructed does not have a run-time context fully established until after these functions complete.

A create function calls the constructor for the thread and ensures that all of the allocations that the thread type required have taken place correctly. If any of the allocations failed, the create function deletes the partially created thread instantiation and returns a null pointer. If the thread has been constructed successfully, the create function returns the pointer to the thread. A create function should not call [DispatchThreadError\(\)](#) because [CreateThread\(\)](#) and [CreateThreadEx\(\)](#) handle error reporting to the calling thread when the create function returns a null pointer.

The create function is exposed completely in C++ source templates. For C or assembly threads, the create function appears only in the thread's header file. If the thread allocates data in `InitFunction()`, you need to modify the create function in the thread's header to verify that the allocations are successful and delete the thread if not.

A thread of a certain thread type can be created at boot time by specifying a boot thread of the given thread type in the development environment. Additionally, if the number of threads in the system is known at build time, all the threads can be boot threads.

Init Function/Constructor

The `InitFunction()` (in C/assembly) and the constructor (in C++) provide a place for a thread to allocate system resources during the dynamic thread creation. A thread uses `malloc` (or `new`) when allocating the thread's local variables. A thread's init function/constructor cannot call any VDK APIs since the function is called from within a different thread's context.

Destructor

The destructor is called by the system when the thread is destroyed. A thread can do this explicitly with a call to `DestroyThread()`. The thread can also be destroyed as a result of an error condition from which it can not recover, or the thread can simply run to completion by reaching the end of its run function and falling out of scope. In all cases, you are responsible for freeing the memory and other system resources that the thread has claimed. Any memory allocated with `malloc` or `new` in the constructor should be released with a corresponding call to `free` or `delete` in the destructor.

A thread is not necessarily destroyed immediately when `DestroyThread()` is called. `DestroyThread()` takes a parameter that provides a choice of priority as to when the thread's destructor is called. If the second parameter, `inDestroyNow`, is `FALSE`, the thread is placed in a queue of threads to be cleaned up by the idle thread, and the destructor is called at a priority lower than that of any user threads. While this scheme has many advantages, it works as, in essence, the background garbage collection. This is not deterministic and presents no guarantees of when the freed resources are available to other threads.

If the `inDestroyNow` argument is passed to `DestroyThread()` with a value of `TRUE`, the destructor is called immediately. This assures the resources are freed when the function returns, but the destructor is effectively called at the priority of the currently running thread even if a lower priority thread is being destroyed. It should be noted that `DestroyThread()` runs on the

Threads

stack of the calling thread. Therefore, if a thread calls [DestroyThread\(\)](#) in order to destroy itself, even if `inDestroyNow` is set to `TRUE`, the freeing of the thread's resources needs to be performed by the [Idle Thread](#).

Writing Threads in Different Languages

The code to implement different thread types may be written in C, C++, or assembly. The choice of language is transparent to the kernel. The development environment generates well commented skeleton code for all three choices.

One of the key properties of threads is that they are separate instances of the thread type templates—each with a unique local state. The mechanism for allocating, managing, and freeing thread local variables varies from language to language.

C++ Threads

C++ threads have the simplest template code of the three supported languages. User threads are derived classes of the abstract base class `VDK::Thread`. C++ threads have slightly different function names and include a `Create()` function as well a constructor.

Since user thread types are derived classes of the abstract base class `VDK::Thread`, member variables may be added to user thread classes in the header as with any other C++ class. The normal C++ rules for object scope apply so that threads may make use of `public`, `private`, and `static` members. All member variables are thread-specific (or instantiation-specific).

Additionally, calls to VDK APIs in C++ are different from C and assembly calls. All VDK APIs are in the `VDK` namespace. For example, a call to [CreateThread\(\)](#) in C++ is `VDK::CreateThread()`. We do not recommend exposing the entire VDK namespace in your C++ threads with the `using` keyword.

C and Assembly Threads

Threads written in C rely on a C++ wrapper in their generated header file but are otherwise ordinary C functions. C thread function implementations are compiled without the C++ compiler extensions.

In C and assembly programming, the state local to the thread is accessed through a handle (a pointer to a pointer) that is passed as an argument to each of the four user thread functions. When more than a single word of state is needed, a block of memory is allocated with `malloc()` in the thread type's `InitFunction()`, and the handle is set to point to the new structure.

Each instance of the thread type allocates a unique block of memory, and when a thread of that type is executing, the handle references the correct memory reference. Note that, in addition to being available as an argument to all functions of the thread type, the handle can be obtained at any time for the currently running thread using the API [GetThreadHandle\(\)](#). The `InitFunction()` and `DestroyFunction()` implementations for a thread should not call [GetThreadHandle\(\)](#) but should instead use the parameter passed to these functions, as they do not execute in the context of the thread being initialized or destroyed.

Global Variables

VDK applications can use global variables as normal variables. In C or C++, a variable defined in exactly one source file is declared as `extern` in other files in which that variable is used. In assembly, the `.GLOBAL` declaration exposes a variable outside a source file, and the `.EXTERN` declaration resolves a reference to a symbol at link time.

You need to plan carefully how global variables are to be used in a multi-threaded system. Limit access to a single thread (a single instantiation of a thread type) whenever possible to avoid reentrancy problems. Critical and/or unscheduled regions should be used to protect operations on global entities that can potentially leave the system in an undefined state if not completed atomically.

Error Handling Facilities

The VDK includes an error-handling mechanism that allows you to define behavior independently for each thread type. Each function call in Chapter 5, “[VDK API Reference](#)”, lists the error codes that may result. For information on the specific error code, refer to [SystemError on page 4-26](#).

The assumption underlying the error-handling mechanism in VDK is that all function calls normally succeed and, therefore, do not require an explicit error code to be returned and verified by the calling code. The VDK’s method differs from common C programming convention in which the return value of every function call must be checked to assure that the call has succeeded without an error. While that model is widely used in conventional systems programming, real time embedded system function calls rarely, if ever, fail. When an error does occur, the system calls the user implemented `ErrorFunction()`. You can call [GetLastThreadError\(\)](#) to obtain an enumeration that describes the error condition. You can also call [GetLastThreadErrorValue\(\)](#) to obtain an additional descriptive value whose definition depends on the enumeration. The thread’s `ErrorFunction()` should check if the value returned by [GetLastThreadError\(\)](#) is one that can be handled intelligently and can perform the appropriate operations. Any enumerated errors that the thread cannot handle must be passed to the default thread error function. For instructions on how to pass an error to the error function, see comments included in the generated thread code.

Scheduling

The scheduler’s role is to ensure that the highest priority ready thread is allowed to run at the earliest possible time. The scheduler is never invoked directly by a thread but is executed whenever a kernel API—called from either a thread or an ISR—changes the highest priority thread. The scheduler is not invoked during critical or unscheduled regions, but can be invoked immediately at the close of either type of protected region.

Ready Queue

The scheduler relies on an internal data structure known as the *ready queue*. The queue holds references to all threads that are not blocked or sleeping. All threads in the ready queue have all resources needed to run; they are only waiting for processor time. The exception is the currently running thread, which remains in the ready queue during execution.

The ready queue is called a queue because it is arranged as a prioritized FIFO buffer. That is, when a thread is moved to the ready queue, it is added as the last entry at its priority. For example, there are four threads in the ready queue at the priorities `kPriority3`, `kPriority5`, and `kPriority7`, and an additional thread is made ready with a priority of `kPriority5` (see [Figure 3-1](#)).




Figure 3-1. Ready Queue

Scheduling

The additional thread is inserted after the old thread with the priority of `kPriority5` but before the thread with the priority of `kPriority7`.

Threads are added to and removed from the ready queue in a fixed number of cycles regardless of the size of the queue.

Scheduling Methodologies

The VDK always operates as a preemptive kernel. However, you can take advantage of a number of modes to expand the options for simpler or more complex scheduling in your applications.

Cooperative Scheduling

Multiple threads may be created at the same priority level. In the simplest scheduling scheme, all threads in the system are given the same priority, and each thread has access to the processor until it manually yields control. This arrangement is called *cooperative multithreading*. When a thread is ready to defer to the next thread at the same priority level, the thread can do so by calling the `Yield()` function, placing the currently running thread at the end of the list. In addition, any system call that causes the currently running thread to block would have a similar result. For example, if a thread pends on a signal that is not currently available, the next thread in the queue at that priority starts running.

Round-robin Scheduling

Round-robin scheduling, also called time slicing, allows multiple threads with the same priority to be given processor time automatically in fixed duration allotments. In the VDK, priority levels may be designated as round-robin mode at build time and their period specified in system ticks. Threads at that priority are run for that duration, as measured by the number of VDK `Ticks`. If the thread is preempted by a higher priority thread for a significant amount of time, the time is not subtracted from the time slice. When a thread's round-robin period completes, it is moved

to the end of the list of threads at its priority in the ready queue. Note that the round-robin period is subject to jitter when threads at that priority are preempted.

Preemptive Scheduling

Full *preemptive scheduling*, in which a thread gets processor time as soon as it is placed in the ready queue if it has a higher priority than the running thread, provides more power and flexibility than pure cooperative or round-robin scheduling.

The VDK allows the use of all three paradigms without any modal configuration. For example, multiple non-time-critical threads can be set to a low priority in the round-robin mode, ensuring that each thread gets processor time without interfering with time critical threads. Furthermore, a thread can yield the processor at any time, allowing another thread to run. A thread does not need to wait for a timer event to swap the thread out when it has completed the assigned task.

Disabling Scheduling

Sometimes it is necessary to disable the scheduler when manipulating global entities. For example, when a thread tries to change the state of more than one signal at a time, the thread can enter an unscheduled region to ensure that all updates occur atomically. Unscheduled regions are sections of code that execute without being preempted by a higher priority thread. Note that interrupts are serviced in an unscheduled region, but the same thread runs on return to the thread domain. Unscheduled regions are entered through a call to [PushUnscheduledRegion\(\)](#). To exit an unscheduled region, a thread calls [PopUnscheduledRegion\(\)](#).

Scheduling

Unscheduled regions (in the same way as critical regions, covered in [“Enabling and Disabling Interrupts” on page 3-33](#)) are implemented with a stack. Using nested critical and unscheduled regions allows you to write code that activates a region without being concerned about the region context when a function is called. For example:

```
void My_UnscheduledFunction()
{
    VDK_PushUnscheduledRegion();
    /* In at least one unscheduled region, but
       this function can be used from any number
       of unscheduled or critical regions */
    /* ... */
    VDK_PopUnscheduledRegion();
}
void MyOtherFunction()
{
    VDK_PushUnscheduledRegion();
    /* ... */
    /* This call adds and removes one unscheduled region */
    My_UnscheduledFunction();
    /* The unscheduled regions are restored here */
    /* ... */
    VDK_PopUnscheduledRegion();
}
```

An additional function for controlling unscheduled regions is [PopNestedUnscheduledRegions\(\)](#). This function completely pops the stack of all unscheduled regions. Although the VDK includes [PopNestedUnscheduledRegions\(\)](#), applications should use the function infrequently and balance regions correctly.

Entering the Scheduler From API Calls

Since the highest priority ready thread is the running thread, the scheduler needs to be called only when a higher priority thread becomes ready. Because a thread interacts with the system through a series of API calls, the points at which the highest priority ready thread may change are well defined. Therefore, a thread invokes the scheduler only at these times, or whenever it leaves an unscheduled region.

Entering the Scheduler From Interrupts

ISRs communicate with the thread domain through a set of APIs that do not assume any context. Depending on the system state, an ISR API call may require the scheduler to be executed. The VDK reserves the lowest priority software interrupt to handle the reschedule process.

If an ISR API call affects the system state, the API raises the lowest priority software interrupt. When the lowest priority software interrupt is scheduled to run by the hardware interrupt dispatcher, the interrupt reduces to subroutine and enters the scheduler. If the interrupted thread is not in an unscheduled region and a higher priority thread has become ready, the scheduler swaps out the interrupted thread and swaps in the new highest priority ready thread. The lowest priority software interrupt respects any unscheduled regions the running thread is in. However, interrupts can still service device drivers, post semaphores, etc. On leaving the unscheduled region, the scheduler is run again, and the highest priority ready thread will become the running thread (see [Figure 3-2 on page 3-14](#)).

Idle Thread

The idle thread is a predefined, automatically created thread that has a priority lower than that of any user threads. Thus, when there are no user threads in the ready queue, the idle thread runs. The only substantial work performed by the idle thread is the freeing of resources of threads that

Scheduling




Figure 3-2. Thread State Diagram

have been destroyed. In other words, the idle thread handles destruction of threads that were passed to `DestroyThread()` with a value of `FALSE` for `inDestroyNow`.

The time spent in threads other than the idle thread is shown plotted as a percentage over time on the **Target Load** tab of the **State History** window in VisualDSP++. See “[VDK State History Window](#)” on page 2-3 and online Help for more information about the **State History** window.

Signals

Threads have four different methods for communication and synchronization:

- “Semaphores” on page 3-15
- “Messages” on page 3-21
- “Events and Event Bits” on page 3-25
- “Device Flags” on page 3-32

Each communication method has a different behavior and use. A thread pends on any of the four types of signals, and if a signal is unavailable, the thread blocks until the signal becomes available or (optionally) a timeout is reached.

Semaphores

Semaphores are protocol mechanisms offered by most operating systems. Semaphores are used to:

- Control access to a shared resource
- Signal a certain system occurrence
- Allow threads to synchronize
- Schedule periodic execution of threads

The maximum number of active semaphores and initial state of the semaphores enabled at boot time are set up when your project is built.

Behavior of Semaphores

A semaphore is a token that a thread acquires so that the thread can continue execution. If the thread pends on the semaphore and it is available (the count value associated with the semaphore is greater than zero), the semaphore is acquired, its count value is decremented by one and the thread continues normal execution. If the semaphore is not available (its count is zero), the thread trying to acquire (pend on) the semaphore blocks until the semaphore is available, or the specified timeout occurs. If the semaphore does not become available in the time specified, the thread continues execution in its error function.

Semaphores are global resources accessible to all threads in the system. Threads of different types and priorities can pend on a semaphore. When the semaphore is posted, the thread with the highest priority that has been waiting the longest is moved to the ready queue. If there are no threads pending on the semaphore, its count value is incremented by one. The count value is limited by the maximum value specified at the time of the semaphore creation. Additionally, unlike many operating systems, VDK semaphores are not owned. In other words, any thread is allowed to post a semaphore (make it available). If a thread has requested (pended on) and acquired a semaphore, and the thread is subsequently destroyed, the semaphore is not automatically posted by the kernel.

Besides operating as a flag between threads, a semaphore can be set up to be periodic. A periodic semaphore is posted by the kernel every n ticks, where n is the period of the semaphore. Periodic semaphores can be used to ensure that a thread is run at regular intervals.

Thread's Interaction With Semaphores

Threads interact with semaphores through the set of semaphore APIs. These functions allow a thread to create a semaphore, destroy a semaphore, pend on a semaphore, post a semaphore, get a semaphore's value, and add or remove a semaphore from the periodic queue.

Pending on a Semaphore

Figure 3-3 illustrates the process of pending on a semaphore.




Figure 3-3. Pending on a Semaphore

Threads can pend on a semaphore with a call to `PendSemaphore()`. When a thread calls `PendSemaphore()`, it performs one of the following:

- Acquires the semaphore, decrements its count by one, and continues execution.
- Blocks until the semaphore is available or the specified timeout occurs.

If the semaphore becomes available before the timeout occurs, the thread continues execution; otherwise, the thread's error function is called and the thread continues execution. You should not call `PendSemaphore()` within an unscheduled or critical region because if the semaphore is not available, then the thread will block, but with the scheduler disabled, execution cannot be switched to another thread. Pending with a timeout of zero on a semaphore pends without timeout.

Posting a Semaphore

Semaphores can be posted from two different scheduling domains: the thread domain and the interrupt domain. If there are threads pending on the semaphore, posting it moves the highest priority thread from the semaphore's list of pending threads to the ready queue. All other threads are left blocked on the semaphore until their timeout occurs, or the semaphore becomes available for them. If there are no threads pending on the semaphore, posting it increments the count value by one. If the maximum count, which is specified when the semaphore is created, is reached, posting the semaphore has no effect.

Posting from the Thread Domain. [Figure 3-4](#) and [Figure 3-5](#) illustrate the process of posting semaphores from the thread domain.

A thread can post a semaphore with a call to the [PostSemaphore\(\)](#) API. If a thread calls [PostSemaphore\(\)](#) from within a scheduled region (see [Figure 3-4](#)), and a higher priority thread is moved to the ready queue, the thread calling [PostSemaphore\(\)](#) is context switched out.

If a thread calls [PostSemaphore\(\)](#) from within an unscheduled region, where the scheduler is disabled, the highest priority thread pending on the semaphore is moved to the ready queue, but no context switch occurs (see [Figure 3-5](#)).

Posting from the Interrupt Domain. Interrupt subroutines can also post semaphores. [Figure 3-6 on page 3-20](#) illustrates the process of posting a semaphore from the interrupt domain.

An ISR posts a semaphore by calling the [VDK_ISR_POST_SEMAPHORE_\(\)](#) macro. The macro moves the highest priority thread to the ready queue and latches the low priority software interrupt if a call to the scheduler is required. When the ISR completes execution, and the low priority software interrupt is run, the scheduler is run. If the interrupted thread is in a scheduled region, and a higher priority thread becomes ready, the interrupted thread is switched out and the new thread is switched in.




Figure 3-4. Thread Domain/Scheduled Region: Posting a Semaphore




Figure 3-5. Thread Domain/Unscheduled Region: Posting a Semaphore




Figure 3-6. Interrupt Domain: Posting a Semaphore

Periodic Semaphores

Semaphores can also be used to schedule periodic threads. The semaphore is posted every n ticks (where n is the semaphore's period). A thread can then pend on the semaphore and be scheduled to run every time the semaphore is posted. A periodic semaphore does not guarantee that the thread pending on the semaphore is the highest priority scheduled to run, or that scheduling is enabled. All that is guaranteed is that the semaphore is posted, and the highest priority thread pending on that semaphore moves to the ready queue.

Periodic semaphores are posted by the kernel during the timer interrupt at system tick boundaries. Periodic semaphores can also be posted at any time with a call to `PostSemaphore()` or `VDK_ISR_POST_SEMAPHORE()`. Calls to these functions do not affect the periodic posting of the semaphore.

Messages

Messages are a protocol mechanism offered by many operating systems. Messages can be used to:

- Communicate information between two threads
- Control access to a shared resource
- Signal a certain occurrence and communicate information about the occurrence
- Allow two threads to synchronize

The maximum number of messages supported in the system is set up when the project is built. When the maximum number of messages is non-zero, a system-owned memory pool is created to support messaging. The properties of this memory pool should not be altered. Further information on memory pools is given in [“Memory Pools” on page 3-55](#).

Behavior of Messages

Messages allow two threads to communicate over logically separate channels. A message is sent on one of 15 possible channels, `kMsgChannel1` to `kMsgChannel15`. Messages are retrieved from these channels in priority order: `kMsgChannel1`, `kMsgChannel2`, ... `kMsgChannel15`. Each message can pass a reference to a data buffer, in the form of a message payload, from the sending thread to the receiving thread.

A thread creates a message (optionally associating a payload) and then posts (sends) the message to another thread. The posting thread continues normal execution unless the posting of the message activates a higher priority thread which is pending on (waiting to receive) the message.

A thread can pend on the receipt of a message on one or more of its channels. If a message is already queued for the thread, it receives the message and continues normal execution. If no suitable message is already queued,

the thread blocks until a suitable message is posted to the thread, or until the specified timeout occurs. If a suitable message is not posted to the thread in the time specified, the thread continues execution in its error function.

Unlike semaphores, each message always has a defined owner at any given time, and only the owning thread can perform operations on the message. When a thread creates a message, it owns the message until it posts the message to another thread. The message ownership changes to the receiving thread following the post, when it is queued on one of the receiving message's channels. The receiving thread is now the owner of the message. The only operation that can be performed on the message at this time is pending on the message, making the message and its contents available to the receiving thread.

A message can only be destroyed by its owner; therefore, a thread that receives a message is responsible for either destroying or reusing a message. Ownership of the associated payload also belongs to the thread that owns the message. The owner of the message is responsible for the control of any memory allocation associated with the payload.

Each thread is responsible for destroying any messages it owns before it destroys itself. If there are any messages left queued on a thread's receiving channels when it is destroyed, then the system destroys the queued messages. As the system has no knowledge of the contents of the payload, the system does not free any resources used by the payload.

Thread's Interaction With Messages

Threads interact with messages through the set of message APIs. The functions allow a thread to create a message, pend on a message, post a message, get and set information associated with a message, and destroy a message.

Pending on a Message

Figure 3-7 illustrates the process of pending on a message.




Figure 3-7. Pending on a Message

Threads can pend on a message with a call to `PendMessage()`, specifying one or more channels that a message is to be received on. When a thread calls `PendMessage()`, it does one of the following:

- Receives a message and continues execution
- Blocks until a message is available on the specified channel(s) or the specified timeout occurs

If messages are queued on the specified channels before the timeout occurs, the thread continues normal execution; otherwise, the thread continues execution in its error function.

Once a message has been received, you can obtain the identity of the sending thread and the channel the message was received on by calling `GetMessageReceiveInfo()`. You can also obtain information about the payload by calling `GetMessagePayload()`, which returns the type and length of the payload in addition to its location. You should not call `PendMessage()` within an unscheduled or critical region because if a message is not avail-

able, then the thread will block, but with the scheduler disabled, execution cannot be switched to another thread. Pending with a timeout of zero on a message pending without timeout.

Posting a Message

Posting a message sends the specified message and the reference to its payload to the specified thread and transfers ownership of the message to the receiving thread. The details of the message payload can be specified by a call on the [SetMessagePayload\(\)](#) function, which allows the thread to specify the payload type, length, and location before posting the message. A thread can send a message it currently owns with a call to [PostMessage\(\)](#), specifying the destination thread and the channel the message is to be sent on.

[Figure 3-8](#) illustrates the process of posting a message from a scheduled region.




Figure 3-8. Posting a Message From a Scheduled Region

If a thread calls [PostMessage\(\)](#) from within a scheduled region, and a higher priority thread is moved to the ready queue on receiving the message, then the thread calling [PostMessage\(\)](#) is context switched out.

Figure 3-9 illustrates the process of posting a message from an unscheduled region.




Figure 3-9. Posting a Message From an Unscheduled Region

If a thread calls `PostMessage()` from within an unscheduled region, even if a higher priority thread is moved to the ready queue to receive the message, the thread that calls `PostMessage()` continues to execute.

Events and Event Bits

Events and event bits are signals used to regulate thread execution based on the state of the system. An event bit is used to signal that a certain system element is in a specified state. An event is a Boolean operation performed on the state of all event bits. When the Boolean combination of event bits is such that the event evaluates to `TRUE`, all threads that are pending on the event are moved to the ready queue and the event remains `TRUE`. Any thread that pends on an event that evaluates as true does not block, but when event bits have changed causing the event to evaluate as `FALSE`, any thread that pends on that event blocks.

The number of events and event bits is limited to a processor's word size minus one. For example, on a 16-bit architecture, there can only be 15 events and event bits; and on a 32-bit architecture, there can be 31 of each.

Behavior of Events

Each event maintains the `VDK_EventData` data structure that encapsulates all the information used to calculate an event's value:

```
typedef struct
{
    bool            matchAll;
    VDK_Bitfield    values;
    VDK_Bitfield    mask;
} VDK_EventData;
```

When setting up an event, you configure a flag describing how to treat a mask and target value:

- `matchAll`: `TRUE` when an event must have an exact match on all of the masked bits. `FALSE` if a match on any of the masked bits results in the event recalculating to `TRUE`.
- `values`: The target values for the event bits masked with the `mask` field of the `VDK_EventData` structure.
- `mask`: The event bits that the event calculation is based on.

Unlike semaphores, events are `TRUE` whenever their conditions are `TRUE`, and all threads pending on the event are moved to the ready queue. If a thread pends on an event that is already `TRUE`, the thread continues to run, and the scheduler is not called. Like a semaphore, a thread pending on an event that is not `TRUE` blocks until the event becomes true, or the thread's timeout is reached. Pending with a timeout of zero on an event pends without timeout.

Global State of Event Bits

The state of all the event bits is stored in a global variable. When a user sets or clears an event bit, the corresponding bit number in the global word is changed. If toggling the event bit affects any events, that event is

recalculated. This happens either during the call to [SetEventBit\(\)](#) or [ClearEventBit\(\)](#) (if called within a scheduled region), or the next time the scheduler is enabled (with a call to [PopUnscheduledRegion\(\)](#)).

Event Calculation

To understand how events use event bits, see the following examples.

Example 1: Calculation for an *All* Event

4	3	2	1	0		event bit number
0	1	0	1	0	<—	bit value
0	1	1	0	1	<—	mask
0	1	1	0	0	<—	target value

Event is FALSE because the global event bit 2 is not the target value.

Example 2: Calculation for an *All* Event

4	3	2	1	0		event bit number
0	1	1	1	0	<—	bit value
0	1	1	0	1	<—	mask
0	1	1	0	0	<—	target value

Event is TRUE.

Example 3: Calculation for an *Any* Event

4	3	2	1	0		event bit number
0	1	0	1	0	<—	bit value
0	1	1	0	1	<—	mask
0	1	1	0	0	<—	target value

Event is TRUE since bits 0 and 3 of the target and global match.

Example 4: Calculation for an *Any* Event

4	3	2	1	0		event bit number
0	1	0	1	1	<—	bit value
0	1	1	0	1	<—	mask
0	0	1	0	0	<—	target value

Event is `FALSE` since bits 0, 2, and 3 do not match.

Effect of Unscheduled Regions on Event Calculation

Each time an event bit is set or cleared, the scheduler is entered to recalculate all dependent event values. By entering an unscheduled region, you can toggle multiple event bits without triggering spurious event calculations that could result in erroneous system conditions. Consider the following code.

```
/* Code that accidentally triggers Event1 trying to set up
   Event2. Assume the prior event bit state = 0x00. */

VDK_EventData data1 = { true, 0x1, 0x3 };
VDK_EventData data2 = { true, 0x3, 0x3 };
VDK_LoadEvent(kEvent1, data1);
VDK_LoadEvent(kEvent2, data2);
VDK_SetEventBit(kEventBit1); /* will trigger Event1 by accident */
VDK_SetEventBit(kEventBit2); /* Event1 is FALSE, Event2 is TRUE */
```

Whenever you toggle multiple event bits, you should enter an unscheduled region to avoid the above loopholes. For example, to fix the above accidental triggering of Event1 in the above code, use the following code:

```
VDK_PushUnscheduledRegion();
VDK_SetEventBit(kEventBit1); /* Event1 has not been triggered */
VDK_SetEventBit(kEventBit2); /* Event1 is FALSE, Event2 is TRUE */
VDK_PopUnscheduledRegion();
```

Thread's Interaction With Events

Threads interact with events by pending on events, setting or clearing event bits, and by loading a new `VDK_EventData` into a given event.

Pending on an Event

Like semaphores, threads can pend on an event's condition becoming `TRUE` with a timeout. Figure 3-10 illustrates the process of pending on an event.

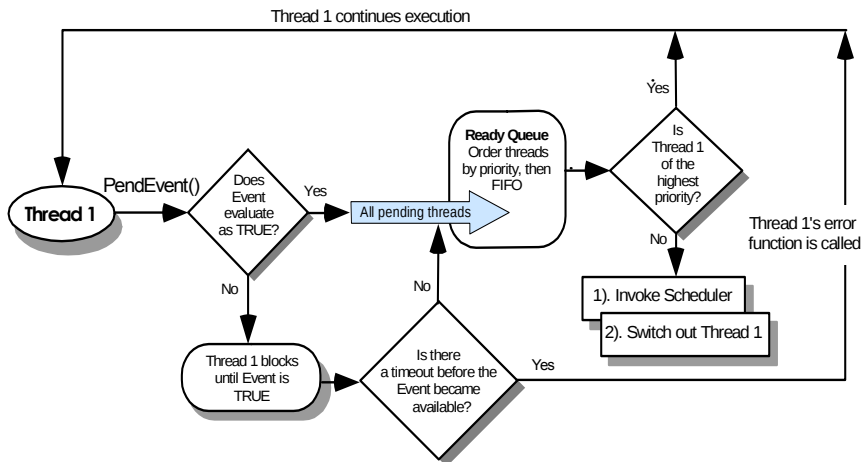


Figure 3-10. Pending on an Event

A thread calls `PendEvent()` and specifies the timeout. If the event becomes `TRUE` before the timeout is reached, the thread (and all other threads pending on the event) is moved to the ready queue. Calling `PendEvent()` with a timeout of zero means that the thread is willing to wait indefinitely.

Setting or Clearing of Event Bits

Changing the status of the event bits can be accomplished in both the interrupt domain and the thread domain. Each domain results in slightly different results.

From the Thread Domain. Figure 3-11 illustrates the process of setting or clearing of an event bit from the thread domain.




Figure 3-11. Thread Domain: Setting or Clearing an Event Bit

A thread can set an event bit by calling `SetEventBit()` and clear it by calling `ClearEventBit()`. Calling either from within a scheduled region recalculates all events that depend on the event bit and can result in a higher priority thread being context switched in.

From the Interrupt Domain. Figure 3-12 on page 3-31 illustrates the process of setting or clearing of an event bit from the interrupt domain.

An Interrupt Service Routine can call `VDK_ISR_SET_EVENTBIT_()` and `VDK_ISR_CLEAR_EVENTBIT_()` to change an event bit values and, possibly, free a new thread to run. Calling these macros *does not*




Figure 3-12. Interrupt Domain: Setting or Clearing an Event Bit

result in a recalculation of the events, but the low priority software interrupt is set and the scheduler entered. If the interrupted thread is in a scheduled region, an event recalculation takes place, and can cause a higher priority thread to be context switched in. If an ISR sets or clears multiple event bits, the calls do not need to be protected with an unscheduled region (since there is no thread scheduling in the interrupt domain); for example:

```

/* The following two ISR calls do not need to be protected: */
VDK_ISR_SET_EVENTBIT(kEventBit1);
VDK_ISR_SET_EVENTBIT(kEventBit2);

```

Loading New Event Data into an Event

From the thread scheduling domain, a thread can get the `VDK_EventData` associated with an event with the `GetEventData()` API. Additionally, a thread can change the `VDK_EventData` with the `LoadEvent()` API. A call to `LoadEvent()` causes a recalculation of the event's value. If a higher priority thread becomes ready because of the call, it starts running if the scheduler is enabled.

Device Flags

Because of the special nature of device drivers, most require synchronization methods that are similar to those provided by events and semaphores, but with different operation. Device flags are created to satisfy the specific circumstances device drivers might require. Much of their behavior cannot be fully explained without an introduction to device drivers, which are covered extensively in [“Device Drivers” on page 3-38](#).

Behavior of Device Flags

Like events and semaphores, a thread can pend on a device flag, but unlike semaphores and events, a device flag is always `FALSE`. A thread pending on a device flag immediately blocks. When a device flag is posted, all threads pending on it are moved to the ready queue.

Device flags are used to communicate to any number of threads that a device has entered a particular state. For example, assume that multiple threads are waiting for a new data buffer to become available from an A/D converter device. While neither a semaphore nor an event can correctly represent this state, a device flag’s behavior can encapsulate this system state.

Thread’s Interaction With Device Flags

A thread accesses a device flag through two APIs: [PendDeviceFlag\(\)](#) and [PostDeviceFlag\(\)](#). Unlike most APIs that can cause a thread to block, [PendDeviceFlag\(\)](#) *must* be called from within a critical region.

[PendDeviceFlag\(\)](#) is set up this way because of the nature of device drivers. See section [“Device Drivers” on page 3-38](#) for a more information about device flags and device drivers.

Interrupt Service Routines

Unlike the Analog Devices standard C implementation of interrupts (using `signal.h`), all VDK interrupts are written in assembly. The VDK encourages users to write interrupts in assembly by giving hand optimized macros to communicate between the interrupt domain and the thread domain. All calculations should take place in the thread domain, and interrupts should be short routines that post semaphores, change event bit values, activate device drivers, and drop tags in the history buffer.

Enabling and Disabling Interrupts

Each processor architecture has a slightly different mechanism for masking and unmasking interrupts. Some architectures require that the state of the interrupt mask be saved to memory before servicing an interrupt or an exception, and the mask be manually restored before returning. Since the kernel installs interrupts (and exception handlers on some architectures), directly writing to the interrupt mask register may produce unintended results. Therefore, VDK provides a simple and platform independent API to simplify access to the interrupt mask.

A call to [GetInterruptMask\(\)](#) returns the actual value of the interrupt mask, even if it has been saved temporarily by the kernel in private storage. Likewise, [SetInterruptMaskBits\(\)](#) and [ClearInterruptMaskBits\(\)](#) set and clear bits in the interrupt mask in a robust and safe manner. Interrupt levels with their corresponding bits set in the interrupt mask are enabled when interrupts are globally enabled. See the *Hardware Reference* manual for your target processor for more information about the interrupt mask.

VDK also presents a standard way of turning interrupts on and off globally. Like unscheduled regions (in which the scheduler is disabled) the VDK supports critical regions where interrupts are disabled. A call to [PushCriticalRegion\(\)](#) disables interrupts, and a call to [PopCriticalRegion\(\)](#) reenables interrupts. These API calls implement a stack style

Interrupt Service Routines

interface, as described in [“Protected Regions” on page 1-7](#). Users are discouraged from turning interrupts off for long sections of code since this increases interrupt latency.

Interrupt Architecture

Interrupt handling can be set up in two ways: support C functions and install them as handlers, or support small assembly ISRs that set flags that are handled in threads or device drivers (which are written in a high level language). Analog Devices standard C model for interrupts uses `signal.h` to install and remove signal (interrupt) handlers that can be written in C. The problem with this method is that the interrupt space requires a C run-time context, and any time an interrupt occurs, the system must perform a complete context save/restore.

VDK's interrupt architecture does not support the `signal.h` strategy for handling interrupts. VDK interrupts should be written in assembly, and their body should set some signal (semaphore, event bit or device flag) that communicates back to the thread or device driver domain. This architecture reduces the number of context saves/restores required, decreases interrupt latency, and still keeps as much code as possible in a high-level language.

The lightweight nature of ISRs also encourages the use of interrupt nesting to further reduce latency. VDK enables interrupt nesting by default on processors that support it.

Vector Table

VDK installs a common header in every entry in the interrupt table. The header disables interrupts and jumps to the interrupt handler. Interrupts are disabled in the header so that you can depend on having access to global data structures at the beginning of their handler. You must remember to reenable interrupts before executing an RTI.

The VDK reserves (at least) two interrupts: the timer interrupt and the lowest priority software interrupt. For a discussion about the timer interrupt, see [“Timer ISR” on page 3-37](#). For information about the lowest priority software interrupt, see [“Reschedule ISR” on page 3-37](#). For information on any additional interrupts reserved by the VDK for particular processors, see [“Processor-Specific Notes” on page A-1](#).

Global Data

Often ISRs need to communicate data back and forth to the thread domain besides semaphores, event bits, and device driver activations. ISRs can use global variables to get data to the thread domain, but you must remember to wrap any access to or from that global data in a critical region and to declare the variable as `volatile` (in C/C++). For example, consider the following:

```
/* MY_ISR.asm */
.EXTERN _my_global_integer;
<REG> = data;
DM(_my_global_integer) = <REG>;
/* finish up the ISR, enable interrupts, and RTI. */
```

And in the thread domain:

```
/* My_C_Thread.c */
volatile int my_global_integer;

/* Access the global ISR data */
VDK_PushCriticalRegion();
if (my_global_integer == 2)
    my_global_integer = 3;
VDK_PopCriticalRegion();
```

Communication with the Thread Domain

The VDK supplies a set of macros that can be used to communicate system state to the thread domain. Since these macros are called from the interrupt domain, they make no assumptions about processor state, available registers, or parameters. In other words, the ISR macros can be called without consideration of saving state or having processor state trampled during a call.

Take for example, the following three equivalent [VDK_ISR_POST_SEMAPHORE_\(\)](#) calls:

```
.VAR/DATA  semaphore_id;

/* Pass the value directly */
VDK_ISR_POST_SEMAPHORE_(kSemaphore1);

/* Pass the value in a register */
<REG> = kSemaphore1;
VDK_ISR_POST_SEMAPHORE_(<REG>);
/* <REG> was not trampled */

/* Post the semaphore one last time using a DM */
DM(semaphore_id) = <REG>;
VDK_ISR_POST_SEMAPHORE_(DM(semaphore_id));
```

Additionally, no condition codes are affected by the ISR macros, no assumptions are made about having space on any hardware stacks (e.g. PC or status), and all VDK internal data structures are maintained.

Most ISR macros raise the low priority software interrupt if thread domain scheduling is required after all other interrupts are serviced. For a discussion of the low priority software interrupt, see section [“Reschedule ISR” on page 3-37](#). Refer to [“Processor-Specific Notes” on page A-1](#) for additional information about ISR APIs.

Within the interrupt domain, every effort should be made to enable interrupt nesting. Nesting is always disabled when an ISR begins. However, leaving it disabled is analogous to staying in an unscheduled region in the thread domain; other ISRs are prevented from executing, even if they have higher priority. Allowing nested interrupts potentially lowers interrupt latency for high priority interrupts.

Timer ISR

The VDK reserves a timer interrupt. The timer is used to calculate round-robin times, sleeping threads' time to keep sleeping, and periodic semaphores. One VDK tick is defined as the time between timer interrupts and is the finest resolution measure of time in the kernel. The timer interrupt can cause a low priority software interrupt (see [“Reschedule ISR” on page 3-37](#)).

Reschedule ISR

The VDK designates the lowest priority interrupt that is not tied to a hardware device as the reschedule ISR. This ISR handles housekeeping when an interrupt causes a system state change that can result in a new high priority thread becoming ready. If a new thread is ready and the system is in a scheduled region, the software ISR saves off the context of the current thread and switches to the new thread. If an interrupt has activated a device driver, the low priority software interrupt calls the dispatch function for the device driver. [For more information, see “Dispatch Function” on page 3-43.](#)

On systems where the lowest priority non-hardware-tied interrupt is not the lowest priority interrupt, all lower priority interrupts must run with interrupts turned off for their entire duration. Failure to do so may result in undefined behavior.

I/O Interface

The I/O interface provides the mechanism for creating an interface between the external environment and VDK applications.

In VisualDSP++ 3.x, only device driver objects can be used to construct the I/O interface.

I/O Templates

I/O templates are analogous to thread types. I/O templates are used to instantiate I/O objects. In VisualDSP++ 3.x, the only types of I/O templates available, and therefore the only classes of I/O objects, are for device drivers. In order to create an instance of a device driver, a boot I/O object must be added to the VDK project using the device driver template.

Device Drivers

The role of a device driver is to abstract the details of the hardware implementation from the software designer. For example, a software engineer designing a finite impulse response (FIR) filter does not need to understand the intricacies of the converters, and is able to concentrate on the FIR algorithm. The software can then be reused on different platforms, where the hardware interface differs.

The Communication Manager controls device drivers in the VDK. Using the Communication Manager APIs, you can maintain the abstraction layers between device drivers, interrupt service routines, and executing threads. This section details how the Communication Manager is organized.



In VisualDSP++ 3.x, device drivers are a part of the I/O interface. Device drivers are added to a VDK project as I/O objects. VisualDSP++ 2.0 device drivers are not compatible with VisualDSP++

3.x device drivers. See [“Migrating Device Drivers” on page B-1](#) for a description of how to convert existing VisualDSP++ 2.0 device drivers for use in VisualDSP++ 3.x projects.

Execution

Device drivers and interrupt service routines are tied very closely together. Typically, DSP developers prefer to keep as much time critical code in assembly as possible. The Communication Manager is designed such that you can keep interrupt routines in assembly (the time critical pieces), and interface and resource management for the device in a high level language without sacrificing speed. The Communication Manager attempts to keep the number of context switches to a minimum, to execute management code at reasonable times, and to preserve the order of priorities of running threads when a thread uses a device. However, you need to thoroughly understand the architecture of the Communication Manager to write your device driver.

There is only one interface to a device driver—through a dispatch function. The dispatch function is called when the device is initialized, when a thread uses a device (open/close, read/write, control), or when an interrupt service routine transfers data to or from the device. The dispatch function handles the request and returns. Device drivers should *not* block (pend) when servicing an initialize request or a request for more data by an interrupt service routine. However, a device driver can block when servicing a thread request and the relevant resource is not ready or available. Device driver initialization and ISR requests are handled within critical regions enforced by the kernel, so their execution does not have to be reentrant, but a thread level request must protect global variables within critical or unscheduled regions.

Parallel Scheduling Domains

This section focuses on a unique role of device drivers in the VDK architecture. Understanding device drivers requires some understanding of the time and method by which device driver code is invoked. VDK applications may be factored into two *domains*, referred to as the thread domain and the ISR domain (see [Figure 3-13](#)). This distinction is not an arbitrary or unnecessary abstraction. The hardware architecture of the processor as well as the software architecture of the kernel reinforces this notion. You should consider this distinction when you are designing your application and apportioning your code.




Figure 3-13. Parallel Scheduling Domains

Threads are scheduled based on their priority and the order in which they are placed in the ready queue. The scheduling portion of the kernel is responsible for selecting the thread to run. However, the scheduler does not have complete control over the processor. It may be preempted by a parallel and higher priority scheduler: the interrupt and exception hardware. While interrupts or exceptions are being serviced, thread priorities

are temporarily moot. The position of threads in the ready queue becomes significant again only when the hardware relinquishes control back to the software based scheduler.

Each of the domains has strengths and weaknesses that dictate the type of code suitable to be executed in that environment. The scheduler in the thread domain is invoked when threads are moved to or from the ready queue. Threads each have their own stack and may be written in a high level language. Threads always execute in “supervisor” or “kernel mode” (if the processor makes this distinction). Threads implement algorithms and are allotted processor time based on the completion of higher priority activity.

In contrast, scheduling in the interrupt domain has the highest system wide priority. Any “ready” ISR takes precedence over any ready thread (outside critical regions), and this form of scheduling is implemented in hardware. ISRs are always written in assembly and must manually restore any registers they use. ISRs execute in “supervisor” or “kernel mode” (if the processor makes this distinction). ISRs respond to asynchronous peripherals at the lowest level only. The routine should perform only activities that are so time critical that data would be lost if the code were not executed as soon as possible. All other activity should occur under the control of the kernel's scheduler based on priority.

Transferring from the thread domain to the interrupt domain is simple and automatic, but returning to the thread domain can be much more laborious. If the ready queue is not changed while in the interrupt domain, then the scheduler need not run when it regains control of the system. The interrupted thread resumes execution immediately. If the ready queue has changed, the scheduler must further determine whether the highest priority thread has changed. If it has changed, the scheduler must initiate a context switch.

Device drivers fill the gap between the two scheduling domains. They are neither thread code nor ISR code, and they are not directly scheduled by either the kernel or the interrupt controller. Device drivers are imple-

mented as C++ objects and run on the stack of the currently running thread. However, they are not “owned” by any thread, and may be used by many threads concurrently.

Using Device Drivers

From the point of view of a thread, there are five functional interfaces to device drivers: `OpenDevice()`, `CloseDevice()`, `SyncRead()`, `SyncWrite()`, and `DeviceIOCtl()`. The names of the APIs are self explanatory since threads mostly treat device drivers as black boxes. Figure 3-14 illustrates the device drivers’ interface. A thread uses a device by opening it, reading and/or writing to it, and closing it. The `DeviceIOCtl()` function is used for sending device-specific control information messages. Each API is a standard C/C++ function call that runs on the stack of the calling thread and returns when the function completes. However, when the device driver does not have a needed resource, one of these functions may cause the thread to be removed from the ready queue and block on a signal, similar to a semaphore or an event, called a device flag.




Figure 3-14. Device Driver APIs

Interrupt service routines have only one API call relating to device drivers: `VDK_ISR_ACTIVATE_DEVICE_()`. This macro is not a function call, and program flow does not transfer from the ISR to the device driver and back. Rather, the macro sets a flag indicating that the device driver's “activate” routine should execute after all interrupts have been serviced.

The remaining two API functions, `PendDeviceFlag()` and `PostDeviceFlag()`, can be called only from within the device driver itself. For example, a call from a thread to `SyncRead()` might cause the device driver to call `PendDeviceFlag()` if there is no data currently available. This would cause the thread to block until the device flag is posted by another code fragment within the device driver that is providing the data.

As another example, when an interrupt occurs because an incoming data buffer is full, the ISR might move a pointer so that the device begins filling an empty buffer before calling `VDK_ISR_ACTIVATE_DEVICE_()`. The device driver's activate routine may respond by posting a device flag and moving a thread to the ready queue so that it can be scheduled to process the new data.

Dispatch Function

The dispatch function is the core of any device driver. It takes two parameters and returns a `void*` (the return value depends on the input values). A dispatch function declaration for a device driver is as follows:

C Driver Code:

```
void* MyDevice_DispatchFunction(VDK_DispatchID    inCode,
                               VDK_DispatchUnion inData);
```

C++ Driver Code:

```
void* MyDevice::DispatchFunction(VDK::DispatchID    inCode,
                                VDK::DispatchUnion &inData);
```

I/O Interface

The first parameter is an enumeration that specifies why the dispatch function has been called:

```
enum VDK_DispatchID
{
    VDK_kIO_Init,
    VDK_kIO_Activate,
    VDK_kIO_Open,
    VDK_kIO_Close,
    VDK_kIO_SyncRead,
    VDK_kIO_SyncWrite,
    VDK_kIO_IOCTL
};
```

The second parameter is a union whose value depends on the enumeration value:

```
union VDK_DispatchUnion
{
    struct OpenClose_t
    {
        void          **dataH;
        char          *flags; /* used for kIO_Open only */
    };
    struct ReadWrite_t
    {
        void          **dataH;
        VDK_Ticks     timeout;
        unsigned int  dataSize;
        int           *data;
    };
    struct IOCTL_t
    {
        void          **dataH;
        VDK_Ticks     timeout;
    };
};
```

```

        int             command;
        char            *parameters;
    };
};

```

The values in the union are only valid when the enumeration specifies that the dispatch function has been called from the thread domain (kIO_Open, kIO_Close, kIO_SyncRead, kIO_SyncWrite, kIO_IOCtl).

A device driver's dispatch function can be structured as follows:

In C:

```

void* MyDevice_DispatchFunction(VDK_DispatchID    inCode,
                               VDK_DispatchUnion inData)
{
    switch(inCode)
    {
        case VDK_kIO_Init:
            /* Init the device */
        case VDK_kIO_Activate:
            /* Get more data ready for the ISR */
        case VDK_kIO_Open:
            /* A thread wants to open the device... */
            /* Allocate memory and prepare everything else */
        case VDK_kIO_Close:
            /* A thread is closing a connection to the device... */
            /* Free all the memory, and do anything else */
        case VDK_kIO_SyncRead:
            /* A thread is reading from the device */
            /* Return an unsigned int of the num. of bytes read */
        case VDK_kIO_SyncWrite:
            /* A thread is writing to the device */
            /* Return an unsigned int of the number of bytes */
            /* written */
        case VDK_kIO_IOCtl:
    }
}

```

I/O Interface

```
        /* A thread is performing device-specific actions: */
default:
        /* Invalid DispatchID code */
        return 0;
    }
}
```

In C++:

```
void* MyDevice::DispatchFunction(VDK::DispatchID    inCode,
                                VDK::DispatchUnion &inData)
{
    switch(inCode)
    {
        case VDK::kIO_Init:
            /* Init the device */
        case VDK::kIO_Activate:
            /* Get more data ready for the ISR */
        case VDK::kIO_Open:
            /* A thread wants to open the device... */
            /* Allocate memory and prepare everything else */
        case VDK::kIO_Close:
            /* A thread is closing a connection to the device...*/
            /* Free all the memory, and do anything else */
        case VDK::kIO_SyncRead:
            /* A thread is reading from the device */
            /* Return an unsigned int of the num. of bytes read */
        case VDK::kIO_SyncWrite:
            /* A thread is writing to the device */
            /* Return an unsigned int of the number of bytes */
            /* written */
        case VDK::kIO_IOCTL:
            /* A thread is performing device-specific actions: */
        default:
            /* Invalid DispatchID code */
    }
```

```

        return 0;
    }
}

```

Each of the different cases in the dispatch function are discussed below.

Init

The device dispatch function is called with the `VDK_kIO_Init` parameter for C style device and `VDK::kIO_Init` for C++ style drivers at system boot time. All device-specific data structures and system resources should be set up at this time. The device driver *should not* call any APIs that throw an error or might block. Additionally, the init function is called within a critical region, and the device driver should not push/pop critical regions, or wait on interrupts.

Open or Close

When a thread opens or closes a device with [OpenDevice\(\)](#) or [CloseDevice\(\)](#), the device dispatch function is called with `VDK_kIO_Open` or `VDK_kIO_Close`. The dispatch function is called from the thread domain, so any stack-based variables are local to that thread. Access to shared data (data that may be accessed by threads and/or interrupts and/or device driver activate functions) should be appropriately protected by the use of unscheduled regions, critical regions, or other means.

When a thread calls the dispatch function attempting to open or close a device, the API passes a union to the device dispatch function whose value is defined with the `OpenClose_t` of the `VDK_DispatchUnion`. The `OpenClose_t` is defined as follows.

```

struct OpenClose_t
{
    void    **dataH;
    char    *flags;    /* used for kIO_Open only */
};

```

I/O Interface

`OpenClose_t.dataH`: A pointer to a thread-specific location that a device driver can use to hold any thread-specific resources. For example, a thread can `malloc` space for a structure that describes the state of a thread associated with a device. The pointer to the structure can be stored in `*dataH`, which is then accessible to every other dispatch call involving this thread. A device driver can free the space when the thread calls `CloseDevice()`.

`OpenClose_t.flags`: The second parameter passed to an `OpenDevice()` call is supplied to the dispatch function as the value of `OpenClose_t.flags`. This is used to pass any device-specific flags relevant to the opening of a device. Note that this part of the union is not used on a call to `CloseDevice()`.

Read or Write

A thread that needs to read or write to a device it has opened calls `SyncRead()` or `SyncWrite()`. The dispatch function is called in the thread domain and on the thread's stack. These functions call the device dispatch function with the parameters passed to the API in the `VDK_DispatchUnion`, and the flags `VDK_kIO_SyncRead` or `VDK_kIO_SyncWrite`. The `ReadWrite_t` is defined as follows:

```
struct ReadWrite_t
{
    void          **dataH;
    VDK::Ticks    timeout;
    unsigned int  dataSize;
    int           *data;
};
```

`ReadWrite_t.dataH`: A thread-specific location, which is passed to the dispatch function on the opening of a device by an `OpenDevice()` call. This variable can be used to store a pointer to a thread-specific data structure detailing what state the thread is in while dealing with the device.

`ReadWrite_t.timeout`: The amount of time in [Ticks](#) that a thread is willing to wait for the completion of a [SyncRead\(\)](#) or [SyncWrite\(\)](#) call. If this timeout behavior is required, it must be implemented by using the value of `ReadWrite_t.timeout` as an argument to an appropriate [PendDeviceFlag\(\)](#) call in the dispatch function.

`ReadWrite_t.dataSize`: The amount of data that the thread reads from or writes to the device.

`ReadWrite_t.data`: A pointer to the location that the thread writes the data to (on a read), or reads from (on a write).

Like calls to the device dispatch function for opening and closing, the calls to read and write are not protected with a critical or unscheduled region. If a device driver accesses global data structures during a read or write, the access should be protected with critical or unscheduled regions. See the discussion in [“Device Drivers” on page 3-38](#) for more information about regions and pending.

IOctl

The VDK supplies an interface for threads to control a device’s parameters with the [DeviceIOctl\(\)](#) API. When a thread calls [DeviceIOctl\(\)](#), the function sets up some parameters and calls the specified device’s dispatch function with the value `VDK_kIO_IOCTL` and the `VDK_DispatchUnion` set up as a `IOctl_t`.

The `IOctl_t` is defined as follows:

```
struct IOctl_t
{
    void          **dataH;
    VDK_Ticks     timeout;
    int           command;
    char          *parameters;
};
```

I/O Interface

`IOctl_t.dataH`: A thread-specific location, which is passed to the dispatch function on the opening of a device by an [OpenDevice\(\)](#) call. This variable can be used to store a pointer to a thread-specific data structure detailing what state the thread is in while dealing with the device.

`IOctl_t.timeout`: The amount of time in ticks that a thread is willing to wait for the completion of a [DeviceIOctl\(\)](#) call. If this timeout behavior is required, it must be implemented by using the value of `IOctl_t.timeout` as an argument to an appropriate [PendDeviceFlag\(\)](#) call in the dispatch function.

`IOctl_t.command`: A device-specific integer (second parameter from the [DeviceIOctl\(\)](#) function).

`IOctl_t.parameters`: A device-specific pointer (third parameter from the [DeviceIOctl\(\)](#) function).

Like read/write and open/close, a device dispatch function call for `IOctl` is not protected by a critical or unscheduled region. If a device accesses global data structures, the device driver should protect them with a critical or an unscheduled region.

Activate

Often a device driver needs to respond to state changes caused by ISRs. The device dispatch function is called with a value `VDK_kIO_Activate` at some point after an ISR has called the macro [VDK_ISR_ACTIVATE_DEVICE_\(\)](#).

When the ISR calls [VDK_ISR_ACTIVATE_DEVICE_\(\)](#), a flag is set indicating that a device has been activated, and the low priority software interrupt is triggered to run (see “[Reschedule ISR](#)” on page 3-37). When the scheduler is entered through the low priority software interrupt, the device’s dispatch function is called with the `VDK_kIO_Activate` value.

The activate part of a device dispatch function should handle posting signals, so that threads waiting on certain device states can continue running. For example, assume that a D/A ISR runs out of data in its buffer. The ISR would call [VDK_ISR_ACTIVATE_DEVICE_\(\)](#) with the `IOID` of the device driver. When the device dispatch function is called with the `VDK_kIO_Activate`, the device posts a device flag, semaphore, or sets an event bit that reschedules any threads that are pending.

Since the activate function is run from the low priority software interrupt, some uncommon circumstances exist. While the dispatch function is executing a `VDK_kIO_Activate` command, it executes in an unscheduled region and runs off the kernel's stack. Since there are no threads executing, the dispatch function must avoid calling APIs that throw an error or can block.

Device Flags

Device flags are synchronization primitives, similar to semaphores, events, and messages, but device flags have a special association with device drivers. Like semaphores and events, a thread can pend on a device flag. This means that the thread waits until the flag is posted by a device driver. The post typically occurs from the activate function of a device driver's dispatch function.

Pending on a Device Flag

When a thread pends on a device flag, unlike with semaphores, events, and messages, the thread always blocks. The thread waits until the flag is posted by another call to the device's dispatch function. When the flag is posted, all threads that are pending on the device flag are moved to the ready queue. Since posting a device flag with the [PostDeviceFlag\(\)](#) API moves an indeterminate number of threads to the ready queue, the call is not deterministic. For more information about posting device flags, see [“Posting a Device Flag” on page 3-52](#).

The rules for pending on device flags are strict compared to other types of signals. The “stack” of critical regions must be exactly one level deep when a thread pends on a device flag. In other words, with interrupts enabled, call `PushCriticalRegion()` exactly once prior to calling `PendDeviceFlag()` from a thread. The reason for this condition becomes clear if you consider the reason for pending. A thread pends on a device flag when it is waiting for a condition to be set from an ISR. However, you must enter a critical region before examining any condition that may be modified from an ISR to ensure that the value you read is valid. Furthermore, `PendDeviceFlag()` pops the critical region stack once, effectively balancing the earlier call to `PushCriticalRegion()`. For example, a typical device driver uses device flags in the following manner.

```
VDK_PushCriticalRegion();
while(should_loop != 0)
{
    /* ... */
    /* access global data structures */
    /* and figure out if we should keep looping */
    /* ... */
    /* Wait for some device state */
    VDK_PendDeviceFlag();
    /* Must reenter the critical region */
    VDK_PushCriticalRegion();
}
VDK_PopCriticalRegion();
```

Figure 3-15 illustrates the process of pending on a device flag.

Posting a Device Flag

Like semaphores and messages, a device flag can be posted. A device dispatch function posts a device flag with a call to `PostDeviceFlag()`. Unlike semaphores and messages, the call moves *all* threads pending on the device

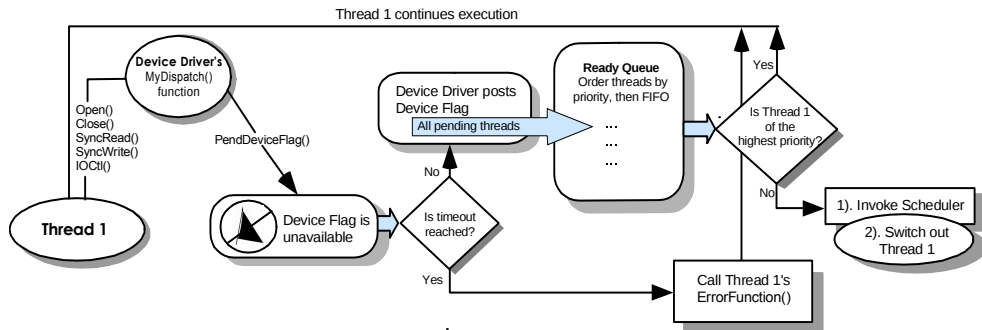


Figure 3-15. Pending on a Device Flag

flag to the ready queue and continues execution. Once `PostDeviceFlag()` returns, subsequent calls to `PendDeviceFlag()` cause the thread to block (as before).

Note that the `PostDeviceFlag()` API does not throw any errors. The reason for this is that the API function is called typically from the dispatch function when the dispatch function has been called with `VDK_kIO_Activate`. This is because the device dispatch function operates on the kernel's stack when it is called with `VDK_kIO_Activate` rather than on the stack of a thread.

Figure 3-16 illustrates the process of posting a device flag.

General Notes

Keep the following tips in mind while writing device drivers. Although many of these topics also apply to threads, they deserve special mention with respect to device drivers.




Figure 3-16. Posting a Device Flag

Variables

Device drivers and ISRs are closely linked. Since ISRs and the dispatch function access the same variables, the easiest way is to declare the variables in the C/C++ device driver routine, and to access them as `extern` from within the assembly ISR. When declaring these variables in the C/C++ source file, you must declare them as `volatile` to ensure that the compiler optimizer is aware that their values may be changed externally to the C/C++ code at any time. Additionally, care must be taken in the ISR to refer to variables defined in C/C++ code correctly, by their decorated/mangled names.

Notes on Critical/Unscheduled Regions

Since many of the data structures and variables associated with a device driver are shared between multiple threads and ISRs, access to them must be protected within critical regions and unscheduled regions. Critical regions keep ISRs from modifying data structures unexpectedly, and unscheduled regions prevent other threads from modifying data structures.

Care must be taken when pending on device flags to remain in the right regions. Device flags must be pending on from within a non-nested critical region and the return from the `PendDeviceFlag()` call pops the critical

region. Therefore, careful region balancing must be adhered to. Take advantage of the errors that are thrown by the APIs when error checking is enabled to detect mismatches.

Memory Pools

Common problems experienced with memory allocation using `malloc` are fragmentation of the heap as well as non-deterministic search times for finding a free area of the heap with the requested size. The memory pool manager uses the defined pools to provide an efficient, deterministic memory allocation scheme as an alternative to `malloc`. The use of memory pools for memory allocation can be advantageous when an application requires significant allocation and deallocation of objects of the same size.

A memory pool is an area of memory subdivided into equally sized memory blocks. Each memory pool contains memory blocks of a single size, but multiple pools can be defined, each with a different block size. Furthermore, on architectures that support the definition of multiple heaps, the heap that a pool is to use can be specified.

Memory Pool Functionality

When creating a memory pool, the block size and number of blocks in the pool are specified. The memory pool manager allocates the memory required for the pool and splits it into blocks at creation time if required. VDK allows the specification that blocks should be created on demand at run time, rather than during the creation of a pool. On demand creation of blocks reduces the overhead at the time the pool is created but increases the runtime overhead when obtaining a new block from the pool. Additionally, allocation and deallocation of a block from a pool is deterministic if the blocks are created when the pool is created.

Memory Pools

In order to conform to memory alignment constraints, the block size specified for a pool is rounded up internally, so that its size is a multiple of the size of a pointer on the architecture in question—all block addresses returned by [MallocBlock\(\)](#) are a multiple of `sizeof(void *)`.

4 VDK DATA TYPES

VDK software comes with the predefined set of data types. This chapter describes the current release of the kernel. Future releases may include more types.

Data Type Summary

VDK data types are summarized in [Table 4-1](#). A description of each type begins [on page 4-3](#).

Table 4-1. VDK Data Types

Data Type	Reference Page
Bitfield	on page 4-3
DeviceDescriptor	on page 4-4
DeviceFlagID	on page 4-5
DispatchID	on page 4-6
DispatchUnion	on page 4-7
DSP_Family	on page 4-9
DSP_Product	on page 4-10
EventBitID	on page 4-12
EventID	on page 4-13
EventData	on page 4-14
HistoryEnum	on page 4-15

Data Type Summary

Table 4-1. VDK Data Types (Cont'd)

Data Type	Reference Page
IMASKStruct	on page 4-17
IOID	on page 4-18
IOTemplateID	on page 4-19
MessageID	on page 4-20
MsgChannel	on page 4-21
PoolID	on page 4-23
Priority	on page 4-24
SemaphoreID	on page 4-25
SystemError	on page 4-26
ThreadCreationBlock	on page 4-29
ThreadID	on page 4-31
ThreadStatus	on page 4-32
ThreadType	on page 4-33
Ticks	on page 4-34
VersionStruct	on page 4-35

Bitfield

The **Bitfield** type is used to store the bit pattern. The size of a **Bitfield** item is the size of a data word:

- 16 bits on ADSP-219x DSPs
- 32 bits on ADSP-21xxx, ADSP-TSxxx, and Blackfin processors

In C:

```
typedef struct
{
    unsigned int    Bit0:1;
    ...
    unsigned int    BitN:1;
} VDK_Bitfield;
```

In C++:

```
typedef struct
{
    unsigned int    Bit0:1;
    ...
    unsigned int    BitN:1;
} VDK::Bitfield;
```

DeviceDescriptor

The [DeviceDescriptor](#) type is used to store the unique identifier of an opened device. The value is obtained dynamically as the return value from [OpenDevice\(\)](#).

In C:

```
typedef unsigned int VDK_DeviceDescriptor;
```

In C++:

```
typedef unsigned int VDK::DeviceDescriptor;
```

DeviceFlagID

The **DeviceFlagID** type is used to store the unique identifier of a device flag.

```
enum DeviceFlagID
{
    /* Defined by IDDE in the vdk.h file. */
};
```

The enumeration in `Vdk.h` will only contain the IDs of the device flags enabled at boot time. Any dynamically created device flags will have an ID of the same type to allow the compiler to do type checking and prevent errors.

In C:

```
typedef enum DeviceFlagID VDK_DeviceFlagID;
```

In C/C++:

```
typedef enum DeviceFlagID VDK::DeviceFlagID;
```

DispatchID

The [DispatchID](#) type enumerates a device driver's dispatch commands.

In C:

```
enum VDK_DispatchID
{
    VDK_kIO_Init,
    VDK_kIO_Activate,
    VDK_kIO_Open,
    VDK_kIO_Close,
    VDK_kIO_SyncRead,
    VDK_kIO_SyncWrite,
    VDK_kIO_IOCTL
};
```

In C++:

```
enum VDK::DispatchID
{
    VDK::kIO_Init,
    VDK::kIO_Activate,
    VDK::kIO_Open,
    VDK::kIO_Close,
    VDK::kIO_SyncRead,
    VDK::kIO_SyncWrite,
    VDK::kIO_IOCTL
};
```

DispatchUnion

A variable of the [DispatchUnion](#) type is passed as a parameter to a device driver dispatch function. Calls to [OpenDevice\(\)](#), [SyncRead\(\)](#), [SyncWrite\(\)](#), and [DeviceIOCtl\(\)](#) sets up the relevant members of this union before calling the device driver's dispatch function.

In C:

```
union VDK_DispatchUnion
{
    struct
    {
        void    **dataH;
        char    *flags; /* used for kIO_Open only */
    } OpenClose_t;
    struct
    {
        void            **dataH;
        VDK_Ticks       timeout;
        unsigned int    dataSize;
        char            *data;
    } ReadWrite_t;
    struct
    {
        void            **dataH;
        void            *command;
        char            *parameters;
    } IOCtl_t;
};
```

In C++:

```
union VDK::DispatchUnion
{
    struct
```

DispatchUnion

```
{
    void      **dataH;
    char      *flags; /* used for kIO_Open only */
} OpenClose_t;
struct
{
    void      **dataH;
    VDK::Ticks timeout;
    unsigned int dataSize;
    char      *data;
} ReadWrite_t;
struct
{
    void      **dataH;
    void      *command;
    char      *parameters;
} IOCtl_t;
};
```

DSP_Family

The [DSP_Family](#) type enumerates the processor families supported by VDK. See also [VersionStruct](#) on page 4-35.

In C:

```
enum VDK_DSP_Family
{
    VDK_kUnknownFamily,
    VDK_kSHARC,
    VDK_k219X,
    VDK_kTSXXX,
    VDK_kBLACKFIN
};
```

In C++:

```
enum VDK::DSP_Family
{
    VDK::kUnknownFamily,
    VDK::kSHARC,
    VDK::k219X,
    VDK::kTSXXX,
    VDK::kBLACKFIN
};
```

DSP_Product

The [DSP_Product](#) type enumerates the devices supported by VDK. See also [VersionStruct](#) on page 4-35.

In C:

```
enum VDK_DSP_Product
{
    VDK_kUnknownProduct,
    VDK_k21060,
    VDK_k21061,
    VDK_k21062,
    VDK_k21065,
    VDK_k21160,
    VDK_k21161,
    VDK_k2191,
    VDK_k2192,
    VDK_k2195,
    VDK_k2196,
    VDK_kBF535,
    VDK_kBF532,
    VDK_kBF531,
    VDK_kBF533,
    VDK_kDM102,
    VDK_kAD6532,
    VDK_kTS101
};
```

In C++:

```
enum VDK::DSP_Product
{
    VDK::kUnknownProduct,
    VDK::k21060,
    VDK::k21061,
```

```
    VDK::k21062,  
    VDK::k21065,  
    VDK::k21160,  
    VDK::k21161,  
    VDK::k2191,  
    VDK::k2192,  
    VDK::k2195,  
    VDK::k2196,  
    VDK::kBF535,  
    VDK::kBF532,  
    VDK::kBF531,  
    VDK::kBF533,  
    VDK::kDM102,  
    VDK::kAD6532,  
    VDK::kTS101  
};
```

EventBitID

The [EventBitID](#) type is used to store the unique identifier of an event bit. The total number of event bits in a system is the size of a data word minus one:

- 15 bits on ADSP-219x DSPs
- 31 bits on ADSP-21xxx, ADSP-TSxxx, and Blackfin processors

```
enum EventBitID
{
    /* Defined by IDDE in the vdk.h file. */
};
```

In C:

```
typedef enum EventBitID VDK_EventBitID;
```

In C/C++:

```
typedef enum EventBitID VDK::EventBitID;
```

EventID

The **EventID** type is used to store the unique identifier of an event. The total number of events in a system is the size of a data word minus one:

- 15 bits on ADSP-219x DSPs
- 31 bits on ADSP-21xxx, ADSP-TSxxx, and Blackfin processors

```
enum EventID
{
    /* Defined by IDDE in the vdk.h file. */
};
```

In C:

```
typedef enum EventID VDK_EventID;
```

In C/C++:

```
typedef enum EventID VDK::EventID;
```

EventData

The `EventData` type has two `Bitfield` (mask and values) and a Boolean (*[Match All| Match Any]* logic) data members.

In C:

```
typedef struct
{
    bool            matchAll;
    VDK_Bitfield    values;
    VDK_Bitfield    mask;
} VDK_EventData;
```

In C++:

```
typedef struct
{
    bool            matchAll;
    VDK::Bitfield   values;
    VDK::Bitfield   mask;
} VDK::EventData;
```

HistoryEnum

The [HistoryEnum](#) type enumerates the events that can be logged with a call to the [LogHistoryEvent\(\)](#) API or [VDK_ISR_LOG_HISTORY_EVENT_\(\)](#) macro.

In C:

```
enum VDK_HistoryEnum {
    VDK_kThreadCreated = INT_MIN,
    VDK_kThreadDestroyed,
    VDK_kSemaphorePosted,
    VDK_kSemaphorePended,
    VDK_kEventBitSet,
    VDK_kEventBitCleared,
    VDK_kEventPended,
    VDK_kDeviceFlagPended,
    VDK_kDeviceFlagPosted,
    VDK_kDeviceActivated,
    VDK_kThreadTimedOut,
    VDK_kThreadStatusChange,
    VDK_kThreadSwitched,
    VDK_kMaxStackUsed,
    VDK_kPoolCreated,
    VDK_kPoolDestroyed,
    VDK_kDeviceFlagCreated,
    VDK_kDeviceFlagDestroyed,
    VDK_kMessagePosted,
    VDK_kMessagePended,
    VDK_kSemaphoreCreated,
    VDK_kSemaphoreDestroyed,
    VDK_kMessageCreated,
    VDK_kMessageDestroyed,
    VDK_kMessageTakenFromQueue,
```

HistoryEnum

```
    VDK_kUserEvent = 1  
};
```

In C++:

```
enum VDK::HistoryEnum  
{  
    VDK::kThreadCreated = INT_MIN,  
    VDK::kThreadDestroyed,  
    VDK::kSemaphorePosted,  
    VDK::kSemaphorePended,  
    VDK::kEventBitSet,  
    VDK::kEventBitCleared,  
    VDK::kEventPended,  
    VDK::kDeviceFlagPended,  
    VDK::kDeviceFlagPosted,  
    VDK::kDeviceActivated,  
    VDK::kThreadTimedOut,  
    VDK::kThreadStatusChange,  
    VDK::kThreadSwitched,  
    VDK::kMaxStackUsed,  
    VDK::kPoolCreated,  
    VDK::kPoolDestroyed,  
    VDK::kDeviceFlagCreated,  
    VDK::kDeviceFlagDestroyed,  
    VDK::kMessagePosted,  
    VDK::kMessagePended,  
    VDK::kSemaphoreCreated,  
    VDK::kSemaphoreDestroyed,  
    VDK::kMessageCreated,  
    VDK::kMessageDestroyed,  
    VDK::kMessageTakenFromQueue,  
    VDK::kUserEvent = 1  
};
```

IMASKStruct

The `IMASKStruct` type is a platform-dependent type used by the `ClearInterruptMaskBits()`, `GetInterruptMask()`, and `SetInterruptMaskBits()` APIs to modify the interrupt mask.

For the TigerSHARC DSP family, this type is defined as:

In C:

```
typedef unsigned long long VDK_IMASKStruct;
```

In C++:

```
typedef unsigned long long VDK::IMASKStruct;
```

On the ADSP-219x, Blackfin, and SHARC processor families, the type is defined as:

In C:

```
typedef unsigned int VDK_IMASKStruct;
```

In C++:

```
typedef unsigned int VDK::IMASKStruct;
```

IOID

The **IOID** type is used to store the unique identifier of an I/O object.

```
enum IOID
{
    /* Defined by IDDE in the vdk.h file. */
};
```

In C:

```
typedef enum IOID VDK_IOID;
```

In C/C++:

```
typedef enum IOID VDK::IOID;
```

IOTemplateID

The **IOTemplateID** type is used to store the unique identifier of an I/O object class.

```
enum IOTemplateID
{
    /* Defined by IDDE in the vdk.h file. */
};
```

In C:

```
typedef enum IOTemplateID VDK_IOTemplateID;
```

In C/C++:

```
typedef enum IOTemplateID VDK::IOTemplateID;
```

MessageID

The [MessageID](#) type is used to store the unique identifier of a message.

```
enum MessageID
{
    /* Defined by IDDE in the vdk.h file. */
};
```

The enumeration in `Vdk.h` will be empty. All the messages are dynamically allocated and have an ID of that type to allow the compiler to do type checking and prevent errors.

In C:

```
typedef enum MessageID VDK_MessageID;
```

In C/C++:

```
typedef enum MessageID VDK::MessageID;
```

MsgChannel

The [MsgChannel](#) type enumerates the channels a message can be posted or pended on.

In C:

```
enum VDK_MsgChannel
{
    VDK_kMsgWaitForAll = 1 << 15,
    VDK_kMsgChannel1   = 1 << 14,
    VDK_kMsgChannel2   = 1 << 13,
    VDK_kMsgChannel3   = 1 << 12,
    VDK_kMsgChannel4   = 1 << 11,
    VDK_kMsgChannel5   = 1 << 10,
    VDK_kMsgChannel6   = 1 << 9,
    VDK_kMsgChannel7   = 1 << 8,
    VDK_kMsgChannel8   = 1 << 7,
    VDK_kMsgChannel9   = 1 << 6,
    VDK_kMsgChannel10  = 1 << 5,
    VDK_kMsgChannel11  = 1 << 4,
    VDK_kMsgChannel12  = 1 << 3,
    VDK_kMsgChannel13  = 1 << 2,
    VDK_kMsgChannel14  = 1 << 1,
    VDK_kMsgChannel15  = 1 << 0
};
```

In C/C++:

```
enum VDK::MsgChannel
{
    VDK::kMsgWaitForAll = 1 << 15,
    VDK::kMsgChannel1   = 1 << 14,
    VDK::kMsgChannel2   = 1 << 13,
    VDK::kMsgChannel3   = 1 << 12,
    VDK::kMsgChannel4   = 1 << 11,
```

MsgChannel

```
    VDK::kMsgChannel5  = 1 << 10,  
    VDK::kMsgChannel6  = 1 << 9,  
    VDK::kMsgChannel7  = 1 << 8,  
    VDK::kMsgChannel8  = 1 << 7,  
    VDK::kMsgChannel9  = 1 << 6,  
    VDK::kMsgChannel10 = 1 << 5,  
    VDK::kMsgChannel11 = 1 << 4,  
    VDK::kMsgChannel12 = 1 << 3,  
    VDK::kMsgChannel13 = 1 << 2,  
    VDK::kMsgChannel14 = 1 << 1,  
    VDK::kMsgChannel15 = 1 << 0  
};
```

PoolID

The **PoolID** type is used to store the unique identifier of a memory pool.

```
enum PoolID
{
    /* Defined by IDDE in the vdk.h file. */
};
```

The enumeration in `Vdk.h` will only contain the IDs for the memory pools enabled at boot time. Any dynamically created memory pools will have an ID of the same type to allow the compiler to do type checking and prevent errors.

In C:

```
typedef enum PoolID VDK_PoolID;
```

In C/C++:

```
typedef enum PoolID VDK::PoolID;
```

Priority

The **Priority** type is used to denote the scheduling priority level of a thread:

- The highest priority is one (zero is reserved)
- The lowest priority is the size of a data word minus two.
For ADSP-219x DSPs—14 bits; for ADSP-21xxx, ADSP-TSxxx, or Blackfin processors—30 bits.

In C:

```
enum VDK_Priority
{
    VDK_kPriority1,
    VDK_kPriority2,
    VDK_kPriority3,
    ...
    VDK_kPriority14/30
};
```

In C++:

```
enum VDK::Priority
{
    VDK::kPriority1,
    VDK::kPriority2,
    VDK::kPriority3,
    ...
    VDK::kPriority14/30
};
```

SemaphoreID

The **SemaphoreID** type is used to store the unique identifier of a semaphore.

```
enum SemaphoreID
{
    /* Defined by IDDE in the vdk.h file. */
};
```

The enumeration in `Vdk.h` will only contain the IDs for the semaphores enabled at boot time. Any dynamically created semaphores will have an ID of the same type to allow the compiler to do type checking and prevent errors.

In C:

```
typedef enum SemaphoreID VDK_SemaphoreID;
```

In C/C++:

```
typedef enum SemaphoreID VDK::SemaphoreID;
```

SystemError

The [SystemError](#) type enumerates system defined errors thrown to the error handler.

In C:

```
enum VDK_SystemError
{
    VDK_kUnknownThreadType = INT_MIN,
    VDK_kUnknownThread,
    VDK_kThreadCreationFailure,
    VDK_kUnknownSemaphore,
    VDK_kUnknownEventBit,
    VDK_kUnknownEvent,
    VDK_kInvalidPriority,
    VDK_kInvalidDelay,
    VDK_kSemaphoreTimeout,
    VDK_kEventTimeout,
    VDK_kBlockInInvalidRegion,
    VDK_kDbgPossibleBlockInRegion,
    VDK_kInvalidPeriod,
    VDK_kAlreadyPeriodic,
    VDK_kNonperiodicSemaphore,
    VDK_kDbgPopUnderflow,
    VDK_kBadIOID,
    VDK_kBadDeviceDescriptor,
    VDK_kOpenFailure,
    VDK_kCloseFailure,
    VDK_kReadFailure,
    VDK_kWriteFailure,
    VDK_kIOCtlFailure,
    VDK_kInvalidDeviceFlag,
    VDK_kDeviceTimeout,
    VDK_kDeviceFlagCreationFailure,
```

```

VDK_kMaxCountExceeded,
VDK_kSemaphoreCreationFailure,
VDK_kSemaphoreDestructionFailure,
VDK_kPoolCreationFailure,
VDK_kInvalidBlockPointer,
VDK_kInvalidPoolParms,
VDK_kInvalidPoolID,
VDK_kErrorPoolNotEmpty,
VDK_kErrorMallocBlock,
VDK_kMessageCreationFailure,
VDK_kInvalidMessageID,
VDK_kInvalidMessageOwner,
VDK_kInvalidMessageChannel,
VDK_kMessageTimeout,
VDK_kMessageInQueue,
VDK_kNoError = 0,
VDK_kFirstUserError,
VDK_kLastUserError = INT_MAX
};

```

In C++:

```

enum VDK::SystemError
{
    VDK::kUnknownThreadType = INT_MIN,
    VDK::kUnknownThread,
    VDK::kThreadCreationFailure,
    VDK::kUnknownSemaphore,
    VDK::kUnknownEventBit,
    VDK::kUnknownEvent,
    VDK::kInvalidPriority,
    VDK::kInvalidDelay,
    VDK::kSemaphoreTimeout,
    VDK::kEventTimeout,
    VDK::kBlockInInvalidRegion,

```

SystemError

```
    VDK::kDbgPossibleBlockInRegion,  
    VDK::kInvalidPeriod,  
    VDK::kAlreadyPeriodic,  
    VDK::kNonperiodicSemaphore,  
    VDK::kDbgPopUnderflow,  
    VDK::kBadIOID,  
    VDK::kBadDeviceDescriptor,  
    VDK::kOpenFailure,  
    VDK::kCloseFailure,  
    VDK::kReadFailure,  
    VDK::kWriteFailure,  
    VDK::kIOCtlFailure,  
    VDK::kInvalidDeviceFlag,  
    VDK::kDeviceTimeout,  
    VDK::kDeviceFlagCreationFailure,  
    VDK::kMaxCountExceeded,  
    VDK::kSemaphoreCreationFailure,  
    VDK::kSemaphoreDestructionFailure,  
    VDK::kPoolCreationFailure,  
    VDK::kInvalidBlockPointer,  
    VDK::kInvalidPoolParms,  
    VDK::kInvalidPoolID,  
    VDK::kErrorPoolNotEmpty,  
    VDK::kErrorMallocBlock,  
    VDK::kMessageCreationFailure,  
    VDK::kInvalidMessageID,  
    VDK::kInvalidMessageOwner,  
    VDK::kInvalidMessageChannel,  
    VDK::kMessageTimeout,  
    VDK::kMessageInQueue,  
    VDK::kNoError = 0,  
    VDK::kFirstUserError,  
    VDK::kLastUserError = INT_MAX  
};
```

ThreadCreationBlock

A variable of the type [ThreadCreationBlock](#) is passed to the [CreateThreadEx\(\)](#) function.

In C:

```
typedef struct VDK_ThreadCreationBlock
{
    VDK_ThreadType    template_id;
    VDK_ThreadID      thread_id;
    unsigned int       thread_stack_size;
    VDK_Priority       thread_priority;
    void               *user_data_ptr;
} VDK_ThreadCreationBlock;
```

In C++:

```
typedef struct VDK::ThreadCreationBlock
{
    VDK::ThreadType    template_id;
    VDK::ThreadID      thread_id;
    unsigned int       thread_stack_size;
    VDK::Priority       thread_priority;
    void               *user_data_ptr;
} VDK::ThreadCreationBlock;
```

- `template_id` corresponds to a [ThreadType](#) defined in the `vdk.h` and `vdk.cpp` files. These files contain the default values for the stack size and initial priority, which may optionally be overridden by the following fields.
- `thread_id` is an output only field. On a successful return, it contains the same value as the function return.
- `thread_stack_size` overrides the default stack size implied by the [ThreadType](#) when it is non-zero.

ThreadCreationBlock

- `thread_priority` overrides the default thread priority implied by the [ThreadType](#) when it is non-zero.
- `user_data_ptr` allows a generic argument to be passed (without interpretation) to the thread creation function and, hence, to the thread constructor. This allows individual thread instances to be parameterized at creation time, without the need to resort to global variables for argument passing.

ThreadID

The **ThreadID** type is used to store the unique identifier of a thread.

```
enum ThreadID
{
    /* Defined by IDDE in the vdk.h file. */
};
```

The enumeration in `Vdk.h` will only contain the IDs for the threads enabled at boot time. Any dynamically created threads will have an ID of the same type to allow the compiler to do type checking and prevent errors.

In C:

```
typedef enum ThreadID VDK_ThreadID;
```

In C/C++:

```
typedef enum ThreadID VDK::ThreadID;
```

ThreadStatus

The [ThreadStatus](#) type is used to enumerate the state of a thread.

In C:

```
enum VDK_ThreadStatus
{
    VDK_kReady,
    VDK_kSemaphoreBlocked,
    VDK_kEventBlocked,
    VDK_kSemaphoreBlockedWithTimeout,
    VDK_kEventBlockedWithTimeout,
    VDK_kDeviceFlagBlocked,
    VDK_kDeviceFlagBlockedWithTimeout,
    VDK_kSleeping,
    VDK_MessageBlocked,
    VDK_kUnknown
};
```

In C++:

```
enum VDK::ThreadStatus
{
    VDK::kReady,
    VDK::kSemaphoreBlocked,
    VDK::kEventBlocked,
    VDK::kSemaphoreBlockedWithTimeout,
    VDK::kEventBlockedWithTimeout,
    VDK::kDeviceFlagBlocked,
    VDK::kDeviceFlagBlockedWithTimeout,
    VDK::kSleeping,
    VDK::MessageBlocked,
    VDK::kUnknown
};
```

ThreadType

The [ThreadType](#) is used to store the unique identifier of a Thread class.

```
enum ThreadType
{
    /* Defined by IDDE in the vdk.h file. */
};
```

In C:

```
typedef enum ThreadType VDK_ThreadType;
```

In C/C++:

```
typedef enum ThreadType VDK::ThreadType;
```

Ticks

Time is measured in system [Ticks](#). A tick is the amount of time between hardware interrupts generated by a hardware timer.

In C:

```
typedef unsigned int VDK_Ticks;
```

In C++:

```
typedef unsigned int VDK::Ticks;
```

VersionStruct

The constant [VersionStruct](#) is used to store four integers that describe the system parameters:

- VDK API version number
- DSP family supported
- base DSP product supported
- build number

In C:

```
typedef struct
{
    int                mAPIVersion;
    VDK_DSP_Family     mFamily;
    VDK_DSP_Product    mProduct;
    long               mBuildNumber;
} VDK_VersionStruct;
```

In C++:

```
typedef struct
{
    int                mAPIVersion;
    VDK::DSP_Family    mFamily;
    VDK::DSP_Product   mProduct;
    long               mBuildNumber;
} VDK::VersionStruct;
```

The [DSP_Family](#) and [DSP_Product](#) types are described [on page 4-9](#) and [on page 4-10](#), respectively.

5 VDK API REFERENCE

The VisualDSP++ Kernel Application Programming Interface is a library of functions and macros that may be called from your application programs. Application programs depend on API functions to perform services that are basic to the VDK. These services include interrupt handling, scheduler management, thread management, semaphore management, memory pool management, events and event bits, device drivers, and message passing.

All of the VDK functions are written in the C++ programming language. You can use the object files of the API library in DSP systems based on ADSP-219x, ADSP-21xxx, ADSP-TSxxx, and ADSP-BF53x processor architectures.

This chapter describes the current release of the API library. Future releases may include additional functions.

This chapter provides information on the following topics:

- [“Calling Library Functions” on page 5-2](#)
- [“Linking Library Functions” on page 5-2](#)
- [“Working With VDK Library Header” on page 5-3](#)
- [“Passing Function Parameters” on page 5-3](#)
- [“Library Naming Conventions” on page 5-3](#)
- [“API Summary” on page 5-4](#)

Calling Library Functions

To use an API function or a macro, call it by name and provide the appropriate arguments. The name and arguments for each library entity appear on its reference page. Note that the function names are C and C++ function names. If you call a C run-time library function from an assembly language program, prefix the function name with an underscore.

Similar to other functions, library functions should be declared. Declarations are supplied in the `vdk.h` header file. For more information about the kernel header file, see [“Working With VDK Library Header” on page 5-3](#).

The reference pages appear in the [“API Summary” on page 5-4](#).

Linking Library Functions

When your code calls an API function, the call creates a reference resolved by the linker when linking your program. One way to direct the linker to the library’s location is to use the default VDK Linker Description File (VDK-*<your_target>*.LDF). The default VDK Linker Description File automatically directs the linker to the `*.DLB` file in the `lib` subdirectory of your VisualDSP++ installation.

If you do not use the default VDK LDF file, add the library file to your project’s LDF. Alternatively, use the compiler’s `-l` (library directory) switch to specify the library to be added to the link line. Library functions are not linked into the `.DXX` unless they are referenced.

Working With VDK Library Header

If one of your program source files needs to call a VDK API library function, you include the `vdk.h` header file with the `#include` preprocessor command. The header file provides prototypes for all VDK public functions. The compiler uses prototypes to ensure each function is called with the correct arguments. The `vdk.h` file also provides declarations for user accessible global variables, macros, type definitions, and enumerations.

Passing Function Parameters

All parameters passed through the VDK library functions listed in “[API Summary](#)” on page 5-4 are either passed by value or as constant objects. This means the VDK does not modify any of the variables passed. Where arguments need to be modified, they are passed by address (pointer).

Library Naming Conventions

[Table 5-1](#) and [Table 5-2](#) show coding style conventions that apply to the entities in the library reference section. By following the library and function naming conventions, you can review VDK sources or documentation and recognize whether the identifier is a function, macro, variable parameter, or a constant.

Table 5-1. Library Naming Conventions

Notation	Description
VDK_Ticks	VDK defined types are written with the first letter uppercase.
kPriority1	Constants are prefixed with a “k”.
inType	Input parameters are prefixed with an “in”.
mDevice	Data members are prefixed with an “m”.

API Summary

Table 5-2. Function and Macro Naming Conventions

Notation	Description
VDK_	C callable function names are prefixed by “VDK_” to distinguish VDK library functions from user functions.
VDK::	C++ callable functions are located in the VDK namespace, thus function names are preceded by “VDK::”.
VDK_Yield(void)	The remaining portion of the function name is written with the first letter of each sub-word in uppercase.
VDK_ISR_SET_EVENTBIT_()	Assembly macros are written in uppercase with words separated by underscores and a trailing underscore.

API Summary

[Table 5-3](#) through [Table 5-18](#) list the VDK library entities included in the current software release. These tables list the library entities grouped by the service each grouping provides. The reference pages beginning [on page 5-15](#) appear in the alphabetic order.

Table 5-3. Interrupt Handling Functions

Function Name	Reference Page
PopCriticalRegion()	on page 5-94
PopNestedCriticalRegions()	on page 5-96
PushCriticalRegion()	on page 5-106

Table 5-4. Interrupt Mask Handling Functions

ClearInterruptMaskBits()	on page 5-17
--	------------------------------

Table 5-4. Interrupt Mask Handling Functions

GetInterruptMask()	on page 5-54
SetInterruptMaskBits()	on page 5-113

Table 5-5. Scheduler Management Functions

PopNestedUnscheduledRegions()	on page 5-98
PopUnscheduledRegion()	on page 5-99
PushUnscheduledRegion()	on page 5-107

Table 5-6. Block Memory Management Functions

CreatePool()	on page 5-24
CreatePoolEx()	on page 5-26
DestroyPool()	on page 5-37
FreeBlock()	on page 5-46
GetNumAllocatedBlocks()	on page 5-61
GetNumFreeBlocks()	on page 5-62
LocateAndFreeBlock()	on page 5-75
MallocBlock()	on page 5-79

Table 5-7. Thread and System Information Functions

GetThreadHandle()	on page 5-65
GetThreadID()	on page 5-66
GetThreadStackUsage()	on page 5-68
GetThreadStatus()	on page 5-69
GetThreadType()	on page 5-70
GetUptime()	on page 5-71

Table 5-7. Thread and System Information Functions (Cont'd)

GetVersion()	on page 5-72
LogHistoryEvent()	on page 5-76

Table 5-8. Thread Creation and Destruction Functions

CreateThread()	on page 5-30
CreateThreadEx()	on page 5-32
DestroyThread()	on page 5-41
FreeDestroyedThreads()	on page 5-48

Table 5-9. Thread Local Storage Functions

AllocateThreadSlot()	on page 5-11
AllocateThreadSlotEx()	on page 5-13
FreeThreadSlot()	on page 5-49
GetThreadSlotValue()	on page 5-67
SetThreadSlotValue()	on page 5-119

Table 5-10. Thread Error Management Functions

DispatchThreadError()	on page 5-45
ClearThreadError()	on page 5-18
GetLastThreadErrorValue()	on page 5-56
GetLastThreadError()	on page 5-55
SetThreadError()	on page 5-118

Table 5-11. Thread Priority Management Functions

GetPriority()	on page 5-63
-------------------------------	------------------------------

Table 5-11. Thread Priority Management Functions (Cont'd)

ResetPriority()	on page 5-110
SetPriority()	on page 5-116

Table 5-12. Thread Scheduling Control Functions

Sleep()	on page 5-120
Yield()	on page 5-126

Table 5-13. Semaphore Management Functions

CreateSemaphore()	on page 5-28
DestroySemaphore()	on page 5-39
GetSemaphoreValue()	on page 5-64
MakePeriodic()	on page 5-77
PendSemaphore()	on page 5-92
PostSemaphore()	on page 5-104
RemovePeriodic()	on page 5-108

Table 5-14. Event and EventBit Functions

ClearEventBit()	on page 5-15
GetEventBitValue()	on page 5-51
GetEventData()	on page 5-52
GetEventValue()	on page 5-53
LoadEvent()	on page 5-73
PendEvent()	on page 5-87
SetEventBit()	on page 5-111

API Summary

Table 5-15. Device Flags Functions

CreateDeviceFlag()	on page 5-21
DestroyDeviceFlag()	on page 5-34
PendDeviceFlag()	on page 5-85
PostDeviceFlag()	on page 5-101

Table 5-16. Device Driver Functions

CloseDevice()	on page 5-19
DeviceIOCtrl()	on page 5-43
OpenDevice()	on page 5-83
SyncRead()	on page 5-122
SyncWrite()	on page 5-124

Table 5-17. Message Functions

CreateMessage()	on page 5-22
DestroyMessage()	on page 5-35
GetMessagePayload()	on page 5-57
GetMessageReceiveInfo()	on page 5-59
MessageAvailable()	on page 5-81
PendMessage()	on page 5-89
PostMessage()	on page 5-102
SetMessagePayload()	on page 5-114

Table 5-18. Assembly Macros

Macro Name	Reference Page
VDK_ISR_ACTIVATE_DEVICE_()	on page 5-129
VDK_ISR_CLEAR_EVENTBIT_()	on page 5-130

Table 5-18. Assembly Macros (Cont'd)

Macro Name	Reference Page
VDK_ISR_LOG_HISTORY_EVENT_()	on page 5-131
VDK_ISR_POST_SEMAPHORE_()	on page 5-132
VDK_ISR_SET_EVENTBIT_()	on page 5-133

API Functions

The following format applies to all of the entries in the library reference section.

C Prototype

Provides the C prototype (as it is found in `vdk.h`) describing the interface to the function

C++ Prototype

Provides the C++ prototype (as it is found in `vdk.h`) describing the interface to the function

Description

Describes the function's operation

Parameters

Describes the function's parameters

Scheduling

Specifies whether the function invokes the scheduler

Determinism

Specifies whether the function is deterministic

Return Value

Describes the function's return value

Errors Thrown

Specifies errors detected by the VDK that can be dealt with by the thread's error-handling routines

AllocateThreadSlot()

C Prototype

```
bool VDK_AllocateThreadSlot( int *ioSlotNum );
```

C++ Prototype

```
bool VDK::AllocateThreadSlot( int *ioSlotNum );
```

Description

Assigns a new slot number if `*ioSlotNum = VDK::kTLSUnallocated` and enters the allocated `*ioSlotNum` into the global slot identifier table.

- Returns `FALSE` immediately if the value of `*ioSlotNum` is not equal to `VDK::kTLSUnallocated (INT_MIN)` to guard against multiple attempts to allocate the same key variable.
- Returns `FALSE` if there are no free slots, in which case `*ioSlotNum` is still `VDK::kTLSUnallocated`.
- Otherwise allocates the first available slot, places the slot number in `*ioSlotNum`, and returns `TRUE`.
- Does not access (change) any thread state.
- Guaranteed to return `TRUE` once only for a given key variable, so the return value may be used to control other one time library initialization.
- May be safely called during system initialization, i.e. before any threads are running.
- Equivalent to calling [AllocateThreadSlotEx\(\)](#) with a `NULL` cleanup function.

API Functions

Parameters

`ioSlotNum` is a pointer to a slot identifier.

Scheduling

Does not invoke the scheduler

Determinism

Not deterministic

Return Value

`TRUE` upon success and `FALSE` upon failure

Errors Thrown

None

AllocateThreadSlotEx()

C Prototype

```
bool VDK_AllocateThreadSlotEx(int *ioSlotNum,
                             void(*cleanupFn)(void*));
```

C++ Prototype

```
bool VDK::AllocateThreadSlotEx(int *ioSlotNum,
                              void(*cleanupFn)(void*));
```

Description

Assigns a new slot number if `*ioSlotNum = VDK::kTLSUnallocated` and enters the allocated `*ioSlotNum` into the global slot identifier table.

- Returns `FALSE` immediately if the value of `*ioSlotNum` is not equal to `VDK::kTLSUnallocated (INT_MIN)` to guard against multiple attempts to allocate the same key variable.
- Returns `FALSE` if there are no free slots, in which case `*ioSlotNum` is still `VDK::kTLSUnallocated`.
- Otherwise allocates the first available slot, places the slot number in `*ioSlotNum`, stores the `cleanupFn` pointer internally, and returns `TRUE`.
- Does not access (change) any thread state.
- Guaranteed to return `TRUE` once only for a given key variable, so the return value may be used to control other one time library initialization.
- May be safely called during system initialization, i.e. before any threads are running.

API Functions

Parameters

`ioSlotNum` is a pointer to a slot identifier.

`cleanupFn` is a pointer to a function to handle cleanup of thread-specific data in the event of thread destruction and:

- may be `NULL`, in which case it does nothing
- is called from within [DestroyThread\(\)](#)
- executes in the context of the calling thread, not the thread that is being destroyed
- is only called when the slot value is not `NULL`
- `free()` may be used as the cleanup function where the slot is used to hold ‘malloced’ data

Scheduling

Does not invoke the scheduler

Determinism

Not deterministic

Return Value

`TRUE` upon success and `FALSE` upon failure

Errors Thrown

None

ClearEventBit()

C Prototype

```
void VDK_ClearEventBit( VDK_EventBitID inEventBitID );
```

C++ Prototype

```
void VDK::ClearEventBit( VDK::EventBitID inEventBitID );
```

Description

Clears the value of the event bit – sets it to FALSE, NULL, or 0. Once the event bit is cleared, the value of each dependent event is recalculated. If several event bits are to be cleared (or set) as a single operation then the [SetEventBit\(\)](#) and/or [ClearEventBit\(\)](#) calls should be made from within an unscheduled region. Event recalculation will not occur until the unscheduled region is popped.

Parameters

`inEventBitID` is the system event bit to clear.

Scheduling

Invokes the scheduler and may result in a context switch

Determinism

Not deterministic

Return Value

No return value

API Functions

Errors Thrown

Full instrumentation and error checking libraries:

`kUnknownEventBit` indicates that `inEventBitID` is not a valid identifier.

No error checking libraries: None

ClearInterruptMaskBits()

C Prototype

```
void VDK_ClearInterruptMaskBits(VDK_IMASKStruct inMask);
```

C++ Prototype

```
void VDK::ClearInterruptMaskBits(VDK::IMASKStruct inMask);
```

Description

Clears bits in the interrupt mask. Any bits set in the parameter is cleared in the interrupt mask. In other words, the new mask is computed as the bitwise AND of the old mask and the one's complement of the `inMask` parameter.

Parameters

`inMask` specifies which bits should be cleared in the interrupt mask.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

No return value

Errors Thrown

None

API Functions

ClearThreadError()

C Prototype

```
void VDK_ClearThreadError(void);
```

C++ Prototype

```
void VDK::ClearThreadError(void);
```

Description

Sets the running thread's error status to `kNoError` and the error value to zero.

Parameters

None

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

No return value

Errors Thrown

None

CloseDevice()

C Prototype

```
void VDK_CloseDevice(VDK_DeviceDescriptor inDD);
```

C++ Prototype

```
void VDK::CloseDevice(VDK::DeviceDescriptor inDD);
```

Description

Closes the specified device. The function calls the dispatch function of the device opened with `inDD`.

Parameters

`inDD` is the [DeviceDescriptor](#) returned from the [OpenDevice\(\)](#) function.

Scheduling

Does not invoke the scheduler, but the user written device driver can call the scheduler

Determinism

Constant time. Note that this function calls user written device driver code which may not be deterministic.

Return Value

No return value

API Functions

Errors Thrown

Full instrumentation and error checking libraries:

`kBadDeviceDescriptor` indicates that `inDD` is not a valid [DeviceDescriptor](#).

No error checking libraries: None

CreateDeviceFlag()

C Prototype

```
VDK_DeviceFlagID VDK_CreateDeviceFlag(void);
```

C++ Prototype

```
VDK::DeviceFlagID VDK::CreateDeviceFlag(void);
```

Description

Creates a new device flag and returns its identifier.

Parameters

None

Scheduling

Does not invoke the scheduler

Determinism

Not deterministic

Return Value

New device flag identifier upon success and `UINT_MAX` upon failure.

Errors Thrown

Full instrumentation and error checking libraries:

`DeviceFlagCreationFailure` indicates that the kernel is not able to allocate and/or initialize memory for the device flag.

No error checking libraries: None

API Functions

CreateMessage()

C Prototype

```
VDK_MessageID VDK_CreateMessage(int          inPayloadType,  
                                unsigned int  inPayloadSize,  
                                void          *inPayloadAddr);
```

C++ Prototype

```
VDK::MessageID VDK::CreateMessage(int          inPayloadType,  
                                unsigned int  inPayloadSize,  
                                void          *inPayloadAddr);
```

Description

Creates and initializes a new message object. The return value is the identifier of the new message. The values passed to [CreateMessage\(\)](#) may be read by calling [GetMessagePayload\(\)](#) and may be reset by calling [SetMessagePayload\(\)](#). The calling thread becomes the owner of the new message.

Parameters

`inPayloadType` is a user defined value that may be used to convey additional information about the message and/or the payload to the receiving thread. This value is not used or modified by the kernel, except that negative values of payload type are reserved for use by VDK. Positive payload types are reserved for use by the application code. It is recommended that the payload address and size are always interpreted in the same way for each distinct message type.

`inPayloadSize` is the length of the payload buffer in the smallest addressable unit on the processor architecture (`sizeof(char)`). When `inPayloadSize` has a value of zero, the kernel assumes `inPayloadAddr` is not a pointer and may contain any user value of the same size.

`inPayloadAddr` is a pointer to the start of the data being passed in the message.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

New message identifier upon success and `UINT_MAX` upon failure.

Errors Thrown

Full instrumentation and error checking libraries:

- `kMaxCountExceeded` indicates that the number of simultaneous messages in the system exceeds the value specified in the GUI.
- `kMessageCreationFailure` indicates that the kernel is not able to allocate and/or initialize memory for the message.

Non error checking libraries: None

API Functions

CreatePool()

C Prototype

```
VDK_PoolID VDK_CreatePool(unsigned int   inBlockSz,  
                           unsigned int   inBlockCount,  
                           bool           inCreateNow);
```

C++ Prototype

```
VDK::PoolID VDK::CreatePool(unsigned int   inBlockSz,  
                             unsigned int   inBlockCount,  
                             bool           inCreateNow);
```

Description

Creates a new memory pool in the system heap and returns the pool identifier.

Parameters

`inBlockSz` specifies the block size in the lowest addressable unit.

`inBlockCount` specifies the total number of blocks in the pool.

`inCreateNow` indicates whether the block construction is done at runtime on an on demand basis (FALSE) or as a part of the creation process (TRUE).

Scheduling

Does not invoke the scheduler

Determinism

Constant time if `inCreateNow` is FALSE. Not deterministic if `inCreateNow` value is TRUE.

Return Value

New pool identifier upon success and `UINT_MAX` upon failure.

Errors Thrown

Full instrumentation and error checking libraries:

- `kInvalidPoolParms` indicates that either `inBlockSz` or `inBlockCount` is zero.
- `kPoolCreationFailure` indicates that the kernel is not able to allocate and/or initialize memory for the pool.

No error checking libraries: None

API Functions

CreatePoolEx()

C Prototype

```
VDK_PoolID VDK_CreatePoolEx(unsigned int   inBlockSz,  
                             unsigned int   inBlockCount,  
                             bool           inCreateNow,  
                             int           inWhichHeap);
```

C++ Prototype

```
VDK::PoolID VDK::CreatePoolEx(unsigned int   inBlockSz,  
                                unsigned int   inBlockCount,  
                                bool           inCreateNow,  
                                int           inWhichHeap);
```

Description

Creates a new memory pool in the specified heap and returns the pool identifier. When architectures do not support multiple heaps, `inWhichHeap` must be initialized to zero. Refer to [“Processor-Specific Notes” on page A-1](#) for architecture-specific information.

Parameters

`inBlockSz` specifies the block size in the lowest addressable units.

`inBlockCount` specifies the total number of blocks in the pool.

`inCreateNow` indicates whether block construction is done at runtime on an done on demand basis (FALSE) or as a part of the creation process (TRUE).

`inWhichHeap` specifies the heap in which the pool is to be created. This parameter is ignored on single heap architectures. Setting the value of `inWhichHeap` to zero specifies the default heap is to be used.

Scheduling

Does not invoke the scheduler

Determinism

Constant time if `inCreateNow` is `FALSE`. Not deterministic if `inCreateNow` value is `TRUE`.

Return Value

New pool identifier upon success and `UINT_MAX` upon failure

Errors Thrown

Full instrumentation and error checking libraries:

- `kInvalidPoolParms` indicates that either `inBlockSz` or `inBlockCount` is zero.
- `kPoolCreationFailure` indicates that the kernel is not able to allocate and/or initialize memory for the pool.

No error checking libraries: None

API Functions

CreateSemaphore()

C Prototype

```
VDK_SemaphoreID VDK_CreateSemaphore(  
                                unsigned int inInitialValue,  
                                unsigned int inMaxCount,  
                                VDK_Ticks inInitialDelay,  
                                VDK_Ticks inPeriod);
```

C++ Prototype

```
VDK::SemaphoreID VDK::CreateSemaphore(  
                                unsigned int inInitialValue,  
                                unsigned int inMaxCount,  
                                VDK::Ticks inInitialDelay,  
                                VDK::Ticks inPeriod);
```

Description

Creates and initializes a dynamic semaphore. If the value of `inPeriod` is non-zero, a periodic semaphore is created.

Parameters

`inInitialValue` is the value the semaphore will have once it is created. A value of zero indicates that the semaphore is unavailable. This value should be between zero and `inMaxCount`.

`inMaxCount` is the maximum number the semaphore's count can reach when posting it. An `inMaxCount` of one creates a binary semaphore, which is equivalent to the semaphores in the VisualDSP++ 2.0 release of VDK.

`inInitialDelay` is the number of ticks before the first posting of a periodic semaphore. `InInitialDelay` must be equal to or greater than one.

`inPeriod` specifies the period property of the semaphore and the number of ticks to sleep at each cycle after the semaphore is first posted. If `inPeriod` is zero, the created semaphore is not periodic.

Scheduling

Does not invoke the scheduler

Determinism

Not deterministic

Return Value

New semaphore identifier upon success and `UINT_MAX` upon failure.

Errors Thrown

Full instrumentation and error checking libraries:

- `kMaxCountExceeded` indicates that `inInitialValue` is greater than `inMaxCount`.
- `kSemaphoreCreationFailure` indicates that the kernel is not able to allocate and/or initialize memory for the semaphore.

No error checking libraries: None

API Functions

CreateThread()

C Prototype

```
VDK_ThreadID VDK_CreateThread(VDK_ThreadType inType);
```

C++ Prototype

```
VDK::ThreadID VDK::CreateThread(VDK::ThreadType inType);
```

Description

Creates a thread of the specified type and returns the new thread.

Parameters

`inType` corresponds to a thread type defined in the `vdk.h` and `vdk.cpp` files. These files contain the default values for the stack size, initial priority, and other properties.

Scheduling

Invokes the scheduler and may result in a context switch

Determinism

Not deterministic

Return Value

New thread identifier upon success and `UINT_MAX` upon failure.

Errors Thrown

Full instrumentation and error checking libraries:

- `kUnknownThreadType` indicates that `inType` is not an element of the `ThreadType` type, as defined in `vdk.h`.
- `kThreadCreationFailure` indicates that the kernel is not able to allocate and/or initialize memory for the thread.

No error checking libraries: None

API Functions

CreateThreadEx()

C Prototype

```
VDK_ThreadID VDK_CreateThreadEx(VDK_ThreadCreationBlock *inOutTCB);
```

C++ Prototype

```
VDK::ThreadID VDK::CreateThreadEx(VDK::ThreadCreationBlock *inOutTCB);
```

Description

Creates a thread with the specified characteristics and returns the new thread.

Parameters

`inOutTCB` is a pointer to a structure of the [ThreadCreationBlock](#) data type.

Scheduling

Invokes the scheduler and may result in a context switch

Determinism

Not deterministic

Return Value

New thread identifier upon success and `UINT_MAX` upon failure.

Errors Thrown

Full instrumentation and error checking libraries:

- `kUnknownThreadType` indicates that `inType` is not an element of the `ThreadType` type, as defined in `vdk.h`.
- `kThreadCreationFailure` indicates that the kernel is not able to allocate and/or initialize memory for the thread.

No error checking libraries: None

API Functions

DestroyDeviceFlag()

C Prototype

```
void VDK_DestroyDeviceFlag(VDK_DeviceFlagID inDeviceFlagID);
```

C++ Prototype

```
void VDK::DestroyDeviceFlag(VDK::DeviceFlagID inDeviceFlagID);
```

Description

Deletes the device flag from the system and releases the associated memory.

Parameters

`inDeviceFlagID` specifies the device flag to be destroyed.

Scheduling

Does not invoke the scheduler

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

`kInvalidDeviceFlag` indicates that `inDeviceFlagID` is not a valid identifier.

No error checking libraries: None

DestroyMessage()

C Prototype

```
void VDK_DestroyMessage(VDK_MessageID inMessageID);
```

C++ Prototype

```
void VDK::DestroyMessage(VDK::MessageID inMessageID);
```

Description

Destroys a message object. Only the thread that is the owner of a message can destroy it. The message payload memory is assumed to be already freed by the user thread. [DestroyMessage\(\)](#) does not free the payload and results in a memory leak if the memory is not freed.

Parameters

`inMessageID` is the identifier of the message to be destroyed.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

No return value

API Functions

Errors Thrown

Full instrumentation and error checking libraries:

- `kInvalidMessageID` indicates that the `inMessageID` is not a valid message identifier.
- `kInvalidMessageOwner` indicates that the thread attempting to destroy the message is not the current owner.
- `kMessageInQueue` indicates that the message has been posted to a thread (the [ThreadID](#) is not known at this point), and it needs to be removed from the message queue by a call to [PendMessage\(\)](#).

No error checking libraries: None

DestroyPool()

C Prototype

```
void VDK_DestroyPool(VDK_PoolID inPoolID);
```

C++ Prototype

```
void VDK::DestroyPool(VDK::PoolID inPoolID);
```

Description

Deletes the pool and cleans up the memory associated with it. If there are any allocated blocks, which are not yet freed, an error is thrown in the fully instrumented and error checking builds, and the pool will not be destroyed. In the non-error checking build, the pool will be destroyed.

Parameters

`inPoolID` specifies the pool to delete.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

No return value

API Functions

Errors Thrown

Full instrumentation and error checking libraries:

- `kInvalidPoolID` indicates that `inPoolID` is not valid.
- `kErrorPoolNotEmpty` indicates that the pool is not empty (there are some blocks that are not freed) and cannot be destroyed.

No error checking libraries: None

DestroySemaphore()

C Prototype

```
void VDK_DestroySemaphore(VDK_SemaphoreID inSemaphoreID);
```

C++ Prototype

```
void VDK::DestroySemaphore(VDK::SemaphoreID inSemaphoreID);
```

Description

Destroys the semaphore associated with `inSemaphoreID`. The destruction does not take place if there is any thread pending on the semaphore, resulting in an error thrown in full instrumentation and error checking builds.

Parameters

`inSemaphoreID` is the semaphore to destroy.

Scheduling

Does not invoke the scheduler

Determinism

Not deterministic

Return Value

No return value

API Functions

Errors Thrown

Full instrumentation and error checking libraries:

- `kSemaphoreDestructionFailure` indicates that the semaphore cannot be destroyed because there are threads pending on it.
- `kUnknownSemaphore` indicates that `inSemaphoreID` is not a valid identifier.

Non error checking libraries: None

DestroyThread()

C Prototype

```
void VDK_DestroyThread(VDK_ThreadID  inThreadID,
                      bool            inDestroyNow);
```

C++ Prototype

```
void VDK::DestroyThread(VDK::ThreadID  inThreadID,
                        bool            inDestroyNow);
```

Description

Initiates the process of removing the specified thread from the system. Although the scheduler never runs the thread again once this function completes, the kernel may optionally defer deallocation of the memory resources associated with the thread to the [Idle Thread](#). Any references to the destroyed thread are invalid and may throw an error. For more information about the low priority thread, see [“Idle Thread” on page 3-13](#).

Parameters

`inThreadID` specifies the thread to remove from the system.

`inDestroyNow` indicates whether the thread’s memory is to be recovered now (TRUE) or in the low priority IDLE thread (FALSE).

Scheduling

Invokes the scheduler and results in a context switch only if a thread passes itself to [DestroyThread\(\)](#).

Determinism

Constant time if `inDestroyNow` is FALSE; otherwise, not deterministic.

API Functions

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

`kUnknownThread` indicates that `inThreadID` is not a valid identifier.

Non error checking libraries: None

DeviceIOctl()

C Prototype

```
int VDK_DeviceIOctl(VDK_DeviceDescriptor inDD,
                   void *inCommand,
                   char *inParameters);
```

C++ Prototype

```
int VDK::DeviceIOctl(VDK::DeviceDescriptor inDD,
                   void *inCommand,
                   char *inParameters);
```

Description

Controls the specified device. The `inCommand` and `inParameters` are passed unchanged to the device driver.

Parameters

`inDD` is the [DeviceDescriptor](#) returned from [OpenDevice\(\)](#).

`inCommand` are the device driver's specific commands.

`inParameters` are the device driver's specific parameters for the above commands.

Scheduling

Does not invoke the scheduler, but the user written device driver can call the scheduler

Determinism

Constant time. Note that this function calls user written device driver code that may not be deterministic.

API Functions

Return Value

Return value of the dispatch function if the device exists and `UINT_MAX` if it does not.

Errors Thrown

Full instrumentation and error checking libraries:

`kBadDeviceDescriptor` indicates that `inDD` is not a valid identifier.

Non error checking libraries: None

DispatchThreadError()

C Prototype

```
int VDK_DispatchThreadError(VDK_SystemError inErr,  
                           const int inVal);
```

C++ Prototype

```
int VDK::DispatchThreadError(VDK::SystemError inErr,  
                             const int inVal);
```

Description

Sets the error and error's value in the currently running thread and calls the thread's error function.

Parameters

`inErr` is the error enumeration. See [“SystemError” on page 4-26](#) for more information about errors.

`inVal` is the value whose meaning determined by the error enumeration.

Scheduling

Does not invoke the scheduler, but the thread exception handler may do so.

Determinism

Not deterministic

Return Value

The current thread's error handler.

Errors Thrown

None

API Functions

FreeBlock()

C Prototype

```
void VDK_FreeBlock(VDK_PoolID inPoolID, void *inBlockPtr);
```

C Prototype

```
void VDK::FreeBlock(VDK::PoolID inPoolID, void *inBlockPtr);
```

Description

Frees the specified block and returns it to the free block list.

Parameters

`inPoolID` specifies the pool from which the block is to be freed.

`inBlockPtr` specifies the block to free.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

- `kInvalidBlockPointer` indicates that `inBlockPtr` is not a valid pointer from `inPoolID`.
- `kInvalidPoolID` indicates that `inPoolID` is not a valid identifier.

Non error checking libraries: None

API Functions

FreeDestroyedThreads()

C Prototype

```
void VDK_FreeDestroyedThreads(void);
```

C++ Prototype

```
void VDK::FreeDestroyedThreads(void);
```

Description

Frees the memory held by the destroyed threads whose resources have not been released by the `IDLE` thread. For more information, see [“Idle Thread” on page 3-13](#).

Parameters

None

Scheduling

Does not invoke the scheduler

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

None

FreeThreadSlot()

C Prototype

```
bool VDK_FreeThreadSlot(int inSlotNum);
```

C++ Prototype

```
bool VDK::FreeThreadSlot(int inSlotNum);
```

Description

Releases and clears the slot table entry in the currently running thread's slot table associated with `inSlotNum` and:

- Returns `FALSE` if (and only if) the key does not identify a currently allocated slot.
- Releases the slot identified by `inSlotNum`, which is previously created with the [AllocateThreadSlot\(\)](#) function.
- The application must ensure that no thread local data is associated with the key at the time it is freed, any specified cleanup functions (see [AllocateThreadSlotEx\(\)](#)) are only called on thread destruction.

Parameters

`inSlotNum` is the static library's preallocated slot number

Scheduling

Does not invoke the scheduler

Determinism

Constant Time

API Functions

Return Value

`TRUE` upon success and `FALSE` upon failure

Errors Thrown

None

GetEventBitValue()

C Prototype

```
bool VDK_GetEventBitValue(VDK_EventBitID inEventBitID);
```

C++ Prototype

```
bool VDK::GetEventBitValue(VDK::EventBitID inEventBitID);
```

Description

Returns the value of the event bit.

Parameters

`inEventBitID` specifies the system event bit to query.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

The value of either the specified event bit (if such a bit exists) or the current thread's error handler (if the specified event bit does not exist).

Errors Thrown

Full instrumentation and error checking libraries:

`kUnknownEventBit` indicates that `inEventBitID` is not a valid identifier.

Non error checking libraries: None

API Functions

GetEventData()

C Prototype

```
VDK_EventData VDK_GetEventData(VDK_EventID inEventID);
```

C++ Prototype

```
VDK::EventData VDK::GetEventData(VDK::EventID inEventID);
```

Description

Returns the `EventData` associated with the queried event. Threads can use this function to get an event's current values. For more information, see [“EventData” on page 4-14](#).

Parameters

`inEventID` is the event to query.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Data associated with the specified event bit if it exists, and a structure filled with zeros if it does not.

Errors Thrown

Full instrumentation and error checking libraries

`UnknownEvent` indicates that `inEventID` is not a valid event identifier.

Non error checking libraries: None

GetEventValue()

C Prototype

```
bool VDK_GetEventValue(VDK_EventID inEventID);
```

C++ Prototype

```
bool VDK::GetEventValue(VDK::EventID inEventID);
```

Description

Returns the value of the specified event.

Parameters

`inEventID` specifies the event to query.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Value of the specified event if it exists, and the return value of the current thread's error handler if it does not.

Errors Thrown

Full instrumentation and error checking libraries:

`kUnknownEventBit` indicates that `inEventBitID` is not a valid identifier.

Non error checking libraries: None

API Functions

GetInterruptMask()

C Prototype

```
VDK_IMASKStruct VDK_GetInterruptMask(void);
```

C++ Prototype

```
VDK::IMASKStruct VDK::GetInterruptMask(void);
```

Description

Returns the current value of the interrupt mask. The function is normally called before setting or clearing bits in the interrupt mask; should be called in an unscheduled region.

Parameters

None

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Current value of the interrupt mask.

Errors Thrown

None

GetLastThreadError()

C Prototype

```
VDK_SystemError VDK_GetLastThreadError(void);
```

C++ Prototype

```
VDK::SystemError VDK::GetLastThreadError(void);
```

Description

Returns the running thread's most recent error. See [“SystemError” on page 4-26](#) for more information about errors.

Parameters

None

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

The running thread's most recent error.

Errors Thrown

None

API Functions

GetLastThreadErrorValue()

C Prototype

```
int VDK_GetLastThreadErrorValue(void);
```

C++ Prototype

```
int VDK::GetLastThreadErrorValue(void);
```

Description

Returns the value parameter of the most recent error thrown by the running thread. See [“SystemError” on page 4-26](#) for more information about errors.

Parameters

None

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Value parameter of the most recent error thrown by the running thread.

Errors Thrown

None

GetMessagePayload()

C Prototype

```
void VDK_GetMessagePayload(VDK_MessageID inMessageID,
                           int            *outPayloadType,
                           unsigned int  *outPayloadSize,
                           void          **outPayloadAddr);
```

C++ Prototype

```
void VDK::GetMessagePayload(VDK::MessageID inMessageID,
                             int            *outPayloadType,
                             unsigned int  *outPayloadSize,
                             void          **outPayloadAddr);
```

Description

Returns the attributes associated with a message payload: type, size and address.

The meaning of these values is application-specific and corresponds to the arguments passed to [CreateMessage\(\)](#). Only the thread that is the owner of a message may examine the attributes of its payload. If other threads call this API, an error will be thrown, and the contents of `outPayloadType`, `outPayloadSize`, and `outPayloadAddress` will remain unchanged.

Parameters

`inMessageID` specifies the message to query.

`*outPayloadType` is an application-specific value that may be used to describe the contents of the payload. Negative values of payload type are reserved for use by VDK.

`*outPayloadSize` is typically the size of the payload in the smallest addressable units of the processor (`sizeof(char)`).

API Functions

`*outPayloadAddr` is typically a pointer to the beginning of the payload buffer. However, if the payload size has a value of zero, then the payload address may contain any user defined data.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

- `kInvalidMessageOwner` indicates that the argument `inMessageID` is not the current owner of the message.
- `kInvalidMessageID` indicates that the argument `inMessageID` is not a valid message identifier.
- `kMessageInQueue` indicates that the message has been posted to a thread (the [ThreadID](#) is not known at this point), and it needs to be removed from the message queue by a call to [PendMessage\(\)](#).

Non error checking libraries: None

GetMessageReceiveInfo()

C Prototype

```
void VDK_GetMessageReceiveInfo(VDK_MessageID    inMessageID,
                               VDK_MsgChannel   *outChannel,
                               VDK_ThreadID     *outSender);
```

C++ Prototype

```
void VDK::GetMessageReceiveInfo(VDK::MessageID    inMessageID,
                                VDK::MsgChannel   *outChannel,
                                VDK::ThreadID     *outSender);
```

Description

Returns the parameters associated with how a message was received.

Only the thread that is the owner of a message should call this API. If a different thread calls the API, there will be an error thrown and the `outChannel` and `outSender` variables will not contain the right information.

Parameters

`inMessageID` specifies message to be queried.

`*outChannel` identifies the channel of the recipient thread's message queue on which the message was posted.

`*outSender` identifies the `ThreadID` of the thread that posted the message.

Scheduling

Does not invoke the scheduler

API Functions

Determinism

Constant time

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

- `kInvalidMessageOwner` indicates that the argument `inMessageID` is not the current owner of the message.
- `kInvalidMessageID` indicates that the argument `inMessageID` is not a valid message identifier.
- `kMessageInQueue` indicates that the message has been posted to a thread (the [ThreadID](#) is not known at this point), and it needs to be removed from the message queue by a call to [PendMessage\(\)](#).

Non error checking libraries: None

GetNumAllocatedBlocks()

C Prototype

```
unsigned int VDK_GetNumAllocatedBlocks(VDK_PoolID inPoolID);
```

C++ Prototype

```
unsigned int VDK::GetNumAllocatedBlocks(VDK::PoolID inPoolID);
```

Description

Gets the number of allocated blocks in the pool.

Parameters

`inPoolID` specifies the pool.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Number of allocated blocks for the specified pool upon success and `UINT_MAX` otherwise.

Errors Thrown

Full instrumentation and error checking libraries:

`kInvalidPoolID` indicates that `inPoolID` is not a valid identifier.

Non error checking libraries: None

API Functions

GetNumFreeBlocks()

C Prototype

```
unsigned int VDK_GetNumFreeBlocks(VDK_PoolID inPoolID);
```

C ++ Prototype

```
unsigned int VDK::GetNumFreeBlocks(VDK::PoolID inPoolID);
```

Description

Gets the number of free blocks in the pool.

Parameters

`inPoolID` specifies the pool to query.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Number of free blocks for the specified pool upon success and `UINT_MAX` otherwise.

Errors Thrown

Full instrumentation and error checking libraries:

`kInvalidPoolID` indicates that `inPoolID` is not a valid identifier.

Non error checking libraries: None

GetPriority()

C Prototype

```
VDK_Priority VDK_GetPriority(VDK_ThreadID inThreadID);
```

C++ Prototype

```
VDK::Priority VDK::GetPriority(VDK::ThreadID inThreadID);
```

Description

Returns the priority of the specified thread.

Parameters

`inThreadID` is the thread whose priority is being queried.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Priority of the specified thread if it exists and `UINT_MAX` if it does not. See [“Priority” on page 4-24](#) for more information.

Errors Thrown

Full instrumentation and error checking libraries:

`kUnknownThread` indicates that `inThreadID` is not a valid identifier.

Non error checking libraries: None

API Functions

GetSemaphoreValue()

C Prototype

```
unsigned int VDK_GetSemaphoreValue(VDK_SemaphoreID inSemaphoreID);
```

C++ Prototype

```
unsigned int VDK::GetSemaphoreValue(VDK::SemaphoreID inSemaphoreID);
```

Description

Returns the value of the specified semaphore.

Parameters

`inSemaphoreID` is the semaphore to query.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Value of the specified semaphore if it exists and `UINT_MAX` if it does not.

Errors Thrown

Full instrumentation and error checking libraries:

`kUnknownSemaphore` indicates that `inSemaphoreID` is not a valid identifier.

Non error checking libraries: None

GetThreadHandle()

C Prototype

```
void** VDK_GetThreadHandle(void);
```

C++ Prototype

```
void** VDK::GetThreadHandle(void);
```

Description

Returns a pointer to a thread's user defined, allocated data pointer. This pointer can be used in C and assembly threads for holding thread local state (for example, member variables).

Parameters

None

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Pointer to a thread's user defined, allocated data pointer.

Errors Thrown

None

API Functions

GetThreadID()

C Prototype

```
VDK_ThreadID VDK_GetThreadID(void);
```

C++ Prototype

```
VDK::ThreadID VDK::GetThreadID(void);
```

Description

Returns the identifier of the currently running thread.

Parameters

None

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Identifier of the currently running thread.

Errors Thrown

None

GetThreadSlotValue()

C Prototype

```
void* VDK_GetThreadSlotValue(int inSlotNum);
```

C++ Prototype

```
void* VDK::GetThreadSlotValue(int inSlotNum);
```

Description

Returns the value in the currently running thread's slot table associated with `inSlotNum`. Returns `NULL` if the key does not identify a currently allocated slot, otherwise returns the current value held in the slot, which may also be `NULL`.

Parameters

`inSlotNum` is the static library's preallocated slot number.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Slot value for the slot number specified.

Errors Thrown

None

API Functions

GetThreadStackUsage()

C Prototype

```
unsigned int VDK_GetThreadStackUsage(VDK_ThreadID inThreadID);
```

C++ Prototype

```
unsigned int VDK::GetThreadStackUsage(VDK::ThreadID inThreadID);
```

Description

Gets the maximum used stack for the specified thread at the time of the call. This API measures the stack usage only in the fully instrumented builds. For any other builds it always returns zero.

Parameters

`inThreadID` specifies the thread whose stack usage we want to query.

Scheduling

Does not invoke the scheduler

Determinism

Not deterministic

Return Value

Maximum used stack till the time of this call or zero for libraries that are not fully instrumented.

Errors Thrown

None

GetThreadStatus()

C Prototype

```
VDK_ThreadStatus VDK_GetThreadStatus(const VDK_ThreadID inThreadID);
```

C++ Prototype

```
VDK::ThreadStatus VDK::GetThreadStatus(const VDK::ThreadID inThreadID);
```

Description

Reports the enumerated status of the specified thread.

Parameters

`inThreadID` is the thread whose status is being queried.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Status of the specified thread if it exists and `VDK::kUnknown` if it does not.
For more information see [“ThreadStatus” on page 4-32](#).

Errors Thrown

None

API Functions

GetThreadType()

C Prototype

```
VDK_ThreadType VDK_GetThreadType(const VDK_ThreadID inThreadID);
```

C++ Prototype

```
VDK::ThreadType VDK::GetThreadType(const VDK::ThreadID inThreadID);
```

Description

Returns the thread type used to create the thread. The thread type is defined in the `vdk.h` and `vdk.cpp` files. [For more information, see “Threads” on page 3-1.](#)

Parameters

`inThreadID` specifies the thread to query.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Type of the specified thread if it exists and `UINT_MAX` if it does not. For more information, see [“ThreadType” on page 4-33.](#)

Errors Thrown

Full instrumentation and error checking libraries:

`UnknowThread` indicates that `inThreadID` is not a valid identifier.

Non error checking libraries: None

GetUptime()

C Prototype

```
VDK_Ticks VDK_GetUptime(void);
```

C++ Prototype

```
VDK::Ticks VDK::GetUptime(void);
```

Description

Returns the time in ticks since the last system reset.

Parameters

None

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Time in ticks since the last system reset.

Errors Thrown

None

API Functions

GetVersion()

C Prototype

```
VDK_VersionStruct VDK_GetVersion(void);
```

C++ Prototype

```
VDK::VersionStruct VDK::GetVersion(void);
```

Description

Returns the current version of VDK, [VersionStruct](#), described on [page 4-35](#).

Parameters

None

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

Current version of VDK.

Errors Thrown

None

LoadEvent()

C Prototype

```
void VDK_LoadEvent(VDK_EventID      inEventID,
                  const VDK_EventData inEventData);
```

C++ Prototype

```
void VDK::LoadEvent(VDK::EventID      inEventID,
                   const VDK::EventData inEventData);
```

Description

Loads the `EventData` associated with the event. For more information, see [“EventData” on page 4-14](#).

Parameters

`inEventID` is the event to be reinitialized.

`inEventData` contains the new values for the event.

Scheduling

Causes the value of the event to be recalculated, invokes the scheduler, and may result in a context switch

Determinism

Not deterministic

Return Value

No return value

API Functions

Errors Thrown

Full instrumentation and error checking libraries:

`kUnknownEvent` indicates that `inEventID` is not a valid identifier.

Non error checking libraries: None

LocateAndFreeBlock()

C Prototype

```
void VDK_LocateAndFreeBlock(void *inBlkPtr);
```

C++ Prototype

```
void VDK::LocateAndFreeBlock(void *inBlkPtr);
```

Description

Determines in which pool the block to be freed resides, frees the block, and returns it to the free block list.

Parameters

`inBlockPtr` specifies the block to be freed.

Scheduling

Does not invoke the scheduler

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

`kInvalidBlockPointer` indicates that `inBlkPtr` does not belong to any of the active memory pools and cannot be freed.

Non error checking libraries: None

API Functions

LogHistoryEvent()

C Prototype

```
void VDK_LogHistoryEvent(VDK_HistoryEnum inEnum, int inValue);
```

C++ Prototype

```
void VDK::LogHistoryEvent(VDK::HistoryEnum inEnum, int inValue);
```

Description

Adds a record to the history buffer. The function does not perform any action if the project is not linked with the fully instrumented libraries.

Parameters

`inEnum` is the enumeration value for this type of event. For more information see [“HistoryEnum” on page 4-15](#).

`inValue` is the value defined by enumeration.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

No return value

Errors Thrown

None

MakePeriodic()

C Prototype

```
void VDK_MakePeriodic(VDK_SemaphoreID  inSemaphoreID,
                     VDK_Ticks         inDelay,
                     VDK_Ticks         inPeriod);
```

C++ Prototype

```
void VDK::MakePeriodic(VDK::SemaphoreID  inSemaphoreID,
                       VDK::Ticks         inDelay,
                       VDK::Ticks         inPeriod);
```

Description

Directs the scheduler to post the specified semaphore after `inDelay` number of ticks. Thereafter, every `inPeriod` ticks, the semaphore is posted and the scheduler is invoked. This allows the running thread to acquire the signal and, if the thread is at the highest priority level, to continue execution.

To be periodic, the running thread must repeat in sequence: perform task and then pend on the semaphore. Note that this differs from sleeping at the completion of activity.

Parameters

`inSemaphoreID` is the semaphore to make periodic.

`inDelay` is the number of ticks before the first posting of the semaphore. `inDelay` must be equal to or greater than one.

`inPeriod` is the number of ticks to sleep at each cycle after the first cycle. `inPeriod` must be equal to or greater than one.

API Functions

Scheduling

Does not invoke the scheduler

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

- `kUnknownSemaphore` indicates that `inSemaphoreID` is not a valid identifier.
- `kInvalidPeriod` indicates that `inPeriod` is zero.
- `kInvalidDelay` indicates that `inDelay` is zero. This is an error because the first posting will not occur until the next tick.
- `kAlreadyPeriodic` indicates that the semaphore is already periodic and can not be made periodic again. If the intention is to change the period, the semaphore has to be made non-periodic first.

Non error checking libraries: None

MallocBlock()

C Prototype

```
void* VDK_MallocBlock(VDK_PoolID inPoolID);
```

C++ Prototype

```
void* VDK::MallocBlock(VDK::PoolID inPoolID);
```

Description

Returns pointer to the next available block from the specified pool.

Parameters:

`inPoolID` specifies the pool from which block is to be allocated.

Scheduling:

Does not invoke the scheduler

Determinism

Constant time if `inCreateNow` was specified to be `true` when the specified pool was created.

Not deterministic if `inCreateNow` was specified to be `FALSE`.

Return Value

Void pointer to a free memory block upon success. `NULL` if the call fails to allocate a block.

API Functions

Errors Thrown

Full instrumentation and error checking libraries:

- `kInvalidPoolID` indicates that `inPoolID` is not a valid identifier.
- `kErrorMallocBlock` indicates that there are no free blocks in the pool, so a new block cannot be allocated.

Non error checking libraries: None

MessageAvailable()

C Prototype

```
bool VDK_MessageAvailable(unsigned int inMessageChannelMask);
```

C++ Prototype

```
bool VDK::MessageAvailable(unsigned int inMessageChannelMask);
```

Description

Enables a thread to use a polling model (rather than a blocking model) to wait for messages in its message queue. This function returns `TRUE` if a subsequent call to [PendMessage\(\)](#) with the same channel mask will not block.

Parameters

`inMessageChannelMask` specifies the receive channels. A set bit corresponds to a receive channel, and a clear bit corresponds to a channel ignored.

If the `VDK::kMsgWaitForAll` flag is set in the channel mask then the query operates with AND logic, rather than the default OR logic. By default, only one message—on any of the receive channels designated in the channel mask—is required for a true result. The `VDK::kMsgWaitForAll` flag requires at least one message to be queued on each of the specified receive channels channel in order for the function to return true.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

API Functions

Return Value

TRUE is there is a message available and FALSE if there is not.

Errors Thrown

Full instrumentation and error checking libraries:

`kInvalidMessageChannel` indicates that `inMessageChannelMask` is not a valid mask.

Non error checking libraries: None

OpenDevice()

C Prototype

```
VDK_DeviceDescriptor VDK_OpenDevice(VDK_IOID inIDNum,
                                     char      *inFlags);
```

C++ Prototype

```
VDK::DeviceDescriptor VDK::OpenDevice(VDK::IOID inIDNum,
                                       char      *inFlags);
```

Description

Opens the specified device.

Parameters

`inIDNum` is the boot IO identifier.

`inFlags` is uninterpreted data passed through to the device being opened.

Scheduling

Does not call the scheduler, but the user written device driver can call the scheduler.

Determinism

Constant time. Note that this function calls user written device driver code that may not be deterministic.

Return Value

Return value of the dispatch function if the device exists and `UINT_MAX` if it does not.

API Functions

Errors Thrown

Full instrumentation and error checking libraries:

- `kBadIOID` indicates that `inIDNum` is not a valid [IOID](#).
- `kOpenFailure` indicates that no more devices can be open simultaneously.

Non error checking libraries: None

PendDeviceFlag()

C Prototype

```
void VDK_PendDeviceFlag(VDK_DeviceFlagID  inFlagID,
                        VDK_Ticks          inTimeout);
```

C++ Prototype

```
void VDK::PendDeviceFlag(VDK::DeviceFlagID  inFlagID,
                          VDK::Ticks          inTimeout);
```

Description

Allows a thread to block on the specified device flag. The thread is blocked and swapped out. Once the device flag is made available via [PostDeviceFlag\(\)](#), *all* threads waiting for this flag are made ready-to-run. If the thread does not resume execution within `inTimeout` ticks, the thread's reentry point is changed to its error function, and the thread is made available for scheduling. If the value of `inTimeout` is passed as zero, then the thread may pend indefinitely.

Note that [PendDeviceFlag\(\)](#) must be called from within a non-nested critical region (a critical region with a stack depth of one), but from outside of any unscheduled regions. On return from [PendDeviceFlag\(\)](#), the critical region is popped.

Parameters

`inFlagID` is the device flag on which the thread pends.

`inTimeout` specifies the maximum duration in ticks for which the thread pends on the device flag.

Scheduling

Invokes the scheduler.

API Functions

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

- `kBlockInInvalidRegion` indicates that `PendDeviceFlag()` is called in an unscheduled region or nested critical region (must be called in a non-nested critical region).
- `kInvalidDeviceFlag` indicates that `inFlagID` is not a valid identifier.
- `kDeviceTimeout` indicates that the timeout value has expired before the device flag was posted.

Non error checking libraries: `kDeviceTimeout` as above.

PendEvent()

C Prototype

```
void VDK_PendEvent(VDK_EventID inEventID, VDK_Ticks inTimeout);
```

C++ Prototype

```
void VDK::PendEvent(VDK::EventID inEventID, VDK::Ticks inTimeout);
```

Description

Provides the mechanism by which threads pend on events. If the named event calculates as being available, execution returns to the running thread. If the event is *not* available, the thread pauses execution until the event is available. When the event becomes available, *all* threads pending on the event are moved to the ready queue.

If the thread does not resume execution within `inTimeout` ticks, the thread's reentry point is changed to its error function, and the thread is made available for scheduling. If the value of `inTimeout` is passed as zero, then the thread may pend indefinitely.

Parameters

`inEventID` is the event on which the thread pends.

`inTimeout` specifies the maximum duration in ticks for which the thread pends on the event.

Scheduling

Invokes the scheduler and may result in a context switch.

Determinism

Constant time if event is available

API Functions

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

- `kEventTimeout` indicates that the timeout value has expired before the event was available.
- `kUnknownEvent` indicates that `inEventID` is not a valid identifier.
- `kBlockInInvalidRegion` indicates that `PendEvent()` is trying to block in an unscheduled region, causing a scheduling conflict.
- `kDbgPossibleBlockInRegion` indicates that `PendEvent()` is being called in an unscheduled region, causing a potential scheduling conflict.

Non error checking libraries: `kEventTimeout` as above.

PendMessage()

C Prototype

```
VDK_MessageID VDK_PendMessage(unsigned int inMessageChannelMask,
                               VDK_Ticks    inTimeout);
```

C++ Prototype

```
VDK::MessageID VDK::PendMessage(unsigned int inMessageChannelMask,
                                  VDK::Ticks    inTimeout);
```

Description

Retrieves a message from a thread's message queue. `PendMessage` is a blocking call: when the specified conditions for a valid message in the queue are not met, the thread suspends execution. The channel mask allows you to specify which channels (`kMsgChannel1` through `kMsgChannel15`) to examine for incoming messages.

In addition, the flag `VDK::kMsgWaitForAll` may be included in (OR-ed into) the channel mask to specify that at least one message must be present on each of the channels specified in the mask. Messages are retrieved from the lowest numbered channels first (`kMsgChannel1`, then `kMsgChannel2`, ...). Once a `MessageID` is returned by `PendMessage`, the message is no longer in the queue and is owned by the calling thread. If the thread does not resume execution within `inTimeout` ticks, the thread's reentry point is changed to its error function, and the thread is made available for scheduling. If the value for `inTimeout` is passed as zero, then the thread may pend indefinitely.

Parameters

`inMessageChannelMask` specifies the receive channels. A set bit corresponds to a receive channel. A clear bit corresponds to a channel that is ignored. The parameter may not be zero.

API Functions

If the `VDK::kMsgWaitForAll` flag is set in the channel mask then the `pend` operates with AND logic, rather than the default OR logic. By default, only one message —on any of the receive channels designated in the channel mask—is required to unblock the pending thread. The `VDK::kMsgWaitForAll` flag requires at least one message to be queued on each of the specified receive channels channel in order to unblock.

`inTimeout` specifies the maximum duration in ticks for which the thread pends on the receipt of the required message(s).

Scheduling

Invokes the scheduler and may result in a context switch.

Determinism

Constant time if there is no need to block.

Return Value

Identifier of the message the thread pended on upon success; `UINT_MAX` otherwise.

Errors Thrown

Full instrumentation and error checking libraries:

- `kDbgPossibleBlockInRegion` indicates that [PendMessage\(\)](#) is being called in an unscheduled region, causing a potential scheduling conflict.
- `kInvalidMessageChannel` indicates that `inMessageChannelMask` does not specify a correct group of channels to mask.

- `kMessageTimeout` indicates that the timeout value has expired before the thread removed the message from its message queue.
- `kBlockInInvalidRegion` indicates that `PendMessage()` is trying to block in an unscheduled region, causing a scheduling conflict.

Non error checking libraries: `kMessageTimeout` as above.

API Functions

PendSemaphore()

C Prototype

```
void VDK_PendSemaphore(VDK_SemaphoreID    inSemaphoreID,  
                      VDK_Ticks          inTimeout);;
```

C++ Prototype

```
void VDK::PendSemaphore(VDK::SemaphoreID  inSemaphoreID,  
                       VDK::Ticks         inTimeout);
```

Description

Provides the mechanism that allows threads to pend on semaphores. If the named semaphore is available (its count is greater than zero), the semaphore's count is decremented by one, and processor control returns to the running thread. If the semaphore is *not* available (its count is zero), the thread pauses execution until the semaphore is posted. If the thread does not resume execution within `inTimeout` ticks, the thread's reentry point is changed to its error function, and the thread is made available for scheduling. If the value of `inTimeout` is passed as zero, then the thread may pend indefinitely.

Parameters

`inSemaphoreID` is the semaphore on which the thread pends.

`inTimeout` specifies the maximum duration in ticks for which the thread pends on the semaphore.

Scheduling

Invokes the scheduler and may result in a context switch.

Determinism

Constant time if semaphore is available

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

- `kSemaphoreTimeout` indicates that the timeout value has expired before the semaphore became available
- `kUnknownSemaphore` indicates that `inSemaphoreID` is not a valid identifier.
- `kBlockInInvalidRegion` indicates that `PendSemaphore()` is called in an unscheduled region, causing a scheduling conflict.
- `kDbgPossibleBlockInRegion` indicates that `PendSemaphore()` may be called in an unscheduled region, causing a potential scheduling conflict.

Non error checking libraries: `kSemaphoreTimeout` as above.

API Functions

PopCriticalRegion()

C Prototype

```
void VDK_PopCriticalRegion(void);
```

C++ Prototype

```
void VDK::PopCriticalRegion(void);
```

Description

Decrements the count of nested critical regions. Use it as a close bracket call to [PushCriticalRegion\(\)](#). A count is maintained to ensure that each entered critical region calls [PopCriticalRegion\(\)](#) before interrupts are reenabled. The kernel ignores additional calls to [PopCriticalRegion\(\)](#) while interrupts are enabled.

Each critical region is also (implicitly) an unscheduled region.

Parameters

None

Scheduling

Invokes the scheduler and may result in a context switch if interrupts are reenabled by this call.

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

`kDbgPopUnderflow` indicates that there were no critical regions to pop.

Non error checking libraries: None

API Functions

PopNestedCriticalRegions()

C Prototype

```
void VDK_PopNestedCriticalRegions(void);
```

C++ Prototype

```
void VDK::PopNestedCriticalRegions(void);
```

Description

Resets the count of nested critical regions to zero, thereby, re-enabling interrupts. The kernel ignores additional calls to [PopNestedCriticalRegions\(\)](#) while interrupts are enabled.

This function does not change the interrupt mask.

Parameters

None

Scheduling

Invokes the scheduler and may result in a context switch (unless interrupts are already enabled).

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

`kDbgPopUnderflow` indicates that there are no critical regions to pop.

Non error checking libraries: None

API Functions

PopNestedUnscheduledRegions()

C Prototype

```
void VDK_PopNestedUnscheduledRegions(void);
```

C++ Prototype

```
void VDK::PopNestedUnscheduledRegions(void);
```

Description

Resets the count of nested unscheduled regions to zero, thereby, re-enabling scheduling. The kernel ignores additional calls to [PopNestedUnscheduledRegions\(\)](#) while scheduling is enabled.

Parameters

None

Scheduling

Invokes the scheduler and may result in a context switch (unless scheduling is already enabled).

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

`kDbgPopUnderflow` indicates that there were no critical regions to pop.

Non error checking libraries: None

PopUnscheduledRegion()

C Prototype

```
void VDK_PopUnscheduledRegion(void);
```

C++ Prototype

```
void VDK::PopUnscheduledRegion(void);
```

Description

Decrements the count of nested unscheduled regions. Use it as a close bracket call to [PushUnscheduledRegion\(\)](#). A nesting count is maintained to ensure that each entered unscheduled region calls [PopUnscheduledRegion\(\)](#) before scheduling is resumed.

The kernel ignores additional calls to [PopUnscheduledRegion\(\)](#) while scheduling is enabled.

Parameters

None

Scheduling

Invokes the scheduler and may result in a context switch if scheduling is reenabled by this call.

Determinism

Not deterministic

Return Value

No return value

API Functions

Errors Thrown

Full instrumentation and error checking libraries:

`kDbgPopUnderflow` indicates that there are no critical regions to pop.

Non error checking libraries: None

PostDeviceFlag()

C Prototype

```
void VDK_PostDeviceFlag(VDK_DeviceFlagID inFlagID);
```

C++ Prototype

```
void VDK::PostDeviceFlag(VDK::DeviceFlagID inFlagID);
```

Description

Posts the specified device flag. Once the device flag is made available, *all* threads waiting for the flag are in the ready-to-run state.

Parameters

`inFlagID` is the [DeviceDescriptor](#) returned from [OpenDevice\(\)](#).

Scheduling

Invokes the scheduler.

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

None

API Functions

PostMessage()

C Prototype

```
void VDK_PostMessage(VDK_ThreadID      inRecipient,  
                    VDK_MessageID     inMessageID,  
                    VDK_MsgChannel     inChannel);
```

C++ Prototype

```
void VDK::PostMessage(VDK::ThreadID    inRecipient,  
                     VDK::MessageID    inMessageID,  
                     VDK::MsgChannel    inChannel);
```

Description

Appends the message `inMessageID` to the message queue of the thread with identifier `inRecipient` on the channel `inChannel`. [PostMessage\(\)](#) is a non-blocking function: returns execution to the calling thread without waiting for the recipient to run or to acknowledge the new message in its queue. The message is considered delivered when [PostMessage\(\)](#) returns. Only the thread that is the owner of a message may post it.

At delivery time, ownership of the message and the associated payload is transferred from the sending thread to the recipient thread. Once delivered, all memory references to the payload, which may be held by the sending thread, are invalid. Memory read and write privileges and the responsibility for freeing the payload memory are passed to the recipient thread along with ownership.

Parameters

`inRecipient` is the `ThreadID` of the thread to receive the message.

`inMessageID` is the [MessageID](#) of the message being sent. A message must be created before it is posted. This parameter is a return value of the call to [CreateMessage\(\)](#).

`inChannel` is the FIFO within the recipient's message queue on which the message is appended. Its value is `kMsgChannel1` through `kMessageChannel15`.

Scheduling

Invokes the scheduler and may result in a context switch.

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

- `kInvalidMessageChannel` indicates that the `inChannel` is not a channel value.
- `kUnknownThread` indicates that `inRecipient` is not a valid thread identifier.
- `kInvalidMessageID` indicates that `inMessageID` is not a valid message identifier.
- `kInvalidMessageOwner` indicates that the thread attempting to post the message is not the current owner. The error value is the [ThreadID](#) of the owner.
- `kMessageInQueue` indicates that the message has been posted to a thread (the [ThreadID](#) is not known at this point), and it needs to be removed from the message queue by a call to [PendMessage\(\)](#).

Non error checking libraries: None

API Functions

PostSemaphore()

C Prototype

```
void VDK_PostSemaphore(VDK_SemaphoreID inSemaphoreID);
```

C++ Prototype

```
void VDK::PostSemaphore(VDK::SemaphoreID inSemaphoreID);
```

Description

Provides the mechanism by which threads post semaphores. Every time a semaphore is posted, its count increases by one until it reaches the maximum value, as specified on creation. Any further posts have no effect. Note that Interrupt Service Routines must use a different interface, as described in [“VDK_ISR_POST_SEMAPHORE_\(\)” on page 5-132](#).

Parameters

`inSemaphoreID` is the semaphore to post.

Scheduling

May invoke the scheduler and may result in a context switch.

Determinism

- No thread pending: Constant time
- Low priority thread pending: Constant time
- High priority thread pending: Constant time plus a context switch

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

`kUnknownSemaphore` indicates that `inSemaphoreID` is not a valid identifier.

Non error checking libraries: None

API Functions

PushCriticalRegion()

C Prototype

```
void VDK_PushCriticalRegion(void);
```

C++ Prototype

```
void VDK::PushCriticalRegion(void);
```

Description

Disables interrupts to enable atomic execution of a critical region of code. Note that critical regions may be nested. A count is maintained to ensure a coequal number of calls to [PopCriticalRegion\(\)](#) are made before restoring interrupts. Each critical region is also (implicitly) an unscheduled region.

Parameters

None

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

No return value

Errors Thrown

None

PushUnscheduledRegion()

C Prototype

```
void VDK_PushUnscheduledRegion(void);
```

C++ Prototype

```
void VDK::PushUnscheduledRegion(void);
```

Description

Disables the scheduler. While in an unscheduled region, the current thread will not be de-scheduled, even if a higher-priority thread becomes ready to run. Note that unscheduled regions may be nested. A count is maintained to ensure a coequal number of calls to [PopUnscheduledRegion\(\)](#) are made before scheduling is reenabled.

Scheduling

Suspends scheduling until a matching [PopUnscheduledRegion\(\)](#) call.

Parameters

None

Determinism

Constant time

Return Value

No return value

Errors Thrown

None

API Functions

RemovePeriodic()

C Prototype

```
void VDK_RemovePeriodic(VDK_SemaphoreID inSemaphoreID);
```

C++ Prototype

```
void VDK::RemovePeriodic(VDK::SemaphoreID inSemaphoreID);
```

Description

Stops periodic posting of the specified semaphore. Trying to stop a non-periodic semaphore has no effect and raises a run-time error.

Parameters

`inSemaphoreID` is the semaphore for which periodic posting is to be halted.

Scheduling

Does not invoke the scheduler

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

- `kUnknownSemaphore` indicates that `inSemaphoreID` is not a valid identifier.
- `kNonperiodicSemaphore` indicates that the semaphore is not periodic.

Non error checking libraries: None

API Functions

ResetPriority()

C Prototype

```
void VDK_ResetPriority(const VDK_ThreadID inThreadID);
```

C++ Prototype

```
void VDK::ResetPriority(const VDK::ThreadID inThreadID);
```

Description

Restores the priority of the named thread to the default value specified in the thread's template.

Parameters

`inThreadID` is the thread whose priority to be reset.

Scheduling

Invokes the scheduler and may result in a context switch

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

`kUnknownThread` indicates that `inThreadID` is not a valid identifier.

Non error checking libraries: None

SetEventBit()

C Prototype

```
void VDK_SetEventBit(VDK_EventBitID inEventBitID);
```

C++ Prototype

```
void VDK::SetEventBit(VDK::EventBitID inEventBitID);
```

Description

Sets the value of the event bit. Once the event bit is set (meaning its value equals to `TRUE`, 1, occurred, etc.), the value of all dependent events is recalculated.

If several event bits are to be set (or cleared) as a single operation, then the [SetEventBit\(\)](#) and/or [ClearEventBit\(\)](#) calls should be made from within an unscheduled region. Event recalculation will not occur until the unscheduled region is popped.

Parameters

`inEventBitID` is the system event bit to set.

Scheduling

Invokes the scheduler and may result in a context switch

Determinism

Not deterministic

Return Value

No return value

API Functions

Errors Thrown

Full instrumentation and error checking libraries:

`kUnknownEventBit` indicates that `inEventBitID` is not a valid identifier.

Non error checking libraries: None

SetInterruptMaskBits()

C Prototype

```
void VDK_SetInterruptMaskBits(VDK_IMASKStruct inMask);
```

C++ Prototype

```
void VDK::SetInterruptMaskBits(VDK::IMASKStruct inMask);
```

Description

Sets bits in the interrupt mask. Any bits set in the parameter is set in the interrupt mask. In other words, the new mask is computed as the bitwise OR of the old mask and the parameter `inMask`.

Parameters

`inMask` specifies which bits should be set in the interrupt mask.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

No return value

Errors Thrown

None

API Functions

SetMessagePayload()

C Prototype

```
void VDK_SetMessagePayload(VDK_MessageID    inMessageID,  
                           int              inPayloadType,  
                           unsigned int     inPayloadSize,  
                           void             *inPayloadAddr);
```

C++ Prototype

```
void VDK::SetMessagePayload(VDK::MessageID inMessageID,  
                           int             inPayloadType,  
                           unsigned int    inPayloadSize,  
                           void            *inPayloadAddr);
```

Description

Sets the values in a message header that describe the payload. The meaning of these values is application-specific and corresponds to the arguments passed to [CreateMessage\(\)](#). This function overwrites the existing values in the message. Only the thread that is the owner of a message may set the attributes of its payload.

Parameters

`inMessageID` specifies the message to be modified.

`inPayloadType` is an application-specific value that may be used to describe the contents of the payload. Negative values of payload type are reserved for use by VDK.

`inPayloadSize` indicates the size of the payload in the smallest addressable units of the processor (`sizeof(char)`).

`inPayloadAddr` is (typically) a pointer to the beginning of the payload buffer.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

- `kInvalidMessageOwner` indicates that the running thread is not the current owner of the message.
- `kInvalidMessageID` indicates that the argument `inMessageID` is not a valid message identifier.
- `kMessageInQueue` indicates that the message has been posted to a thread (the [ThreadID](#) is not known at this point), and it needs to be removed from the message queue by a call to [PendMessage\(\)](#).

Non error checking libraries: None

API Functions

SetPriority()

C Prototype

```
void VDK_SetPriority(const VDK_ThreadID    inThreadID,  
                   const VDK_Priority    inPriority);
```

C++ Prototype

```
void VDK::SetPriority(const VDK::ThreadID    inThreadID,  
                   const VDK::Priority    inPriority);
```

Description

Dynamically sets the priority of the named thread, overriding the default value. All threads are given an initial priority level at creation time. The thread's template specifies the priority initial value.

Parameters

`inPriority` is the new priority level. The [Priority](#) data type is described on [on page 4-24](#).

`inThreadID` is the thread to modify.

Scheduling

Invokes the scheduler and may result in a context switch

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

- `kUnknownThread` indicates that `inThreadID` is not a valid identifier.
- `kInvalidPriority` indicates that `inPriority` is not a priority of the [Priority](#) type.

Non error checking libraries: None

API Functions

SetThreadError()

C Prototype

```
void VDK_SetThreadError(VDK_SystemError inErr, int inVal);
```

C++ Prototype

```
void VDK::SetThreadError(VDK::SystemError inErr, int inVal);
```

Description

Sets the running thread's error value.

Parameters

`inErr` is the error enumeration. See [“SystemError” on page 4-26](#) for more information about errors.

`inVal` is the value whose meaning is determined by the error enumeration.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

No return value

Errors Thrown

None

SetThreadSlotValue()

C Prototype

```
bool VDK_SetThreadSlotValue(int inSlotNum, void *inValue);
```

C++ Prototype

```
bool VDK::SetThreadSlotValue(int inSlotNum, void *inValue);
```

Description

Sets the value in the currently running thread's slot table associated with `inSlotNum`. Returns `FALSE` if (and only if) `inSlotNum` does not identify a currently allocated slot. Otherwise, stores `inValue` in the thread slot identified by `inSlotNum` and returns `TRUE`.

Parameters

`inSlotNum` is the static library's preallocated slot number.

`inValue` is the value to store in thread's slot table (at `inSlotNum` entry).

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Return Value

`FALSE` if `inSlotNum` does not identify a currently allocated slot and `TRUE` otherwise.

Errors Thrown

None

API Functions

Sleep()

C Prototype

```
void VDK_Sleep(VDK_Ticks inSleepTicks);
```

C++ Prototype

```
void VDK::Sleep(VDK::Ticks inSleepTicks);
```

Description

Causes a thread to pause execution for at least the given number of clock ticks. Once delay ticks have elapsed, the calling thread is in the ready-to-run state. The thread resumes execution only if it is the highest priority thread with ready status. The minimum delay is one.

Parameters

`inSleepTicks` is the duration in ticks for which the thread should sleep.

Scheduling

Invokes the scheduler and will result in a context switch

Determinism

Not deterministic

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

- `kBlockInInvalidRegion` indicates that `Sleep()` is being called in an unscheduled region, causing a scheduling conflict.
- `kInvalidDelay` indicates that `inSleepTicks` is zero.

Non error checking libraries: None

API Functions

SyncRead()

C Prototype

```
unsigned int VDK_SyncRead(VDK_DeviceDescriptor inDD,  
                          char *outBuffer,  
                          const unsigned int inSize,  
                          VDK_Ticks inTimeout);
```

C++ Prototype

```
unsigned int VDK::SyncRead(VDK::DeviceDescriptor inDD,  
                           char *outBuffer,  
                           const unsigned int inSize,  
                           VDK::Ticks inTimeout);
```

Description

Invokes the read functionality of the driver.

Parameters

`inDD` is the [DeviceDescriptor](#) returned from [OpenDevice\(\)](#).

`outBuffer` is the address of the data buffer to be filled by the device.

`inSize` is the number of words to be read from the device.

`inTimeout` is the number of ticks to wait before timeout occurs.

Scheduling

Does not call the scheduler, but the user written device driver can call the scheduler.

Determinism

Constant time. Note that this function calls user written device driver code that may not be deterministic.

Return Value

Return value of the dispatch function if the device exists and `UINT_MAX` if it does not.

Errors Thrown

Full instrumentation and error checking libraries:

- `kBadDeviceDescriptor` indicates that `inDD` is not a valid [DeviceDescriptor](#).
- `kDeviceTimeout` indicates that the timeout value has expired before the read was finished.

Non error checking libraries: None

API Functions

SyncWrite()

C Prototype

```
unsigned int VDK_SyncWrite(VDK_DeviceDescriptor inDD,  
                           char *outBuffer,  
                           const unsigned int inSize,  
                           VDK_Ticks inTimeout);
```

C++ Prototype

```
unsigned int VDK::SyncWrite(VDK::DeviceDescriptor inDD,  
                             char *outBuffer,  
                             const unsigned int inSize,  
                             VDK::Ticks inTimeout);
```

Description

Invokes the write functionality of the driver.

Return Value

Return value of the dispatch function if the device exists and `UINT_MAX` if it does not.

Parameters

`inDD` is the [DeviceDescriptor](#) returned from [OpenDevice\(\)](#).

`outBuffer` is the address of the data buffer to be read by the device.

`inSize` is the number of words to be written to the device.

`inTimeout` is the number of ticks to wait before timeout occurs.

Scheduling

Does not call the scheduler, but the user written device driver can call the scheduler.

Determinism

Constant time. Note that this function calls user written device driver code that may not be deterministic.

Errors Thrown

Full instrumentation and error checking libraries:

- `kBadDeviceDescriptor` indicates that `inDD` is not a valid [DeviceDescriptor](#).
- `kDeviceTimeout` indicates that the timeout value has expired before the write was finished.

Non error checking libraries: None

API Functions

Yield()

C Prototype

```
void VDK_Yield(void);
```

C++ Prototype

```
void VDK::Yield(void);
```

Description

Yields control of the processor and moves the thread to the end of the wait queue of threads at its priority level. When [Yield\(\)](#) is called from a thread at a priority level using round-robin multithreading, the call also yields the remainder of the thread's time slice.

Parameters

None

Scheduling

Invokes the scheduler and may result in a context switch

Determinism

Constant time and conditional context switch.

Return Value

No return value

Errors Thrown

Full instrumentation and error checking libraries:

`kBlockInInvalidRegion` indicates that [Yield\(\)](#) is called in an unscheduled region, causing a scheduling conflict.

Non error checking libraries: None

Assembly Macros

This section describes the assembly-language macros that allow interrupt service routines (ISRs) to communicate with the VDK kernel. These macros are known collectively as the ISR API and are the only part of VDK that can be safely called from interrupt level.

In VDK applications, interrupt service routines should execute quickly and should leave as much of the processing as possible to be performed either by a thread or by a device driver activation. The principle purpose of the ISR API is therefore to provide the means to initiate such (thread or driver) activity.

In order to eliminate the overhead of saving and restoring the processor state, and of setting up a C run-time environment for each ISR entry, interrupt service routines for VDK are always written in assembly language and are responsible for saving and restoring any registers that they use. No assumptions can be made about the processor state at the time of entry to an ISR.

Each ISR API macro saves and restores all of the registers that it uses, and the macros are safe to use with nested interrupts enabled.

The ISR assembly macros are:

- [“VDK_ISR_ACTIVATE_DEVICE_\(\)” on page 5-129](#)
- [“VDK_ISR_CLEAR_EVENTBIT_\(\)” on page 5-130](#)

Assembly Macros

- “[VDK_ISR_LOG_HISTORY_EVENT_\(\)](#)” on page 5-131
- “[VDK_ISR_POST_SEMAPHORE_\(\)](#)” on page 5-132
- “[VDK_ISR_SET_EVENTBIT_\(\)](#)” on page 5-133

VDK_ISR_ACTIVATE_DEVICE_()

Prototype

```
VDK_ISR_ACTIVATE_DEVICE_(VDK_I0ID inID);
```

Description

Executes the named device driver prior to execution returning to the thread domain.

Parameters

`inID` is the device driver to run.

Scheduling

Invokes the scheduler prior to returning to the thread domain.

Determinism

Constant time

Errors Thrown

None

VDK_ISR_CLEAR_EVENTBIT_()

Prototype

```
VDK_ISR_CLEAR_EVENTBIT_(VDK_EventBitID inEventBit);
```

Description

Clears the value of `inEventBit` by setting it to 0 (FALSE). *All* event bit clears which occur in the interrupt domain are processed immediately prior to returning to the thread domain.

Parameters

`inEventBit` specifies the event bit to clear.

Scheduling

If `inEventBit` is currently set (1), the macro invokes the scheduler prior to returning to the thread domain. This allows the value of all dependent events to be recalculated and may cause a thread context switch.

Determinism

Constant time

Errors Thrown

None

VDK_ISR_LOG_HISTORY_EVENT_()

Prototype

```

VDK_ISR_LOG_HISTORY_EVENT_(VDK_HistoryEnum  inEnum,
                             int              inVal,
                             VDK_ThreadID    inThreadID);

```

Description

Adds a record to the history buffer. It is NULL if the `VDK_INSTRUMENTATION_LEVEL_` macro is set to zero or one (the value of two indicates that fully instrumented libraries are in use). See online Help for more information.

Parameters

`inEnum` is the enumeration value for this type of event. For more information, see [“HistoryEnum” on page 4-15](#).

`inVal` is the information defined by the enumeration.

`inThreadID` is the ThreadID to be stored with History Event.

Scheduling

Does not invoke the scheduler

Determinism

Constant time

Errors Thrown

None

VDK_ISR_POST_SEMAPHORE_()

Prototype

```
VDK_ISR_POST_SEMAPHORE_(VDK_SemaphoreID inSemaphoreID);
```

Description

Posts the named semaphore. Every time a semaphore is posted, its count increases until it reaches its maximum value, as specified on creation. Any further posts have no effect. All semaphore posts which occur in the interrupt domain are processed immediately prior to returning to the thread domain.

Parameters

`inSemaphoreID` is the semaphore to post.

Scheduling

If a thread is pending on `inSemaphoreID`, the macro invokes the scheduler prior to returning to the thread domain.

Determinism

Constant time

Errors Thrown

None

VDK_ISR_SET_EVENTBIT_()

Prototype

```
VDK_ISR_SET_EVENTBIT_(VDK_EventBitID inEventBit);
```

Description

Sets the value of `inEventBit` by setting it to 1 (TRUE). *All* event bit sets which occur in the interrupt domain are processed immediately prior to returning to the thread domain.

Parameters

`inEventBit` specifies the event bit to set.

Scheduling

If `inEventBit` is currently clear (zero), the macro invokes the scheduler prior to returning to the thread domain. This allows the value of all dependent events to be recalculated and may cause a thread context switch.

Determinism

Constant time

Errors Thrown

None

A PROCESSOR-SPECIFIC NOTES

This appendix provides processor-specific information for ADSP-BF53x, ADSP-219x, ADSP-TSxxx, ADSP-2106x, and ADSP-2116x processors.

Information will be added to this appendix to cover DSP architectures for which subsequent releases of VisualDSP++ 3.x take place.

VDK for Blackfin Processors (AD6532, ADSP-BF531, ADSP-BF532, ADSP-BF533, ADSP-BF535, and DM102)

User and Supervisor Modes

The Blackfin processor's architecture makes a distinction between execution in user and supervisor modes. Unlike VDK of VisualDSP++ 2.0, which supports both modes and switches back and forth as necessary, VDK of VisualDSP++3.0 runs entirely in supervisor mode, including all user thread code.

Since supervisor mode provides a superset of the capabilities of user mode, applications using VDK no longer need to be aware of the processor mode and do not need to raise exceptions in order to access protected resources.

Thread, Kernel, and Interrupt Execution Levels

VDK reserves execution level 15 as the run-time execution level for most user and VDK API code, and execution level 14 for internal VDK operations. In the following text, these are referred to as “Thread Level” and “Kernel Level”, respectively. Execution levels 13–6 are collectively referred to as “Interrupt Level”. All of these levels (15–6) execute in supervisor mode.

All thread functions execute at Thread Level (execution level 15), including `Run()` and `ErrorHandler()`. Conversely, all Interrupt Service Routines (ISRs) execute at higher priority (lower numbered) execution levels, according to the interrupt source that invoked them. The implementation function for device drivers (their single entry point) may be called by the kernel at either Thread Level (execution level 15) or at Kernel Level (execution level 14), depending on the purpose of the call.

Device driver ‘activate’ (`kIO_Activate`) functionality is the only user code that executes at Kernel Level. All other device driver code executes at Thread Level. Entry to Kernel Level is, in this case, initiated by an ISR calling [VDK_ISR_ACTIVATE_DEVICE_\(\)](#) and is, therefore, asynchronous with respect to Thread Level, except within a critical region. Thus, care must be taken to synchronize access to shared data between Thread level and Kernel Level, as well as between Thread Level and Interrupt Level (and also between Kernel Level and Interrupt Level). Critical regions may be used for both of these purposes.

Compared with VisualDSP++ 2.0, there are fewer limits on the operations that can be performed in threads. System memory mapped registers may be accessed directly and at any time. It is, however, the user’s responsibility to ensure the operations are performed correctly and at appropriate times.

Critical and Unscheduled Regions

Because VDK now executes entirely in supervisor mode, the execution-time cost of entering or leaving a critical region is reduced, compared to that in VisualDSP++ 2.0. However, because VDK now disables interrupts for much shorter periods of time, it is more likely that the worst-case interrupts-off time will be set by critical regions in user code. Therefore, care must be taken that such usage does not impact the interrupt latency of the system to an unacceptable degree.

Exceptions

VDK reserves service exception ID 0 (EXCPT 0) for internal use. Additionally, the Integrated Development and Debugging Environment (IDDE) automatically generates a source file for all VDK projects for Blackfin processors. The source file defines an entry point for any service or error exceptions you wish to trap. When an exception occurs, VDK intercepts the exception; if it is not the VDK exception, the user defined exception handler executes.

Do not manipulate the IMASK system register from within your exception handler. If you need to mask or unmask an interrupt in response to an exception, raise an interrupt and change the value of IMASK in the ISR after the exception handler.

ISR APIs

The Blackfin processor's assembly syntax requires the use of separate API macros, depending on whether the arguments are constants (immediate values, enumerations) or data registers (R0 through R7). The arguments to

VDK for Blackfin Processors (AD6532, ADSP-BF531, ADSP-BF532, ADSP-BF533, ADSP-BF535, and DM102)

the default APIs, described in Chapter 5, “[VDK API Reference](#)”, must be constants. When passing data registers as arguments, append “REG_” to the macro name as follows.

```
VDK_ISR_POST_SEMAPHORE_REG_(semaphore_num_);  
VDK_ISR_ACTIVATE_DEVICE_REG_(dev_num_);  
VDK_ISR_SET_EVENTBIT_REG_(eventbit_num_);  
VDK_ISR_CLEAR_EVENTBIT_REG_(eventbit_num_);  
VDK_ISR_LOG_HISTORY_REG_(enum_, value_, threadID_);
```

The API macros, as defined in “[Assembly Macros](#)” on page 5-127 (without “REG_”), only accept constants as arguments. Passing a register name results in an assembler error.

Interrupts

The following hardware events (interrupts) are reserved for use by VDK on Blackfin processors.

- EVT_EVSW (EVT_EVX) – the software exception handler. User code handles software exceptions by modifying the source file created by the VisualDSP++ file `ExceptionHandler-BF532.asm` or `ExceptionHandler-BF535.asm`, depending on the processor.
- EVT_TMR – the interrupt associated with the timer integral to the processor core. This timer generates the interrupts for system ticks and provides all VDK timing services. Disabling this timer stops sleeping, round-robin scheduling, pending with timeout, and periodic semaphores.
- EVT_IVG14 – general interrupt #14. This interrupt is reserved for use by VDK and may not be used in any other manner.
- EVT_IVG15 – general interrupt #15. This interrupt is reserved to provide a supervisor mode runtime for user and VDK code and may not be used in any other manner.

The ADSP-BF53x processor designates hardware events seven (`IVG7`) through thirteen (`IVG13`) as ‘general interrupts’ and maps each to more than one physical peripheral. The IDDE generates source code templates with a single entry point per interrupt level, rather than for each peripheral. Therefore, you are responsible for dispatching interrupts when more than one peripheral is used at the same level. The technique used depends on the application, but, typically, either chaining or a jump table is used. When chaining, a handler will check to see whether the interrupt was caused by the specific peripheral it knows how to handle. If not, execution is passed to the next handler in the chain. A jump table uses an identifier or a constant as index to a table of function pointers when searching for the appropriate interrupt handler.

Timer

`VDK_Ticks` are derived from the timer implemented in the inner processor core of Blackfin processors and are synchronized to the main core clock `CCLK`. However, this timer is disabled when an ADSP-BF53x processor enters low power mode. Thus, all VDK timing services (such as sleeping, timeouts, and periodic semaphores) do not operate while the core is in `IDLE` or low power mode.

The ADSP-BF535 and AD6532 processor simulators shipped with VisualDSP++ 3.x exhibit an anomaly in the manner in which the timer is clocked. The simulator decrements the time once per instruction without taking into account the number of cycles required to dispatch the instruction; for example, stall cycles introduced by access `L2` memory. Thus, the period of a `Tick` is a larger number of cycles than on the actual hardware. This behavior differs from execution in hardware.

ADSP-BF531, ADSP-BF532, ADSP-BF533, DM102 Processor Memory

By default, the `VDK-BF532.LDF`, `VDK-BF533.LDF`, and `VDK-DM102.LDF` files place all user code into a section called `program`, and all static and global data into a section called `data`. The `VDK-BF531.LDF` file places all user code into sections called `program_ram1` and `program_ram2` and all static and global data into a section called `data`.

The `program` section (or `program_ram1` and `program_ram2` sections) occupy most of the processor's L1 SRAM code area, together with the startup and interrupt vectors. The `data` section occupies the lower part of the L1 Data Bank A SRAM. The upper part is occupied by the system (i.e. kernel) stack.

The L1 Data Bank B SRAM is entirely given over to the heap. The stack regions for individual threads are allocated from the heap.

These assignments may be overridden by editing the LDF file. Refer to the *VisualDSP++ 3.x Linker and Utilities Manual for Blackfin Processors* for information on segmenting your code and data.

ADSP-BF535 and AD6532 Processor Memory

By default, the `VDK-BF535.LDF` and `VDK-AD6532.LDF` files place all user code and data into a single section called `program`, which is placed in L2 memory. Refer to the *VisualDSP++ 3.x Linker and Utilities Manual for Blackfin Processors* for information on segmenting your code and data.

By default, the L1 memory regions are not used in the `.LDF` file and are not configured to be used as Level 1 cache. If you choose to use the processor in a L1 cache mode, you must write an initialization device driver for configuration, as described in [“Device Drivers” on page 3-38](#).

A cache-control device driver for the ADSP-BF535 processor is provided as one of the VDK examples and is located in the VisualDSP++3.x installation under:

```
...VisualDSP\Blackfin\Examples\VDK Examples\BF535\Cache_Cntrl_BF535
```

Interrupt Nesting

In VisualDSP++ 3.x, VDK fully supports nested interrupts:

- The skeleton interrupt service routines, which are generated by the VisualDSP++ 3.x IDDE for user defined interrupt handlers, enable nesting by pushing the contents of the RETI register onto the stack on entry and popping it immediately prior to exit.
- The ISR API macros are fully reentrant.

Note that the skeleton ISRs in VisualDSP++ 2.0 do not enable nesting. When converting existing projects to VisualDSP++ 3.x, manually add the `[--SP] = RETI;` and `RETI = [SP++];` instructions to the existing ISRs to obtain the benefits of interrupt nesting. Conversely, if nesting is not required in VisualDSP++ 3.x projects, it is acceptable to manually delete these instructions from the ISRs.

Thread Stack Usage by Interrupts

Because all thread code executes in supervisor mode, there is no automatic switching between user and system stack pointers. Therefore, all ISRs execute using the stack of the current thread, which is the thread that is executing at the time the interrupt is serviced. This means each thread stack must—in addition to the thread's own requirements—reserve sufficient space for the requirements of ISRs. When interrupt nesting is enabled (as it is by default), the worst-case space requirement is the total of the requirements of the individual ISRs. When nesting is disabled, the requirement is only the largest of the individual ISR requirements, which is one possible reason for disabling interrupt nesting.

Interrupt Latency

Every effort has been made to minimize the duration of the intervals in which interrupts are disabled by VDK. Interrupts are disabled only where necessary for synchronization with Interrupt Level and for the shortest feasible number of instructions. The instruction sequences executed during these interrupts-off periods are deterministic.

Within VDK itself, synchronization between Thread Level and Kernel Level is achieved by selectively masking the Kernel Level interrupt, while leaving the higher priority interrupts unmasked.

VDK for ADSP-219x DSPs (ADSP-2191, ADSP-2192-12, ADSP-2195, and ADSP-2196)

Thread, Kernel, and Interrupt Execution Levels

VDK runs most user and VDK API code at the normal (non-interrupt) execution level but also reserves one of the low priority interrupt levels for internal VDK operations. In the following text, these are referred to as “Thread Level” and “Kernel Level”, respectively. The remaining interrupt levels are collectively referred to as “Interrupt Level”.

All thread functions execute at Thread Level, including `Run()` and `ErrorHandler()`. Conversely, all Interrupt Service Routines execute at higher priority execution levels, according to the interrupt source that invoked them. The implementation function for device drivers (their single entry point) may be called by the kernel at either Thread Level or at Kernel Level, depending on the purpose of the call.

Device driver ‘activate’ (`kIO_Activate`) functionality is the only user code which executes at Kernel Level (all other device driver code executes at Thread Level). Entry to Kernel Level is, in this case, initiated by an ISR calling `VDK_ISR_ACTIVATE_DEVICE_()` and is, therefore, asynchronous with respect to Thread Level, except within a critical region. For this reason, care must be taken to synchronize access to shared data between Thread Level and Kernel Level, as well as between Thread Level and Interrupt Level (and also between Kernel Level and Interrupt Level). Critical regions may be used for both of these purposes.

Critical and Unscheduled Regions

VDK in VisualDSP++ 3.x disables interrupts internally for much shorter periods of time than in VisualDSP++ 2.0. It is, therefore, more likely that the worst-case interrupts-off time will now be set by the use of critical regions in user code. Care must be taken that such usage does not impact the interrupt latency of the system to an unacceptable degree.

Interrupts on ADSP-2192 DSPs

The following interrupts are reserved for use by VDK on the ADSP-2192 DSP.

- `IRPTL[5]` – the timer interrupt. `Timer2` generates the interrupts for system ticks and provides all VDK timing services. Disabling this timer stops sleeping, round robin scheduling, pending with time-out, and periodic semaphores. This interrupt is reserved for use by the scheduler and may not be used in any other manner.
- `IRPTL[14]` – the kernel interrupt. This interrupt is reserved for use by VDK and may not be used in any other manner.

When using VDK, there are some restrictions placed on the ISRs associated with the AC'97 codec port. This is due to the AC'97 interrupt being hardwired to `IRPTL[15]` and, therefore, executing at a lower priority than

VDK for ADSP-219x DSPs (ADSP-2191, ADSP-2192-12, ADSP-2195, and ADSP-2196)

VDK's Kernel Level interrupt. This means that your entire ISR at priority level 15 must execute with interrupts disabled. You should never enable interrupts while servicing an AC'97 Frame, and a nested interrupt should never be allowed to occur once an AC'97 Frame ISR begins.

Interrupts on ADSP-2191 DSPs

The following interrupts are reserved for use by VDK on the ADSP-2191 DSP.

- `IRPTL[4]` – the timer interrupt. `Timer2` generates the interrupts for system ticks and provides all VDK timing services. Disabling this timer stops sleeping, round robin scheduling, pending with time-out, and periodic semaphores. This interrupt is reserved for use by the scheduler and may not be used in any other manner.
- `IRPTL[15]` – the kernel interrupt. This interrupt is reserved for use by VDK and may not be used in any other manner.

The ADSP-2191 DSP interrupt controller allows Interrupt Levels five (`IRPTL[5]`) through fourteen (`IRPTL[14]`) to be mapped in software to the various on-chip peripherals. VDK uses this facility itself to map `Timer2` to priority level four, as mentioned above. User code may set up other mappings but must neither change the mapping of `Timer2` nor map any other peripheral onto priority four.

The VisualDSP++ 3.x IDDE generates ISR source code templates with a single entry point per Interrupt Level, rather than one for each peripheral. Therefore, if several peripherals share the same Interrupt Level, then user code (within the ISR for that level) is responsible for determining the true source of the interrupt and dispatching to the appropriate handler.

Timer

The ADSP-2192 DSP has a single timer, which is reserved by VDK for its internal timekeeping functions and may not be used in any other manner.

The ADSP-2191 DSP has three identical timers (`Timer0`-`Timer2`), one of which (`Timer2`) is reserved by VDK for its internal timekeeping functions and may not be used in any other manner.

Memory

By default, the VDK LDF file places all user and VDK code into a single 24-bit section called `program`. Most global and static data is placed in the 16-bit section `data1`. A separate memory region is used for the heap.

These default arrangements can be customized by editing the project's LDF file. Refer to the *VisualDSP++ 3.x Linker and Utilities Manual for ADSP-21xx DSPs* for information on segmenting your code.

Interrupt Nesting

In VisualDSP++ 3.x, VDK fully supports nested interrupts (with the exception, as mentioned [on page A-9](#), of the AC'97 interrupt on ADSP-2192 DSPs):

- The skeleton Interrupt Service Routines, which are generated by the VisualDSP++ 3.x IDDE for user defined interrupt handlers, enable nesting by re-enabling interrupts after entry.
- The ISR API macros are re-entrant between different interrupt levels.

Note that the skeleton ISRs in VisualDSP++ 2.0 do not enable nesting. When converting existing projects, it is, therefore, necessary to add the “`ENA INT;`” instructions by hand to the existing ISRs in order to obtain the benefits of interrupt nesting. Conversely, if nesting is not required in VisualDSP++ 3.x projects, it is acceptable to manually delete these instructions from the ISRs.

Interrupt Latency

Every effort has been made to minimize the duration of the intervals where interrupts are disabled by VDK. Interrupts are disabled only where necessary for synchronization with Interrupt Level, and then for the shortest feasible number of instructions. The instruction sequences executed during these interrupts-off periods are deterministic.

Within VDK itself, synchronization between Thread Level and Kernel Level is achieved by selectively masking the Kernel Level interrupt, while leaving the higher priority interrupts unmasked.

VDK for ADSP-TS101 TigerSHARC Processors

Thread, Kernel, and Interrupt Execution Levels

VDK runs most user and VDK API code at the normal (non-interrupt) execution level but also reserves one of the low-priority interrupt levels for internal VDK operations. In the following text, these are referred to as “Thread Level” and “Kernel Level”, respectively. The remaining interrupt levels are collectively referred to as “Interrupt Level”.

All thread functions execute at Thread Level, including `Run()` and `ErrorHandler()`. Conversely, all Interrupt Service Routines execute at higher priority execution levels, according to the interrupt source that invoked them. The implementation function for device drivers (their single entry point) may be called by the kernel at either Thread Level or at Kernel Level, depending on the purpose of the call.

Device driver ‘activate’ (`kIO_Activate`) functionality is the only user code which executes at Kernel Level (all other device driver code executes at Thread Level). Entry to Kernel Level is, in this case, initiated by an ISR calling [VDK_ISR_ACTIVATE_DEVICE_\(\)](#) and is, therefore, asynchronous with respect to Thread Level (except within a critical region). For this reason, care must be taken to synchronize access to shared data between Thread Level and Kernel Level, between Thread Level and Interrupt Level, as well as between Kernel Level and Interrupt Level. Critical regions may be used for both of these purposes.

Critical and Unscheduled Regions

VDK in VisualDSP++ 3.x disables interrupts internally for much shorter periods of time than in VisualDSP++ 2.0. It is, therefore, more likely that the worst-case interrupts-off time will now be set by the use of critical regions in user code. Care must be taken that such usage does not impact the interrupt latency of the system to an unacceptable degree.

Interrupts on ADSP-TS101 Processors

The following interrupts are reserved for use by VDK on the ADSP-TS101 processor.

- `Timer1HP` – the timer interrupt. `Timer1` generates the interrupts for system ticks and provides all VDK timing services. Disabling this timer stops sleeping, round robin scheduling, pending with time-out and periodic semaphores. This interrupt is reserved for use by the scheduler and may not be used in any other manner.
- `Timer0LP` – the kernel interrupt. This interrupt is reserved for use by VDK and may not be used in any other manner.

The corresponding priority levels for each timer, `Timer1LP` and `Timer0HP`, are also reserved and may not be used in any other manner.

Timer

On ADSP-TS101 processors, both `Timer0` and `Timer1` are reserved by VDK for its internal functions and may not be used in any other manner.

ADSP-TS101 Processor Memory

By default, the VDK LDF file places all user and VDK code into a single section named `program`, which is mapped into memory block 0. Most global and static data is placed in the section `data1`, which is mapped into memory block 1. A section named `data2` is mapped into block 2 and is available for variables with an explicit segment assignment.

A separate memory region in block 1 is used for the system heap, and another such region in block 2 is used for a secondary heap. VDK thread stacks are allocated from heap space, and the use of two heaps allows VDK to allocate the thread J and K stacks in different memory blocks.

‘System’ J and K stacks also exist in memory blocks 1 and 2, respectively. This stack pair is used by VDK itself for Kernel-level processing. This means that when device-driver ‘activate’ code is invoked, it is these system stacks which are in use. Therefore, you need to ensure that the size of these stacks is sufficient for the requirements of the user supplied activate code.

These default arrangements can be customized by editing the project’s LDF file. Refer to the *VisualDSP++ 3.x Linker and Utilities Manual for TigerSHARC Processors* for information on segmenting your code. Refer to the *VisualDSP++ 3.x C/C++ Compiler and Library Manual for TigerSHARC Processors* for information on defining additional heaps.

Interrupt Nesting

In VisualDSP++ 3.x, VDK fully supports nested interrupts:

- The skeleton Interrupt Service Routines, which are generated by VisualDSP++ for user defined interrupt handlers, enable nesting by storing the `RETIB` register to memory after entry.
- The ISR API macros are re-entrant between different interrupt levels.

VDK for SHARC DSPs (ADSP-21060, ADSP-21061, ADSP-21062, ADSP-21065L, ADSP-21160, and ADSP-21161)

Note that the skeleton ISRs in VisualDSP++ 2.0 do not enable nesting. Therefore, when converting existing projects, it is necessary to add the instructions to store and reload RETIB to the existing ISRs in order to obtain the benefits of interrupt nesting. Conversely, if nesting is not required in VisualDSP++ 3.x projects, it is acceptable to manually delete these instructions from the ISRs.

Interrupt Latency

Every effort has been made to minimize the duration of the intervals where interrupts are disabled by VDK. Interrupts are disabled only where necessary for synchronization with Interrupt Level and then, for the shortest feasible number of instructions. The instruction sequences executed during these interrupts-off periods are deterministic.

Within VDK itself, synchronization between Thread Level and Kernel Level is achieved by selectively masking the Kernel Level interrupt, while leaving the higher priority interrupts unmasked.

VDK for SHARC DSPs (ADSP-21060, ADSP-21061, ADSP-21062, ADSP-21065L, ADSP-21160, and ADSP-21161)

Thread, Kernel and Interrupt Execution Levels

VDK runs most user and VDK API code at the normal (non-interrupt) execution level but also reserves one of the low-priority interrupt levels for internal VDK operations. In the following text, these are referred to as “Thread Level” and “Kernel Level”, respectively. The remaining interrupt levels are collectively referred to as “Interrupt Level”.

All thread functions execute at Thread Level, including `Run()` and `ErrorHandler()`. Conversely, all Interrupt Service Routines execute at higher-priority execution levels according to the interrupt source that invoked them. The implementation function for device drivers (their single entry point) may be called by the kernel at either Thread Level or at Kernel Level, depending on the purpose of the call.

Device driver ‘activate’ (`kIO_Activate`) functionality is the only user code which executes at Kernel Level (all other device driver code executes at Thread Level). Entry to Kernel Level is, in this case, initiated by an ISR calling `VDK_ISR_ACTIVATE_DEVICE_()` and is, therefore, asynchronous with respect to Thread Level, except within a critical region. For this reason, care must be taken to synchronize access to shared data between Thread level and Kernel Level, as well as between Thread Level and Interrupt Level (and also between Kernel Level and Interrupt Level). Critical regions may be used for both of these purposes.

Critical and Unscheduled Regions

VDK in VisualDSP++ 3.x disables interrupts internally for much shorter periods of time than in VisualDSP++ 2.0. It is, therefore, more likely that the worst-case interrupts-off time will now be set by the use of critical regions in user code. Care must be taken that such usage does not impact the interrupt latency of the system to an unacceptable degree.

Interrupts on ADSP-2106x DSPs

The following interrupts are reserved for use by VDK on the ADSP-21060, ADSP-21061, ADSP-21062 and ADSP-21065L DSPs.

- `TMZHI` – the timer interrupt (`IRPTL/IMASK` bit 4). The core timer (`Timer0` on ADSP-21065L DSPs) generates the interrupts for system ticks and provides all VDK timing services. Disabling this

VDK for SHARC DSPs (ADSP-21060, ADSP-21061, ADSP-21062, ADSP-21065L, ADSP-21160, and ADSP-21161)

timer stops sleeping, round robin scheduling, pending with timeout, and periodic semaphores. This interrupt is reserved for use by the scheduler and may not be used in any other manner.

- **SFT3I** – the kernel interrupt (**IRPTL/IMASK** bit 31). This interrupt is reserved for use by VDK and may not be used in any other manner.

Note that the low priority interrupt for the timer (**TMZLI**) is available for use, either as a software interrupt, or to support a second timer where one is present in the hardware (as with ADSP-21065L DSPs). By default, VDK assigns `Timer1` to the low-priority timer interrupt (**IRPTL/IMASK** bit 23), so using it only requires an interrupt handler to be defined for **TMZLI** and for `Timer1` to be initialized.

Interrupts on ADSP-2116x DSPs

The following interrupts are reserved for use by VDK on the ADSP-21160 and ADSP-21161 DSPs.

- **TMZHI** – the timer interrupt (**IRPTL/IMASK** bit 4). `Timer0` generates the interrupts for system ticks and provides all VDK timing services. Disabling this timer stops sleeping, round robin scheduling, pending with timeout, and periodic semaphores. This interrupt is reserved for use by the scheduler and may not be used in any other manner.
- **SFT3I** – the kernel interrupt (**IRPTL/IMASK** bit 30). This interrupt is reserved for use by VDK and may not be used in any other manner.

There is no API for setting or clearing bits in the **LIRPTL** register, so selective masking and unmasking of individual link port interrupts must be done using inline assembly code. If the link port interrupts are set to be “enabled at boot”, then this should not be necessary. Masking and

unmasking of all link port interrupts together can be done via

```
VDK::SetInterruptMaskBits(LPISUMI) and
```

```
VDK::ClearInterruptMaskBits(LPISUMI).
```

On ADSP-21160 DSPs, if the `VIRPT` interrupt is being used, then it cannot be enabled at boot (due to a bit-position clash with link port 5).

Instead, it must be enabled explicitly via

```
VDK::SetInterruptMaskBits(VIRPTI).
```

Timer

Where only one timer is provided by the hardware, it is reserved by VDK for its internal timekeeping functions and may not be used in any other manner.

Where more than one timer is provided (i.e. on ADSP-21065L DSPs), `Timer0` is reserved by VDK for its internal timekeeping functions and may not be used in any other manner. `Timer1` may be used by user code.

Memory

By default, the VDK LDF file places all user and VDK code into a single section called “program”, which is mapped into `PM` memory. Most global and static data is placed in the `data1` section, which is mapped into `DM` memory. A separate memory region in `DM` memory is used for the system heap.

A “System” stack also exists in `DM` memory. This stack is used by VDK itself for Kernel Level processing. This means that whenever device driver ‘activate’ code is invoked, it is this system stack which is in use. Care should, therefore, be taken that the size of this stack is sufficient for requirements of the user-supplied activate code.

VDK for SHARC DSPs (ADSP-21060, ADSP-21061, ADSP-21062, ADSP-21065L, ADSP-21160, and ADSP-21161)

These default arrangements can be customized by editing the project's LDF file. Refer to the *VisualDSP++ 3.x Linker and Utilities Manual for SHARC DSPs* for information on segmenting your code. Refer to the *VisualDSP++ 3.x C/C++ Compiler and Libraries Manual for SHARC DSPs* for information on defining additional heaps.

Register Usage

The ADSP-21x6x DSP devices provide a set of alternate (or background) registers for both the computation units and the DAGs. VDK does not save and restore these alternate registers during context switches but instead employs them to accelerate entry to and exit from the scheduler interrupt by providing it with a dedicated C run-time environment.

This means that the following alternate registers contain fixed values and must not be changed by user code:

Table A-1. ADSP-21x6x DSP Alternate Registers

Register	Value
M5', M13'	0
M6', M14'	1
M7', M15'	-1
B6', B7'	System Stack base
L6', L7'	System Stack length
L2' - L5', L10' - L15'	0

Since these are used at Interrupt Level (which may be entered at any time), it is not sufficient to save these registers and restore them after use, except within a critical region.

The following alternate registers are not used by VDK and are available for use by user code:

I0', M0', B0', L0',
I1', M1', B1', L1',
I8', M8', B8', L8',
I9', M9', B9', L9'

Note that these registers do not form part of the thread context and, hence, are not context-switched between threads. Therefore, their most appropriate use is in implementing fast I/O interrupt handlers for devices where DMA is not available.

The remaining alternate registers:

R0'-R15',
I2'-I7', I10'-I15',
M2'-M4', M10'-M12',
B2'-B5', B10'-B15'

are used by the scheduler ISR and, hence, may change at any time (unless inside a critical region). If these registers are to be used in a user-defined Interrupt Service Routine, then their contents must be saved and restored before returning from the interrupt, as for the primary registers.

Note that the background multiplier result register, MRB, is not considered to be an alternate register in the above discussion, and that both MRB and MRF are context-switched between threads.

40-bit Register Usage

The ALU registers of the ADSP-21x6x DSPs devices are 40 bits in width in order to support the 40-bit extended precision floating-point mode. However, the C/C++ run-time environment does not support this mode

VDK for SHARC DSPs (ADSP-21060, ADSP-21061, ADSP-21062, ADSP-21065L, ADSP-21160, and ADSP-21161)

and, instead, runs with 32-bit rounding enabled (bit `RND32=1` in the `MODE1` register). The eight additional Least Significant Bits (LSBs) of the mantissa are, therefore, not used in C/C++ applications.

VDK only preserves the 32-bit form of the data registers during a context switch. In addition, both the VDK `Timer` ISR and any user ISR which calls the VDK ISR API macros, may cause the eight additional LSBs of `F0–F3` to be cleared. Neither of these restrictions cause any difficulty for C/C++ code or for assembly code which runs with `RND32=1` (the default), but code that requires the use of the 40-bit extended precision mode must observe the following restrictions:

- It must be written in assembly
- It must clear `RND32` explicitly and set it again when the 40-bit computation is complete.
- It must execute within an unscheduled region.
- It must avoid the use of `F0`, `F1`, `F2`, and `F3`, unless all interrupts are disabled (`F4–F15` are safe to use with interrupts enabled).

Interrupt Nesting

In VisualDSP++ 3.x, VDK fully supports nested interrupts:

- The skeleton Interrupt Service Routines, which are generated by the VisualDSP++ 3.x IDDE for user-defined interrupt handlers, enable nesting by re-enabling interrupts after entry.
- The ISR API macros are re-entrant between different interrupt levels.

Note that the skeleton ISRs in VisualDSP++ 2.0 do not enable nesting, but a comment indicates that this should be done before returning from the ISR. When converting existing projects, it is, therefore, necessary to manually move the instruction to re-enable interrupts to an earlier point

in existing ISRs in order to obtain the benefits of interrupt nesting. Conversely, if nesting is not required in VisualDSP++ 3.x projects, it is acceptable to manually delete these instructions from the ISRs.

Interrupt Latency

Every effort has been made to minimize the duration of the intervals where interrupts are disabled by VDK. Interrupts are disabled only where necessary for synchronization with Interrupt Level, and then for the shortest feasible number of instructions. The instruction sequences executed during these interrupts-off periods are deterministic.

Within VDK itself, synchronization between Thread Level and Kernel Level is achieved by selectively masking the Kernel Level interrupt, while leaving the higher priority interrupts unmasked.

**VDK for SHARC DSPs (ADSP-21060, ADSP-21061, ADSP-21062,
ADSP-21065L, ADSP-21160, and ADSP-21161)**

B MIGRATING DEVICE DRIVERS

The device driver architecture has fundamentally changed between VisualDSP++ 2.0 and VisualDSP++ 3.x. Device drivers have become part of the I/O interface, and as a part of that move are now class based. While the underlying changes to device drivers have a minimal effect from the usage perspective, device drivers created under VisualDSP++ 2.0 are incompatible with the device drivers of VisualDSP++ 3.x.

This appendix describes how to convert device drivers of the previous release for use in projects built using VisualDSP++ 3.x.

The converting procedure includes:

- [“Step 1: Saving Existing Sources” on page B-1](#)
- [“Step 2: Revising Properties” on page B-2](#)
- [“Step 3: Revising Sources” on page B-3](#)
- [“Step 4: Creating Boot Objects” on page B-4](#)

Step 1: Saving Existing Sources

Make backup copies of all of the existing VisualDSP++2.0 device driver sources and header files.

Step 2: Revising Properties

Open the existing VisualDSP++ 2.0 device driver project in VisualDSP++ 3.x IDDE. The following changes pertaining to device drivers can be observed in the kernel window:

- a. The VisualDSP++ 2.0 **Device drivers** node can now be found under the **I/O Interface** node in the kernel window. Due to the class based model of device drivers in VisualDSP++ 3.x, there is also the **Boot I/O Objects** node under the **I/O Interface** node.
- b. Under the **I/O Interface** node, there is a newly added property **Max number of I/O Objects**. Set this property to the number of device drivers to be used.
- c. Device flags are no longer associated directly with a specific device driver. In VisualDSP++ 3.x, there is a separate **Device Flags** node in the kernel window.

Create a device flag with the same name under the **Device Flags** node for each device flag present in the original VisualDSP++ 2.0 project.

Step 3: Revising Sources

Changing the existing device driver source and header files to the new model:

- a. Delete all the existing **Device Drivers** that have been imported from the original VisualDSP++ 2.0 project under the **I/O Interface, Device Drivers** node in the kernel window. Remove the original device driver source and header files from the project directory.
- b. Create new device drivers from the **I/O Interface, Device Drivers** node with the same names as those in the original VisualDSP++ 2.0 project. The automatically generated source files use the VisualDSP++ 3.x templates.
- c. Copy any code added to the VisualDSP++ 2.0 device driver header files to the analogous location in the equivalent header files generated from the VisualDSP++3.0 templates. Changes may include additional variable declarations and include files.
- d. Copy any code added to the VisualDSP++ 2.0 device driver source files to the analogous location in the equivalent header files generated from the VisualDSP++3.0 templates.

Note that the dispatch function cases are renamed from `kDD::xxx` to `kIO::xxx` in C++ device drivers and `kDD_xxx` to `kIO_xxx` in C device drivers.

- e. Check all project sources and header files for references to the dispatch functions `VDK::kDD_xxx` (C++ device drivers) or `VDK_kDD_xxx` (C device drivers) and replace them with `VDK::kIO_xxx` or `VDK_kIO_xxx`, respectively.

Step 4: Creating Boot Objects

For C++ device drivers, the dispatch function is scoped under the device driver class in VisualDSP++ 3.x. For C device drivers, the dispatch function is identical to that in VisualDSP++ 2.0, apart from the renaming of cases.

- f. Modify the error checking code.

The errors thrown by VDK functions have changed from `VDK::kDDxxx` (C++ device drivers) or `VDK_kDDxxx` (C device drivers) to `VDK::kxxx` or `VDK_kxxx`, respectively.

- g. The `DeviceDispatchUnion` data type has been renamed `DispatchUnion` in VisualDSP++ 3.x and any variables of that type must be renamed accordingly.

Step 4: Creating Boot Objects

Create boot I/O objects for each device driver:

In VisualDSP++ 3.x the created device drivers are merely templates. In order to instantiate a device driver so it can be used at boot, an I/O object must be created using the device driver template.

For each device driver created under the **I/O Interface, Device Drivers** node, create a boot I/O object of that type from under **I/O Interface, Boot I/O Objects**.

Note that the name of the boot I/O object must be different from that of the device driver template.

I INDEX

Symbols

.EXTERN, assembly directive [3-7](#)
.GLOBAL, assembly directive [3-7](#)
.LDF files [2-2](#), [5-2](#), [A-6](#), [A-11](#), [A-15](#),
[A-19](#)

A

abstract base thread class [3-6](#)
abstracting hardware implementation
[3-38](#)
AC'97 codec port [A-9](#)
ADSP-219x DSPs [A-8](#)
 enabling/disabling interrupts [A-9](#)
 execution levels [A-8](#)
 interrupt latency [A-12](#)
 interrupt nesting [A-11](#)
 ISRs [A-9](#), [A-10](#)
 memory [A-11](#)
 timer [A-9](#), [A-10](#)
ADSP-21xxx DSPs, See SHARC DSPs
ADSP-BF53x processors, See Blackfin
 processors
ADSP-TSxxx DSPs, See TigerSHARC
 DSPs
alignment constraints [3-56](#)
AllocateThreadSlot() [5-11](#)
AllocateThreadSlotEx() [5-13](#)

allocating

 memory, malloc/new [3-5](#), [3-7](#)
 system resources, threads [3-5](#)
alternate registers (ADSP-21x6x DSPs)
[A-20](#)

API

 function parameters [5-3](#)
 functions list [5-4](#)
 header (vdk.h) [5-3](#)
 library, reference format [5-10](#)
 linking library functions [5-2](#)
 naming conventions [5-3](#)
 See also assembly macros
architecture, ISRs [3-34](#)
assembly threads [3-7](#)

B

Bitfield data type [4-3](#)
Blackfin DSPs
 enabling/disabling interrupts [A-3](#)
 exceptions [A-3](#)
 execution levels [A-2](#)
 interrupt latency [A-8](#)
 interrupt nesting [A-7](#)
 ISR APIs [A-3](#)
 ISRs [A-4](#)
 memory [A-5](#)

INDEX

thread stacks [A-7](#)
timer [A-5](#)
blocking on semaphores [3-16](#), [3-18](#)
blocks of memory [3-55](#)
boot I/O objects [B-4](#)
boot threads [3-4](#)

C

C threads [3-7](#)
C++ threads [3-6](#)
C/C++ run-time environment [A-21](#)
calling library functions [5-2](#)
channels [3-21](#)
circular buffers [2-3](#)
ClearEventBit() [3-27](#), [3-30](#), [5-15](#)
ClearInterruptMaskBits() [3-33](#), [5-17](#)
ClearThreadError() [5-18](#)
CloseDevice() [3-47](#), [3-48](#), [5-19](#)
code reuse [1-3](#)
Communication Manager [3-38](#)
configuring VDK projects [2-1](#)
constructors, C++ threads [3-4](#), [3-5](#)
context switches [3-18](#), [3-30](#), [3-31](#),
[3-34](#), [3-37](#)
conventions of this manual [xxvi](#)
cooperative
 multithreading [3-10](#)
 scheduling [3-10](#)
create functions, threads [3-4](#)
CreateDeviceFlag() [3-6](#), [5-21](#)
CreateMessage() [5-22](#)
CreatePool() [5-24](#)
CreatePoolEx() [5-26](#)
CreateSemaphore() [5-28](#)

CreateThread() [3-4](#), [5-30](#)
CreateThreadEx() [5-32](#)
critical regions [3-8](#), [3-32](#), [3-33](#), [3-35](#),
[3-39](#), [3-47](#), [3-54](#)
customer support [xx](#)

D

data types
 Bitfield [4-3](#)
 DeviceDescriptor [4-4](#)
 DeviceFlagID [4-5](#)
 DispatchID [4-6](#)
 DispatchUnion [4-7](#)
 DSP_Family [4-9](#)
 DSP_Product [4-10](#)
 EventBitID [4-12](#)
 EventData [4-14](#)
 EventID [4-13](#)
 HistoryEnum [4-15](#)
 IMASKStruct [4-17](#)
 IOID [4-18](#)
 IOTemplateID [4-19](#)
 MessageID [4-20](#)
 MsgChannel [4-21](#)
 PoolID [4-23](#)
 Priority [4-24](#)
 SemaphoreID [4-25](#)
 SystemError [4-26](#)
 ThreadID [4-31](#)
 ThreadStatus [4-32](#)
 Ticks [4-34](#), [A-5](#)
 VersionStruct [4-35](#)
debugged control structures [1-2](#)
debugging VDK projects [2-3](#)

- general tips 2-5
 - DestroyDeviceFlag() 5-34
 - DestroyMessage() 5-35
 - DestroyPool() 5-37
 - DestroySemaphore() 5-39
 - DestroyThread() 3-5, 3-14, 5-41
 - destructor functions, threads 3-5
 - device drivers 1-10, 3-38, A-2, A-13
 - activate functions 3-51
 - dispatch functions 3-37, 3-39, 3-43, 3-52
 - execution of 3-39
 - general notes 3-53
 - init functions 3-47
 - initializing 3-39
 - interacting with ISRs 3-39
 - interfacing with threads 3-42
 - migrating to VisualDSP++ 3.0 B-1
 - open/close by threads 3-47
 - using 3-42
 - variables in 3-54
 - device flags 3-32, 3-51, B-2
 - pending on 3-51
 - posting 3-32, 3-52
 - DeviceDescriptor data type 4-4
 - DeviceFlagID data type 4-5
 - DeviceIOCtl() 3-42, 3-49, 5-43
 - disabling
 - interrupts 1-8, 3-33
 - scheduling 1-7, 3-11
 - dispatch functions
 - posting device flags 3-52
 - See also device drivers
 - DispatchID data type 4-6
 - DispatchThreadError() 3-4, 5-45
 - DispatchUnion data type 4-7, B-4
 - DispatchUnion, dispatch function
 - parameter 3-44
 - domains
 - interrupt domain 3-18
 - thread domain 3-18
 - DSP_Family data type 4-9
 - DSP_Product data type 4-10
 - dynamic threads, creating 3-4
- E
- entering scheduler
 - from API calls 3-13
 - from interrupts 3-13
 - enumeration
 - error codes 3-8
 - error functions 3-3
 - priorities of threads 3-3
 - error codes 3-8
 - error functions, threads 3-3
 - error handling facilities 3-8
 - ErrorFunction() 3-8
 - ErrorHandler() A-8, A-13
 - event bits 3-25
 - changing status from interrupt domain 3-30
 - changing status from thread domain 3-30
 - global state 3-26
 - EventBitID data type 4-12
 - EventData data type 4-14
 - EventID data type 4-13
 - events

INDEX

- behavior of [3-26](#)
- calculating [3-26](#), [3-27](#)
- combination of event bits [3-25](#)
- loading new event data [3-31](#)
- mask flag [3-26](#)
- matchAll flag [3-26](#)
- number of in a system [3-25](#)
- recalculating [3-31](#)
- setting up [3-26](#)
- state (TRUE or FALSE) [3-25](#)
- triggering [3-31](#)
- values flag [3-26](#)
- VDK_EventData [3-26](#)
- EVT_EVSW, hardware interrupt [A-4](#)
- EVT_IVG14, hardware interrupt [A-4](#)
- EVT_IVG15, hardware interrupt [A-4](#)
- EVT_TMR, hardware interrupt [A-4](#)
- execution levels [A-2](#), [A-8](#), [A-13](#)
 - Interrupt Level [A-2](#), [A-8](#), [A-13](#), [A-16](#)
 - Kernel Level [A-2](#), [A-8](#), [A-13](#), [A-16](#)
 - Thread Level [A-2](#), [A-8](#), [A-13](#), [A-16](#)
- execution modes
 - supervisor [A-1](#)
 - user [A-1](#)
- EXTERN, assembly directive [3-7](#)

F

- FreeBlock() [5-46](#)
- FreeDestroyedThreads() [5-48](#)
- freeing
 - memory with free/delete [3-5](#)
 - thread resources [3-5](#)
- FreeThreadSlot() [5-49](#)

G

- GetEventBitValue() [5-51](#)
- GetEventData() [3-31](#), [5-52](#)
- GetEventValue() [5-53](#)
- GetInterruptMask() [3-33](#), [5-54](#)
- GetLastThreadError() [3-8](#)
- GetLastThreadErrorValue() [3-8](#), [5-56](#)
- GetMessagePayload() [5-57](#)
- GetMessageReceiveInfo() [3-23](#), [5-59](#)
- GetNumAllocatedBlocks() [5-61](#)
- GetNumFreeBlocks() [5-62](#)
- GetPriority() [5-63](#)
- GetSemaphoreValue() [5-64](#)
- GetThreadHandle() [3-7](#), [5-65](#)
- GetThreadID() [3-2](#), [5-66](#)
- GetThreadSlotValue() [5-67](#)
- GetThreadStackUsage() [5-68](#)
- GetThreadStatus() [5-69](#)
- GetThreadType() [5-70](#)
- GetUptime() [5-71](#)
- GetVersion() [5-72](#)
- global
 - data [3-35](#)
 - variables [3-7](#)
- global data [A-6](#)
- GLOBAL, assembly directive [3-7](#)

H

- hardware abstraction [1-3](#)
- header file (vdk.h) [2-2](#)
- header files for C/assembly threads [3-7](#)
- history
 - buffers [2-3](#)
 - logs [2-3](#)

HistoryEnum data type [4-15](#)

I

I/O

- interface [3-38](#), [B-1](#)
- migrating device drivers [3-38](#)
- templates [3-38](#)

idle thread [2-4](#), [3-5](#), [3-13](#)

IMASKStruct data type [4-17](#)

init functions

- C++ constructor [3-5](#)
- C/assembly [3-4](#)
- device drivers [3-47](#)
- threads [3-4](#)

InitFunction(), C/assembly threads [3-5](#)

instrumented build info [2-3](#)

Interrupt [A-16](#)

interrupt service routines (ISRs) [1-10](#)

interrupt vector table [3-34](#)

interrupts [1-10](#), [3-33](#)

- ADSP-2106x DSPs [A-17](#)
- ADSP-2116x DSPs [A-18](#)
- ADSP-2191 DSPs [A-10](#)
- ADSP-2192 DSPs [A-9](#)
- ADSP-TS101 DSPs [A-14](#)
- domain [3-18](#), [3-40](#)
- handler [3-34](#)
- handling [3-34](#)
- latching [3-18](#)
- latency [3-34](#), [A-8](#), [A-12](#), [A-16](#), [A-23](#)
- nesting [3-34](#), [A-7](#), [A-11](#), [A-15](#), [A-22](#)

invoking scheduler

- from APIs [3-13](#)
- from ISRs [3-13](#)

IOCtl_t struct [3-49](#)

- command [3-50](#)
- data [3-50](#)
- parameters [3-50](#)
- timeout [3-50](#)

IOID data type [3-51](#), [4-18](#)

IOTemplateID data type [4-19](#)

ISRs

- activating device drivers [3-50](#)
- architecture [3-34](#)
- assembly implemented [3-33](#)
- communicating with thread domain [3-36](#)
- disabling (masking) [3-33](#)
- enabling (unmasking) [3-33](#)
- re-enabling [3-33](#)

K

kernel interrupt [A-9](#), [A-10](#), [A-14](#)

kIO_Activate [3-44](#), [3-53](#), [A-2](#), [A-9](#), [A-13](#)

kIO_Close [3-44](#), [3-45](#)

kIO_Init [3-44](#)

kIO_IOCTL [3-44](#), [3-45](#)

kIO_Open [3-44](#), [3-45](#)

kIO_SyncRead [3-44](#), [3-45](#)

kIO_SyncWrite [3-44](#), [3-45](#)

L

LDF (Linker Description File) [2-2](#)

LDFs [A-6](#)

library

- documented functions [5-4](#)
- format [5-10](#)

INDEX

- working with headers [5-3](#)
- linking
 - API functions [5-2](#)
 - thread safe libraries [2-2](#)
- LoadEvent() [3-31](#), [5-73](#)
- LocateAndFreeBlock() [5-75](#)
- LogHistoryEvent() [5-76](#)
- lowest priority interrupts [3-35](#), [3-36](#),
[3-37](#), [3-50](#)

M

- MakePeriodic() [5-77](#)
- MallocBlock() [5-79](#)
- memory
 - AD6532 processor [A-6](#)
 - ADSP-219x DSP [A-11](#)
 - ADSP-21xxx DSPs [A-19](#)
 - ADSP-BF531 processor [A-6](#)
 - ADSP-BF532 processor [A-6](#)
 - ADSP-BF533 processor [A-6](#)
 - ADSP-BF535 processor [A-6](#)
 - ADSP-TS101 processor [A-15](#)
 - DM102 processor [A-6](#)
- memory allocations
 - alternative to malloc [3-55](#)
 - with malloc/new [3-5](#)
- memory blocks
 - addresses of [3-56](#)
 - on demand creation [3-55](#)
- memory pools [3-55](#)
- MessageAvailable() [5-81](#)
- MessageID data type [4-20](#)
- messages [3-21](#)
 - behavior of [3-21](#)

- interacting with threads [3-22](#)
- ownership [3-22](#)
- pending on [3-23](#)
- posting from scheduled regions [3-24](#)
- posting from unscheduled regions
[3-25](#)
- priority of [3-21](#)
- motivation for using VDK [1-1](#)
- MsgChannel data type [4-21](#)
- multitasking [3-10](#)

N

- namespaces
 - VDK [3-6](#), [5-4](#)
- notation conventions, API library [5-3](#)

O

- open/close functions, threads [3-47](#)
- OpenClose_t struct [3-47](#)
 - dataH [3-48](#)
 - flags [3-48](#)
- OpenDevice() [3-47](#), [3-48](#), [5-83](#)

P

- parallel scheduling domains [3-40](#)
- partitioning applications [1-4](#)
- passing function parameters [5-3](#)
- payloads, message associated [3-21](#)
- PendDeviceFlag() [3-32](#), [5-85](#)
- PendEvent() [3-29](#), [5-87](#)
- pending
 - on device flags [3-51](#)
 - on semaphores [3-17](#)
- PendMessage() [3-23](#), [5-89](#)

- PendSemaphore() 3-17, 5-92
 - periodic semaphores 3-16, 3-20
 - PoolID data type 4-23
 - PopCriticalRegion() 3-33, 5-94
 - PopNestedCriticalRegions() 5-96
 - PopNestedUnscheduledRegions() 3-12, 5-98
 - PopUnscheduledRegion() 3-11, 3-27, 5-99
 - porting existing device drivers B-2
 - PostDeviceFlag() 3-32, 3-51, 5-101
 - PostMessage() 5-102
 - PostSemaphore() 3-18, 3-20, 5-104
 - preemption 1-7
 - preemptive scheduling 3-11
 - priorities of threads 1-5, 3-3
 - changing dynamically 3-3
 - default 3-3
 - highest 3-3
 - Priority data type 4-24
 - protected regions 1-7, 3-7, 3-8, 3-49
 - nested 3-12
 - notes on 3-54
 - PushCriticalRegion() 3-33, 5-106
 - PushUnscheduledRegion() 3-11, 5-107
- Q**
- queue of ready threads, See ready queue
 - queue of waiting threads 3-3
- R**
- RAD (Rapid Application Development) 1-2
 - ReadWrite_t struct 3-48
 - data 3-49
 - data size 3-49
 - handle 3-48
 - timeout 3-49
 - ready queue 1-5, 3-9, 3-11, 3-18, 3-20, 3-25, 3-32
 - recalculating events 3-27
 - re-enabling interrupts 3-33
 - reentrancy 3-7
 - RemovePeriodic() 5-108
 - reschedule ISRs 3-37
 - ResetPriority() 3-3, 5-110
 - round-robin scheduling 3-10
 - run functions, threads 3-3
 - Run(), C++ threads 3-3
 - RunFunction(), C/assembly threads 3-3
- S**
- scheduled regions 3-24, 3-31
 - scheduler 3-18, 3-31
 - scheduling 1-5, 3-8
 - disable 3-11
 - enable 3-20
 - pending on events 3-26
 - pending on semaphores 3-18, 3-26
 - periodic semaphores 3-20
 - scheduling methods/modes 3-10
 - cooperative 3-10
 - preemptive 3-11
 - round-robin 3-10
 - scope, thread data members 3-6
 - sections

INDEX

- data1 [A-11](#), [A-15](#)
- data2 [A-15](#)
- program [A-11](#), [A-15](#)
- SemaphoreID data type [4-25](#)
- semaphores [3-15](#)
 - behavior of [3-16](#)
 - pending on [3-16](#)
 - periodic [3-16](#)
 - posting from interrupt domains [3-18](#)
 - posting from thread domains [3-18](#)
 - posting of [3-16](#)
- SetEventBit() [3-27](#), [3-30](#), [5-111](#)
- SetInterruptMaskBits() [3-33](#), [5-113](#)
- SetMessagePayload() [5-114](#)
- SetPriority() [3-3](#), [5-116](#)
- SetThreadError() [5-118](#)
- SetThreadSlotValue() [5-119](#)
- SHARC DSPs
 - 40-bit registers [A-21](#)
 - ADSP-2106x DSP ISRs [A-17](#)
 - ADSP-2116x DSP ISRs [A-18](#)
 - alternate registers [A-20](#)
 - enabling/disabling interrupts [A-17](#)
 - execution levels [A-16](#)
 - interrupt latency [A-23](#)
 - interrupt nesting [A-22](#)
 - memory [A-19](#)
 - system stack [A-19](#)
 - timer [A-19](#)
- signal.h [3-33](#), [3-34](#)
- signals [3-15](#)
 - device flags, See device flags
 - events, See events and event bits
 - semaphores, See semaphores
- Sleep() [5-120](#)
- software interrupts [3-18](#)
- source templates [3-4](#)
- stack
 - overflows [3-2](#)
 - size of threads [3-2](#)
- stack system (ADSP-21xxx DSPs) [A-19](#)
- standard C/C++ libraries, thread safe
 - versions [2-2](#)
- state of the system [3-25](#)
- storing event bits state in global
 - variables [3-26](#)
- synchronization, threads [3-15](#)
 - device flags [3-32](#)
 - events and event bits [3-25](#)
 - semaphores [3-15](#)
- SyncRead() [3-48](#), [5-122](#)
- SyncWrite() [3-48](#), [5-124](#)
- SystemError data type [4-26](#)
-
- T
 - thread domain [3-18](#), [3-36](#), [3-40](#)
 - calling dispatch function [3-45](#)
 - software scheduling [1-9](#)
 - thread safe libraries [2-2](#)
 - ThreadId [3-2](#)
 - ThreadId data type [4-31](#)
 - threads [3-1](#)
 - assembly implemented [3-7](#)
 - blocking [1-7](#), [3-32](#)
 - C implemented [3-7](#)
 - C++ implemented [3-6](#)
 - C++ template [3-6](#)
 - controlling device parameters [3-49](#)

- create functions [3-3](#)
- creating a message [3-21](#)
- destructor functions [3-5](#)
- dynamic creation with malloc/new [3-5](#)
- entry point [3-3](#)
- error functions [3-3](#), [3-17](#)
- hardware interaction [1-8](#)
- header file [3-4](#)
- interacting with device flags [3-32](#)
- interacting with events [3-29](#)
- interacting with semaphores [3-16](#)
- lowest priority (idle) [2-4](#)
- open/close device drivers [3-47](#)
- pending on device flags [3-32](#)
- pending on events [3-25](#), [3-29](#)
- pending on messages [3-22](#)
- posting messages [3-21](#), [3-24](#)
- posting semaphores [3-18](#)
- priority, See priorities of threads
- reading to /writing from [3-48](#)
- required functionality [3-3](#)
- Run() functions [3-3](#)
- scope of [3-5](#)
- scope of data members [3-6](#)
- setting/clearing event bits [3-30](#)
- stack size [3-2](#)
- state and status data [2-4](#)
- types of [3-2](#)
- ThreadStatus data type [4-32](#)
- Ticks data type [4-34](#)
- TigerSHARC DSPs
 - enabling/disabling interrupts [A-14](#)
 - execution levels [A-13](#)
 - interrupt latency [A-16](#)
 - interrupt nesting [A-15](#)
 - memory [A-15](#)
 - memory blocks [A-15](#)
- time slicing [3-10](#)
- timed system events [3-37](#)
- timeout [3-16](#), [3-17](#)
- timer
 - ADSP-219x DSPs [A-10](#)
 - ADSP-21xxx DSPs [A-19](#)
 - ADSP-BF53x processor [A-5](#)
 - ADSP-TSxxx DSPs [A-14](#)
- Timer ISR [A-22](#)
- timer ISR [3-35](#), [3-37](#), [A-9](#), [A-10](#), [A-14](#)
- triggering multiple events [3-28](#)
- U
- unscheduled regions [3-8](#), [3-11](#), [3-25](#), [3-28](#), [3-31](#), [3-33](#), [3-47](#), [3-54](#)
- user code [A-6](#)
- using, C++ keyword [3-6](#)
- V
- variables [3-54](#)
- VDK [3-50](#)
 - Communication Manager [3-38](#)
 - device drivers conversion [B-2](#)
 - namespace [3-6](#), [5-4](#)
 - State History window [2-3](#)
 - target load graph [2-4](#)
- VDK library
 - assembly macros [5-127](#)
 - function list [5-4](#)
 - header (vdk.h) [5-3](#)

INDEX

- linking library functions [5-2](#)
 - naming conventions [5-3](#)
 - reference format [5-10](#)
 - VDK project
 - Linker Description File [2-2](#)
 - vdk.h header file [2-2](#)
 - VDK Status window [2-4](#)
 - VDK_DispatchUnion [3-48](#), [3-49](#)
 - VDK_EventData [3-26](#), [3-29](#), [3-31](#)
 - VDK_ISR_ACTIVATE_DEVICE_[3-50](#), [5-129](#), [A-2](#), [A-9](#)
 - VDK_ISR_CLEAR_EVENTBIT_[3-30](#), [5-130](#)
 - VDK_ISR_LOG_HISTORY_EVENT_[5-131](#)
 - VDK_ISR_POST_SEMAPHORE_[3-18](#), [3-20](#), [3-36](#), [5-132](#)
 - VDK_ISR_SET_EVENTBIT_ [3-30](#), [5-133](#)
 - VDK_kIO_Activate [3-50](#)
 - VDK_kIO_Close [3-50](#)
 - VDK_kIO_Init [3-47](#)
 - VDK_kIO_IOCTL [3-49](#)
 - VDK_kIO_Open [3-47](#)
 - VDK_kIO_SyncRead [3-48](#)
 - VDK_kIO_SyncWrite [3-48](#)
 - vector table [3-34](#)
 - VersionStruct data type [4-35](#)
 - VisualDSP++
 - migrating device drivers [B-1](#)
 - state history graph [2-3](#)
 - System State History window [2-3](#)
 - Target Load graph [2-4](#)
 - Thread History window [3-14](#)
 - volatile variables [3-54](#)
- W**
- working with library headers [5-3](#)
 - writing threads in different languages [3-6](#)
- Y**
- Yield() [3-10](#), [5-126](#)