



Reconfigurable Application-Specific Computing User's Guide

007-4718-007

COPYRIGHT

© 2004, 2005, 2006, 2007, 2008, SGI. All rights reserved; provided portions may be copyright in third parties, as indicated elsewhere herein. No permission is granted to copy, distribute, or create derivative works from the contents of this electronic documentation in any manner, in whole or in part, without the prior written permission of SGI.

LIMITED RIGHTS LEGEND

The software described in this document is “commercial computer software” provided with restricted rights (except as to included open/free source) as specified in the FAR 52.227-19 and/or the DFAR 227.7202, or successive sections. Use beyond license provisions is a violation of worldwide intellectual property laws, treaties and conventions. This document is provided with limited rights as defined in 52.227-14.

TRADEMARKS AND ATTRIBUTIONS

SGI, the SGI logo, and Altix are registered trademarks and NUMAflex, NUMAlink, RASC, and SGI ProPack are trademarks of SGI in the United States and/or other countries worldwide.

Impulse is a trademark of Impulse Accelerated Technologies, Inc. Intel and Itanium are registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. Linux is a registered trademark of Linus Torvalds, used with permission by Silicon Graphics, Inc. Mitrionics is a trademark of Mitrionics, Inc. Red Hat and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries. Synplicity and Synplify Pro are registered trademarks of Synplicity, Inc. UNIX is a registered trademarks of the Open Group in the United States and other countries. Verilog is a registered trademark of Cadence Design Systems, Inc. Windows is a registered trademark of Microsoft Corporation in the United States and/or other countries. Xilinx is a registered trademark of Xilinx, Inc.

All other trademarks mentioned herein are the property of their respective owners.

Record of Revision

Version	Description
001	September 2004 Original publication
002	November 2004 Updated to support latest version of RASC software.
003	July 2005 Updated to support a new version of the RASC-brick with increased SRAM memory and the RASC 1.0 software release.
004	March 2006 Updated to support the blade version of SGI RASC Technology.
005	January 2007 Updated to support the RASC 2.1 software release.
006	March 2007 Updated with additional information in support of the RASC 2.1 software release.
007	January 2008 Updated with additional information in support of the RASC 2.20 software release.

New Features in This Guide

This revision of the *Reconfigurable Application-Specific Computing User's Guide* supports the RASC 2.20 software release.

Major Technical Changes in the RASC 2.20 Release

With the RASC 2.20 release, the following features of Core Services are now configurable:

- Streaming Direct Memory Access (DMA)

SGI Core Services provide four input and four output streaming DMA engines that allow data to transfer directly into the FPGA. These are individually selectable and you may select them in any combination. A new `strm_in_<n>_complete` signal also allows you to input data in chunks so you do not have to read or write an entire DMA length. You can use the new SGI RASC Algorithm Configuration tool described in Chapter 8, “SGI RASC Algorithm Configuration Tool” to configure these streams. For more information, see “Streaming DMA Configuration” on page 158.

- Memory

Three memory configurations instead of just one are available. As in prior releases, two 64-bit memory ports are paired up to be used as a 128-bit memory with a fifth 64-bit memory port available. The RASC 2.20 release provides a new configuration that enables five 64-bit memory ports. All five of these ports are independent and can be configured individually. Each memory port can be configured to only allow access by the algorithm to reduce logic overhead. You can also choose to have no external memory ports at all. Memory configuration is described in “Memory Configuration” on page 159.

- Algorithm Defined Registers (ADRs)

The SGI Core Services provides up to sixty-four (64) 64-bit registers to handle the communication between the host program and the algorithm. The number of registers is now selectable. For more information, see “Algorithm Defined Registers” on page 161. The algorithm defined registers can also be further specified for usage, as described in “Algorithm Defined Register Configuration” on page 161.

- Debug Registers

Core Services provides up to sixty-four (64) 64-bit debug registers. The number of registers is now selectable. For more information, see “Debug Registers” on page 162.

Core Services also now provide the following:

- SGI generated implementation scripts that support two synthesis tools: Synplify Pro and Xilinx XST, as described in “Synthesis Tool Specification” on page 163.
- A supplemental algorithm clock, as described in “Supplemental Algorithm Clock” on page 164.
- The new SGI RASC Algorithm Configuration tool described in Chapter 8, “SGI RASC Algorithm Configuration Tool” that enables you to configure memory, Streaming DMA, ADRs, and debug registers.

Major Documentation Changes

Changes in this guide for this release include the following:

- Added a second DMA engine in each direction for SRAMs, as described in “Core Services Architecture Overview” on page 23.
- Added information about the application algorithm secondary clock feature in “Supplemental Algorithm Clock” on page 63.
- Added information about a new `RASCLIB_IGNORE_PARITY` flag in Chapter 4, “RASC Abstraction Layer”.
- Added descriptions for the `rasclib_algorithm_malloc`, `rasclib_algorithm_mfree`, `rasclib_cop_malloc`, and `rasclib_cop_mfree` functions in Chapter 4, “RASC Abstraction Layer”.
- Added information about scaling applications with RASC in Chapter 6, “Scaling Applications with RASC”.

- Added information about a new SGI RASC algorithm configuration tool in Chapter 8, “SGI RASC Algorithm Configuration Tool”.
- More user control of Core Services local memory configurations
- Addition of stream completion signalling mechanisms
- Improved error correction/detection for local memories
- Padding out incomplete cachelines from FPGA to host

Contents

Figures	xix
Tables	xxi
Related Publications	xxiv
SGI Documentation	xxiv
Additional Documentation Sites and Useful Reading	xxvi
Obtaining SGI Publications	xxviii
Conventions	xxviii
Reader Comments	xxviii
1. RASC Introduction	1
An Overview of Reconfigurable Computing	1
Silicon Graphics ccNUMA Systems	3
Silicon Graphics Reconfigurable Application-Specific Computing (RASC)	4
Getting Started with RASC Programming	5
Overview of FPGA Programming	5
SGI FPGA Programming Approach	6
Bitstream Development Overview	6
Helpful SGI Tools	9
RASC Hardware Overview	9
RASC Software Overview	12
2. Altix System Overview	15
SGI Altix 350 System Overview	15
SGI Altix 450 and Altix 4700 System Overview	17
3. RASC Algorithm FPGA Hardware Design Guide	21
RASC Core Services Overview	22
Core Services Features	22
Core Services Architecture Overview	23

SRM23
SXM24
Programmed Input/Output Request Engine24
TNUM Tracker24
Request Gate24
Memory Mapped Registers24
Interrupt Generator24
Input Direct Memory Access24
Output Direct Memory Access24
Memory Controller25
Algorithm Block25
Algorithm Run Modes25
Algorithm Interfaces26
Algorithm Control Interface26
Algorithm Defined Registers27
Algorithm Defined Register Configuration28
Algorithm Defined Register Usage29
Streaming Direct Memory Access29
Input Streaming DMA Engine30
Start Input Stream30
End input stream31
Output Stream32
Start Output Stream33
Flushing Output Stream34
Streaming DMA Extractor Statements34
Debug Register Interface35
SRAM Interface35
Arbitration35
Handshaking Methodologies37
Using Busy Signal37
Using SRAM Crediting Scheme37
SRAM Configuration37
Address Offsets38

SRAM Read Operation	39
SRAM Read with Busy Operation	40
SRAM Write with Busy Operation	40
Algorithm Design Details.	42
Basic Algorithm Control	42
Recommendations for Memory Distribution	45
Input and Output Placement	45
Implementation Options for Debug Mode.	46
Clock Cycle Step Size Mode	46
Variable Step Size Mode.	46
External Memory Write Transaction Control	47
Example Write Transaction Timing Diagram.	47
External Memory Read Transaction Control	48
Example Read Transaction Timing Diagram	48
Designing an Algorithm for Multibuffering	49
Purpose	49
Definitions	50
Hardware Support	50
Software Responsibilities	52
Passing Parameters to Algorithm Block	52
Small Parameters.	52
Parameter Arrays	53
Recommended Coding Guidelines for Meeting Internal Timing Requirements	54
Connecting Internal Signals to the Debugger	54
RASC FPGA Design Integration	55
Design Hierarchy	56
FPGA Clock Domains	57
Core Clock Domain	57
Algorithm Clock Domain	57
SSP Clock Domain	59
QDR-II SRAM Clock Domains	59
supplemental Algorithm Clock.	60
Resets	60

Algorithm Synthesis-time Parameters 60
Setting the Primary Algorithm Clock Speed 61
SRAM Port Usage. 61
Supplemental Algorithm Clock 63
Setting the supplemental Algorithm Clock Speed 63
Special Case Frequencies. 64
General Case Frequencies 65
Simulating the Design 66
Intent of the Sample Test Bench 66
Sample Test Bench Setup 66
Compiling the Sample Test Bench 67
Running a Diagnostic. 68
Viewing Waveform Results 70
Writing a Diagnostic 71
Sample Test Bench Constants and Dependencies 76
Sample Test Bench Utilities 78
. 79
4. RASC Abstraction Layer 81
RASC Abstraction Layer Overview 81
RASC Abstraction Layer Calls 83
raslib_resource_reserve Function 86
raslib_resource_release Function 88
raslib_resource_configure Function 89
raslib_resource_return Function 90
raslib_resource_alloc Function 91
raslib_resource_free Function 92
raslib_algorithm_malloc Function 93
raslib_algorithm_mfree Function 94
raslib_algorithm_open Function 94
raslib_algorithm_send Function 96
raslib_algorithm_get_num_cops Function. 97
raslib_algorithm_receive Function 97
raslib_algorithm_go Function 99

rasclib_algorithm_commit Function	99
rasclib_algorithm_wait Function	100
rasclib_algorithm_close Function	102
rasclib_algorithm_reg_multi_cast Function	103
rasclib_algorithm_reg_read Function	104
rasclib_algorithm_reg_write Function	106
rasclib_algorithm_exception_handler_register Function	108
rasclib_cop_open Function	109
rasclib_cop_send Function	110
rasclib_cop_receive Function	111
rasclib_cop_malloc Function	113
rasclib_cop_mfree Function	113
rasclib_cop_go Function	114
rasclib_cop_commit Function	116
rasclib_cop_wait Function	118
rasclib_cop_close Function	119
rasclib_cop_reg_read Function	120
rasclib_cop_reg_write Function	121
rasclib_cop_exception_handler_register Function	123
rasclib_error Function	124
rasclib_error_dump Function	125
How the RASC Abstraction Layer Works	125
5. Direct I/O	129
Direct I/O Overview	129
Direct I/O Driver Requirements	129
Using Hugepages with rasclib Functions	130
6. Scaling Applications with RASC	133
Data Buffer Size	133
Achieving Maximal Bandwidth	134
Conclusion	135
7. RASC Algorithm FPGA Implementation Guide	137
Implementation Overview	138

Summary of the Implementation Flow	138
Supported Tools and OS Versions	139
Installation and Setup	140
SGI Altix System Installation	140
PC Installation	140
Implementation Constraint Files	142
Synthesis Using Synplify Pro	142
Synthesis Using XST	142
ISE (User Constraint File)	143
Adding Extractor Directives to the Source Code	144
Inserting Extractor Comments	145
Example of Comments in a Verilog, VHDL, or Header File	146
Implementation with Pre-synthesized Core	148
Makefile.local Customizations	149
Synthesis Project Customization	149
Synplify Pro	150
XST	151
Makefile Targets	151
Full-chip Implementation	152
Makefile.local Customizations	152
Synthesis Project Customization	153
Implementation File Descriptions	153
8. SGI RASC Algorithm Configuration Tool	155
Running the RASC Algorithm Configuration Tool	156
Algorithm Name	156
Algorithm Clock Rate	157
Streaming DMA Configuration	158
Memory Configuration	159
Memory Port Specification	160
Algorithm Defined Registers	161
Algorithm Defined Register Configuration	161
Debug Registers	162
Synthesis Tool Specification	163

Supplemental Algorithm Clock164
Output File Locations164
9. Running and Debugging Your Application167
Loading the Bitstream.167
RASC Device Manager168
RASC Device Manager Overview169
RASC Device Manager Structure170
Using the Device Manager Command (devmgr)170
Add a Bitstream To the Bitstream registry171
Delete a Bitstream From the Bitstream registry171
List the Contents of a Bitstream registry171
Update an Algorithm in the Bitstream registry172
List the FPGAs in the Inventory172
Mark an FPGA as Available or Unavailable173
Turn Debugging On or Off174
Device Manager Load FPGA Command174
Device Manager Reload FPGA Command175
Device Manager Version Information176
Device Manager Server Command176
Using the Device Manager Server (devmgr_server) Command177
Device Manager Logging Facility177
Using the GNU Project Debugger (GDB)178
GDB Commands179
fpgaactive179
set fpga fpganum.180
fpgaregisters180
info fpga180
fpgastep181
fpgacont181
fpgatrace181
Registers181
Values and Stepping182
FPGA Run Status182

	Changes To GDB Commands	183
10.	RASC Examples and Tutorials	185
	System Requirements	185
	Prerequisites	186
	Tutorial Overview	186
	Simple Algorithm Tutorial	187
	Overview	187
	Application.	188
	Coding Techniques: Verilog	189
	Overview	189
	Integrating with Core Services	189
	Extractor Comments	191
	Coding Techniques: VHDL Algorithm	192
	Overview	192
	Integrating with Core Services	193
	Extractor Comments	193
	Compiling for Simulation.	194
	Building an Implementation	195
	Transferring to the Altix Platform	196
	Verification using GDB	197
	Data Flow Algorithm Tutorial	200
	Application.	200
	Loading the Tutorial	200
	Integrating with Core Services	202
	Extractor Comments	202
	Compiling for Simulation.	202
	Building an Implementation	203
	Transferring to the Altix Platform	204
	Verification Using GDB	204
	Streaming DMA Algorithm Tutorial	206
	Application.	207
	Loading the Tutorial	207
	Integrating with Core Services	209

Extractor Comments209
Compiling for Simulation209
Building an Implementation.211
Transferring to the Altix Platform211
Verification Using GDB212
A. Device Driver213
FPGA Core Services213
Control and Status Registers.213
Interrupts213
Driver Application Programming Interface (API)214
Input Direct Memory Access214
Output Direct Memory Access215
Function Control215
Example Use of Device Driver215
B. SSP Stub User's Guide223
Introduction to SSP Stub223
Recommended Reading223
Verification Environment and Testbench224
Verification Environment224
Sample Test Bench224
SSP Stub File Descriptions225
Compiling and Running a Test227
SSP Stub Commands229
Packet Commands230
Command Fields.230
Send Commands.232
Receive Commands233
Other Commands.235
Command Summary236
Comments237
Sample Diagnostic.237
Using the Stub243

C.	How Extractor Works	245
	Extractor Script	245
	Core Services Configuration File	246
	Algorithm Configuration File	251
	Index	255

Figures

Figure 1-1	Reconfigurable Computer	2
Figure 1-2	Bitstream Development	7
Figure 1-3	Bitstream Development with High-level Tools	8
Figure 1-4	RASC FPGA Functional Block Diagram	10
Figure 1-5	RASC Blade Hardware	12
Figure 1-6	RASC Software Overview	13
Figure 2-1	Example of SGI Altix 350 Rack Systems	16
Figure 2-2	Altix 4700 Blade, Individual Rack Unit, and Rack	19
Figure 3-1	Block Diagram of the RASC Core Services	23
Figure 3-2	Algorithm Defined Registers (ADR) Interface Usage	29
Figure 3-3	Start of Input Stream Transaction	31
Figure 3-4	End of Input Stream Transaction	31
Figure 3-5	Start of Output Stream Transaction	33
Figure 3-6	End of Output Stream Transaction	34
Figure 3-7	SRAM Read Operation.	40
Figure 3-8	SRAM Read with Busy.	40
Figure 3-9	SRAM Write with Busy Assertion	41
Figure 3-10	Example of a Continuous, Normal Mode Algorithm Run.	43
Figure 3-11	Hardware Accelerated Algorithm Design Flow	44
Figure 3-12	Clock Cycle Stepping Mode Example	46
Figure 3-13	Variable Step Size Mode Example	47
Figure 3-14	Single, and Multiple Write Commands	47
Figure 3-15	Single Read Transaction	48
Figure 3-16	Multiple Read Transaction	49
Figure 3-17	Instance Hierarchy of the RASC FPGA Design	56
Figure 3-18	Core Clock and Algorithm Clock Source	59
Figure 3-19	Sample vcdplus.vpd Waveform in Virsim	71

Figure 4-1	Abstraction Layer Software Block Diagram	82
Figure 7-1	RASC FPGA Implementation Flow	138
Figure 8-1	SGI RASC Algorithm Configuration Tool Welcome Screen . .	156
Figure 8-2	Algorithm Name Configuration	157
Figure 8-3	Algorithm Clock Rate Configuration	158
Figure 8-4	Streaming DMA Configuration	159
Figure 8-5	Memory Configuration Selection	160
Figure 8-6	Memory Port Specification	160
Figure 8-7	Number of Algorithm Defined Register Selection	161
Figure 8-8	Algorithm Defined Register Specification	162
Figure 8-9	Number of Debug Registers Selection	163
Figure 8-10	Synthesis Tool Specification	163
Figure 8-11	Supplemental Algorithm Clock Configuration	164
Figure 8-12	Overwrite of <code>alg_block_top.v</code> Verification	165
Figure 8-13	Write-Protect File Notification	166
Figure 10-1	Simple Algorithm for Verilog	188
Figure 10-2	Simple Algorithm for Verilog and VHDL	190
Figure 10-3	Simple Algorithm for Verilog and VHDL	192
Figure 10-4	Data Flow Algorithm	201
Figure 10-5	Streaming DMA Algorithm	208

Tables

Table 3-1	Algorithm Control Interface Signals	26
Table 3-2	ADR Signal Definitions	27
Table 3-3	Input Streaming DMA Engine Signal Definitions	30
Table 3-4	Output Stream Signal Definitions	32
Table 3-5	Debug Register Definition	35
Table 3-6	SRAM Interface Signal Definition	38
Table 3-7	Sample Testbench Algorithms and Commands.	70
Table 3-8	Files in the <code>sample_tb</code> directory	77
Table 4-1	Abstraction Layer Function Definitions - Summary	83
Table 7-1	Supported Implementation Tools	139
Table 7-2	Environment Variables Required for Bundle Environment	141
Table 7-3	Top Level Directory Descriptions	141
Table 7-4	Synthesis Constraint Files Provided	142
Table 7-5	Implementation Constraint Files Provided	143
Table 7-6	Extractor Comment Fields	146
Table 7-7	Common Makefile.local Variable Settings	149
Table 7-8	Makefile Targets	151
Table 7-9	Required Full Chip Makefile Variable Settings	153
Table 7-10	Commonly Referenced Files	153
Table A-1	Device Driver API System Calls	214
Table B-1	Packet Types used by SSP	230
Table C-1	Core Services Configuration File Fields	246
Table C-2	Algorithm Configuration File Fields	251

About This Guide

The SGI reconfigurable application-specific software computing (RASC) program delivers scalable, configurable computing elements for the SGI Altix family of servers and superclusters.

Note: Make sure that you read the section called “Major Technical Changes in the RASC 2.20 Release” on the “New Features in This Guide” page in the frontmatter of this manual.

This guide provides information about RASC and covers the following topics:

- Chapter 1, “RASC Introduction”
- Chapter 2, “Altix System Overview”
- Chapter 3, “RASC Algorithm FPGA Hardware Design Guide”
- Chapter 4, “RASC Abstraction Layer”
- Chapter 5, “Direct I/O”
- Chapter 6, “Scaling Applications with RASC”
- Chapter 7, “RASC Algorithm FPGA Implementation Guide”
- Chapter 8, “SGI RASC Algorithm Configuration Tool”
- Chapter 9, “Running and Debugging Your Application”
- Chapter 10, “RASC Examples and Tutorials”
- Appendix A, “Device Driver”
- Appendix B, “SSP Stub User’s Guide”
- Appendix C, “How Extractor Works”

Related Publications

Documents listed in this section contain additional information that might be helpful, as follows:

- “SGI Documentation” on page xxiv
- “Additional Documentation Sites and Useful Reading” on page xxvi

SGI Documentation

The following documentation is available for the SGI Altix family of servers and superclusters and is available from the online SGI Technical Publications Library:

- *SGI ProPack 5 for Linux Start Here*
Provides information about the SGI ProPack for Linux release including information about major new features, software installation, and product support.
- *SGI ProPack 5 for Linux Release Notes*
Provide the latest information about software and documentation in this release. The release notes are on the SGI ProPack for Linux Documentation CD in the `root` directory, in a file named `README.TXT`.
- *Linux Device Driver Programmer’s Guide- Porting to SGI Altix Systems*
Provides information on programming, integrating, and controlling drivers.
- *Porting IRIX Applications to SGI Altix Platforms: SGI ProPack for Linux*
Provides information about porting an application to the SGI Altix platform.
- *Message Passing Toolkit (MPT) User’s Guide*
Describes industry-standard message passing protocol optimized for SGI computers.
- *Performance Co-Pilot for IA-64 Linux User’s and Administrator’s Guide*
Describes the Performance Co-Pilot (PCP) software package of advanced performance tools for SGI systems running the Linux operating system.
- *Linux Configuration and Operations Guide*
Provides information on how to perform system configuration and operations for SGI ProPack servers.
- *Linux Resource Administration Guide*

Provides a reference for people who manage the operation of SGI ProPack servers and contains information needed in the administration of various system resource management features such as Comprehensive System Accounting (CSA), Array Services, CPU memory sets and scheduling, and the Cpuset System.

- *SGI Altix 350 System User's Guide*
Provides an overview of the Altix 350 system components, and it describes how to set up and operate this system.
- *SGI Altix 350 Quick Start Guide*
Guides a knowledgeable user through the installation, setup, and simple configuration of most SGI Altix 350 systems.
- *SGI Altix 3700 Bx2 User's Guide*
This guide provides an overview of the architecture and descriptions of the major components that compose the SGI Altix 3700 Bx2 family of servers. It also provides the standard procedures for powering on and powering off the system, basic troubleshooting information, and important safety and regulatory specifications.
- *SGI Altix 450 System User's Guide*
This guide provides an overview of the architecture and descriptions of the major components that compose the SGI Altix 450 system. It also provides the standard procedures for powering on and powering off the system, basic troubleshooting information, and important safety and regulatory specifications.
- *SGI Altix 4700 System User's Guide*
This guide provides an overview of the architecture and descriptions of the major components that compose the SGI Altix 4700 family of servers. It also provides the standard procedures for powering on and powering off the system, basic troubleshooting information, and important safety and regulatory specifications.
- *Silicon Graphics Prism Visualization System User's Guide*
Provides an overview of the Silicon Graphics Prism Visualization System components and it describes how to set up and operate this system.
- *SGIconsole 2.1 Start Here*
Provides an introduction to SGIconsole and information about setting up and configuring SGIconsole hardware and software.
- *Console Manager for SGIconsole Administrator's Guide*
Provides information about the Console Manager software graphical interface allows you to control multiple SGI servers, SGI partitioned systems, and large single-system image servers.
- *SGI L1 and L2 Controller Software User's Guide*

Describes how to use the L1 and L2 controller commands at your system console to monitor and manage the SGI Altix 3000 and SGI Altix 4000 family of servers and superclusters.

- *XFS for Linux Administration*
Describes XFS, an open-source, fast recovery, journaling filesystem that provides direct I/O support, space preallocation, access control lists, quotas, and other commercial file system features.
- *XVM Volume Manager Administrator's Guide*
Describes the configuration and administration of XVM logical volumes using the XVM Volume Manager.
- *Event Manager User Guide*
Provides information about the Event Manger application that collects event information from other applications. This document describes the Event Manager application, the application programming interface that you can use to access it, the procedures that you can use to communicate with it from another application, and the commands that you can use to control it.
- *Embedded Support Partner User Guide*
Provides information about using the Embedded Support Partner (ESP) software suite to monitor events, set up proactive notification, and generate reports. This revision of the document describes ESP version 3.0, which is the first version of ESP that supports the Linux operating system.
- *Linux Application Tuning Guide*
Provides information about tuning application programs on SGI Altix systems. Application programs include Fortran and C programs written on SGI Linux systems with the compilers provided by Intel.
- *SCSL User's Guide*
Provides information about the scientific libraries on SGI Altix systems and SGI IRIX systems. Topics include discussions about BLAS, LAPACK, and FFT routines.

Additional Documentation Sites and Useful Reading

The following sites and books may be useful:

- *Xilinx Development System Reference Guide*

This manual provides a bitstream generation workflow diagram and detailed description of all the files generated and used in the workflow and the tools that create and use these files. It is available at <http://www.xilinx.com/>. Click on the **Documentation** link. Under **Design Tools Documentation**, select **Software Manuals**. For the RASC 2.20 release, reference the **9.x Software Manuals**.

- *Verilog Quick Start, A Practical Guide to Simulation and Synthesis in Verilog*, Third Edition, James M. Lee, Copyright 2002 by Kluwer Academic Publishers
- *Verilog HDL Synthesis, A Practical Primer*, J. Bhasker, Copyright 1998 by Lucent Technologies and published by Star Galaxy Publishing.
- *A Verilog HDL Primer*, J. Bhasker, Second Edition, Copyright 1997, 1999 by Lucent Technologies and published by Star Galaxy Publishing.

Obtaining SGI Publications

You can obtain SGI documentation in the following ways:

- See the SGI Technical Publications Library at <http://docs.sgi.com>. Various formats are available. This library contains the most recent and most comprehensive set of online books, release notes, man pages, and other information.
- You can view man pages by typing `man <title>` on a command line.

Conventions

The following conventions are used throughout this publication:

Convention	Meaning
<code>command</code>	This fixed-space font denotes literal items such as commands, files, routines, path names, signals, messages, and programming language structures. Angle brackets <> are sometimes used in tables to save space.
<i>variable</i>	Italic typeface denotes variable entries and words or concepts being defined.
user input	This bold, fixed-space font denotes literal items that the user enters in interactive sessions. (Output is shown in nonbold, fixed-space font.)
<code>manpage(x)</code>	Man page section identifiers appear in parentheses after man page names.
GUI element	This font denotes the names of graphical user interface (GUI) elements such as windows, screens, dialog boxes, menus, toolbars, icons, buttons, boxes, fields, and lists.

Reader Comments

If you have comments about the technical accuracy, content, or organization of this document, contact SGI. Be sure to include the title and document number of the manual with your comments. (Online, the document number is located in the front matter of the manual. In printed manuals, the document number is located at the bottom of each page.)

You can contact SGI in any of the following ways:

- Send e-mail to the following address:
techpubs@sgi.com
- Contact your customer service representative and ask that an incident be filed in the SGI incident tracking system.
- Send mail to the following address:
Technical Publications
SGI
1140 East Arques Avenue
Sunnyvale , CA 94085-4602

SGI values your comments and will respond to them promptly.

RASC Introduction

This chapter provides an introduction to Reconfigurable Computing (RC) including challenges in design and implementation of RC systems, an introduction to SGI's system platform and an overview of SGI's Reconfigurable Application Specific Computing (RASC). It covers the following topics:

- "An Overview of Reconfigurable Computing" on page 1
- "Silicon Graphics ccNUMA Systems" on page 3
- "Silicon Graphics Reconfigurable Application-Specific Computing (RASC)" on page 4
- "Getting Started with RASC Programming" on page 5
- "RASC Hardware Overview" on page 9
- "RASC Software Overview" on page 12

Note: Make sure that you read the section called "Major Technical Changes in the RASC 2.20 Release" on the "New Features in This Guide" page in the frontmatter of this manual.

An Overview of Reconfigurable Computing

Reconfigurable computing is defined as a computer having hardware that can be reconfigured to implement application-specific functions (see Figure 1-1). The basic concept of a reconfigurable computer (RC) was proposed in the 1960s, but only in the last decade has an RC been feasible. RC systems combine microprocessors and programmable logic devices, typically field programmable gate arrays (FPGAs), into a single system. Reconfigurable computing allows applications to map computationally dense code to hardware. This mapping often provides orders of magnitude improvements in speed while decreasing power and space requirements.




Figure 1-1 Reconfigurable Computer

As defined above, RC uses programmable FPGAs. Current FPGA technology provides more than 10 million logic elements, internal clock rates over 200MHz, and pin toggle rates approaching 10Gb/s. With these large devices, the scope of applications that can target an FPGA has dramatically increased. The challenges to using FPGAs effectively fall into two categories: ease of use and performance. Ease of use issues include the following:

- Methodology of generating the “program” or bitstream for the FPGA
- Ability to debug an application running on both the microprocessor and FPGA
- Interface between the application and the system or Application Programming Interface (API)

Performance issues include the following:

- Data movement (bandwidth) between microprocessors and FPGAs
- Latency of communication between microprocessors and FPGAs
- Scalability of the system topology

Historically, programming an FPGA required a hardware design engineer. (Non-hardware designers can program FPGAs with the newer tools described in “Getting Started with RASC Programming” on page 5). Typically, an algorithm was hand translated into a Hardware Description Language (HDL), verified by human-generated HDL tests, synthesized into logic elements, physically placed in the FPGA and then analyzed for speed of operation. If errors occurred, these design and verification steps were repeated. This iterative process is conducive for semiconductor chip design, but it impedes the rapid proto-typing and solution oriented goals of application-specific programming.

Debugging FPGAs requires specialized logic to be inserted into register transfer level (RTL) code. Then, FPGA-specific tools, along with typical microprocessor debug methods, are used to analyze anomalies in application behavior.

Lastly, users are hampered by the lack of standardized application interfaces. This deficiency forces recoding when hardware or platform upgrades become available--a time-consuming and error-inducing process.

Performance is the fundamental reason for using RC systems. By mapping algorithms to hardware, designers can tailor not only the computational components, but also perform data flow optimization to match the algorithm. Today's FPGAs provide over a terabyte per second of memory bandwidth from small on-chip memories as well as tens of billions of operations per second. Transferring data to the FPGA and receiving the results poses a difficult challenge to RC system designers. In addition to bandwidth, efficient use of FPGA resources requires low latency communication between the host microprocessor and the FPGAs. When low latency is achieved, scaling and optimization across multiple computational elements can occur, often called load balancing.

Although challenges abound, RC systems allow users to explore solutions that are not viable in today's limited computing environments. The benefits in size, speed, and power alone make RC systems a necessity.

Silicon Graphics ccNUMA Systems

SGI was founded in 1982 based on Stanford University's research in accelerating a specific application, three dimensional graphics. SGI pioneered acceleration of graphics through hardware setting records and providing technological capabilities that were impossible without specialized computational elements. Tackling difficult problems required a supercomputer with capabilities that were not available from other computer vendors. SGI chose to develop its own large scale supercomputer with the features needed to drive graphics. The development of large scale single system image (SSI) machines was also pioneered at SGI. From the early days systems such as Power Series, Challenge, and Power Challenge defined large shared bus SSI systems, but the focus continued to be on providing the high bandwidth and scaling that was needed to drive graphics applications. To transcend the 36 microprocessor Challenge Series systems, conventional single backplane architectures were not sufficient. SGI returned again to its roots at Stanford University and the Distributed Architecture for Shared (DASH) memory project. The original concepts for cache coherent non-uniform memory access (ccNUMA) architecture are based on the DASH memory project. The ccNUMA

architecture allows for distributed memory through the use of a directory-based coherency scheme, removing the need for large common busses like those in the Challenge systems. Without the restrictions of a single bus, bandwidth for the system increased by orders of magnitude, while latency was reduced. This architecture has allowed SGI to set new records for system scalability including 1024 CPU SSI, 1TB/s Streams benchmark performance, and many others.

The SGI Altix system is the only fourth generation Distributed Shared Memory (DSM) machine using a NUMA architecture that is connected by a high-bandwidth, low-latency interconnect. In keeping with ever increasing demands, Altix allows independent scaling for CPUs, memory, Graphics Processing Units (GPU), I/O interfaces, and specialized processors. The NUMalink interconnect allows Altix to scale to thousands of CPUs, terabytes of memory, hundreds of I/O channels, hundreds of graphics processors and thousands of application-specific devices.

SGI uses NUMalink on all of its ccNUMA systems. NUMalink 4 is a third generation fabric that supports topologies starting at basic rings. With the addition of routers, meshes, hypercubes, and modified hypercubes, full fat tree topologies can be built. The protocol and fabric allows the topology to be matched to the workload as needed. By using the high bandwidth, low latency interconnect, SGI has a flexible and powerful platform to deliver optimized solutions for HPC.

Silicon Graphics Reconfigurable Application-Specific Computing (RASC)

The RASC program leverages more than 20 years of SGI experience accelerating algorithms in hardware. Rather than using relatively fixed implementations, such as graphics processing units (GPUs), RASC uses FPGA technology to develop a full-featured reconfigurable computer. The RASC program also addresses the ease of use and performance issues present in typical RC environments.

To address performance issues, RASC connects FPGAs into the NUMalink fabric making them a peer to the microprocessor and providing both high bandwidth and low latency. By attaching the FPGA devices to the NUMalink interconnect, RASC places the FPGA resources inside the coherency domain of the computer system. This placement allows the FPGAs extremely high bandwidth (up to 6.4GB/s/FPGA), low latency, and hardware barriers. These features enable both extreme performance and scalability. The RASC product also provides system infrastructure to manage and reprogram the contents of the FPGA quickly for reuse of resources.

RASC defines a set of APIs through the RASC Abstraction Layer (RASCAL). The abstraction layer can abstract the hardware to provide deep and wide scaling or direct and specific control over each hardware element in the system. In addition RASC provides a FPGA-aware version of GNU Debugger (GDB) that is based on the standard Linux version of GDB with extensions for managing multiple FPGAs. The RASC debug environment does not require learning new tool sets to quickly debug an accelerated application.

RASC supports the common hardware description languages (HDLs) for generating algorithms. RASC provides templates for Verilog- and VHDL-based algorithms. Several 3rd-party high-level language tool vendors are developing RASC interfaces and templates to use in place of traditional hardware design languages.

Getting Started with RASC Programming

This section provides an overview of RASC programming and covers these topics:

- “Overview of FPGA Programming” on page 5
- “SGI FPGA Programming Approach” on page 6
- “Bitstream Development Overview” on page 6
- “Helpful SGI Tools” on page 9

Overview of FPGA Programming

FPGA programming is a fairly complex task when using the main FPGA programming languages, VHDL and Verilog, directly. They require an electrical engineering background and the understanding of timing constraints; that is, the time it takes for an electrical signal to travel on the chips, delays introduced by buffers, and so on. For example, the blank FPGA has physical connection from the memory pins to FPGA I/O pins. You need a protocol like QDR-II that specifies memory transfer rates for the dual in-line memory modules connected to the memory pins.

Low-level abstractions allow an application to read a memory location without understanding the underlying hardware. SGI RASC calls this functionality Core Services (for more information, see “Algorithm Interfaces” on page 26). These low-level abstractions can almost be thought of as the basic input/output system (BIOS) of the RASC unit.

Currently, FPGAs are running at clock speeds of about 200 MHz. This may seem slow compared to an Itanium processor; however, the FPGA can be optimized for specific algorithms and potentially performs several hundreds or thousands of operations in parallel.

Programming high-performance computing applications in VHDL and/or Verilog is extremely time-consuming and resource-intensive and probably best left to very advanced users. However, high-level tools provided by vendors such as Mitronics, Inc. or Impulse Accelerated Technologies, Inc. are available.

These tools produce VHDL or Verilog code (potentially thousands of lines for even a small code fragment). This code then has to be synthesized (compiled) into a netlist. This netlist then is used by a place and route program to implement the physical layout on the FPGAs.

SGI FPGA Programming Approach

A summary of the SGI approach to FPGA programming is, as follows: (see Figure 1-2 on page 7 and Figure 1-3 on page 8)

- Write an application in C programming language for system microprocessor
- Identify computation intense routine(s)
- Generate a bitstream using Core Services and language of choice
- Replace routines with RASC abstraction layer (`rasc.lib`) calls that support both a C and Fortran90 interface
- Run your application and debug with GDB (see “Helpful SGI Tools” on page 9)

The RASC tutorial (see “Tutorial Overview” on page 186) steps you through the entire RASC design flow: integrating the algorithm with Core Services; simulating behavior on the algorithm interfaces; synthesizing the algorithm code; generating a bitstream; transferring that bitstream and metadata to the Altix platform; executing an application; and using GDB to debug an application on the Altix system and FPGA simultaneously.

Bitstream Development Overview

In this guide, *implementation flow* (bitstream development) refers to the comprehensive run of the extractor, synthesis, and Xilinx ISE tools that turn the Verilog or VHDL source

into a binary bitstream and configuration file that can be downloaded into the RASC Algorithm FPGA (for more information, see “Summary of the Implementation Flow” on page 138).

Figure 1-2 shows bitstream development on an x86 Linux platform for an Altix RASC hardware. The RASC Abstraction Layer (`rasc.lib`) provides an application programming interface (API) for the kernel device driver and the RASC hardware. It is intended to provide a similar level of support for application development as the standard `open/close/read/write/ioctl` calls for IO peripheral. For more on the RASC Abstraction Layer, see Chapter 4, “RASC Abstraction Layer”.

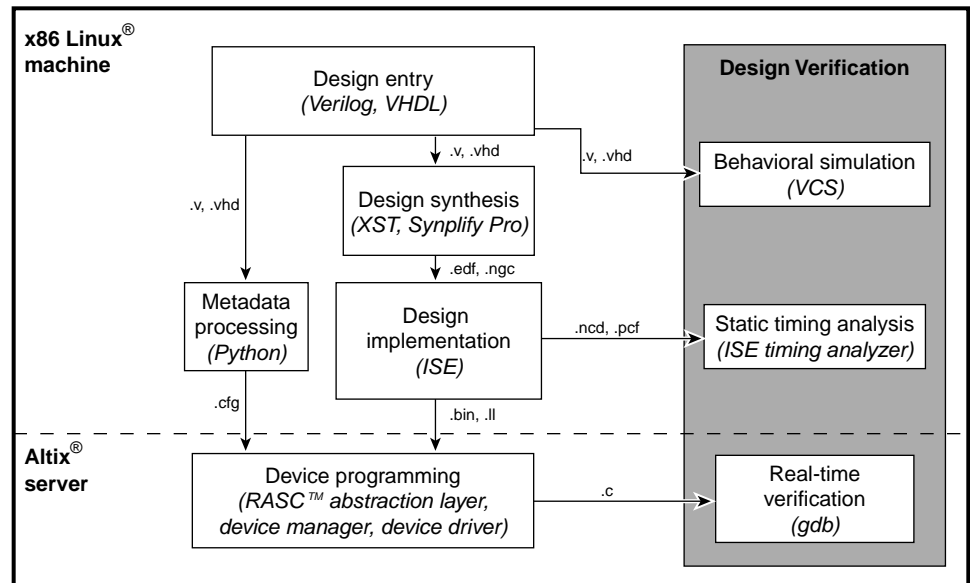


Figure 1-2 Bitstream Development

Figure 1-3 shows bitstream development using third-party development tools. (see “SGI FPGA Programming Approach” on page 6)

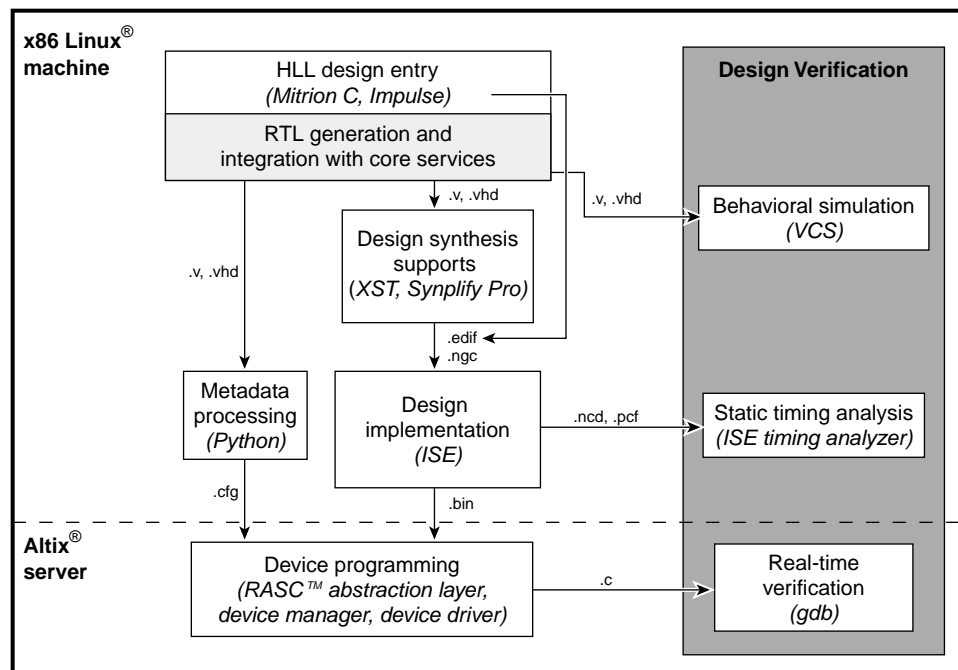


Figure 1-3 Bitstream Development with High-level Tools

Synplify Pro is a synthesis product developed by Synplicity, Inc. For more information, click on the **Literature** link on the top of the homepage at www.synplicity.com.

Xilinx Synthesis Technology (XST) is a synthesis product developed by Xilinx, Inc. Information on XST is available at <http://www.xilinx.com/>.

The *Xilinx Development System Reference Guide* provides a bitstream generation workflow diagram and detailed description of all the files generated and used in the workflow and the tools that create and use these files. From the Xilinx, Inc. homepage, Click on the **Documentation** link. Under **Design Tools Documentation**, select **Software Manuals**.

For the RASC 2.20 release, reference the **9.x Software Manuals**.

For additional documentation that you may find helpful, see "Additional Documentation Sites and Useful Reading" on page xxvi.

Helpful SGI Tools

SGI provides a Device Manager for loading bitstreams. It maintains a registry of algorithm bitstreams that can be loaded and executed using the RASC abstraction layer (`rasc.lib`). The `devmgr` user command is used to add, delete, and query algorithms in the registry. For more information on the Device Manager, see “RASC Device Manager” on page 168.

The SGI RASC environment has a product called `gdbfpga` that provides extensions to the GNU Debugger (GDB) command set to handle debugging of one or more FPGAs. An overview of GDB is, as follows:

- Based on Open Source GNU Debugger (GDB)
- Uses extensions to current command set
- Can debug host application and FPGA
- Provides notification when FPGA starts or stops
- Supplies information on FPGA characteristics
- Can "single-step" or "run N steps" of the algorithm
- Dumps data regarding the set of "registers" that are visible when the FPGA is active

For more information on debugging in the RASC environment, see “Using the GNU Project Debugger (GDB)” on page 178.

RASC Hardware Overview

The initial RASC hardware implementation used SGI’s first generation peer attached I/O brick for the base hardware. The RASC hardware module is based on an application-specific integrated circuit (ASIC) called TIO. TIO attaches to the Altix system NUMALink interconnect directly instead of being driven from a compute node using the XIO channel. TIO supports two PCI-X busses, an AGP-8X bus, and the Scalable System Port (SSP) port that is used to connect the Field Programmable Gate Array (FPGA) to the rest of the Altix system for the RASC program. The RASC module contains a card with the co-processor (COP) FPGA device as shown Figure 1-4.

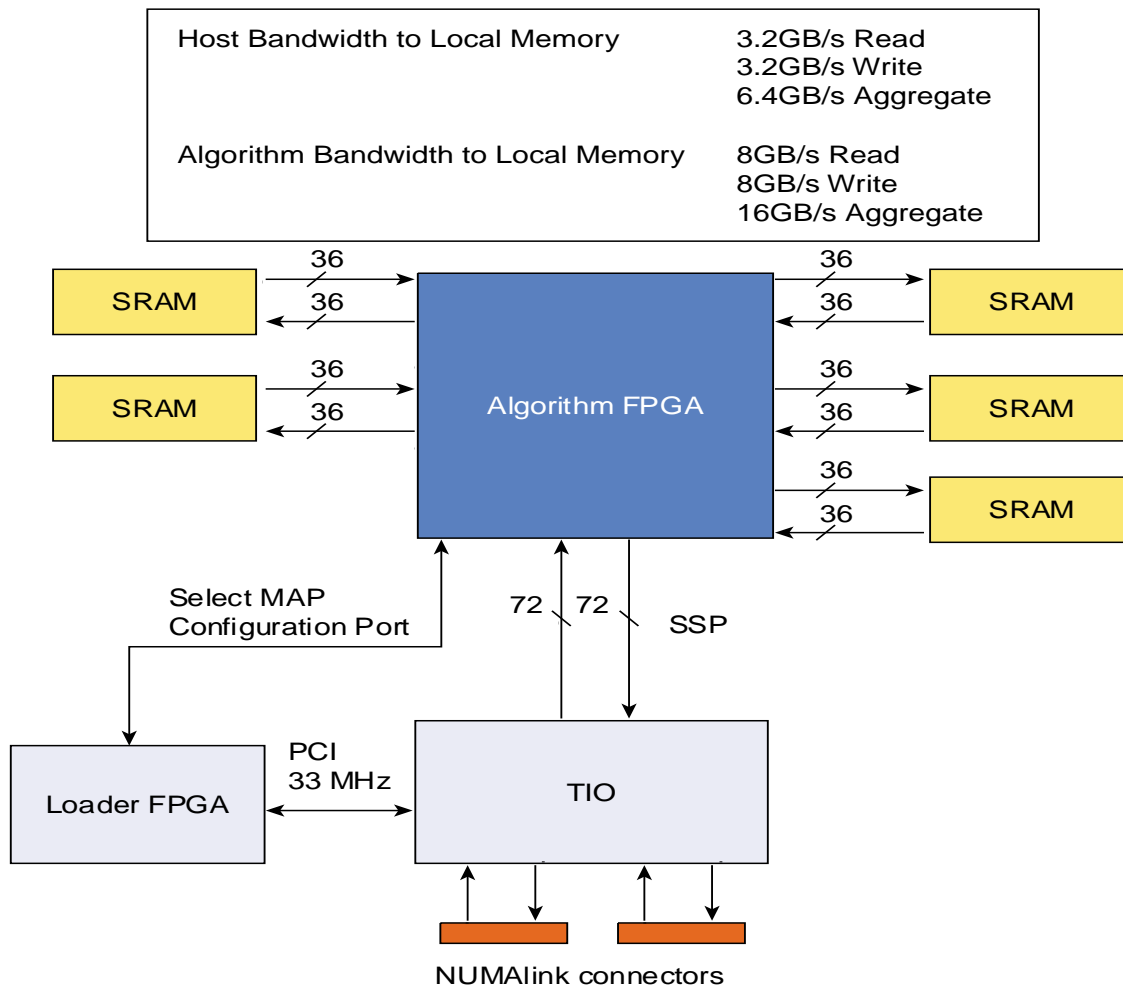


Figure 1-4 RASC FPGA Functional Block Diagram

The FPGA is connected to an SGI Altix system via the SSP port on the TIO ASIC. It is loaded with a bitstream that contains two major functional blocks:

- The reprogrammable algorithm
- The Core Services that facilitate running the algorithm

For more information on the RASC FPGA, see Chapter 3, “RASC Algorithm FPGA Hardware Design Guide”.

For more information on the Altix system topology and a general overview of the Altix 350 system architecture, see Chapter 2, “Altix System Overview”.

RASC hardware implementation for SGI Altix 4700 systems is based on blade packaging as shown in Figure 1-5. For an overview of the Altix 4700 system, see “SGI Altix 450 and Altix 4700 System Overview” on page 17.

The RASC hardware blade contains two computational FPGAs, two TIO ASICs, and a loader FPGA for loading bitstreams into the computational FPGAs. The computational FPGAs connect directly into the NUMalink fabric via SSP ports on the TIO ASICs. The new RASC blade has high-performance FPGAs with 200K logic cells and increased

memory resources with 10 synchronous static RAM dual in-line memory modules (SSRAM DIMMs).




Figure 1-5 RASC Blade Hardware

For legacy systems, optional brick packaging is available for the latest RASC hardware.

RASC Software Overview

Figure 1-6 shows an overview of the RASC software.

Major software components are, as follows:

- Standard Linux GNU debugger with FPGA extensions
- Download utilities
- Abstraction layer library
- Device Manager
- Algorithm device driver
- Download driver
- Linux kernel
- COP (TIO, Algorithm FPGA, memory, download FPGA)

This software is described in detail in this manual in the chapters that follow.




Figure 1-6 RASC Software Overview

Altix System Overview

This chapter provides an overview of the physical and architectural aspects of your SGI Altix 350 system or SGI Altix 450 or SGI Altix 4700 system. This chapter includes the following sections:

- “SGI Altix 350 System Overview” on page 15
- “SGI Altix 450 and Altix 4700 System Overview” on page 17

Note: This chapter provides an overview of the SGI Altix 350 system. For more information on this system, see the *SGI Altix 350 System User’s Guide* available on the SGI Technical Publications Library. It provides a detailed overview of the SGI Altix 350 system components and it describes how to set up and operate the system. For an overview of SGI ProPack software and installation and upgrade information, see the *SGI ProPack 5 for Linux Start Here*.

SGI Altix 350 System Overview

The Altix 350 system advances the SGI NUMAflex approach to mid-range modular computing. It is designed to deliver maximum sustained performance in a compact system footprint. Independent scaling of computational power, I/O bandwidth, and in-rack storage lets you configure a system to meet your unique computational needs. The small footprint and highly modular design of the Altix 350 system makes it ideal for high computational throughput, media streaming, or complex data management.

The Altix 350 system can be expanded from a standalone single-module system with 2GB of memory and 4 PCI/PCI-X slots to a high-performance system that contains 32 processors, one or two routers, up to 192 GB of memory, and 64 PCI/PCI-X slots. For most configurations, the Altix 350 system is housed in one 17U rack or one 39U rack as shown in Figure 2-1; however, for small system configurations, the Altix 350 system can be placed on a table top.

Systems that are housed in 17U racks have a maximum weight of approximately 610 lb (277 kg). The maximum weight of systems that are housed in 39U racks is approximately 1,366 lb (620 kg). The racks have casters that enable you to remove the system from the shipping container and roll it to its placement at your site.

See Chapter 1, “Installation and Operation,” in the *SGI Altix 350 System User’s Guide* for more information about installing your system. Check with your SGI service representative for additional physical planning documentation that may be available.




Figure 2-1 Example of SGI Altix 350 Rack Systems

The Altix 350 system is based on the SGI NUMAflex architecture, which is a shared-memory system architecture that is the basis of SGI HPC servers and supercomputers. The NUMAflex architecture is specifically engineered to provide technical professionals with superior performance and scalability in a design that is easy to deploy, program, and manage. It has the following features:

Shared access of processors, memory, and I/O. The Super Hub (SHub) ASICs and the NUMAlink-4 interconnect functions of the NUMAflex architecture enable applications to share processors, memory, and I/O devices.

- Each SHub ASIC in the system acts as a memory controller between processors and memory for both local and remote memory references.
- The NUMALink interconnect channels information between all the modules in the system to create a single contiguous memory in the system of up to 384 GB and enables every processor in a system direct access to every I/O slot in the system.

Together, the SHub ASICs and the NUMALink interconnect enable efficient access to processors, local and remote memory, and I/O devices without the bottlenecks associated with switches, backplanes, and other commodity interconnect technologies.

System scalability. The NUMAflex architecture incorporates a low-latency, high-bandwidth interconnect that is designed to maintain performance as you scale system computing, I/O, and storage functions. For example, the computing dimension in some system configurations can range from 1 to 32 processors in a single system image (SSI).

Efficient resource management. The NUMAflex architecture is designed to run complex models and, because the entire memory space is shared, large models can fit into memory with no programming restrictions. Rather than waiting for all of the processors to complete their assigned tasks, the system dynamically reallocates memory, resulting in faster time to solution.

SGI Altix 450 and Altix 4700 System Overview

Note: This chapter provides a brief overview of the SGI Altix 4700 series system. For more information on this system, see the *SGI Altix 4700 System User's Guide* available on the SGI Technical Publications Library. It provides a detailed overview of the SGI Altix 4700 system components and it describes how to set up and operate the system. For an overview of the SGI Altix 450 system, see Chapter 3, "System Overview" in the SGI Altix 450 System User's Guide.

The Altix 4700 series is a family of multiprocessor distributed shared memory (DSM) computer systems that currently scales from 8 to 512 CPU sockets (up to 1,024 processor cores) and can accommodate up to 6TB of globally shared memory in a single system while delivering a teraflop of performance in a small-footprint rack.

The SGI Altix 450 currently scales from 2 to 76 cores as a cache-coherent single system image (SSI). For an overview of the SGI Altix 450 system, see Chapter 3, “System Overview” in the *SGI Altix 450 System User’s Guide*.

Future releases will scale to larger processor counts for single system image (SSI) applications. Contact your SGI sales or service representative for the most current information on this topic.

In a DSM system, each processor board contains memory that it shares with the other processors in the system. Because the DSM system is modular, it combines the advantages of low entry-level cost with global scalability in processors, memory, and I/O. You can install and operate the Altix 4700 series system in a rack in your lab or server room. Each 42U SGI rack holds from one to four 10U high enclosures that support up to ten processor and I/O sub modules known as “blades.” These blades are single printed circuit boards (PCBs) with ASICs, processors, and memory components mounted on a mechanical carrier. The blades slide directly in and out of the Altix 4700 IRU enclosures. Each individual rack unit (IRU) is 10U in height (see Figure 2-2).




Figure 2-2 Altix 4700 Blade, Individual Rack Unit, and Rack

Note: An Altix 4700 system can support up to eight RASC blades per single system image. Current configuration rules require two compute blades for every RASC blade in your system.

The Altix 4700 computer system is based on a distributed shared memory (DSM) architecture. The system uses a global-address-space, cache-coherent multiprocessor that

scales up to sixty four Intel 64-bit processors in a single rack. Because it is modular, the DSM combines the advantages of lower entry cost with the ability to scale processors, memory, and I/O independently to a maximum of 512 processors on a single-system image (SSI). Larger SSI configurations may be offered in the future, contact your SGI sales or service representative for information.

The system architecture for the Altix 4700 system is a fourth-generation NUMAflex DSM architecture known as NUMALink-4. In the NUMALink-4 architecture, all processors and memory are tied together into a single logical system with special crossbar switches (routers). This combination of processors, memory, and crossbar switches constitute the interconnect fabric called NUMALink. There are four router switches in each 10U IRU enclosure.

The basic expansion building block for the NUMALink interconnect is the processor node; each processor node consists of a Super-Hub (SHub) ASIC and one or two 64-bit processors with three levels of on-chip secondary caches. The Intel 64-bit processors are connected to the SHub ASIC via a single high-speed front side bus.

The SHub ASIC is the heart of the processor and memory node blade technology. This specialized ASIC acts as a crossbar between the processors, local SDRAM memory, the network interface, and the I/O interface. The SHub ASIC memory interface enables any processor in the system to access the memory of all processors in the system. Its I/O interface connects processors to system I/O, which allows every processor in a system direct access to every I/O slot in the system.

Another component of the NUMALink-4 architecture is the router ASIC. The router ASIC is a custom designed 8-port crossbar ASIC. Using the router ASICs with a highly specialized backplane or NUMALink-4 cables provides a high-bandwidth, extremely low-latency interconnect between all processor, I/O, and other option blades within the system.

RASC Algorithm FPGA Hardware Design Guide

This chapter describes how to implement the algorithm that you have identified as a candidate for acceleration and provides a hardware reference for the RASC Algorithm Field Programmable Gate Array (FPGA) hardware. It covers the following topics:

- “RASC Core Services Overview” on page 22
- “Algorithm Interfaces” on page 26
- “Algorithm Design Details” on page 42
- “RASC FPGA Design Integration” on page 55
- “Simulating the Design” on page 66

Note: Make sure that you read the section called “Major Technical Changes in the RASC 2.20 Release” on the “New Features in This Guide” page in the frontmatter of this manual.

RASC Core Services Overview

Figure 3-1 shows a block diagram of RASC Core Services. The Core Services is made up of all the functional blocks excluding the algorithm block. Core Services is user independent and is pre-synthesized. The algorithm block is specific to the user algorithm and evolves through the RASC design flow. In the RASC 2.20 release, Core Services has new features that increase the bandwidth and reduce the latency as perceived by the algorithm. The components of Core Services are described in “Core Services Architecture Overview” on page 23.

The RASC algorithm FPGA is a Xilinx Virtex 4 LX200 part (XC4VLX200-FF1513-10). It is connected to an SGI Altix system via the SSP port on the TIO ASIC and loaded with a bitstream that contains two major functional blocks:

- The reprogrammable Algorithm Block
- The Core Services Block that facilitates running the algorithm.

Core Services Features

Core Services is the key component of RASC which facilitates execution of the user algorithm in the algorithm block. It helps in synchronizing the I/O, memory, and Algorithm Block operations. This section describes the services provided by the Core Services Block of the RASC FPGA. These features include:

- Scalable System Port (SSP) Implementation: physical interface and protocol
- Global Clock Generation and Control
- Independent read and write ports to each of the three logical or five physical random access memories (SRAMs)
- Single-step and multi-step control of the Algorithm Logic
- Independent direct memory access (DMA) engines for read and write data (two read DMA in the Input DMA block and two write DMA engines in the Output DMA block)
- Both block DMA and stream DMA are supported (for more information on stream DMA see Chapter 5, “Direct I/O”)
- Programmed Input/Output (PIO) access to algorithm’s Debug port and algorithm defined registers

- Control and Status registers
- Host and FPGA process synchronization capabilities, including interrupts and atomic memory operations (AMOs).

Core Services Architecture Overview

The Core Services in the RASC 2.20 release have new direct memory access streaming capabilities that increases bandwidth and reduce latency as perceived by the algorithm. This describes the components of Core Services as shown in Figure 3-1.

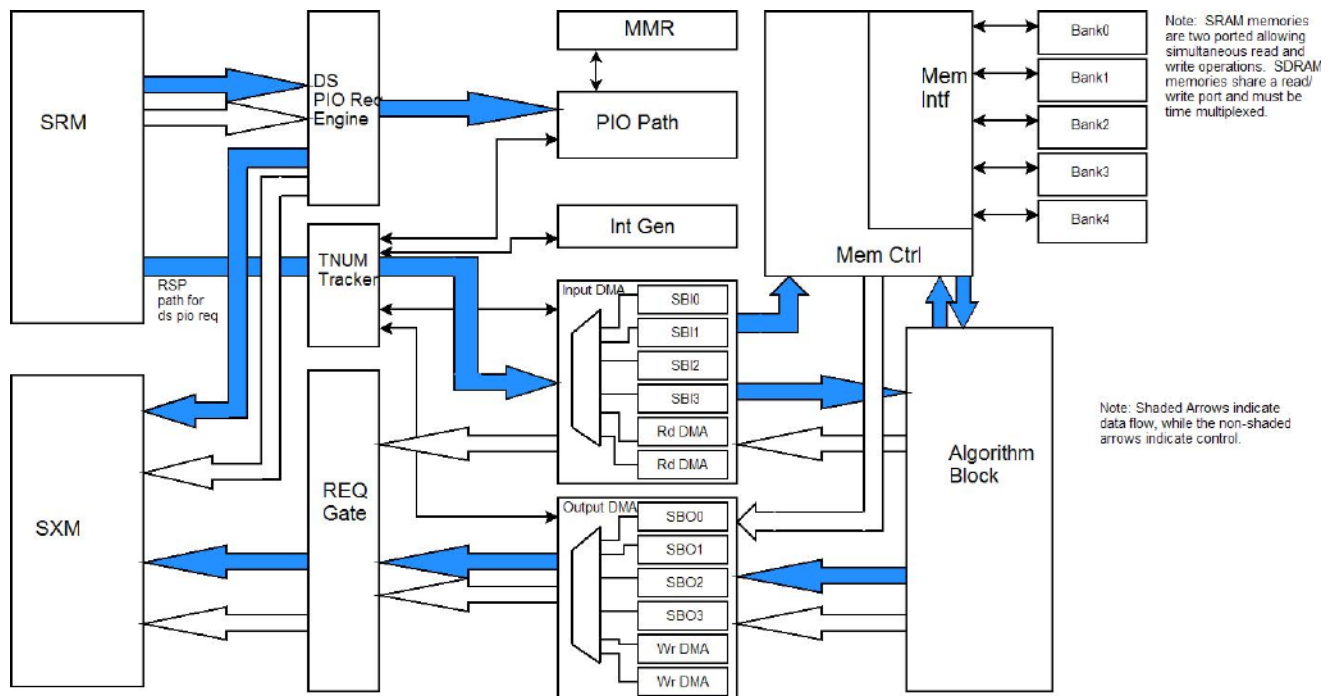


Figure 3-1 Block Diagram of the RASC Core Services

SRM

The SSP Receive Module (SRM) is the receive logic for the Scalable System Port (SSP).

SXM

The SSP Transmit Module (SXM) is the transmit logic for the SSP.

Programmed Input/Output Request Engine

The Programmed Input/Output (PIO) request engine handles the host write and read requests for individual 64-bit values.

TNUM Tracker

The TNUM tracker assigns and tracks transaction IDs within the FPGA.

Request Gate

This logic assembles the SSP packet for FPGA write and read from main memory requests.

Memory Mapped Registers

The memory mapped register (MMR) block houses and handles the registers used to control and use the FPGA.

Interrupt Generator

The interrupt generator assembles the proper packet data to interrupt the host in the case of a direct memory access (DMA) done or `alg_done`, for example.

Input Direct Memory Access

The input direct memory access (DMA) block is comprised of up to four stream DMA engines that target the algorithm for the data coming in from main memory and a block DMA read engine that targets SRAM for the data coming in from main memory.

Output Direct Memory Access

The output DMA block is comprised of up to four stream DMA engines that take data directly from the algorithm and target main memory and a block DMA write engine that takes data from SRAM and targets main memory.

Memory Controller

The Memory Controller handles the memory interfaces and is responsible for arbitrating between sources and destinations. It provides a generic interface to the algorithm for accessing on-board memory.

Algorithm Block

The algorithm block is the application specific, user defined code that performs the desired calculation.

Algorithm Run Modes

The Core Services Logic provides the mechanism for application and debugger software to run the Algorithm Block in one of two modes: *Normal Mode* or *Debug Mode*. In Normal Mode, the algorithm is enabled to run and allowed to run to completion uninterrupted. In Debug Mode, the algorithm is enabled to run but forced to execute in blocks of one or more steps, allowing the user to stall the algorithm at intermediate points in the execution and query intermediate internal values. A step could be as small as one Algorithm Block clock cycle, or an adhoc size defined by Algorithm Block logic. Core Services Logic continues to run more or less normally regardless if in Normal or Debug Mode. For a description of how the Algorithm Block can use the debug mode hooks, see the “Algorithm Interfaces” on page 26.

Algorithm Interfaces

This section defines the interface signals between the Algorithm Block and the Core Services Block and algorithm defined registers (ADRs). It covers the following topics:

- “Algorithm Control Interface” on page 26
- “Algorithm Defined Registers” on page 27
- “Streaming Direct Memory Access” on page 29
- “Debug Register Interface” on page 35
- “SRAM Interface” on page 35

Algorithm Control Interface

There are control signals to enable the smooth transition of control from Core Services to the algorithm and back. These signals are listed and explained in Table 3-1.

Table 3-1 Algorithm Control Interface Signals

Signals	Direction	Description
<code>clk</code>	input	The algorithm clock is available at frequencies of 50, 66, 100, and 200 MHz. This clock signal is gated, synchronous and phase aligned to the <code>core_clk</code> used in Core Services. The clock frequency is determined by a macro defined at synthesis of the FPGA bitstream. The clock is driven on global clock buffers for low skew routing.
<code>clk_s</code>	input	The optional supplemental clock has a frequency that is determined by macro definitions at synthesis of the FPGA bitstream to any frequency $n/m * 200$ MHz, where n and m are five-bit integers. The clock is driven on global clock buffers for low skew routing. Note that a supplemental clock, whose cycles are not of a multiple of 5 nanosecond (ns), needs to be considered asynchronous to <code>clk</code> and appropriate logic must be added to handle properly the clock domain crossings between <code>clk</code> and <code>clk_s</code> .
<code>rst</code>	input	Resets the algorithm logic. The algorithm logic is always put into reset for four clock cycles synchronous to <code>clk</code> before triggering it to execute.

Table 3-1 Algorithm Control Interface Signals

Signals	Direction	Description
step_flag_out	output	Step boundary indicator flag. For every clock cycle that the <code>step_flag</code> signal is asserted, it signals to Core Services that one step has been completed. The Core Services logic ignores this signal when the algorithm is not run in debug mode.
alg_done	output	Set when the algorithm has finished processing. This can either be set and held until the algorithm is reset or pulsed. When <code>alg_done</code> is asserted, <code>clk</code> will become inactive the following clock cycle. The signal <code>alg_done</code> must clear when <code>rst</code> is asserted.

Algorithm Defined Registers

The Algorithm Defined Registers (ADRs) are bidirectional, that is, they can be read or written to by the algorithm. ADRs can be read and/or written to by both the host and algorithm. This gives a simple mechanism for an algorithm to forward information to future invocations of the algorithm. The write capability overlaps the debug registers up to RASC 2.20. All ADR signals are synchronous to `clk`.

There are hooks to implement from one to 64 ADRs. By default, eight ADRs are included. If a different number is required, change the `user_space_wrapper` and `alg_block_top` configuration files. The ADR Verilog module is parameterized, which should minimize the need to make changes when changing the quantity of these registers.

The ADR signal definitions are provided in Table 3-2.

Table 3-2 ADR Signal Definitions

Signals	Direction	Description
alg_def_reg<n>[63:0]	input	Current value of ADR <i>n</i>
alg_def_reg_updated[<n>]	input	Asserted one clock cycle when <code>alg_def_reg<n></code> is updated by the host. Can be left unconnected if not used.

Table 3-2 ADR Signal Definitions

Signals	Direction	Description
<code>alg_def_reg_polled[<n>]</code>	input	Asserted one clock cycle when <code>alg_def_reg<n></code> is read (polled) by the host. Can be left unconnected if not used.
<code>alg_def_reg_write_[<n>]</code>	output	Assert this signal to write to <code>alg_def_reg<n></code> . The signal <code>alg_def_reg<n>_wr_data</code> must be valid. The value will appear in <code>alg_def_reg<n></code> the next clock cycle. A simultaneous write from the host supersedes this operation and this write will be ignored. In the event this happens, <code>alg_def_reg<n>_updated</code> will be asserted the following cycle.
<code>alg_def_reg<n>_wr_data[63:0]</code>	output	Value to write to <code>alg_def_reg<n></code> when <code>alg_def_reg<n>_write</code> is asserted.

Algorithm Defined Register Configuration

The following Verilog macros must be defined in the `alg.h` file. By setting the first macro, up to 64 ADRs can be defined. Note that a setting other than 8 requires changes to `user_space_wrapper` and `alg_block_top.v` configuration files. The remaining macros are bit masks that allow you to optimize by restricting usage, for example, making the ADR read-only by the host. The macros are, as follows:

```
// number of ADRs to implement
`define ADR_REG_NUM 8

// bit mask to specify which ADRs to implement
`define ADR_REG_IN_USE      `ADR_REG_NUM'hff

// bit mask to specify which ADRs are read-only by alg
`define ALG_ADR_READ_ONLY  `ADR_REG_NUM'hff

// bit mask to specify which ADRs are read-only by host
`define HOST_ADR_READ_ONLY `ADR_REG_NUM'h00

// bit mask to specify which ADRs are write-only by host
`define HOST_ADR_WRITE_ONLY `ADR_REG_NUM'h00
```

Algorithm Defined Register Usage

The host first writes to `alg_def_reg<n>`, then reads it. Subsequently, the algorithm does a write (see Figure 3-2). These signals represent what the algorithm sees as a result of both host initiated actions and algorithm initiated actions.




Figure 3-2 Algorithm Defined Registers (ADR) Interface Usage

Streaming Direct Memory Access

This section describes streaming direct memory access (DMA) and covers the following topics:

- “Input Streaming DMA Engine” on page 30
- “Start Input Stream” on page 30
- “End input stream” on page 31
- “Output Stream” on page 32
- “Start Output Stream” on page 33
- “Flushing Output Stream” on page 34
- “Streaming DMA Extractor Statements” on page 34

In streaming DMA, the direct memory access transfers data back to host memory automatically after it receives the data from the algorithm. When the algorithm completes, it asserts the flush signal and waits for the return of the flushed signal before it terminates itself. See the descriptions of the `strm_out_<n>_flush` and `strm_out_<n>_flushed` signals in Table 3-4 on page 32.

Input Streaming DMA Engine

Table 3-3 shows the signal definitions of the input streaming DMA engine.

Table 3-3 Input Streaming DMA Engine Signal Definitions

Signals	Direction	Description
<code>strm_in_<n>_data[127:0]</code>	input	Input stream data. Valid when <code>strm_in_<n>_data_vld</code> is asserted.
<code>strm_in_<n>_data_rdy</code>	input	When asserted, data in the input stream is available. <code>strm_in_<n>_data</code> is valid.
<code>strm_in_<n>_data_vld</code>	input	When asserted, <code>strm_in_<n>_data</code> is valid.
<code>strm_in_<n>_complete</code>	input	A one cycle pulse means that the transfer of data for this stream is complete. This can be asserted at the same time as the last assertion of <code>strm_in_<n>_data_vld</code> or sometime after.
<code>strm_in_<n>_rd_en</code>	output	Assert this signal to read from the input stream buffer. If there is no data in the buffer, this signal has no effect. If the input stream is to be treated as a FIFO, this signal can be used as a pop signal. The logic, when it sees an asserted <code>strm_in_<n>_rd_en</code> signal, will deliver the next datum if and only if data is available. Otherwise, the signal is ignored.

Start Input Stream

This timing diagram shows the sequence to start an input stream (see Figure 3-3). In prior RASC releases, there was a delay of a couple cycles between assertion of the `rd_en` and `data_vld` signals. The new `strm_in_<n>_data_rdy` signal eliminates this delay or the need to poll to determine if data is available to the algorithm.




Figure 3-3 Start of Input Stream Transaction

End input stream

This timing diagram shows the end of an input stream (see Figure 3-4).




Figure 3-4 End of Input Stream Transaction

Output Stream

Table 3-4 contains signal definitions for the output stream DMA engines. There are four instantiations of the engine.

Table 3-4 Output Stream Signal Definitions

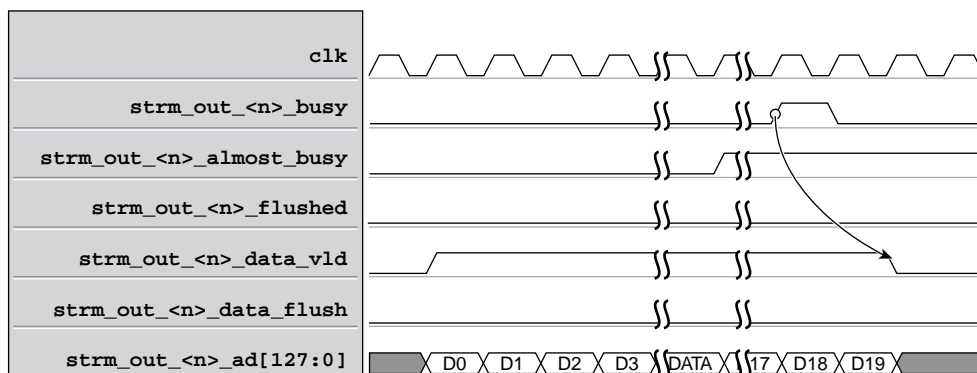
Signals	Direct	Description
strm_out_<n>_busy	input	When asserted, the buffer in output stream logic is full. <code>strm_out_<n>_data_vld</code> must be deasserted within two cycles.
strm_out_<n>_almost_busy	input	When first asserted, it may become necessary to deassert <code>strm_out_<n>_data_vld</code> within 256 cycles.
strm_out_<n>_flushed	input	<p>All previously sent data has been flushed to memory. The algorithm must assert <code>strm_out_<n>_flush</code> if it expects <code>strm_out_<n>_flushed</code> to return.</p> <p>The flushed signal is intended to give the algorithm a guarantee that all outstanding writes to host memory have been completed prior to an <code>alg_done</code> interrupt.</p> <p>When using SRAM for multi-buffering, the algorithm knows all the data has been transferred to SRAM with the last write to SRAM, after which it can signal <code>alg_done</code>. In this case, it is guaranteed that all the data will be in SRAM by the time the DMA to read the SRAM starts.</p> <p>In the case of streaming, there is potential for a race condition where the interrupt following an <code>alg_done</code>, without a flush signal asserted, arrives before all of the data is written to host memory.</p>
strm_out_<n>_data_vld	output	Assert when <code>strm_out_<n>_ad_out</code> is valid. This signal should not be asserted if <code>strm_out_<n>_flush</code> was previously asserted and <code>strm_out_<n>_flushed</code> has not yet been asserted.
strm_out_<n>_flush	output	Assert at the end of an output stream. This signal is used to guarantee data previously delivered to the output stream has been sent to host memory. <code>strm_out_<n>_data_vld</code> should not be asserted until <code>strm_out_<n>_flushed</code> has been asserted.

Table 3-4 Output Stream Signal Definitions

Signals	Direct	Description
<code>strm_out_<n>_ad[127:0]</code>	output	Three possible types of data for this signal are, as follows: system address, byte count, or data. Currently, only the data is supported.
<code>strm_out_<n>_data_last</code>	output	Assert at the end of an output stream. Assertion of this signal obviates the need for the algorithm to pad the output data. Assertion of <code>strm_out_<n>_data_last</code> also signals the corresponding DMA engine to terminate the current DMA after all outstanding writes to host have been completed. This allows the algorithm to terminate the DMA without the need to deliver all the data specified for the transfer.

Start Output Stream

This timing diagram shows the sequence to start an output stream (see Figure 3-5) and the flow control necessary when the output stream becomes busy.

**Figure 3-5** Start of Output Stream Transaction

Flushing Output Stream

Figure 3-6 shows the end of an output stream. Note that the flush signal can be asserted any time during or after the final assertion of `strm_out_<n>_data_vld`.




Figure 3-6 End of Output Stream Transaction

Streaming DMA Extractor Statements

The extractor statements for DMA streaming are in the form of `// extractor stream:<stream_name> <stream #> <stride>`, as follows:

```
// extractor STREAM_IN:in_a 0 0
// extractor STREAM_IN:in_b 1 0
```

A *stride* is the minimum size (in bytes) of a piece of data that is guaranteed to **not** have dependencies with other same-sized blocks of data in the input stream. A stride of 0 means that no such data exists and that the input stream may **not** be broken apart and spread across FPGAs.

There are four DMA stream engines. Four data stream pointers are defined in the following extractor statement:

```

| * stream numbers
V
// extractor STREAM_IN:in_a 0 0
// extractor STREAM_IN:in_b 1 0
// extractor STREAM_IN:in_c 2 0
// extractor STREAM_IN:in_d 3 0
// extractor STREAM_OUT:out_a 0 0
```



```
// extractor STREAM_OUT:out_b 1 0
// extractor STREAM_OUT:out_c 2 0
// extractor STREAM_OUT:out_d 3 0
```

The arrow above the extractor statement points to the stream number. You can use multiple stream DMAs simultaneously, but total transfer rate remains the same. The bandwidth is evenly divided among active streams.

Debug Register Interface

There are hooks to implement from zero to 64 debug registers. By default, sixty-four debug registers are included. If you require a different number, you must make changes to `user_space_wrapper` and `alg_block_top`.

Table 3-5 Debug Register Definition

Signal Name	Direction	Description
<code>debug<n>[63:0]</code>	output	Current Value of Debug <i>n</i>

This Verilog module is also parameterized, so the needed changes should be minimal.

SRAM Interface

This section describes the SRAM interface and covers the following topics:

- “Arbitration” on page 35
- “Handshaking Methodologies” on page 37
- “SRAM Configuration” on page 37
- “Address Offsets” on page 38

Arbitration

The algorithm may run concurrent to the Core Services. This means that you must be careful when assigning the ports at the SRAM interface. When designing an algorithm,

you can no longer rely on inherent arbitration that comes from running the master (algorithm or Core Services), exclusively.

If a situation is created, where a particular external SRAM port is being accessed by Core Services (via some host action such as a DMA or PIO) and also by the algorithm, simultaneously, causing an access conflict. Core Services and the algorithm must go through arbitration to gain access to the port when an access conflict exists. If you schedule an implementation where an access conflict to a port can occur, you must design your algorithm implementation to handle the arbitration activity that can occur. If you schedule your implementation where no access conflict can occur, no accommodations must be implemented to deal with the arbitration.

You can still assign the ports to avoid any arbitration by using extractor statements and directives.

You may assign ports in the `alg.h` file through 'define statements. An example of allowing the algorithm to read from SRAM Bank 0 and write to SRAM Bank 1 is, as follows:

```
////////////////////////////////////  
// Specify SRAM ports for algorithm //  
////////////////////////////////////  
'define alg_uses_sram0_rd  
'define alg_uses_sram1_wr
```

You may also define the port usage for Core Services. This is done through extractor statements in the Verilog. An example of assigning the Core Services the write port of SRAM Bank 0 and the read port of SRAM Bank 1 is, as follows:

```
// A multiple-source error is generated if a destination output is  
targeted  
// by more than one input source. (captured in the bit 0 of debug  
register 9)  
//  
// extractor CS: 1.0  
// extractor VERSION:19.1  
// extractor SRAM:op_in0 1024 64 sram[0] 0 in u stream  
// extractor SRAM:res0 1024 64 sram[1] 0 out u stream
```

For a more detailed discussion of the extractor statements, please see "Adding Extractor Directives to the Source Code" on page 144.

Handshaking Methodologies

There are two handshaking techniques available. You may use either the busy signal methodology or the crediting scheme. Only one method is needed on each interface and you may select one method for one interface and the other method for a second interface.

If, however, it is guaranteed that the algorithm and DMA will never use the same SRAM port, neither of the schemes described below are required.

Using Busy Signal

The busy signal method is a straight forward back pressure method to indicate to the algorithm that it must stall its SRAM write queue until the SRAM is available. This is signaled by the SRAM interface by asserting the `mem_<n>_rd_busy` or `mem_<n>_wr_busy` signal with a high value.

The algorithm must deassert the `mem_<n>_rd_cmd_vld` and `mem_<n>_wr_cmd_vld` signal two clocks after the interface asserts busy and may not assert it until the busy signal is no longer asserted.

Using SRAM Crediting Scheme

The interface indicates that space is available for a 128-bit transfer by pulsing the `mem_<n>_rd_fifo_cred` or `mem_<n>_wr_fifo_cred` signal high for one clock period. For example, after reset if the SRAM can handle sixteen (16) 128-bit transfers, the interface asserts the proper credit signal for sixteen (16) clocks. Furthermore, as a previously full buffer slot is freed, a credit is sent back to the algorithm indicating that there is room for another transfer.

In the crediting scheme, the algorithm must maintain a credit counter and it must not send data when there are no credits available. Credit overflow and underflow are the responsibility of the algorithm designer.

SRAM Configuration

There are three acceptable and mutually exclusive memory configurations allowed in RASC Core Services. The first is the same as the configuration found in the prior RASC 2.1 release. This configuration pairs up two 64-bit memory ports to be used as a 128-bit memory. There is a fifth 64-bit memory port available in this configuration.

A new configuration is available in the RASC 2.20 release (and later) that enables five 64-bit memory ports. All five of these ports are independent and can be configured individually.

You may also decide that you do not need memory ports and all data will flow in and out through the streaming DMA engines. This enables an area optimization by removing the memory interfaces from the integrated design.

Each memory configuration can also be defined as algorithm access only (no DMA).

The SGI RASC Algorithm Configuration tool allows you to configure memory resources. For more information, see “Memory Configuration” on page 159.

Address Offsets

Note that the address offsets of the `alg_mem_<n>_offset` macro were changed from bits 9:0 to 23:14 in the RASC 2.1 release. Bits 31:24 are ignored.

Table 3-6 SRAM Interface Signal Definition

Signals	Direction	Descriptions
<code>mem_<n>_rd_cmd_vld</code>	output	When asserted, <code>mem_<n>_rd_addr</code> is valid. Some number of cycles later <code>mem_<n>_rd_data_vld</code> and <code>mem_<n>_rd_data</code> will be asserted with valid data. After the assertion of <code>mem_<n>_rd_busy</code> , this signal must be deasserted within two cycles.
<code>mem_<n>_wr_cmd_vld</code>	output	When asserted, <code>mem_<n>_wr_addr</code> , <code>mem_<n>_wr_be</code> , and <code>mem_<n>_wr_data</code> are valid. After the assertion of <code>mem_<n>_wr_busy</code> , this signal must be deasserted within two cycles.
<code>mem_<n>_rd_busy</code>	input	When asserted, this signal indicates the <code>mem_<n>_rd_cmd_vld</code> signal must be deasserted within two cycles.
<code>mem_<n>_wr_busy</code>	input	When asserted, this signal indicates the <code>mem_<n>_wr_cmd_vld</code> signal must be deasserted within two cycles.

Table 3-6 SRAM Interface Signal Definition

Signals	Direction	Descriptions
<code>mem_<n>_rd_fifo_cred</code>	input	This signal is asserted for one clock period for each available location in the command FIFO. After reset, the core services will assert <code>mem_<n>_rd_fifo_cred</code> for the number of locations available. As commands enter the FIFO, the algorithm should decrement its credit counter. The core services will assert <code>mem_<n>_rd_fifo_cred</code> following the processing of the command, and the algorithm will increment its credit counter acknowledging the available location.
<code>mem_<n>_wr_fifo_cred</code>	input	Identical to the <code>mem_<n>_rd_fifo_cred</code> signal, but for writes.
<code>mem_<n>_wr_be[15:0 or 7:0]</code>	output	This signal is used to indicate a write enable of 64 bits of <code>mem_<n>_wr_data</code> . For a 64-bit port, bit 0 is used for bits [63:0], and for a 128-bit port, bits 0 and 8 are used for bits [63:0] and [127:64], respectively.
<code>mem_<n>_rd_addr[23:0 or 22:0]</code>	output	For a 128-bit port, this signal indicates the byte address for a memory read (or write) request. The lower four bits must be zero. For a 64-bit port, this signal has a 23 bit address and only the last 3 bits need to be zero.
<code>mem_<n>_wr_addr[23:0 or 22:0]</code>	output	For a 128-bit port, this signal indicates the byte address for a memory read (or write) request. The lower four bits must be zero. For a 64-bit port, this signal has a 23 bit address and only the last 3 bits need to be zero.
<code>mem_<n>_wr_data[127:0 or 63:0]</code>	output	This signal is the data for a memory write access.
<code>mem_<n>_rd_data_vld</code>	input	When asserted, <code>mem_<n>_rd_data</code> is valid.
<code>mem_<n>_rd_data[127:0 or 63:0]</code>	input	Data from a previous read request. Data is returned in the order requested.
<code>mem_<n>_error</code>	input	When asserted, this signal indicates an uncorrectable read data error occurred during a memory access.

SRAM Read Operation

Figure 3-7 shows a read of four cycles, no busy.

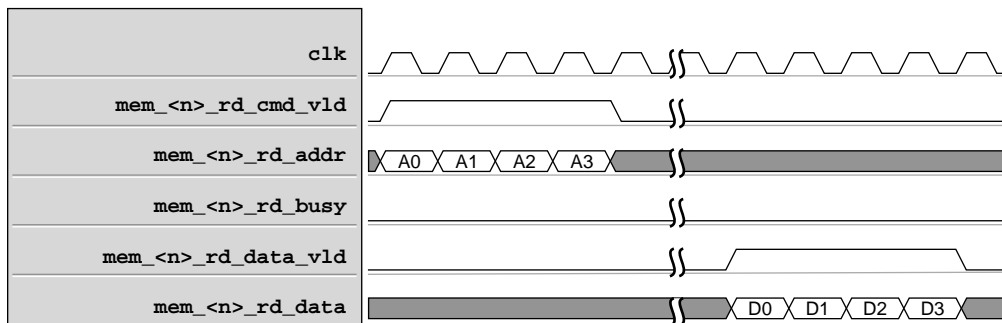


Figure 3-7 SRAM Read Operation

SRAM Read with Busy Operation

Figure 3-8 shows a read of eight cycles with busy. Note that `mem_<n>_rd_cmd_vld` must be deasserted two cycles after assertion of busy.

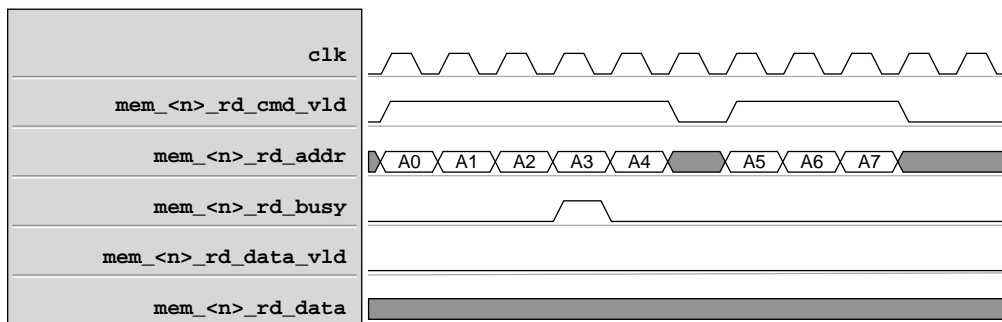


Figure 3-8 SRAM Read with Busy

SRAM Write with Busy Operation

Figure 3-9 shows a write of eight cycles with busy. Note that `mem_<n>_wr_cmd_vld` must be deasserted two cycles after assertion of busy.

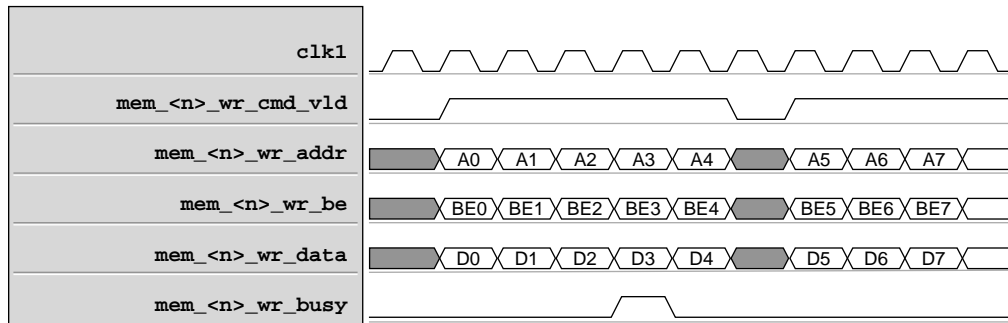


Figure 3-9 SRAM Write with Busy Assertion

Algorithm Design Details

This section provides information that the algorithm designer needs to implement the hardware accelerated algorithm within the FPGA and the Core Services Block framework. It covers the following topics:

- “Basic Algorithm Control” on page 42
- “Recommendations for Memory Distribution” on page 45
- “Implementation Options for Debug Mode” on page 46
- “External Memory Write Transaction Control” on page 47
- “External Memory Read Transaction Control” on page 48
- “Designing an Algorithm for Multibuffering” on page 49
- “Passing Parameters to Algorithm Block” on page 52
- “Recommended Coding Guidelines for Meeting Internal Timing Requirements” on page 54
- “Connecting Internal Signals to the Debugger” on page 54

Basic Algorithm Control

This section covers the general algorithm control sequence during a normal algorithm run (a run without breakpoints). Figure 3-10 illustrates such a sequence. The algorithm clock begins to toggle and the algorithm is put into reset for 4 algorithm clock cycles before triggering a new iteration of the algorithm.

When the algorithm is done, `alg_done` can either be set and held until the algorithm is reset or pulsed. Once asserted, one more algorithm clock cycle will be generated. At this point, the algorithm is done and no further algorithm clock pulses will be generated for this iteration. The user can then read the state of the algorithm block via the debug registers. The next time the algorithm is triggered by software, the reset sequence will start all over again. When there are no breakpoints, the activity of the `step_flag` signal is ignored. There are two resets, as follows:

- `core_rst` for Core Services (used only in Core Services) controlled by Programmed Input/Output requests (PIOs) to the TIO ASIC (see Figure 1-4 on page 10 and Figure 1-5 on page 12)

- algorithm reset (`rst`) in the Algorithm Block controlled by writing `alg_go` and `alg_clr` in `CM_CONTROL`

For more information on resets, see “Resets” on page 60.

This particular example (Figure 3-10) holds `step_flag` high, the method for clock based stepping.




Figure 3-10 Example of a Continuous, Normal Mode Algorithm Run

An overview of the hardware algorithm design steps is presented in Figure 3-11.

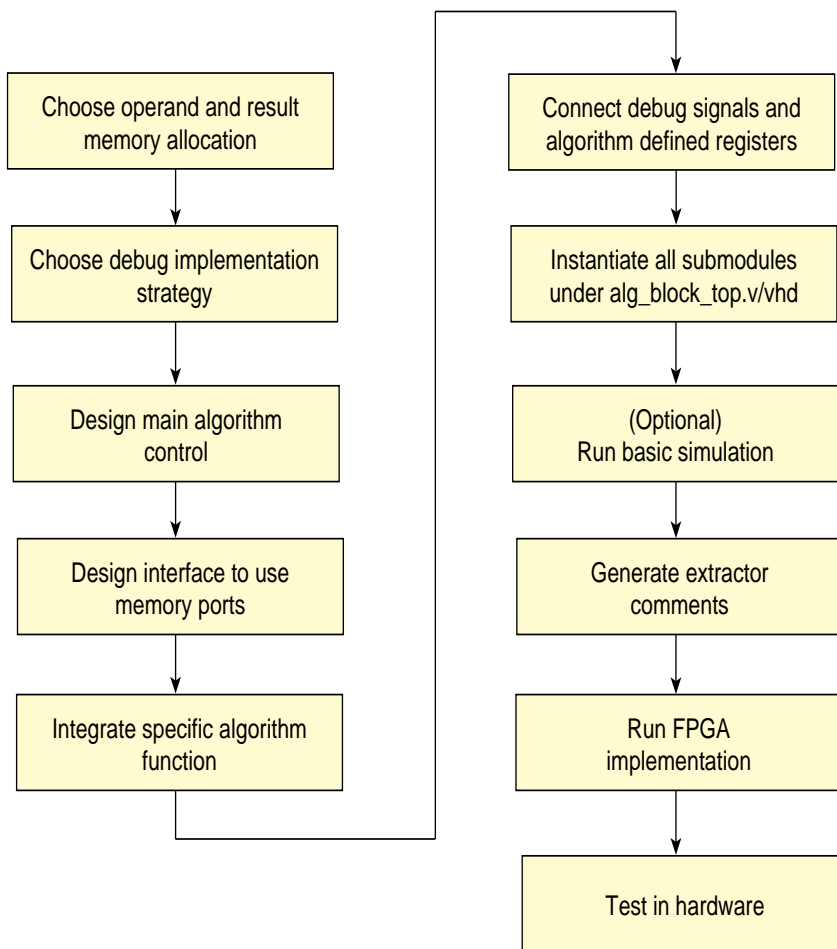


Figure 3-11 Hardware Accelerated Algorithm Design Flow

Recommendations for Memory Distribution

The RASC FPGA gives the algorithm access to five banks of up to 40MB SRAM (see Figure 1-5 on page 12). This section discusses the considerations for algorithm designers when deciding how to distribute input operands and parameters and output results among the available SRAM banks.

Input and Output Placement

The primary recommendation for data distribution is to organize algorithm inputs and outputs on separate SRAM banks. In other words, if bank 0 is used for input data, it should not also be used for output data and vice versa (by splitting the SRAM into two logical halves, for example). The motivation for this guideline comes from the fact that when an algorithm accesses a particular bank's read or write port, it blocks access to the DMA engine that wants to unload finished results or load new operands.

To avoid multiple arbitration cycles that add to read and write latency when designing an algorithm, see "Arbitration" on page 35, that describes how arbitration changed in the RASC 2.1 release from prior releases.

In order for the hardware accelerated algorithm to run efficiently on large data sets, it is recommended to overlap data loading and unloading with algorithm execution. To do this successfully, the algorithm designer needs to start with an SRAM layout of algorithm operands and results that allows each bank to be simultaneously accessed by the algorithm for one direction (read or write) and a DMA engine for the other direction (write or read).

Implementation Options for Debug Mode

The Algorithm can implement two different forms of Debug Mode, based on convenience or the desired granularity of step size: clock cycle based stepping or variable (ad hoc) stepping. The differences between the two determine the step size, or how long the algorithm will run when triggered to step once. There are also implementation differences for the step size variants. Currently, only one type of debug mode is supported at a time.

Clock Cycle Step Size Mode

Clock cycle based stepping means that the step size is one clock (`clk`) cycle. This method is easily implemented in RTL-based algorithms by tying the `step_flag` output (`set_flag_out`) to one (logic high). The step counter used by the debugger is 16 bits, so the maximum number of clock cycles that can be stepped in one debugger command is $2^{16}-1 = 65,535$. An example of this mode is shown in Figure 3-12. This is for one `alg_clk` cycle. The figure shows `rst` asserted for four rising `clk` edges. Note that you can rely on reset (`rst`) being asserted for longer than one cycle.

Note that since the Algorithm Block cannot detect when `clk` has stopped, the effect of stepping is transparent to the Algorithm Block.




Figure 3-12 Clock Cycle Stepping Mode Example

Variable Step Size Mode

Another approach for implementing debugging is to assert `step_flag` at points of interest rather than every clock cycle, which makes the step size a variable number of clock cycles. One example of this method would be to use `step_flag` as an output of the last state of the FSM. Another example would be for the user to put in a “trigger” for when an internal counter or state machine reaches a specific value (with an indeterminate number of clock cycle steps in between). In this case, `step_flag` is tied to the trigger so that the algorithm can break at a designated point.

The ad hoc nature of this approach requires the algorithm to define and notify Core Services of step boundaries with the `step_flag` signal. The `clk` signal will not stop toggling during the same clock cycle that the `step_flag` signal is asserted; it will turn off on the following clock (`clk`) cycle. See the timing diagram in Figure 3-13.




Figure 3-13 Variable Step Size Mode Example

External Memory Write Transaction Control

The process of using a write port involves the following step (example given for SRAM<n> alone):

When the address, data and byte enables are valid for a write, assert `mem_<n>_wr_cmd_vld` (reoccurring phase)

Example Write Transaction Timing Diagram

Figure 3-14 shows single and back-to-back write commands.

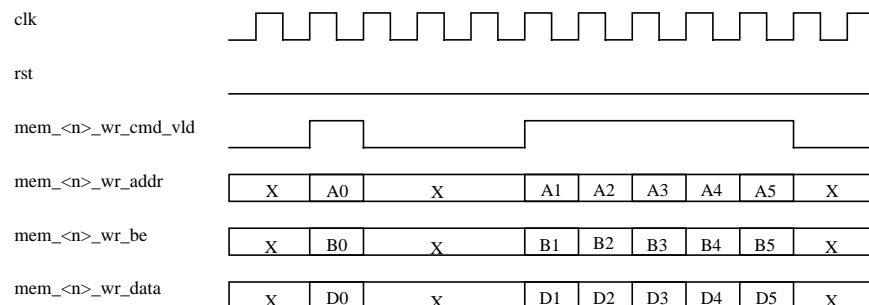


Figure 3-14 Single, and Multiple Write Commands

External Memory Read Transaction Control

The process of using a read port involves the following steps (example given for SRAM<n> alone):

1. When the address is valid for a read, assert `mem_<n>_rd_cmd_vld` (reoccurring phase). This step can be repeated while waiting for data (back-to-back burst reads). The algorithm can issue one quad-word (16 byte) read command every `clk` cycle per bank.
2. The read data will return on the bus `mem_<n>_rd_data[127:0]` several clock cycles later in the order it was requested (`mem_<n>_rd_cmd_vld` indicates that the read data is valid). The observed read latency from the algorithm's perspective will vary based on the clock period ratio between `alg_clk` and `core_clk`. Read latency is nominally 13 `core_clk` cycles; burst read commands are recommended for optimal read performance.

The algorithm should use the provided read data valid signal to determine when read data is valid and not attempt to count data cycles for expected latency.

Example Read Transaction Timing Diagram

Figure 3-15 and Figure 3-16 show single and back-to-back read commands.




Figure 3-15 Single Read Transaction

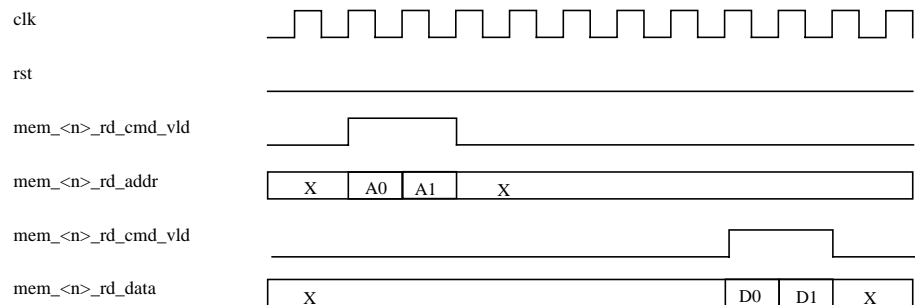


Figure 3-16 Multiple Read Transaction

Designing an Algorithm for Multibuffering

Purpose

Many applications targeted for hardware acceleration have large input and output data sets that will need to be segmented to fit subsets into the RASC brick's or RASC blade's external SRAM banks at any given time (a total of 40 MB per FPGA or 80 MB per blade) are available). Certain applications have large input data sets that have the same processing performed on subsets and only require that new input data be available in order to maintain continuous processing. These requirements bring up the notion of multibuffering of a large data set through an algorithm. Multibuffering data provides a continuous and parallel flow of data to and from the algorithm. Multibuffering provides the means to sequentially load data, execute, and unload a block of data in a loop one at a time.

In order for the hardware accelerated algorithm to run efficiently on large data sets, it is recommended to overlap data loading and unloading with algorithm execution. To do this successfully, start with an SRAM layout of algorithm operands and results that complies with the recommendations for memory distribution. The input and output data needs to be segmented into at least two segments per SRAM bank so that the Algorithm Block can execute on one segment while the next data segment is loaded by the Read DMA Engine and the previous output data segment is unloaded by the Write DMA Engine (ping-pong buffer effect). The SRAM bank can be segmented into any number of segments, as the algorithm and application designers best see fit.

Definitions

- **algorithm iteration**

One run of the algorithm; the operation that the algorithm performs between the time `alg_rst` is deasserted until the `alg_done` flag is asserted. Successive iterations require software to retrigger the algorithm.

- **segment / segment size**

The amount of memory needed on a particular SRAM bank for an algorithm iteration, round up to the nearest power of 2. One segment could include multiple input operands or multiple output operands, with spaces of unused memory within the segment if desired. A segment does not include fixed parameters that are applied to multiple algorithm iterations.

A multibuffered segment must be at most 1/2 of the SRAM bank's size. The minimum segment size is 16KB (but not all 16KB need to be used). Different SRAM banks can have different segment sizes according to the sizes and number of operands / results that reside on a particular SRAM bank.

```
alg_mem_0_offset[23:14] and alg_mem_1_offset[23:14]
```

Inputs to the Algorithm Block that specify for each SRAM the starting address of the current data segment. When used, these inputs get mapped to the Algorithm Block's SRAM address bits [23:14] (byte aligned address). The actual number of bits that are used for the mapping is determined by the segment size. If the segment size is 16KB, all 10 bits are mapped to SRAM address bits [23:14]. If the segment size is 512 KB, only bits [23:19] are mapped to SRAM address bits [23:19]. If the segment size is 8MB, only bit [23] is mapped to SRAM address bit [23].

Hardware Support

In order to support multibuffering, an algorithm should allow the upper bit(s) of its read and write SRAM addresses to be programmable via the `alg_mem*_offset` inputs. The offset inputs come from an internal FPGA register within the Core Services Block, accessible by the software layer. The offset inputs are 10 bits each, and can map to bits [23:14] of the corresponding SRAM address. Only the bits that correspond to the segment offset are used for a particular algorithm/application. For example, if the segment size is 32 KB, which leads to 512 segments in the 16MB bank, only the upper 9 of the 10 offset bits are used. Example Register Transfer Level (RTL) code for this configuration is shown below:

```
reg [23:15] mem0_rd_segment_addr;  
reg [14:0] mem0_rd_laddr;
```



```

wire [23:0] mem0_rd_addr;

// Read pointer
always @(posedge clk)
begin
// Fixed upper address bits [23:15] per iteration
if (rst) mem0_rd_segment_addr <= alg_mem0_offset[23:15];
// Counter for lower address bits [14:0]
if (rst)
mem0_rd_laddr <= 15'h000;
else if (rd_advance_en)
mem0_rd_laddr <= mem0_rd_laddr + 15'h10;
end

assign mem0_rd_addr = {mem0_rd_segment_addr, mem0_rd_laddr};

```

The algorithm must define a legal segment size between the minimum and maximum allowable sizes, and only operate on and generate one segment worth of data per iteration.

Special extractor directives are required to pass information about the algorithm's data layout and multibuffering capabilities to the software layer. The software layer requires a declaration of the input and output data arrays on each SRAM bank, with attributes size, type, and buffer-ability defined. The declarations are provided as extractor directives, or comments in the Verilog or VHDL source code. The following example comments are used to declare two 16KB input data arrays located on SRAM 0, and one non-multibuffering input parameter array, also located on SRAM 0:

```

##Array name
## # of elements in array
## Bit width of element
## SRAM location
## Byte offset (within given SRAM)
## Direction
## Type
## Stream flag
// extractor SRAM:input_a 2048 64 sram[0] 0x000000 in unsigned stream
// extractor SRAM:input_b 2048 64 sram[0] 0x004000 in unsigned stream
// extractor SRAM:param_i 512 64 sram[0] 0xffc000 in unsigned fixed

```

Note: In the code example above, `stream` is in the context of multibuffering. The `stream` data type is an historical anachronism and thus still appears in extractor directive code.

For arrays that are defined as buffered, the byte offset provided in the extractor comments is used to establish the data placement within a particular segment. In the declaration, the byte offset is given at the lowest segment address. For fixed arrays, the byte offset is the absolute placement of the data within the SRAM.

Further details on extractor comments can be found in Chapter 7, “RASC Algorithm FPGA Implementation Guide”

Software Responsibilities

Software uses configuration information derived from extractor directives to move the current active segment on each run of the algorithm. Software changes the values of `alg_mem_0_offset[23:14]` and `alg_mem_1_offset[23:14]` for each run of the algorithm and similarly moves the DMA engines’ input and output memory regions based on the current active segment.

Passing Parameters to Algorithm Block

This section describes the ways that variable parameters can get passed to the algorithm. For the purposes of this document, a parameter is distinguished from the general input data in that it can be fixed over multiple runs of the algorithm and does not require reloading. It is assumed that input data changes more often than parameters.

Small Parameters

The method used to pass variable parameters depends on the size and number of the required parameters. For a small number of 1-8 byte-sized parameters, the Algorithm Block can associate parameters with up to 64 Algorithm Defined Registers. The Algorithm Defined Registers are 64 optional registers that exist within the SSP memory mapped register region whose current values are inputs to the Algorithm Block (by default only the first eight of the following are defined: `alg_def_reg<n>[63:0]` where `<n>` is an integer 0-63). The Algorithm Block can assign reset (default) values for the parameters by tying the output signals `alg_def_reg<n>_wr_data[63:0]`, and allow the host application to change them.

Parameter Arrays

When an algorithm requires larger fixed parameters, portions of the SRAM banks can be used to hold the parameter data. This portion of the SRAM needs to be reserved for parameter data and kept unused by input data, so parameters need to be considered in the initial memory allocation decisions. Just as with small parameters, the mapping of parameter data to SRAM addresses is specified with extractor comments. The template and an example is provided below; further details are in the “Adding Extractor Directives to the Source Code” on page 144. In the provided example, a 1024-element parameter matrix (8KB) is mapped to the upper 8KB of SRAM0, which starts at address 0xFFE000. The type is unsigned and the array is fixed, which denotes that it is a parameter array and not as variable as an input data array (the other option is “stream”).

```
// extractor SRAM:<parameter array name> <number of elements in array>
<bit width of elements> <sram bank> <offset byte address into sram> in
<data type of array> fixed

// extractor SRAM:param_matrix0 1024 64 sram[0] 0xFFE000 in u fixed
```

Another use of declaring a fixed array in one of the SRAMs could be for a dedicated scratch pad space. The only drawback to using SRAM memory for scratch pad space is that an algorithm normally writes and then reads back scratch pad data. This usage model violates the multibuffering rule requiring an algorithm to dedicating each SRAM bank for either inputs or outputs. If you have a free SRAM bank that you do not need for inputs or outputs, this violation can be avoided and the multibuffering model can be maintained. If you have a free SRAM that is not being used for anything else, then you do not even have to add an extractor directive. An extractor directive is necessary if the SRAM bank is being used for other purposes so that software does not overwrite your scratch pad space. An extractor directive is also necessary to be able to access the scratch pad space from the debugger (reads and writes), so in general, an extractor directive is recommended.

Note that if you violate the multibuffering model for SRAM direction allocation, data will not be corrupted but the benefit of multibuffering will not occur because data transfer and algorithm execution cannot be overlapped. A template and an example is provided below for writing an extractor comment for a scratch pad space is, as follows:

```
// extractor SRAM:<scratch_pad_array_name>
    <number of elements in array>
    <bit width of elements>
    <sram bank>
    <offset byte address into sram>
    inout
```

```
        <data type of array>  
        <signed / unsigned>  
        fixed  
  
    // extractor SRAM:scratch1 1024 64 sram[2] 0x000000 inout u fixed
```

Recommended Coding Guidelines for Meeting Internal Timing Requirements

These guidelines are suggestions for achieving 200 MHz timing, when possible (not including floating point operations or use of multipliers).

1. Flop all outputs of the algorithm block, especially critical outputs, such as `step_flag`.
2. Flop and replicate the `rst` input if needed to distribute it as a high fanout signal.
3. Flop the input `strm_in_<n>_data_vld` data, `mem_<n>_rd_data_vld` and `mem_<n>_rd_data` before performing combinatorial logic on the data or data valid signals.
4. The general rule to abide by when trying to code a design that passes timing at 200 MHz is this: do not give PAR (the place and route tool) any tough decisions on placement where it would be difficult to find a good location. If a critical signal loads logic in multiple blocks, replicate it so that PAR does not have to try to optimize placement of the driving flop relative to the various loading blocks. You may have to add synthesis directives to prevent the synthesis tool from “optimizing out” your manually replicated flops. As far as possible, do not have a flop drive combinatorial logic in one block that then loads additional combinatorial logic in another block (such as Core Services), unless they can be physically grouped to adjacent locations, or in the worst case, minimize the total number of logic levels.

Connecting Internal Signals to the Debugger

This section shows how to make signals internal to the Algorithm Block viewable by the FPGA-enabled GNU Debugger (GDB) software. The Algorithm Block has up to 64 debugger output ports, each 64-bits wide. In order to make internal signals visible, the algorithm code should connect signals of interest to these outputs ports. To ease the timing issues on paths coming from the Algorithm Block, it is suggested to feed reregistered outputs to the debug outputs. Several examples are shown below:

```
assign debug0 = 64'h0000_000c_0000_0003; //[63:32] alg#, [31:0] rev#
```

In the above example, the outputs are tied, so it is not important to register the outputs.

```
always @(posedge clk)
    debug1 <= {32'h0, running_pop_count};
```

Since the intermediate value `running_pop_count` is also loaded by internal Algorithm Block logic, it is recommended to flop debug register 1 rather than use a wire connection. This helps isolate the loads of `running_pop_count` and reduce the number of constraints on the place and route program.

Besides connections to the debug port, the algorithm has to contain extractor comments that will pass the debug information to the software layer. Debug outputs use the `REG_OUT` type of extractor comment. The extractor comment tells the software layer what the mapping will be for internal signals to the corresponding debug address location. Examples are as follows:

```
// extractor REG_OUT:rev_id 32 u debug_port[0][31:0]

and

// extractor REG_OUT:running_pop_count 32 u debug_port[1][31:0]
```

The general format is:

```
REG_OUT:<signal name> <signal bit width> <type:unsigned/signed> <debug
port connection>[<bit range>]
```

RASC FPGA Design Integration

This section discusses additional details, including locations of the Algorithm Block in the design hierarchy, and global FPGA logic, such as clocks and resets. It covers the following topics:

- “Design Hierarchy” on page 56
- “FPGA Clock Domains” on page 57
- “Resets” on page 60

Design Hierarchy

Figure 3-17 shows the instance hierarchy of the RASC FPGA design. The top-of-chip is a wrapper module called `acs_top`. The instances in the top level include I/O buffer wrappers, clock resources wrappers, and the two major subdivisions of the logic design: `acs_core`, the pre-synthesized Core Services logic, and the `user_space_wrapper`, the top-level wrapper for the user/algorithm logic. As the algorithm application writer, you should begin the algorithm design using `alg_block_top` as the top level of the algorithm. The other instances within `user_space_wrapper` are small parts of Core Services resources that are left to be resynthesized based on their use and fanout within the algorithm logic. These include reset fanout logic and the debug port multiplexor.

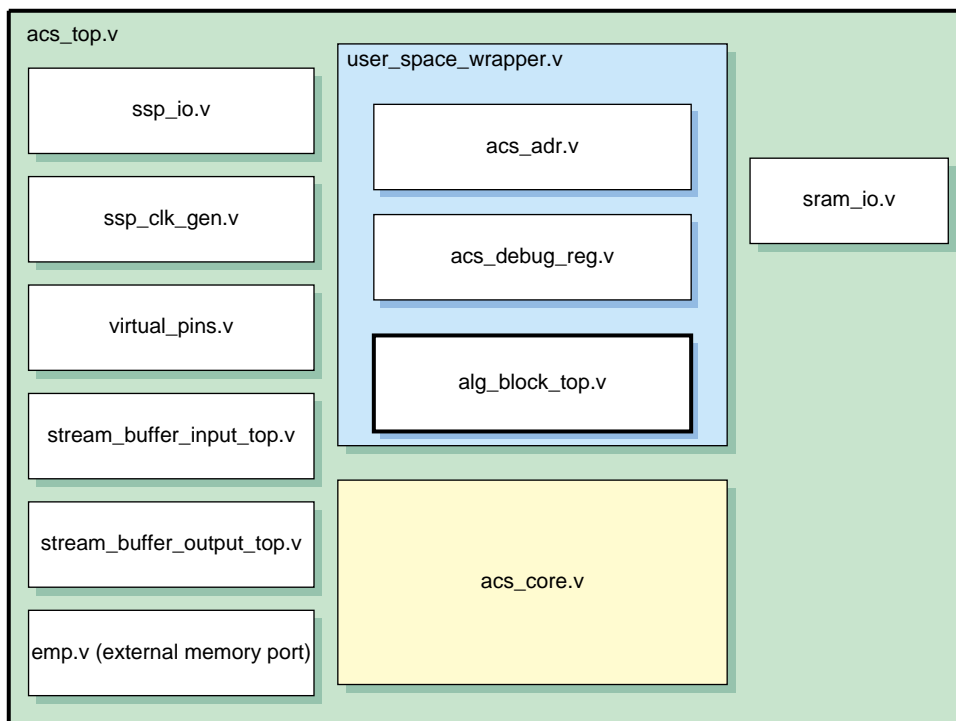


Figure 3-17 Instance Hierarchy of the RASC FPGA Design

The Algorithm / Core Services interface as defined in the section entitled “Algorithm Interfaces” on page 26, consist of the input and output signals defined for the module `alg_block_top`.

FPGA Clock Domains

This section describes the clock domains within the RASC Field Programmable Gate Array (FPGA), with a focus on the algorithm clock domain used by the Algorithm Block logic. There are two major domains: core clock and algorithm clock. However, the two domains are not completely asynchronous. They may either both be 200 MHz and phase aligned, or the algorithm clock can have a 50, 66, or 100 MHz and the clocks will be phase / edge-aligned (that is, a rising edge of the algorithm clock will correspond to a rising edge of the core clock).

This section covers the following topics:

- “Core Clock Domain” on page 57
- “Algorithm Clock Domain” on page 57
- “SSP Clock Domain” on page 59
- “QDR-II SRAM Clock Domains” on page 59
- “supplemental Algorithm Clock” on page 60

Core Clock Domain

The core clock has a fixed clock rate of 200 MHz. It is the main clock used within the Core Services Block. It is derived from the input control clock on the Scalable System Port (SSP) interface. The input clock is used as the clock input to the core clock digital clock manager (DCM) module. The `clk0` output of the DCM is driven onto a low-skew global clock buffer (BUFG) and from there is routed to `core_clk` domain registers as well as the feedback of the DCM for fine phase adjustments. In the place and route report, this clock is labelled `core_clk`.

Algorithm Clock Domain

The algorithm clock rate is selectable through the use of macro definitions. Speeds of 50, 66, 100, and 200 MHz can be selected. For speeds slower than 200 MHz, the DCM clock divider in the Xilinx Virtex 4 FPGA is used to create the `alg_block_top` module input called `clk`, the specified clock signal. Figure 3-18 shows a block diagram of the logic used

to create the clock signals. The RASC 2.20 release supports a supplemental algorithm clock. The additional clock signal is `alg_clk_s` that can be driven by one of the following outputs of DCM: `clk0`, `clkdv`, or `clkfx`. For more information, see “Supplemental Algorithm Clock” on page 63.

Macro definitions are used to select between the `clk0` and `clkdv` outputs of the DCM to drive the BUFGCE clock buffer, which in turn drives the `clk` signal in the Algorithm Block via the `alg_clk` signal.

The signal `alg_clk_enable` (not shown) generated by Core Services gates `alg_clk`. By gating the clock, the signal toggles only when the Algorithm Block is active. In the place and route clock report, this clock is labelled `alg_clk`.

When `clkdv` is used to drive the algorithm clock, the phase relationship between the core clock and the algorithm clock is determined by the Virtex 4 parameter `CLKOUT_PHASE`, which specifies the phase relationship between the DCM outputs of the same DCM. For Virtex 4 FPGAs, this parameter is specified as +/- 140 ps. Although the Xilinx timing tools do not take the `CLKOUT_PHASE` into account directly during analysis, an additional 140 ps has been added as input clock jitter to force the tools to correctly check paths crossing the `core_clk` and `alg_clk` domains. Any phase difference that is derived from the varying routes between the DCM outputs and the BUFG / BUFGCE elements as well as the clock tree fanouts are automatically considered by the Xilinx timing tools.

Core clock domain signals in the Core Services Block that communicate with the Algorithm Block, which is entirely in the `alg_clk` domain, have been designed to transition on the rising edge of `alg_clk`, even when `alg_clk` is run slower than the rate of `core_clk`.

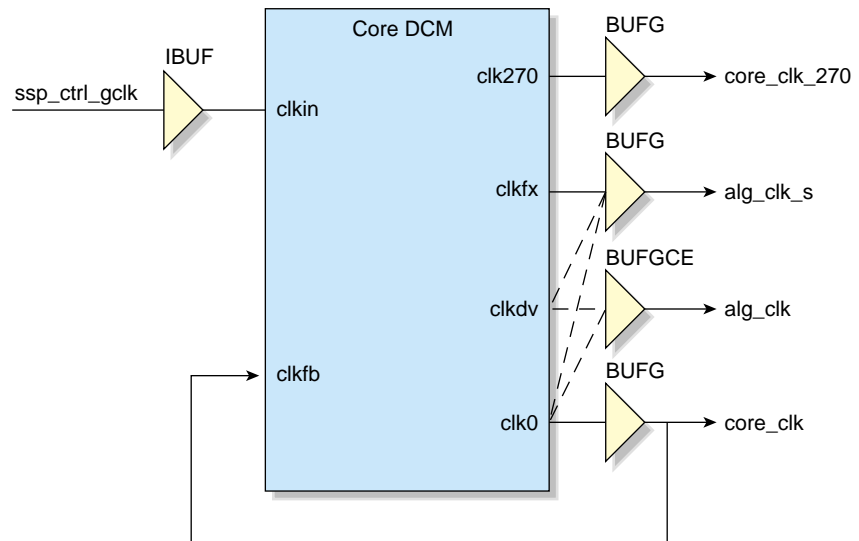


Figure 3-18 Core Clock and Algorithm Clock Source

SSP Clock Domain

Besides the core clock domain, which is equivalent to the Scalable System Port (SSP) control group domain, there are four data group clock domains within the SSP source synchronous receiver and transmitter logic. These four groups have a determined phase relationship between one another: each group is separated by a nominal 250 ps shift to reduce simultaneous switching noise on the SSP interface. In the place and route clock report, these clocks are labelled `int_grp1_clk` - `int_grp4_clk`.

QDR-II SRAM Clock Domains

The QDR-II SRAM module interfaces, a part of Core Services, uses five additional clock domains: one receive clock for each external SRAM (for a total of five, one for each physical SRAM component), and one common data transmit clock, which is a quarter clock cycle early relative to the `core_clk`. In the place and route clock report, these clocks are labelled `bank0_sram_rcv_clk`, `bank1_sram_rcv_clk`, `bank2_sram_rcv_clk`, `bank3_sram_rcv_clk`, and `core_clk270`.

supplemental Algorithm Clock

The RASC 2.20 software release supports a new supplemental clock feature on the RC100 blade. For more information, see “Supplemental Algorithm Clock” on page 63.

Resets

This section describes the hardware resets in the RASC Field Programmable Gate Array. The primary reset input to the FPGA is the `CM_RESET` signal on the SSP port. It is used as both a clock reset and as the basis of a full-chip control logic reset. When used as a clock reset, it is used as an asynchronous reset to the DCMs in the design. A full-chip logic reset is generated based on the assertion of `CM_RESET` or the de-assertion of any of the DCM locked signals. There is a circuit in Core Services that generates a synchronous `core_rst` signal, synchronous to the core clock domain. This reset is used throughout the control logic in Core Services.

The Algorithm Block receives a different reset generated by the Core Services’ Algorithm Controller sub-module. The Algorithm Block receives a reset that is synchronous to the algorithm clock. After a full-chip reset, the Algorithm Block’s reset input will remain asserted even after `core_rst` is removed. When the algorithm is started by software for the first time, the algorithm reset is removed. This is to allow for debug testing of internal signals after the algorithm completes. When software restarts the algorithm for another run (by setting the `ALG_GO` bit in the `CM_CONTROL` Register), the Algorithm Block will be held in reset for a total of 4 algorithm clock cycles, and then the algorithm reset will be released to allow the Algorithm Block to execute.

To summarize: the Algorithm Block is held in reset until its first use. Each time the algorithm is triggered to execute, the Algorithm Block will be held in reset for 4 clock cycles.

Algorithm Synthesis-time Parameters

This section describes the parameters to be specified by the algorithm designer in order to set the clock rate of the algorithm clock and to enable access to specific SRAM banks.

These synthesis-time parameters are specified in a Verilog include file called `alg.h` and are used by the top-level of design hierarchy. An example of this header file can be found in `$RASC/example/alg_simple_v/alg.h`.

Setting the Primary Algorithm Clock Speed

The primary algorithm clock rate can be set at 50, 66, 100, or 200 MHz.

The following is a portion from the above example include file that selects the clock speed. This shows the four macros used to specify an algorithm clock speed.

```

////////////////////////////////////
//   Specify clock speed of algorithm   //
////////////////////////////////////
// Only one of the below four should be uncommented

// For 200 MHz
`define alg_clk_5ns

// For 100 MHz
// `define alg_clk_10ns

// For 66 MHz
// `define alg_clk_15ns

// For 50 MHz
// `define alg_clk_20ns

```

One and only one of the above four macros must be defined. Here the macro to set the algorithm clock rate to 200 MHz is defined and uncommented whereas the other macros are left undefined.

SRAM Port Usage

In order for the Algorithm Block to have access to a given SRAM port, the corresponding macro must be defined in `alg.h`.

The following portion from the include file enables access to SRAM ports and shows one particular memory configuration. You can also use the SGI RASC Algorithm Configuration Wizard to configure memory, as described in “Memory Configuration” on page 159.

```

////////////////////////////////////
//   Specify SRAM ports for algorithm use //
////////////////////////////////////

```

```
`define alg_uses_sram0_rd
  `define alg_uses_sram0_wr
  `define alg_uses_sram1_rd
  `define alg_uses_sram1_wr
```

All, some, or none of these macros can be defined. In this example, the read port for SRAM bank 0 and the write port for SRAM bank 1 are enabled, while the write port for SRAM bank 0 and the read port for SRAM bank 1 cannot be used by the Algorithm Block.

The DMA engines and the Algorithm block can operate concurrently, but the algorithm must deal with arbitration. For more information on arbitration, see “Arbitration” on page 35.

The system memory configuration is defined in the `Makefile.local` file, as follows:

```
## Specify memory configuration
# Configuration no memory interfaces
#MEM_CFG_NONE=1

# Configuration 0: max 2x double-width SRAM QDR2 +
#                   1x single-width SRAM QDR
MEM_CFG_0=1
# Comment out unused memories below
USE_MEM_0=1
USE_MEM_1=1
#USE_MEM_2=1

# Configuration 1: max 5x single-width SRAM QDR2
#MEM_CFG_1=1
# Comment out unused memories below
#USE_MEM_0=1
#USE_MEM_1=1
#USE_MEM_2=1
#USE_MEM_3=1
#USE_MEM_4=1

# Configuration 2: max 5x double-width DRAM DDR2
#MEM_CFG_2=1
# Comment out unused memories below
#USE_MEM_0=1
#USE_MEM_1=1
#USE_MEM_2=1
#USE_MEM_3=1
#USE_MEM_4=1
```

The macro for `USE_MEM_2` in `alg.h` needs to be uncommented, as in:

```
// Comment out unused memories below
#define USE_MEM_0
#define USE_MEM_1
#define USE_MEM_2
```

The first two define statements, of course, can be commented out.

Supplemental Algorithm Clock

The RASC 2.20 software release supports a new supplemental clock feature on the RC100 blade.

Note: You can use this feature running RASC 2.20 if you update the `design/clk_gen/ssp_clk_gen.v` file.

The supplemental clock feature provides an application algorithm a second independent clock. It provides a finer granularity in choosing the algorithm clock frequency, relative to the four provided by the primary clock (see “FPGA Clock Domains” on page 57). It also allows you to implement an algorithm where sections of the algorithm are clocked at different rates. Regardless of the algorithm clocking frequency choices you make, the `core_services` block is clocked at a 200 MHz rate. The phase relationship of the supplemental clock to the primary clock is dependent upon the core clock digital clock manager (DCM) module of the Xilinx FPGA (see “Core Clock Domain” on page 57).

It is important to understand that you, as the algorithm designer, are responsible for assuring reliable data transfer across any clock domain boundaries created by the use of the supplemental clock. Potential issues involved and possible solutions are discussed in Xilinx documents such as *fifo_generator_ds317.pdf* and Chapter 4, “Block RAM” of the *Virtex-4 User Guide* available at <http://www.xilinx.com/>.

Setting the supplemental Algorithm Clock Speed

The supplemental clock is enabled and controlled by define statements in the `alg.h` file, as follows:

1. Add the supplemental clock ports to modules.

```
\define USE_ALG_CLK_S
```

This adds the `alg_block_top` module input called `clk_s` to which you can connect sections or all of the algorithm block.

This also adds the `clk_s` port to the `alg_block_top` instantiation in the `user_space_wrapper` module. In the `user_space_wrapper` module, this `clk_s` port is connected to the `alg_clk_s` input port of `user_space_wrapper`.

2. Select the supplemental clock frequency. See “Special Case Frequencies” on page 64 and “General Case Frequencies” on page 65.

Special Case Frequencies

Some frequency selections are treated as special cases. These are the cases where the selected frequency has a simple relationship to the 200 MHz `core_services` clock frequency and does not use the general frequency synthesis capability of the DCM. These are the supported special cases. The special cases only use the `CLK0` and `CLKDV` outputs of the DCM for the primary and supplemental clocks.

Case #1 - Primary clock = 200 MHz, supplemental clock = 100 MHz

```
`define alg_clk_5ns
```

Do not define `alg_clk_s_5ns`

Do not define `alg_clk_s_fx`

Case #2 - Primary clock = 100 MHz, supplemental clock = 200 MHz

```
`define alg_clk_s_5ns
```

```
`define alg_clk_10ns
```

Do not define `alg_clk_s_fx`

Case #3 - Primary clock = 66 MHz, supplemental clock = 200 MHz

```
`define alg_clk_s_5ns
```

```
`define alg_clk_15ns
```

Do not define `alg_clk_s_fx`

Case #4 - Primary clock = 50 MHz, supplemental clock = 200 MHz

```
`define alg_clk_s_5ns
```

```
`define alg_clk_20ns
```

Do not define `alg_clk_s_fx`

There are a few uninteresting cases when none of the supplemental clock parameters are defined and the primary clock is not 200 MHz, as follows:

If `alg_clk_s_5ns` is not defined,

and `alg_clk_s_fx` is not defined,

and `alg_clk_10ns`, `alg_clk_15ns` or `alg_clk_20ns` is defined

In these cases, the primary and supplemental clocks are the same frequency, therefore, the supplemental clock offers no benefits.

General Case Frequencies

This mode provides a large selection of discrete supplemental clock frequencies. This mode uses DCM frequency synthesis, so the supplemental clock frequency is essentially independent of the `core_services` clock frequency. In this mode, the supplemental clock is connected to the DCM `clkfx` output.

This mode is enabled when the `alg.h` file has the following defines:

```
`define alg_clk_s_fx
```

Do not define `alg_clk_s_5ns`

The synthesized frequency is controlled by two additional defines in `alg.h`. See the *Xilinx Virtex-4 datasheet ds302.pdf* for restrictions upon the frequency synthesis capability. They are listed under the "DCM and PMCD Switching Characteristics" and "Output Clock Phase Alignment" sections of this document (Tables 42, 43 and 48 in a recent revision).

The synthesized frequency is determined by applying multiplication and division factors to the 200 MHz `core_services` clock. Each factor has its own define. The following is an example for synthesizing a 300 MHz supplemental clock:

```
`define ALG_CLK_S_MULT 3
`define ALG_CLK_S_DIV 2
```

For information on the algorithm primary clock and how to set it, see "Setting the Primary Algorithm Clock Speed" on page 61.

Simulating the Design

This section provides a reference on how to simulate the Algorithm using the provided SSP Stub, sample test bench, and VCS simulator. It covers the following topics:

- “Intent of the Sample Test Bench” on page 66
- “Sample Test Bench Setup” on page 66
- “Running a Diagnostic” on page 68
- “Writing a Diagnostic” on page 71
- “Sample Test Bench Constants and Dependencies” on page 76
- “Sample Test Bench Utilities” on page 78

Intent of the Sample Test Bench

The Sample Test Bench (also called `sample_tb`) is a basic simulation environment for users to do sandbox testing of their algorithm code. The Sample Test Bench is provided as an optional intermediate step between writing an algorithm and loading the algorithm into RASC hardware. It is intended to help insure that the algorithm will function on a basic level (for example, a single algorithm iteration) prior to debugging a bitstream in hardware.

Sample Test Bench Setup

The sample test bench is designed for use with VCS. For use with other simulators, the user should modify the sample test bench along with associated scripts and makefiles.

A primary component of the sample test bench, the SSP Stub, consists predominantly of Verilog modules, although it also includes PLI calls to functions written in C code. The stub is instantiated in a sample Verilog test bench along with the Algorithm FPGA. The files for this test bench are in the directory, `$RASC/dv/sample_tb/`. In this directory you will find the following Verilog modules and other files:

- `top.v`: The top level of the sample test bench containing the Algorithm FPGA design (Core Services and the user’s algorithm), SSP Stub, SRAM simulation models, and clock generator.

- `ssp_stub.v`: Top level Verilog of the SSP Stub which passes signals to and from conversion modules. More information on submodules, PLI calls, and C functions that comprise the SSP Stub can be found in the “SSP Stub User’s Guide” section of this document.
- `init_sram0_good_parity.dat`, `init_sram1_good_parity.dat`, `init_sram2_good_parity.dat`, `init_sram3_good_parity.dat`, `init_sram4_good_parity.dat`: These SRAM initialization files contain data which is automatically loaded into the respective SRAM simulation models at the beginning of simulation. The data is in a format which the SRAM simulation model uses with the ECC bits embedded with the data. These default files can be overridden by the user on the command line at runtime.
- `final_sram0.dat`, `final_sram1.dat`, `final_sram2.dat`, `final_sram3.dat`, `final_sram4.dat`: These files contain data extracted from the respective SRAM simulation models at the end of simulation. These default files can be overridden by the user on the command line at runtime.
- `timescale.v`: This file contains the Verilog timescale of each of the components of the SSP Stub, as well as the algorithm FPGA design files. It is required that the algorithm being simulated makes use of the same timescale as the rest of the design.

In order to use the sample test bench, your VCS environment variables should be set up as follows:

```
### Environment Variables for VCS ###
setenv VCS_HOME <your_vcs_install_directory>
setenv VCSPLIDIR $VCS_HOME/<your_vcs_pli_directory>
setenv PATH $PATH\: $VCS_HOME/bin
```

Compiling the Sample Test Bench

Compiling the sample test bench is done using the Makefile provided. In order to compile the sample testbench including the SSP Stub and the algorithm Core Services logic, an algorithm must be specified (See the following note).

Note: The Makefile in the `sample_tb` directory uses the `$ALG_DIR` environment variable. This defaults to `$RASC/examples` though it can be modified by the user. The design files of the algorithm you specify must be in a directory under the `$ALG_DIR` path.

If the algorithm is written in VHDL, set the HDL_LANG environment to vhd1, as follows:

```
% setenv HDL_LANG vhd1
```

Otherwise, the algorithm is considered to be written in Verilog.

The algorithm you are building is specified on the command line. To compile the design with your algorithm, change directory to \$RASC/dv/sample_tb and enter:

```
% make ALG=<your_algorithm>
```

where <your_algorithm> is the directory name where the algorithm design files are. When no algorithm is specified, the default is ALG=alg_simple_v.

To remove older compiled copies of the testbench, type:

```
% make clean
```

Running a Diagnostic

To run a diagnostic on your algorithm, call the Makefile in the sample_tb directory using the “run” target and specifying which diag to run. The following is the usage and options of the “run” target:

```
% make run DIAG=diag_filename ALG=your_algorithm  
SRAM0_IN=sram0_input_file SRAM1_IN=sram1_input_file  
SRAM2_IN=sram2_input_file SRAM3_IN=sram3_input_file  
SRAM4_IN=sram4_input_file  
SRAM0_OUT=sram0_output_file SRAM1_OUT=sram1_output_file  
SRAM2_OUT=sram2_output_file SRAM3_OUT=sram3_output_file  
SRAM4_OUT=sram4_output_file
```

The `diag_file` specifies the diagnostic to be run and should be relative to the current directory. Again, the algorithm must be specified using the `ALG=your_algorithm` command line option. If none is specified, the runtime command uses same default as above (ALG=alg_simple_v). Specifying ALG this way allows the user to reuse the same diagnostic for multiple algorithms. The contents of each SRAM at the end of simulation will be dumped into .dat files that can be user-specified. If they are not specified, they default to:

```
init_sram0_good_parity.dat  
init_sram1_good_parity.dat  
init_sram2_good_parity.dat  
init_sram3_good_parity.dat  
init_sram4_good_parity.dat
```

```

final_sram0.dat
final_sram1.dat
final_sram2.dat
final_sram3.dat
final_sram4.dat

```

Note that there are five input and five output SRAM data files while the design is implemented for two logical SRAMs. Each of the logical SRAMs 0-4 are implemented as two separate physical SRAMs in the sample testbench.

The sram0* and sram1* files correspond to the first logical SRAM while sram2*, sram3*, correspond to the second logical SRAM. The sram4* files correspond to the 5th physical SRAM.

By specifying the SRAM input and output files the user can skip the DMA process for quick verification of the algorithm. This shortens the diagnostic run time, makes for less complex diagnostics, and allows the user to ignore core services as it has already been verified by SGI. The option of utilizing the DMA engines in simulation is included for completeness but should not be necessary for typical algorithm verification.

The association of SRAM0_IN and SRAM0_OUT with physical memory is, as follows:

```

mem0[127:64] -> qdr_sram_bank1.SMEM -> init_sram1_*.dat, final_sram1.dat
mem0[63:0   ] -> qdr_sram_bank0.SMEM -> init_sram0_*.dat, final_sram0.dat
mem1[127:64] -> qdr_sram_bank3.SMEM -> init_sram3_*.dat, final_sram3.dat
mem1[63:0   ] -> qdr_sram_bank4.SMEM -> init_sram2_*.dat, final_sram2.dat
mem2[63:0   ] -> qdr_sram_bank2.SMEM -> init_sram4_*.dat, final_sram4.dat

```

As the diagnostic runs, it will output status to the screen and to an output file named `<diag_filename>.<your_algorithm>.run.log`. When the stub receives an unexpected packet, it will output the following information in order: the command for the next expected packet, SSP fields of the expected packet, the command translation (if one exists) for the received packet, and the SSP fields of the received packet. This log file will appear in the same directory in which that the diagnostic is located.

Table 3-7 shows a summary of the algorithms, diagnostics, and commands provided with the sample testbench.

Table 3-7 Sample Testbench Algorithms and Commands

Algorithm Name	Diagnostic	Compile and Run Commands
alg_simple_v	diags/alg_simple_v	make ALG=alg_simple_v make run DIAG=diags/alg_simple_v ALG=alg_simple_v
alg_data_flow_v	diags/alg_data_flow_v	make ALG=alg_data_flow_v make run DIAG=diags/alg_data_flow_v ALG=alg_data_flow_v

Viewing Waveform Results

Each time a diagnostic is run, a file named `vcdplus.vpd` is generated in the `sample_tb` directory. This file can be input to Virsim for viewing the waveform. Since this file is generally large, it is overwritten for each diagnostic run. To save the waveform for a given diagnostic, copy the corresponding `vcdplus.vpd` file to a new name.

To view the waveform saved in the `vcdplus.vpd` file, use the following command:

```
% vcs -RPP vcdplus.vpd
```

A sample configuration file `sample_tb/basic.cfg` is provided for use when viewing waveforms in Virsim. It contains a limited number of relevant signals on the SSP interface, SRAM interfaces, and inside the design. Figure 3-19 shows a sample `vcdplus.vpd` waveform in Virsim.

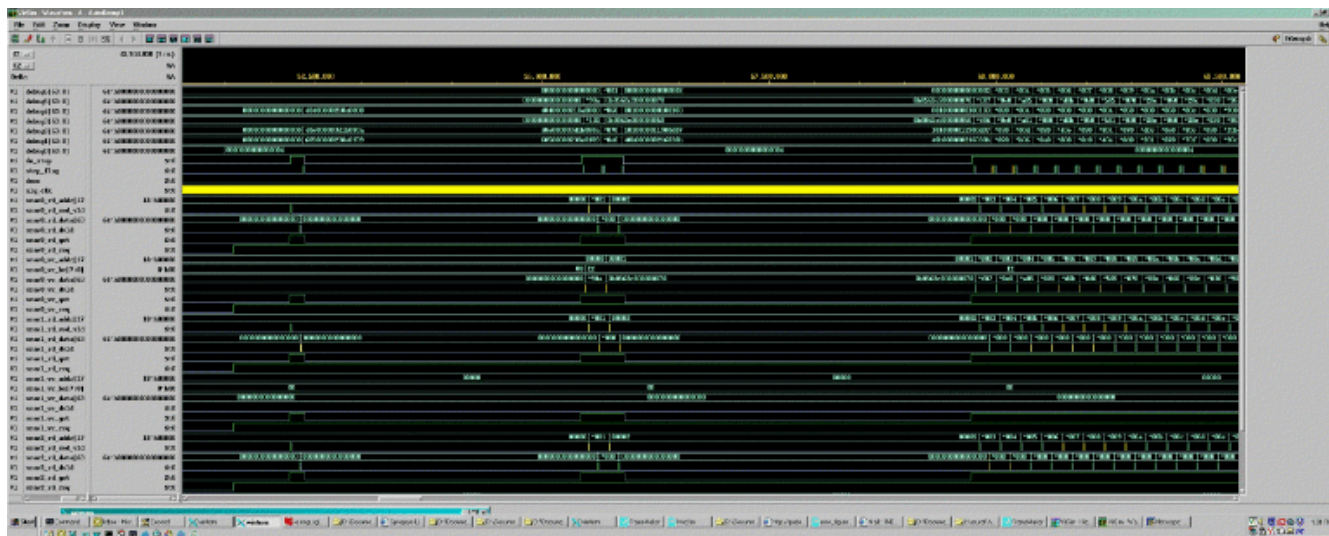


Figure 3-19 Sample vcdplus.vpd Waveform in Virsim

Writing a Diagnostic

The SSP Stub retrieves instructions through a text input file, the diagnostic. The SSP Stub parses this file at each semicolon to extract commands that the Stub executes. Many of the allowed commands in a diagnostic correspond to SSP packet types. There are other commands that the SSP Stub supports for diagnostic writing and debugging. The primary components of the diagnostic file are: packet commands, debugging commands, and comments.

It is important to note that most SSP packets come in pairs: a request and a response. For these types of packets, the request command and response command must be listed sequentially in a diagnostic. This method of keeping requests and response paired is used by the stub to associate request and response packets with the corresponding transaction number (TNUM). For more information on SSP packet types, see the *Scalable System Port Specification*. Also, when running the DMA engines, all transactions related to that sequence of events should be grouped together. See Appendix B, “SSP Stub User’s Guide” for more details on diagnostic writing and using the SSP stub.

The code listed below comprises a diagnostic that exercises the basic functionality of the algorithm FPGA outlined in the following steps:

- Initializes the algorithm FPGA Core Services (primarily MMR Writes)
- Executes DMA Reads to send data to the FPGA (stored in SRAM)
- Starts the Algorithm ($d = a \& b \mid c$) and polls the memory mapped registers (MMRs) to see when the Algorithm is done
- Executes DMA Writes to retrieve the Algorithm's results
- Checks the error status in the MMRs to verify that no errors were flagged.

The example diagnostic provided below is intended as a template that may be edited to match the user's algorithm.

```
##### Initialization packets. #####
# Arm regs by setting the REARM_STAT_REGS bit in the CM_CONTROL reg
snd_wr_req ( PIO, DW, ANY, 0x00000000000020, 0x0000000600f0003 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# Clear the CM_ERROR_STATUS register by writing all zeroes.
snd_wr_req ( PIO, DW, 3, 0x00000000000060, 0x000000000000000 );
rcv_wr_rsp ( PIO, DW, 3, 0 );

# Enable CM_ERROR_DETAIL_* regs by writing all zeroes to CM_ERROR_DETAIL_1.
snd_wr_req ( PIO, DW, ANY, 0x00000000000010, 0x000000000000000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# Enable desired interrupt notification in the CM_ERROR_INTERRUPT_ENABLE
register.
snd_wr_req ( PIO, DW, 4, 0x00000000000070, 0xFFFFFFFFFFFFFFFF );
rcv_wr_rsp ( PIO, DW, 4, 0 );

# Set up the Interrupt Destination Register.
snd_wr_req ( PIO, DW, ANY, 0x00000000000038, 0x000000000000000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

print "\n\nWait for SRAM bank calibration to complete.\n\n";

# Poll CM_STATUS for bits 51, 57, 58, 59, & 63.
poll (0x8, 63, 300);

print "\n\n*****Initialization finished\n\n";

##### Configure DMA Engines and Algorithm. #####

##### Configure the Read DMA Engine Registers. #####
print "\n\n*****Configure Read DMA Engine. Tell it to fill 32 cache lines
of data.\n\n";

# RD_DMA_CTRL register.
snd_wr_req ( PIO, DW, ANY, 0x00000000000110, 0x0000000000100020 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# RD DMA addresses.
snd_wr_req ( PIO, DW, ANY, 0x00000000000100, 0x0000000000100000 );
```

```
rcv_wr_rsp ( PIO, DW, ANY, 0 );
snd_wr_req ( PIO, DW, ANY, 0x000000000000108, 0x0000000000000000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# RD AMO address.
snd_wr_req ( PIO, DW, ANY, 0x000000000000118, 0x0000000000000000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# RD_DMA_DEST_INT
snd_wr_req ( PIO, DW, ANY, 0x000000000000120, 0x0000000200002000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

##### Configure the Write DMA Engine Registers. #####
print "\n\n*****Configure Write DMA Engine.\n\n";

# Write to the WR_DMA_CTRL register.
snd_wr_req ( PIO, DW, ANY, 0x000000000000210, 0x0000000000100020 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# WR_DMA_SYS_ADDR
snd_wr_req ( PIO, DW, ANY, 0x000000000000200, 0x0000000000100000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# WR_DMA_LOC_ADDR
snd_wr_req ( PIO, DW, ANY, 0x000000000000208, 0x0000000000000000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# WR_DMA_AMO_DEST
snd_wr_req ( PIO, DW, ANY, 0x000000000000218, 0x0000000000000000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# WR_DMA_INT_DEST
snd_wr_req ( PIO, DW, ANY, 0x000000000000220, 0x0000000400004000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

##### Configure the Algorithm Registers. #####
print "\n\n*****Configure Algorithm Registers\n\n";

snd_wr_req ( PIO, DW, ANY, 0x000000000000300, 0x0000000000000000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

snd_wr_req ( PIO, DW, ANY, 0x000000000000308, 0x0000000600006000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

##### Start Read DMA Engine for Read DMA 1 #####
print "\n\n*****Start Read DMA Engine for SRAM0\n\n";

# Set Bit 36 of the CM_CONTROL Reg to 1.
snd_wr_req ( PIO, DW, ANY, 0x000000000000020, 0x0000001400f00003 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# 1 of 32
rcv_rd_req ( MEM, FCL, ANY, 0x00000000100000 );
snd_rd_rsp ( MEM, FCL, ANY, 0, 0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
```



```

print "\n\n*****Polling for DMA RD-SRAM1 done (bit 42 of CM_STATUS).\n\n";
poll (0x8, 42, 200);
print "\n\n*****Done storing data in SRAM 1.\n\n";

##### Reconfigure DMA Engine for Read DMA 3 #####

# RD DMA addresses.
snd_wr_req ( PIO, DW, ANY, 0x00000000000100, 0x0000000000100000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );
snd_wr_req ( PIO, DW, ANY, 0x00000000000108, 0x0000000000400000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

##### Start Read DMA Engine for Read DMA 3 #####
print "\n\n*****Start Read DMA Engine for SRAM2\n\n";

# Set Bit 36 of the CM_CONTROL Reg to 1.
snd_wr_req ( PIO, DW, ANY, 0x00000000000020, 0x0000001400f00003 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# 1 of 32
rcv_rd_req ( MEM, FCL, ANY, 0x00000000100000 );
snd_rd_rsp ( MEM, FCL, ANY, 0, 0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C );

# Other Read DMA Transactions omitted here

# 32 of 32
rcv_rd_req ( MEM, FCL, ANY, 0x00000000100F80 );
snd_rd_rsp ( MEM, FCL, ANY, 0, 0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C );

print "\n\n*****Polling for DMA RD-SRAM2 done (bit 42 of CM_STATUS).\n\n";
poll (0x8, 42, 200);
print "\n\n*****Done storing data in SRAM 2.\n\n";

##### Start the Algorithm #####

# Set bit 38 of CM Control Register to 1 to start algorithm.
snd_wr_req ( PIO, DW, ANY, 0x00000000000020, 0x0000004400f00003 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

print "\n\n*****Started Algorithm.\n\n";

# Poll for ALG_DONE bit in CM_STATUS.
poll ( 0x8, 48, 2000);
print "\n\n*****Algorithm Finished.\n\n";

##### Start Write DMA Engine. #####

```

```

# Set bit 37 of CM Control Register to 1 to start Write DMA Engine.
snd_wr_req ( PIO, DW, ANY, 0x000000000000020, 0x0000002400f0003 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

print "\n\n*****Started Write DMA Engine.\n\n";

# 1 of 32
rcv_wr_req ( MEM, FCL, ANY, 0x00000000100000, 0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC );

snd_wr_rsp ( MEM, FCL, ANY, 0 );

# Other Write DMA Transactions omitted here

# 32 of 32
rcv_wr_req ( MEM, FCL, ANY, 0x00000000100F80, 0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC,
                                                    0xDCACBCECDCACBCEC,0xDCACBCECDCACBCEC );

snd_wr_rsp ( MEM, FCL, ANY, 0 );

print "\n\n*****Polling for DMA WR-SRAM0 done (bit 45 of CM_STATUS).\n\n";
poll (0x8, 45, 200);
print "\n\n*****Done retrieving data from SRAM 0.\n\n";

##### Finish Up #####

# dma_clear(). Set bits 39, 40, and 41 to 1 in CM_CONTROL.
snd_wr_req ( PIO, DW, ANY, 0x000000000000020, 0x0000038400f0003 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# finalcheck_ccc() Check CACHE_RD_DMA_FSM.
snd_rd_req ( PIO, DW, ANY, 0x000000000000130 );
rcv_rd_rsp ( PIO, DW, ANY, 0, 0x00000000000400000 );

print "Reading the Error Status Register to insure no errors were
logged.\n";
snd_rd_req ( PIO, DW, ANY, 0x000000000000060 );
rcv_rd_rsp ( PIO, DW, ANY, 0, 0x0000000000000000 );

```

Sample Test Bench Constants and Dependencies

Various constants and definitions for the sample test bench are contained within the following files:

- `ssp_defines.h` (internal stub variables)

- `user_const.h` (user modifiable)

Table 3-8 lists the files in the `sample_tb` directory with their function and calls dependencies.

Table 3-8 Files in the `sample_tb` directory

File	Functions	Dependent On
<code>start_ssp.c</code>	<code>start_ssp()</code>	<code>queue_pkt.h</code> , <code>setup_pkt.h</code> , <code>send_rcv_flits.h</code>
<code>send_rcv_flits.h</code>	<code>send_rcv_flits()</code> , <code>send_flit()</code> , <code>rcv_flit()</code> , <code>snd_poll()</code> , <code>rcv_poll()</code> , <code>finish_ssp()</code>	<code>setup_pkt.h</code> , <code>process_pkt.h</code> , <code>get_fields.h</code> , <code>snd_rcv_fns.h</code>
<code>queue_pkt.h</code>	<code>queue_pkt(string)</code> , <code>q_string_it(token)</code> , <code>pkt_string</code> , <code>strtok_checked(s1, s2)</code>	--
<code>setup_pkt.h</code>	<code>setup_pkt(snd_rcv)</code>	<code>snd_rcv_fns.h</code>
<code>snd_rcv_fns.h</code>	<code>snd_wr_req(pio_mem_n, size, tnum, addr,</code> <code>data, pkt)</code> , <code>snd_rd_req(pio_mem_n, size, tnum,</code> <code>addr, pkt)</code> , <code>snd_wr_rsp(pio_mem_n, size,</code> <code>tnum, error, pkt)</code> , <code>snd_rd_rsp(pio_mem_n,</code> <code>size, tnum, error, data)</code> , <code>snd_amo_rsp(tnum,</code> <code>error, pkt)</code> , <code>inv_flush(tnum, pkt)</code> , <code>rcv_wr_rsp(pio_mem_n, size, tnum, error,</code> <code>pkt)</code> , <code>rcv_rd_rsp(pio_mem_n, size, tnum, error,</code> <code>data, pkt)</code> , <code>rcv_wr_req(pio_mem_n, size, tnum,</code> <code>addr, data, pkt)</code> , <code>rcv_rd_req(pio_mem_n, size,</code> <code>tnum, addr, pkt)</code> , <code>rcv_amo_req(tnum, addr,</code> <code>data, pkt)</code>	<code>construct_pkt.h</code>
<code>construct_pkt.h</code>	<code>construct_pkt(type, tnum, address, data, error,</code> <code>pkt, to_from_n)</code> , <code>pkt_size(type)</code>	<code>make_command.h</code>
<code>make_command.h</code>	<code>make_command(type, tnum, error, to_from_n)</code>	--
<code>get_fields.h</code>	<code>get_fields(type)</code> , <code>f_string_it(token)</code>	--
<code>process_pkt.h</code>	<code>process_pkt(type)</code>	--

Sample Test Bench Utilities

The sample test bench includes utilities that help in generating and interpreting diagnostic data. To compile these files into executables, run the following command:

```
% gcc -c file_name -o executable name
```

The utilities provided include the following:

- `convert_sram_to_dw.c`

This program takes a standard SRAM input/output file (e.g. `final_sram0.dat`), and converts it to a more readable version consisting of one SGI double word (64-bits) of data per line. It assumes that the input file is made up of 36-bit words containing ECC bits. This utility is helpful when trying to interpret results from the stub output files.

Use: `convert_sram_to_dw input_file [output_file]`

Default output file: `convert_sram_to_dw_output.dat`

- `convert_dw_to_sram_good_parity.c`

This program takes a file containing one SGI double word (64-bits) of data per line, calculates ECC and outputs a file that can be loaded into SRAM for simulation (36-bits of data with ECC per line).

It assumes the input file contains the correct number of lines to fill the SRAM. This utility is useful when you want to input specific data to an SRAM and skip the DMA process in simulation.

Use: `convert_dw_to_sram_good_parity input_file [output_file]`

Default output file: `convert_dw_to_sram_good_parity_output.dat`.

- `command_fields.c`

This program takes an SSP command word, splits it into its SSP fields and outputs the SSP field information to the screen. The utility provides this data in the same format as the `get_fields.h` function in the SSP stub. This feature is potentially useful in debugging from the Virsim viewer.

Use: `command_fields 32-bit_hex_value`

- `check_alg_data_flow.c`

This program uses the SRAM output file to check data against input data. It takes data input from `init_sram0_good_parity.dat` and `init_sram1_good_parity.dat`, removes ECC, and performs a byte-sort on this data (byte-sorts each 8 byte quantity as done by `alg_data_flow_v`).

The program compares the result to the data in the `final_sram2.dat` and `final_sram3.dat` files (with ECC removed). As the comparisons are done, the program prints the byte-sorted input data on the left and the results from the final data on the right. If there are differences in the data, it prints "ERROR" on the corresponding line. The program exits after finding a specific number of errors (adjustable in the source code).

- `check_alg_simple.c`

This file is an example of a program that uses the SRAM output data file to check data against input data. It takes the SRAM0 and SRAM1 input data in `init_sram0_good_parity.dat` and `init_sram1_good_parity.dat` and calculates the results of $A \& B \mid C$. It then compares its expected data to the data in the files `final_sram2.dat` and `final_sram3.dat` (the default output for SRAM2 and SRAM3, respectively). As it proceeds, it prints out the results of each double word result D: on the left, what it expects based on the input data, on the right, what it is seeing in the final results. If there are discrepancies in the data, it prints "ERROR" on the line in question.

The code is set to exit after it finds 64 errors, but this number can be raised or lowered easily (line 60).

RASC Abstraction Layer

This section describes the Reconfigurable Application-Specific Computing (RASC) Abstraction Layer and covers the following topics:

- “RASC Abstraction Layer Overview” on page 81
- “RASC Abstraction Layer Calls” on page 83
- “How the RASC Abstraction Layer Works” on page 125

RASC Abstraction Layer Overview

The RASC Abstraction Layer provides an application programming interface (API) for the kernel device driver and the RASC hardware. It provides a similar level of support for application development as the standard open/close/read/write/ioctl calls for IO peripherals.

The Abstraction Layer is actually implemented as two layers. The lowest level is the COP (Co-Processor) level. This provides calls to function individual devices. The upper level, which is built on top of the COP level, is the algorithm level. Here, the application treats a collection of devices as a single, logical algorithm implementation. The Abstraction Layer manages data movement to and from the devices, spreading work across multiple devices to implement scaling.

As an application develops, you must decide which level you are programming to. They are mutually exclusive. You can use one or the other, but **never** both.

Figure 4-1 illustrates the interactions between the pieces and layers of the system.

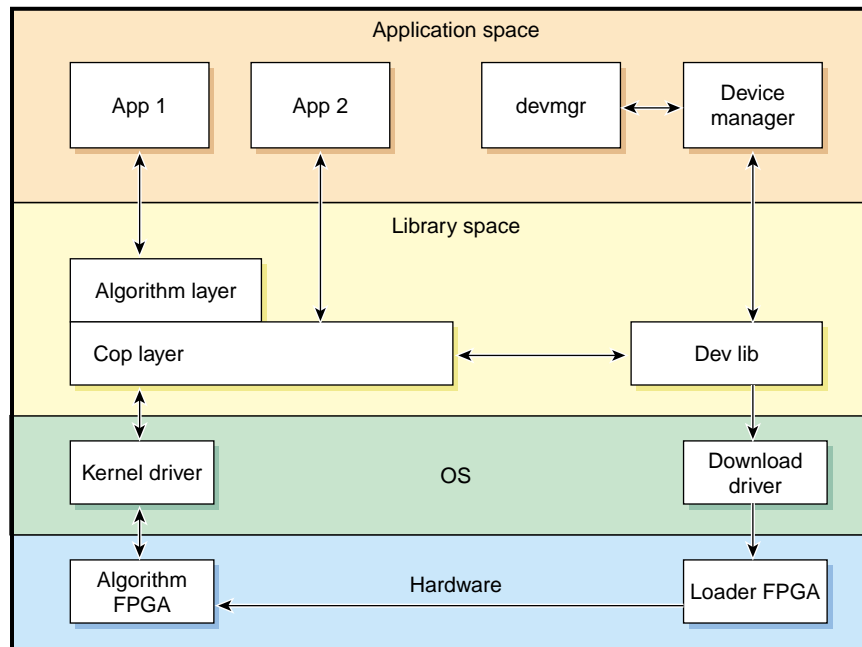


Figure 4-1 Abstraction Layer Software Block Diagram

RASC Abstraction Layer Calls

Function entry points can be called from either C or Fortran90 software language modules. The same basic entry point names are used for both languages; except that, the Fortran90 interfaces have the `_F` suffix appended to the C name. Fortran90 names are, of course, case insensitive. The same information must be passed to either the C or Fortran90 functions, however, basic differences in the languages may require slightly different formats for the argument lists. Both languages return the same values from a function call.

All function prototypes, data structures, and constants are defined in the `rasclib.h` file for C calls or in the `rasclib_definitions.mod` file for Fortran90 calls.

Table 4-1 contains a summary of the abstraction layer function definitions.

Table 4-1 Abstraction Layer Function Definitions - Summary

Function	Description
<code>rasclib_resource_reserve</code>	Sends a request to the device manager to reserve a number of devices for the calling application.
<code>rasclib_resource_release</code>	Releases the reservation on a number of devices and then moves the devices from the reservation pool to the free pool.
<code>rasclib_resource_configure</code>	Loads the given bitstream into a device and marks the device as in use.
<code>rasclib_resource_return</code>	Returns a configured device to its reservation pool.
<code>rasclib_resource_alloc</code>	Sends a request to the device manager to allocate/reserve devices for the application. This routine has been deprecated and will be removed in a future RASC release. You should use <code>rasclib_resource_reserve</code> and <code>rasclib_resource_configure</code> in its place.

Table 4-1 Abstraction Layer Function Definitions - Summary (**continued**)

Function	Description
<code>rasclib_resource_free</code>	Returns devices to the device manager's free pool. This routine has been deprecated and will be removed in a future RASC release. You should use <code>rasclib_resource_return</code> and <code>rasclib_resource_release</code> in its place.
<code>rasclib_algorithm_open</code>	Notifies the <code>rasclib</code> library that the user wants to use all of the reserved devices loaded with the bitstream associated with <code>al_id</code> as a single logical device.
<code>rasclib_algorithm_send</code>	Queues up a command that will move data from host memory to the device input buffer identified by <code>algo_array_id</code> .
<code>rasclib_algorithm_get_num_cops</code>	Returns the number of physical devices that are participating in the algorithm.
<code>rasclib_algorithm_receive</code>	Queues up a command that will move data from the device output buffer to host memory.
<code>rasclib_algorithm_go</code>	Queues up a command to tell the bitstream compute engine to begin computation.
<code>rasclib_algorithm_commit</code>	Causes the <code>rasclib</code> library to send all queued commands to the kernel driver for execution on the device.
<code>rasclib_algorithm_wait</code>	Blocks until all commands that have been sent to the driver have completed execution.
<code>rasclib_algorithm_close</code>	Releases all host resources associated with the algorithm.
<code>rasclib_algorithm_reg_multi_cast</code>	Writes the same data word to all FPGAs that constitute the algorithm.
<code>rasclib_algorithm_reg_read</code>	Reads data from the registers of the FPGAs.
<code>rasclib_algorithm_reg_write</code>	Writes data to the registers of the FPGAs.
<code>rasclib_algorithm_exception_handler_register</code>	Registers an application exception handler for application-defined FPGA logic exception events

Table 4-1 Abstraction Layer Function Definitions - Summary (**continued**)

Function	Description
<code>rasclib_algorithm_malloc</code>	Allocates contiguous memory for an algorithm.
<code>rasclib_algorithm_mfree</code>	Frees contiguous memory for an algorithm.
<code>rasclib_cop_open</code>	Notifies <code>rasclib</code> that the user wants to use one of the reserved devices loaded with the bitstream associated with <code>al_id</code> .
<code>rasclib_cop_send</code>	Queues up a command that will move data from host memory to the device input buffer identified by <code>algo_array_id</code> .
<code>rasclib_cop_receive</code>	Queues up a command that will move data from the device output buffer to host memory.
<code>rasclib_cop_go</code>	Queues up a command to tell the bitstream compute engine to begin computation.
<code>rasclib_cop_commit</code>	Causes <code>rasclib</code> to send all queued commands to the kernel driver for execution on the device
<code>rasclib_cop_wait</code>	Blocks until all commands that have been sent to the kernel driver have completed execution.
<code>rasclib_cop_close</code>	Releases all host resources associated with the algorithm.
<code>rasclib_cop_reg_read</code>	Reads data from the FPGA register.
<code>rasclib_cop_reg_write</code>	Writes data to the FPGA register.
<code>rasclib_cop_exception_handler_register</code>	Registers an application exception handler for application-defined FPGA logic exception events.
<code>rasclib_perror</code>	Prints error message.
<code>rasclib_error_dump</code>	Prints all error messages.
<code>rasclib_cop_malloc</code>	Allocates contiguous memory for a COP.
<code>rasclib_cop_mfree</code>	Frees contiguous memory for a COP.

rasclib_resource_reserve Function

The `rasclib_resource_reserve` function is, as follows:

SYNOPSIS

C:

```
int rasclib_resource_reserve(int ncops, const char *reserve_name);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_RESOURCE_RESERVE_F(NCOPS, RESERVE_NAME)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_RESOURCE_RESERVE_F
    INTEGER(KIND=4), INTENT(IN) :: NCOPS
    CHARACTER(LEN=*), INTENT(IN) :: RESERVE_NAME
  END FUNCTION RASCLIB_RESOURCE_RESERVE_F
END INTERFACE
```

DESCRIPTION

`rasclib_resource_reserve()` sends a request to the device manager to reserve a number of devices for the calling application. The application must make only a single call to `rasclib_resource_alloc()` (though this is not enforced) to avoid deadlock conditions.

`ncops` is the number of devices to reserve. `reserve_name` is the name that the device manager will give to the reservation pool that contains the reserved devices. If `reserve_name` is `NULL`, `rasclib` will generate a reservation pool name with the following call:

```
snprintf(resrv_name, sizeof(resrv_name), "app_%d", getpid());
```

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_MULTI_ALLOC` if `rasclib_resource_reserve()` is called when there is a current reservation with devices still un-released.

Returns `RASCLIB_NODEVICE_REQUESTED` if `ncops == 0`

Returns `RASCLIB_RESERVE_FAILED` if the device manager cannot satisfy the reservation request.

rasclib_resource_release Function

The `rasclib_resource_release` is, as follows:

SYNOPSIS

C:

```
int rasclib_resource_release(int ncops, const char *reserve_name);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_RESOURCE_RELEASE_F(NCOPS, RESERVE_NAME)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_RESOURCE_RELEASE_F
    INTEGER(KIND=4), INTENT(IN) :: NCOPS
    CHARACTER(LEN=*), INTENT(IN) :: RESERVE_NAME
  END FUNCTION RASCLIB_RESOURCE_RELEASE_F
END INTERFACE
```

DESCRIPTION

`rasclib_resource_release()` releases the reservation on a number of devices. That is, it moves devices from the reservation pool to the free pool.

`ncops` is the number of devices to release.

`reserve_name` is a character string and is the name of the reservation pool from which to release. If `reserve_name` is `NULL`, then a reserve name is constructed by `rasclib` using the following function call:

```
snprintf(resrv_name, sizeof(resrv_name), "app_%d", getpid());
```

There must be a reservation pool with this name for release to work.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_RELEASE_FAILED` if the device manager cannot satisfy the request.

rasclib_resource_configure Function

The `rasclib_resource_configure` function is, as follows:

SYNOPSIS

C:

```
int rasclib_resource_configure(const char * al_id, int ncops, const
char *reserve_name);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_RESOURCE_CONFIGURE_F(AL_ID, NCOPS, RESERVE_NAME)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_RESOURCE_CONFIGURE_F
    CHARACTER(LEN=*) , INTENT(IN) :: AL_ID
    INTEGER(KIND=4) , INTENT(IN) :: NCOPS
    CHARACTER(LEN=*) , INTENT(IN) :: RESERVE_NAME
  END FUNCTION RASCLIB_RESOURCE_CONFIGURE_F
END INTERFACE
```

DESCRIPTION

`rasclib_resource_configure()` loads the given bitstream into a device and marks the device as in use.

`al_id` is the algorithm identifier as presented to the device manager at bitstream registration time.

`ncops` is the number of devices to be loaded with the given bitstream.

`reserve_name` is a char pointer that gives the name of the reservation pool from which to choose the device to be loaded. If `reserve_name` is NULL, `rasclib` will build a reservation pool name using the following call:

```
snprintf(resrv_name, sizeof(resrv_name), "app_%d", getpid());
```

There must be a reservation pool with the given name for the configure to succeed.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_ALLOCATION_FAILED` if the device manager cannot satisfy the request.

rasclib_resource_return Function

The `rasclib_resource_return` function is, as follows:

SYNOPSIS

C:

```
int rasclib_resource_return(const char *al_id, int ncops);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_RESOURCE_RETURN_F(AL_ID, NCOPS)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_RESOURCE_RETURN_F
    CHARACTER(LEN=*) , INTENT(IN) :: AL_ID
    INTEGER(KIND=4) , INTENT(IN) :: NCOPS
  END FUNCTION RASCLIB_RESOURCE_RETURN_F
END INTERFACE
```

DESCRIPTION

`rasclib_resource_return()` returns a configured device to its reservation pool. It is still reserved, but can no longer be used by the application until it is again configured.

`al_id` is the algorithm identifier as presented to the device manager at bitstream registration time.

`ncops` is the number of devices that are loaded with the given bitstream that will be returned to their reservation pool.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_FREE_FAILED` if the device manager cannot satisfy the request.

rasclib_resource_alloc Function

Note: This routine has been deprecated and will be removed in a future RASC release. You should use `rasclib_resource_reserve` and `rasclib_resource_configure` in its place.

The `rasclib_resource_alloc` function is, as follows:

SYNOPSIS

C:

```
int rasclib_resource_alloc(rasclib_algorithm_request_t *al_request, int
num_cops);
```

Fortran90:

No Fortran interface declaration exists.

Use `rasclib_resource_reserve_f` and `rasclib_resource_configure_f` in place of this function.

DESCRIPTION

`rasclib_resource_alloc()` sends a request to the device manager to allocate/reserve devices for the application. The application must make only a single call to `rasclib_resource_alloc()` (though this is not enforced) to avoid deadlock conditions.

`al_request` is a pointer to an array of `rasclib_algorithm_request_t` requests. A `rasclib_algorithm_request_t` is defined as:

```
typedef struct rasclib_algorithm_request_s {
char * algorithm_id;
int num_devices;
}rasclib_algorithm_request_t;
```

The `algorithm_id` identifies which bitstream should be loaded. It is given to the device manager at bitstream registration and can be obtained from the device manager user interface.

The `num_devices` argument tells the device manager how many devices should be loaded with the bitstream.

`num_cops` is the number of elements in the array of requests.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_NODEVICE_REQUESTED` if no devices are requested.

Returns `RASCLIB_RESERVE_FAILED` if the device reservation fails.

Returns `RASCLIB_ALLOCATION_FAILED` if the device allocation fails.

`rasclib_resource_free` Function

Note: This routine has been deprecated and will be removed in a future RASC release. You should use `rasclib_resource_return` and `rasclib_resource_release` in its place.

The `rasclib_resource_free` function is, as follows:

SYNOPSIS

C:

```
int rasclib_resource_free(rasclib_algorithm_request_t *al_request, int
num_cops);
```

Fortran90:

No Fortran interface declaration exists.

Use `rasclib_resource_return_f` and `rasclib_resource_release_f` in place of this function.

DESCRIPTION

`rasclib_resource_free()` returns devices to the device manager's free pool. The arguments and their meaning are the same as the `rasclib_resource_alloc()` function.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

rasclib_algorithm_malloc Function

The `rasclib_algorithm_malloc` function is, as follows:

SYNOPSIS

C:

```
void *rasclib_algorithm_malloc(int algo_desc, long size, unsigned flags)
```

DESCRIPTION

`rasclib_algorithm_malloc()` allocates a block of contiguous memory for each COP participating in the algorithm and returns a pointer to those blocks of memory. It is intended that the user application use this pointer to send and receive data to/from the algorithm.

`algo_desc` is the algorithm descriptor returned by `rasclib_algorithm_open`. `size` is the size in bytes of the data area. `flags` can be one of the following:

- `RASCLIB_HUGE_ALLOC` tells `rasclib` to use huge pages.
- `RASCLIB_STANDARD_ALLOC` tells `rasclib` to use standard virtual memory pages.

DIAGNOSTICS

`RASCLIB_HUGE_MAP`: `rasclib` was unable to map a huge page.

`RASCLIB_MALLOC_FAILED`: `rasclib` was unable to allocate an internal data structure.

`NULL ptr`: `rasclib` was unable to malloc memory.

rasclib_algorithm_mfree Function

The `rasclib_algorithm_mfree` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_mfree(int algo_desc, void *ptr);
```

DESCRIPTION

`rasclib_algorithm_mfree()` returns the memory pointed to by `ptr` to the system. If the memory is normal memory, it is returned with the `libc free()` function. If the memory is a huge page, it is returned by unmapping the page. Any references to an unmapped page result in a `SEGV`.

`algo_desc` is an algorithm descriptor, returned from `rasclib_algorithm_open()`.
`ptr` is the pointer to the memory to be freed.

This routine is thread safe.

DIAGNOSTICS

`RASCLIB_SUCCESS` for successful completion.

rasclib_algorithm_open Function

The `rasclib_algorithm_open` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_open(rasclib_algorithm_id_t al_id, io_control);
```

Fortran90:

INTERFACE

FUNCTION RASCLIB_ALGORITHM_OPEN_F(AL_ID, IO_CONTROL)

IMPLICIT NONE

INTEGER(KIND=4) :: RASCLIB_ALGORITHM_OPEN_F

```

CHARACTER(LEN=*) , INTENT(IN) :: AL_ID
INTEGER(KIND=4) , INTENT(IN) :: IO_CONTROL
END FUNCTION RASCLIB_ALGORITHM_OPEN_F
END INTERFACE

```

DESCRIPTION

`rasclib_algorithm_open()` notifies `rasclib` that the user wants to use all of the reserved devices loaded with the bitstream associated with `al_id` as a single logical device. The `rasclib` library initializes all necessary data structures and returns a small integer referred to as an algorithm descriptor. This is similar in concept to a UNIX file descriptor. It is used internally in the library to identify a particular algorithm.

`al_id` is the same as the algorithm id in the `allocate` and `free` routines.

`io_control` indicates the mode of data transfer that is performed. Permitted values are `RASCLIB_BUFFERED_IO`, `RASCLIB_DIRECT_IO`, or `RASCLIB_IGNORE_PARITY`.

If the `RASCLIB_IGNORE_PARITY` flag is set, the driver will not kill the process if a parity error is sensed. When the flag is **not** set (the default case), if a process attempts an SRAM memory access that results in a parity error, the process is killed.

On the screen, you will only see "**Killed**" printed on their screen. The syslog will have a fairly detailed record of what happened, what the error is classified as (correctable or uncorrectable) and if the process was killed, similar to the following:

```

Sep 14 16:05:03 kernel: RASC Error Exception Event:
Sep 14 16:05:03 kernel:     nasid                9
Sep 14 16:05:03 kernel:     PID                21025
Sep 14 16:05:03 kernel:     cm_status          0x9e08040f20200000
Sep 14 16:05:03 kernel:     cm_error_status    0x0000000000800000
Sep 14 16:05:03 kernel:     cm_error_detail    0x0000000000000000
Sep 14 16:05:03 kernel:     cm_error_detail1   0x0000000000000000
Sep 14 16:05:03 kernel: RASC Error Exception - uncorrectable error:
process 21025 killed

```

DIAGNOSTICS

Returns `RASCLIB_FAIL` on failure, an algorithm descriptor (which is greater than or equal to zero) on success.

rasclib_algorithm_send Function

The `rasclib_algorithm_send` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_send(int algo_desc, rasclib_algo_array_id_t
algo_array_id, void *buf, int count);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_ALGORITHM_SEND_F(ALGO_DESC, ALGO_ARRAY_ID, BUF, COUNT)
  IMPLICIT NONE
  INTEGER(KIND=4) :: RASCLIB_ALGORITHM_SEND_F
  INTEGER(KIND=4), INTENT(IN) :: ALGO_DESC
  CHARACTER(LEN=*), INTENT(IN) :: ALGO_ARRAY_ID
  CHARACTER(LEN=*), INTENT(IN) :: BUF
  INTEGER(KIND=4), INTENT(IN) :: COUNT
  END FUNCTION RASCLIB_ALGORITHM_SEND_F
END INTERFACE
```

DESCRIPTION

`rasclib_algorithm_send()` queues up a command that will move data from host memory to the device input buffer identified by `algo_array_id`.

`algo_desc` is the value returned by the `rasclib_algorithm_open()` function .

`algo_array_id` identifies the input data area and is specified in the configuration file associated with the bitstream.

`buf` is a pointer to the data in host memory.

`count` is the number of bytes in the data. Count must be an integral multiple of the size of the data area on the device.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_ARRAY_NOTFOUND` if the input array name cannot be found in the configuration data.

Returns `RASCLIB_READONLY_SRAM` when it is asked to send data to a read-only SRAM.

Returns `RASCLIB_ASYMETRIC_DATA` when the data size of the send and receive do not match up.

rasclib_algorithm_get_num_cops Function

The `rasclib_algorithm_get_num_cops` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_get_num_cops(int al_desc);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_ALGORITHM_GET_NUM_COPS_F(AL_DESC)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_ALGORITHM_GET_NUM_COPS_F
    INTEGER(KIND=4), INTENT(IN) :: AL_DESC
  END FUNCTION RASCLIB_ALGORITHM_GET_NUM_COPS_F
END INTERFACE
```

DESCRIPTION

`rasclib_algorithm_get_num_cops` gets the number of COPs that are participating in the algorithm referred to by `al_desc`.

DIAGNOSTICS

Returns the number of COPs involved in the algorithm or `RASCLIB_FAIL` on failure.

rasclib_algorithm_receive Function

The `rasclib_algorithm_receive` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_receive(int algo_desc, rasclib_algo_array_id_t
algo_array_id, void *buf, int count);
```

Fortran90:

```
INTERFACE
FUNCTION RASCLIB_ALGORITHM_RECEIVE_F(ALGO_DESC, ALGO_ARRAY_ID, BUF, COUNT)
IMPLICIT NONE
INTEGER(KIND=4) :: RASCLIB_ALGORITHM_RECEIVE_F
INTEGER(KIND=4), INTENT(IN) :: ALGO_DESC
CHARACTER(LEN=*), INTENT(IN) :: ALGO_ARRAY_ID
CHARACTER(LEN=*), INTENT(OUT) :: BUF
INTEGER(KIND=4), INTENT(IN) :: COUNT
END FUNCTION RASCLIB_ALGORITHM_RECEIVE_F
END INTERFACE
```

DESCRIPTION

`rasclib_algorithm_receive()` queues up a command that will move data from the device output buffer to host memory.

`algo_desc` is the value returned by the `rasclib_algorithm_open()` function.

`algo_array_id` identifies the output data area and is specified in the configuration file associated with the bitstream.

`buf` is a pointer to the data buffer in host memory that will receive the data.

`count` is the number of bytes in the data. Count must be an integral multiple of the size of the data area on the device.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_ARRAY_NOTFOUND` if the input array name cannot be found in the configuration data.

Returns `RASCLIB_WRITEONLY_SRAM` when it is asked to receive data from a write-only SRAM.

Returns `RASCLIB_ASYMETRIC_DATA` when the data size of the send and receive do not match up.

rasclib_algorithm_go Function

The `rasclib_algorithm_go` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_go(int al_desc);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_ALGORITHM_GO_F(AL_DESC)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_ALGORITHM_GO_F
    INTEGER(KIND=4), INTENT(IN) :: AL_DESC
  END FUNCTION RASCLIB_ALGORITHM_GO_F
END INTERFACE
```

DESCRIPTION

`rasclib_algorithm_go()` issues a command to tell the bitstream compute engine to begin computation.

`al_desc` is the value returned by `rasclib_algorithm_open()`.

DIAGNOSTICS

Returns `RASCLIB_FAIL` on failure, `RASCLIB_SUCCESS` on success.

rasclib_algorithm_commit Function

The `rasclib_algorithm_commit` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_commit(int al_desc,  
void(*rasclib_algo_callback)(int, char*, unsigned long));
```

Fortran90:

```
INTERFACE  
FUNCTION RASCLIB_ALGORITHM_COMMIT_F(AL_DESC, RASCLIB_ALGO_CALLBACK) RESULT (RES)  
IMPLICIT NONE  
INTEGER(KIND=4), INTENT(IN) :: AL_DESC  
EXTERNAL :: RASCLIB_ALGO_CALLBACK  
OPTIONAL :: RASCLIB_ALGO_CALLBACK  
INTEGER(KIND=4) :: RES  
END FUNCTION RASCLIB_ALGORITHM_COMMIT_F  
END INTERFACE
```

DESCRIPTION

`rasclib_algorithm_commit()` causes `rasclib` to send all queued commands to the kernel driver for execution on the device. The commands are sent as a single command list. Until a commit is done, all user buffers are still in use and may not be reused without the (almost certain) risk of data corruption.

`al_desc` is the value returned by `rasclib_algorithm_open()`.

`rasclib_algo_callback` is a function provided by the user to handle situations such as data overruns or data underruns. It is called by `rasclib` when user intervention is required. Currently, it is not called, therefore, this parameter should always be `NULL`.

DIAGNOSTICS

Returns `RASCLIB_FAIL` on failure, `RASCLIB_SUCCESS` on success.

`rasclib_algorithm_wait` Function

The `rasclib_algorithm_wait` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_wait(int al_desc);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_ALGORITHM_WAIT_F(AL_DESC)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_ALGORITHM_WAIT_F
    INTEGER(KIND=4), INTENT(IN) :: AL_DESC
  END FUNCTION RASCLIB_ALGORITHM_WAIT_F
END INTERFACE
```

DESCRIPTION

`rasclib_algorithm_wait()` blocks until all commands that have been sent to the driver have completed execution. There is an implied commit associated with the wait. That is, before the thread of execution waits, it sends all queued commands to the kernel driver for execution (see `rasclib_algorithm_commit`). After the wait call returns, all outstanding data transfers have occurred and the application may do any buffer management that it may require.

`al_desc` is the value returned by the `rasclib_algorithm_open()` function.

DIAGNOSTICS

Returns `RASCLIB_FAIL` on failure, `RASCLIB_SUCCESS` on success.

rasclib_algorithm_close Function

The `rasclib_algorithm_close` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_close(int al_desc);
```

Fortran90:

```
INTERFACE  
  FUNCTION RASCLIB_ALGORITHM_CLOSE_F(AL_DESC)  
    IMPLICIT NONE  
    INTEGER(KIND=4) :: RASCLIB_ALGORITHM_CLOSE_F  
    INTEGER(KIND=4), INTENT(IN) :: AL_DESC  
  END FUNCTION RASCLIB_ALGORITHM_CLOSE_F  
END INTERFACE
```

DESCRIPTION

`rasclib_algorithm_close()` releases all host resources associated with the algorithm. It does not release any allocated devices (see `rasclib_resource_free`).

`al_desc` is the value returned by the `rasclib_algorithm_open()` function.

DIAGNOSTICS

Returns the value of `al_desc`.

rasclib_algorithm_reg_multi_cast Function

The `rasclib_algorithm_reg_multi_cast` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_reg_multi_cast(int al_desc, char *alg_reg,
unsigned long *data, int ndata_elements, unsigned flags);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_ALGORITHM_REG_MULTI_CAST_F(AL_DESC, ALG_REG, DATA, NDATA_ELEMENTS, FLAGS)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_ALGORITHM_REG_MULTI_CAST_F
    INTEGER(KIND=4), INTENT(IN) :: AL_DESC
    CHARACTER(LEN=*), INTENT(IN) :: ALG_REG
    INTEGER(KIND=8), DIMENSION(*), INTENT(IN) :: DATA
    INTEGER(KIND=4), INTENT(IN) :: NDATA_ELEMENTS
    INTEGER(KIND=4), INTENT(IN) :: FLAGS
  END FUNCTION RASCLIB_ALGORITHM_REG_MULTI_CAST_F
END INTERFACE
```

DESCRIPTION

`rasclib_algorithm_reg_multi_cast` sends the same data to all of the FPGAs that implement the algorithm.

`al_desc` is the algorithm descriptor returned from the `rasclib_algorithm_open` function.

`alg_reg` is the name of the register.

`data` is a pointer to the data to be written to the register.

`flags` can be either `RASCLIB_QCMD` or `RASCLIB_IMMED_CMD`.

If `flags` is `RASCLIB_QCMD`, the command is queued in `rasclib` to be sent to the driver to be executed in order within the command list.

If `flags` is `RASCLIB_IMMED_CMD`, the register writes are done before `rasclib_algorithm_reg_multi_cast` returns to the user.

The `rasclib_algorithm_reg_multi_cast` function supports fat registers. A *fat register* is register that spans multiple FPGA registers. Their length in bits must be an integral multiple of 64 and they must begin at the beginning of a physical register. Therefore, a fat register exactly spans an integral number of physical registers. It is assumed that the data array supplied is large enough to hold all the data for a fat register. The `rasclib` library does not support partial reads or writes to fat registers.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS`. Exits with exit status `RASCLIB_MALLOC_FAILED` if `rasclib` cannot allocate memory.

rasclib_algorithm_reg_read Function

The `rasclib_algorithm_reg_read` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_reg_read(int al_desc, char *alg_reg, unsigned
long *data, int ndata_elements, unsigned flags);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_ALGORITHM_REG_READ_F(AL_DESC, ALG_REG, DATA, NDATA_ELEMENTS, FLAGS)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_ALGORITHM_REG_READ_F
    INTEGER(KIND=4), INTENT(IN) :: AL_DESC
    CHARACTER(LEN=*), INTENT(IN) :: ALG_REG
    INTEGER(KIND=8), DIMENSION(*), INTENT(OUT) :: DATA
    INTEGER(KIND=4), INTENT(IN) :: NDATA_ELEMENTS
    INTEGER(KIND=4), INTENT(IN) :: FLAGS
  END FUNCTION RASCLIB_ALGORITHM_REG_READ_F
END INTERFACE
```

DESCRIPTION

`rasclib_algorithm_reg_read` reads a register from each of the FPGAs that implement an algorithm.

It is assumed that there is enough space in the array pointed to by `data` to handle all of the data. See “`rasclib_algorithm_get_num_cops` Function” on page 97.

`al_desc` is the algorithm descriptor returned from the `rasclib_algorithm_open` function.

`alg_reg` is the name of the register.

`data` is a pointer to the data where the data read from the register will be placed.

`flags` can be either `RASCLIB_QCMD` or `RASCLIB_IMMED_CMD`.

If `flags` is `RASCLIB_QCMD`, the command is queued in `rasclib` to be sent to the driver to be executed in order within the command list.

If `flags` is `RASCLIB_IMMED_CMD`, the register reads are done before `rasclib_algorithm_reg_read` returns to the user.

The `rasclib_algorithm_reg_read` function supports fat registers. A fat register is a register that spans multiple FPGA registers. Their length in bits must be an integral multiple of 64 and they must begin at the beginning of a physical register. Therefore, a fat register exactly spans an integral number of physical registers. It is assumed that the data array supplied is large enough to hold all the data for a fat register. The `rasclib` library does not support partial reads or writes to fat registers

DIAGNOSTICS

Returns `RASCLIB_SUCCESS`.

Returns `RASCLIB_NO_REG_DATA`: Couldn't find the named register in config file

Returns `RASCLIB_DATA_SIZE_ERROR`: Size mismatch in register access request

Returns `RASCLIB_INVALID_FLAG`: Invalid flag passed to `rasclib` routine

rasclib_algorithm_reg_write Function

The `rasclib_algorithm_reg_write` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_reg_write(int al_desc, char *alg_reg, unsigned
long *data, int ndata_elements, unsigned flags);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_ALGORITHM_REG_WRITE_F(AL_DESC, ALG_REG, DATA, NDATA_ELEMENTS, FLAGS)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_ALGORITHM_REG_WRITE_F
    INTEGER(KIND=4), INTENT(IN) :: AL_DESC
    CHARACTER(LEN=*), INTENT(IN) :: ALG_REG
    INTEGER(KIND=8), DIMENSION(*), INTENT(IN) :: DATA
    INTEGER(KIND=4), INTENT(IN) :: NDATA_ELEMENTS
    INTEGER(KIND=4), INTENT(IN) :: FLAGS
  END FUNCTION RASCLIB_ALGORITHM_REG_WRITE_F
END INTERFACE
```

DESCRIPTION

`rasclib_algorithm_reg_write` writes the supplied data to the FPGAs that implement the algorithm. It is assumed that there is enough data supplied to satisfy one register write for each FPGA in the algorithm (see “`rasclib_algorithm_get_num_cops` Function” on page 97).

`al_desc` is the algorithm descriptor returned from the `rasclib_algorithm_open` function.

`alg_reg` is the name of the register.

`data` is a pointer to the data to be written to the register.

`flags` can be either `RASCLIB_QCMD` or `RASCLIB_IMMED_CMD`.

If `flags` is `RASCLIB_QCMD`, the command is queued in `rasclib` to be sent to the driver to be executed in order within the command list.

If `flags` is `RASCLIB_IMMED_CMD`, the register writes are done before `rasclib_algorithm_reg_write` returns to the user.

The `rasclib_algorithm_reg_write` function supports fat registers. A fat register is a register that spans multiple FPGA registers. Their length in bits must be an integral multiple of 64 and they must begin at the beginning of a physical register. Therefore, a fat register exactly spans an integral number of physical registers. It is assumed that the data array supplied is large enough to hold all the data for fat register. The `rasclib` library does not support partial reads or writes to fat registers

DIAGNOSTICS

Returns `RASCLIB_SUCCESS`.

Returns `RASCLIB_NO_REG_DATA`: Couldn't find the named register in config file

Returns `RASCLIB_DATA_SIZE_ERROR`: Size mismatch in register access request

Returns `RASCLIB_INVALID_FLAG`: Invalid flag passed to `rasclib` routine

rasclib_algorithm_exception_handler_register Function

The `rasclib_algorithm_exception_handler_register` function is, as follows:

SYNOPSIS

C:

```
int rasclib_algorithm_exception_handler_register(int algo_desc, int
(*exception_handler)(void))
```

Fortran90:

```
INTERFACE
FUNCTION RASCLIB_ALGORITHM_EXCEPTION_REGISTER_F(AL_DESC, EXCEPTION_HANDLER) RESULT (RES)
IMPLICIT NONE
INTEGER(KIND=4), INTENT(IN) :: AL_DESC
EXTERNAL :: EXCEPTION_HANDLER
INTEGER(KIND=4) :: RES
END FUNCTION RASCLIB_ALGORITHM_EXCEPTION_REGISTER_F
END INTERFACE
```

DESCRIPTION

`rasclib_algorithm_exception_handler_register` function registers an application exception handler that will be invoked when the FPGA logic drives the `ALG_EVENT0` signal.

It is intended for application-defined FPGA logic exception events.

The FPGA logic is paused for the duration of the exception processing. The exception handler returns 0 if the FPGA logic can continue or 1 to stop the current logic run.

The algorithm registration function registers the handler for each FPGA represented by `algo_desc`.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_UNSUPPORTED` if Core Services does not support exception processing. This is the case for releases prior to the RASC 2.1 release.

rasclib_cop_open Function

The `rasclib_cop_open` function is, as follows:

SYNOPSIS

C:

```
int rasclib_cop_open(rasclib_algorithm_id_t as_id,io_control);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_COP_OPEN_F(AS_ID, IO_CONTROL)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_COP_OPEN_F
    CHARACTER(LEN=*) , INTENT(IN) :: AS_ID
    INTEGER(KIND=4) , INTENT(IN) :: IO_CONTROL
  END FUNCTION RASCLIB_COP_OPEN_F
END INTERFACE
```

DESCRIPTION

`rasclib_cop_open()` notifies `rasclib` that the user wants to use one of the reserved devices loaded with the bitstream associated with `al_id`. The `rasclib` library initializes all necessary data structures and returns a small integer referred to as an COP descriptor. This is similar in concept to a UNIX file descriptor. It is used internally in the library to identify a particular COP.

`al_id` is the same as the algorithm ID in the `allocate` and `free` routines.

`io_control` indicates the mode of data transfer that is performed. Permitted values are `RASCLIB_BUFFERED_IO` or `RASCLIB_DIRECT_IO`.

DIAGNOSTICS

Returns a `cop_desc` ≥ 0 on success.

Returns `RASCLIB_NO_HANDLE` if no cop handle can be found for `alg_id`.

Returns `RASCLIB_NO_DESC_NUM` if no descriptor can be found.

Returns `RASCLIB_DEV_OPEN_FAILED` if an open system call failed.

Returns `RASCLIB_MALLOC_FAILED` if a `malloc` call failed.

rasclib_cop_send Function

The `rasclib_cop_send` function is, as follows:

SYNOPSIS

C:

```
int rasclib_cop_send(int cop_desc, rasclib_algo_array_id_t
algo_array_id, void *buf, int count);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_COP_SEND_F(COP_DESC, ALGO_ARRAY_ID, BUF, COUNT)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_COP_SEND_F
    INTEGER(KIND=4), INTENT(IN) :: COP_DESC
    CHARACTER(LEN=*), INTENT(IN) :: ALGO_ARRAY_ID
    CHARACTER(LEN=*), INTENT(IN) :: BUF
    INTEGER(KIND=4), INTENT(IN) :: COUNT
  END FUNCTION RASCLIB_COP_SEND_F
END INTERFAC
```

DESCRIPTION

`rasclib_cop_send()` queues up a command that will move data from host memory to the device input buffer identified by `algo_array_id`.

`cop_desc` is the value returned by `rasclib_cop_open()`.

`algo_array_id` identifies the input data area and is specified in the configuration file associated with the bitstream.

`buf` is a pointer to the data in host memory.

`count` is the number of bytes in the data. Count must be equal to the size of the data area on the device.

DIAGNOSTICS

Returns RASCLIB_SUCCESS when all goes well.

Returns RASCLIB_NO_DESC if cop_desc cannot be translated into a descriptor pointer.

Returns RASCLIB_MALLOC_FAILED if a malloc call failed.

Returns RASCLIB_NO_ARRAY if the input array name cannot be found in the configuration data.

Returns RASCLIB_READONLY_SRAM when it is asked to send data to a read-only SRAM.

Returns RASCLIB_SIZE_MISMATCH if the input size does not match the SRAM buffer size.

rasclib_cop_receive Function

The rasclib_cop_receive function is, as follows:

SYNOPSIS

C:

```
int rasclib_cop_receive(int cop_desc, rasclib_algo_array_id_t
algo_array_id, void *buf, int count);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_COP_RECEIVE_F(COP_DESC, ALGO_ARRAY_ID, BUF, COUNT)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_COP_RECEIVE_F
    INTEGER(KIND=4), INTENT(IN) :: COP_DESC
    CHARACTER(LEN=*), INTENT(IN) :: ALGO_ARRAY_ID
    CHARACTER(LEN=*), INTENT(OUT) :: BUF
    INTEGER(KIND=4), INTENT(IN) :: COUNT
  END FUNCTION RASCLIB_COP_RECEIVE_F
END INTERFACE
```

DESCRIPTION

`rasclib_cop_receive()` queues up a command that will move data from the device output buffer to host memory.

`cop_desc` is the value returned by `rasclib_cop_open()`.

`algo_array_id` identifies the output data area and is specified in the configuration file associated with the bitstream.

`buf` is a pointer to the data buffer in host memory that will receive the data.

`count` is the number of bytes in the data. Count must be equal to the size of the data area on the device.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_NO_DESC` if `cop_desc` cannot be translated into a descriptor pointer.

Returns `RASCLIB_MALLOC_FAILED` if a `malloc` call failed.

Returns `RASCLIB_NO_ARRAY` if the input array name cannot be found in the configuration data.

Returns `RASCLIB_WRITEONLY_SRAM` when it is asked to receive data from a write-only SRAM.

Returns `RASCLIB_SIZE_MISMATCH` if the input size does not match the SRAM buffer size.

rasclib_cop_malloc Function

The `rasclib_cop_malloc` function is as follows:

SYNOPSIS

C:

```
void *rasclib_cop_malloc(int cop_desc, long size, unsigned flags);
```

DESCRIPTION

`rasclib_cop_malloc()` allocates a block of memory on the node that is closest (in a topological sense) to the COP in question. If allocating huge pages for contiguous memory, it will allocate exactly one huge page for the COP.

`cop_desc` is the COP descriptor return from `rasclib_cop_open()`. `size` is the size of the data area requested. `flags` can be one of the following:

- `RASCLIB_HUGE_ALLOC` allocate a huge page.
- `RASCLIB_STANDARD_ALLOC` allocate normal memory pages.

DIAGNOSTICS

`RASCLIB_HUGE_MAP` `rasclib` cannot mmap a huge page.

`RASCLIB_MALLOC_FAILED` `rasclib` cannot malloc memory for an internal data structure.

`NULL ptr` `rasclib` cannot malloc the requested space.

rasclib_cop_mfree Function

The `rasclib_cop_mfree` function is, as follows:

SYNOPSIS

C:

```
int rasclib_cop_mfree(int cop_desc, void *ptr);
```

DESCRIPTION

`rasclib_cop_mfree()` returns memory allocated by `rasclib_cop_malloc()` to the system. If the memory is regular memory, it is freed by the libc `free()` call. If the memory is a huge page, it is unmapped. If unmapped memory is referenced, it results in a SEGV.

`cop_desc` is a COP descriptor returned by `rasclib_cop_open()`. `ptr` is a pointer to memory that has been allocated by `rasclib_cop_malloc`. This routine is thread safe.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS`

rasclib_cop_go Function

The `rasclib_cop_go` function is, as follows:

SYNOPSIS

rasclib_algorithm_mfreeC:

```
int rasclib_cop_go(int cop_desc);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_COP_GO_F(COP_DESC)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_COP_GO_F
    INTEGER(KIND=4), INTENT(IN) :: COP_DESC
  END FUNCTION RASCLIB_COP_GO_F
END INTERFACE
```

DESCRIPTION

`rasclib_cop_go()` issues a command to tell the bitstream compute engine to begin computation.

`cop_desc` is the value returned by `rasclib_cop_open()`.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_MALLOC_FAILED` if a `malloc` call failed.

rasclib_cop_commit Function

The `rasclib_cop_commit` function is, as follows:

SYNOPSIS

C:

```
int rasclib_cop_commit (int cop_desc, int(*rasclib_cop_callback)(int,
char *, unsigned long, char *buf));
```

Fortran90

```
INTERFACE
  FUNCTION RASCLIB_COP_COMMIT_F(COP_DESC, RASCLIB_COP_CALLBACK) RESULT (RES)
  IMPLICIT NONE
  INTEGER(KIND=4), INTENT(IN) :: COP_DESC
  EXTERNAL :: RASCLIB_COP_CALLBACK
  OPTIONAL :: RASCLIB_COP_CALLBACK
  INTEGER(KIND=4) :: RES
  END FUNCTION RASCLIB_COP_COMMIT_F
END INTERFACE
```

DESCRIPTION

`rasclib_cop_commit()` causes the `rasclib` library to send all queued commands to the kernel driver for execution on the device. The commands are sent as a single command list. Until a commit call is done, all user buffers are still in use and may not be reused without the (almost certain) risk of data corruption.

`cop_desc` is the value returned by `rasclib_cop_open()`.

`rasclib_cop_callback` is a function provided by the user to handle situations such as data overruns or data underruns. This callback is **ONLY** called for streaming DMA engines, never for block DMAs. The status argument contains either `RASCLIB_STREAM_SEND_COMPLETE` or `RASCLIB_STREAM_RECEIVE_COMPLETE`. The `RASCLIB_STREAM_SEND_COMPLETE` argument says that a stream, identified by the second argument, has completed a SEND call. The `RASCLIB_STREAM_RECEIVE_COMPLETE` argument says that a stream, identified by the second argument, has completed a RECEIVE call. The function should always return either `RASCLIB_CALLBACK_NORMAL` (for a normal return) or `RASCLIB_CALLBACK_UNREGISTER` (to unregister the callback so that no further calls

will be made). Below is a simple example, for informational purposes only, that illustrates how the callback is used:

```
pthread_mutex_t callback_lock = PTHREAD_MUTEX_INITIALIZER;

int alg12_callback(int cop_desc, char *name, char *buf, unsigned long
status) {
    int i;
    char *p = buf;
    static done_wr = 0, done_rd = 0;

    pthread_mutex_lock(&callback_lock);
    if (done_rd && done_wr){
        pthread_mutex_unlock(&callback_lock);
        return RASCLIB_CALLBACK_UNREGISTER;
    }
    buf += (SIZE/2 * sizeof(unsigned long));
    if (status == RASCLIB_STREAM_SEND_COMPLETE) {
        rasclib_cop_send(cop_desc, name, buf, (SIZE/2 *
sizeof(unsigned long)));
        rasclib_cop_commit(cop_desc, alg12_callback);
        done_rd = 1;
    } else {
        rasclib_cop_receive(cop_desc, name, buf, (SIZE/2 *
sizeof(unsigned long)));
        rasclib_cop_commit(cop_desc, alg12_callback);
        done_wr = 1;
    }
    pthread_mutex_unlock(&callback_lock);
    return RASCLIB_CALLBACK_NORMAL;
}
```

DIAGNOSTICS

Returns RASCLIB_SUCCESS when all goes well.

Returns RASCLIB_NO_DESC if cop_desc cannot be translated into a descriptor pointer.

Returns RASCLIB_MALLOC_FAILED if a malloc call failed.

rasclib_cop_wait Function

The `rasclib_cop_wait` function is, as follows:

SYNOPSIS

C:

```
int rasclib_cop_wait(int cop_desc);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_COP_WAIT_F(COP_DESC)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_COP_WAIT_F
    INTEGER(KIND=4), INTENT(IN) :: COP_DESC
  END FUNCTION RASCLIB_COP_WAIT_F
END INTERFACE
```

DESCRIPTION

`rasclib_cop_wait()` blocks until all commands that have been sent to the kernel driver have completed execution. There is an implied commit associated with the wait. That is, before the thread of execution waits, it sends all queued commands to the kernel driver for execution (see `rasclib_cop_commit`). After the wait call returns, all outstanding data transfers have occurred and the application may do any buffer management that it may require.

`cop_desc` is the value returned by `rasclib_cop_open()`.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_NO_DESC` if `cop_desc` cannot be translated into a descriptor pointer.

Returns `RASCLIB_MALLOC_FAILED` if a `malloc` call failed.

rasclib_cop_close Function

The `rasclib_cop_close` function is, as follows:

SYNOPSIS

C:

```
int rasclib_cop_close(int cop_desc);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_COP_CLOSE_F(COP_DESC)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_COP_CLOSE_F
    INTEGER(KIND=4), INTENT(IN) :: COP_DESC
  END FUNCTION RASCLIB_COP_CLOSE_F
END INTERFACE
```

DESCRIPTION

`rasclib_cop_close()` releases all host resources associated with the algorithm. It does not release any allocated devices (see `rasclib_resource_free`).

DIAGNOSTICS

`cop_desc` is the value returned by `rasclib_cop_open()`.

rasclib_cop_reg_read Function

The `rasclib_cop_reg_read` function is, as follows:

SYNOPSIS

C:

```
int rasclib_cop_reg_read(int cop_desc, char *alg_reg, unsigned long
*data, int ndata_elements, unsigned flags);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_COP_REG_READ_F(COP_DESC, ALG_REG, DATA, NDATA_ELEMENTS, FLAGS)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_COP_REG_READ_F
    INTEGER(KIND=4), INTENT(IN) :: COP_DESC
    CHARACTER(LEN=*), INTENT(IN) :: ALG_REG
    INTEGER(KIND=8), DIMENSION(*), INTENT(OUT) :: DATA
    INTEGER(KIND=4), INTENT(IN) :: NDATA_ELEMENTS
    INTEGER(KIND=4), INTENT(IN) :: FLAGS
  END FUNCTION RASCLIB_COP_REG_READ_F
END INTERFACE
```

DESCRIPTION

`rasclib_cop_reg_read()` reads a register from a single FPGA, as specified by `cop_desc`.

It is assumed that there is enough space in the array pointed to by `data` to handle all of the data.

`cop_desc` is the cop descriptor returned from the `rasclib_cop_open` function.

`alg_reg` is the name of the register.

`data` is a pointer to the data where the data read from the register will be placed.

`flags` can be either `RASCLIB_QCMD` or `RASCLIB_IMMED_CMD`.

If `flags` is `RASCLIB_QCMD`, the command is queued in `rasclib` to be sent to the driver to be executed in order within the command list.

If `flags` is `RASCLIB_IMMED_CMD`, the register reads are done before `rasclib_cop_reg_read` returns to the user.

The `rasclib_cop_reg_read` function supports fat registers. A fat register is register that spans multiple FPGA registers. Their length in bits must be an integral multiple of 64 and they must begin at the beginning of a physical register. Therefore, a fat register exactly spans an integral number of physical registers. It is assumed that the data array supplied is large enough to hold all the data for a fat register. The `rasclib` library does not support partial reads or writes to fat registers

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_INVALID_FLAG` if the flag value is invalid.

rasclib_cop_reg_write Function

The `rasclib_cop_reg_write` function is, as follows:

SYNOPSIS

C:

```
int rasclib_cop_reg_write(int cop_desc, char *alg_reg, unsigned long
*data, int ndata_elements, unsigned flags);
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_COP_REG_WRITE_F(COP_DESC, ALG_REG, DATA, NDATA_ELEMENTS, FLAGS)
    IMPLICIT NONE
    INTEGER(KIND=4) :: RASCLIB_COP_REG_WRITE_F
    INTEGER(KIND=4), INTENT(IN) :: COP_DESC
    CHARACTER(LEN=*), INTENT(IN) :: ALG_REG
    INTEGER(KIND=8), DIMENSION(*), INTENT(IN) :: DATA
    INTEGER(KIND=4), INTENT(IN) :: NDATA_ELEMENTS
    INTEGER(KIND=4), INTENT(IN) :: FLAGS
  END FUNCTION RASCLIB_COP_REG_WRITE_F
END INTERFACE
```

DESCRIPTION

`rasclib_cop_reg_write()` writes to a register of a single FPGA, as specified by `cop_desc`.

It is assumed that there is enough space in the array pointed to by `data` to fill the register.

`cop_desc` is the cop descriptor returned from the `rasclib_cop_open` function.

`alg_reg` is the name of the register.

`data` is a pointer to the data area where the data to be written resides.

`flags` can be either `RASCLIB_QCMD` or `RASCLIB_IMMED_CMD`.

If `flags` is `RASCLIB_QCMD`, the command is queued in `rasclib` to be sent to the driver to be executed in order within the command list.

If `flags` is `RASCLIB_IMMED_CMD`, the register reads are done before `rasclib_cop_reg_write` returns to the user.

The `rasclib_cop_reg_write` function supports fat registers. A fat register is register that spans multiple FPGA registers. Their length in bits must be an integral multiple of 64 and they must begin at the beginning of a physical register. Therefore, a fat register exactly spans an integral number of physical registers. It is assumed that the data array supplied is large enough to hold all the data for fat register. The `rasclib` library does not support partial reads or writes to fat registers

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_INVALID_FLAG` if the flag value is invalid.

rasclib_cop_exception_handler_register Function

The `rasclib_cop_exception_handler_register` function is, as follows:

SYNOPSIS

C:

```
int rasclib_cop_exception_handler_register(int cop_desc, int
(*exception_handler)(void))
```

Fortran90:

```
INTERFACE
  FUNCTION RASCLIB_COP_EXCEPTION_REGISTER_F(COP_DESC, EXCEPTION_HANDLER) RESULT (RES)
    IMPLICIT NONE
    INTEGER(KIND=4), INTENT(IN) :: COP_DESC
    EXTERNAL :: EXCEPTION_HANDLER
    INTEGER(KIND=4) :: RES
  END FUNCTION RASCLIB_COP_EXCEPTION_REGISTER_F
END INTERFACE
```

DESCRIPTION

`rasclib_cop_exception_handler_register` function registers an application exception handler that is invoked when the FPGA logic drives the `ALG_EVENT0` signal.

It is intended for application-defined FPGA logic exception events.

The FPGA logic is paused for the duration of the exception processing. The exception handler returns 0 if the FPGA logic can continue or 1 to stop the current logic run.

The algorithm registration function registers the handler for each FPGA represented by `cop_desc`.

DIAGNOSTICS

Returns `RASCLIB_SUCCESS` when all goes well.

Returns `RASCLIB_UNSUPPORTED` if Core Services does not support exception processing. This is the case for releases prior to the RASC 2.1 release.

rasclib_perror Function

The `rasclib_perror` function is, as follows:

SYNOPSIS

C:

```
void rasclib_perror(char *string, int ecode);
```

Fortran90:

```
INTERFACE
  SUBROUTINE RASCLIB_PERROR_F(STRING, ECODE)
    IMPLICIT NONE
    CHARACTER(LEN=*) , INTENT(IN) :: STRING
    INTEGER(KIND=4) , INTENT(IN) :: ECODE
  END SUBROUTINE RASCLIB_PERROR_F
ENDINTERFACE
```

DESCRIPTION

`rasclib_perror` prints an error string and a description of the error code.

`string` is a descriptive character string supplied by the caller.

`ecode` is an error code returned by a `rasclib` routine.

rasclib_error_dump Function

The `rasclib_error_dump` function is, as follows:

SYNOPSIS

C:

```
void rasclib_error_dump(void);
```

Fortran90:

```
INTERFACE
  SUBROUTINE RASCLIB_ERROR_DUMP_F
    IMPLICIT NONE
  END SUBROUTINE RASCLIB_ERROR_DUMP_F
END INTERFACE
```

DESCRIPTION

`rasclib_error_dump` prints the last three errors detected by `rasclib`. The `rasclib` library maintains a number of three-deep circular buffer of errors encountered by `rasclib`. One is a global buffer, one is associated with each algorithm, and one is associated with each co-processor (`cop`). `rasclib_error_dump` first prints the global errors, then the algorithm errors, and then the `cop` errors.

How the RASC Abstraction Layer Works

The `rasclib` library provides an asynchronous interface to the FPGA devices. That is, each call into the interface to perform some action on the device (`send`, `go`, `receive`) queues the command to the RASC abstraction layer. Commands are not actually sent to the driver (and, thus, to the device) until a `commit` call is done. At `commit` call time, all queued commands are sent to the driver as a list of commands. The driver sends the commands to the device. This takes advantage of any parallelism that may be available in the DMA engines or compute section. The user application cannot make any assumptions as to when the committed commands have finished executing. The user application must execute a `wait` call to make sure all committed commands have finished. The `wait` call blocks until all outstanding commands have completed. When the `wait` call returns, all user data buffers have either been completely read from or completely filled, depending on their use.

The algorithm layer also exists for ease of use. The application can reserve as many devices as it can get with the `allocate` call, then allow the abstraction layer to use them as efficiently as it can, spreading the data across as many devices as there are available and looping over the devices until all the data is consumed. The primary purpose is related to the above. By spreading the data across all the devices, the abstraction layer implements wide scaling. That is, multiple devices participate in the calculation in parallel. Deep scaling, or chaining the output of one device to the input of another device without having to stage the data through host memory, is left for later.

The input/output data buffers are assumed to be at least 16 KBs in size. Currently, user data requests to the algorithm layer are assumed to be an integral multiple of the data buffer in size and user data requests to the COP layer are assumed to be exactly the size of the data buffer.

So, given the above, here is how it all works. Assume a C-like program:

```
unsigned long A[SIZE],B[SIZE],C[SIZE];
main() {
    comput(A,B,C);
}
```

Where A and B are inputs to some function and C is the output. If the developer were to implement the function “`comput`” in an FPGA, then he would transform the above to:

```
unsigned long A[SIZE],B[SIZE],C[SIZE];
main() {

    int al_desc,
    int num_devices = N;
    int res;

    res = rasclib_resource_reserve(num_devices, NULL);
    if (res != RASCLIB_SUCCESS) {
        fprintf(stderr,"reserve failed at %d: %d\n", __LINE__, res);
        rasclib_perror("reserve", res);
        return 1;
    }

    res = rasclib_resource_configure("compute", num_devices, NULL);
    if (res != RASCLIB_SUCCESS) {
        fprintf(stderr,"configure failed at %d: %d\n", __LINE__, res);
        rasclib_perror("configure", res);
        return 1;
    }
}
```

```

al_desc = rasclib_algorithm_open("compute");
if (al_desc == RASCLIB_FAIL) {
    exit(1);
}

if (rasclib_algorithm_send(al_desc, "input_one", A, SIZE *
sizeof(unsigned long)) != RASCLIB_SUCCESS) {
    exit(1);
}

    if (rasclib_algorithm_send(al_desc, "input_two", B, SIZE *
sizeof(unsigned long)) != RASCLIB_SUCCESS) {
        exit(1);
    }

    if (rasclib_algorithm_go(al_desc) != RASCLIB_SUCCESS) {
        exit(1);
    }

    if (rasclib_algorithm_receive(al_desc, "output", C, SIZE *
sizeof(unsigned long)) != RASCLIB_SUCCESS) {
        exit(1);
    }

    if (rasclib_algorithm_commit(al_desc) != RASCLIB_SUCCESS) {
        exit(1);
    }

    if (rasclib_algorithm_wait(al_desc) != RASCLIB_SUCCESS) {
        exit(1);
    }
    rasclib_resource_return("compute", num_devices);
    rasclib_resource_release(num_devices, NULL);
}

```

The `rasclib_resource_reserve` and `rasclib_resource_configure` calls reserve and configure devices through the device manager. The application must request all of the devices it needs for the life of the application in one call or deadlocks may result. The `rasclib_algorithm_open` tells the abstraction layer to allocate all necessary internal data structures for a logical algorithm. The two `rasclib_algorithm_send` calls pull data down to the input data areas on the COP. The `rasclib_algorithm_go` starts execution of the COP. The `rasclib_algorithm_receive` call pushes the result back out to host memory. The `rasclib_algorithm_commit` causes all of the commands that have been queued up by the previous four calls to be sent to the devices.

The data is broken up on data area size boundaries and spread out over all available COPs. The `rasclib_algorithm_wait` blocks until all the command that were sent to all of the devices are complete, then returns.

The `rasclib_resource_return` and `rasclib_resource_release` calls tell the device manager that the application is done with the devices and that they can be allocated to other applications.

Direct I/O

This section describes RASC direct I/O capabilities and covers the following topics:

- “Direct I/O Overview” on page 129
- “Direct I/O Driver Requirements” on page 129
- “Using Hugepages with rasclib Functions” on page 130

Direct I/O Overview

Direct I/O capabilities in the Reconfigurable Application-Specific Computing (RASC) library (`rasclib`) allows the driver to send data directly from the user space buffers to the RASC coprocessor, without an intervening kernel copy. The main advantage of direct I/O is improved performance.

Direct I/O Driver Requirements

The direct I/O driver requires that all user buffers be aligned on a 128 byte (128B) cacheline boundary. Furthermore, the buffers must be an even multiple of a 128B cacheline in size. The buffers must be mapped onto contiguous (or as contiguous as possible) physical memory. The physical memory that implements the buffer must have no more than 64 instances of memory discontinuity. All send and receive operations must be direct or all operations must be indirect. Mixed I/O mode is not supported.

The `rasclib_open` call has a new argument, as follows:

```
int rasclib_cop_open(char * as_id, unsigned flags)
```

or

```
int rasclib_algorithm_open(char * al_id, unsigned flags)
```

The flags argument can be passed either the `RASCLIB_BUFFERED_IO` or `RASCLIB_DIRECT_IO` flag.

If `rasclib_open` call is passed `RASCLIB_BUFFERED_IO`, everything works, as previously, with the kernel space copy of the user buffer to a kernel direct memory access (DMA) buffer.

If `rasclib_open` call is passed `RASCLIB_DIRECT_IO`, the driver performs direct I/O, enforcing the above restrictions on all send /receive operations associated with this device or algorithm.

Using Hugepages with `rasclib` Functions

In addition, `rasclib` now supports a memory allocation utility that allocates memory that satisfies the driver restrictions. It does this by allocating memory from the hugepages.

In order to use these routines, you must have your system administrator enable the `hugetlbfs` configuration option. The following commands can be run very shortly after boot (or as part of the `init` process):

```
echo 16 > /proc/sys/vm/nr_hugepages
echo 1 > /proc/sys/fs/shm-use-hugepages
echo 1 > /proc/sys/fs/mmap-use-hugepages
mount none -t hugetlbfs /mnt/hugetlb
chmod 0777 /mnt/hugetlb
```

The first line sets the number of hugepages. The value of 16 is an example. In a production environment, you need to change that value to suit your system. The last two lines mount the `hugetlbfs` filesystem and enable general access.

In the example above, it is mounted at `/mnt/hugetlb`. You can mount it anywhere `rasclib` can find it. If it is **not** mounted, the allocation fails.

For more information on hugepages, see `/usr/src/linux-2.6.x.x/Documentation/vm/hugetlbpage.txt`.

The `rasclib` function calls for hugepages are, as follows:

```
void * rasclib_huge_alloc(long size)
void rasclib_huge_free(void *p)
void rasclib_huge_clear()
```

The `rasclib_huge_alloc()` function allocates `size` bytes and returns a `void *` (pointer) to the resulting pages. The `size` argument is rounded up to the next 128B cacheline. It returns a number less than zero (which can be passed to `rasclib_perror()`) on failure.

The `rasclib_huge_free()` function frees the pointer argument it is passed. The space is returned to the internal free list, it is not unmapped. It is coalesced, if possible.

The `rasclib_huge_clear()` function unmaps all hugepages that have been allocated, whether they have been freed or not. Use this routine only when you are finished using hugepages.

For more information on `rasclib`, see Chapter 4, “RASC Abstraction Layer”.

Scaling Applications with RASC

Getting an application to scale well with RASC can require some effort. If the application does not require large bandwidth, it is just a matter of determining what, if any, data dependencies exist and using the algorithm layer to automatically scale across multiple FPGAs. If the application requires maximal bandwidth (it is bandwidth limited), then care should be taken when determining buffer sizes and placement.

The `rasclib_cop_malloc()`, `rasclib_algorithm_malloc()` routines and the associated `*_mfree()` routines aid in this process. This chapter describes some things you need to consider when using these routines.

Data Buffer Size

The buffer size you choose to do sends and receives should match up as closely as possible with the `Hugepagesize` as shown by the `/proc/meminfo` file. An example `/proc/meminfo` file is, as follows:

```
system 56% cat meminfo
MemTotal:      48034320 kB
MemFree:       6848464 kB
Buffers:       208 kB
Cached:        34947536 kB
SwapCached:    0 kB
Active:        25110112 kB
Inactive:      13881616 kB
HighTotal:     0 kB
HighFree:      0 kB
LowTotal:      48034320 kB
LowFree:       6848464 kB
SwapTotal:     69824176 kB
SwapFree:      69824176 kB
Dirty:         32736 kB
Writeback:     0 kB
AnonPages:     3982080 kB
Mapped:        123472 kB
```

```
Slab:          1648864 kB
CommitLimit:  93841328 kB
Committed_AS: 6141376 kB
PageTables:   67520 kB
VmallocTotal: 137416910048 kB
VmallocUsed:  233168 kB
VmallocChunk: 137416675984 kB
HugePages_Total: 0
HugePages_Free: 0
HugePages_Rsvd: 0
Hugepagesize: 262144 kB
```

To avoid wasting space, the buffer sizes should be as close to `Hugepagesize` as possible without exceeding it. The RASC abstraction layer will allocate the memory associated with a given COP as close (in a topology sense) to the COP as it can. It will also make an effort to spread the memory allocation across memory nodes as evenly as possible to avoid over-subscribing the memory bandwidth on a memory node. The `rasclib_cop_malloc()` routine allocates a single huge page for a single COP. Requesting a size that is larger than the huge page size will result in an error.

Requesting a size that is smaller than the huge page size will result in wasted space.

The `rasclib_algorithm_malloc()` routine will allocate a single huge page for each COP participating in the algorithm. So, it is **VERY** important to match the data buffer size to the huge page size. If this cannot be done, then using `rasclib*_malloc()` may not result in optimal behavior.

Achieving Maximal Bandwidth

For absolute, maximal bandwidth, there should be one memory node with processors for every FPGA.

The more removed from this one-to-one ratio that a topology is, the lower the aggregate bandwidth. The memory nodes need to have processors (not memory only nodes) because each FPGA requires processors for the threads that drive the DMA to run on.

High bandwidth also requires the use of direct I/O (see Chapter 5, “Direct I/O”).

Conclusion

As mentioned above, the size of the application I/O buffers should match the size of the system's huge page size. The memory allocation routines provided by the RASC abstraction layer will only allocate, at most, one huge page per request.

RASC Algorithm FPGA Implementation Guide

This chapter describes how to generate a bitstream and the associated configuration files for the RASC Algorithm FPGA once the Algorithm Block design has been completed. Instructions for installation, setup, and implementation are included along with examples to demonstrate different implementation flows. The implementation flow begins with a completed HDL design and ends with generating the `<alg_name>.bin` file, which is used by the Loader FPGA to configure the Algorithm FPGA. Other implementation flow outputs are the files `user_space.cfg` and `core_services.cfg`, which are used by the software abstraction layer to recognize the algorithm.

This chapter contains the following sections:

- “Implementation Overview” on page 138
- “Installation and Setup” on page 140
- “Adding Extractor Directives to the Source Code” on page 144
- “Implementation with Pre-synthesized Core” on page 148
- “Full-chip Implementation” on page 152
- “Implementation File Descriptions” on page 153

Implementation Overview

This section provides an overview of the RASC Algorithm FPGA implementation process. It covers the following topics:

- “Summary of the Implementation Flow” on page 138
- “Supported Tools and OS Versions” on page 139

Summary of the Implementation Flow

In this document, *implementation flow* refers to the comprehensive run of the extractor, synthesis, and Xilinx ISE tools that turn the Verilog or VHDL source into a binary bitstream and configuration file that can be downloaded into the RASC Algorithm FPGA. There are several stages to this process, but the implementation flow has been encapsulated in a Makefile flow for ease of use.

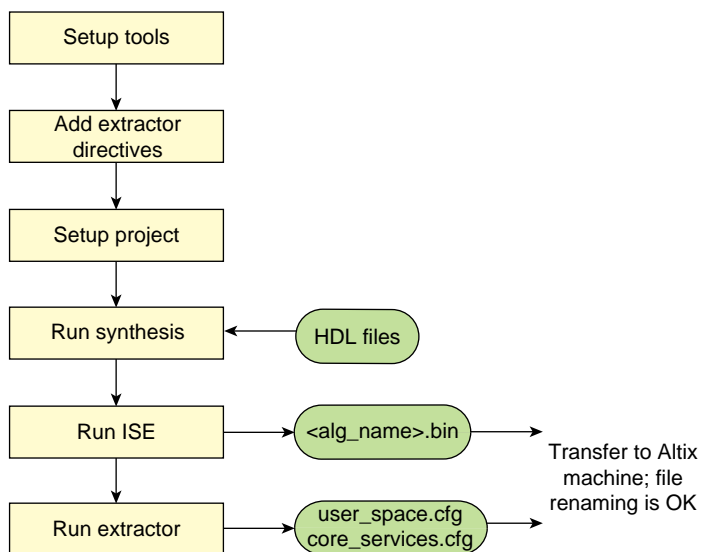


Figure 7-1 RASC FPGA Implementation Flow

There are two supported ways to create a bitstream. The first method, and recommended for most cases, is to implement the Algorithm FPGA using our pre-synthesized Core

Services Block EDIF netlist. This allows faster synthesis time and easy trials of multiple algorithms.

The second method, which allows more optimization between the Core Services and the Algorithm Block, is to synthesize from the top-down, including re-synthesis of the Core Services Block source codes.

Supported Tools and OS Versions

The supported flow has been tested using the tool versions and operating system shown in Table 7-1.

Table 7-1 Supported Implementation Tools

Tool	Version
OS	Red Hat Enterprise Linux 3.0 (or later)
Shell	Bash or tcsh
Python	Python 2.4 or higher
Synthesis	Synplicity Synplify Pro 8.9
ISE	Xilinx ISE 9.2, Service Pack 2

SGI suggests that you consult the reference documentation associated with the supported tools of your choice for the details on tool setup.

- ISE Software Manuals
http://www.xilinx.com/support/sw_manuals/xilinx9/index.htm
- Synplify Pro Documentation
<http://www.synplicity.com/literature/index.html>

Installation and Setup

SGI Altix System Installation

RASC software is distributed either on a CD-ROM or through file transfer protocol (FTP). If an International Standards Organization (ISO) file is received through FTP, put it in a convenient directory in the SGI Altix system. If the directory `/mnt/cdrom` does not exist, create that directory with the following command:

```
mkdir /mnt/cdrom
```

At the directory where the ISO file is saved, execute the following command:

```
mount -o loop iso_file_name /mnt/cdrom
```

The CD-ROM image will be visible in the directory `/mnt/cdrom`.

On the CD-ROM image, there will be one or more files with the extension `.rpm` and a `readme.txt` file. Read the `readme.txt` file for further instructions for installing the software.

PC Installation

This section describes the installation and setup of the RASC Algorithm FPGA Developer's Bundle. Please note that the Xilinx ISE tools and the synthesis tool of choice must be installed on an x86 Linux system and each user must have the associated environment variables defined for their use.

The RASC Algorithm FPGA Developer's Bundle comes as a `tar.gz` file. This `tar.gz` file is available in the Altix system after the `rasc_cs-xxx.ia64.rpm` is installed. It is in the directory `/usr/share/rasc/x86_dev_env`. Run the executable `x86_dev_env_install` and read the license agreement. Agree to the agreement to extract the file `x86_dev_env`. Transfer this file to a directory in the PC. If there is an older version of the bundle software, remove it. Unzip the `tar.gz` file using the following command:

```
tar -zxvf file_name
```

The Developer's Bundle requires setting one additional environment variable, `RASC`, to identify the Bundle installation directory to the supporting Makefiles (see Table 7-2).

Table 7-2 Environment Variables Required for Bundle Environment

Name	Description
RASC	RASC Algorithm FPGA Developer's Bundle Installation Directory
XILINX	Xilinx tools installation directory
SYNPLIFY	Synplify Installation directory
LIB	Required Synplify variable

The directory hierarchy within the RASC Algorithm FPGA Developer's Bundle is shown in the README file at the top level. A description of the top level directories included in the Bundle is shown in Table 7-3.

Table 7-3 Top Level Directory Descriptions

Directory	Description
examples	Contains subdirectories with example algorithm source code and implementation files.
design	Contains Verilog source codes for the top-level FPGA wrapper and the Core Services Block modules.
pd	Additional files needed for physical design / implementation: synthesis and ISE constraint files, the pre-synthesized Core Services Block netlist and the macros needed for re-synthesis of the Core Services Block.
implementations	Contains template files for running the bitstream implementation flow (Makefiles and project files).
dv	Functional simulation components: testbench with SSP stub and QDR2-SRAM models, example diagnostic stimulus code, and scripts for VCS simulator.

When the Developer's Bundle is installed and the `$RASC` variable is set, the user can proceed to experiment with the example implementations.

Implementation Constraint Files

Synthesis Using Synplify Pro

A synthesis constraint file (.sdc) is provided for use with the Synplify Pro tool. It is used to provide additional inputs to the synthesis tool on desired mapping, netlist formats, and post-synthesis timing estimates.

The synthesis constraint file specified in the <alg_name>.prj is added to the generated full_<alg_name>.prj synthesis project file created by Makefile build process. If a user desires to add constraints for portions of the Algorithm Block, this can be done by editing the .sdc file manually or through the associated synthesis GUI.

The synthesis constraint files are located under the install directory at \$RASC/pd/constraints/synthesis.

Synthesis Using XST

A script file (.scr) is provided for use with the XST tool. An XST script is a set of various options used to run synthesis. It is utilized to control various aspects of the synthesis process to meet your design goals or obtain the best implementation of your design.

The synthesis script is used along with the synthesis project file <alg_name>_xst.prj. If a user desires to add constraints for portions of the Algorithm Block, this can be done by creating an .xcf file. To specify the XCF file in the script, use the -uc switch with run command.

The XST script file should be located under the synthesis project directory of each algorithm.

Table 7-4 Synthesis Constraint Files Provided

File	Synthesis Tool	Description
acs_top.sdc	Synplify Pro	Basic synthesis constraint file
<alg_name>_xst.scr	XST	General (global) constraints

ISE (User Constraint File)

A top level user constraint file (ASCII) needs to be used with the ISE tools to guide placement and routing and to achieve timing-driven placement results. The UCF file is generated from separate constraint files (see Table 7-5) located in the directory `$RASC/pd/constraints/ise`.

Table 7-5 Implementation Constraint Files Provided

File	Description
<code>alg.ucf</code>	Algorithm specific UCF constraints, optional
<code>ssp.ucf</code>	Timing constraints and I/O assignments
<code>virtual_gnd.ucf</code>	Virtual ground I/O assignments
<code>prohibit.ucf</code>	Prohibit constraints for specific sites
<code>qdr_bank0.ucf</code> , <code>qdr_bank1.ucf</code> , ... , <code>qdr_bank4.ucf</code> , and <code>qdr_misc.ucf</code>	Location and specific constraints for QDR banks

The constraint file should always be included in the design, and the provided constraints should not be modified. However, a user can add to the constraint file to put in constraints related to the Algorithm Block. See the last line of the `Makefile.local`. Uncomment the last line, as follows:

```
# Algorithm dependent UCF constraints file
#ALG_UCF=alg.ucf
```

The constraint file provided in the Bundle provides the following types of constraints:

- Period specifications for all internal clock signals, both for the Core Services Block and the Algorithm Block
- I/O Pad placement constraints and I/O standard settings
- Setup and hold constraints for input pads
- DCM attributes settings

- Cross clock domain path constraints
- BUFG and DCM placement constraints

All constraints are written in the Xilinx-specified User Constraint File format.

An example of a constraint file is, as follows:

```
--# Set the current user constraint file for ISE
UCF_FILE_LIST=./alg_cmu_dft_n16k_w32_local.ucf
${UCF_DIR}/virtual_gnd.ucf ${UCF_DIR}/ssp.ucf ${UCF_DIR}/prohibit.ucf
${UCF_DIR}/qdr_bank0.ucf ${UCF_DIR}/qdr_bank1.ucf
${UCF_DIR}/qdr_bank2.ucf ${UCF_DIR}/qdr_bank3.ucf
${UCF_DIR}/qdr_bank4.ucf ${UCF_DIR}/qdr_misc.ucf
```

The `./alg_cmu_dft_n16k_w32_local.ucf` file contains these statements that are currently commented-out:

```
#INST
"u_user_space_wrapper/u_alg_block_top/u_calc_control/u_dft/*/FIFOinst*/fifo*"
LOC=RAMB16_X2Y0:RAMB16_X6Y47;

# INST
"u_user_space_wrapper/u_alg_block_top/u_calc_control/u_dft/instList699/instList2221/instrc222
6/instList2235/instParallelWithControl2243/seu2313/FIFOinst2318/fifo*"
LOC=RAMB16_X3Y0:RAMB16_X6Y47;

# INST
"u_user_space_wrapper/u_alg_block_top/u_calc_control/u_dft/instList699/instrc2222/instList234
6/instParallelWithControl2347/seu2362/FIFOinst2367/fifo*"
LOC=RAMB16_X3Y0:RAMB16_X6Y47;
```

Adding Extractor Directives to the Source Code

Symbol table information is required for GDB to correctly display internal FPGA values and for the RASC Abstraction Layer (RASCAL) to communicate parameters and data that should be written to and read from the SRAM banks and internal registers. A python script called *extractor* parses all the Verilog, VHDL, and header files in an arbitrary algorithm directory that is passed as an argument. The user specifies the top level algorithm directory by setting the `ALG_DIR` variable in the `Makefile.local` file, to be discussed later. When the extractor script is called, it generates two configuration files that communicate the necessary information for the RASCAL and GNU source-level debugger (GDB) to communicate with the hardware correctly.

The extractor script parses comment lines that contain special extractor directives from within source and header files in the algorithm directory. There is a template in the `examples` directory, in addition to full examples for each of the example algorithms. The comment fields can be located in one or more files below the directory specified in the second argument to the extractor call.

The command to run the extractor has been automated in the Makefile flow discussed in detail later. The user does not need to run extractor at the command line, but can issue the `make extractor` command after the Makefiles have been set up.

Inserting Extractor Comments

The comment fields that need to be inserted within the algorithm source or header code are: *core services version*, *algorithm version*, *sram* (denoting where data will be read from and written to on the sram interface), *register in* (for parameters set through an application writer's code), and *register out* (for a section of code that needs to be mapped to a debug register).

The comments can be located anywhere within a Verilog, VHDL or header file. They can have Verilog or VHDL syntax ("`//`" or "`--`" respectively). Extractor comments begin with the tag "`//extractor`" in the source file. The general format is shown below.

```
//extractor <type>:[<value/specifier1> ... <value/specifierN>] (Verilog)
```

```
-- extractor <type>:[<value/specifier1> ... <value/specifierN>] (VHDL)
```

For details on each extractor comment, see Table 7-6.

Example of Comments in a Verilog, VHDL, or Header File

Table 7-6 Extractor Comment Fields

Type	Description	Value	Example
CS	Core Services Version: Tag detailing which iteration of core services was used in this bitstream.	<version>.<revision>	//extractor CS:0.7
VERSION	Algorithm Version: Tag for the type and version number of the algorithm in this bitstream. This is user defined but it should match the first register in the debug space (see Chapter 3, “RASC Algorithm FPGA Hardware Design Guide” for more details.)	<version>.<revision>	//extractor VERSION:0.12
SRAM	Array: Tag for the SRAM inputs and outputs to the algorithms.	<array name> <number of elements in array> <bit-width of each element> <SRAM bank for array> <starting byte offset> <array direction: in out inout> <data type for array elements: u=unsigned> <array use attribute: stream fixed>	//extractor SRAM:a_in 512 64 sram[0] 0x0000 in u stream
REG_IN	register in: Tag for the registers that are PIO written as specified by the application and where the value will be mapped in the address space.	<register name> <bit width of register> <data type for element> <alg_def_reg mapping> <bit range mapping of alg_def_reg>	//extractor REG_IN:match 32 u alg_def_reg[0][31:0]

Table 7-6 Extractor Comment Fields (continued)

Type	Description	Value	Example
REG_OUT	register out: Tag for the registers that are PIO written or read by the application or the debugger and where the value will be mapped in the address space.	<register name> <bit width of register> <data type for element> <debug port mapping> <bit range mapping of debug_port>	//extractor REG_OUT: a 16 u debug_port[1] debug_port[1] [15:0]
STREAM_IN	stream in tag	<stream_name> <stream number> <stride>	// extractor stream:input_stream in 8
STREAM_OUT	stream out tag	<stream_name> <stream number> <stride>	// extractor stream:output_stream out 8

A sample of the comments that might be included within a source file is:

```
// extractor VERSION: 6.3
// extractor CS: 1.0
// extractor SRAM:a_in 2048 64 sram[0] 0x0000 in u stream
// extractor SRAM:b_in 2048 64 sram[0] 0x4000 in u stream
// extractor SRAM:c_in 2048 64 sram[0] 0x8000 in u stream
// extractor SRAM:d_out 2048 64 sram[1] 0x0000 out u stream
// extractor REG_OUT:alg_id 32 u debug_port[0][63:32]
// extractor REG_OUT:alg_rev 32 u debug_port[0][31:0]
// extractor REG_OUT:rd_addr 64 u debug_port[1]
// extractor REG_OUT:rd_data_sram0_lo 64 u debug_port[2]
// extractor REG_OUT:rd_data_sram0_hi 64 u debug_port[3]
// extractor REG_OUT:wr_addr 64 u debug_port[4]
// extractor REG_OUT:wr_data_sram1_lo 64 u debug_port[5]
// extractor REG_OUT:wr_data_sram1_hi 64 u debug_port[6]
// extractor REG_OUT:cntl_sigs 64 u debug_port[7]
```

Implementation with Pre-synthesized Core

The first recommended implementation flow uses the pre-synthesized version of the Core Services Block. The bulk of the Core Services Block (excluding clocking, resets, and I/O) has been pre-synthesized and has soft-fixed source code (we do not support customer modifications). The implementation time could be significantly reduced by using the pre-synthesized core, but results vary depending on the loading of signals between the Algorithm Block and the Core Services Block. We suggest that you first try this method, and if timing is not met due to loading issues on the Algorithm Block / Core Services interface, re-synthesize the full design from scratch (discussed later).

Each algorithm project will have its own associated implementation project directory, which can be located anywhere in the directory space of the user. The user customizes template Makefiles and synthesis project files to run the desired bitstream generation flow. Each project directory must contain the following three files, the first two provided in the `$RASC/implementations/templates` directory:

- **Makefile** (linked or copied from the templates directory).

This Makefile contains general definitions and targets necessary to run the synthesis and ISE tools. It includes the `Makefile.local` file. This file is meant to be linked into each new algorithm project directory. Customizations should be added to `Makefile.local`; it is recommended to keep this file common between all projects.

- **Makefile.local**

This included Makefile allows custom definitions and targets for each algorithm project. Here you can define the synthesis project name, specify different constraint files, and override tool options.

- **Synthesis Project File**

– **Synplify Pro** `<alg_name>.prj`.

This file defines the algorithm source files and synthesis tool options. *This is where you need to add / modify the paths to the algorithm code.* The Synplify Pro project file is a TCL script, which can also be opened and run by the GUI version of the synthesis tools. Read the synthesis tool documentation for more information on the TCL commands available. You will notice that the template synthesis project file does not contain any source files for the Core Services Block. All Core Services files are stored in another file in the templates directory, `acs.files`. The Makefile flow combines `acs.files` with the synthesis project file of the user to create a complete synthesis project (renamed as

full_<alg_name>.prj). If you are creating a synthesis project from scratch or not using the Makefile flow, you will need to copy acs.files into your synthesis project.

- XST <alg_name>_xst.prj.

This file only defines the algorithm source files. *This is where you need to add / modify the pointers to the algorithm code.* The synthesis options are included in a separate script file <alg_name>_xst.scr which can be opened and run by the synthesis tools.

Examples of each of these required files are located in the templates directory, as well as the included example algorithm implementation project directories:

```
$RASC/implementations/<alg_name>
```

Makefile.local Customizations

Any Makefile variables you want to redefine can be changed using the Makefile.local file. This allows each project you create to use the same base Makefile, but have project-specific customizations in a separate file. The key Makefile variables that are commonly redefined for each project are listed here. Default settings are provided with the Makefile.local version in the templates directory.

Synthesis Project Customization

Table 7-7 shows common Makefile.local variable settings.

Table 7-7 Common Makefile.local Variable Settings

Variable	Example	Definition
ALG_DIR	\$RASC/examples /alg_simple	Pointer to top level algorithm directory. This is a full path, not a relative path.
SYNTHESIS_PROJ	<any name>, e.g., alg_simple_hc	Name of the synthesis project; the base filename to Synplify's *.prj file.

Table 7-7 Common Makefile.local Variable Settings (**continued**)

Variable	Example	Definition
SYNTHESIS_DIR	<i><any relative path directory></i> , e.g. rev_1	Output directory for synthesis and implementation flow, relative to project's implementation directory.
DEFAULT_SYNTHESIS_TOOL	synplify_pro, ise_xst	Specify synthesis tool to use for default make flow
SYN_RESULT_EXT	edf, ngc	Specify the synthesis tool's output file extension
LOCAL_SOURCE_DIR	<i><any full path list of directories></i>	List of directories of additional edf/sources related to the algorithm that need to be read by ngdbuild (default: none).
MAP_OPTIONS, PAR_OPTIONS	Any ISE allowed options	Set ISE tool options for a specific run.
UCF_DIR, CURR_UCF_FILE, etc.		Additional variables to override the Makefile settings

Synplify Pro

The Synplify Pro tool uses a *<alg_name>.prj* file to specify the HDL files for synthesis and other compiler, mapping, and output options. When creating a new project for a new algorithm, the algorithm related files need to be added to the project file so that the synthesis tool knows which algorithm to compile. The user should specify the location of the files *alg_block_top.v* and *alg.h*, along with any submodules. For example, if you are compiling an algorithm whose two source files *alg_block_top.v* and *sort_byte.v* are located under *\$RASC/examples/alg_simple_v/*, you would need to add the following lines to the project file:

```
add_file -verilog "$RASC/examples/alg_simple_v/alg_block_top.v"
add_file -verilog "$RASC/examples/alg_simple_v/sort_byte.v"
```

Please note that for some projects it is useful to read the EDIF netlist of Core Services Block during the synthesis. To do that add the following line to the project file:

```
add_file -edif "$RASC/pd/acs_core/acs_core.edf"
```

XST

The XST tool uses a `<alg_name>_xst.prj` file to specify the HDL files for synthesis and creates Xilinx specific netlist files called `ngc` files. All other compiler (for example, verilog include directories), synthesis, and output options are included in a separate script file `<alg_name>_xst.scr`. When creating a new project for a new algorithm, the algorithm related files need to be added to the project file so that the synthesis tool knows which algorithm to compile. The user should specify the location of the files `alg_block_top.v` and `alg.h`, along with any submodules. For example, if you are compiling an algorithm whose source file `alg_block_top.v` is located under `$/RASC/examples/alg_simple_v/`, you would need to add its full or relative path to the project file:

```
verilog work "../../../../example/alg_simple_v/alg_block_top.v"
```

Makefile Targets

After the `Makefile`, `Makefile.local`, and synthesis project file are set up and customized, the next step is to run the `make` command to generate the bitstream and configuration files for download to the RASC hardware.

FPGA implementation comprises many sub-steps such as synthesis, `ngdbuild`, `map`, `place-and-route`, `bitgen`, and timing analysis. The user can choose to run the entire flow or individual subsections based on the `Makefile` target. If no target is specified, the full implementation flow is executed:

```
% cd $RASC/implementations/alg_simple_v
% make all
```

Other available targets are described in Table 7-8. The default command is `make all`.

Table 7-8 Makefile Targets

Command Line	Description
<code>make all</code>	Runs extractor, synthesis and ISE tools, results in the <code><alg_name>.bin</code> and related <code>.cfg</code> files, ready to download to FPGA. This is the default make target.
<code>make extractor</code>	Run the Extractor script to generate the configuration files based on the source code comments,

Table 7-8 Makefile Targets (**continued**)

Command Line	Description
<code>make synplify_pro</code>	Run Synplify Pro tool alone. Results in EDF file.
<code>make amplify</code>	Run Amplify tool alone. Results in EDF file.
<code>make ise_xst</code>	Run XST tool alone. Results in NGC file.
<code>make oneloop</code>	Run ISE tools, does not rerun synthesis. Performs 1 PAR iteration. Results in NCD output
<code>make ngdbuild</code>	Run NGDBUILD alone.
<code>make map</code>	Run MAP alone.
<code>make par</code>	Run PAR alone.
<code>make trce</code>	Run TRACE (static timing analysis) alone.
<code>make bitgen</code>	Run BITGEN alone.
<code>make clean</code>	Remove all intermediate and output files.

Full-chip Implementation

A full-chip implementation refers to synthesis of the entire FPGA design, including a re-synthesis of the Core Services Block. The implementation steps for this flow are similar to those for the flow using the pre-synthesized Core Services Block netlist, with differences specified in the `Makefile.local` and the synthesis project file.

Makefile.local Customizations

When synthesizing the full FPGA, including a re-synthesis of the Core Services Block, we only support the use of Synplify Pro. The Core Services Block design has been optimized for use with these tools and results from using another synthesis tool will not likely achieve the same quality-of-results and timing closure. Therefore, the `Makefile.local` for the full chip project must use the following Makefile variable settings (see Table 7-9).

Table 7-9 Required Full Chip Makefile Variable Settings

Variable	Possibilities
DEFAULT_SYNTHESIS_TOOL	synplify_pro
SYN_RESULT_EXT	edf

Synthesis Project Customization

A synthesis option must be set to allow the Core Services Block to be re-synthesized. This option is in the form of a Verilog compiler directive, which can be added to the synthesis project file with the following TCL command:

```
set_option -hdl_define -set RESYNTHESIZE_CS=1
```

This option is commented in the <alg_name>.prj and must be uncommented by the user desiring to re-synthesize the Core Services Block. The rest of the synthesis project should be the same as described earlier.

Implementation File Descriptions

Table 7-10 includes the commonly referenced files. The synthesis and ISE tools use a number of intermediate files. To find out more information about files not discussed here, please consult the on-line reference guides mentioned above.

Table 7-10 Commonly Referenced Files

Name	Type	Produced By	Description
.bin	Binary	Bitgen	Contains only configuration data for the FPGA. The .bin has no header like the .bit file. Meant for input to non-Xilinx tools such as the FPGA Loader Device Driver.
.bit	Binary	Bitgen	Contains proprietary header information as well as configuration data. Meant for input to other Xilinx tools such as PROMGen and Impact.

Table 7-10 Commonly Referenced Files **(continued)**

Name	Type	Produced By	Description
.cfg	ASCII	Extractor	Configuration file used by the Abstraction Layer software to recognize features of the Algorithm and Core Services.
.edf	ASCII	Synplify Pro Synthesis Tool	EDIF Netlist.
.ngc	ASCII	XST Synthesis Tool	NGC Netlist.
.ncd	Data	Map, PAR	Flat physical design database. After PAR, this is the database file with the full placed-and-routed design.
.ncf	ASCII	Synplify Pro Synthesis Tool	Constraint file produced by the Synplify Pro synthesis tool to be associated with the .edf file of the same name.
.prj	ASCII	User or Synthesis Tool	Synthesis Project definition file.
.sdc	ASCII	User (text editor) / Synthesis tools GUI (Synplify Pro or Amplify)	Provides directives and attributes for the synthesis tool to use while it is compiling the design
.twr	ASCII	Trace	Timing report file. Provides details on failing timing paths.
.twx	XML	Trace	Timing report file, with the same information as TWR but can be opened by the Timing Analyzer tool.
.ucf	ASCII	User (text editor) / ISE Constraints Editor or PACE	User-specified logic constraints file. Constrains pinout, placement and timing constraints.

SGI RASC Algorithm Configuration Tool

The SGI Core Services, described in Chapter 3, “RASC Algorithm FPGA Hardware Design Guide”, are designed to interact with the algorithm that you are creating. Certain configuration files are necessary to ensure that the Core Services and the algorithm are both aware of the resources in use. Prior to the RASC 2.20 release, these configuration files needed to be hand-edited based on templates provided with the RASC software source code.

The SGI RASC Algorithm Configuration Wizard enables you to select the resources to be used in a given algorithm through a GUI. Core Services features, such as, streaming direct memory access (DMA), memory resources, algorithm defined registers (ADRs), and debug registers can be configured.

The tool then creates the customized files required to integrate the algorithm design with Core Services and the implementation scripts used to generate a bitstream from the source code.

This chapter covers the following topics:

- “Running the RASC Algorithm Configuration Tool” on page 156
- “Algorithm Name” on page 156
- “Algorithm Clock Rate” on page 157
- “Streaming DMA Configuration” on page 158
- “Memory Configuration” on page 159
- “Memory Port Specification” on page 160
- “Algorithm Defined Registers” on page 161
- “Algorithm Defined Register Configuration” on page 161
- “Debug Registers” on page 162
- “Synthesis Tool Specification” on page 163
- “Supplemental Algorithm Clock” on page 164

- “Output File Locations” on page 164

Running the RASC Algorithm Configuration Tool

The SGI RASC Algorithm Configuration Wizard is a Tool Command Language (TCL) script located in the `$RASC/design/tools` directory. You must run it in the same filesystem as the source code. To run the script, perform the following:

```
.* ./alg_cfg_tool.tcl
```

The welcome screen appears, as shown in Figure 8-1.



Figure 8-1 SGI RASC Algorithm Configuration Tool Welcome Screen

Every screen in the Wizard allows you to cancel your configuration and exit the tool without updating the output design and implementation files. The following sections in this chapter will walk you through using the Wizard. Click **Start** to begin the configuration steps.

Algorithm Name

The algorithm name is used to create the directories for the design files and the implementation files. It is also used as a way of attaching constraint files to the design during the implementation steps of the bitstream. Enter the name of your algorithm, as shown in Figure 8-2.



Figure 8-2 Algorithm Name Configuration

Since the algorithm name is used to create directories in which to place the output files, it is highly recommended that you use an algorithm name that conforms to Unix/Linux directory names.

Algorithm Clock Rate

The SGI Core Services provide an algorithm clock that enables communication between the Core Services and the algorithm interfaces. There are four acceptable frequencies for this clock. For more information on the algorithm clock rate, see “FPGA Clock Domains” on page 57.

You can select one of the available frequencies, as shown in Figure 8-3. A supplemental clock is available at a later configuration stage.




Figure 8-3 Algorithm Clock Rate Configuration

Streaming DMA Configuration

SGI Core Services provide four input and four output streaming DMA engines that allow data to transfer directly into the FPGA. These are individually selectable and you may select them in any combination using the interface, as shown in Figure 8-4.

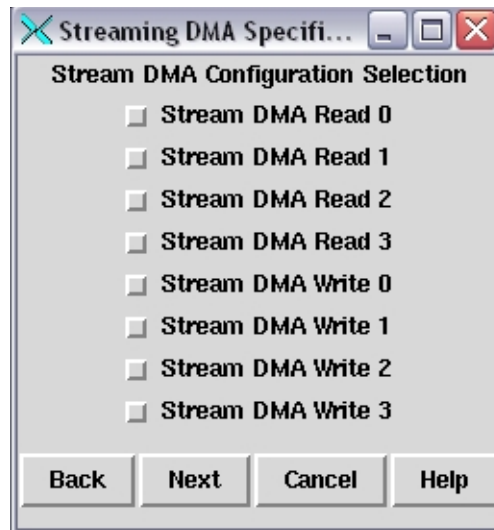


Figure 8-4 Streaming DMA Configuration

Memory Configuration

There are three acceptable and mutually exclusive memory configurations allowed in RASC Core Services. The first is the same as the configuration found in the prior RASC 2.1 release. This configuration pairs up two 64-bit memory ports to be used as a 128-bit memory. There is a fifth 64-bit memory port available in this configuration.

A new configuration is available in the RASC 2.20 release (and later) to enable five 64-bit memory ports. All five of these ports are independent and can be configured individually.

You may also decide that you do not need memory ports and all data will flow in and out through the streaming DMA engines. This enables an area optimization by removing the memory interfaces from the integrated design.

Select the appropriate memory configuration, as shown in Figure 8-5.




Figure 8-5 Memory Configuration Selection

Memory Port Specification

Once your memory configuration has been selected you can make further optimizations and usage directions, as shown in Figure 8-6. This screen has four columns. The first, leftmost, column has checkboxes that, when checked, enable the use of the port. If the checkbox is clear, the port will be configured out of the design, and the area will be optimized.

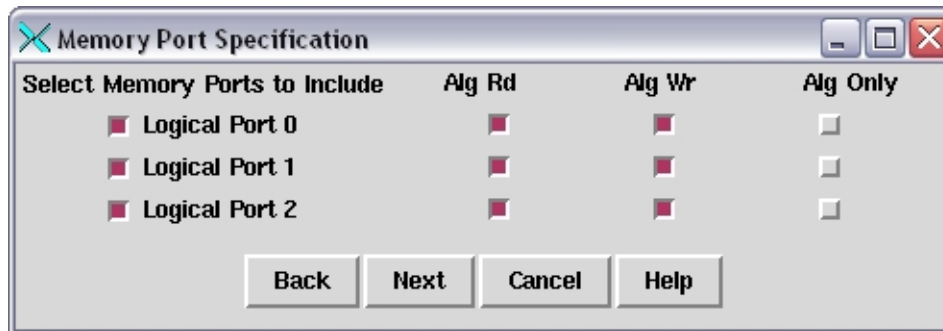


Figure 8-6 Memory Port Specification

The second and third columns define the ports of the memory that the algorithm intends to make use of. The algorithm can read and/or write each of the memory ports based on the checkboxes in these columns.

The rightmost column enables further area optimization. If the algorithm intends to make use of the memory as a scratchpad or temporary storage, and the algorithm intends to take responsibility to move the data, you can check the box in the **Alg Only** column and thus eliminate the core services DMA connections and arbitration to that memory port.

Algorithm Defined Registers

The SGI Core Services provides up to 64 registers to handle the communication between the host program and the bitstream. The number of registers is selected using the slider, as shown in Figure 8-7.

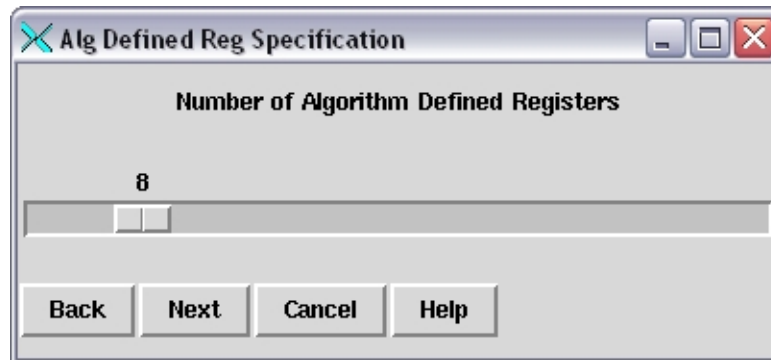


Figure 8-7 Number of Algorithm Defined Register Selection

Algorithm Defined Register Configuration

The algorithm defined registers can be further specified for usage, as shown in Figure 8-8.

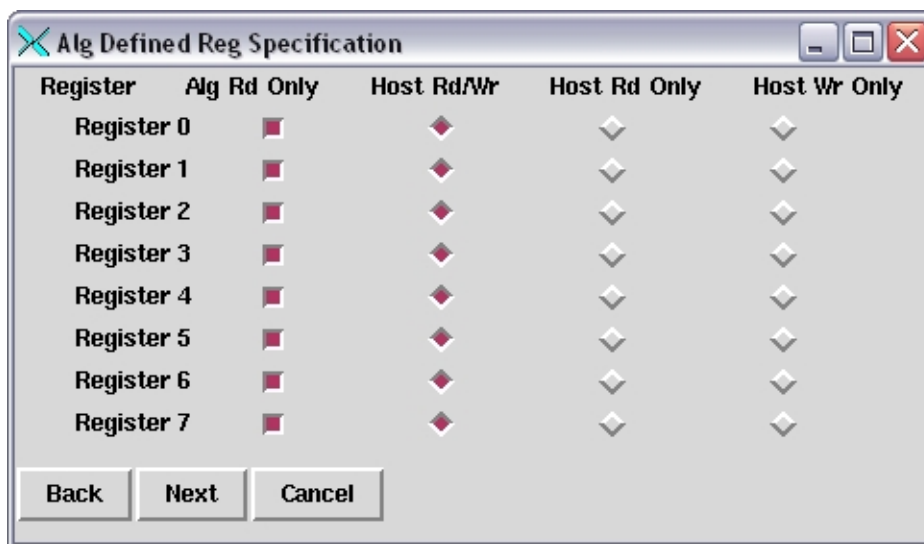


Figure 8-8 Algorithm Defined Register Specification

Debug Registers

Core Services provides up to sixty-four (64) 64-bit debug registers. These registers are up to the algorithm author to specify, and they can be read through GNU Debugger (GDB). For more information about the debug registers and the GDB, see “Using the GNU Project Debugger (GDB)” on page 178. These registers are also specified using a slider, as shown in Figure 8-9.




Figure 8-9 Number of Debug Registers Selection

Synthesis Tool Specification

The SGI generated implementation scripts support two synthesis tools: Synplify Pro and Xilinx XST. Since Xilinx XST is part of the Xilinx ISE Foundation, it is selected as default, as shown in Figure 8-10. Xilinx and SGI recommend Synplify Pro for the highest quality of results, and if it is available, checking this box will ensure that the proper Synplify Pro scripts will be generated.

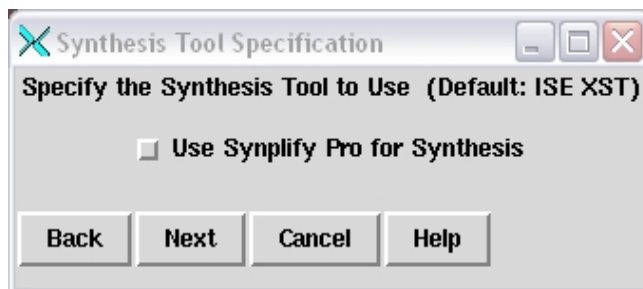


Figure 8-10 Synthesis Tool Specification

Supplemental Algorithm Clock

A supplemental algorithm clock is provided by the Core Services. This clock is configured according to the *Xilinx Virtex 4 User's Guide*. It involves the specification of a numerator and a divisor for a scaling factor, as shown in Figure 8-11.

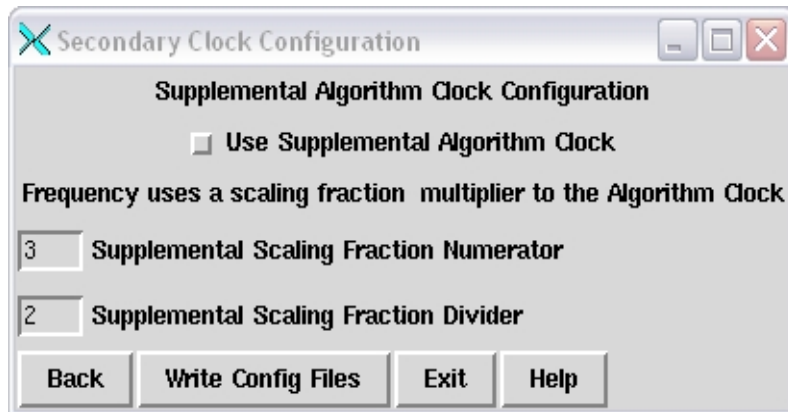


Figure 8-11 Supplemental Algorithm Clock Configuration

The checkbox at the top of the window enables the use of the supplemental clock. The two lower boxes enable the specification of the numerator and divisor.

This is the last step of the configuration. Clicking **Write Config Files** will generate the proper scripts.

Output File Locations

The output files are split into two categories. The first is the design category. These are Verilog files that are placed in the design tree. They are located at:

```
$RASC/design/alg_core/{ALG_NAME}
```

The second category is implementation scripts. These scripts are used to feed the synthesis and implementation tools, and they can be found at:

```
$RASC/implementation/{ALG_NAME}
```

If the algorithm name specified already exists, there may already be files in these locations that can be overwritten (see Figure 8-12).

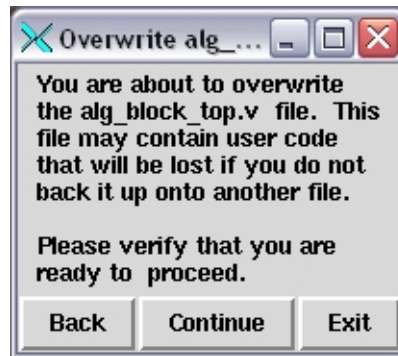


Figure 8-12 Overwrite of `alg_block_top.v` Verification

Note: The `alg_block_top.v` file may contain the user algorithm code that will be overwritten during this reconfiguration stage. It is important to save any of the logic in this file elsewhere and replace it in the newly generated `alg_block_top.v` file. Save the original file to `alg_block_top.v.backup`.

If the files are writable, clicking **Continue** will generate the proper files and exit the Wizard.

If this is a reconfiguration, and the files are not writable, you will see the write-protect file notification, as shown in Figure 8-13.

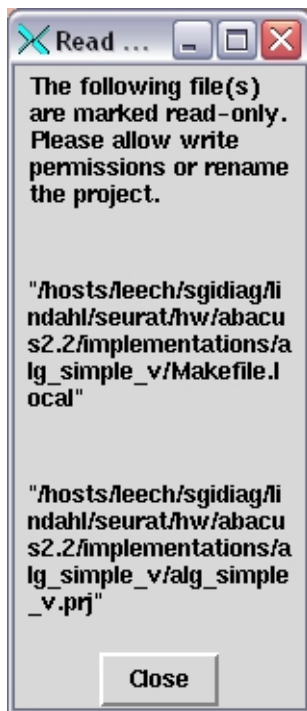


Figure 8-13 Write-Protect File Notification

This screen displays the filenames for files that exist and are write-protected. A change in the permissions of these files, or a change in the algorithm name is required before the files can be written. The screen can be closed, and you are sent back to the final stage of configuration. If you have changed permissions, you may select **Write Config Files** again, and the files will be generated and the program will exit. If you wish to change the algorithm name instead, you may use the **Back** buttons to return to the algorithm name configuration, or you may exit the Wizard and start over.

Running and Debugging Your Application

This chapter describes how to run and debug your application and covers the following topics:

- “Loading the Bitstream” on page 167
- “RASC Device Manager” on page 168
- “Using the GNU Project Debugger (GDB)” on page 178

Loading the Bitstream

The Device Manager (`devmgr`) maintains a bitstream registry of algorithm bitstreams that can be loaded and executed using the RASC abstraction layer. The `devmgr` user command is used to add, delete, and query algorithms in the bitstream registry. An algorithm is identified by its user supplied name and consists of a set of files, as follows:

- Algorithm bitstream
- Configuration file for the algorithm bitstream
- Core services configuration file

All of these files must be accessible by the `devmgr` user command when the algorithm is added to the bitstream registry.

The `devmgr` user command executes on the target Altix system, but the algorithm bitstream and configuration files are built on a PC using one of several FPGA development tools.

Note: For detailed information using the `devmgr` user command, see “RASC Device Manager” on page 168.

To make the algorithm files available to the `devmgr` command, perform the following:

1. Prepare to FTP three files from the PC to the Altix system.
 - The files can reside anywhere on the Altix system that can be accessed by the `devmgr` command. A convenient location is, as follows:
`/usr/share/rasc/bitstream/your_algorithm_directory`
 - Create the bitstream and optionally a user specific subdirectory in `/usr/src/RASC`, if they do not already exist on the Altix system.
2. Execute the FTP command from the PC to transfer the algorithm bitstream, its configuration file, and the core services configuration file to the user generated subdirectory.
3. Login to the Altix system and execute the `devmgr` user command to add the algorithm to the bitstream registry giving the algorithm a user supplied name. All references to the bitstream in the bitstream registry is made using this name. Note that the bitstream file suffix must be `*.bin` and the configuration files suffix must be `*.cfg`.

For example, assume that an algorithm bitstream and configuration files have been FTP'd to the Altix system in subdirectory `dog`. A list of the files in `dog` could look like the following:

```
/usr/share/rasc/bitstream/dog/spot.bin
/usr/share/rasc/bitstream/dog/spot.cfg
/usr/share/rasc/bitstream/dog/core_services.cfg
```

4. This `devmgr` command could be executed on the Altix system to add the algorithm to the bitstream registry with the name `bark`, as follows:

```
devmgr -a -n bark -b /usr/share/rasc/bitstream/dog/spot.bin
```

Notice that the full path name of the bitstream must be specified. If the algorithm bitstream configuration file is not called `spot.cfg`, the `devmgr -c` option must also specify the full path name of the algorithm configuration file. Similarly, if the core services file is not `core_services.cfg`, the `-s` option must specify its full path name.

RASC Device Manager

This section provides detailed information on using the Device Manager (`devmgr`) software and covers the following topics:

- “RASC Device Manager Overview” on page 169

- “RASC Device Manager Structure” on page 170
- “Using the Device Manager Command (`devmgr`)” on page 170
- “Device Manager Server Command” on page 176
- “Using the Device Manager Server (`devmgr_server`) Command” on page 177
- “Device Manager Logging Facility” on page 177

RASC Device Manager Overview

The RASC Device Manager (`devmgr`) maintains a bitstream registry and a FPGA hardware inventory. You can add, delete, and list bitstreams in the bitstream registry using the `devmgr` command and can allocate available FPGAs by making requests using the RASC abstraction layer software. You can also use the `devmgr` command to list the FPGAs in the inventory.

You can manually load an algorithm into an FPGA on demand using the `devmgr user` command `-l` option.

You can use the `devmgr -r` option to direct the Device Manager to always reload an algorithm into an FPGA, even when that FPGA is already loaded with the current version of the requested algorithm.

When an FPGA is allocated to a user, the requested bitstream is loaded into those FPGAs. Allocation requests must be made using abstraction software (see Chapter 4, “RASC Abstraction Layer” for specific API calls) and are passed to the device manager.

Each entry in the bitstream registry is identified by the name, which is limited to 63 characters, that is assigned by the user when the bitstream is added to the bitstream registry. The entry also includes copies of the bitstream binary file, the bitstream configuration file, and the cores services configuration file. The location of these files on the local machine is also specified by the user when the entry is added.

You can use the `devmgr -m unavail` option to mark an FPGA device node as unavailable for application allocation. You can use the `devmgr -m avail` option to mark an FPGA device node as available for application allocation.

The `devmgr -v` option shows the `devmgr` command version, and build date and time. This information can be helpful when diagnosing Device Manager software problems.

RASC Device Manager Structure

The Device Manager is structured as client and server components that communicate over a TCP/IP connection. Both the `devmgr` user command and the services requested by the abstraction layer interface with the client side of the device manager. The bitstream registry management and FPGA allocation and loader services occur on the server side. Before a bitstream can be added to the bitstream registry, all of the bitstream files must exist on the machine where the server is running.

Using the Device Manager Command (`devmgr`)

The `devmgr` user command requires one action argument that may require additional arguments, and accepts an optional client debugging argument. The command actions are, as follows:

- “Add a Bitstream To the Bitstream registry” on page 171
- “Delete a Bitstream From the Bitstream registry” on page 171
- “List the Contents of a Bitstream registry” on page 171
- “Update an Algorithm in the Bitstream registry” on page 172
- “List the FPGAs in the Inventory” on page 172
- “Mark an FPGA as Available or Unavailable” on page 173
- “Turn Debugging On or Off” on page 174
- “Device Manager Load FPGA Command” on page 174
- “Device Manager Reload FPGA Command” on page 175
- “Device Manager Version Information” on page 176

Guidelines for using the `devmgr` command are, as follows:

- The add and delete actions require a bitstream registry entry name; the add action also requires the binary bitstream file name, which must have the “.bin” suffix.
- If the bitstream configuration file name is not specified, it defaults to the bitstream file basename with a “.cfg” suffix in the same directory as the binary bitstream file.

- If the cores services file name is not specified, it defaults to "core_services.cfg" in the same directory as the binary bitstream file.

Add a Bitstream To the Bitstream registry

The devmgr command syntax to add a bitstream to the bitstream registry is, as follows:

```
devmgr -a -n algorithm name -b bitstream file name [-c bitstream config file name]
[-s core services config file name] [-y on | off]
```

The devmgr command accepts the following options:

-a	Adds an entry to the bitstream registry.
-n <i>algorithm name</i>	Names an entry in the bitstream registry.
-b <i>bitstream file name</i>	Indicates the location of the bitstream file.
-c <i>bitstream config file name</i>	Indicates the location of the configuration file.
-s <i>core services config file name</i>	Indicates the location of the core services config file.
-y on off	Turns client debugging on or off.

Delete a Bitstream From the Bitstream registry

The devmgr command syntax to delete a bitstream from the bitstream registry is, as follows:

```
devmgr -d -n algorithm name [-y on | off]
```

The devmgr command accepts the following options:

-d	Deletes an entry from the bitstream registry.
-n <i>algorithm name</i>	Identifies the name of an entry in the bitstream registry.
-y on off	Turns client debugging on or off.

List the Contents of a Bitstream registry

The devmgr command syntax to list a specific entry or all entries in the bitstream registry is, as follows:

```
devmgr -q [-n algorithm name][-y on | off]
```

The devmgr command accepts the following options:

<code>-q</code>	Lists all entries in the bitstream registry.
<code>-n <i>algorithm name</i></code>	List a specific algorithm in the bitstream registry.
<code>-y on off</code>	Turns client debugging on or off.

Update an Algorithm in the Bitstream registry

You can use the `-u` option to update an existing algorithm in the Device Manager bitstream registry with a new version of the bitstream file, bitstream configuration file, and core services configuration file. All three files are replaced in the bitstream registry with the new files specified with the command. The `-u` option is logically equivalent to a `-d` option followed by the `-a` option, where both options specify the same algorithm name.

After a successful update of an algorithm in the bitstream registry, any FPGAs currently loaded with the same algorithm name will be reloaded with the new version of the algorithm the next time that FPGA is allocated.

The `devmgr` command syntax to update an algorithm in the bitstream registry is, as follows:

```
devmgr -u -n algorithm name -b bitstream file name [-c bitstream config file name]  
[-s core services config file name] [-y on | off]
```

The `devmgr` command accepts the following options:

<code>-u</code>	Updates the entry in the bitstream registry.
<code>-n <i>algorithm name</i></code>	Names an entry in the bitstream registry.
<code>-b <i>bitstream file name</i></code>	Indicates the location of the bitstream file.
<code>-c <i>bitstream config file name</i></code>	Indicates the location of the configuration file.
<code>-s <i>core services config file name</i></code>	Indicates the location of the core services config file.
<code>-y on off</code>	Turns client debugging on or off.

List the FPGAs in the Inventory

You can use the `devmgr` command to list the FPGAs in the inventory. The list indicates whether the FPGA is available or in use, the system node where the FPGA resides (its device node ID), part number, and manufacture (mfg) number. FPGAs that are in use also

list the Linux process ID (*pid*) and Pthread thread ID (*thread id*) of the current FPGA user, and the name of the algorithm currently loaded in to the FPGA.

The `devmgr` command syntax to list a specific entry or all entries in the FPGA inventory is, as follows:

```
devmgr -i [-N odd integer][-y on | off]
```

The `devmgr` command accepts the following options:

<code>-i</code>	Lists all entries in the inventory.
<code>-N <i>odd integer</i></code>	List the specific FPGA device node ID.
<code>-y on off</code>	Turns client debugging on or off.

Mark an FPGA as Available or Unavailable

You can mark an FPGA device node as unavailable for use by an application with the `devmgr -m unavail` option. If the FPGA is currently in use, the mark action will remain pending until the FPGA is freed; then the FPGA is marked as unavailable.

The `devmgr -m avail` option marks an FPGA device node as available for allocation by an application.

You must specify the target FPGA device node ID using the `-N` option whenever you use the `devmgr -m` option.

You can use the `devmgr -i` option to display the availability of an FPGA device (see “List the FPGAs in the Inventory” on page 172).

The `devmgr` command syntax to to mark an FPGA device as available or unavailable is, as follows:

```
devmgr -m avail | unavail -N odd integer [-y on | off]
```

The `devmgr` command accepts the following options:

<code>-m avail unavail</code>	Marks an FPGA device node as available or unavailable for use.
<code>-N <i>odd integer</i></code>	Specifies the target FPGA device node ID.
<code>-y on off</code>	Turns client debugging on or off.

Turn Debugging On or Off

The `devmgr` command syntax to turn server or client debugging on or off is, as follows:

```
devmgr -x on | off [-y on | off]
```

The `devmgr` command accepts the following options:

`-x on | off` Turns server debugging on or off.

`-y on | off` Turns client debugging on or off.

Device Manager server and client debugging can also be turned on or off using environment variables. This is the only way to control Device Manager client debugging for RASC applications. The `devmgr_server` command `-x` and `-y` options override any environment variable settings. The `devmgr_server` command is used to start the server component of the device manager. For more information on the `devmgr_server` command, see “Device Manager Server Command” on page 176.

To turn client debugging on or off, set the `RASC_DEVMGR_CLIENT_DEBUG` variable, as follows (for the Korn shell):

```
export RASC_DEVMGR_CLIENT_DEBUG=on
export RASC_DEVMGR_CLIENT_DEBUG=ON

export RASC_DEVMGR_CLIENT_DEBUG=off
export RASC_DEVMGR_CLIENT_DEBUG=OFF
```

To turn server debugging on or off set the `RASC_DEVMGR_SERVER_DEBUG` variable, as follows (for the Korn shell):

```
export RASC_DEVMGR_SERVER_DEBUG=on
export RASC_DEVMGR_SERVER_DEBUG=ON

export RASC_DEVMGR_SERVER_DEBUG=off
export RASC_DEVMGR_SERVER_DEBUG=OFF
```

Device Manager Load FPGA Command

You can manually load an algorithm into an FPGA on demand using the `devmgr` user command `-l` option. Usually the algorithm is automatically loaded into the FPGA by the RASC abstraction layer as part of the co-processor (COP) allocation request. FPGAs loaded by the abstraction layer are marked as in use until the RASC abstraction layer frees the FPGA. FPGAs manually loaded using the `devmgr` command are not marked in

use and their algorithm can be immediately changed by another user command or the RASC abstraction layer after the `devmgr` command terminates.

The `-l` option requires an algorithm name specified with the `-n` option and a system node ID identifying the FPGA location with the `-N` option.

Note: Note that I/O nodes, including FPGA nodes, are always an odd number. Use the `devmgr` command with the `-i` option to view the FPGA inventory and their node IDs.

The `devmgr` command syntax to manually load an FPGA with an algorithm is, as follows:

```
devmgr -l -N FPGA odd numbered node ID -n algorithm name [-y on | off]
```

The `devmgr` command accepts the following options:

<code>-l</code>	Loads this FPGA with this algorithm.
<code>-N <i>FPGA odd numbered node ID</i></code>	FPGA device odd numbered node ID.
<code>-y on off</code>	Turns client debugging on or off.

Device Manager Reload FPGA Command

By default, the Device Manager only loads an algorithm into an FPGA under the following conditions:

- The FPGA is **not** already loaded with the requested algorithm, that is, the name of the algorithm loaded in the FPGA is not the same as the requested algorithm name.
- A newer version of the algorithm is available in the bitstream registry.

Otherwise, the FPGA is not loaded, eliminating the latency needed to load the FPGA.

You can use the `-r` option to direct the Device Manager to always reload an algorithm into an FPGA, even when that FPGA is already loaded with the current version of the requested algorithm. Always reloading an FPGA can be useful when debugging a suspected hardware or software FPGA load problem. Specify `on` to enable always reload an FPGA and `off` to only reload an FPGA when needed.

To determine whether the always reload option is on or off, use the `devmgr -i` command.

The `devmgr` command syntax to always reload an FPGA or only reload an FPGA when it is needed is, as follows:

```
devmgr -r on | off [-y on | off]
```

The `devmgr` command accepts the following options:

<code>-r on off</code>	Turns the always reload an FPGA action on or off.
<code>-y on off</code>	Turns client debugging on or off.

Device Manager Version Information

You can use the `devmgr -v` option to show the version, and build date and time of the `devmgr` command. This information can help you diagnose Device Manager software problems.

The `devmgr` command syntax to show version information is, as follows:

```
devmgr -v [-y on | off]
```

The `devmgr` command accepts the following options:

<code>-v</code>	Shows <code>devmgr</code> version information.
<code>-y on off</code>	Turns client debugging on or off.

Device Manager Server Command

The Device Manager includes a server component that manages the algorithm bitstream registry, allocates and frees FPGA devices from its inventory, and loads algorithm bitstreams into the FPGAs. The `devmgr_server` command is used to start the server component of the device manager.

Both `devmgr` user command and abstraction layer requests are sent over a TCP connection to the server to be processed, and the response is sent back over the same connection to the requester.

By default, the server listens for incoming connection requests on port number 9999. The port number can be changed by the system administrator by adding a port definition to the TCP/IP services file. Usually this file is located, as follows:

```
/etc/services
```

An example of how to change the listening port to 9998 is, as follows:

```
rasc_devmgr      9998/tcp      # RASC Device Manager Server
```

Note: After you make a port number change, you must restart the server (see the “Using the Device Manager Server (devmgr_server) Command”).

Using the Device Manager Server (devmgr_server) Command

After changing a port number as described in “Device Manager Server Command”, the server must be started before any requests can be processed. The `devmgr_server` command is used to start the server component of the device manager. The `debug` option enables server initialization debugging, which cannot be done using the `devmgr` command with the `-x on` action. By the time a `devmgr` command can request that server debugging be turned on, the server initialization has already completed.

The `devmgr_server` command syntax is, as follows:

```
devmgr_server [-x on | off]
```

where the `-x` option turns server debugging on or off. The default setting is off.

Device Manager Logging Facility

The Device Manager client and server provides an optional logging facility that tracks Device Manager requests and responses. Client or server logging is independently enabled by specifying the log file name using an environment variable. Log entries are time stamped and marked with the owning process and thread identifiers. The server log entries also include the process and thread identifiers of the client.

A summary of client application Device Manager activity can be obtained by enabling client logging for the application. Each client should use its own log file by specifying a unique file name for its log. This makes finding and browsing the log for a particular application easier.

The current working directory of the device manager’s server is always root. Unless a path is specified, the server log file is created in the root directory.

To enable client logging set the `RASC_DEVMGR_CLIENT_LOG` variable to the name of the file that is to contain the client log entries, as shown below for the Korn shell:

```
export RASC_DEVMGR_CLIENT_LOG=client_log_file_name
```

To disable client logging, unset the `RASC_DEVMGR_CLIENT_LOG` variable, as shown below for the Korn shell:

```
unset RASC_DEVMGR_CLIENT_LOG
```

To enable server logging set the `RASC_DEVMGR_SERVER_LOG` variable to the name of the file that is to contain the server log entries, as shown below for the Korn shell:

```
export RASC_DEVMGR_SERVER_LOG=server_log_file_name
```

To disable server logging, unset the `RASC_DEVMGR_SERVER_LOG` variable, as shown below for the Korn shell:

```
unset RASC_DEVMGR_SERVER_LOG
```

Using the GNU Project Debugger (GDB)

Note: Access to memories configured as 64-bit wide is not supported at this time in GDB FPGA (gdbfpga).

This document describes extensions to the GNU Debugger (GDB) command set to handle debugging of one or more FPGAs. Normal GDB commands and facilities are unchanged.

Brackets [] are used in this chapter to indicate that the value is optional. The [] here is not something a GDB user types; it is a syntactic convention used in this document to express an optional command field.

We use `N` here to indicate a number such as 0 or 1. The FPGA number is assigned by GDB when the FPGA is opened (loaded into the GDB command session). The first FPGA opened is assigned number zero, the next is assigned one, and so on. Numbers are not reused. GDB interacts with the RASC abstraction layer to implement several of the GDB features documented below. The `N` FPGA number is not the same as the `cop_desc` integer assigned in a user's application (that `cop_desc` integer is printed in the 'info fpga' output).

Some FPGA data is treated as if it were in registers. Such things use GDB normal register syntax, so one types a prefix \$ (and GDB shows a prefix \$) to the data name. Other FPGA data is treated as variables or arrays and normal C-like language syntax is used.

Note: For this release, GDB FPGA (gdbfpga) does NOT support DMA `stream_in` and `stream_out` functionality described in “Streaming Direct Memory Access” on page 29.

GDB Commands

The commands added to GDB unique to RASC are, as follows:

- `fpgaactive [on/off]`
- `set fpga fpganum = N`
- `info fpgaregisters [regname]`
- `info fr` (alias for ‘`info fpgregisters`’)
- `info fpga`
- `fpgastep`
- `fpgacont`
- `fpgatrace [on/off]`

Examples of standard commands and expressions with special meanings for FPGAs are, as follows:

```
print $a_0
print a_in[12]
```

where the RASC configuration has specified `a_0` as being in the Debug Port Interface of the Algorithm Block, and `a_in` is an array in an SRAM (possibly N buffered).

These are FPGA-specific and visible only when actively stepping an FPGA. For more information, see “Adding Extractor Directives to the Source Code” on page 144.

fpgaactive

The syntax for the `fpgaactive` command is, as follows:

```
fpgaactive [on,off]
```

With no operands, shows the activity-state (on or off) of the current FPGA. Defaults to `off`, meaning FPGA is inactive and registers are invisible. The RASC Abstraction layer normally maintains the off/on state when you load GDB FPGA with an executable that allocates an FPGA. The state can be set on or off manually (in GDB, not the device) with an optional operand of `on` or `off`, however, doing so is not useful, and could break an application.

set fpga fpganum

The syntax for the `fpganum` command is, as follows:

```
set fpga fpganum = N
```

This sets the current `fpga` number. This value provides an identity so when there are multiple FPGAs one knows which one is being acted/reported on. The `N` is the FPGA number assigned by GDB, not the `cop_desc` integer set in the application executable. The FPGA number is not 'recycled' in a GDB session.

fpgaregisters

The syntax for the `fpgaregisters` command is, as follows:

```
info fpgaregisters [regname]
```

Prints the list of register names and values for the current FPGA. If `regname` is given prints just that register name and value. It ignores normal processor registers.

info fpga

The syntax for the `info fpga` command is, as follows:

```
info fpga [full] [N]
```

This reports information about the FPGA set in use. That is, it displays overview information about each FPGA. The `full` option expands the information shown. The `N` option restricts output to a specific FPGA (`N` being the GDB-assigned FPGA number).

fpgastep

The syntax for the `fpgastep` command is, as follows:

```
fpgastep [N]
```

This command does a single step by default. The `N` option tells GDB to step the FPGA `N` steps. Because hardware does the `N` steps, `N` steps go as fast as a single step.

fpgacont

The syntax for the `fpgacont` command is, as follows:

```
fpgacont
```

Sets the current FPGA running to completion.

fpgatrace

The syntax for the `fpgatrace` command is, as follows:

```
fpgatrace [on/off]
```

You can use the `fpgatrace` to dynamically report FPGA use.

Turn `fpga start/stop` tracking (reporting) on or off. With no operands this reports the current state of tracking. With tracking on, each start/stop of an FPGA prints a short line giving the FPGA configuration file name, GDB number, and the ID assigned by the application. It has not escaped our attention that the mechanism could be elaborated to provide detailed information live or to a postprocessing tool to aid in understanding when multiple FPGAs are in use or the FPGA use is otherwise complicated.

Registers

FPGA values that have been mapped to debug ports and labelled by extractor directives are visible as if in the processor register set. If the user space configuration file says it is an output register (using a `REG_OUT` extractor directive), then (when the FPGA is active) the command `'print $a'` will print that register.

You must specify a register as `REG_IN` or `REG_OUT`. The register can be defined as R,W, or RW, but on a bit basis, a direction must be defined.

The requirements for this operation are that the register name is unique per FPGA register and across the entire set of register names in the processor. The FPGA registers are an extended set appended to the processor registers. It is important to note that these registers are only visible when the FPGA is active. This is a normal GDB construct with an FPGA aware extension and it adheres to all typical register syntax.

Values and Stepping

When an FPGA is running one may want to see the internal values if one is stepping the FPGA. Internal values that are exposed depends on the following:

- FPGA programming
- What the FPGA version of core-services supports
- What the device can do
- Linux device driver support

As with normal language programming, there is 'stepping' as known at the hardware level, and 'stepping' at a higher level. This is particularly relevant for a device like an FPGA because whether a given value is even 'current' or 'valid' depends on details of the FPGA state. The device program exposes 'step points' that are meaningful at a software level. Programming to the interface ensures values are meaningful at those step points, and those step points are what is iterated `fgastep`. Some device programs may not expose such logical-step points and for those device programs values printed via GDB will be more difficult to interpret sensibly (since it is not possible to print the entire FPGA internal state).

FPGA Run Status

When the GDB break command is used on an application executable that is enabled with RASC API calls, `rasclib_brkpt_start()` and `rasclib_brkpt_done()` functions, the GDB inserts break points when an FPGA is started or done, respectively.

To be able to see the FPGA state when `alg_done` occurs, you would expect to set `rasclib_brkpt_done`. However, you actually have to first set `rasclib_brkpt_start`. If you do not do so, you will not be able to see any of the FPGA registers when you hit `rasclib_brkpt_done`. The sequence is similar to the following:

```
break rasclib_brkpt_start
break rasclib_brkpt_done
handle SIGUSR1 nostop pass noprint
run
-> stop at rasclib_brkpt_start and probe registers as desired
fpgacont
-> stop at rasclib_brkpt_done and probe registers as desired
delete
cont
-> success
```

An alternative method of determining FPGA run status is the `fpgatrace` on command that reports details of FPGA starts and stops with more detail than the `break` command.

Changes To GDB Commands

The `gdb cont` command has a new side effect in `gdbfpga` of implying `fpgacont` on all FPGAs on which you can apply the `fpgastep` or `fpgacont` commands. The side effect makes `cont` get the entire application running as you would expect no matter what state FPGAs are in.

RASC Examples and Tutorials

The chapter contains Reconfigurable Application-Specific Computing Software (RASC) examples and tutorials and covers these topics:

- “System Requirements” on page 185
- “Prerequisites” on page 186
- “Tutorial Overview” on page 186
- “Simple Algorithm Tutorial” on page 187
- “Data Flow Algorithm Tutorial” on page 200
- “Streaming DMA Algorithm Tutorial” on page 206

System Requirements

For design, synthesis and bitstream generation you need the following

- PC with 1 GHz or greater clock speed
- At least 8 Gbytes random access memory (RAM)
- Red Hat Linux Enterprise version 3.0 or later
- Xilinx ISE development tools (version 9.2i, Service Pack 3 or higher)
- Optional: High-level language compiler.
- Optional: 3rd party FPGA synthesis software supporting Xilinx FPGAs (such as, Synplicity Synplify Pro 8.9 or later)
- Optional: 3rd party HDL simulation software

For bitstream download, algorithm acceleration, and real-time verification you need the following:

- One Altix system

- One or more RASC bricks or blades
- SGI ProPack 5 for Linux
- RASC software module
- A network connection to the PC detailed earlier

Prerequisites

The information and tutorials in this Examples and Tutorials section of the User Guide assume that you have previously installed and familiarized yourself with the Xilinx ISE tools and all optional software. It is also assumed that you have read the Chapter 7, “RASC Algorithm FPGA Implementation Guide” and Chapter 9, “Running and Debugging Your Application” and that you have some experience with Verilog and/or VHDL.

Additional background information (not from SGI) is available in the following:

- *Xilinx ISE Software Manuals and Help*
<http://toolbox.xilinx.com/docsan/xilinx6/books/manuals.pdf>
- *Synplicity's Synplify Pro User Guide and Tutorial*
http://www.synplicity.com/literature/pdf/synpro_ug_1001.pdf

Tutorial Overview

The following tutorials illustrate the implementation details of the algorithm programming interface using two different algorithms. During the following sections you will learn how to integrate algorithms into the RASC brick or RASC blade that are written in hardware description languages (HDLs). You will also see a subset of the optimizations that can be made for RASC implementations.

For both algorithms we will step through the entire RASC design flow: integrating the algorithm with core service; simulating behavior on the algorithm interfaces; synthesizing the algorithm code; generating a bitstream; transferring that bitstream and metadata to the Altix platform; executing an application; and using GDB to debug an application on the Altix system and FPGA simultaneously.

These tutorials only illustrate a subset of the options available for implementing an algorithm on RASC. For more details, see Chapter 3, “RASC Algorithm FPGA Hardware Design Guide” and Chapter 7, “RASC Algorithm FPGA Implementation Guide”.

The Verilog example codes are on your x86 system because Verilog codes are compiled by using Xilinx XST which runs on the x86 system. The example application C codes are compiled by the C compiler which runs on your Alitx system.

Simple Algorithm Tutorial

Overview

The first algorithm we will use to describe the interfaces and various programming templates for RASC is $(d = a \& b \mid c)$. This simple algorithm allows you to compare coding options and analyze optimization techniques. This section steps through integrating an algorithm written in Verilog and VHDL. Then it will demonstrate simulation, synthesis, bitstream generation, platform transfer, running and debugging the application. These steps are the same for this algorithm regardless of the coding technique used.

Figure 10-1 contains a diagram of the algorithm and its memory patterns.

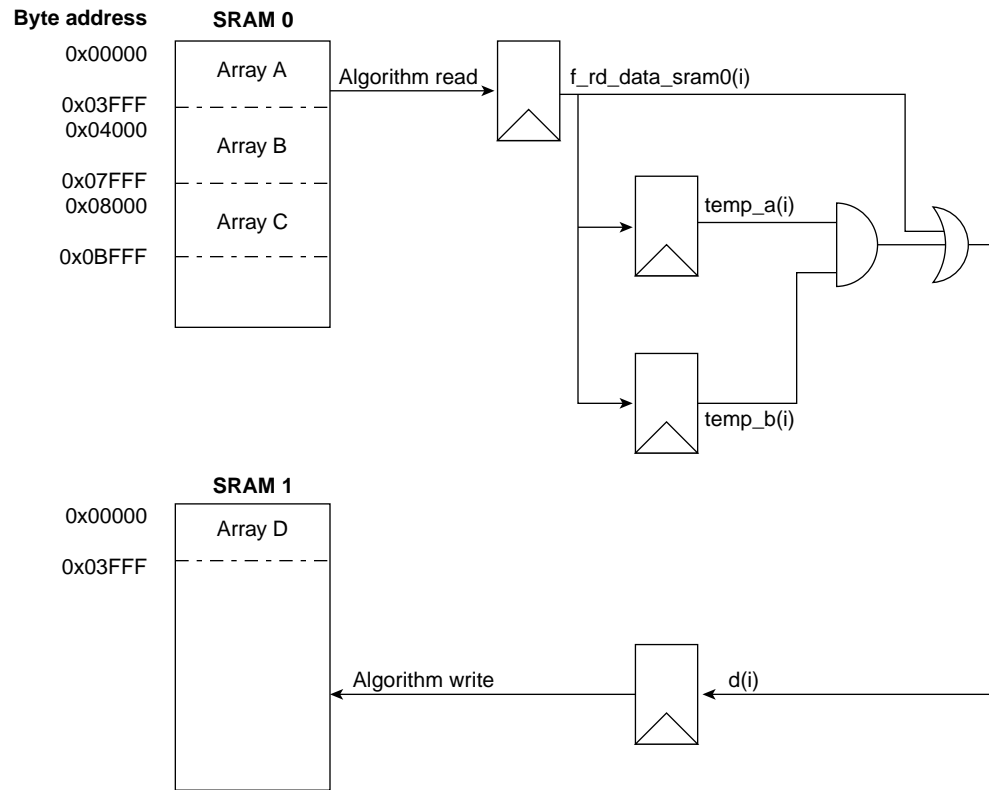


Figure 10-1 Simple Algorithm for Verilog

Application

The application that runs $(d = a \& b | c)$ on the Altix platform is fairly simple. The following code demonstrates the RASC Abstraction Layer calls that are required to utilize the RASC-brick as a Co-Processing (COP) unit. The application also runs the algorithm on the Altix box to verify its results. The application C code is on your Altix system at the following location:

```
/usr/share/rasc/examples/alg6.c
```

This simple application runs quickly on an Altix machine, although it is not optimized C code. Please note that this application is not chosen to emphasize the optimizations

available from acceleration in hardware, but rather to compare and contrast the various programming methods for RASC. As you work through the tutorials for the different languages, there will be similarities and differences that highlight advantages of one methodology versus another. For a more computationally intensive example, please see the Data Flow Algorithm in Verilog later in this chapter.

Coding Techniques: Verilog

Overview

First we will analyze how to write a Verilog version of $(d = a \& b \mid c)$ for RASC. It is important to note that the source code for this example allows for multi-buffering.

Integrating with Core Services

Begin by loading the hardware description file for the Verilog algorithm, `alg_block_top.v`, into your text editor. Change directory to `$RASC/examples/alg_simple_v` and select `alg_block_top.v`.

This file is the top level module for the algorithm. The other file that is required for this and all other implementation is the file `alg.h`. This Verilog version of $(d = a \& b \mid c)$ reads `a` from the first address in SRAM 0, `b` from the 16384th address of SRAM 0, `c` from the 32768th address in SRAM 0, and then it writes out the resulting value, `d`, to the first address of SRAM 1. Arrays `a`, `b`, `c`, and `d` are 2048 elements long where each element is a 64-bit unsigned integer, and all the arrays are enabled for multi-buffering by the RASC Abstraction Layer. The version of the algorithm, read data, write data, read address, write address, and control signals are all brought out to debug mux registers.

Figure 10-2 contains a diagram of the algorithm and its memory access patterns.

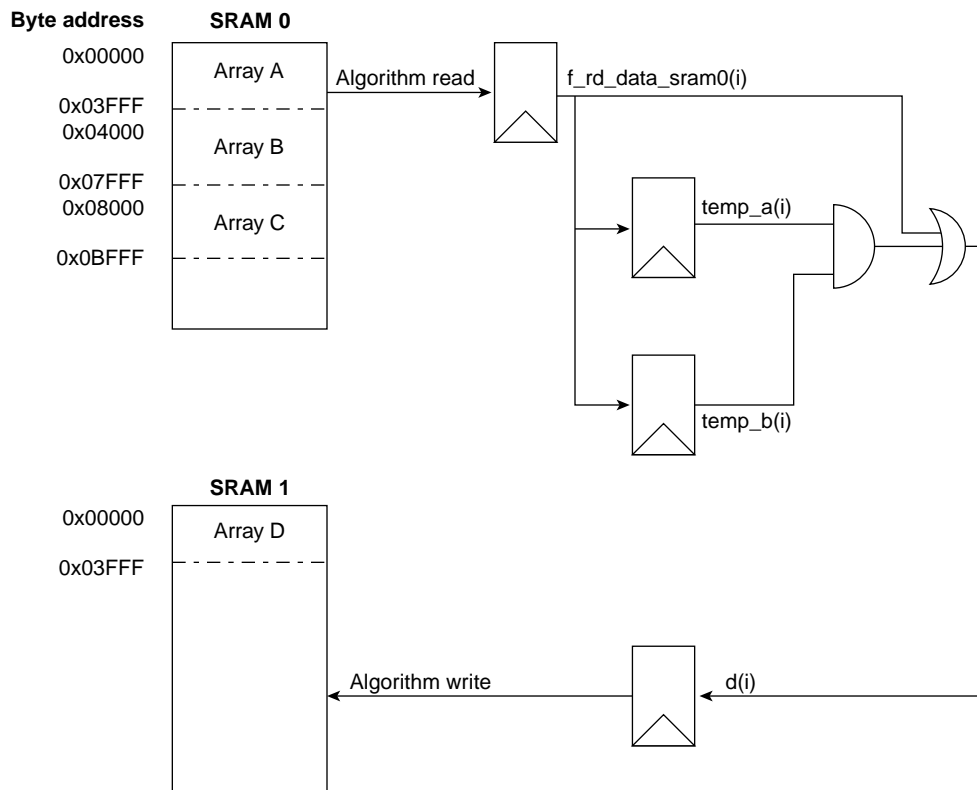


Figure 10-2 Simple Algorithm for Verilog andVHDL

The source code is available on your x86 system at the following location:

`$RASC/examples/alg_simple_v/alg_block_top.v`

Note: The handshaking methodologies (see “Handshaking Methodologies” on page 37) are not used here since it is assumed SRAM port access conflicts will never occur between the algorithm and DMA.

Extractor Comments

Other important source code considerations include adding the extractor comments that are required for accurate data movement and debugger control by RASClib. A python script called extractor parses all the Verilog, VHDL, and header files in your algorithm directory to generate the symbol tables required by GDB and to communicate to the abstraction layer the data that should be written and read from the SRAM banks.

Comment fields to generate the configuration files for the algorithm are inserted in these examples. There is a template in the `alg_core` directory, and several examples. The comment fields can be located in any file in or below the directory specified as the second argument to the extractor call (see Chapter 7, “RASC Algorithm FPGA Implementation Guide” for more detail on how to specify the makefile target). The fields are core services version, algorithm version, SRAM denoting where data will be read from and written to on the SRAM interface, register in for parameters set through an application writer’s code, or register out for a section of code that needs to be mapped to a debug register.

The debug comments for metadata parsing in this file are:

```
// extractor VERSION: 6.3
// extractor CS: 2.1
// extractor SRAM:a_in 2048 64 sram[0] 0x0000 in u stream
// extractor SRAM:b_in 2048 64 sram[0] 0x4000 in u stream
// extractor SRAM:c_in 2048 64 sram[0] 0x8000 in u stream
// extractor SRAM:d_out 2048 64 sram[1] 0x0000 out u stream
// extractor REG_IN:op_length1 10 u alg_def_reg[0][9:0]
// extractor REG_OUT:alg_id 32 u debug_port[0][63:32]
// extractor REG_OUT:alg_rev 32 u debug_port[0][31:0]
// extractor REG_OUT:rd_addr 64 u debug_port[1]
// extractor REG_OUT:rd_data_sram0_lo 64 u debug_port[2]
// extractor REG_OUT:rd_data_sram0_hi 64 u debug_port[3]
// extractor REG_OUT:wr_addr 64 u debug_port[4]
// extractor REG_OUT:wr_data_sram1_lo 64 u debug_port[5]
// extractor REG_OUT:wr_data_sram1_hi 64 u debug_port[6]
// extractor REG_OUT:cntl_sigs 64 u debug_port[7]
// extractor REG_OUT:dummy_param0_out 16 u debug_port[8][15:0]
// extractor REG_OUT:dummy_param1_out 16 u debug_port[8][31:16]
// extractor REG_OUT:dummy_param2_out 16 u debug_port[8][47:32]
// extractor REG_OUT:dummy_param3_out 16 u debug_port[8][63:48]
```

These comments are located within `alg_block_top.v` in this case, but they can be anywhere within the algorithm hierarchy as a header or source file. The core services tag helps describe which version of core services was used in generating a bitstream, this is useful with debugging. The version tag allows the user to understand from their GDB session which algorithm and revision he or she has loaded. The register out tag (REG_OUT) specifies registers that are pulled out to the debug mux. The SRAM tag is to describe arrays that are written to or read from the SRAM banks by the algorithm. For more information, see Chapter 7, “RASC Algorithm FPGA Implementation Guide”.

Coding Techniques: VHDL Algorithm

Overview

Now we will analyze how to write a VHDL version of $(d = a \& b \mid c)$ for RASC. This source code also allows for multi-buffering.

Figure 10-3 contains a diagram of the algorithm and its memory patterns.

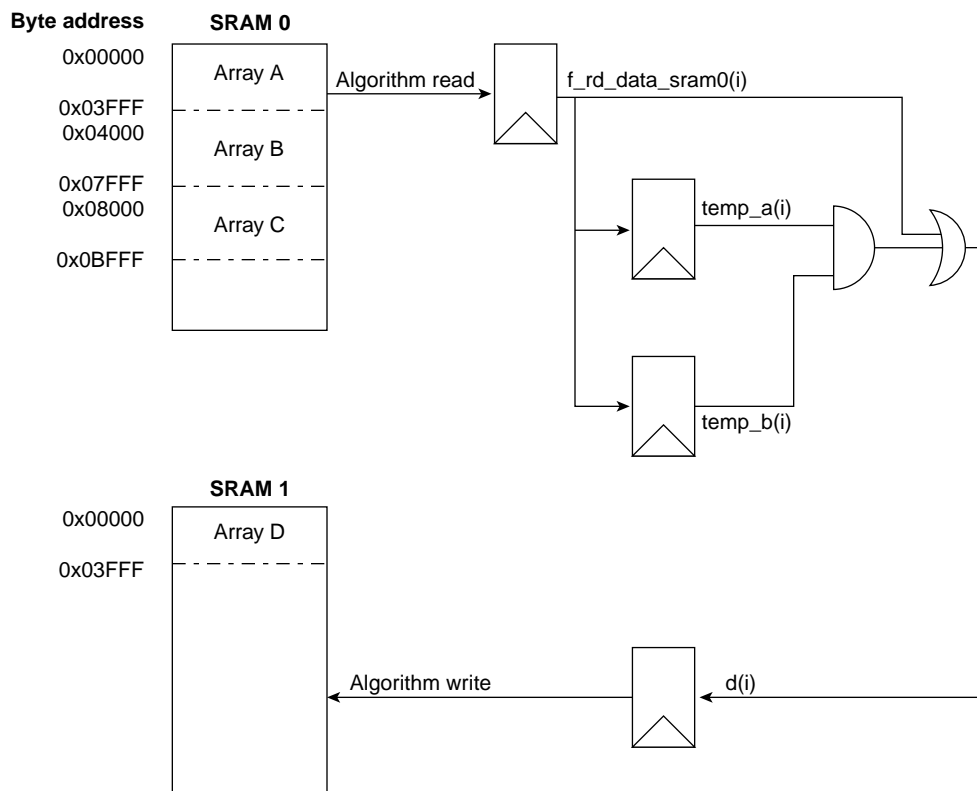


Figure 10-3 Simple Algorithm for Verilog and VHDL

Integrating with Core Services

Begin by loading the hardware description file for the VHDL algorithm, `alg_block.vhd`, into your text editor. Change directory to `$RASC/examples/alg_simple_vhd` where you will see `alg_block_top.v` and `alg_block.vhd`.

These files are the top level module and computation block, respectively. This VHDL version of $(d = a \& b \mid c)$ reads `a` from the first address in SRAM 0, `b` from the 16384th address of SRAM 0, `c` from the 32768th address in SRAM 0, and then it writes out the resulting value, `d`, to the first address of SRAM 1. Arrays `a`, `b`, `c`, and `d` are 2048 elements long where each element is a 64-bit unsigned integer, and all the arrays are enabled for multi-buffering by the RASC Abstraction Layer. The version of the algorithm, read data, write data, read address, and write address for the algorithm are all brought out to debug mux registers.

The source for the `alg_block_top.v` file is similar in this instance to the verilog version of $(d = a \& b \mid c)$, except that it performs no calculation. Instead, it wraps the `alg_block.vhd` so it can speak to the `user_space_wrapper.v` that instantiates it. There is no reason why a computation block needs to be wrapped in a Verilog module. In this case it was done for convenience.

Note: The handshaking methodologies (see “Handshaking Methodologies” on page 37) are not used here since it is assumed SRAM port access conflicts will never occur between the algorithm and DMA.

Extractor Comments

Other important source code considerations include adding the extractor comments that are required for accurate data movement and debugger control. A python script called `extractor` parses all the Verilog, VHDL, and header files in your algorithm directory to generate the symbol tables required by GDB and to communicate to the abstraction layer the data that should be written and read from the SRAM banks.

Comment fields to generate the configuration files for the algorithm are provided in this example for `alg_block.vhd`. There is a template in the `alg_core` directory, and several examples. The comment fields that can be located anywhere below the directory specified in the second argument to the extractor call (see Chapter 7, “RASC Algorithm FPGA Implementation Guide” for more detail on how to specify the makefile target). The

fields are core services version, algorithm version, SRAM denoting where data will be read from and written to on the SRAM interface, register in for parameters set through an application writer's code, or register out for a section of code that needs to be mapped to a debug register.

The debug comments for metadata parsing in this file are embedded in the VHDL code. They appear as:

```
-- extractor VERSION: 9.1
-- extractor SRAM:a_in 2048 64 sram[0] 0x0000 in u stream
-- extractor SRAM:b_in 2048 64 sram[0] 0x4000 in u stream
-- extractor SRAM:c_in 2048 64 sram[0] 0x8000 in u stream
-- extractor SRAM:d_out 2048 64 sram[1] 0x0000 out u stream
-- extractor REG_OUT:version 64 u debug_port[0]
-- extractor REG_OUT:rd_addr 64 u debug_port[1]
-- extractor REG_OUT:rd_data_sram0_lo 64 u debug_port[2]
-- extractor REG_OUT:rd_data_sram0_hi 64 u debug_port[3]
-- extractor REG_OUT:wr_addr 64 u debug_port[4]
-- extractor REG_OUT:wr_data_sram1_lo 64 u debug_port[5]
-- extractor REG_OUT:wr_data_sram1_hi 64 u debug_port[6]
```

These comments are located within `alg_block.vhd` in this case, but they can be anywhere within the algorithm hierarchy as a header or source file. The core services tag helps describe what version of core services was used in generating a bitstream. This information is useful when debugging. The version tag allows the user to understand from his GDB session which algorithm and revision he has loaded. The register out tag (REG_OUT) specifies registers that are mapped to the debug mux. The SRAM tag is to describe arrays that are written to or read from the SRAM banks by the algorithm.

Compiling for Simulation

To build the generic SSP stub test bench for the Verilog version of the simple algorithm, change to the `$RASC/dv/sample_tb` directory. At the prompt, enter the following:

```
% make ALG=alg_simple_v
```

If you do not enter the ALG tag, the makefile will default to compiling for `alg_simple_v`.

To run the diagnostic `alg_simple_v`, at the prompt, enter the following:

```
% make run DIAG=diags/alg_simple_v ALG=alg_simple_v
```

As with the case in building the test bench, the ALG tag default is `alg_simple` and it needs to be overwritten for the data flow algorithm.

The SRAM mapping with physical memory is, as follows:

```
mem0[127:64] -> qdr_sram_bank1.SMEM -> init_sram1_*.dat, final_sram1.dat
mem0[63:0 ] -> qdr_sram_bank0.SMEM -> init_sram0_*.dat, final_sram0.dat
mem1[127:64] -> qdr_sram_bank3.SMEM -> init_sram3_*.dat, final_sram3.dat
mem1[63:0 ] -> qdr_sram_bank4.SMEM -> init_sram2_*.dat, final_sram2.dat
mem2[63:0 ] -> qdr_sram_bank2.SMEM -> init_sram4_*.dat, final_sram4.dat
```

These can all be overwritten on the command line.

By specifying the SRAM input files, the user can skip the DMA process for the purposes of testing the algorithm, providing a fast check for the algorithm without verifying the DMA engines of core services.

You can also use a simple C program called `check_alg_simple.c` to verify the test results; build and run it to analyze the initial and final SRAM simulation contents.

Running a diag through this test bench produces results in four formats:

- `*vcdplus.vpd*` - this file contains the simulation results for the run.
- `*terminal output*` - the status of the test is output to the screen as it runs, notifying the user of packets sent/received.
- `*log file*` - the output to the screen is also stored in the logfile:`<diag_name>. <alg_name>.run.log` (for example, `dma_alg_simple_v.alg_simple_v.run.log`)
- `*sram output files*` - at the end of simulation (when the diag finishes because it has been successful, an incorrect packet has been received, or time-out has occurred), the contents of of all SRAMS are dumped to the corresponding `.dat` output files (the defaults or user-specified files).

Building an Implementation

When the algorithm has been integrated and verified, it is time to build an implementation.

Change directories to `$RASC/implementations/alg_simple_*/`

To synthesize the design type, use one of the following commands:

```
make synplify
```

or

```
make xst.
```

This is set up to utilize the black-boxed version of core services and should synthesize faster.

To generate the required metadata information for the abstraction layer and the debugger, you need to run the extractor script on your file. The physical design makefile includes a `make extractor` target for this purpose. When it is executed, it will generate two configuration files--one describing core services, and one describing the algorithm behavior.

To execute the ISE foundation tools and run the extractor script on the file enter the following command:

```
make all
```

This will take approximately one to two hours due to the complex mapping and place and route algorithms executed by the ISE tools. Please note that the details of setting up your own project are described in Chapter 7, "RASC Algorithm FPGA Implementation Guide", specifically in "Installation and Setup" on page 140.

Transferring to the Altix Platform

To transfer to the Altix platform, you must add your RASC design implementation into the Device Manager registry. This transfer must occur regardless of the algorithm generation method.

1. Use FTP to move the algorithm files from the PC to the `/usr/share/rasc/bitstreams` directory on the Altix machine:

```
$RASC/implementations/alg_simple_*/rev_1/alg_simple_v.bin
```

```
$RASC/implementations/alg_simple_*/<core_services>.cfg
```

```
$RASC/implementations/alg_simple_*/<user_space>.cfg
```

2. Log into the Altix machine and execute the Device Manager user command `devmgr`

```
devmgr -a -n alg6 -b alg_simple_v.bin -c <user_space>.cfg -s  
<core_services>.cfg
```

The script will default the bitstream and configuration files to these names, although the device manager can add files of any name to the registry, so users should feel free to rename project files as convenient.

Verification using GDB

To run a debug session on this bitstream, you must start the application from a GDB session window. GDB is enabled with all versions of this algorithm. To run an application using RASClib, you must execute the extended GDB on the application detailed at the beginning of this example.

```

% gdbfpga /usr/share/rasc/examples/alg6
GNU gdb 6.3.50.20050510
Copyright 2004 Free Software Foundation, Inc.
GDB is free software, covered by the GNU General Public License, and you are
welcome to change it and/or distribute copies of it under certain conditions.
Type "show copying" to see the conditions.
There is absolutely no warranty for GDB. Type "show warranty" for details.
This GDB was configured as "ia64-unknown-linux-gnu"...Using host libthread_db library
"/lib/libthread_db.so.1".

(gdb) break rasclib_brkpt_start
Function "rasclib_brkpt_start" not defined.
Make breakpoint pending on future shared library load? (y or [n]) y
Breakpoint 1 (rasclib_brkpt_start) pending.
(gdb) break rasclib_brkpt_done
Function "rasclib_brkpt_done" not defined.
Make breakpoint pending on future shared library load? (y or [n]) y
Breakpoint 2 (rasclib_brkpt_done) pending.
(gdb) handle SIGUSR1 nostop pass noprint
Signal      Stop      Print     Pass to program  Description
SIGUSR1     No       No        Yes              User defined signal 1
(gdb) run
Starting program: /usr/share/rasc/examples/alg6
Failed to read a valid object file image from memory.
[Thread debugging using libthread_db enabled]
[New Thread 2305843009292443296 (LWP 24946)]
Breakpoint 3 at 0x2000000000085bb1: file rasclib_debug.c, line 86.
Pending breakpoint "rasclib_brkpt_start" resolved
Breakpoint 4 at 0x2000000000085b82: file rasclib_debug.c, line 104.
Pending breakpoint "rasclib_brkpt_done" resolved[New Thread 2305843009303867984 (LWP 24949)]
[New Thread 2305843009312551504 (LWP 24951)]
fpga config file def reg count (64) of /var/rasc/rasc_registry/alg6/bitstream.cfg exceeds
gdb current maximum of 8, excess ignored
[Switching to Thread 2305843009292443296 (LWP 24946)]

Breakpoint 3, rasclib_brkpt_start (cop_desc=547920) at rasclib_debug.c:86
86      rasclib_debug.c: No such file or directory.
      in rasclib_debug.c
(gdb) info fpga
fpga 0
  Active      : on
  State       : ready-to-fpgastep-fpgacont
  Algorithm id : alg6
  Core svc ver : 2.100000
  Algorithm ver: 6.300000
  Algorithm src: v
  Alg. dev    :
  Alg. config : /var/rasc/rasc_registry/alg6/bitstream.cfg
  CS version  : 2.100000
  CS config   : /var/rasc/rasc_registry/alg6/core_services.cfg
  prev step ct : 0
  step ct     : 0
(gdb) info fpgaregisters
alg_id      0x6      6
alg_rev     0x3      3
rd_addr     0x0      0
rd_data_sram0_lo0xa5a4a3a2a1a09f9e 11935844831330344862
rd_data_sram0_hi0xadacabaaa9a8a7a6 12514566214034958246
wr_addr     0x0      0
wr_data_sram1_lo0xf7f4f7f2f3f0fffe 17867178244532404222
wr_data_sram1_hi0xffffcfffafbf8fff6 18445899627237015542

```

```

cntl_sigs      0x0      0
dummy_param0_out0x0      0
dummy_param1_out0x100    256
dummy_param2_out0x20     32
dummy_param3_out0x3303   13059
op_length1     0x3ff     1023
dummy_param0_in0x0      0
dummy_param1_in0x100    256
dummy_param2_in0x20     32
dummy_param3_
(gdb) fpgastep 55
(gdb) info fpgaregisters
alg_id         0x6        6
alg_rev        0x3        3
rd_addr        0x811     2065
rd_data_sram0_lo0xd5d4d3d2d1d0cfce      15408173127558025166
rd_data_sram0_hi0xdddcdbdad9d8d7d6      15986894510262638550
wr_addr        0xd        13
wr_data_sram1_lo0xd7d6d5d4d3d2d1d0      15552853473234178512
wr_data_sram1_hi0xdfdeddcdbdad9d8      16131574855938791896
cntl_sigs      0x3        3
dummy_param0_out0x0      0
dummy_param1_out0x100    256
dummy_param2_out0x20     32
dummy_param3_out0x3303   13059
op_length1     0x3ff     1023
dummy_param0_in0x0      0
dummy_param1_in0x100    256
dummy_param2_in0x20     32
dummy_param3_in0x3303   13059
(gdb) print $rd_addr
$1 = 2065
(gdb) print $wr_addr
$2 = 13
(gdb) print /x $rd_data_sram0_lo
$3 = 0xd5d4d3d2d1d0cfce
(gdb) print /x a_in[8]
[Switching to Thread 2305843009292443296 (LWP 24946)]
$4 = 0x4746454443424140
(gdb) fpgastep 3
(gdb) print $rd_addr
$5 = 2066
(gdb) print $wr_addr
$6 = 14
(gdb) print /x $rd_data_sram0_lo
$7 = 0xe5e4e3e2e1e0dfde
(gdb) print /x a_in[8]
[Switching to Thread 2305843009292443296 (LWP 24946)]
$8 = 0x4746454443424140
(gdb) fpgacont
(gdb) print $rd_addr
$9 = 3071
(gdb) print $wr_addr
$10 = 0
(gdb) print /x $rd_data_sram0_lo
$11 = 0xa5a4a3a2a1a09f9e
(gdb) print /x b_in[3]
[Switching to Thread 2305843009292443296 (LWP 24946)]
$12 = 0x1e1d1c1b1a191817
(gdb) delete
Delete all breakpoints? (y or n) y
(gdb) cont

```

```
Continuing.  
[Thread 2305843009312551504 (LWP 24951) exited]  
success  
Program exited normally.  
(gdb) quit
```

Many other commands are available. For more information on these commands, see “Using the GNU Project Debugger (GDB)” on page 178.

Data Flow Algorithm Tutorial

This example algorithm illustrates the optimization considerations for multi-buffering a complex algorithm on the RASC platform. This section steps you through the design process for this algorithm with source code in Verilog that steps by clocks.

Application

The application for the data flow algorithm is slightly more complex. This example creates a 16 KB array, sorts it from most-significant byte to least-significant byte, runs a string search on the sorted data against a match tag, and then performs a pop count. The location of application code to perform this operation on both the Altix system and the FPGA is provided below. The results are compared to verify the algorithm implementation. The application C code is on your Altix system at the following location:

```
/usr/share/rasc/examples/alg10.c
```

Loading the Tutorial

Begin by loading the hardware description files into your text editor. Change directory to `$RASC/examples/alg_data_flow_v/`

and you will see several files:

```
alg_block_top.v, sort_byte.v, string_search.v and pop_cnt.v.
```

If you look through the files you will see that the data flow algorithm reads 16K bytes of data from SRAM 0. Then it sorts the bytes of each double-word of the input data from most significant to least significant byte order. The algorithm writes those results out to SRAM 1, and then it performs a string search on the sorted data with a 16-bit match string that is provided by the application writer. The match tags resulting from the string search

are written out to SRAM 1 and a population count is then run on the data. The resulting population count is written to debug register 1.

Figure 10-4 contains a diagram of the major computational blocks and the memory access patterns for the data flow algorithm.

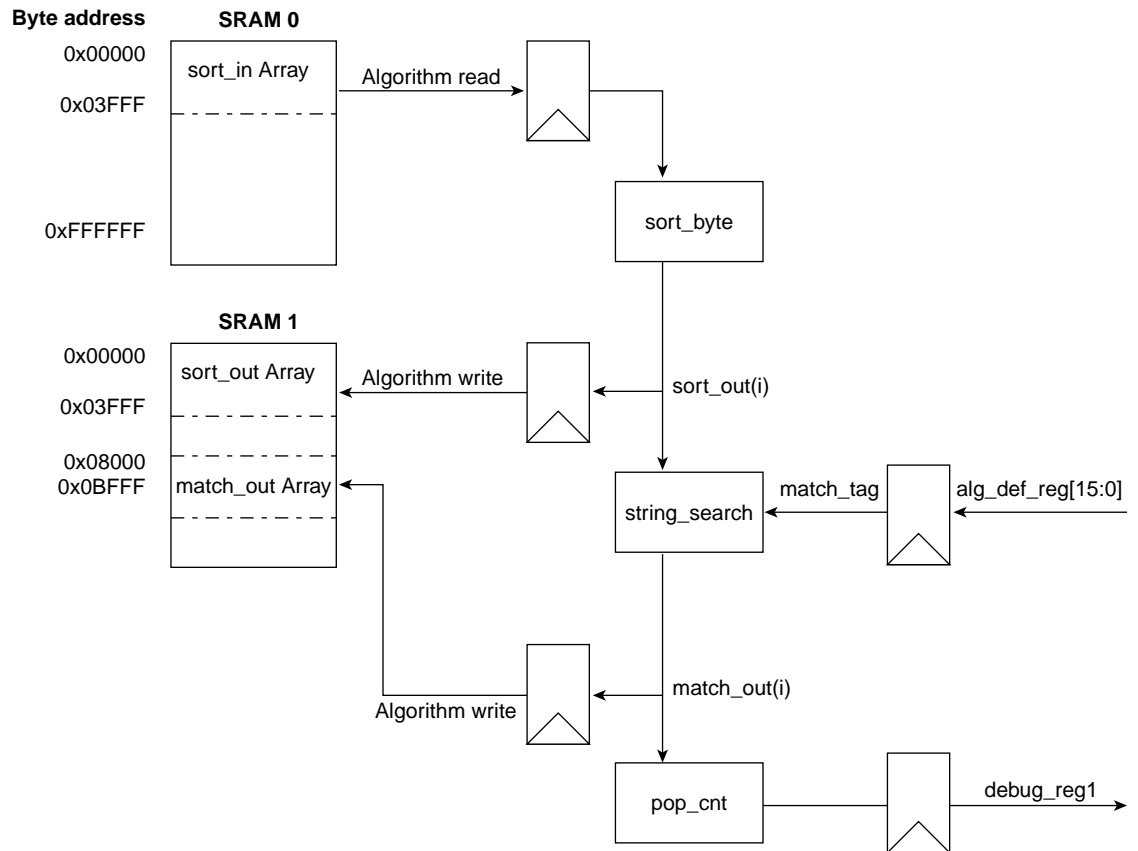


Figure 10-4 Data Flow Algorithm

Integrating with Core Services

Extractor Comments

Extractor comments are inserted in the hierarchy to describe the algorithm. In the `alg_block_top.v` file for the data flow algorithm the following comments exist:

```
// extractor CS:1.0
// extractor VERSION:10.4
// extractor SRAM:data_in          2048 64 sram[0] 0x0000 in  u stream
// extractor SRAM:sort_output      2048 64 sram[1] 0x0000 out u stream
// extractor SRAM:bitsearch_match_vector 2048 64 sram[1] 0x8000 out u stream
// extractor REG_IN:match_string 16 u alg_def_reg[0][15:0]
// extractor REG_IN:multi_iter_rst 64 u alg_def_reg[2]
// extractor REG_IN:op_length1 11 u alg_def_reg[1][10:0]
// extractor REG_OUT:version 64 u debug_port[0]
// extractor REG_OUT:running_pop_count 64 u debug_port[1]
// extractor REG_OUT:sort_input 64 u debug_port[2]
// extractor REG_OUT:sort_output 128 u debug_port[3]
// extractor REG_OUT:match_vector 64 u debug_port[5]
// extractor REG_OUT:pipe_vld 64 u debug_port[6]
// extractor REG_OUT:dw_pcnt 64 u debug_port[8]
// extractor REG_OUT:f_pcnt_63_0 64 u debug_port[9]
```

The core services tag helps describe what version of core services was used in generating a bitstream which is useful with debugging. The version tag allows the user to understand from his GDB session which algorithm and revision he has loaded. The register out tag (`REG_OUT`) specifies registers that are pulled out to the debug mux. The SRAM tag is to describe arrays that are written to or read from the SRAM banks by the algorithm.

Note: The handshaking methodologies (see “Handshaking Methodologies” on page 37) are not used here since it is assumed SRAM port access conflicts will never occur between the algorithm and DMA.

Compiling for Simulation

To build the Generic SSP stub test bench for the data flow algorithm change to the `$RASC/dv/sample_tb` directory. At the prompt, enter the following:

```
% make ALG=alg_data_flow_v
```

If you do not enter the `ALG` tag, the makefile will default to compiling for `alg_simple_v`.

To run the diagnostic `alg_data_flow_v`, at the prompt, enter the following:

```
% make run DIAG=diags/alg_data_flow_v ALG=alg_data_flow_v
```

As with the case in building the test bench, the ALG tag default is `alg_simple` and it needs to be overwritten for the data flow algorithm.

The SRAM mapping with physical memory is, as follows:

```
mem0[127:64] -> qdr_sram_bank1.SMEM -> init_sram1_*.dat, final_sram1.dat
mem0[63:0 ] -> qdr_sram_bank0.SMEM -> init_sram0_*.dat, final_sram0.dat
mem1[127:64] -> qdr_sram_bank3.SMEM -> init_sram3_*.dat, final_sram3.dat
mem1[63:0 ] -> qdr_sram_bank4.SMEM -> init_sram2_*.dat, final_sram2.dat
mem2[63:0 ] -> qdr_sram_bank2.SMEM -> init_sram4_*.dat, final_sram4.dat
```

These can all be overwritten on the command line.

By specifying the SRAM input files, the user can skip the DMA process for the purposes of testing the algorithm, providing a fast check for the algorithm without verifying the DMA engines of core services. You can also use a simple C program called `check_alg_data_flow.c` to verify the test results.

Running a diag through this test bench produces results in four formats:

- `*vcdplus.vpd*` - this file contains the simulation results for the run.
- `*terminal output*` - the status of the test is output to the screen as the it runs, notifying the user of packets sent/received.
- `*log file*` - the output to the screen is also stored in the logfile:`<diag_name>.<alg_name>.run.log` (for example, `dma_alg_data_flow_v.alg_data_flow_v.run.log`)
- `*sram output files*` - at the end of simulation (when the diag finishes because it has been successful, an incorrect packet has been received, or time-out has occurred), the contents of the 4 SRAMs are dumped to `.dat` files (the defaults or user-specified files).

Building an Implementation

When the algorithm has been integrated and verified, it is time to build an implementation.

Change directories to `$RASC/implementations/alg_data_flow_v/`

To synthesize the design type make `synplify` or `make xst`. This is set up utilize the black-boxed version of core services and should synthesize faster.

To generate the required metadata information for the abstraction layer and the debugger, you need to run the extractor script on your file. The physical design makefile includes a `make extractor` target for this purpose. When it is executed, it will generate two configuration files--one describing core services, and one describing the algorithm behavior.

To execute the ISE foundation tools and run the extractor script on the file type `make all`. This will take approximately one to two hours due to the complex mapping and place and route algorithms executed by the ISE tools. Please note that the details of setting up your own project are described in Physical Implementation chapter.

Transferring to the Altix Platform

To transfer to the Altix platform, you must add your RASC design implementation into the Device Manager Registry by performing following steps:

1. Use FTP to move the algorithm files from the PC to the `/usr/share/rasc/bitstreams/` directory on the Altix machine:

```
$RASC/implementations/alg_data_flow_v/rev_1/alg_data_flow_v.bin  
$RASC/implementations/alg_data_flow_v/<core_services>.cfg  
$RASC/inplementations/alg_data_flow_v/<user_space>.cfg
```
2. Log into the Altix machine and execute the Device Manager user command `devmgr`

```
devmgr -a -n alg10 -b alg_data_flow_v.bin -c <user_space>.cfg -s  
<core_services>.cfg
```

The script will default the bitstream and configuration files to these names, although the device manager can add files of any name to the registry, so users should feel free to rename project files as convenient.

Verification Using GDB

To run a debug session on this bitstream, you must start the application from a GDB session window. To do that, you must execute the extended GDB on the application detailed at the beginning of this example.

```

% gdbfpga /usr/share/rasc/examples/alg10
GNU gdb 6.3.50.20050510
Copyright 2004 Free Software Foundation, Inc.
GDB is free software, covered by the GNU General Public License, and you are
welcome to change it and/or distribute copies of it under certain conditions.
Type "show copying" to see the conditions.
There is absolutely no warranty for GDB.  Type "show warranty" for details.
This GDB was configured as "ia64-unknown-linux-gnu"...Using host libthread_db library
"/lib/libthread_db.so.1".

(gdb) break rasclib_brkpt_start
Function "rasclib_brkpt_start" not defined.
Make breakpoint pending on future shared library load? (y or [n]) y
Breakpoint 1 (rasclib_brkpt_start) pending.
(gdb) handle SIGUSR1 nostop pass noprint
Signal      Stop      Print     Pass to program  Description
SIGUSR1     No       No        Yes              User defined signal 1
(gdb) run
Starting program: /usr/share/rasc/examples/alg10
Failed to read a valid object file image from memory.
[Thread debugging using libthread_db enabled]
[New Thread 2305843009292443296 (LWP 1157)]
Breakpoint 2 at 0x2000000000085bb1: file rasclib_debug.c, line 86.
Pending breakpoint "rasclib_brkpt_start" resolved
[New Thread 2305843009303867984 (LWP 1160)]
[New Thread 2305843009312551504 (LWP 1169)]
fpga config file def reg count (64) of /var/rasc/rasc_registry/alg10/bitstream.cfg exceeds
gdb current maximum of 8, excess ignored
[Switching to Thread 2305843009292443296 (LWP 1157)]

Breakpoint 2, rasclib_brkpt_start (cop_desc=547920) at rasclib_debug.c:86
86      rasclib_debug.c: No such file or directory.
      in rasclib_debug.c
(gdb) info fpga
fpga 0
  Active      : on
  State       : ready-to-fpgastep-fpgacont
  Algorithm id : alg10
  Core svc ver : 2.100000
  Algorithm ver: 10.400000
  Algorithm src: v
  Alg. dev    :
  Alg. config : /var/rasc/rasc_registry/alg10/bitstream.cfg
  CS version  : 2.100000
  CS config   : /var/rasc/rasc_registry/alg10/core_services.cfg
  prev step ct : 0
  step ct     : 0
(gdb) info fpgaregisters
version      0xa00000001      42949672961
running_pop_count0x0    0
sort_input   0x0          0
sort_output  0x0          0
match_vector 0x0          0
pipe_vld     0x0          0
dw_pcmt      0x0          0
f_pcmt_63_0  0x0          0
test_cnt     0x64590000000    110333414866944
match_string 0x0          0
multi_iter_rst 0x0        0
op_length1   0x0          0
(gdb) fpgastep 5

```

```
(gdb) info fpgaregisters
version          0xa00000001      42949672961
running_pop_count0x0  0
sort_input       0x0            0
sort_output      0x0            0
match_vector     0x0            0
pipe_vld         0x0            0
dw_pcnt          0x0            0
f_pcnt_63_0     0x0            0
test_cnt         0x800000004      34359738372
match_string     0x0            0
multi_iter_rst  0x0            0
op_length1      0x0            0
(gdb) fpgastep 3
(gdb) print /x $test_cnt
$1 = 0xb00000007
(gdb) print /x data_in[8]
[Switching to Thread 2305843009292443296 (LWP 1157)]
$2 = 0x67fce141a13ee970
(gdb) fpgastep 6
(gdb) print /x $test_cnt
$3 = 0x110000000d
(gdb) print /x data_in[14]
[Switching to Thread 2305843009292443296 (LWP 1157)]
$4 = 0xbb5cf98961bed875
(gdb) delete
Delete all breakpoints? (y or n) y
(gdb) fpgacont
(gdb) cont
Continuing.
[Thread 2305843009312551504 (LWP 1169) exited]
success sorted
success match_list
popcnts = 10c
HW POP COUNT = 268
SW POP COUNT = 268

Program exited normally.
(gdb) quit
```

The above commands would execute the application, and then hit the breakpoint inserted by the `rasclib_breakpoint_start` function call. At that stage you would be able to query generic data about the FPGA that is configured in the system. Turning stepping on, you can view internal registers and arrays within the session at different steps. Many other commands are available. For more information, see “GDB Commands” on page 179.

Streaming DMA Algorithm Tutorial

This example algorithm illustrates the use of the streaming DMA feature available to an algorithm running on the RASC platform. This section steps you through the design process for this algorithm with source code in Verilog.

Application

The application for the streaming DMA algorithm sends data to the FPGA using an input stream and receives data from the FPGA using an output stream. This example creates an array of 512K 8-byte integers and increments each integer by a constant value. The application sets the number of integers and the increment value in the algorithm using extractor symbols `op_length1` and `alg_inc_val`, runs the algorithm, and verifies the results. The location of application code to perform this operation on both the Altix system and the FPGA is provided below. The results are compared to verify the algorithm implementation. The application C code is on your Altix system at the following location:

```
/usr/share/rasc/examples/alg12_strm.c
```

A second application for the streaming DMA algorithm is identical to the preceding example, except it uses an end-of-stream notification to signal the algorithm when all of the data has been sent. This application indicates to the algorithm that it will use end-of-stream notification instead of the count of integers by setting the algorithm's `use_strm_in_complete` variable; the `op_length1` variable is not used.

```
/usr/share/rasc/examples/alg12_strm_eos.c
```

Loading the Tutorial

Begin by loading the hardware description files into your text editor. Change directory to `$RASC/examples/alg_dma_stream_v/`

and you will see several files:

```
alg_block_top.v, acs_adr.v, acs_debug_reg.v and  
user_space_wrapper.v
```

If you look through the files, you will see that the streaming DMA algorithm reads the specified number of integers from the input stream, two values at a time. The algorithm then increments each integer by the specified value and writes the results, two values at a time, to the output system.

Figure 10-5 contains a diagram of the major computational blocks and the data access for the streaming DMA algorithm.

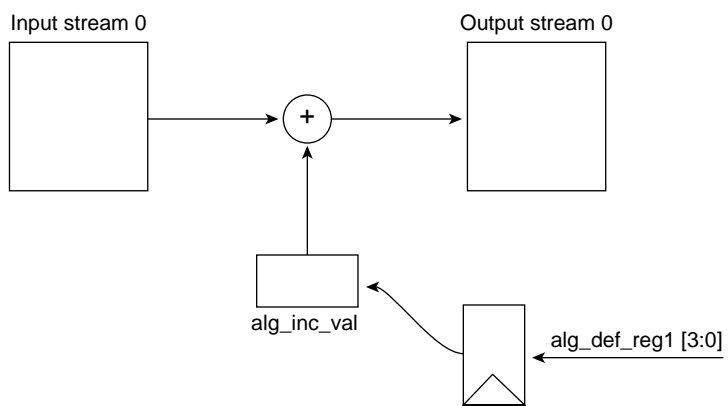


Figure 10-5 Streaming DMA Algorithm

Integrating with Core Services

Extractor Comments

Extractor comments are inserted in the hierarchy to describe the algorithm. In the `alg_block_top.v` file for the streams DMA algorithm the following comments exist:

```
// extractor CS: 2.1
// extractor VERSION:12.3
// extractor STREAM_IN:op_in_strm 0 0
// extractor STREAM_OUT:result_out_strm 0 0 //
// extractor REG_IN:op_length1 18 u alg_def_reg[0][17:0]
// extractor REG_IN:alg_inc_val 4 u alg_def_reg[1][3:0]
// extractor REG_OUT:version 64 u debug_port[0]
// extractor REG_OUT:alg_def0 64 u debug_port[1]
// extractor REG_OUT:alg_def1 64 u debug_port[2]
// extractor REG_OUT:cntl_sigs 64 u debug_port[9]
// extractor REG_OUT:rd_data_lo 64 u debug_port[10]
// extractor REG_OUT:rd_data_hi 64 u debug_port[11]
// extractor REG_OUT:wr_data_lo 64 u debug_port[12]
// extractor REG_OUT:wr_data_hi 64 u debug_port[13]
```

The core services tag helps describe what version of core services was used in generating a bitstream which is useful with debugging. The version tag allows the user to track the algorithm revision using the output from the `devmgr -q` command. The register out tag (`REG_OUT`) specifies registers that are pulled out to the debug mux. The `STREAM_IN` and `STREAM_OUT` tags describe data written to or read from DMA streams.

Compiling for Simulation

To build the Generic SSP stub test bench for the data flow algorithm change to the `$RASC/dv/sample_tb` directory. At the prompt, enter the following:

```
% make ALG=alg_dma_stream_v
```

If you do not enter the `ALG` tag, the makefile will default to compiling for `alg_simple_v`.

To run the diagnostic `alg_dma_stream_v`, at the prompt, enter the following:

```
% make run DIAG=diags/alg_dma_stream_v ALG=alg_dma_stream_v
```

As with the case in building the test bench, the ALG tag default is `alg_simple` and it needs to be overwritten for the streaming DMA algorithm.

The algorithm input stream data is coded as cache line packets in the `alg_dma_stream_v` diagnostic. The input data is identified by this leading comment:

```
print "\n\n*****Start Stream Read DMA Engine 0\n\n";
```

and cache lines of data are coded as a receive packet request and a packet response with 128 bytes of input stream data. In this diagnostic a total of 32 cache lines of input stream data is sent to the algorithm. An example is, as follows:

```
# 1 of 32
rcv_rd_req ( MEM, FCL, ANY, 0x00000000100000 ); snd_rd_rsp ( MEM, FCL, ANY, 0,
0x7654321076543210, 0x7654321076543210,
                                0x7654321076543210, 0x7654321076543210,
                                0x7654321076543210, 0x7654321076543210,
                                0x7654321076543210, 0x7654321076543210,
                                0x7654321076543210, 0x7654321076543210,
                                0x7654321076543210, 0x7654321076543210,
                                0x7654321076543210, 0x7654321076543210 );
```

The diagnostic also handles output stream data from the algorithm in cache line units. The output packet request contains the output stream data and the response acknowledges receipt of that request. The receive packet request is coded with the expected values in that cache line and a data miscompare between expected and simulated output causes the diagnostic to fail. The expected output stream data is identified by a leading comment, as follows:

```
print "\n\n*****Starting Stream 0 Write DMA Engine.\n\n";
```

In this diagnostic a total of 32 cache lines of output stream data is expected from the algorithm. An example is, as follows:

```
# 1 of 32
rcv_wr_req ( MEM, FCL, ANY, 0x00000000100000, 0x7654321076543215, 0x7654321076543215,
0x7654321076543215, 0x7654321076543215,
0x7654321076543215, 0x7654321076543215,
0x7654321076543215, 0x7654321076543215,
0x7654321076543215, 0x7654321076543215,
0x7654321076543215, 0x7654321076543215,
0x7654321076543215, 0x7654321076543215 );
snd_wr_rsp ( MEM, FCL, ANY, 0 );
```

Running a diag through this test bench produces results in four formats:

- `*vcdplus.vpd*` - this file contains the simulation results for the run.

- **terminal output** - the status of the test is output to the screen as the it runs, notifying the user of packets sent/received.
- **log file** - the output to the screen is also stored in the logfile:<diag_name>.<alg_name>.run.log (for example, alg_dma_stream_v.alg_dma_stream_v.run.log)

Building an Implementation

When the algorithm has been integrated and verified, it is time to build an implementation.

Change directories to \$RASC/implementations/alg_dma_stream_v/

To synthesize the design type make `synplify` or make `xst`. This is set up to utilize the black-boxed version of core services and should synthesize faster.

To generate the required metadata information for the abstraction layer, you need to run the extractor script on your file. The physical design makefile includes a make extractor target for this purpose. When it is executed, it will generate two configuration files--one describing core services, and one describing the algorithm behavior.

To execute the ISE foundation tools and run the extractor script on the file type make `all`. This will take approximately one to two hours due to the complex mapping and place and route algorithms executed by the ISE tools. Please note that the details of setting up your own project are described in Physical Implementation chapter.

Transferring to the Altix Platform

To transfer to the Altix platform, you must add your RASC design implementation into the Device Manager Registry by performing following steps:

1. Use FTP to move the algorithm files from the PC to the `/usr/share/rasc/bitstreams/` directory on the Altix machine:


```
$RASC/implementations/alg_dma_stream_v/rev_1/alg_dma_stream_v.bin
$RASC/implementations/alg_dma_stream_v/<core_services>.cfg
$RASC/implementations/alg_dma_stream_v/<user_space>.cfg
```
2. Log into the Altix machine and execute the Device Manager user command `devmgr`

```
devmgr -a -n alg12_strm -b alg_dma_stream_v.bin -c <user_space>.cfg -s  
<core_services>.cfg
```

The script will default the bitstream and configuration files to these names, although the device manager can add files of any name to the registry, so users should feel free to rename project files as convenient.

Verification Using GDB

The GDB debugger does not support the streaming DMA feature and cannot be used to debug the `alg_dma_stream_v` algorithm.

Device Driver

This section describes the Field Programmable Gate Arrays (FPGA) core services device driver and covers the following topics:

- “FPGA Core Services” on page 213
- “Driver Application Programming Interface (API)” on page 214
- “Example Use of Device Driver” on page 215

FPGA Core Services

All FPGAs connected to a TIO in a RASC-brick or a RASC-blade will contain a fixed set of services or core logic from SGI that control data movement, function initialization and initiation. These services are spelled out in detail in Chapter 3, “RASC Algorithm FPGA Hardware Design Guide”.

The FPGA core services device driver is implemented as a character special device driver. It provides open, close, read, write, ioctl and mmap entry points that allow access to the core services direct memory access (DMA) engines and registers.

Control and Status Registers

Specific knowledge of the core services memory-mapped registers is not required to use the core services device driver. The layout of those registers is provide in Chapter 3, “RASC Algorithm FPGA Hardware Design Guide”.

Interrupts

The device driver uses interrupts for DMA and Algorithm completion. The atomic memory operation (AMO) features of the RASC-brick will be implemented in a later revision. All system calls block until complete.

Driver Application Programming Interface (API)

The device driver API implements these system calls (see Table A-1):

Table A-1 Device Driver API System Calls

System Call	Description
<code>open()</code>	Opens the character-special file. Only one application may have the file opened at any one time.
<code>close()</code>	Closes the character-special file.
<code>read()</code>	Reads from the FPGA SDRAM to host memory.
<code>write()</code>	Writes from host memory to the FPGA SDRAM.
<code>lseek()</code>	Seeks to a specific location in FPGA SDRAM.
<code>mmap()</code>	Maps the core services registers or SDRAM into user space.

The character special files used with the device driver are, as follows:

- `/dev/RASC/acs/<nasid>/gscr`
Used for memory mapping core services memory mapped registers.
- `/dev/RASC/acs/<nasid>/sram`
Used for all other system calls, for example. reading, writing, and memory mapping SRAM.

Input Direct Memory Access

The input direct memory access (DMA) block is comprised of four stream DMA engines that target the algorithm for the data coming in from main memory and two block DMA read engines that targets SRAM for the data coming in from main memory. For more information, see “RASC Core Services Overview” on page 22.

Output Direct Memory Access

The output DMA block is comprised of four stream DMA engines that take data directly from the algorithm and target main memory and two block DMA write engines that takes data from SRAM and targets main memory. For more information, see “RASC Core Services Overview” on page 22.

Function Control

The `ioctl` command `COP_IOCTL_ALGO_START` is used to start the algorithm in the FPGA. Use `COP_IOCTL_ALGO_STEP` instead of `COP_IOCTL_ALGO_START` to start the algorithm but only run a given number of clocks. Successive `COP_IOCTL_ALGO_STEP` `ioctl` calls will move the clock. The `COP_IOCTL_ALGO_CONT` call is used to drive the algorithm to completion after stepping.

Upon successful completion, zero is returned. Otherwise -1 is returned and `errno` is set.

Example Use of Device Driver

This section provides an example of using the FPGA core services device driver.

```
#include "sys/types.h"
#include "stdio.h"
#include "stdlib.h"
#include "unistd.h"
#include "string.h"
#include "fcntl.h"
#include "sys/ioctl.h"

#include "acs.h"

/*
 * FPGA algorithm info.
 *
 * (This is a simple test-only algorithm built largely for software
 * testing.)
 *
 *      FPGA algorithm is 'a = a & b | c'.
 *      Max number of each operands is 512 64-bit words.
 *      Operand 'a' input starts at word zero of sram0
 *      Operand 'b' input starts at word zero of sram1
```

```
*      Operand 'c' input starts at word zero of sram2
*      The results are placed at word zero of sram0 (overwriting
operand a)
*/

typedef unsigned long long uint64_t;

#define RAM_SIZE (2*1024*1024)          // individual SRAM size
#define SRAM0 0
#define SRAM1 RAM_SIZE
#define SRAM2 (RAM_SIZE*2)

#define OPERAND_COUNT 512

/* 64-bit words to bytes */
#define W2B(wc) (wc*sizeof(uint64_t))

uint64_t operand_a[OPERAND_COUNT];
uint64_t operand_b[OPERAND_COUNT];
uint64_t operand_c[OPERAND_COUNT];
uint64_t hard_results[OPERAND_COUNT];
uint64_t soft_results[OPERAND_COUNT];

/*
 * host processor version of algorithm
 */
void
soft_algo(uint64_t *r, uint64_t *a, uint64_t *b, uint64_t *c, int
count)
{
    int i;

    for (i=0; i < count; i++)
        r[i] = (a[i] & b[i]) | c[i];
}

/*
 * FPGA version of algorithm
 */
int
hard_algo(uint64_t *r, uint64_t *a, uint64_t *b, uint64_t *c, int
count)
{
    int fd;
```

```
int bufsize = W2B(count);
int n, rv = 0;
char *path = "/dev/RASC/acs/1/sram";

/*
 * open core services
 */
fd = open(path, O_RDWR);
if (fd < 0) {
    perror("open");
    return -1;
}

/*
 * move operand 'a'
 */
lseek(fd, SRAM0, SEEK_SET);
n = write(fd, a, bufsize);
if (n != bufsize) {
    fprintf(stderr, "wrote %d of %d bytes\n", n, bufsize);
    perror("write");
    rv = -2;
    goto exit;
}

/*
 * move operand 'b'
 */
lseek(fd, SRAM1, SEEK_SET);
n = write(fd, b, bufsize);
if (n != bufsize) {
    fprintf(stderr, "wrote %d of %d bytes\n", n, bufsize);
    perror("write");
    rv = -3;
    goto exit;
}

/*
 * move operand 'c'
 */
lseek(fd, SRAM2, SEEK_SET);
n = write(fd, c, bufsize);
if (n != bufsize) {

    fprintf(stderr, "wrote %d of %d bytes\n", n, bufsize);
```

```
        perror("write");
        rv = -4;

        goto exit;
    }

    /*
     * start the FPGA algorithm and wait for it to complete
     */
    if (ioctl(fd, COP_IOCTL_ALGO_START, 0) < 0) {
        perror("ioctl");
        rv = -5;
        goto exit;
    }

    /*
     * get the results
     */
    lseek(fd, SRAM0, SEEK_SET);
    n = read(fd, r, bufsize);
    if (n != bufsize) {
        fprintf(stderr, "read %d of %d bytes\n", n, bufsize);
        perror("read");
        rv = -6;
        goto exit;
    }

exit:
    close(fd);

    return rv;
}

void
compare_results(uint64_t *hard_results, uint64_t *soft_results, int
count)
{
    int i;
    int misses = 0;

    for (i=0; i < count; i++)
        if (hard_results[i] != soft_results[i]) {
            if (!misses)
                printf("miscompare at %d: got: 0x%llx expected:
0x%llx\n",
```



```
                i, hard_results[i], soft_results[i]);
                misses++;
            }

            printf("%d result mismatches between algorithms\n", misses);
        }

void
operand_set(uint64_t *operand, int operand_count, uint64_t value)
{
    int i;

    for (i=0; i < operand_count; i++)
        *operand++ = value;
}

int
main(int argc, char **argv)
{
    int operand_count, bufsize;
    int n, c;

    uint64_t val_a = 'a';
    uint64_t val_b = 'b';
    uint64_t val_c = 'c';
    int verify = 0;

    operand_count = OPERAND_COUNT;

    while ((c = getopt(argc, argv, "a:b:c:s:v")) != EOF) {
        switch(c) {
            case 'a':
                val_a = strtoull(optarg, 0, 0);
                break;
            case 'b':
                val_b = strtoull(optarg, 0, 0);
                break;
            case 'c':
                val_c = strtoull(optarg, 0, 0);
                break;
            case 's':
                operand_count = strtoull(optarg, 0, 0);
                break;
        }
    }
}
```

```
        case 'v':
            verify++;
            break;
    }
}

if (operand_count > OPERAND_COUNT || operand_count <= 0) {
    printf("invalid operand_count. ");
    printf("must be: 0 > operand_count <=
OPERAND_COUNT\n");
    return -1;
}

/* syscalls & lib functions usually require a byte count */
bufsize = W2B(operand_count);

/*
 * initialize input and output buffers
 */
operand_set(operand_b, operand_count, val_b);
operand_set(operand_c, operand_count, val_c);
operand_set(operand_a, operand_count, val_a);
memset(hard_results, 0x0, bufsize);

/*
 * optionally generate results on host processor
 */
if (verify) {
    printf("Starting the software algorithm\n");
    soft_algo(soft_results, operand_a, operand_b,
operand_c,
                                                    operand_count);
}

/*
 * run the FPGA algorithm
 */
printf("Starting the hardware algorithm\n");
n = hard_algo(hard_results, operand_a, operand_b, operand_c,
                                                    operand_count);
if (n < 0) {
    printf("hard_algo failed: %d\n", n);
    return -1;
}
}
```

```
        printf("example results:\n    0x%llx & 0x%llx | 0x%llx =\n\n",\n               operand_a[0], operand_b[0], operand_c[0],\n               hard_results[0]);\n\n        if (verify) {\n            compare_results(hard_results, soft_results,\n                           operand_count);\n        }\n\n        printf("done\n");\n\n        return 0;\n    }
```


SSP Stub User's Guide

Introduction to SSP Stub

The Scalable System Port (SSP) Stub is a verification tool intended to help simulate and verify an algorithm designed for a RASC Field Programmable Gate Array (FPGA). This tool has been created to assist in initial testing and debugging of an algorithm in simulation prior to loading the algorithm onto the FPGA. Used with VCS/Virsim, the Stub allows the user to simulate sending/receiving SSP packets to/from the RASC FPGA to transfer data, start and stop the algorithm, and check status. This appendix covers the following topics:

- “Recommended Reading” on page 223
- “Verification Environment and Testbench” on page 224
- “SSP Stub Commands” on page 229
- “Sample Diagnostic” on page 237
- “Using the Stub” on page 243

Recommended Reading

Related and recommended documents include the following:

- *VCS User Guide*
- *GNU Make Manual* (or equivalent text on the Make utility)

The *VCS User Guide* and *GNU Make Manual* are available through the world-wide web.

Verification Environment and Testbench

Provided with the SSP Stub is a verification environment and sample testbench. This consists of the files which make up the SSP Stub as well as supporting files located in a directory tree. The directory tree, testbench, and sample tests are intended to help you quickly gain proficiency in creating tests for a specific algorithm. This section covers the following topics:

- “Verification Environment” on page 224
- “Sample Test Bench” on page 224
- “Compiling and Running a Test” on page 227

Verification Environment

The SSP Stub is intended for use on Verilog designs and is customized for use with the VCS/Virsim simulator. To insure that SSP Stub and sample testbench function properly, the following environment variables must be set to the correct directories:

- VCS_HOME (*your_vcs_install_directory*)
- VCSPLIDIR `$VCS_HOME/gui/virsim/linux/vcdplus/vcs` (*your_vcs_version*)
- PATH `$PATH\:$VCS_HOME/bin`

Sample Test Bench

The SSP Stub consists of Verilog modules as well as PLI calls to functions written in C code. The stub is instantiated in a sample Verilog testbench along with the RASC FPGA. The C-code and top-level stub modules are located in the *your_root/dv/sample_tb* directory. Here, *your_root* is the installation directory of the RASC Core Services design and verification tree. The RASC Core Services modules are located in the design branch of the tree under the *your_root/design* directory. The algorithm being tested must be created in the design branch of the tree, located in the directory *your_root/design/alg_core/alg_name*. See existing sample algorithms for exact file locations. The sample test bench relies on the algorithm to be in this location for compiling.

The primary Verilog modules in the SSP Stub and other files are listed below. They can be found in the *sample_tb* directory.

- `top.v`: Top level of the Sample Test Bench containing the SSP Stub, Algorithm FPGA, clock generator, and SRAM instances. Note that there are five physical SRAM instances in `top.v` while there are three logical SRAMs. The SRAMs are used in pairs to form a single logical SRAM. This structure mirrors the real hardware.
- `ssp_stub.v`: Top level Verilog of the SSP Stub which passes signals to and from conversion modules.
- `init_sram0_good_parity.dat`, `init_sram2_good_parity.dat`, `init_sram1_good_parity.dat`, `init_sram3_good_parity.dat`, and `init_sram3_good_parity.dat`: These SRAM initialization files contain data which is automatically loaded into the respective SRAM simulation models at the beginning of simulation. The data is in a format which the SRAM simulation model uses (one bit of parity per byte of data is shifted in with the data). These default files can be overridden by the user on the command line at runtime.
- `final_sram0.dat`, `final_sram1.dat`, `final_sram2.dat`, `final_sram3.dat`, and `final_sram4.dat`: These files contain data extracted from the respective SRAM simulation models at the end of simulation. These default files can be overridden by the user on the command line at runtime.
- `timescale.v`: This file contains the Verilog timescale of each of the components of the SSP Stub, as well as the algorithm FPGA design files. It is required that the algorithm being simulated makes use of the same timescale as the rest of the design.

For simulation, the Verilog timescale of each of the components of the SSP Stub, as well as the RASC FPGA design files, is set by the `timescale.v` file in the `sample_tb` directory. It is required that the algorithm being simulated makes use of the same timescale as the rest of the design.

SSP Stub File Descriptions

The SSP Stub consists of the following files:

- `ssp_stub.v`: Top level Verilog of the SSP Stub which passes signals to and from conversion modules.
- `cx_stub_ssr.v`: A component of the SSP Stub which converts the raw 64-bit Double Data Rate (DDR) flits from the SSP link into 128-bit Single Data Rate (SDR) flits for use in the stub.
- `ssp_rcv_recap.v`: A component of the SSP Stub which captures data (from `cx_stub_ssr.v`) on rising edge.

- `cx_stub_ssd.v`: A component of the SSP Stub which converts the 128-bit SDR data from the stub to 64-bit DDR data for transmitting over the SSP link.
- `ssp_sdr_stub.v`: A component of the SSP Stub which passes 128-bit data to and from C-code.
- `tio_shrd_ecc_generate.v`: A component of the SSP Stubh which generates ECC bits in accordance with SSP protocol.
- `pli.tab`: This PLI file links the `$start_ssp` call from the Verilog to the `start_ssp()` routine in C code.
- `start_ssp.c`: This file directs the initialization of the stub. It opens the diag file, parses the commands, sets up the packet queues, initializes counters, and calls `send_rcv_flits()` to enter the main loop.
- `user_const.h`: A header file containing user specified constants (e.g. simulation timeout, maximum size of the input file, maximum number of commands). Additional constants may be added to this file as needed by the user.
- `ssp_defines.h`: This header file contains macros for the program, such as the names of packet types.
- `send_rcv_flits.h`: This header file contains the main loop function, `send_rcv_flits()`, which calls itself once every clock cycle through the `tf_setdelay(5)` PLI call [`tf_setdelay(time)` waits for the specified amount of time, then calls `misc_tf` with input "reason_reactivate"; in the `pli.tab` file, the `misc` function is specified to be sent to "send_rcv_flits"].
`send_rcv_flits()` runs through the various conditions for sending, receiving, timing out, printing, delaying, and polling. It uses the function `process_pkt` to get information about error packets (in the syntax used for the input diag file) and the function `get_fields` to get a breakdown of the command word for logging purposes. It uses the sub-functions `send_flit()`, `rcv_flit()`, `snd_poll()`, and `rcv_poll()` to place data on the outgoing lines and read data off of the incoming lines. It calls `finish_ssp()` to print success and exit. `send_flit()` puts data for the next flit on the outgoing lines. `rcv_flit()` takes data from the incoming lines and stores it after comparing it with expected data. `snd_poll()` sends out a PIO 8-Byte Read Request packet for a given address. `rcv_poll()` takes the read response to a poll request, and checks the bit in question to see if it has been set.
- `queue_pkt.h`: This file contains the mechanism for reading an input command line from the diag and converting it into data values representing the packet that will need to be sent/received OR the information for a `print/delay/poll`. `queue_pkt(string)` goes to tokenizes the input string and extracts data from

each field. `q_string_it(token, pkt_string)` appends the latest token to the stored packet string (for logging purposes). `strtok_checked(s1, s2)` calls `strtok` on the inputs and throws an error if the resulting token is NULL when it should not be.

- `setup_pkt.h`: This file contains the function `setup_pkt(snd_rcv)`, which takes care of setting up packets for regular sends and receives (not poll packets), based on the information stored by `queue_pkt(string)`. It uses the functions in `snd_rcv_fns.h` to get the data for each flit, and sets up the remaining fields individually (head, tail, error, `req_dval`, `rsp_dval`).
- `snd_rcv_fns.h`: This file contains the equivalent of the eleven packet functions available to the diag, but with an extra field for the `pkt[]` array which will contain their results. Each function passes the appropriate type and fields to the `construct_pkt(type, tnum, addr, data, error, pkt, snd_rcv_n)` function.
- `construct_pkt.h`: This file takes the packet information from `snd_rcv_fns.h` and turns it into 64-bit segments of data. `construct_pkt(type, tnum, address, data, error, pkt, to_from_n)` takes the input information and creates appropriate data, storing this data in `pkt`. `pkt_size(type)` takes the type and returns a value for the number of flits in that packet (1, 2, 9, or 10).
- `make_command.h`: This file takes the type, transaction number, error bit, and direction of a packet and constructs an SSP command word, which it returns as an int.
- `get_fields.h`: This file takes input of the type of packet to deconstruct (i.e., send, receive, incoming, poll request, poll response, poll expected) and returns the command word with its decomposed fields as well as any address or data fields. It also contains a helper function, `f_string_it(token)`, which appends information to its destination string (`data_fields[]`) as it proceeds.
- `process_pkt.h`: This file takes the type of packet (either incoming or poll response) and turns the data from that receive packet into a string following the syntax of the diag file. It has a helper function `p_string_it(token)` that behaves the same as `f_string_it(token)`, but stores its data in `processed_string[]`.

Compiling and Running a Test

Compiling the sample testbench is done using the provided `Makefile`. In order to compile the sample testbench including the SSP Stub and the RASC Core Services logic,

an algorithm must be specified. This algorithm is specified to the make utility on the command line, (change directory to the `sample_tb` directory and enter the following command) as follows:

```
% cd your_root/dv/sample_tb
% make ALG=your_algorithm
```

The Makefile uses *your_algorithm* to find the directory in the tree where you algorithm design files exist (for example, `ALG=alg0`). For the case where no `ALG=your_algorithm` is specified, the default algorithm is used from `your_root/examples/alg_simple_v`. As mentioned earlier, it is required that user generated algorithms are saved in a corresponding directory (for example, `your_root/examples/my_alg/*.v`).

To run a test, remain in the `sample_tb` directory and enter the following command:

```
% make run DIAG=diag_filename ALG=your_algorithm SRAM0_IN=sram0_input_filename
SRAM1_IN=sram1_input_filename SRAM2_IN=sram2_input_filename
SRAM3_IN=sram3_input_filename SRAM0_OUT=sram0_output_filename
SRAM1_OUT=sram1_output_filename SRAM2_OUT=sram2_output_filename
SRAM3_OUT=sram3_output_filename
```

The *diag_filename* specifies the diag to be run and should be relative to the current directory. Again, the algorithm must be specified using the `ALG=your_algorithm` command line option. If none is specified, the runtime command uses same default as above. The ALG option allows the user to reuse the same diag for multiple algorithms. The input and output data files for the three SRAMs may be specified; if they are not specified, they default to:

```
SRAM0_IN=init_sram0_good_parity.dat
SRAM1_IN=init_sram1_good_parity.dat
SRAM2_IN=init_sram2_good_parity.dat
SRAM3_IN=init_sram3_good_parity.dat
SRAM4_IN=init_sram4_good_parity.dat
SRAM0_OUT=final_sram0.dat
SRAM1_OUT=final_sram1.dat
SRAM2_OUT=final_sram2.dat
SRAM3_OUT=final_sram3.dat
SRAM4_OUT=final_sram4.dat
```

As the test runs, it will output status to the screen and to an output file named *diag_filename.your_algorithm.run.log*. This log file will appear in the same directory that the diagnostic is located in. The contents of each SRAM at the end of simulation will be dumped into `.dat` files which can be either the default filenames or user-specified. Each time a diagnostic is run, the file `vcdplus.vpd` is generated in the `sample_tb`

directory. This file can be input to Virsim for viewing the waveform. Since the file `vcdplus.vpd` is generally large, it is overwritten for each diagnostic run. To save the waveform for a given diagnostic, copy corresponding `vcdplus.vpd` file to a new name.

To view the waveform saved in the `vcdplus.vpd` file, use the following command:

```
% vcs -RPP vcdplus.vpd
```

When the stub receives an incorrect packet, it will output, in order, the following information: the command for the next expected packet, SSP fields of the expected packet, the command translation (if one exists) for the received packet, and the SSP fields of the received packet.

SSP Stub Commands

The Stub retrieves instructions from a test through a text input file, the diagnostic. The SSP Stub parses this file at each semicolon to extract commands which the Stub executes. Many of the allowed commands in a diagnostic correspond to SSP packets. The SSP Stub supports other commands that are used for diagnostic writing and debugging. The primary components of the diagnostic file are, packet commands, other commands, and comments.

It is important to note that most SSP packets come in pairs, a request and a response. For these types of packets, the request command and response command **must** be listed sequentially in a diagnostic. This method of keeping requests and response paired is used by the stub to associate request and response packets with the corresponding trnums. Also, when running the DMA engines, all transactions related to that sequence of events should be grouped together. See the sample diagnostic included later in this document for an example of how this is done.

Packet Commands

RASC makes use of a subset of the available SSP packet types. The packet types used by the SSP Stub are included in the table below.

Table B-1 Packet Types used by SSP

Type	Transaction	Valid Transactions Stub to FPGA	Valid Transactions FPGA to Stub
Reads	PIO 8-Byte Read	Yes	No
	PIO Partial 128-Byte Read	Yes	No
	Memory Full 128-Byte Read	No	Yes
Writes	PIO 8-Byte Write	Yes	Yes
	PIO Partial 128-Byte Write	Yes	No
	Memory Full 128-Byte Write	No	Yes
Invalidate-Cache Flush	Invalidate-Cache Flush	Yes	No

Of the standard SSP packet types listed in the SSP Specification, this stub does **NOT** support FPGA-to-Stub Processor Input/Output (PIO) 8-Byte Reads, FPGA-to-Stub Memory Partial 128-Byte Writes, FPGA-to-Stub PIO Partial 128-Byte Reads or Writes, or Graphics Writes. It does, however, support all SSP packets that are used or accepted by the RASC Core Services.

Packet commands begin with the name "snd_" or "rcv_", contain fields within parentheses separated by commas, and end with a semicolon, for example:

```
snd_wr_req(PIO, DW, ANY, 0x0000000000000000, 0x000000000000FFFF);
```

There are two categories of stub commands: sends and receives. The diagnostic should specify both the requests and responses it wishes to send as well as the requests and responses it expects to receive from the FPGA.

Command Fields

Below is a list of the various fields used by packet commands. A list of allowable values and their corresponding meanings is also included.

- Field: PIO/MEM
PIO (Processor Input/Output): Use for read/write requests initiated by the stub, also used for interrupt packets from the FPGA.
MEM (Memory): Use for read/write requests initiated by the FPGA (except interrupts).
- Field: Size
DW (Double Word): Use for 8-Byte transactions (MMRs and Interrupts)
FCL (Full Cache-Line): Use for Full 128-Byte transactions (usually DMA's).
- Field: Tnum
0x0 - 0xff: Correspond to the tnum of a packet.
ANY: Used when any tnum is allowable or when responding with the tnum of the corresponding request.
For outgoing requests, the tnum may be specified (though overlapping tnums for the same type of transactions should not be used per the SSP Specification). If "ANY" is specified for an outgoing request, the Stub will assign a tnum automatically. For incoming responses, the tnum may be specified if the tnum of the corresponding request was specified; in all other cases, specifying "ANY" will allow the stub to accept any tnum for the incoming packet.
Also, note that the RASC FPGA will send out PIO Write Requests to the stub with tnums always in the range 0xf0 - 0xff. The RASC FPGA should not use tnums in this range for any other purpose.
- Field: Address
Any 64-bit value: Used to specify the intended or expected address of a packet (note that actual addresses are 56-bit values). For all requests, this field specifies the SRAM, MMR, or Memory address location.
- Field: Data
A list of 64-bit values, dependent on the size specified. Each 64-bit data segment (from least-significant to most-significant) is separated by a comma. For DWs, there is just one data segment. For FCLs, there are 16 segments. The sample diagnostic included later in this document contains examples.
- Field: Error
0: No error is reported in the header of the SSP packet.

1: An error is reported/expected in the corresponding packet.

The error field is required for all responses, whether received or sent. The error bit is one when there has been some error in the transaction; for example, the type of request may not be valid or the address may be invalid.

Send Commands

snd_wr_req

The syntax of the `snd_wr_req` function is, as follows:

```
snd_wr_req(pio_mem_n, size, tnum, addr, data):
```

The `snd_wr_req` function is used to send PIO writes to the memory-mapped registers (MMRs) (8-Byte).

- `pio_mem_n` = PIO
- `size` = DW
- `tnum` = ANY or user-specified
- `addr` = address of MMR or in SRAM
- `data` = data to be written

snd_rd_req

The syntax of the `snd_rd_req` function is, as follows:

```
snd_rd_req(pio_mem_n, size, tnum, addr):
```

The `snd_rd_req` function is used to request PIO reads from the MMRs (8-Byte) and SRAMs (Full 128-Byte).

- `pio_mem_n` = PIO
- `size` = DW or FCL
- `tnum` = ANY or user-specified
- `addr` = address of MMR or in SRAM

snd_wr_rsp

The syntax of the `snd_wr_rsp` function is, as follows:

```
snd_wr_rsp(pio_mem_n, size, tnum, error):
```

The `snd_wr_rsp` function is used to respond to write requests from the FPGA during DMAs (Full 128-Byte) and interrupts (8-Byte).

- `pio_mem_n` = PIO (interrupts) or MEM (DMAs)
- `size` = DW (interrupts) or FCL (DMAs)
- `tnum` = ANY or user-specified
- `error` = 0 or 1

snd_rd_rsp

The syntax of the `snd_rd_rsp` function is, as follows:

```
snd_rd_rsp(pio_mem_n, size, tnum, error, data):
```

The `snd_rd_rsp` function is used to respond to read requests from the FPGA: DMAs (Full 128-Byte).

- `pio_mem_n` = MEM
- `size` = FCL
- `tnum` = ANY or user-specified
- `error` = 0 or 1
- `data` = 0 if this is an error response, otherwise a list of 16 64-bit data values, separated by commas.
- `tnum` = ANY or user-specified

inv_flush

The syntax of the `inv_flush` function is, as follows:

```
inv_flush(tnum):
```

- `tnum` = ANY or user-specified

Receive Commands

rcv_wr_rsp

The syntax of the `rcv_wr_rsp` function is, as follows:

```
rcv_wr_rsp(pio_mem_n, size, tnum, error):
```

The `rcv_wr_rsp` function is used to receive write responses from the FPGA (after writes to MMRs).

- `pio_mem_n` = PIO
- `size` = DW
- `tnum` = ANY or user-specified
- `error` = 0 or 1

rcv_rd_rsp

The syntax of the `rcv_rd_rsp` function is, as follows:

```
rcv_rd_rsp(pio_mem_n, size, tnum, error, data):
```

The `rcv_rd_rsp` function is used to receive read responses from the FPGA (after read requests to the MMRs).

- `pio_mem_n` = PIO
- `size` = DW
- `tnum` = ANY or user-specified
- `error` = 0 or 1
- `data` = 0 if an error response, one 64-bit value otherwise

rcv_wr_req

The syntax of the `rcv_wr_req` function is, as follows:

```
rcv_wr_req(pio_mem_n, size, tnum, addr, data):
```

The `rcv_wr_req` function is used to receive write request from the FPGA (interrupts and DMA writes).

- `pio_mem_n` = PIO (interrupts) or MEM (DMAs)
- `size` = DW (interrupts) or FCL (DMAs)
- `tnum` = ANY or user-specified
- `addr` = address specified for interrupts or system memory address (64-bits)

- data = one 64-bit value for interrupts; 16 64-bit values for DMAs

rcv_rd_req

The syntax of the `rcv_rd_req` function is, as follows:

```
rcv_rd_req(pio_mem_n, size, tnum, addr):
```

The `rcv_rd_req` function is used to receive read requests from the FPGA (DMA reads).

- `pio_mem_n` = MEM
- `size` = FCL
- `tnum` = ANY or user-specified
- `addr` = system memory address

Other Commands

```
print "text";
```

The user can output information from the diagnostic through the `print` command. These are indicated by the keyword `print` followed by the output text enclosed in quotation marks and followed by a semicolon:

```
print "Done initializing registers.\n";
```

The stub will output the characters in between the quotation marks verbatim, with three exceptions: `'\n'`, `'\t'`, and `'\v'` are treated as "newline," "horizontal tab," and "vertical tab," respectively. The `print` command will execute immediately after the command above it is executed. The SSP Stub will not wait for a `print` command to be executed before proceeding with the rest of the diagnostic.

```
delay(cycles);
```

The user can specify a specific delay between sending packets. The value specified in the `delay` command is the number of 5 ns clock cycles to wait before continuing to execute commands in the diagnostic. This can be used to wait for the algorithm to execute steps or finish executing entirely.

```
poll(address, bit, interval);
```

The user can specify a bit to read at a specified interval to see if a process is finished (for example, if the algorithm, DMA read engine, or DMA write engine is complete). The stub will read the MMR of *address* on the specified *interval* until the *bit* is set to 1

Note: NOTE: If the interval chosen is smaller than the time necessary to complete the read transaction (approximately 44 clock cycles), the stub will wait until it receives a response to its last request before proceeding with the next read request (that is, the actual time between reads may be greater than the specified interval).

- *address* = address of MMR to be read
- *bit* = bit in register to be checked
- *interval* = number of clock cycles in between register reads

Command Summary

The table below provides a summary of packet commands and their possible usages.

Command	Packet Type	PIO/ MEM	Size	Tnum	Error	Address	Data
snd_wr_req	PIO 8-Byte Write Request	PIO	DW	User-specified or ANY	-	MMR or SRAM Address	1 64-bit value
snd_rd_req	PIO 8-Byte Read Request	PIO	DW	User-specified or ANY	-	MMR Address	-
	Memory Full 128-Byte Read Request	MEM	FCL	User-specified or ANY	-	SRAM Address	-
snd_wr_rsp	PIO 8-Byte Write Response	PIO	DW	User-specified or ANY	0 or 1	-	-
	Memory Full 128-Byte Write Response	MEM	FCL	User-specified or ANY	0 or 1	-	-
snd_rd_rsp	Memory Full 128-Byte Read Response	MEM	FCL	User-specified or ANY	0 or 1	-	16 64-bit values

inv_flush	Invalidate Cache-Flush Request	-	-	User-specified or ANY	-	-	-
rcv_wr_rsp	PIO 8-Byte Write Response	PIO	DW	User-specified or ANY	0 or 1	-	-
rcv_rd_rsp	PIO 8-Byte Read Response	PIO	DW	User-specified or ANY	0 or 1	-	1 64-bit value
rcv_wr_req	PIO 8-Byte Write Request	PIO	DW	User-specified or ANY	-	Interrupt Destination Address	1 64-bit value
	Memory 128-Byte Write Request	MEM	FCL	User-specified or ANY	-	System Memory Address	16 64-bit values
rcv_rd_req	Memory 128-Byte Read Request	MEM	FCL	User-specified or ANY	-	System Memory Address	-

Comments

Comments are indicated by a '#' sign. The stub will ignore sections of the diagnostic from the '#' indicator to the end of the line (excluding '#'s contained in print statements).

Sample Diagnostic

The code listed below comprises a diagnostic which exercises the basic functionality of the RASC FPGA outlined in the following steps:

1. Initializes the RASC FPGA Core Services (primarily MMR Writes)
2. Executes DMA Reads to send data to the FPGA
3. Starts the Algorithm (A & B | C) and polls the MMR's to see when the Algorithm is done
4. Executes DMA Writes to retrieve the Algorithm's results
5. Checks the error status in the MMR's to verify that no errors were flagged.

```
##### Initialization packets. #####
# Arm registers by setting the REARM_STAT_REGS bit in the CM_CONTROL register.
snd_wr_req ( PIO, DW, ANY, 0x00000000000020, 0x0000000600f00003 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );
```

```
# Clear the CM_ERROR_STATUS register by writing all zeroes.
snd_wr_req ( PIO, DW, 3, 0x000000000000060, 0x0000000000000000 );
rcv_wr_rsp ( PIO, DW, 3, 0 );

# Enable CM_ERROR_DETAIL_* regs by writing all zeroes to CM_ERROR_DETAIL_1.
snd_wr_req ( PIO, DW, ANY, 0x000000000000010, 0x0000000000000000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# Enable desired interrupt notification in the CM_ERROR_INTERRUPT_ENABLE
register.
snd_wr_req ( PIO, DW, 4, 0x000000000000070, 0xFFFFFFFFFFFFFFFF );
rcv_wr_rsp ( PIO, DW, 4, 0 );

# Set up the Interrupt Destination Register.
snd_wr_req ( PIO, DW, ANY, 0x000000000000038, 0x0000000000000000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

print "\n\n*****Initialization finished\n\n";

##### Configure DMA Engines and Algorithm. #####

##### Configure the Read DMA Engine Registers. #####
print "\n\n*****Configure Read DMA Engine. Tell it to fill 32 cache lines of
data.\n\n";

# RD_DMA_CTRL register.
snd_wr_req ( PIO, DW, ANY, 0x000000000000110, 0x0000000000100020 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# RD DMA System Address.
snd_wr_req ( PIO, DW, ANY, 0x000000000000100, 0x0000000000100000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# RD DMA Local Address.
snd_wr_req ( PIO, DW, ANY, 0x000000000000108, 0x0000000000000000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# RD_DMA_DEST_INT
snd_wr_req ( PIO, DW, ANY, 0x000000000000120, 0x0000000200002000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

##### Configure the Write DMA Engine Registers. #####
print "\n\n*****Configure Write DMA Engine.\n\n";

# Write to the WR_DMA_CTRL register.
snd_wr_req ( PIO, DW, ANY, 0x000000000000210, 0x0000000000100020 );
```

```

rcv_wr_rsp ( PIO, DW, ANY, 0 );

# WR_DMA_SYS_ADDR
snd_wr_req ( PIO, DW, ANY, 0x00000000000200, 0x000000000100000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# WR_DMA_LOC_ADDR
snd_wr_req ( PIO, DW, ANY, 0x00000000000208, 0x000000000000000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# WR_DMA_INT_DEST
snd_wr_req ( PIO, DW, ANY, 0x00000000000220, 0x0000000400004000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

##### Configure the Algorithm Registers. #####
print "\n\n*****Configure Algorithm Registers\n\n";

snd_wr_req ( PIO, DW, ANY, 0x00000000000300, 0x0000000000000000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

snd_wr_req ( PIO, DW, ANY, 0x00000000000308, 0x0000000600006000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

##### Start Read DMA Engine for Read DMA 1 #####
print "\n\n*****Start Read DMA Engine for SRAM0\n\n";

# Set Bit 36 of the CM_CONTROL Reg to 1.
snd_wr_req ( PIO, DW, ANY, 0x0000000000020, 0x0000001400f00003 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# 1 of 32
rcv_rd_req ( MEM, FCL, ANY, 0x00000000100000 );
snd_rd_rsp ( MEM, FCL, ANY, 0, 0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF );

# Other Read DMA Transactions omitted here

# 32 of 32
rcv_rd_req ( MEM, FCL, ANY, 0x00000000100F80 );
snd_rd_rsp ( MEM, FCL, ANY, 0, 0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF, 0xDEADBEEFDEADBEEF,
            0xDEADBEEFDEADBEEF );

```



```

0xF0F0F0F0F0F0F0F0, 0xF0F0F0F0F0F0F0F0,
0xF0F0F0F0F0F0F0F0, 0xF0F0F0F0F0F0F0F0,
0xF0F0F0F0F0F0F0F0, 0xF0F0F0F0F0F0F0F0,
0xF0F0F0F0F0F0F0F0, 0xF0F0F0F0F0F0F0F0 );

print "\n\n*****Done storing 32 cache lines of data in SRAM1.\n\n";

print "\n\n*****Polling for DMA RD-SRAM1 done (bit 42 of CM_STATUS).\n\n";
poll (0x8, 42, 200);

##### Reconfigure DMA Engine for Read DMA 3 #####

# RD DMA addresses.
snd_wr_req ( PIO, DW, ANY, 0x00000000000100, 0x0000000000100000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );
snd_wr_req ( PIO, DW, ANY, 0x00000000000108, 0x0000000000400000 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

##### Start Read DMA Engine for Read DMA 3 #####
print "\n\n*****Start Read DMA Engine for SRAM2\n\n";

# Set Bit 36 of the CM_CONTROL Reg to 1.
snd_wr_req ( PIO, DW, ANY, 0x00000000000020, 0x0000001400f00003 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# 1 of 32
rcv_rd_req ( MEM, FCL, ANY, 0x00000000100000 );
snd_rd_rsp ( MEM, FCL, ANY, 0, 0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C );

# Other Read DMA Transactions omitted here

# 32 of 32
rcv_rd_req ( MEM, FCL, ANY, 0x00000000100F80 );
snd_rd_rsp ( MEM, FCL, ANY, 0, 0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C,
0x0C0C0C0C0C0C0C0C, 0x0C0C0C0C0C0C0C0C );

```

```
print "\n\n*****Done storing 32 cache lines of data in SRAM 3.\n\n";

print "\n\n*****Polling for DMA RD-SRAM2 done (bit 42 of CM_STATUS).\n\n";
poll (0x8, 42, 200);

##### Start the Algorithm #####

# Set bit 38 of CM Control Register to 1 to start algorithm.
snd_wr_req ( PIO, DW, ANY, 0x00000000000020, 0x0000004400f00003 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

print "\n\n*****Started Algorithm.\n\n";

# Poll for ALG_DONE bit in CM_STATUS.
poll ( 0x8, 48, 200 );

##### Start Write DMA Engine. #####

# Set bit 37 of CM Control Register to 1 to start Write DMA Engine.
snd_wr_req ( PIO, DW, ANY, 0x00000000000020, 0x0000002400f00003 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

print "\n\n*****Started Write DMA Engine.\n\n";

# 1 of 32
rcv_wr_req ( MEM, FCL, ANY, 0x00000000100000, 0xDCACBCECDCACBCEC,
0xDCACBCECDCACBCEC,
    0xDCACBCECDCACBCEC, 0xDCACBCECDCACBCEC,
    0xDCACBCECDCACBCEC, 0xDCACBCECDCACBCEC,
    0xDCACBCECDCACBCEC, 0xDCACBCECDCACBCEC,
    0xDCACBCECDCACBCEC, 0xDCACBCECDCACBCEC,
    0xDCACBCECDCACBCEC, 0xDCACBCECDCACBCEC );
rcv_wr_rsp ( MEM, FCL, ANY, 0 );

# Other Write DMA Transactions omitted here

# 32 of 32
rcv_wr_req ( MEM, FCL, ANY, 0x00000000100F80, 0xDCACBCECDCACBCEC,
0xDCACBCECDCACBCEC,
    0xDCACBCECDCACBCEC, 0xDCACBCECDCACBCEC,
    0xDCACBCECDCACBCEC, 0xDCACBCECDCACBCEC,
    0xDCACBCECDCACBCEC, 0xDCACBCECDCACBCEC,
    0xDCACBCECDCACBCEC, 0xDCACBCECDCACBCEC,
    0xDCACBCECDCACBCEC, 0xDCACBCECDCACBCEC,
    0xDCACBCECDCACBCEC, 0xDCACBCECDCACBCEC );
```



```

snd_wr_rsp ( MEM, FCL, ANY, 0 );

print "\n\n*****Polling for DMA WR-SRAM0 done (bit 45 of CM_STATUS).\n\n";
poll (0x8, 45, 200);

##### Finish Up #####

# dma_clear(). Set bits 39, 40, and 41 to 1 in CM_CONTROL.
snd_wr_req ( PIO, DW, ANY, 0x00000000000020, 0x0000038400f00003 );
rcv_wr_rsp ( PIO, DW, ANY, 0 );

# finalcheck_ccc() Check CACHE_RD_DMA_FSM.
snd_rd_req ( PIO, DW, ANY, 0x000000000000130 );
rcv_rd_rsp ( PIO, DW, ANY, 0, 0x0000000000040000 );

print "Reading the Error Status Register to insure no errors were logged.\n";
snd_rd_req ( PIO, DW, ANY, 0x000000000000060 );
rcv_rd_rsp ( PIO, DW, ANY, 0, 0x0000000000000000 );

```

Using the Stub

Because the SSP Stub is a very simple verification tool intended only for sandbox testing of new algorithm designs, capabilities and support for this stub are minimal. The main intent is to provide the user with a means to simulate loading data, running their algorithm, and reading the results. More comprehensive testing should be done on real hardware. Due to limitations of the stub, below are some guidelines/rules for writing a diagnostic which will help the user quickly create functional tests.

- Commands in the stub **MUST** alternate between request and responses. This allows the stub to associate transaction pairs and calculate transaction numbers accordingly.
- Comments operate as follows: when the stub encounters a '#' character, it ignores all text from that character until it reaches the end of that line. Print statements are an exception to this rule; if the stub encounters a '#' in between the quotation marks of a print statement, it does **NOT** recognize this as a comment.
- A command should **ALWAYS** be followed by a semicolon.
- The parser checks for the correct number of inputs to a function, but it does **NOT** check the validity of these inputs. Mislaced semicolons may also lead to unpredictable behavior.

- Because the input file is in text format, no mathematical operations, functions, or variables are available. The diagnostic writer must know the full data value to be written to a register.
- Diagnostic files must **ALWAYS** end with a newline character.
- Always follow a long sequence of DMA transactions with a `poll()` command of the `WR_DMA_DONE` bit or `RD_DMA_DONE` bit. Enabling Interrupts will help test for the completion of a long set of transactions.
- Other typical errors include the following:
 - diagnostic is too long. If you are running into memory-related errors, try increasing the `FILE_MAX` or `COMMAND_MAX` constants in `user_const.h` (number of characters allowed in diagnostic file).
 - Timed out while waiting for an operation to complete. If you are timing out before you receive your algorithm or interrupt, increase `TIME_OUT` in `user_const.h`. If you are timing out during a poll, increase `POLL_MAX` in `user_const.h`.
- Read the Error Status Register at the end of a diagnostic to insure that no unreported errors were logged during the diagnostic.

How Extractor Works

This section describes how metadata is extracted from source files and how to interpret the generated configuration files. It covers the following topics:

- “Extractor Script” on page 245
- “Core Services Configuration File” on page 246
- “Algorithm Configuration File” on page 251

Extractor Script

The extractor script extracts necessary configuration information from a user’s register transfer level (RTL) code.

To use the extractor script, enter the following:

```
python $RASC/pd/shrd/extractor $RASC/design/top/acs_cm_id.h $ALG_DIR
```

where `$RASC` is the installation path to the top of the RASC Core Services tree, and `$ALG_DIR` is the top of the algorithm design file tree. These variables can be set in the physical design makefile. Please see the physical design specifications for further details.

When called from a directory, extractor creates two configuration files: `core_services.cfg` and `user_space.cfg`.

The `core_services.cfg` file is created from the file that is passed in as the first command line argument. In the file that is passed as the first argument, extractor searches the header for the glob that begins with `'core_services_version'` and ends with `'end:'`.

The `user_space.cfg` file is created from the files in the directory that is passed as the second argument to extractor. The search includes all `*.v`, `*.h`, or `*.vhd` files in and below the directory that is passed as the second argument. In those files, extractor searches for any comments that start with: extractor. Then it searches for the particular

tag, either CS, VERSION, REG_IN, REG_OUT, or SRAM. The line following these tags is then stored, manipulated in some cases, and written out to the `user_space.cfg` configuration file.

Both configuration files are written to the directory in which the script is called.

Please note that the system extractor is run upon requires a Python 2 interpreter. You can obtain the interpreter at no cost from www.python.org.

Core Services Configuration File

The format for entries in the core services configuration file is:

`<type>:<value>`

The various fields for the configuration file are described in detail in Table C-1.

Table C-1 Core Services Configuration File Fields

Type	Description	Value	Example
core_services_version	Tag for the version and revision of core services.	<version>.<revision>	core_services_version:0.7
part_num	Tag for the part number that corresponds to the CM_ID register in the hardware.	<hexadecimal part number>	part_num:0xF001
mfg_id	Tag for the manufacturer's identification number that corresponds to the CM_ID register in the hardware.	<hexadecimal manufacturer's ID>	mfg_id: 0x1
dma	Tag for the number and direction of the DMA engines.	<number> <direction[0] ... direction[n-1]>.	For two read and one write DMA engine, dma:3 rd rd wr
stream_in	Tag for the number and direction of the DMA engines.	<stream_name> <identifier0-3>	stream_in: 0 0x001000000
stream_out	Tag for the number and direction of the DMA engines.	<stream_name> <identifier0-3>	stream_out: 0 0x001800000

Table C-1 Core Services Configuration File Fields (**continued**)

Type	Description	Value	Example
sram[n]	Tag for the SRAMs on the board The address field is from the DMA perspective; it is not the PIO space for that SRAM.	<DMA address for sram 'n'> <size of sram 'n' in MB>	sram[0]:0x00000000000000 2MB
amo	Tag for whether or not this core services is AMO capable	1 is yes, 0 is no	amo:1
interrupt	Tag for whether or not this core services is interrupt capable	1 is yes, 0 is no	interrupt:1
step	Tag for whether or not this core services is algorithm step capable.	1 is yes, 0 is no	step:1
mmr	Tag for the MMR space.	<base register for PIO> <length of space> <alignment>	mmr:0x00000000000000 2KB 64
debug	Tag for the Debug space.	<base register for PIO> <length of space> <alignment>	debug:0x00000000100000 1MB 64
debug_port[n]	Tag for the PIO address of the debug register that is associated with debug port 'n'.	<PIO address for the register> <bit width of the register>	debug_port[0]:0x0000 64
alg_def_reg[n]	Tag for the PIO address for a register 'n' that the RASC application wants to write to the algorithm	<PIO address for the register> <bit width of the register>	alg_def_reg[0]:0x0320 64
end	Tag for the end of the file	It has no value	end:

The core services configuration file should have the name `core_services.cfg`. This file does not use the ability for the application to write an internal register, so the file should appear as:

```
part_num:0xF005
mfg_id:0x1
sram[0]: 0x00000200000000 16MB
sram[1]: 0x00000300000000 16MB
sram[2]: 0x00000400000000 8MB
stream_in:0 0x001000000
stream_in:1 0x001100000
stream_in:2 0x001200000
stream_in:3 0x001300000
stream_out:0 0x001800000
stream_out:1 0x001900000
stream_out:2 0x001a00000
stream_out:3 0x001b00000
amo:1
interrupt:1
step:1
mmr:0x00000000000000 2KB 64
debug:0x00000000100000 1MB 64
debug_port[0]: 0x0000000001000 64
debug_port[1]: 0x0000000001008 64
debug_port[2]: 0x0000000001010 64
debug_port[3]: 0x0000000001018 64
debug_port[4]: 0x0000000001020 64
debug_port[5]: 0x0000000001028 64
debug_port[6]: 0x0000000001030 64
debug_port[7]: 0x0000000001038 64
debug_port[8]: 0x0000000001040 64
debug_port[9]: 0x0000000001048 64
debug_port[10]: 0x0000000001050 64
debug_port[11]: 0x0000000001058 64
debug_port[12]: 0x0000000001060 64
debug_port[13]: 0x0000000001068 64
debug_port[14]: 0x0000000001070 64
debug_port[15]: 0x0000000001078 64
debug_port[16]: 0x0000000001080 64
debug_port[17]: 0x0000000001088 64
debug_port[18]: 0x0000000001090 64
debug_port[19]: 0x0000000001098 64
debug_port[20]: 0x00000000010A0 64
debug_port[21]: 0x00000000010A8 64
debug_port[22]: 0x00000000010B0 64
debug_port[23]: 0x00000000010B8 64
debug_port[24]: 0x00000000010C0 64
debug_port[25]: 0x00000000010C8 64
debug_port[26]: 0x00000000010D0 64
```

```
debug_port[27]: 0x000000000010D8 64
debug_port[28]: 0x000000000010E0 64
debug_port[29]: 0x000000000010E8 64
debug_port[30]: 0x000000000010F0 64
debug_port[31]: 0x000000000010F8 64
debug_port[32]: 0x00000000001100 64
debug_port[33]: 0x00000000001108 64
debug_port[34]: 0x00000000001110 64
debug_port[35]: 0x00000000001118 64
debug_port[36]: 0x00000000001120 64
debug_port[37]: 0x00000000001128 64
debug_port[38]: 0x00000000001130 64
debug_port[39]: 0x00000000001138 64
debug_port[40]: 0x00000000001140 64
debug_port[41]: 0x00000000001148 64
debug_port[42]: 0x00000000001150 64
debug_port[43]: 0x00000000001158 64
debug_port[44]: 0x00000000001160 64
debug_port[45]: 0x00000000001168 64
debug_port[46]: 0x00000000001170 64
debug_port[47]: 0x00000000001178 64
debug_port[48]: 0x00000000001180 64
debug_port[49]: 0x00000000001188 64
debug_port[50]: 0x00000000001190 64
debug_port[51]: 0x00000000001198 64
debug_port[52]: 0x000000000011A0 64
debug_port[53]: 0x000000000011A8 64
debug_port[54]: 0x000000000011B0 64
debug_port[55]: 0x000000000011B8 64
debug_port[56]: 0x000000000011C0 64
debug_port[57]: 0x000000000011C8 64
debug_port[58]: 0x000000000011D0 64
debug_port[59]: 0x000000000011D8 64
debug_port[60]: 0x000000000011E0 64
debug_port[61]: 0x000000000011E8 64
debug_port[62]: 0x000000000011F0 64
debug_port[63]: 0x000000000011F8 64
alg_def_reg[0]: 0x00000000000000 64
alg_def_reg[1]: 0x00000000000008 64
alg_def_reg[2]: 0x00000000000010 64
alg_def_reg[3]: 0x00000000000018 64
alg_def_reg[4]: 0x00000000000020 64
alg_def_reg[5]: 0x00000000000028 64
alg_def_reg[6]: 0x00000000000030 64
alg_def_reg[7]: 0x00000000000038 64
```

```
alg_def_reg[8]: 0x00000000000040 64
alg_def_reg[9]: 0x00000000000048 64
alg_def_reg[10]: 0x00000000000050 64
alg_def_reg[11]: 0x00000000000058 64
alg_def_reg[12]: 0x00000000000060 64
alg_def_reg[13]: 0x00000000000068 64
alg_def_reg[14]: 0x00000000000070 64
alg_def_reg[15]: 0x00000000000078 64
alg_def_reg[16]: 0x00000000000080 64
alg_def_reg[17]: 0x00000000000088 64
alg_def_reg[18]: 0x00000000000090 64
alg_def_reg[19]: 0x00000000000098 64
alg_def_reg[20]: 0x000000000000a0 64
alg_def_reg[21]: 0x000000000000a8 64
alg_def_reg[22]: 0x000000000000b0 64
alg_def_reg[23]: 0x000000000000b8 64
alg_def_reg[24]: 0x000000000000c0 64
alg_def_reg[25]: 0x000000000000c8 64
alg_def_reg[26]: 0x000000000000d0 64
alg_def_reg[27]: 0x000000000000d8 64
alg_def_reg[28]: 0x000000000000e0 64
alg_def_reg[29]: 0x000000000000e8 64
alg_def_reg[30]: 0x000000000000f0 64
alg_def_reg[31]: 0x000000000000f8 64
alg_def_reg[32]: 0x00000000000100 64
alg_def_reg[33]: 0x00000000000108 64
alg_def_reg[34]: 0x00000000000110 64
alg_def_reg[35]: 0x00000000000118 64
alg_def_reg[36]: 0x00000000000120 64
alg_def_reg[37]: 0x00000000000128 64
alg_def_reg[38]: 0x00000000000130 64
alg_def_reg[39]: 0x00000000000138 64
alg_def_reg[40]: 0x00000000000140 64
alg_def_reg[41]: 0x00000000000148 64
alg_def_reg[42]: 0x00000000000150 64
alg_def_reg[43]: 0x00000000000158 64
alg_def_reg[44]: 0x00000000000160 64
alg_def_reg[45]: 0x00000000000168 64
alg_def_reg[46]: 0x00000000000170 64
alg_def_reg[47]: 0x00000000000178 64
alg_def_reg[48]: 0x00000000000180 64
alg_def_reg[49]: 0x00000000000188 64
alg_def_reg[50]: 0x00000000000190 64
alg_def_reg[51]: 0x00000000000198 64
alg_def_reg[52]: 0x000000000001a0 64
```



```

alg_def_reg[53]: 0x0000000000001a8 64
alg_def_reg[54]: 0x0000000000001b0 64
alg_def_reg[55]: 0x0000000000001b8 64
alg_def_reg[56]: 0x0000000000001c0 64
alg_def_reg[57]: 0x0000000000001c8 64
alg_def_reg[58]: 0x0000000000001d0 64
alg_def_reg[59]: 0x0000000000001d8 64
alg_def_reg[60]: 0x0000000000001e0 64
alg_def_reg[61]: 0x0000000000001e8 64
alg_def_reg[62]: 0x0000000000001f0 64
alg_def_reg[63]: 0x0000000000001f8 64
end:

```

Algorithm Configuration File

The format for entries in the algorithm configuration file is:

```
<type>:<value>
```

The various fields for the configuration file are described in detail in Table C-2..

Table C-2 Algorithm Configuration File Fields

Type	Description	Value	Example
core_services_version	Tag detailing which iteration of core services was used in this bitstream.	<version> <revision>	core_services_version:0.7
algorithm_version	Tag for the type and version number of the algorithm in this bitstream. This is user defined but it should match the first register in the debug space (see Hardware Reference for more details.)	<version> <revision>	algorithm_version:0.12
src	Tag for source type that generated the algorithm.	Verilog by default and VHDL if ANY .vhd file exists in the hierarchy	src:v

Table C-2 Algorithm Configuration File Fields (**continued**)

Type	Description	Value	Example
reg	Tag for the registers that are PIO written or read by the application or the debugger and where the value will be mapped in the address space.	<register name> <bit width of register> <data type of the element mapped> <port mapping to the application or the debug relevant debug register>	reg:a 64 u debug_port[1]
array	Tag for the SRAM inputs and outputs to the algorithm.	<array name> <number of elements in array> <width of an element in bits> <associated SRAM bank for array> <byte offset into the SRAM bank> <direction of array: in, out, or inout> <data type for the elements in array> <whether the array is streaming or fixed>	array:a_in 512 64 sram[0] 0x0000 in u stream
stream_in	Tag for streaming input to the algorithm	<stream name> <stream number> <stride>	stream_in:strm_a 0 0
stream_out	Tag for streaming output to the algorithm	<stream name> <stream number> <stride>	stream_out:strm_b 0 0
end	Tag for the end of the file.	It has no value field.	end:

The name of the algorithm configuration file that is generated by the extractor script will be `user_space.cfg`. This name can be changed after the script is run at the discretion of the user. The algorithm configuration file should appear as follows:

```
core_services_version:0.7
```

```
algorithm_version:0.6
src:v
reg:a_0 64 u debug_port[1]
reg:b_0 64 u debug_port[2]
reg:c_0 64 u debug_port[3]
reg:tmp 64 u debug_port[4]
reg:d_0 64 u debug_port[5]
reg:i_0 64 u debug_port[6]
array:a 512 64 sram[0] 0x00 in u stream
array:b 512 64 sram[1] 0x00 in u stream
array:c 512 64 sram[2] 0x00 in u stream
array:d 512 64 sram[0] 0x00 out u stream
end:
```

Index

A

- Adding extractor directives to source code, 144
- Algorithm Block
 - debug mode, 25
- Algorithm block
 - passing parameters, 52
- Algorithm Block modes
 - normal mode, 25
- Algorithm control interface, 26
- Algorithm debug mode, 46
- Algorithm design details, 42
- algorithm inputs and outputs, 45
- Algorithm interfaces, 26
- Algorithm run modes, 25
- Algorithm streaming
 - algorithm iteration, 49
 - segment/segment size, 50
- Algorithms, diagnostics , and commands, 70

B

- Basic algorithm control, 42
- Bitstream development overview, 6

C

- ccNUMA systems, 3
- Clock cycle based stepping, 46

- Coding guidelines for timing, 54
- Core Services architecture overview, 23
- Core Services Block, 22
- Core services, FPGA, 213

D

- Designing an algorithm for multibuffering, 49
- Determining FPGA run status, 182
- Determining if an FPGA has run, 182
- Device driver, 213
 - control and status registers, 213
- Device driver API, 214
- Device driver DMA, 214
- Device driver example, 215
- Device driver, function control, 215
- Device driver, interrupts, 213
- Device manager, 168
 - overview, 169
 - structure, 170
 - using, 170
- Device manager load command, 174
- Driver API, 213

F

- FPGA clock domains, 57, 60
- FPGA core services, 213

FPGA design integration, 55
FPGA device values and stepping, 182
FPGA programming, 5
FPGA programming approach summary, 6
FPGA registers, 182
FPGA run status, 182

G

GDB commands, 179
GNU Debugger
 connecting to internal signals, 54
GNU Debugger (GDB), 178
 overview, 9

H

Hardware resets, 60

I

Internal timing requirements, 54

L

Loading the bitstream, 167

M

Manually loading an FPGA, 174
Memory distribution recommendations, 45

P

Passing parameters to the algorithm block, 52

R

RASC Abstraction Layer, 7, 81
 functions, 83
 how it works, 125
RASC Algorithm Field Programmable Gate Array
 (FPGA) hardware, 21
RASC Algorithm FPGA
 implementation guide, 137
RASC Algorithm FPGA implementation
 flow, 138
 Full-chip implementation, 152
 installation and setup, 140
 Makefile targets, 151
 Makefile.local customizations, 149
 overview, 138
 pre-compiled cores, 148
 supported tools, 139
RASC hardware overview, 9
 reconfigurable algorithm, 9
RASC overview, 4
RASC Programming, getting started, 5
RASC software overview, 12
RASC tutorial
 data flow algorithm, 200, 206
 overview, 186
 simple algorithm, 187
 system requirements, 185
Reconfigurable computing, 1
Run status of an FPGA, 182
Running a diagnostic, 66

S

- Sample test bench constants and definitions, 76
- Sample test bench setup, 66
- Scalable System Port (SSP) Stub, 223
 - compiling and running a test, 227
 - packet commands summary, 236
 - sample testbench, 224
 - stub commands, 229
 - verification environment, 224
- Simulating the design, 66
- SRAM
 - external memory read transaction control, 48
- SRAM external memory write transaction control, 47
- Streaming direct memory access (DMA), 29
- Streaming DMA, 29

U

- Using the device manager, 170
- Using the device manager command, 177
- Using the GNU Project debugger, 178

V

- Variable step size mode, 46
- Verilog, VHDL or header file
 - adding comments, 145
- VHDL and Verilog programming languages, 5

W

- Writing a diagnostic, 71

