



Uniform Driver Interface

UDI Physical I/O Specification Version 1.0



UDI Physical I/O Specification

Abstract

The UDI Physical I/O Specification defines the required interfaces and semantics for UDI environments that support Programmed I/O (PIO), Direct Memory Access (DMA), and Interrupts. This is an optional extension to the UDI Core Specification, which is defined in a separate book. The intended audience for this book includes driver writers, environment implementors, and metalanguage implementors.

UDI drivers that require physical I/O services must be written to conform to this specification, and can assume that all services described herein are available. Environments that don't need such drivers may choose not to support the physical I/O extensions, but any environment that supports UDI physical I/O drivers must conform to this specification, as well as to the UDI Core Specification.

Environments that support UDI drivers for devices that use any physical I/O bus types covered by other UDI specifications must support and conform to those Bus Binding specifications.

See the Document Organization chapter in the UDI Core Specification for a description of other books in the UDI Specification collection, as well as references to additional tutorial materials.

Status of This Document

This document has been reviewed by Project UDI Members and other interested parties and has been endorsed as a Final Specification. It is a stable document and may be used as reference material or cited as a normative reference from another document. This version of the specification is intended to be ready for use in product design and implementation. Every attempt has been made to ensure a consistent and implementable specification. Implementations should ensure compliance with this version.

Copyright Notice

Copyright © 1999 Adaptec, Inc; Compaq Computer Corporation; Hewlett-Packard Company; International Business Machines Corporation; Interphase Corporation; Lockheed Martin Corporation; The Santa Cruz Operation, Inc; SBS Technologies, Inc; Sun Microsystems (“copyright holders”). All Rights Reserved.

This document and other documents on the Project UDI web site (www.project-UDI.org) are provided by the copyright holders under the following license. By obtaining, using and/or copying this document, or the Project UDI document from which this statement is linked, you agree that you have read, understood, and will comply with the following terms and conditions:

Permission to use, copy, and distribute the contents of this document, or the Project UDI document from which this statement is linked, in any medium for any purpose and without fee or royalty is hereby granted, provided that you include all of the following on *ALL* copies of the document, or portions thereof, that you use:

1. A link or URI to the original Project UDI document.
2. The pre-existing copyright notice of the original author, or, if it doesn't exist, a Project UDI copyright notice of the form shown above.
3. *If it exists*, the STATUS of the Project UDI document.

When space permits, inclusion of the full text of this **NOTICE** should be provided. In addition, credit shall be attributed to the copyright holders for any software, documents, or other items or products that you create pursuant to the implementation of the contents of this document, or any portion thereof.

No right to create modifications or derivatives is granted pursuant to this license.

THIS DOCUMENT IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, NON-INFRINGEMENT, OR TITLE; THAT THE CONTENTS OF THE DOCUMENT ARE SUITABLE FOR ANY PURPOSE; NOR THAT THE IMPLEMENTATION OF SUCH CONTENTS WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

COPYRIGHT HOLDERS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE DOCUMENT OR THE PERFORMANCE OR IMPLEMENTATION OF THE CONTENTS THEREOF.

The names and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to this document or its contents without specific, written prior permission. Title to copyright in this document will at all times remain with copyright holders.

Preface

Acknowledgements

The authors would like to thank everyone who reviewed working drafts of the specification and submitted suggestions and corrections.

The authors would especially like to thank their significant others for putting up with the many hours of overtime put into the development of this specification over long periods.

Thanks to the following folks who contributed significant amounts of time, ideas, or authoring in support of the development of this specification:

Richard Arndt (IBM), Bob Barned (Lockheed Martin), Mark Bradley (Adaptec), Darren Busing (Adaptec), Thomas Clark (Sun), Mark Evenson (HP), Barry Feild (SCO), Kurt Gollhardt (SCO), Jim Heidbrink (Lockheed Martin), Chris Ilnicki (HP), Bret Indrelee (SBS Technologies), David Kahn (Sun), John Lee (Sun), Robert Lipe (SCO), Lynne McCue (IBM), Mark Parenti (DEC), James Partridge (IBM), Kevin Quick (Interphase), Larry Robinson (Adaptec), James Smart (Compaq), Pete Smoot (HP), David Stoff (HP), Wolfgang Thaler (Sun), Linda Wang (Sun), Mike Wenzel (HP).

In addition, thanks to everyone who worked on the prototype implementations that helped us validate the specification, including:

Betty Dall (HP), Don Dugger (Intel), Bob Goudreau (Data General), James Hall (SCO), Matt Kaufman (SCO), Mike Lyons (IBM), Alex Malone (DEC), Guru Pangal (Starcom), Scott Popp (SCO), Richard Schall (Intel), Sam Shteingart (HP), Ajmer Singh (SCO), Ramaswamy Tummala (Starcom).

Finally, thanks to David Roberts (Certek Software Designs) for designing the Project UDI logo.



Table of Contents

Abstract	i
Copyright Notice	ii
Acknowledgements.....	iii
Table of Contents.....	v
List of Reference Pages by Chapter.....	vii
Alphabetical List of Symbols.....	ix

Section 1: Physical I/O Services

1 Introduction to Physical I/O	1-1
1.1 Overview	1-1
1.2 General Requirements	1-1
1.2.1 Versioning	1-1
1.2.2 Header Files	1-2
1.3 Normative References	1-2
1.4 Extensions to Static Driver Properties	1-2
1.5 Bus Bindings	1-3
2 Direct Memory Access (DMA).....	2-1
2.1 Overview	2-1
2.2 DMA Block Vector Segment	2-1
2.3 DMA Service Calls and Structures	2-5
3 Programmed I/O (PIO)	3-1
3.1 Overview	3-1
3.2 PIO Handle Allocation and Initialization	3-2
3.3 PIO Access Service Calls and Structures	3-11
3.4 PIO Handle Transferability	3-28

Section 2: Bus Bridge Metalanguage

4	Bus Bridge Metalanguage.....	4-1
4.1	Overview	4-1
4.1.1	Versioning	4-1
4.2	Binding/Unbinding Operations	4-2
4.3	Interrupt Registration Operations	4-10
4.4	Interrupt Event Operations	4-19
4.5	Instance Attribute Bindings	4-27
4.6	Bus Bridge Trace Events	4-28
4.7	Bus Bridge Metalanguage States	4-28
4.8	Bus Bridge Status Codes	4-28

Section 3: Bus Bindings

5	Introduction to Bus Bindings	5-1
5.1	Normative References	5-1
5.2	Header Files	5-1
5.3	PIO Bindings	5-1
5.3.1	udi_pio_map	5-1
5.4	Interrupt Bindings	5-1
5.4.1	Interrupt Index Values	5-1
5.4.2	Event Info	5-2
5.5	Instance Attribute Bindings	5-2
5.5.1	Enumeration Attributes	5-2
5.5.2	Filter Attributes	5-2
5.5.3	Generic Attributes	5-2
5.5.4	Parent-Visible Attributes	5-2

Section 4: Appendices

A	Glossary.....	A-1
	Index.....	X-1



List of Reference Pages by Chapter

Chapter 2 Direct Memory Access (DMA)

udi_dma_handle_t	DMA handle type	2-6
udi_dma_limits	Platform-specific allocation and access limits	2-7
udi_busaddr64_t	64-bit bus address data type	2-9
udi_scgth_t	I/O Bus scatter/gather structure	2-10
udi_constraints_attr_t (DMA)	DMA constraints attributes	2-14
udi_dma_prepare	Prepare for DMA mapping	2-24
udi_dma_buf_map	Map a buffer for DMA	2-25
udi_dma_buf_unmap	Release a buffer's DMA mapping	2-28
udi_dma_mem_alloc	Allocate shared control structure memory	2-29
udi_dma_sync	Sync host & device views of DMA-able memory	2-32
udi_dma_scgth_sync	Sync host & device views of scatter/gather list	2-34
udi_dma_mem_barrier	Ordering barrier for accesses to DMA-able memory	2-35
udi_dma_free	Free DMA resources	2-36
udi_dma_mem_to_buf	Convert DMA-mapped control memory into a buffer	2-37

Chapter 3 Programmed I/O (PIO)

udi_pio_handle_t	PIO handle type	3-3
udi_pio_map	Map device memory/registers for access	3-4
udi_pio_unmap	Unmap a PIO handle and free associated resources	3-8
udi_pio_atomic_sizes	Retrieve supported PIO operation atomicity	3-9
udi_pio_abort_sequence	Register a PIO abort sequence	3-10
udi_pio_trans_t	PIO transaction descriptor	3-13
UDI_PIO_REP_ARGS	Parameters for repeated PIO transactions	3-22
udi_pio_trans	Generate PIO transactions	3-24
udi_pio_probe	Probe a PIO device that might not be present	3-26
udi_layout_t (PIO)	Data layout specifier for PIO	3-29

Chapter 4 Bus Bridge Metalanguage

udi_bus_device_ops_t	Device driver entry point ops vector	4-3
udi_bus_bridge_ops_t	Bridge driver entry point ops vector	4-4
udi_bus_bind_cb_t	Control block for bus bridge binding operations	4-5
udi_bus_bind_req	Request a binding to a bridge driver	4-6
udi_bus_bind_ack	Acknowledge a bus bridge binding	4-7
udi_bus_unbind_req	Request a bridge driver unbinding (child to bridge)	4-8
udi_bus_unbind_ack	Acknowledge a bus bridge unbinding	4-9
udi_intr_attach_cb_t	Control block for interrupt registration operations	4-11
udi_intr_attach_req	Request an interrupt attachment	4-13

List of Reference Pages by Chapter

udi_intr_attach_ack - - - - -	<i>Acknowledge an interrupt attachment</i>	4-15
udi_intr_attach_ack_unused - - - - -	<i>Proxy for udi_intr_attach_ack</i>	4-15
udi_intr_detach_cb_t - - - - -	<i>Control block for interrupt detachment operations</i> ...	4-16
udi_intr_detach_req - - - - -	<i>Request an interrupt detachment</i>	4-17
udi_intr_detach_ack - - - - -	<i>Acknowledge an interrupt detachment</i>	4-18
udi_intr_detach_ack_unused - - - - -	<i>Proxy for udi_intr_detach_ack</i>	4-18
udi_intr_handler_ops_t - - - - -	<i>Interrupt handler ops vector</i>	4-20
udi_intr_dispatcher_ops_t - - - - -	<i>Interrupt dispatcher ops vector</i>	4-21
udi_intr_event_cb_t - - - - -	<i>Control block for interrupt event ops</i>	4-22
udi_intr_event_ind - - - - -	<i>Interrupt event indication</i>	4-24
udi_intr_event_res - - - - -	<i>Acknowledge an interrupt event</i>	4-26



Alphabetical List of Symbols

udi_bus_bind_ack	4-7
UDI_BUS_BIND_CB_NUM	4-5
udi_bus_bind_cb_t	4-5
udi_bus_bind_req	4-6
UDI_BUS_BRIDGE_OPS_NUM	4-4
udi_bus_bridge_ops_t	4-4
UDI_BUS_DEVICE_OPS_NUM	4-3
udi_bus_device_ops_t	4-3
UDI_BUS_INTR_ATTACH_CB_NUM	4-11
UDI_BUS_INTR_DETACH_CB_NUM	4-16
UDI_BUS_INTR_DISPATCH_OPS_NUM	4-21
UDI_BUS_INTR_EVENT_CB_NUM	4-22
UDI_BUS_INTR_HANDLER_OPS_NUM	4-20
udi_bus_unbind_ack	4-9
udi_bus_unbind_req	4-8
udi_busaddr64_t	2-9
udi_constraints_attr_t (DMA)	2-14
UDI_DL_PIO_HANDLE_T	3-29
UDI_DMA_ADDR_FIXED_BITS	2-14
UDI_DMA_ADDR_FIXED_TYPE	2-14
UDI_DMA_ADDR_FIXED_VALUE_HI	2-14
UDI_DMA_ADDR_FIXED_VALUE_LO	2-14
UDI_DMA_ADDRESSABLE_BITS	2-14
UDI_DMA_ALIGNMENT_BITS	2-14
UDI_DMA_ANY_ENDIAN	4-7
UDI_DMA_BIG_ENDIAN	2-15
udi_dma_buf_map	2-25
udi_dma_buf_unmap	2-28
UDI_DMA_DATA_ADDRESSABLE_BITS	2-14
UDI_DMA_ELEMENT_ALIGNMENT_BITS	2-14
UDI_DMA_ELEMENT_GRANULARITY_BITS	2-14
UDI_DMA_ELEMENT_LENGTH_BITS	2-14
UDI_DMA_FIXED_ELEMENT	2-15
UDI_DMA_FIXED_LIST	2-15
UDI_DMA_FIXED_VALUE	2-15
udi_dma_free	2-36
udi_dma_handle_t	2-6
UDI_DMA_IN	2-24
udi_dma_limits	2-7
UDI_DMA_LITTLE_ENDIAN	2-15
udi_dma_mem_alloc	2-29
udi_dma_mem_barrier	2-35

Alphabetical List of Symbols

udi_dma_mem_to_buf	2-37
UDI_DMA_MIN_ALLOC_LIMIT	2-7
UDI_DMA_NEVERSWAP	2-29
UDI_DMA_NO_PARTIAL	2-14
UDI_DMA_OUT	2-24
udi_dma_prepare	2-24
UDI_DMA_REWIND	2-25
UDI_DMA_SCGTH_ADDRESSABLE_BITS	2-14
UDI_DMA_SCGTH_ALIGNMENT_BITS	2-14
UDI_DMA_SCGTH_BOTH_VISIBLE	2-14
UDI_DMA_SCGTH_DMA_VISIBLE	2-14
UDI_DMA_SCGTH_DRIVER_VISIBLE	2-14
UDI_DMA_SCGTH_ENDIANNESS	2-14
UDI_DMA_SCGTH_MAX_EL_PER_SEG	2-14
UDI_DMA_SCGTH_MAX_ELEMENTS	2-14
UDI_DMA_SCGTH_MAX_SEGMENTS	2-14
UDI_DMA_SCGTH_PREFIX_BYTES	2-14
udi_dma_scgth_sync	2-34
UDI_DMA_SCGTH_VISIBILITY	2-14
UDI_DMA_SEQUENTIAL	2-14
UDI_DMA_SLOP_BARRIER_BITS	2-14
UDI_DMA_SLOP_IN_BITS	2-14
UDI_DMA_SLOP_OUT_BITS	2-14
UDI_DMA_SLOP_OUT_EXTRA	2-14
udi_dma_sync	2-32
udi_intr_attach_ack	4-15
udi_intr_attach_ack_unused	4-15
udi_intr_attach_cb_t	4-11
udi_intr_attach_req	4-13
udi_intr_detach_ack	4-18
udi_intr_detach_ack_unused	4-18
udi_intr_detach_cb_t	4-16
udi_intr_detach_req	4-17
udi_intr_dispatcher_ops_t	4-21
udi_intr_event_cb_t	4-22
udi_intr_event_ind	4-24
udi_intr_event_res	4-26
udi_intr_handler_ops_t	4-20
UDI_INTR_MASKING_NOT_REQUIRED	4-24
UDI_INTR_NO_EVENT	4-22
UDI_INTR_UNCLAIMED	4-22
udi_layout_t (PIO)	3-29
UDI_NULL_DMA_HANDLE	2-6
UDI_NULL_PIO_HANDLE	3-3
UDI_PHYSIO_VERSION	1-1
UDI_PIO_16BYTE	3-13
UDI_PIO_1BYTE	3-13
UDI_PIO_2BYTE	3-13
UDI_PIO_32BYTE	3-13
UDI_PIO_4BYTE	3-13
UDI_PIO_8BYTE	3-13

Alphabetical List of Symbols

udi_pio_abort_sequence	3-10
UDI_PIO_ADD	3-14
UDI_PIO_ADD_IMM	3-14
UDI_PIO_AND	3-13
UDI_PIO_AND_IMM	3-13
udi_pio_atomic_sizes	3-9
UDI_PIO_BARRIER	3-14
UDI_PIO_BIG_ENDIAN	3-4
UDI_PIO_BRANCH	3-14
UDI_PIO_BUF	3-13
UDI_PIO_CSKIP	3-13
UDI_PIO_DELAY	3-14
UDI_PIO_DIRECT	3-13
UDI_PIO_END	3-14
udi_pio_handle_t	3-3
UDI_PIO_IN	3-13
UDI_PIO_IN_IND	3-13
UDI_PIO_LABEL	3-14
UDI_PIO_LITTLE_ENDIAN	3-4
UDI_PIO_LOAD	3-13
UDI_PIO_LOAD_IMM	3-13
UDI_PIO_LOADCACHING_OK	3-4
udi_pio_map	3-4
UDI_PIO_MEM	3-13
UDI_PIO_MERGING_OK	3-4
UDI_PIO_NEVERSWAP	3-4
UDI_PIO_OR	3-14
UDI_PIO_OR_IMM	3-14
UDI_PIO_OUT	3-13
UDI_PIO_OUT_IND	3-13
udi_pio_probe	3-26
UDI_PIO_R0	3-13
UDI_PIO_R1	3-13
UDI_PIO_R2	3-13
UDI_PIO_R3	3-13
UDI_PIO_R4	3-13
UDI_PIO_R5	3-13
UDI_PIO_R6	3-13
UDI_PIO_R7	3-13
UDI_PIO_REP_ARGS	3-22
UDI_PIO_REP_IN_IND	3-14
UDI_PIO_REP_OUT_IND	3-14
UDI_PIO_SCRATCH	3-13
UDI_PIO_SHIFT_LEFT	3-13
UDI_PIO_SHIFT_RIGHT	3-13
UDI_PIO_STORE	3-13
UDI_PIO_STORECACHING_OK	3-4
UDI_PIO_STRICTORDER	3-4
UDI_PIO_SUB	3-14
UDI_PIO_SYNC	3-14
UDI_PIO_SYNC_OUT	3-14

Alphabetical List of Symbols

udi_pio_trans	3-24
udi_pio_trans_t	3-13
UDI_PIO_UNALIGNED	3-4
udi_pio_unmap	3-8
UDI_PIO_UNORDERED_OK	3-4
UDI_PIO_XOR	3-14
UDI_SCGTH_32	2-10
UDI_SCGTH_64	2-10
UDI_SCGTH_DMA_MAPPED	2-10
UDI_SCGTH_DRIVER_MAPPED	2-10
UDI_SCGTH_EXT	2-10
udi_scgth_t	2-10



UDI Physical I/O Specification

Section 1: Physical I/O Services



Introduction to Physical I/O

1

1.1 Overview

The UDI Physical I/O Specification specifies the services and metalanguage interfaces required to support a UDI physical I/O driver. A physical I/O driver is one that controls a directly-attached device (aka a *physical device*), which can be programmed without indirect access through another device. Physical devices typically interface to the host system via an I/O bridge (aka I/O bus) and have registers that are accessed by the host CPU via Programmed I/O (PIO). Physical devices may also generate interrupts to the I/O bridge and/or support direct access to system memory via the I/O bridge (aka DMA).

Section 1 introduces the general requirements, static driver property extensions, and additional UDI services required in order to support physical I/O. Section 2 defines the Bus Bridge Metalanguage, which is used by a physical I/O driver to communicate with its parent bus bridge driver. Section 3 defines the requirements for Bus Binding specifications.

1.2 General Requirements

Certain basic rules apply to all UDI Physical I/O drivers. In order to be UDI-compliant, a driver must follow all of these rules. UDI Physical I/O drivers must also follow the rules specified in the UDI Core Specification. Rules specific to Physical I/O drivers are listed here.

1.2.1 Versioning

All functions and structures defined in the UDI Physical I/O Specification, except for those defined in Chapter 4, “*Bus Bridge Metalanguage*”, are part of the “`udi_physio`” interface, currently at version “0x100”. A driver that conforms to and uses the UDI Physical I/O Specification, Version 1.0, must include the following declaration in its `udiprops.txt` file (see Chapter 31, “*Static Driver Properties*”, of the UDI Core Specification):

```
requires udi_physio 0x100
```

In each UDI physical I/O driver source file, before including any UDI header files, the driver must define the preprocessor symbol, `UDI_PHYSIO_VERSION`, to indicate the version of the UDI Physical I/O Specification to which it conforms, which must be the same as the interface version defined above:

```
#define UDI_PHYSIO_VERSION 0x100
```

As defined in Section 31.4.6, “Requires Declaration,” on page 31-6 of the UDI Core Specification, the two least-significant hexadecimal digits of the interface version represent the minor number; the rest of the hex digits represent the major number. Versions that have the same “major version number” as an earlier version shall be backward compatible with that earlier version (i.e. a strict superset).

1.2.2 Header Files

Each device driver source file must include the file “`udi_physio.h`” after it includes “`udi.h`”, as follows:

```
#include <udi.h>
#include <udi_physio.h>
```

These header files contain environment-specific definitions of standard UDI structures and types, as well as all function prototypes and other definitions needed to use the core and physical I/O UDI interfaces and services. Additional include files may be needed for other non-core services and metalanguages as defined in other UDI Specifications.

To maintain portability across UDI supportive platforms, device driver writers shall not assume any knowledge of the contents of these header files with respect to implementation-dependent aspects of the UDI interfaces (such as the definition of handles or abstract types). Similarly, drivers shall not access any functions or objects external to the driver except those defined in the UDI Specifications to which they conform.

1.3 Normative References

The UDI Physical I/O Specification references the UDI specifications listed below. These specifications contain provisions that, through reference in this document, constitute provisions of the UDI Physical I/O Specification.

1. UDI Core Specification, Version 1.0.

The UDI Physical I/O Specification also references the IEEE 1212.1 standard for Direct Memory Access Devices, but these references are informational and do not constitute provisions of the UDI Physical I/O Specification.

1.4 Extensions to Static Driver Properties

The UDI Physical I/O Specification defines the following extensions to the Static Driver Properties file, `udiprops.txt` (see Chapter 31, “*Static Driver Properties*” in the UDI Core Specification). These extensions are enabled by the presence of the following declaration anywhere in the file:

```
requires udi_physio 0x100
```

Any driver that uses functions or structures from the UDI Physical I/O Specification must include the above declaration in its Static Driver Properties.

The “`region`” declaration is extended to support the following additional region attributes:

Table 1-1 Physical I/O Region Attributes

<code><region_attribute></code>	<code><value></code>	Meaning
<code>type</code>	<code>interrupt</code>	An interrupt region. See <code>udi_intr_attach_req</code> on page 4-13.
<code>pio_probe</code>	<code>no</code>	Regions of this type never use <code>udi_pio_probe</code> . This is the default value for this attribute.

Table 1-1 Physical I/O Region Attributes

<region_attribute>	<value>	Meaning
pio_probe	yes	Regions of this type may use <code>udi_pio_probe</code> . (See <code>udi_pio_probe</code> on page 3-26.)

A “nonsharable_interrupt” declaration is added:

```
nonsharable_interrupt <msgnum> [ <intr_idx> ]
```

At most one “nonsharable_interrupt” declaration must be included for each interrupt source of each supported device. If present, this indicates that the specified interrupt source does not support sharing, so must not be configured on the same interrupt line as another interrupt source. The `<msgnum>` must match `<msgnum>` in a “device” declaration, and indicates that this declaration applies to that device. The `<intr_idx>`, if present, selects one of several interrupt sources for the device, starting from zero; if not present, all interrupt sources for the device are affected. This declaration may indicate that the interrupt source is electrically non-sharable, or it may indicate a logically non-sharable interrupt source (one for which there is no way for the driver to determine if the interrupt has actually been asserted).

1.5 Bus Bindings

The UDI Physical I/O Specification must be used in conjunction with one or more Bus Binding definitions for the type(s) of physical I/O bus on which a driver’s device is supported. A UDI Bus Binding definition provide the usage details that are specific to a particular bus type.

Chapter 5, “Introduction to Bus Bindings”, defines the general requirements on UDI Bus Binding definitions. The actual Bus Binding definitions are found in other specifications, such as the UDI PCI Bus Binding Specification.



Direct Memory Access (DMA)

2

2.1 Overview

Direct Memory Access (DMA) refers to data transfers initiated by a device to access system memory. This is contrasted with Programmed I/O (PIO), which involves transfers initiated by a CPU under driver software control to access registers or memory on a device.

UDI DMA services can be used to set up a UDI buffer or a piece of system memory for a DMA data transfer, to build scatter/gather lists, and to synchronize DMA caches. UDI DMA service calls use an opaque DMA handle to refer to DMA resources managed by the environment. DMA handles are allocated by `udi_dma_prepare` or `udi_dma_mem_alloc` and use the `udi_dma_handle_t` opaque handle type.

In order for a buffer or system memory to be usable for DMA, it must first be mapped for DMA access and a scatter/gather list describing the relevant address ranges must be built. This is accomplished using `udi_dma_buf_map` for UDI buffers or `udi_dma_mem_alloc` for shared control structures in system memory. These functions provide the driver with an IEEE 1212.1 compatible DMA scatter/gather list (see “DMA Block Vector Segment” below). The scatter/gather list contains a set of bus address-length pairs, referencing the DMA target, which the driver may use as-is, translate in place or copy to the scatter/gather entries of its device’s DMA engine, as is appropriate for its device.

If a buffer is too big to handle in one DMA transfer (for example, due to constraints on the number of scatter/gather elements that the DMA engine can support), `udi_dma_buf_map` will set up a partial transfer. The driver will then perform multiple transfers in order to complete the original buffer request, calling `udi_dma_buf_map` repeatedly to refill the scatter/gather list after each partial transfer.

2.2 DMA Block Vector Segment

As mentioned in the UDI Core Specification, Chapter 14, “*Buffer Management*”, buffers are logically contiguous but possibly virtually and physically scattered. Such discontinuities are not visible to drivers, however, except in the way they are mapped for DMA so they can be accessed by a driver’s device. The data structures in this section are used to represent the individual address-length pairs for each contiguous (from the card’s viewpoint) block of data. These structures are intended to be as simple and convenient as possible for communicating between UDI and the driver. However, when mapping buffers for DMA, it would be even more advantageous if the I/O card could traverse these structures directly. To this end, UDI uses IEEE-1212.1 compatible scatter/gather structures.

I/O Block Vector Segments are IEEE-1212.1 compatible data structures that allow the I/O Unit to perform reads (writes) to (from) more than one physical segment (e.g., page) in a single request. The IEEE-1212.1 DMA Framework defines standard structures for representing physically-scattered but logically-contiguous data. There are two formats defined, one for buses with 32-bit addresses and one for 64-bit addresses, as shown in Figure 2-1.

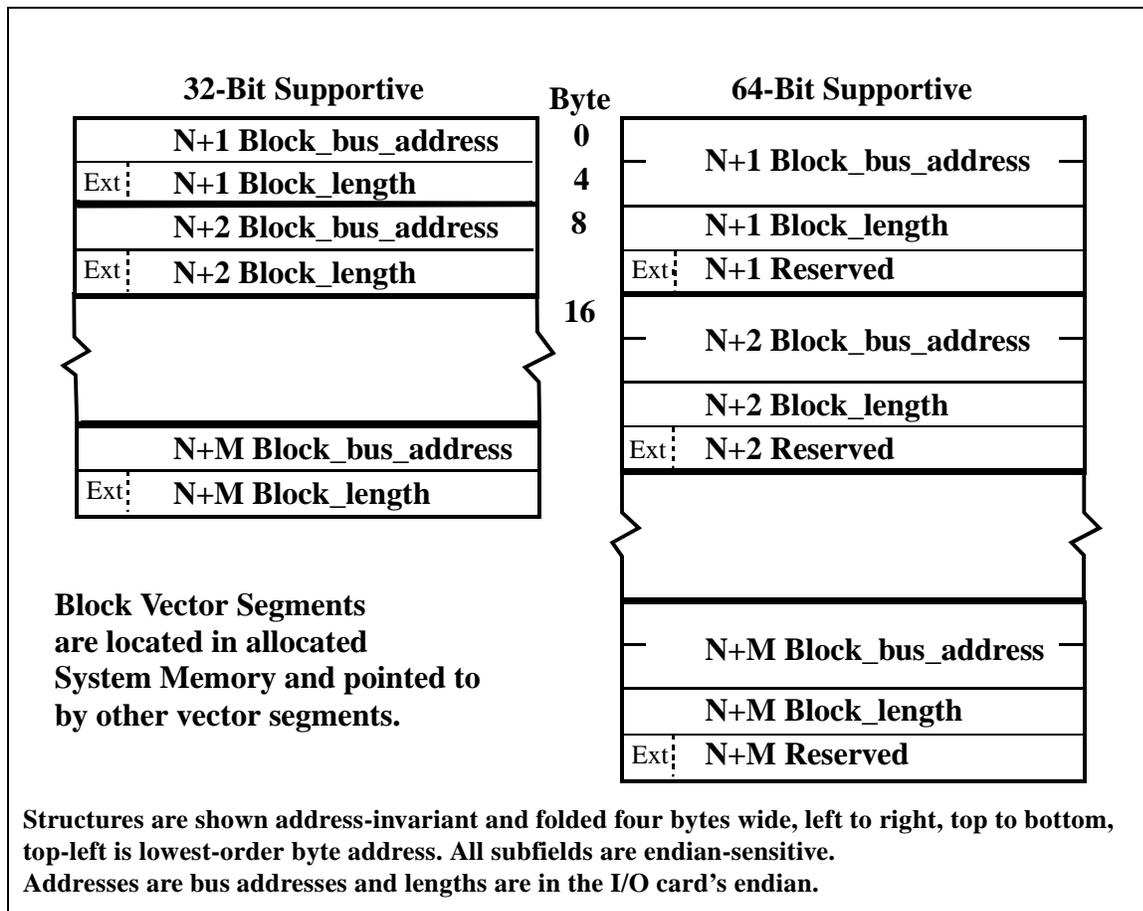


Figure 2-1 IEEE-1212.1 Block Vector Structure

As described in IEEE-1212.1, each address-length pair in a block vector segment points to and delimits a block of data bytes that is contiguous from the perspective of the I/O card. If the *Ext* bit is 0 for a pair, the address references payload data. If the *Ext* bit is 1, the address references an extension of the Block Vector structure itself (i.e., another array of address-length pairs) and the length indicates the physical size of this extension segment in bytes (which facilitates pre-fetch). Each block vector segment can contain at most one extension address; trees of vector segments are not allowed. Any address-length pairs that may follow an extension address are ignored. UDI limits its use of the flexibility allowed by IEEE-1212.1 in that for any Block Vector segment created by the environment, only the last address-length pair in a segment will ever be an indirect pointer to another Block Vector segment. All other address-length pairs will only reference data. This allows drivers for devices that are only partially IEEE-1212.1 compatible to manipulate address-length pairs without searching for indirect Block Vector references.

Environments must construct block vector segments according to the following rules (adapted from IEEE-1212.1) and the driver-specified DMA constraints (see `udi_constraints_attr_t` (DMA) on page 2-14).

1. Environments should use as few segments as possible, for space and time efficiency.

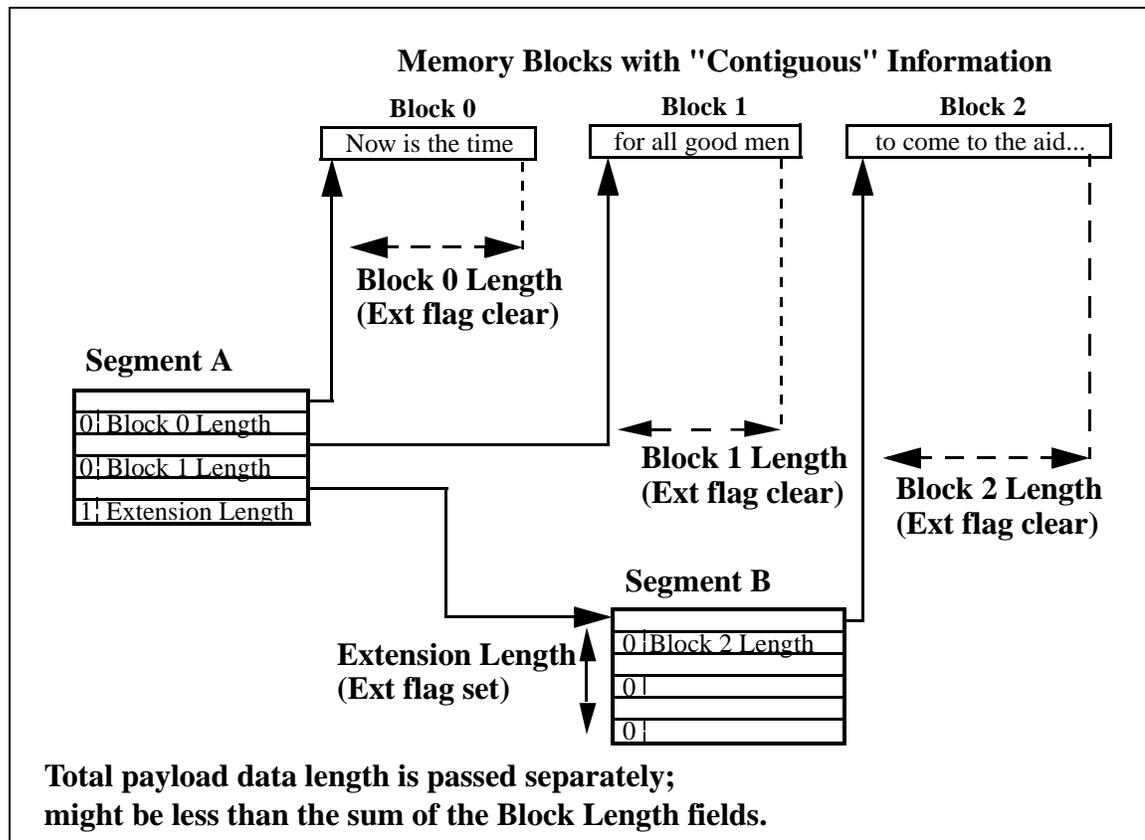


Figure 2-2 Example of Vector Use and Extension

2. When extension segments are used, the `UDI_SCGTH_EXT` flag bit shows whether an element references payload data (0) or the next Block Vector extension segment (1). If an extension segment is referenced, the *block_length* field gives the size of the next segment (making it easier for the I/O Unit to handle the segment in one block-copy operation, if desired); this size includes the indirect element, if any. The length of the last segment may extend beyond the end of the last element. The driver (or device) can always realize it has reached the last element in a segment by either encountering an indirect element, or by reaching the required total data length or total number of direct elements (see *scgth_num_elements* in `udi_scgth_t` on page 2-10).
3. When a data segment is referenced by *block_busaddr* (i.e., `UDI_SCGTH_EXT` is 0), the byte referenced by this I/O address will be aligned as specified by the associated DMA constraints; the *block_length* field, which may also be odd and will be \geq zero, gives the number of contiguous bytes of data to be sent, or space available to receive data.
4. The first indirect element in a vector segment causes the I/O Unit to follow to the next segment, ignoring any following address-length pairs in the current Block Vector segment. Therefore, for speed and simplicity, Block Vector Segments are single-threaded rather than stacked to build trees, arrays, or circular lists of extension segments. Only the last address-length pair in a segment will ever contain an indirect element; the last address-length pair in a Block Vector segment will only contain an indirect element, or be unused at the end of the data transfer.

5. The total amount of data contained by the logically contiguous buffer represented by these structures is carried by some other vehicle (normally a metalanguage control block) and a transfer must be terminated when this number of bytes have been handled, regardless of whether additional bytes of buffer space are indicated by the address-length pairs. This is true for both inbound and outbound data transfers. All elements in a vector are usable.
6. Null address values are not defined since I/O bus addresses are used and all values are valid addresses. Zero-length data segments are allowed for flexibility in use by higher software layers (however, they are not encouraged because of the slight performance impact). These shall be ignored. **Block_length** must not be zero if the UDI_SCGTH_EXT flag is set.
7. All address-length pairs of a given block vector must have the same structure, 32- versus 64-bit supportive (therefore, one **scgth_format** value applies to the entire scatter/gather list). Also, for card/host compatibility, block vector segments will be aligned according to the address size of their format; i.e., 32-bit supportive block vector segments shall be at least 4-byte aligned, and 64-bit supportive segments shall be at least 8-byte aligned.
8. All fields of block vector segments are endian-sensitive, including the UDI_SCGTH_EXT flag.
9. The reserved field for the 64-bit format shall be set to zero when initialized and must be ignored when read.

2.3 DMA Service Calls and Structures

NAME	udi_dma_handle_t <i>DMA handle type</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> typedef <HANDLE> udi_dma_handle_t; /* Null handle value for udi_dma_handle_t */ #define UDI_NULL_DMA_HANDLE <NULL_HANDLE></pre>
DESCRIPTION	<p>The DMA handle type, <code>udi_dma_handle_t</code>, holds an opaque handle that refers to DMA resources managed by the environment.</p> <p>DMA handles are not transferable between regions.</p>

NAME	udi_dma_limits <i>Platform-specific allocation and access limits</i>
SYNOPSIS	<pre> #include <udi.h> #include <udi_physio.h> typedef struct { udi_size_t max_legal_contig_alloc; udi_size_t max_safe_contig_alloc; udi_size_t cache_line_size; } udi_dma_limits_t; void udi_dma_limits (udi_dma_limits_t *dma_limits); /* Guaranteed Minimum Allocation Size */ #define UDI_DMA_MIN_ALLOC_LIMIT 4000 </pre>
MEMBERS	<p>max_legal_contig_alloc is the maximum legal size of a single DMA-able memory element (one element in a <code>udi_scgth_t</code> scatter/gather list) that can be requested from <code>udi_dma_mem_alloc</code> or <code>udi_dma_buf_map</code>. Any request for a larger size will produce indeterminate results, which could include termination of the driver or region, or even a complete system abort.</p> <p>max_safe_contig_alloc is the maximum size of a single DMA-able memory element that should be requested from <code>udi_dma_mem_alloc</code> or <code>udi_dma_buf_map</code> without being prepared to cancel an unsuccessful allocation.</p> <p>cache_line_size is the size, in bytes, of the largest cache line that affects DMA-able memory.</p>
DESCRIPTION	<p><code>udi_dma_limits_t</code> reflects the DMA memory allocation limits available on a particular system, for a particular region. These limits may vary from region to region, but will remain constant for the life of a region.</p> <p>The <code>udi_dma_limits_t</code> structure is passed back to a driver by a call to <code>udi_dma_limits</code>.</p> <p>Since UDI can be implemented on a wide variety of systems from small embedded systems to large server systems, the ability to provide contiguous DMA-able memory through <code>udi_dma_mem_alloc</code> and <code>udi_dma_buf_map</code> can vary widely. <code>udi_dma_limits</code> allows drivers to adjust their allocation algorithms to best fit their environment.</p> <p>There are two types of allocation limits: <i>legal</i> limits and <i>safe</i> limits. Legal limits represent the absolute upper bound on a single allocation. Drivers must not make requests that would exceed the legal limits.</p>

Safe limits represent the maximum amount that a driver may safely request without arranging to deal with unsuccessful allocations. For any size greater than the safe limit (but not exceeding the legal limit), drivers must cancel the request (using `udi_cancel`) after a reasonable amount of time has expired. To do this, they must set a timer using `udi_timer_start` or `udi_timer_start_repeating`. Drivers may also cancel allocations below the safe limit, but they are not expected to do so.

The `max_legal_contig_alloc` and `max_safe_contig_alloc` limits affect the size of a DMA-able memory element (one element in a `udi_scgth_t` scatter/gather list), which must be contiguous in bus address space. In most cases, DMA constraints allow the environment sufficient flexibility to use smaller pieces; in such cases, these limits wouldn't matter. However, the following DMA and transfer constraints attributes could force the environment to use larger element sizes:

`UDI_DMA_ELEMENT_GRANULARITY_BITS`

`UDI_DMA_SCGTH_MAX_ELEMENTS`

`UDI_DMA_NO_PARTIAL`

`UDI_XFER_GRANULARITY`

All of the above allocation limits are guaranteed to be greater than or equal to `UDI_DMA_MIN_ALLOC_LIMIT` (4000 bytes). This means drivers don't need to check these limits for requests that don't exceed 4000 bytes.

The `cache_line_size` value may be used to set appropriate DMA constraints for devices that need data or "slop" aligned on cache line boundaries.

REFERENCES

`udi_limits_t`, `udi_dma_mem_alloc`, `udi_dma_buf_map`,
`udi_cancel`

NAME	udi_busaddr64_t <i>64-bit bus address data type</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> typedef <OPAQUE> udi_busaddr64_t;</pre>
DESCRIPTION	<p>The <code>udi_busaddr64_t</code> type used for bus addresses in the <code>UDI_SCGTH_64</code> format is a self-contained opaque type that is guaranteed to be 64 bits in size and hold a bus address in the appropriate endianness. Drivers may copy <code>udi_busaddr64_t</code> values using assignment statements, or pass them by value or by reference to function calls, but may not use arithmetic operations or otherwise make assumptions about the internal structure of this data type. When swapping endianness of a <code>udi_busaddr64_t</code> value, drivers must use an 8-byte (64-bit) transaction size with a utility function to ensure proper endianness conversion (see Section 23.2.3, “Endian-Swapping Utilities,” on page 23-11 of the UDI Core Specification).</p>
WARNINGS	<p>The <code>udi_busaddr64_t</code> type is not transferable between regions.</p>

NAME	udi_scgth_t	<i>I/O Bus scatter/gather structure</i>
SYNOPSIS	<pre> #include <udi.h> #include <udi_physio.h> typedef struct { udi_ubit32_t block_busaddr; udi_ubit32_t block_length; } udi_scgth_element_32_t; typedef struct { udi_busaddr64_t block_busaddr; udi_ubit32_t block_length; udi_ubit32_t el_reserved; } udi_scgth_element_64_t; /* Extension Flag */ #define UDI_SCGTH_EXT 0x80000000 typedef struct { udi_ubit16_t scgth_num_elements; udi_ubit8_t scgth_format; udi_boolean_t scgth_must_swap; union { udi_scgth_element_32_t *e132p; udi_scgth_element_64_t *e164p; } scgth_elements; union { udi_scgth_element_32_t e132; udi_scgth_element_64_t e164; } scgth_first_segment; } udi_scgth_t; /* Values for scgth_format */ #define UDI_SCGTH_32 (1U<<0) #define UDI_SCGTH_64 (1U<<1) #define UDI_SCGTH_DMA_MAPPED (1U<<6) #define UDI_SCGTH_DRIVER_MAPPED (1U<<7) </pre>	
MEMBERS	<p>block_busaddr is the I/O card-relative bus address of a contiguous block of data bytes (if UDI_SCGTH_EXT bit is cleared), or an extension array of udi_scgth_element_t elements (if UDI_SCGTH_EXT is set). These blocks are contiguous from the I/O card's perspective. See the DESCRIPTION below for comments on device endianness.</p> <p>block_length is the length of the related block. On indirect elements this is at least the size, in bytes, of the direct elements plus one additional indirect element, if such an indirect element is included. In the UDI_SCGTH_32 format, the UDI_SCGTH_EXT bit is part of this field.</p>	

el_reserved includes the UDI_SCGTH_EXT bit in the UDI_SCGTH_64 format. Other bits are unused.

UDI_SCGTH_EXT is a mask value, to be applied to the **block_length** field of a `udi_scgth_element_32_t` or the **el_reserved** field of a `udi_scgth_element_64_t`. If this bit is set, then the scatter/gather element is an indirect reference to a possibly physically-discontiguous continuation of the scatter/gather list.

scgth_num_elements gives the number of direct `udi_scgth_element_t` elements included in the entire scatter/gather list (indirect elements are not included, but elements from all segments are included). This may be used to estimate I/O driver control resources.

scgth_format indicates the format of the scatter/gather elements. Exactly one of the following flags will be set in **scgth_format**, corresponding to the 32-bit and 64-bit scatter/gather formats, respectively:

UDI_SCGTH_32

UDI_SCGTH_64

In addition, the **UDI_SCGTH_DMA_MAPPED** flag will be set if the **UDI_DMA_SCGTH_VISIBILITY** constraints attribute used when the scatter/gather list was created was set to **UDI_DMA_SCGTH_DMA_VISIBLE** or **UDI_DMA_SCGTH_BOTH_VISIBLE**. If so, the scatter/gather list is “DMA-mapped”, meaning that the list elements themselves are made readable from the adapter via DMA and are in an endianness appropriate for access by the device (as indicated by the **UDI_DMA_SCGTH_ENDIANNES**s constraints attribute). The driver must ensure that its DMA device does not write to the scatter/gather list memory.

The **UDI_SCGTH_DRIVER_MAPPED** flag will be set if **UDI_DMA_SCGTH_VISIBILITY** was set to **UDI_DMA_SCGTH_DRIVER_VISIBLE** or **UDI_DMA_SCGTH_BOTH_VISIBLE**. If so, the scatter/gather list is (also) “driver-mapped”, meaning that the list elements are readable from the driver via the **e132p** or **e164p** pointer. Otherwise, these pointers are unused and their value is unspecified. The list elements will be in the driver’s endianness if and only if driver-mapped and not DMA-mapped.

scgth_must_swap is a flag indicating that the driver must swap endianness when accessing driver-mapped list elements via **e132p** or **e164p**. This flag will always be **FALSE** if the scatter/gather list is not DMA-mapped. If not driver-mapped, the value of the flag is unspecified and must not be used. When the scatter/gather list is both DMA-mapped and driver-mapped, the driver must check

this flag; it must not assume it knows whether or not to swap, based on its specified `UDI_DMA_SCGTH_ENDIANNESS`, since interposed bus bridges may change the endianness.

See Section 23.2, “Endianness Management,” on page 23-2 of the UDI Core Specification for details on how to construct C structure definitions for proper endianness handling, and the endian swapping utilities that are available.

`e132p`, `e164p` are members of the `scgth_elements` union. One of these, depending on the value of `scgth_format`, will be set to be a pointer to the scatter/gather list in the driver’s address space if `UDI_SCGTH_DRIVER_MAPPED` is set in `scgth_format`.

`scgth_first_element` is only used for DMA-mapped scatter/gather lists. Depending on the value of `scgth_format`, `e132` or `e164` contains the I/O-card relative address of the beginning of the scatter/gather list and the length in bytes of the first segment in this list. `scgth_first_element` is similar to an IEEE-1212.1 indirect element, except that the `UDI_SCGTH_EXT` bit is never set.

DESCRIPTION

The `udi_scgth_t` structure provides a way to access the block vector array segment(s) that describe the memory addresses of a DMA-mapped data buffer or shared control structure memory. Scatter/gather lists are built by the environment as a result of calls to `udi_dma_buf_map` or `udi_dma_mem_alloc`.

Scatter/gather lists are presented in either 32-bit or 64-bit format, as requested by the driver through its DMA constraints. The format used is reported back to the driver through `scgth_format`. (See `UDI_DMA_SCGTH_ADDRESSABLE_BITS` in `udi_constraints_attr_t` (DMA) on page 2-14.)

Depending on the value of the `UDI_DMA_SCGTH_VISIBILITY` constraints attribute, the scatter/gather elements may be DMA-mapped, driver-mapped, or both. The driver chooses the appropriate visibility depending on how closely its device’s scatter/gather format resembles the UDI scatter/gather list format (which is based on IEEE-1212.1). If they are identical, the driver should choose `UDI_DMA_SCGTH_DMA_VISIBLE` and let its device access the scatter/gather list directly. If the formats are similar, but not identical, the driver may choose to use `UDI_DMA_SCGTH_BOTH_VISIBLE` and modify the scatter/gather elements in place (being careful to observe `scgth_must_swap` if set) before having the device access the list. Otherwise, the driver should use `UDI_DMA_SCGTH_DRIVER_VISIBLE` and use separately allocated DMA-able memory or PIO to build a scatter/gather list in the appropriate format, based on the values it reads from the `scgth_elements` array.

If DMA-mapped, the memory for the scatter/gather segments themselves will conform to the DMA constraints related to scatter-gather memory, and the associated `udi_scgth_t` structure will contain an indirect reference to the

first element of the scatter/gather list in *scgth_first_segment*. The members of the *e132* or *e164* element within this union are in the driver's endianness. The *block_busaddr* member contains the I/O card-relative address of the beginning of the scatter/gather list, and the *block_length* member contains the length of the first segment of this list, including an indirect element, if any exists.

If not DMA-mapped, the *scgth_first_segment* field of *udi_scgth_t* is not valid, and the scatter/gather list will not contain extension elements (UDI_SCGTH_EXT); the entire list is provided in a single block vector segment.

The environment's use of the IEEE-1212.1 structures is such that all the *udi_scgth_element_t* elements in a block vector segment are used before chaining to the next segment. Therefore, the driver can use the length of the block vector segment to determine where the next indirect *udi_scgth_element_t* pair is, if any. Therefore, it is not necessary for the driver to check for UDI_SCGTH_EXT for other elements in the segment. Note, though, that the last segment may contain fewer elements than covered by the length of the segment; the total number of direct elements (*scgth_num_elements*) or the total size of the mapped memory must be used to determine that the end of the list has been reached.

WARNINGS

The *udi_scgth_t* type is not transferable between regions.

NAME	udi_constraints_attr_t (DMA)	<i>DMA constraints attributes</i>
SYNOPSIS	<pre> #include <udi.h> #include <udi_physio.h> typedef udi_ubit8_t udi_constraints_attr_t; /* DMA Convenience Attribute Codes */ #define UDI_DMA_ADDRESSABLE_BITS 100 #define UDI_DMA_ALIGNMENT_BITS 101 /* DMA Constraints on the Entire Transfer */ #define UDI_DMA_DATA_ADDRESSABLE_BITS 110 #define UDI_DMA_NO_PARTIAL 111 /* DMA Constraints on the Scatter/Gather List */ #define UDI_DMA_SCGTH_MAX_ELEMENTS 120 #define UDI_DMA_SCGTH_VISIBILITY 121 #define UDI_DMA_SCGTH_ENDIANNESNESS 122 #define UDI_DMA_SCGTH_ADDRESSABLE_BITS 123 #define UDI_DMA_SCGTH_MAX_SEGMENTS 124 /* DMA Constraints on Scatter/Gather Segments */ #define UDI_DMA_SCGTH_ALIGNMENT_BITS 130 #define UDI_DMA_SCGTH_MAX_EL_PER_SEG 131 #define UDI_DMA_SCGTH_PREFIX_BYTES 132 /* DMA Constraints on Scatter/Gather Elements */ #define UDI_DMA_ELEMENT_ALIGNMENT_BITS 140 #define UDI_DMA_ELEMENT_LENGTH_BITS 141 #define UDI_DMA_ELEMENT_GRANULARITY_BITS 142 /* DMA Constraints for Special Addressing */ #define UDI_DMA_ADDR_FIXED_BITS 150 #define UDI_DMA_ADDR_FIXED_TYPE 151 #define UDI_DMA_ADDR_FIXED_VALUE_LO 152 #define UDI_DMA_ADDR_FIXED_VALUE_HI 153 /* DMA Constraints on DMA Access Behavior */ #define UDI_DMA_SEQUENTIAL 160 #define UDI_DMA_SLOP_IN_BITS 161 #define UDI_DMA_SLOP_OUT_BITS 162 #define UDI_DMA_SLOP_OUT_EXTRA 163 #define UDI_DMA_SLOP_BARRIER_BITS 164 /* Values for UDI_DMA_SCGTH_VISIBILITY */ #define UDI_DMA_SCGTH_DMA_VISIBLE 1 #define UDI_DMA_SCGTH_DRIVER_VISIBLE 2 #define UDI_DMA_SCGTH_BOTH_VISIBLE 3 </pre>	

DESCRIPTION

```

/* Values for UDI_DMA_SCGTH_ENDIANNESS */
#define UDI_DMA_LITTLE_ENDIAN          (1U<<6)
#define UDI_DMA_BIG_ENDIAN             (1U<<5)

/* Values for UDI_DMA_ADDR_FIXED_TYPE */
#define UDI_DMA_FIXED_ELEMENT          1
#define UDI_DMA_FIXED_LIST             2
#define UDI_DMA_FIXED_VALUE           3

```

This page lists additional attribute codes that can be used with the `udi_constraints_attr_t` type and `udi_constraints_attr_set`, defined in the UDI Core Specification, to modify DMA-related constraints attributes of a constraints object.

Constraints objects, referenced by constraints handles (`udi_constraints_t`), are used to constrain transfer properties as well as memory placement of `udi_buf_t` data, memory allocated by `udi_dma_mem_alloc`, and scatter/gather lists, in order to make them addressable via DMA from a particular device.

Any driver that will be using DMA, or drives a bus bridge that has an effect on DMA transfers, must apply its DMA constraints attributes to a constraints object by calling `udi_constraints_attr_set`. The driver must then pass the constraints handle to `udi_dma_prepare` or `udi_dma_mem_alloc` to allow the environment to allocate and/or copy the memory and mapping registers used for DMA so that they meet the specified constraints.

A list of supported DMA constraints attribute codes is given below, along with the range of valid values for each attribute, which of these values are considered least and most restrictive, and the default value for the attribute. This is presented in the form of a table for each attribute category; each table is followed by detailed descriptions of each attribute.

In addition to the individual attributes specified for DMA operations there are a set of convenience attributes, which are used to specify a grouping of individual attributes. For most DMA situations, there are sets of related attributes which can all be set to identical values because the DMA engine does not distinguish between the conditions represented by the individual attributes. UDI provides convenience attributes that refer to a group of individual attributes and when used, causes all the individual attributes which are part of the group to be set at once without requiring explicit specification of each individual attribute.

Table 2-1 DMA Convenience Attributes

Convenience Attribute	Related Individual Attributes
UDI_DMA_ADDRESSABLE_BITS	UDI_DMA_DATA_ADDRESSABLE_BITS UDI_DMA_SCGTH_ADDRESSABLE_BITS
UDI_DMA_ALIGNMENT_BITS	UDI_DMA_ELEMENT_ALIGNMENT_BITS UDI_DMA_SCGTH_ALIGNMENT_BITS

UDI_DMA_ADDRESSABLE_BITS is used to simultaneously set **UDI_DMA_DATA_ADDRESSABLE_BITS** and **UDI_DMA_SCGTH_ADDRESSABLE_BITS**. This specifies the number of bits of bus address that the DMA engine can generate for access to either data or scatter/gather elements. If ≤ 32 , scatter/gather lists will be created in **UDI_SCGTH_32** format; otherwise they will be in **UDI_SCGTH_64** format.

UDI_DMA_ALIGNMENT_BITS is used to simultaneously set **UDI_DMA_ELEMENT_ALIGNMENT_BITS** and **UDI_DMA_SCGTH_ALIGNMENT_BITS**. This specifies the # of LSB bits that must be zero in the starting bus addresses of data elements and scatter/gather segments; i.e. the starting bus addresses of data elements and scatter/gather segments must be multiples of $2^{\text{UDI_DMA_ALIGNMENT_BITS}}$.

Table 2-2 DMA Constraints on the Entire Transfer

Valid Range	Least Restrictive Value	Most Restrictive Value	Default Value	Special Case Behavior for 0
UDI_DMA_DATA_ADDRESSABLE_BITS				
16..255	255	16	255	N/A
UDI_DMA_NO_PARTIAL				
0..1	0	1	0	N/A

UDI_DMA_DATA_ADDRESSABLE_BITS is the # of bits of bus address that the DMA engine can generate for access to data elements. If ≤ 32 , scatter/gather lists will be created in **UDI_SCGTH_32** format; otherwise, they will be in **UDI_SCGTH_64** format.

UDI_DMA_NO_PARTIAL is a flag indicating (if non-zero) that the device and/or driver cannot handle partial DMA mappings from **udi_dma_buf_map**. Either the entire request must be mapped in one call, or it must be failed.

Table 2-3 DMA Constraints on the Whole Scatter/Gather List

Valid Range	Least Restrictive Value	Most Restrictive Value	Default Value	Special Case Behavior for 0
UDI_DMA_SCGTH_MAX_ELEMENTS				
0..65535	0	1	0	no restriction
UDI_DMA_SCGTH_VISIBILITY				
see below	N/A	N/A	see below	N/A
UDI_DMA_SCGTH_ENDIANNES				
see below	N/A	N/A	see below	N/A
UDI_DMA_SCGTH_ADDRESSABLE_BITS				
16..255	255	16	255	N/A
UDI_DMA_SCGTH_MAX_SEGMENTS				
0..255	0	1	0	no restriction

UDI_DMA_SCGTH_MAX_ELEMENTS is the maximum # of elements that can be handled in one scatter/gather list. For DMA engines without scatter/gather support, this should be set to 1.

UDI_DMA_SCGTH_VISIBILITY selects one of three styles of visibility of the scatter/gather list elements. The elements may be made visible to the device via DMA, to the driver via a direct pointer, or both. Legal values for this attribute are:

UDI_DMA_SCGTH_DMA_VISIBLE
UDI_DMA_SCGTH_DRIVER_VISIBLE
UDI_DMA_SCGTH_BOTH_VISIBLE

The default value for this attribute is **UDI_DMA_SCGTH_DMA_VISIBLE**.

The following attributes are relevant only to DMA-mapped scatter/gather lists and need not be set in the **UDI_DMA_SCGTH_DRIVER_VISIBLE** case:

UDI_DMA_SCGTH_ENDIANNES
UDI_DMA_SCGTH_ADDRESSABLE_BITS
UDI_DMA_SCGTH_ALIGNMENT_BITS
UDI_DMA_SCGTH_MAX_SEGMENTS
UDI_DMA_SCGTH_MAX_EL_PER_SEG
UDI_DMA_SCGTH_PREFIX_BYTES

UDI_DMA_SCGTH_ENDIANNES determines the desired device endianness of DMA-mapped scatter/gather elements. This attribute may be set to one of the following values:

UDI_DMA_LITTLE_ENDIAN

UDI_DMA_BIG_ENDIAN

The default value for this attribute is unspecified. If UDI_DMA_SCGTH_VISIBILITY is not UDI_DMA_SCGTH_DRIVER_VISIBLE, this attribute must be explicitly set before the constraints object can be used with udi_dma_prepare or udi_dma_mem_alloc.

Ignored if UDI_DMA_SCGTH_VISIBILITY is set to UDI_DMA_SCGTH_DRIVER_VISIBLE.

UDI_DMA_SCGTH_ADDRESSABLE_BITS is the # of bits of bus address that the DMA engine can generate for accesses to scatter/gather list elements.

Ignored if UDI_DMA_SCGTH_VISIBILITY is set to UDI_DMA_SCGTH_DRIVER_VISIBLE.

UDI_DMA_SCGTH_MAX_SEGMENTS is the maximum # of scatter/gather segments that the DMA engine can handle. (One segment references another through indirect scatter/gather elements.)

Ignored if UDI_DMA_SCGTH_VISIBILITY is set to UDI_DMA_SCGTH_DRIVER_VISIBLE.

Table 2-4 DMA Constraints on Scatter/Gather Segments

Valid Range	Least Restrictive Value	Most Restrictive Value	Default Value	Special Case Behavior for 0
UDI_DMA_SCGTH_ALIGNMENT_BITS				
0..255	0	255	0	N/A
UDI_DMA_SCGTH_MAX_EL_PER_SEG				
0..65535	0	1	0	no restriction
UDI_DMA_SCGTH_PREFIX_BYTES				
0..65535	0	65535	0	N/A

UDI_DMA_SCGTH_ALIGNMENT_BITS is the # of LSB bits that must be zero in the starting bus address of each scatter/gather segment); i.e. the starting bus address of each segment must be a multiple of 2^UDI_DMA_SCGTH_ALIGNMENT_BITS.

Ignored if UDI_DMA_SCGTH_VISIBILITY is set to UDI_DMA_SCGTH_DRIVER_VISIBLE.

UDI_DMA_SCGTH_MAX_EL_PER_SEG is the maximum # of scatter/gather elements that can be handled in one segment, not including chain pointers to the next segment if any.

Ignored if UDI_DMA_SCGTH_VISIBILITY is set to UDI_DMA_SCGTH_DRIVER_VISIBLE.

UDI_DMA_SCGTH_PREFIX_BYTES is the # bytes of extra uninitialized DMA-mapped memory that will be allocated preceding the actual scatter/gather element array in each scatter/gather segment. This is typically used to provide additional shared control structures for communicating with the device, or as space for conversion to alternate (larger) scatter/gather formats.

Alignment constraints on the starting address of scatter/gather segments will be applied to the beginning of this prefix area rather than the beginning of the element array.

From the DMA device, the prefix bytes will be accessible at bus addresses immediately preceding the bus address of the start of the segment.

Driver access to the memory, if needed, is via the same mechanism used for access to the scatter/gather elements, through the *scgth_elements* pointers. In this case, UDI_DMA_SCGTH_VISIBILITY must be set to UDI_DMA_SCGTH_BOTH_VISIBLE.

Ignored if UDI_DMA_SCGTH_VISIBILITY is set to UDI_DMA_SCGTH_DRIVER_VISIBLE.

Table 2-5 DMA Constraints on Individual Scatter/Gather Elements

Valid Range	Least Restrictive Value	Most Restrictive Value	Default Value	Special Case Behavior for 0
UDI_DMA_ELEMENT_ALIGNMENT_BITS				
0..255	0	255	0	no restriction
UDI_DMA_ELEMENT_LENGTH_BITS				
0..32	0	1	0	no restriction
UDI_DMA_ELEMENT_GRANULARITY_BITS				
0..32	0	1	0	no restriction

UDI_DMA_ELEMENT_ALIGNMENT_BITS is the # of LSB bits that must be zero in the starting bus address of each individual scatter/gather element; i.e. the starting bus address must be a multiple of $2^{\text{UDI_DMA_ELEMENT_ALIGNMENT_BITS}}$.

UDI_DMA_ELEMENT_LENGTH_BITS is the # bits supported for the length (in bytes) of a single scatter/gather element (*block_length* in *udi_scgth_t*). No matter what this parameter is set to, no more than UDI_DMA_ADDRESSABLE_BITS will be used.

UDI_DMA_ELEMENT_GRANULARITY_BITS is the granularity of each individual scatter/gather element. The length in bytes of each element, except the last one in the whole transfer, must be a multiple of $2^{\text{UDI_DMA_ELEMENT_GRANULARITY_BITS}}$.

Table 2-6 DMA Constraints for Special Addressing Restrictions

Valid Range	Least Restrictive Value	Most Restrictive Value	Default Value	Special Case Behavior for 0
UDI_DMA_ADDR_FIXED_BITS				
0..255	0	1	0	no fixed bits
UDI_DMA_ADDR_FIXED_TYPE				
see below	see below	see below	see below	N/A
UDI_DMA_ADDR_FIXED_VALUE_LO				
0.. $2^{32}-1$	N/A	N/A	0	see below
UDI_DMA_ADDR_FIXED_VALUE_HI				
0.. $2^{32}-1$	N/A	N/A	0	see below

UDI_DMA_ADDR_FIXED_BITS if non-zero, indicates that some number of MSB bits are "fixed"; that is, they must have the same value for all addresses in a range. Depending on the value of **UDI_DMA_ADDR_FIXED_TYPE** the specified bits may be "fixed" across an entire scatter/gather list or just within each element. In most cases, it doesn't matter what the "fixed" value is, as long as it is unchanging across the required range.

The value of **UDI_DMA_ADDR_FIXED_BITS** indicates the number of "variable" bits less significant than the least significant "fixed" bit. (This can also be thought of as the bit number of the first fixed bit, counting from the LSB, starting with 0.) All bits more significant than the first "fixed" bit, up to the most significant bit according to **UDI_DMA_ADDRESSABLE_BITS**, are also "fixed".

This somewhat counter-intuitive encoding is unfortunately necessary to prevent the value from being invalidated by changes to **UDI_DMA_ADDRESSABLE_BITS**.

The following attributes are relevant only when **UDI_DMA_ADDR_FIXED_BITS** is non-zero and need not be set otherwise:

- UDI_DMA_ADDR_FIXED_TYPE
- UDI_DMA_ADDR_FIXED_VALUE_LO
- UDI_DMA_ADDR_FIXED_VALUE_HI

UDI_DMA_ADDR_FIXED_TYPE specifies the type of "fixed" address restriction, if any. It takes one of the following values:

UDI_DMA_FIXED_ELEMENT:

Each scatter/gather element may have different "fixed" values. This would typically be for a case where the DMA engine's current address register has upper bits that don't cascade, such as ISA motherboard DMAC "page" registers.

UDI_DMA_FIXED_LIST:

The whole scatter/gather list must have the same "fixed" value. Seems an unusual case, but it's here for completeness.

UDI_DMA_FIXED_VALUE:

The "fixed" bits must be equal to the value given by the pair of UDI_DMA_ADDR_FIXED_VALUE_LO and UDI_DMA_ADDR_FIXED_VALUE_HI for the whole scatter/gather list. This would only be used by special bridge drivers for some case like multiple OS instances sharing the same bus, in which the bridge driver "knows better" than the OS since it knows it was configured to allow only a sub-portion of the bus address space to be used by this instance.

UDI_DMA_FIXED_ELEMENT is least restrictive and the default. UDI_DMA_FIXED_VALUE is the most restrictive.

Ignored if UDI_DMA_ADDR_FIXED_BITS is zero.

UDI_DMA_ADDR_FIXED_VALUE_LO is the least significant 32 bits of the required "fixed" bits value in the UDI_DMA_FIXED_VALUE case; the LSB of UDI_DMA_ADDR_FIXED_VALUE_LO corresponds to the first "fixed" bit, not to the LSB of the address as a whole.

Ignored if UDI_DMA_ADDR_FIXED_BITS is zero or UDI_DMA_ADDR_FIXED_TYPE is not UDI_DMA_FIXED_VALUE.

UDI_DMA_ADDR_FIXED_VALUE_HI is the most significant 32 bits of the required "fixed" bits value in the UDI_DMA_FIXED_VALUE case. This should rarely ever need to be set.

Ignored if UDI_DMA_ADDR_FIXED_BITS is zero or UDI_DMA_ADDR_FIXED_TYPE is not UDI_DMA_FIXED_VALUE or UDI_DMA_ADDRESSABLE_BITS minus UDI_DMA_FIXED_BITS is not greater than 32.

Table 2-7 DMA Constraints on DMA Access Behavior

Valid Range	Least Restrictive Value	Most Restrictive Value	Default Value	Special Case Behavior for 0
UDI_DMA_SEQUENTIAL				
0..1	0	1	0	N/A
UDI_DMA_SLOP_IN_BITS				
0..8	0	8	0	N/A
UDI_DMA_SLOP_OUT_BITS				
0..8	0	8	0	N/A
UDI_DMA_SLOP_OUT_EXTRA				
0..65535	0	65535	0	N/A
UDI_DMA_SLOP_BARRIER_BITS				
0..255	1	0	1	no slop barrier

UDI_DMA_SEQUENTIAL is a flag indicating (if non-zero) that the DMA engine accesses memory sequentially for a data transfer, in the order listed in the scatter/gather list and in sequential address order within each element. Some environment implementations can optimize DMA transfers if they get this guarantee from the driver.

UDI_DMA_SLOP_IN_BITS is the worst-case # bits of "input slop" alignment. This indicates that DMA cycles may cause accesses to additional physical memory locations besides those explicitly listed in scatter/gather elements. Specifically, the starting (inclusive) and ending (exclusive) byte addresses of each scatter/gather element may be rounded out to boundaries that are multiples of $2^{\text{UDI_DMA_SLOP_IN_BITS}}$ for "inbound" bus transactions (i.e. from device to memory).

UDI_DMA_SLOP_OUT_BITS is the worst-case # bits of "output slop" alignment. This indicates that DMA cycles may cause accesses to additional physical memory locations besides those explicitly listed in scatter/gather elements. Specifically, the starting (inclusive) and ending (exclusive) byte addresses of each scatter/gather element may be rounded out to boundaries that are multiples of $2^{\text{UDI_DMA_SLOP_OUT_BITS}}$ for "outbound" bus transactions (i.e. from memory to device).

UDI_DMA_SLOP_OUT_EXTRA is the worst-case # bytes of "extra slop". This represents additional bytes of prefetch beyond the end of each scatter/gather element after rounding up based on **UDI_DMA_SLOP_OUT_BITS**. This applies during outbound bus transactions only.

UDI_DMA_SLOP_BARRIER_BITS is the # bits of "slop barrier". This represents a boundary across which UDI_DMA_SLOP_OUT_EXTRA prefetching will not occur; i.e. prefetching will not cross address boundaries that are multiples of $2^{\text{UDI_DMA_SLOP_BARRIER_BITS}}$. This is useful to alleviate the worst-case assumptions of UDI_DMA_SLOP_OUT_EXTRA.

Ignored if UDI_DMA_SLOP_OUT_EXTRA is zero.

REFERENCES

udi_constraints_attr_t, udi_constraints_t, udi_buf_t,
udi_dma_mem_alloc, udi_dma_prepare,
udi_constraints_attr_set

NAME	udi_dma_prepare <i>Prepare for DMA mapping</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_dma_prepare (udi_dma_prepare_call_t *callback, udi_cb_t *gcb, udi_constraints_t constraints, udi_ubit8_t flags); typedef void udi_dma_prepare_call_t (udi_cb_t *gcb, udi_dma_handle_t new_dma_handle); /* Values for flags */ #define UDI_DMA_OUT (1U<<2) #define UDI_DMA_IN (1U<<3)</pre>
ARGUMENTS	<p>callback, gcb are standard arguments described in the “Asynchronous Service Calls” section of “Standard Calling Sequences” in the UDI Core Specification.</p> <p>constraints is a constraints handle for this device.</p> <p>flags is a bitmask of flags indicating the direction(s) of transfers for which the handle will most likely be used:</p> <pre>UDI_DMA_OUT data transfer from memory to device UDI_DMA_IN data transfer from device to memory</pre> <p>new_dma_handle is an opaque handle to the newly allocated DMA object.</p>
DESCRIPTION	<p>udi_dma_prepare allocates a DMA handle that can be used to map UDI buffers for DMA transfer. In some cases DMA resources such as mapping registers will be pre-allocated at this time. The new DMA handle is passed to the driver with the callback.</p> <p>It is intended that drivers avoid using udi_dma_prepare in the main I/O path. Where possible, it should be used at bind time, with many calls to udi_dma_buf_map being made for one call to udi_dma_prepare.</p>
REFERENCES	udi_dma_buf_map, udi_dma_free, udi_constraints_t

NAME	udi_dma_buf_map	<i>Map a buffer for DMA</i>
SYNOPSIS	<pre> #include <udi.h> #include <udi_physio.h> void udi_dma_buf_map (udi_dma_buf_map_call_t *callback, udi_cb_t *gcb, udi_dma_handle_t dma_handle, udi_buf_t *buf, udi_size_t offset, udi_size_t len, udi_ubit8_t flags); typedef void udi_dma_buf_map_call_t (udi_cb_t *gcb, udi_scgth_t *scgth, udi_buf_t *new_buf, udi_boolean_t complete, udi_status_t status); /* Values for flags */ #define UDI_DMA_OUT (1U<<2) #define UDI_DMA_IN (1U<<3) #define UDI_DMA_REWIND (1U<<4) </pre>	
ARGUMENTS	<p>callback, gcb are standard arguments described in the “Asynchronous Service Calls” section of “<i>Standard Calling Sequences</i>” in the UDI Core Specification.</p> <p>dma_handle is the UDI DMA handle.</p> <p>buf is a pointer to the UDI buffer containing the data request.</p> <p>offset is the offset, in bytes, from the first valid data byte in buf at which to begin the DMA mapping. This must be less than or equal to the current valid data size of the buffer.</p> <p>len is the number of bytes to map for DMA. If flags includes UDI_DMA_IN and there are fewer than len bytes of valid data in the buffer, starting at offset, then the buffer will be extended but the initial values of the new data bytes are unspecified. If flags does not include UDI_DMA_IN, offset plus len must be less than or equal to the valid data length of the buffer.</p> <p>flags is a bitmask of one or more flags indicating the direction(s) and range of transfers for which the mapping applies:</p> <ul style="list-style-type: none"> UDI_DMA_OUT data transfer from memory to device UDI_DMA_IN data transfer from device to memory UDI_DMA_REWIND rewind to the beginning of the buffer <p>scgth is a pointer to a DMA scatter/gather structure.</p>	

DESCRIPTION

new_buf is a possibly new UDI buffer that the driver must use in place of the original **buf** after the callback is called, since the environment may need to re-allocate the buffer to accommodate new data added to the buffer as a result of this service call.

complete is a boolean flag indicating whether the buffer was completely or partially mapped for DMA.

status is a UDI status code indicating the success or failure of the service call. If not UDI_OK, the **scgth** and **complete** values are unspecified and must be ignored.

udi_dma_buf_map allocates any additional DMA resources required to prepare the specified buffer for a DMA transfer and passes back a pointer to a DMA scatter/gather structure. The udi_scgth_t contains a set of bus address/ length pairs, referencing the DMA target, which the driver may use as is, translate in place or copy to the scatter/gather entries of its device's DMA engine, as is appropriate for its device. This scatter/gather list covers the **len** mapped bytes of the buffer, starting at **offset**.

The memory referenced by the **scgth** is guaranteed to meet the DMA constraints indicated by the **constraints** argument passed to udi_dma_prepare when **dma_handle** was allocated; these constraints take precedence over those associated with the buffer. Any data modified through the DMA mapping will be visible through the buffer handle once it is unmapped with udi_dma_buf_unmap.

While DMA-mapped, the buffer data must not be accessed via the buffer handle. The driver may change the buffer's buf_size value, however, but it will not take effect until the next service call.

The first time udi_dma_buf_map is called for a particular buffer using this DMA handle, the data is mapped starting at the beginning of the buffer. If the full size cannot be mapped in one piece, the **complete** flag is set to FALSE, indicating a partial transfer. Once this piece is complete, the driver can call udi_dma_buf_map again with the same arguments. The environment will "remember" where it left off (typically using information in the DMA handle) and will map the next section of the buffer.

To rewind to the beginning of the buffer after a partial transfer, set the UDI_DMA_REWIND flag. The first call to udi_dma_buf_map for a particular buffer always starts at the beginning of the buffer, whether or not the UDI_DMA_REWIND flag is set.

At least one of UDI_DMA_IN or UDI_DMA_OUT must be specified in the **flags** argument, optionally combined with UDI_DMA_REWIND.

udi_dma_buf_map performs an implicit udi_dma_sync for the entire mapped range with the same UDI_DMA_IN and/or UDI_DMA_OUT **flags** as passed to udi_dma_buf_map.

STATUS VALUES

UDI_STAT_RESOURCE_UNAVAIL is returned if UDI_DMA_NO_PARTIAL was set in the constraints that were used to allocate the DMA handle and the environment is unable to map the entire request at one time.

REFERENCES

udi_dma_buf_unmap, udi_dma_prepare, udi_scgth_t

NAME	udi_dma_buf_unmap	<i>Release a buffer's DMA mapping</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_dma_buf_unmap (udi_dma_handle_t dma_handle);</pre>	
ARGUMENTS	dma_handle is a DMA handle previously mapped via udi_dma_buf_map.	
DESCRIPTION	<p>udi_dma_buf_unmap frees any resources associated with a DMA handle by a previous udi_dma_buf_map request. It should be used when a DMA transfer completes and its DMA handle is not going to be reused with the associated buffer.</p> <p>If the flags passed to udi_dma_buf_map for this handle included UDI_DMA_IN, udi_dma_buf_unmap performs an implicit inbound udi_dma_sync for the entire mapped range and ensures that any modified data is visible through the original buffer handle.</p> <p>If dma_handle is equal to UDI_NULL_DMA_HANDLE, explicitly or implicitly (zeroed by initial value or by using udi_memset), this function acts as a no-op. Otherwise, dma_handle must have been allocated by udi_dma_prepare.</p> <p>Even if the buffer's buf_size value was changed while the buffer was mapped, the entire buffer will be unmapped. Any buffer data beyond buf_size at the time of the udi_dma_buf_unmap will be discarded (though the memory may not).</p>	
WARNINGS	The driver must make sure that its device is no longer accessing the buffer or control structure memory before it calls udi_dma_buf_unmap.	
REFERENCES	udi_dma_buf_map, udi_dma_prepare	

NAME	udi_dma_mem_alloc	<i>Allocate shared control structure memory</i>
SYNOPSIS	<pre> #include <udi.h> #include <udi_physio.h> void udi_dma_mem_alloc (udi_dma_mem_alloc_call_t *callback, udi_cb_t *gcb, udi_constraints_t constraints, udi_ubit8_t flags, udi_size_t size); typedef void udi_dma_mem_alloc_call_t (udi_cb_t *gcb, udi_dma_handle_t new_dma_handle, void *mem_ptr, udi_scgth_t *scgth, udi_boolean_t must_swap); /* Values for flags */ #define UDI_DMA_OUT (1U<<2) #define UDI_DMA_IN (1U<<3) #define UDI_DMA_BIG_ENDIAN (1U<<5) #define UDI_DMA_LITTLE_ENDIAN (1U<<6) #define UDI_DMA_NEVERSWAP (1U<<7) </pre>	
ARGUMENTS	<p>callback, gcb are standard arguments described in the “Asynchronous Service Calls” section of “<i>Standard Calling Sequences</i>” in the UDI Core Specification.</p> <p>constraints is a UDI constraints handle that defines the memory constraints of the DMA engine that will be accessing the allocated memory.</p> <p>flags is a bitmask of flags, that control the operation of this function.</p> <p>These flags must include one or both of the following direction flags, which indicate the direction(s) of DMA transfers that may be used:</p> <ul style="list-style-type: none"> UDI_DMA_IN - transfers from device to memory. UDI_DMA_OUT - transfers from memory to device. <p>The flags must also include exactly one of:</p> <ul style="list-style-type: none"> UDI_DMA_BIG_ENDIAN - access data in big endian format. UDI_DMA_LITTLE_ENDIAN - access data in little endian format. UDI_DMA_NEVERSWAP - access data with no byte swapping (appropriate for character data). 	

Finally, **flags** may optionally include UDI_MEM_NOZERO.

size is the number of bytes of memory to allocate. The size selected must conform to the allocation limits for the current environment; see `udi_dma_limits` on page 2-7 for system-specific limits.

new_dma_handle is an opaque handle to the newly allocated DMA object.

mem_ptr is a driver-mapped pointer to the allocated DMA-able memory.

scgth is a pointer to a DMA scatter/gather structure for the newly allocated memory.

must_swap is a flag indicating that the driver must swap endianness when it accesses the DMA-able memory via **mem_ptr**. This will be computed by the environment as a function of the driver's endianness, the device endianness specified in **flags**, and any interceding bus bridges. Because bridges may introduce additional endianness changes, drivers must always check this flag rather than assuming swapping or not swapping.

See Section 23.2, "Endianness Management," on page 23-2 of the UDI Core Specification for details on how to construct C structure definitions for proper endianness handling, and the endian swapping utilities that are available.

DESCRIPTION

Allocates control structure memory that is to be shared between a UDI device driver and its corresponding DMA device. This memory will conform to the requirements expressed in **constraints**. The device accesses the memory via DMA. The driver accesses the memory using the **mem_ptr** pointer, and must handle endian conversions if so indicated by the **must_swap** flag.

The newly allocated memory will be zeroed unless UDI_MEM_NOZERO is set, in which case the initial values are unspecified.

The newly allocated memory will be aligned on the most restrictive alignment of the platform's natural alignments for *long* and pointer data types, allowing the allocated memory to be directly accessed as C structures.

At the time of the callback, the memory is mapped for DMA access by the device. The DMA mapped memory remains allocated and mapped until **new_dma_handle** is freed by a call to `udi_dma_free` or `udi_dma_mem_to_buf`, at which point the memory will be automatically unmapped and deallocated. `udi_dma_mem_to_buf` allows a range of data from the control structure memory to be placed into a UDI buffer.

Note – Unlike `udi_dma_buf_map`, `udi_dma_mem_alloc` will always produce a complete mapping.

This call is typically used for allocating memory that is contiguous from the device's perspective, but some devices may support discontinuous control memory. Whether or not contiguous device addresses are used is under control of the UDI_DMA_SCGTH_MAX_ELEMENTS property of the constraints handle, but the memory will always be virtually contiguous when accessed through the *mem_ptr* pointer. See *udi_constraints_attr_t* (DMA) for details on DMA constraints.

Note – In order to use the byte-by-byte structure layout technique for mixed endianness access to shared control structure memory, and the utility macros defined in Section 23.2, “Endianness Management,” on page 23-2 of the UDI Core Specification, the driver must declare two versions of the relevant structure(s), type-casting appropriately, and using the version that matches the device's endianness if *must_swap* is FALSE, or the “anti-endian” version if *must_swap* is TRUE.

REFERENCES

udi_dma_limits, *udi_constraints_attr_t* (DMA),
udi_dma_prepare, *udi_dma_buf_map*, *udi_dma_free*,
udi_dma_sync, *udi_dma_mem_barrier*, *udi_scgth_t*,
udi_dma_mem_to_buf

NAME	udi_dma_sync	<i>Sync host & device views of DMA-able memory</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_dma_sync (udi_dma_sync_call_t *callback, udi_cb_t *gcb, udi_dma_handle_t dma_handle, udi_size_t offset, udi_size_t length, udi_ubit8_t flags); typedef void udi_dma_sync_call_t (udi_cb_t *gcb); /* Values for flags */ #define UDI_DMA_OUT (1U<<2) #define UDI_DMA_IN (1U<<3)</pre>	
ARGUMENTS	<p>callback, gcb are standard arguments described in the “Asynchronous Service Calls” section of “<i>Standard Calling Sequences</i>” in the UDI Core Specification.</p> <p>dma_handle is a DMA handle previously mapped via <code>udi_dma_buf_map</code> or <code>udi_dma_mem_alloc</code>.</p> <p>offset is a logical offset into the mapped buffer data. That is, it is relative to the offset value provided to <code>udi_dma_buf_map</code>.</p> <p>length is the length, in bytes, of the area to synchronize. When length is zero, it applies to the entire data object referenced by dma_handle, and offset must be zero.</p> <p>flags indicates which view of the buffer to synchronize. The flags argument may be set to one or more of the following values:</p> <p style="padding-left: 20px;">UDI_DMA_OUT - sync before starting an outbound DMA to the device</p> <p style="padding-left: 20px;">UDI_DMA_IN - sync after completing an inbound DMA from the device</p>	
DESCRIPTION	<p><code>udi_dma_sync</code> is used to synchronize the host and device views of a data object that has been loaded for DMA. This may involve flushes of CPU or I/O caches, or assuring that hardware write buffers have drained.</p> <p>The required direction flags depend on the direction(s) of I/O transactions since the last synchronization points. If the buffer has been modified by the CPU, and is going to be read by the device’s DMA engine, then <code>udi_dma_sync</code> must be called with <code>UDI_DMA_OUT</code> set. This ensures that the device’s DMA engine sees the changes previously made to the buffer memory by the driver. If the device’s DMA engine has written to the buffer,</p>	

and it is going to be read by the driver, `udi_dma_sync` must be called with `UDI_DMA_IN` set. This makes sure the CPU's view of the memory includes any changes previously made by the device's DMA engine.

Note that the `UDI_DMA_OUT` and/or `UDI_DMA_IN` flags must have been set in the call to `udi_dma_prepare` or `udi_dma_mem_alloc` if the same flag is to be set in `udi_dma_sync`.

REFERENCES

`udi_dma_buf_map`, `udi_dma_buf_unmap`,
`udi_dma_mem_alloc`, `udi_dma_mem_barrier`,
`udi_dma_scgth_sync`

NAME	udi_dma_scgth_sync	<i>Sync host & device views of scatter/gather list</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_dma_scgth_sync (udi_dma_scgth_sync_call_t *callback, udi_cb_t *gcb, udi_dma_handle_t dma_handle); typedef void udi_dma_scgth_sync_call_t (udi_cb_t *gcb);</pre>	
ARGUMENTS	<p><i>callback</i>, <i>gcb</i> are standard arguments described in the “Asynchronous Service Calls” section of “<i>Standard Calling Sequences</i>” in the UDI Core Specification.</p> <p><i>dma_handle</i> is a DMA handle previously mapped via <code>udi_dma_buf_map</code> or <code>udi_dma_mem_alloc</code>.</p>	
DESCRIPTION	<p><code>udi_dma_scgth_sync</code> is used to synchronize the host and device views of the scatter/gather list memory for a data object that has been loaded for DMA, since <code>udi_dma_sync</code> only affects the actual data memory. This may involve flushes of CPU or I/O caches, or assuring that hardware write buffers have drained.</p> <p>The entire set of scatter/gather elements, in all segments, as well as any prefix bytes (<code>UDI_DMA_SCGTH_PREFIX_BYTES</code> in <code>udi_constraints_attr_t</code> (DMA) on page 2-14) included in the synchronization. It is assumed that the DMA device did not write to this memory unless <code>UDI_DMA_SCGTH_PREFIX_BYTES</code> was greater than zero, but the driver may have read and/or written.</p> <p>This function is only needed when the scatter/gather list visibility is <code>UDI_DMA_SCGTH_BOTH_VISIBLE</code> (see <code>udi_constraints_attr_t</code> (DMA) on page 2-14), since this is the only case in which the driver will write to the scatter/gather segment memory and the device will read from it, and must not be used in other cases. In the <code>UDI_DMA_SCGTH_BOTH_VISIBLE</code> case, <code>udi_dma_scgth_sync</code> must be called before the device reads from the scatter/gather list.</p>	
REFERENCES	<code>udi_dma_buf_map</code> , <code>udi_dma_mem_alloc</code> , <code>udi_dma_sync</code>	

NAME	udi_dma_mem_barrier	<i>Ordering barrier for accesses to DMA-able memory</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_dma_mem_barrier (udi_dma_handle_t dma_handle);</pre>	
ARGUMENTS	dma_handle is a DMA handle previously allocated by <code>udi_dma_mem_alloc</code> .	
DESCRIPTION	<code>udi_dma_mem_barrier</code> is used to impose ordering constraints between driver accesses to shared control structure memory allocated by <code>udi_dma_mem_alloc</code> . This ensures that no loads or stores to the memory associated with dma_handle that are executed after the call to <code>udi_dma_mem_barrier</code> will be visible to the device until all prior loads or stores are visible.	
REFERENCES	<code>udi_dma_mem_alloc</code> , <code>udi_dma_sync</code>	

NAME	udi_dma_free <i>Free DMA resources</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_dma_free (udi_dma_handle_t <i>dma_handle</i>);</pre>
ARGUMENTS	<i>dma_handle</i> is a DMA handle allocated by udi_dma_prepare or udi_dma_mem_alloc.
DESCRIPTION	<p>udi_dma_free frees a DMA handle and any associated resources allocated by udi_dma_prepare or udi_dma_mem_alloc.</p> <p>If the <i>dma_handle</i> is associated with shared control structure memory allocated by udi_dma_mem_alloc, the allocated memory will be freed, as will the scatter/gather list.</p> <p>If the <i>dma_handle</i> was mapped to a buffer with udi_dma_buf_map, it must be unmapped (using udi_dma_buf_unmap) before udi_dma_free is called.</p> <p>If <i>dma_handle</i> is equal to UDI_NULL_DMA_HANDLE, explicitly or implicitly (zeroed by initial value or by using udi_memset), this function acts as a no-op. Otherwise, <i>dma_handle</i> must have been allocated by udi_dma_prepare or udi_dma_mem_alloc.</p> <hr/> <p>Note – udi_dma_free does not do a udi_dma_sync operation.</p> <hr/>
WARNINGS	The driver must make sure that its device is no longer accessing the buffer or control structure memory before it calls udi_dma_free.
REFERENCES	udi_dma_prepare, udi_dma_buf_map, udi_dma_buf_unmap, udi_dma_mem_alloc

NAME	udi_dma_mem_to_buf <i>Convert DMA-mapped control memory into a buffer</i>
SYNOPSIS	<pre> #include <udi.h> #include <udi_physio.h> void udi_dma_mem_to_buf (udi_dma_mem_to_buf_call_t *callback, udi_cb_t *gcb, udi_dma_handle_t dma_handle, udi_size_t offset, udi_size_t size); typedef void udi_dma_mem_to_buf_call_t (udi_cb_t *gcb, udi_buf_t *new_buf); </pre>
ARGUMENTS	<p><i>callback</i>, <i>gcb</i> are standard arguments described in the “Asynchronous Service Calls” section of “Standard Calling Sequences” in the UDI Core Specification.</p> <p><i>dma_handle</i> is a DMA handle previously allocated by <code>udi_dma_mem_alloc</code>.</p> <p><i>offset</i> is the beginning offset from the first byte of the shared control structure memory to include in the buffer.</p> <p><i>size</i> is the number of bytes to include in the buffer.</p> <p><i>new_buf</i> is a pointer to a UDI buffer containing the indicated data.</p>
DESCRIPTION	<p><code>udi_dma_mem_to_buf</code> frees a DMA handle and any associated resources allocated by <code>udi_dma_mem_alloc</code>, except the data from the indicated range of the memory block are used as the initial contents of a newly-allocated buffer.</p> <p>The constraints that were used to allocate <i>dma_handle</i> shall also be used to allocate <i>new_buf</i>.</p> <p><code>udi_dma_mem_to_buf</code> is typically used for inbound data from devices that store both shared control structures and data in the same piece of memory.</p>
REFERENCES	<code>udi_dma_mem_alloc</code> , <code>udi_buf_write</code>



3.1 Overview

Programmed I/O (PIO) refers to data transfers initiated by a CPU under driver software control to access registers or memory on a device. This is contrasted with Direct Memory Access (DMA), which involves transfers initiated by a device to access system memory. On some hardware platforms PIO is handled via normal memory loads and stores (“memory-mapped I/O”); on others it requires special I/O instructions. UDI hides this difference from drivers.

In UDI, Programmed I/O (PIO) operations are performed through environment service calls coded as function calls rather than by direct memory references or I/O instructions in the drivers. This abstraction is necessary for driver portability across different machine architectures for (at least) the following reasons:

1. Different machines support different endian architectures and some provide hardware-assisted endianness swapping.
2. Some machine architectures may restrict direct access.
3. Some bus bridges are non-transparent and require software intervention (via a bus bridge driver) for each PIO access.

To reduce the overhead associated with using function calls, UDI allows multiple PIO transactions to be performed with a single function call.

PIO access properties are encapsulated with a PIO handle (`udi_pio_handle_t`). A PIO handle is an opaque data type containing the addressing, data translation and access constraint information required to access a device or memory address in a particular address space. For example, for PCI devices the address space could be memory space, I/O space or configuration space. Information specifying endianness, required atomicity, and data ordering constraints is also contained in the PIO handle.

Also associated with each PIO handle is a transaction list, that specifies the PIO operations to be invoked when that handle is passed to `udi_pio_trans`.

For each transaction that accesses the device, a PIO offset is specified, which indicates the offset into the space referenced by a PIO handle at which an I/O operation is to occur. The space covered by a particular PIO handle is contiguous with respect to the addresses seen by the device.

Synchronization and ordering requirements for PIO operations, with respect to system operations such as cache or I/O buffer flushes, are device and driver specific. Drivers control when such synchronization and ordering operations are performed through the use of `UDI_PIO_SYNC` and `UDI_PIO_BARRIER` transactions. Additional data ordering and synchronization requirements may be associated with a PIO handle via the *pio_attributes* parameter to `udi_pio_map`.

3.2 PIO Handle Allocation and Initialization

The following functions are used to allocate and initialize handles to PIO-accessible memory. The `udi_pio_map` function allocates a handle for a range of device memory/registers. The `udi_pio_unmap` function deallocates the PIO handle and any associated resources when the driver no longer needs them. The `udi_pio_atomic_sizes` function returns a bitmask of transaction sizes that can be handled atomically.

Additionally, the driver must register a special PIO handle with the environment via the `udi_pio_abort_sequence` function, which the environment can use when “killing” a faulting region to stop the corresponding device from initiating further actions.

NAME	udi_pio_handle_t <i>PIO handle type</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> typedef <HANDLE> udi_pio_handle_t; /* Null handle value for udi_pio_handle_t */ #define UDI_NULL_PIO_HANDLE <NULL_HANDLE></pre>
DESCRIPTION	<p>The PIO handle type, <code>udi_pio_handle_t</code>, holds an opaque handle that refers to an environment object which contains addressing, data translation and access information, as well as a PIO transaction list.</p> <p>PIO handles are transferable between regions.</p> <p>Drivers will often have multiple PIO handles for different address spaces on the same device (e.g. one handle for configuration space, another for I/O registers, etc.). Drivers might also have multiple PIO handles for the same address space if different translation requirements or constraints exist within that address space, or if different types of operations are to be used at different times.</p>
REFERENCES	<code>udi_pio_trans_t</code>

NAME	udi_pio_map	<i>Map device memory/registers for access</i>
SYNOPSIS	<pre> #include <udi.h> #include <udi_physio.h> void udi_pio_map (udi_pio_map_call_t *callback, udi_cb_t *gcb, udi_ubit32_t regset_idx, udi_ubit32_t base_offset, udi_ubit32_t length, const udi_pio_trans_t *trans_list, udi_ubit16_t list_length, udi_ubit16_t pio_attributes, udi_ubit32_t pace); typedef void udi_pio_map_call_t (udi_cb_t *gcb, udi_pio_handle_t new_pio_handle); /* Values for pio_attributes */ #define UDI_PIO_STRICTORDER (1U<<0) #define UDI_PIO_UNORDERED_OK (1U<<1) #define UDI_PIO_MERGING_OK (1U<<2) #define UDI_PIO_LOADCACHING_OK (1U<<3) #define UDI_PIO_STORECACHING_OK (1U<<4) #define UDI_PIO_BIG_ENDIAN (1U<<5) #define UDI_PIO_LITTLE_ENDIAN (1U<<6) #define UDI_PIO_NEVERSWAP (1U<<7) #define UDI_PIO_UNALIGNED (1U<<8) </pre>	
ARGUMENTS	<p>callback, gcb are standard arguments described in the “Asynchronous Service Calls” section of “<i>Standard Calling Sequences</i>” in the UDI Core Specification.</p> <p>regset_idx is the index to the specified register set within a given address space. The definition of regset_idx is bus and device dependent, and must be determined by the driver via the <i>bus_type</i> instance attribute (see Chapter 16, “<i>Instance Attribute Management</i>” in the UDI Core Specification) in conjunction with the corresponding UDI bus bindings (see “Section 3: Bus Bindings”). For example, if <i>bus_type</i> is “pci” then the settings would be based on the definitions in the <i>UDI PCI Bus Binding Specification</i>.</p> <p>base_offset is the offset into the selected register address space at which this mapping is to start. This offset must be a multiple of all transaction sizes used to access device addresses through this handle, unless UDI_PIO_UNALIGNED is set in pio_attributes.</p>	

length is the length to be mapped in bytes.

trans_list is a list of one or more PIO transactions to apply to this device when using `udi_pio_trans` with this PIO handle. The memory pointed to by **trans_list** must be in a module-global (and thus read-only) variable.

list_length is the number of elements in the **trans_list** list.

pio_attributes are attribute flags that specify the data translation and ordering requirements of accesses to the device memory.

Data Translation Flags

Attribute flags specifying the data translation requirements of the device are mutually exclusive. At most one must be used in a given handle allocation. The default is `UDI_PIO_NEVERSWAP` if no data translation flags are specified.

The following data translation flags are defined:

`UDI_PIO_BIG_ENDIAN` - access data in big endian format.

`UDI_PIO_LITTLE_ENDIAN` - access data in little endian format.

`UDI_PIO_NEVERSWAP` - access data with no byte swapping (appropriate for character data). Device transactions greater than one byte in size are illegal if this flag is set.

Data Ordering Flags

Attribute flags specifying the data ordering requirements for the device may be used in combination. With the exception of `UDI_PIO_STRICTORDER` these values are advisory, not mandatory. For example, data can be ordered without being merged or cached, even though a driver permits unordered, merged and cached operation. Strict ordering may be required for certain I/O operations but can be very costly on high performance computers. For unordered execution, merging or caching, the `UDI_PIO_SYNC` operation (see `udi_pio_trans`) can provide more specific and efficient synchronization.

The default is `UDI_PIO_STRICTORDER` if no ordering flags are specified. If **pace** is non-zero, **pio_attributes** must not include any data ordering flags other than `UDI_PIO_STRICTORDER`.

The following data ordering flags are defined:

UDI_PIO_STRICTORDER - the CPU must issue the references in order, as the programmer specified. Strict ordering is the default if no other data ordering flags are set. UDI_PIO_STRICTORDER must not be combined with any other data ordering flags.

UDI_PIO_UNORDERED_OK - the CPU may reorder both load and store references to the mapped area.

UDI_PIO_MERGING_OK - merging and batching: the CPU may merge individual stores to consecutive locations (for example, turn two consecutive byte stores into one halfword store), and it may batch individual loads (for example, turn two consecutive byte loads into one halfword load). If UDI_PIO_MERGING_OK is set, UDI_PIO_UNORDERED_OK is treated as if it were set, even if it isn't actually set.

UDI_PIO_LOADCACHING_OK - load caching: the CPU may cache the data it fetches and reuse it until another store occurs. The default is to fetch new data on every load. If UDI_PIO_LOADCACHING_OK is set, UDI_PIO_MERGING_OK is treated as if it were set, even if it isn't actually set.

UDI_PIO_STORECACHING_OK - store caching: the CPU may keep the data in the cache and push it to the device (perhaps with other data) at a later time. The default is to push data immediately to the device. If UDI_PIO_STORECACHING_OK is set, UDI_PIO_LOADCACHING_OK is treated as if it were set, even if it isn't actually set.

Data Alignment Flags

UDI_PIO_UNALIGNED - unaligned accesses allowed: if this flag is set, there are no restrictions on *base_offset* and on PIO offsets for individual transactions; otherwise, both of these must be multiples of all transaction sizes used in *trans_list*. If this flag is set, none of the transactions using this handle are guaranteed to be atomic.

pace

is the PIO pacing time, in microseconds, for this handle. The environment will guarantee that any device access that occurs using this PIO handle will be followed by at least *pace* microseconds before another access occurs to the same device register set (via any handle). If non-zero, the UDI_PIO_STRICTORDER attribute must be specified.

PIO transactions may be posted/cached for future completion, so the apparent execution time of PIO operation calls may vary from the driver's perspective, but there will be at least *pace* microseconds between them from the hardware's perspective¹.

DESCRIPTION

new_pio_handle is an opaque data type containing the addressing, endian translation, and access constraint information required to access the associated device memory.

`udi_pio_map` initializes a handle through which a driver instance may access its device using PIO access routines. The **reset_idx**, **base_offset**, and **length** arguments select the range of device memory to be made accessible; the **pio_attributes** argument specifies the ways in which it may be accessed.

Note that the driver is allowed to map the same piece of device memory multiple times for different access requirements. For example, if the driver wants to do both strongly-ordered and weakly ordered accesses to a given register set it can map the register set once with `UDI_PIO_STRICTORDER` (strongly ordered accesses) and once with `UDI_PIO_UNORDERED_OK` (weakly ordered accesses). The driver would then use the “strong” PIO handle when it wants to do a strongly ordered access, and the “weak” handle otherwise. Similarly, multiple handles to a given register set could be used for zero pacing vs. various nonzero pacing values, and for other variations in attributes, offset, length, or trans list.

REFERENCES

`udi_pio_handle_t`, `udi_pio_unmap`, `udi_pio_trans`

1. The PIO pacing delay may be implemented in hardware or software by the environment, but represents a potential inline delay in region execution. The driver writer is strongly advised to use only the delays required for their device, as excessive use can have adverse effects on the driver and the rest of the system.

NAME	udi_pio_unmap	<i>Unmap a PIO handle and free associated resources</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_pio_unmap (udi_pio_handle_t pio_handle);</pre>	
ARGUMENTS	pio_handle is the PIO handle to be freed.	
DESCRIPTION	<p>This service call unmaps a PIO access handle, freeing any associated resources. Upon return pio_handle is no longer valid.</p> <p>If pio_handle is equal to UDI_NULL_PIO_HANDLE, either explicitly or implicitly (zeroed by initial value or by using <code>udi_memset</code>), this function acts as a no-op. Otherwise, pio_handle must have been allocated by <code>udi_pio_map</code>.</p>	
REFERENCES	<code>udi_pio_handle_t</code> , <code>udi_pio_map</code>	

NAME	udi_pio_atomic_sizes	<i>Retrieve supported PIO operation atomicity</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> udi_ubit32_t udi_pio_atomic_sizes (udi_pio_handle_t pio_handle);</pre>	
ARGUMENTS	pio_handle is a PIO handle previously acquired via <code>udi_pio_map</code> .	
DESCRIPTION	<p>This routine retrieves an encoding of the PIO atomic transaction sizes (see <i>atomic transaction</i> in the Glossary) for which atomicity is supported when doing PIO accesses through pio_handle. The set of supported atomic sizes is platform and bus specific, and may also be dependent on device address space, and register set.</p> <p>The driver will typically use this routine to verify up front that the atomic sizes it requires are supported, and then fail its initialization if they're not supported. If such a check is not done and the platform does not support the device's atomicity requirements then the results of corresponding PIO accesses will be indeterminate.</p> <p>Atomicity is not guaranteed, regardless of transaction size, for any handle mapped with the <code>UDI_PIO_UNALIGNED</code> flag, or when using <code>udi_pio_probe</code>.</p>	
RETURN VALUES	A bit encoding of the PIO operation sizes for which atomicity is supported is returned. Each bit position corresponds to a power of two byte PIO operation size. If a bit is set PIO operations of the corresponding size are supported atomically through the bus hierarchy to the device. For example, a configuration supporting atomic PIO operations of 1, 2, and 4 bytes would return a binary 0111.	
REFERENCES	<code>udi_pio_map</code>	

NAME	udi_pio_abort_sequence <i>Register a PIO abort sequence</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_pio_abort_sequence (udi_pio_handle_t pio_handle, udi_size_t scratch_requirement);</pre>
ARGUMENTS	<p>pio_handle is a special PIO handle, previously acquired from <code>udi_pio_map</code>, whose associated trans list is designed to stop the associated device from initiating (mastering) transactions on its bus in case the driver fails.</p> <p>scratch_requirement is the number of bytes of scratch space needed during processing of the associated trans list.</p>
DESCRIPTION	<p>When a driver does something illegal, causes a system fault, or otherwise misbehaves, the UDI environment can “kill” the corresponding driver region. However, since the faulting driver region can no longer be trusted to execute correctly, the environment needs some way to stop the corresponding device from proceeding without having to re-execute the faulting region. <code>udi_pio_abort_sequence</code> provides the environment a PIO handle (with an associated trans_list) which it can use for this purpose.</p> <p>The driver calls <code>udi_pio_abort_sequence</code> to register pio_handle with the environment. pio_handle must be mapped with a trans list that can be used to stop the corresponding device from initiating (mastering) transactions on the bus, including the generation of DMA transactions and interrupts. The PIO sequence specified by the associated trans list does not need to flush data or preserve device state, and should do the simplest sequence possible to stop the device, such as resetting it or otherwise stopping its ability to do bus mastering.</p> <p>If needed, the registered trans list will be executed as if by <code>udi_pio_trans</code>. The <code>UDI_PIO_SCRATCH</code> addressing mode may be used to access up to scratch_requirement bytes of scratch space. <code>UDI_PIO_BUF</code> and <code>UDI_PIO_MEM</code> must not be used.</p> <p>To facilitate handling faults in as wide a portion of the driver as possible, the driver should call <code>udi_pio_abort_sequence</code> as early as possible in its per-instance initialization sequence. If the device changes state in such a way that a different procedure is needed to shut it down, the driver may call <code>udi_pio_abort_sequence</code> again to replace the previously-registered sequence, but this should only be done if absolutely necessary.</p> <p>The PIO handle passed to this service call is “given away” (as if with <code>udi_pio_unmap</code>). The driver must no longer access this handle.</p>
REFERENCES	<code>udi_pio_map</code> , <code>udi_pio_unmap</code> , <code>udi_pio_trans_t</code>

3.3 PIO Access Service Calls and Structures

PIO devices are accessed by passing a list of PIO transactions to `udi_pio_trans`. Each of the transactions are processed in turn. When all the transactions in the list are completed, the driver's callback routine is called.

There are three types of PIO transactions that affect PIO devices: input, output, and synchronization. Input transactions read data from a device; output transactions send data to a device; synchronization transactions are used to control ordering of input/output transactions on weakly-ordered architectures.

The following rules and restrictions apply to all PIO transactions:

1. The ***tran_size*** parameter of each transaction descriptor must be from 0 to 5 (which may be represented as `UDI_PIO_1BYTE`, `UDI_PIO_2BYTE`, ...). Since it encodes size as a power of two, this represents transaction sizes from 1 to 32 bytes (256 bits). For transaction descriptors to which transaction size does not apply (e.g. `UDI_PIO_BRANCH`), ***tran_size*** must be zero.
2. Endian translation is performed, when needed, on each $2^{\text{tran_size}}$ byte quantity. To transfer a character or byte array to or from device memory the driver must use a ***tran_size*** of `UDI_PIO_1BYTE` with an appropriate repeat count rather than using `UDI_PIO_NEVERSWAP` with a single basic transaction.
3. If the ***pio_attributes*** parameter to `udi_pio_map` did not include `UDI_PIO_UNALIGNED`, then the ***base_offset*** in `udi_pio_map`, the base address alignment of the register set itself, and any PIO offsets accessed by PIO transactions must all be transaction-size aligned (i.e. multiples of $2^{\text{tran_size}}$).
4. If copied to driver memory, the datum must be transaction-size aligned relative to a `udi_mem_alloc'd` piece of memory, a control block's scratch area, or a variable whose data type is $2^{\text{tran_size}}$ bytes in size, even if `UDI_PIO_UNALIGNED` was set in ***pio_attributes***.
5. PIO pacing, when indicated, occurs between each $2^{\text{tran_size}}$ byte transfer that accesses the PIO device (once per repetition). The only PIO transaction types that access the PIO device are `UDI_PIO_IN`, `UDI_PIO_OUT`, `UDI_PIO_IN_IND`, `UDI_PIO_OUT_IND`, `UDI_PIO_REP_IN_IND`, and `UDI_PIO_REP_OUT_IND`.
6. If a PIO device access uses a transaction size that is one of the atomic transaction sizes indicated by `udi_pio_atomic_sizes`, all $2^{\text{tran_size}}$ bytes are guaranteed to be transferred together "atomically" (as defined in "atomic transaction" in the Glossary), unless `UDI_PIO_UNALIGNED` was set in ***pio_attributes***.

Since odd sizes (eg. 3 bytes) are not allowed, access to, for example, a 24-bit register must either be encapsulated within a larger access and the unused bits appropriately tossed (as described below), or it must be split up into smaller accesses (eg. 1- and 2-byte accesses) as appropriate for the IO card or shared control structure.

For example, a 24-bit register at PIO offsets 1 through 3 in little endian device memory can be read in an endian-neutral manner by doing a 32-bit read at PIO offset 0 and shifting the result right by 8 bits.

A 24-bit register at offsets 0 through 2 could be read by doing the 32-bit read and then extracting the low-order 3 bytes using a mask of `0xFFFFF`.

A 24-bit register at offsets 2 through 4 could be read by doing a 16-bit read at PIO offset 2, and adding this to the result of an 8-bit read at offset 4 shifted left by 16 bits.

NAME	udi_pio_trans_t	<i>PIO transaction descriptor</i>
SYNOPSIS	<pre> #include <udi.h> #include <udi_physio.h> typedef struct { udi_ubit8_t pio_op; udi_ubit8_t tran_size; udi_ubit16_t operand; } udi_pio_trans_t; /* Values for tran_size */ #define UDI_PIO_1BYTE 0 #define UDI_PIO_2BYTE 1 #define UDI_PIO_4BYTE 2 #define UDI_PIO_8BYTE 3 #define UDI_PIO_16BYTE 4 #define UDI_PIO_32BYTE 5 /* Values for register numbers in pio_op */ #define UDI_PIO_R0 0 #define UDI_PIO_R1 1 #define UDI_PIO_R2 2 #define UDI_PIO_R3 3 #define UDI_PIO_R4 4 #define UDI_PIO_R5 5 #define UDI_PIO_R6 6 #define UDI_PIO_R7 7 /* Values for addressing modes in pio_op */ #define UDI_PIO_DIRECT 0x00 #define UDI_PIO_SCRATCH 0x08 #define UDI_PIO_BUF 0x10 #define UDI_PIO_MEM 0x18 /* Values for Class A opcodes in pio_op */ #define UDI_PIO_IN 0x00 #define UDI_PIO_OUT 0x20 #define UDI_PIO_LOAD 0x40 #define UDI_PIO_STORE 0x60 /* Values for Class B opcodes in pio_op */ #define UDI_PIO_LOAD_IMM 0x80 #define UDI_PIO_CSKIP 0x88 #define UDI_PIO_IN_IND 0x90 #define UDI_PIO_OUT_IND 0x98 #define UDI_PIO_SHIFT_LEFT 0xA0 #define UDI_PIO_SHIFT_RIGHT 0xA8 #define UDI_PIO_AND 0xB0 #define UDI_PIO_AND_IMM 0xB8 </pre>	

```

#define UDI_PIO_OR                0xC0
#define UDI_PIO_OR_IMM           0xC8
#define UDI_PIO_XOR              0xD0
#define UDI_PIO_ADD              0xD8
#define UDI_PIO_ADD_IMM          0xE0
#define UDI_PIO_SUB              0xE8

/* Values for Class C opcodes in pio_op */
#define UDI_PIO_BRANCH           0xF0
#define UDI_PIO_LABEL            0xF1
#define UDI_PIO_REP_IN_IND       0xF2
#define UDI_PIO_REP_OUT_IND      0xF3
#define UDI_PIO_DELAY            0xF4
#define UDI_PIO_BARRIER        0xF5
#define UDI_PIO_SYNC             0xF6
#define UDI_PIO_SYNC_OUT        0xF7
#define UDI_PIO_END              0xFF

```

MEMBERS

pio_op is the type of operation to use for this transaction. **pio_op** encodes the type of transaction and the selection of source and destination operands.

tran_size is the power-of-2 size of a basic PIO transaction. The actual size is $2^{\text{tran_size}}$ bytes. The **tran_size** value must be from 0 to 5, corresponding to a transaction size of 1 to 32 bytes (256 bits). The mnemonic constants UDI_PIO_1BYTE through UDI_PIO_32BYTE are available for use in setting **tran_size**.

operand is an operand value or address for the transaction. Its interpretation depends on the type of operation given by **pio_op**.

DESCRIPTION

A PIO transaction descriptor (`udi_pio_trans_t`) describes in most cases a single PIO register/memory access transaction or an operation using a set of temporary registers. Multiple such transaction descriptors may be combined into a single list, called a *trans list*, to perform more complex sequences with a single call to `udi_pio_trans`. Advanced operations in the trans list allow for bit manipulations, repeat counts, serialization, and conditional branching. All PIO transactions are subject to the rules and procedural information listed on page 3-11.

Each element in the trans list array represents a basic transaction, and operates on a single transaction-sized piece of data. When transferring data to or from a device (for example, with UDI_PIO_IN or UDI_PIO_OUT), the entire transaction-size bytes will be transferred as a single atomic bus transaction if possible. See `udi_pio_atomic_sizes` on page 3-9 for a description of how to determine the atomic sizes supported on a particular platform.

While simple transactions can be handled with simple operations, UDI supports a rich set of PIO operation types, so that even fairly sophisticated manipulation of device data can be performed with a single call to `udi_pio_trans`. To allow these operations to be efficiently expressed, `udi_pio_trans` uses an abstract register load/store model.

During execution of `udi_pio_trans`, the UDI environment maintains a set of 8 temporary “registers”, numbered 0 through 7, each of which can hold one data item up to 32 bytes (256 bits) in size. Most PIO operations use one or more of these registers. If the size of a value loaded into a register (determined by `tran_size`) is less than 32 bytes, the “upper” (most-significant) bytes are treated as zero if the register value is subsequently used with a larger transaction size. If a register value is used in an operation with a smaller transaction size than when the register was last loaded, the upper bytes will be ignored. Arithmetic operations do not generate underflows or overflows, they simply wrap around; in other words, all arithmetic is modulo $2^{(2^{\text{tran_size}} * 8)}$.

In addition to using the temporary registers, some PIO operations allow access to permanent memory associated with one of the arguments to `udi_pio_trans`. These operations can access control block scratch space, `udi_buf_t` buffer contents, or driver memory.

PIO operations are divided into 3 classes, based on the type of operands used:

Class A operations can use register values or permanent memory values. The `pio_op` value for a class A operation selects a register number (0-7) as well as an addressing mode from the following table.

Table 3-1 PIO Addressing Modes

Mnemonic	Value	Description
UDI_PIO_DIRECT	0x00	Register contents are used directly
UDI_PIO_SCRATCH	0x08	Register holds 32-bit offset into scratch space
UDI_PIO_BUF	0x10	Register holds 32-bit offset into buffer data
UDI_PIO_MEM	0x18	Register holds 32-bit offset into memory block

To set `pio_op` for a Class A operation, the register number and exactly one of the addressing mode mnemonics must be added to the opcode mnemonic:

$$\text{pio_op} = \text{operation_code} + \text{addressing_mode} + \text{register}$$

The mnemonic constants `UDI_PIO_R0` through `UDI_PIO_R7` are available for use in selecting registers for all opcodes and operands that use register numbers.

Class A `pio_op` values are encoded with the following bit patterns: `000aarrrr` - `011aarrrr`, where `aa` selects the addressing mode according to Table 3-1 and `rrrr` selects a register number.

Class B operations can use register values but not permanent memory values. The *pio_op* value for a class B operation selects a register number (0-7), which is used as if the UDI_PIO_DIRECT addressing mode were specified.

To set *pio_op* for a Class B operation, the register number must be added to the opcode mnemonic:

$$pio_op = operation_code + register$$

Class B *pio_op* values are encoded with the following bit patterns: 10000rrrr - 11101rrrr, where rrr selects a register number to be used with that addressing mode.

Class C operations use neither register values nor permanent memory values. The *pio_op* value for a class C operation consists only of the opcode:

$$pio_op = operation_code$$

Class C *pio_op* values are encoded with the following bit patterns: 11110000 - 11111111.

The following tables list the opcodes for each class and their encoded values, along with the corresponding operation definitions, including an indication of how the *operand* value is used in each case. In these tables, *addr* means the contents of the location referred to by the selected addressing mode and register, *reg(operand)* means a direct register value where the low-order 3 bits of *operand* are used to select the register number (unless otherwise specified, the remaining bits must be zero), *PIO(operand)* means a location on the PIO device where *operand* indicates the desired PIO offset, and *reg* means the selected register contents.

Table 3-2 Class A PIO Operation Codes

Operation Code	Value	Operation
UDI_PIO_IN	0x00	<i>addr</i> <- PIO(<i>operand</i>)
UDI_PIO_OUT	0x20	PIO(<i>operand</i>) <- <i>addr</i>
UDI_PIO_LOAD	0x40	<i>reg(operand)</i> <- <i>addr</i>
UDI_PIO_STORE	0x60	<i>addr</i> <- <i>reg(operand)</i>

Table 3-3 Class B PIO Operation Codes

Operation Code	Value	Operation
UDI_PIO_LOAD_IMM	0x80	<i>reg</i> <- <i>operand</i> (multi-part)
UDI_PIO_CSKIP	0x88	Conditional skip. The value in <i>reg</i> is compared with zero. The <i>operand</i> value selects a condition code from Table 3-5 to determine the type of comparison. If the condition is TRUE, the following trans list element will be skipped.
UDI_PIO_IN_IND	0x90	<i>reg</i> <- PIO(<i>reg(operand)</i>)
UDI_PIO_OUT_IND	0x98	PIO(<i>reg(operand)</i>) <- <i>reg</i>

Table 3-3 Class B PIO Operation Codes

Operation Code	Value	Operation
UDI_PIO_SHIFT_LEFT	0xA0	reg <- reg << operand
UDI_PIO_SHIFT_RIGHT	0xA8	reg <- reg >> operand
UDI_PIO_AND	0xB0	reg <- reg & reg(operand)
UDI_PIO_AND_IMM	0xB8	reg <- reg & operand (zero-extended)
UDI_PIO_OR	0xC0	reg <- reg reg(operand)
UDI_PIO_OR_IMM	0xC8	reg <- reg operand (zero-extended)
UDI_PIO_XOR	0xD0	reg <- reg ^ reg(operand)
UDI_PIO_ADD	0xD8	reg <- reg + reg(operand)
UDI_PIO_ADD_IMM	0xE0	reg <- reg + operand (sign-extended)
UDI_PIO_SUB	0xE8	reg <- reg - reg(operand)

Table 3-4 Class C PIO Operation Codes

Operation Code	Value	Operation
UDI_PIO_BRANCH	0xF0	Unconditional branch to the UDI_PIO_LABEL with the same operand value. tran_size is unused and must be set to zero.
UDI_PIO_LABEL	0xF1	Destination of a UDI_PIO_BRANCH. tran_size is unused and must be set to zero.
UDI_PIO_REP_IN_IND	0xF2	Repeated indirect PIO input transactions. operand is encoded according to the UDI_PIO_REP_ARGS macro described below.
UDI_PIO_REP_OUT_IND	0xF3	Repeated indirect PIO output transactions. operand is encoded according to the UDI_PIO_REP_ARGS macro described below.
UDI_PIO_DELAY	0xF4	Delay for operand microseconds during a PIO wait loop.
UDI_PIO_BARRIER	0xF5	Place an ordering barrier between PIO transactions. tran_size is unused and must be set to zero. See below for details on the behavior of this operation and its use of operand .
UDI_PIO_SYNC	0xF6	Synchronize with respect to outstanding PIO transactions. See below for details on the behavior of this operation and its use of tran_size and operand .
UDI_PIO_SYNC_OUT	0xF7	Synchronize with respect to outstanding PIO output transactions. See below for details on the behavior of this operation and its use of tran_size and operand .
UDI_PIO_END	0xFF	Terminate processing. result_code <- operand tran_size must be UDI_PIO_2BYTE.

Detailed description of Addressing Modes:

UDI_PIO_DIRECT — In this mode, the contents of the selected register, *reg*, are used for the *addr* value itself. In this case, *addr* is unaffected by any stride values from **UDI_PIO_REP_IN_IND** or **UDI_PIO_REP_OUT_IND** operations.

UDI_PIO_SCRATCH — In this mode, the contents of the selected register, *reg*, are used as an offset into the scratch space of the control block passed to *udi_pio_trans*. The scratch space data bytes starting at this offset and extending for the selected transaction size are used as the *addr* value. Only the low-order 32 bits of the register value are used; any higher order bits are ignored; smaller values previously loaded into the register are zero-extended to 32 bits. The offset must be a multiple of the transaction size.

UDI_PIO_BUF — In this mode, the contents of the selected register, *reg*, are used as an offset into the valid data area of the *udi_buf_t* buffer passed to *udi_pio_trans*. The buffer data bytes starting at this offset and extending for the selected transaction size are used as the *addr* value. Only the low-order 32 bits of the register value are used; any higher order bits are ignored; smaller values previously loaded into the register are zero-extended to 32 bits. The offset must be a multiple of the transaction size. All offsets accessed with **UDI_PIO_BUF** must be less than the buffer's **buf_size** value at the time *udi_pio_trans* was called.

UDI_PIO_MEM — In this mode, the contents of the selected register, *reg*, are used as an offset into the auxiliary memory block pointed to by the **mem_ptr** argument passed to *udi_pio_trans*. The memory block data bytes starting at this offset and extending for the selected transaction size are used as the *addr* value. Only the low-order 32 bits of the register value are used; any higher order bits are ignored; smaller values previously loaded into the register are zero-extended to 32 bits. The offset must be a multiple of the transaction size.

When PIO device registers/memory are used as the source (**UDI_PIO_IN**, **UDI_PIO_IN_IND**, or **UDI_PIO_REP_IN_IND**) or destination (**UDI_PIO_OUT**, **UDI_PIO_OUT_IND**, or **UDI_PIO_REP_OUT_IND**) of an operation, the specified PIO offset is interpreted relative to the register set and base offset specified through arguments to *udi_pio_map* and passed to *udi_pio_trans* via the **pio_handle** argument. **UDI_PIO_IN** and **UDI_PIO_OUT** only support PIO offsets up to 65535. **UDI_PIO_IN_IND**, **UDI_PIO_OUT_IND**, **UDI_PIO_REP_IN_IND** and **UDI_PIO_REP_OUT_IND** use a register value to provide the PIO offset,

allowing for indirections and larger offsets; the 32 low-order bits of the register value are used for the PIO offset; any higher order bits are ignored; smaller values previously loaded into the register are zero extended to 32 bits.

UDI_PIO_LOAD_IMM:

The immediate load operation takes a value directly from the **operand** portion of a transaction descriptor and loads it into the selected temporary register. The transaction size for this operation must be at least 2 bytes.

If UDI_PIO_LOAD_IMM is used with a transaction size greater than 2 bytes, multiple trans list elements are used to hold all the bytes of the immediate value. Starting from the least significant 2 bytes, pairs of bytes are loaded from each successive **operand** value. The **pio_op** and **tran_size** members of the extra trans list elements are unused. This allows arbitrarily long immediate values to be loaded into temporary registers (up to the maximum transaction size of 32 bytes).

Shift operations:

UDI_PIO_SHIFT_LEFT and UDI_PIO_SHIFT_RIGHT use the **operand** value as a shift count for shifting the entire transaction-size byte register value. The shift count must be from 1 to 32. Any bits shifted out of the low order transaction-size bytes will be discarded. All bits shifted in to the low order transaction-size bytes will be zero. Any previous more significant bytes shall effectively be zeroed.

Repeat operations:

A UDI_PIO_REPEAT_IN_IND or UDI_PIO_REPEAT_OUT_IND operation can be used to repeat a basic PIO transaction up to $2^{32}-1$ times. The basic PIO transaction is like a UDI_PIO_IN_IND or UDI_PIO_OUT_IND, in that it performs a PIO transaction between a PIO offset indicated in one register and a permanent memory location addressed via another register. A third register holds the repeat count, which must be from zero to $2^{32}-1$.

The **operand** value for a repeat operation holds a number of parameters, initialized via the UDI_PIO_REPEAT_ARGS macro. See UDI_PIO_REPEAT_ARGS on page 3-22 for more details on repeat operations.

Branching operations:

UDI_PIO_BRANCH is an unconditional branch operation. The **operand** value is used to find the next trans list element to execute instead of continuing in order. The operand value is matched against UDI_PIO_LABEL trans list elements; execution continues after the UDI_PIO_LABEL with the same **operand**

value as the UDI_PIO_BRANCH. It is illegal to have two UDI_PIO_LABEL elements with the same **operand** value in one trans list. It is illegal to have a UDI_PIO_BRANCH that does not have a corresponding UDI_PIO_LABEL.

If a UDI_PIO_LABEL operation is executed sequentially, it acts as a no-op.

Conditional operation:

UDI_PIO_CSKIP is a conditional skip operation. It compares a register value with zero and skips the following operation if the comparison meets the required conditions. The **operand** specifies the condition code, according to the following table:

Table 3-5 PIO Condition Codes

Mnemonic	Value	Condition Description
UDI_PIO_Z	0	reg == 0
UDI_PIO_NZ	1	reg != 0
UDI_PIO_NEG	2	reg < 0 [signed]
UDI_PIO_NNEG	3	reg >= 0 [signed]

Busy-wait operation:

A UDI_PIO_DELAY pauses execution of the trans list for a specified time. Transaction lists that loop waiting for device events (such as ready bits being set in status registers) must include a UDI_PIO_DELAY in the loop to avoid consuming excessive CPU resources. The **operand** value provides the desired delay time in microseconds. This is a **hint** to the environment, and must not be used to perform precise timings.

Synchronization operations:

The synchronization operations (UDI_PIO_BARRIER, UDI_PIO_SYNC and UDI_PIO_SYNC_OUT) do not actually transfer any new data. Instead, they make sure that any PIO transactions that have already been generated are made visible to the device before any subsequent PIO transactions to/from the same device register set become visible (thus acting as a “barrier”). These operations only affect accesses to the device or memory referenced by the PIO handle, not buffer or driver memory.

The **operand** value for UDI_PIO_BARRIER must be either 0 or UDI_PIO_OUT. If UDI_PIO_OUT, the barrier is only guaranteed to act with respect to device output transactions; this may be faster than a full barrier.

The UDI_PIO_SYNC and UDI_PIO_SYNC_OUT operations provide even stronger synchronization: they also wait for the affected transactions to reach the device. UDI_PIO_SYNC affects all PIO transactions; UDI_PIO_SYNC_OUT is only required to affect output transactions. **operand** and **tran_size** must be set to a PIO offset range of the device that causes no side effects if input from. On some platforms, PIO transactions can only be synchronized by performing an additional PIO input transaction. In such cases, the environment will read $2^{\text{tran_size}}$ bytes from offset **operand** of the device (and discard the result).

Synchronization operations are not needed if UDI_PIO_STRICTORDER was specified in the PIO attributes for the PIO handle. With UDI_PIO_STRICTORDER, each basic transaction that accesses the PIO device is followed by an implicit UDI_PIO_SYNC.

Upon completion of a udi_pio_trans sequence, via execution of a UDI_PIO_END operation, an implicit UDI_PIO_BARRIER with **operand** of zero is performed.

Environment implementations may choose to impose stronger ordering and synchronization than required by the synchronization operations used, up to and including strict ordering.

UDI_PIO_END:

Upon reaching a UDI_PIO_END, processing of the trans list is terminated and the result code (an arbitrary 8-bit code that is defined by the driver) in the udi_pio_trans_call_t callback is set to the low byte of the **operand** value.

The last element of a PIO trans list must be either a UDI_PIO_END operation or a UDI_PIO_BRANCH.

REFERENCES

udi_pio_map, udi_pio_trans, UDI_PIO_REP_ARGS, udi_buf_t, udi_mem_alloc

NAME	UDI_PIO_REP_ARGS	Parameters for repeated PIO transactions
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> #define \ UDI_PIO_REP_ARGS (\ mode, mem_reg, mem_stride, \ pio_reg, pio_stride, cnt_reg) \ ((mode) (mem_reg) ((mem_stride)<<5) \ ((pio_reg)<<7) ((pio_stride)<<10) \ ((cnt_reg)<<13))</pre>	
ARGUMENTS	<p>mode is the addressing mode used to access the memory operand, according to Table 3-1, “PIO Addressing Modes,” on page 3-15.</p> <p>mem_reg is the number of the register used to access the memory operand, according to Table 3-1.</p> <p>mem_stride is the stride value used to increment the memory offset between repeats.</p> <p>pio_reg is the number of the register used to hold the initial PIO offset. The 32 low-order bits of the register value are used for the PIO offset; any higher order bits are ignored; smaller values previously loaded into the register are zero extended to 32 bits.</p> <p>pio_stride is the stride value used to increment the PIO offset between repeats.</p> <p>cnt_reg is the number of the register used to hold the repeat count. The 32 low-order bits of the register value are used for the count; any higher order bits are ignored; smaller values previously loaded into the register are zero extended to 32 bits.</p>	
DESCRIPTION	<p>The UDI_PIO_REP_ARGS macro is used to construct the operand value for UDI_PIO_REP_IN_IND and UDI_PIO_REP_OUT_IND repeating PIO operations (see udi_pio_trans_t on page 3-13).</p> <p>A repeat operation repeats a basic PIO transaction the number of times indicated by the repeat count from the cnt_reg register. The memory location and PIO offset for the first repetition are determined by mode, mem_reg, and pio_reg. For subsequent repetitions, the memory offset (if mode is not UDI_PIO_DIRECT) and PIO offset are incremented according to the corresponding stride values. The values in the original registers (mem_reg, pio_reg, and cnt_reg) are not affected by stride increments or by repeat count decrements.</p>	

The stride values *mem_stride* and *pio_stride* indicate a (possibly zero) multiple of the target operation's transaction size, according to the following table.

Table 3-6 Stride Values for PIO Repeat Operations

Stride Code	Stride Size in Bytes	Multiple of Transaction Size
0	0	0
1	$2^{\text{tran_size}}$	1
2	$2^{(\text{tran_size}+1)}$	2
3	$2^{(\text{tran_size}+2)}$	4

Some examples of how to use stride values are shown in the following table.

Table 3-7 Example Uses of PIO Stride Parameters

Function	mem_stride	pio_stride
Copy array to/from a contiguous range of device memory	1	1
Copy array to/from single device register	1	0
Fill range of device memory w/single value	0	1
Fill lower 2 bytes of a series of 4-byte aligned device registers (<i>tran_size</i> = 2) with a single value	0	2
Fill lower 2 bytes of a series of 4-byte aligned device registers (<i>tran_size</i> = 2) from an array of 2-byte values	1	2

NAME	udi_pio_trans <i>Generate PIO transactions</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_pio_trans (udi_pio_trans_call_t *callback, udi_cb_t *gcb, udi_pio_handle_t pio_handle, udi_buf_t *buf, void *mem_ptr); typedef void udi_pio_trans_call_t (udi_cb_t *gcb, udi_buf_t *new_buf, udi_status_t status, udi_ubit16_t result);</pre>
ARGUMENTS	<p>callback, gcb are standard arguments described in the “Asynchronous Service Calls” section of “Standard Calling Sequences” in the UDI Core Specification.</p> <p>pio_handle is the PIO handle associated with the desired device register set and transaction list, previously acquired from <code>udi_pio_map</code>.</p> <p>buf is a pointer to a data buffer that can be accessed by the PIO transactions using <code>UDI_PIO_BUF</code>. If buf is set to <code>NULL</code>, <code>UDI_PIO_BUF</code> transactions are illegal for this call.</p> <p>mem_ptr is a pointer to an auxiliary memory block that can be accessed by PIO transactions that specify <code>UDI_PIO_MEM</code>. If mem_ptr is set to <code>NULL</code>, such transactions are illegal for this call.</p> <p>new_buf is a possibly new UDI buffer that the driver must use in place of the original buf after the callback is called, since the environment may need to re-allocate the buffer to accommodate new data added to the buffer as a result of this service call.</p> <p>status is a UDI status code. If the PIO transactions all completed without error, it will be set to <code>UDI_OK</code>; otherwise it will be set to <code>UDI_STAT_HW_PROBLEM</code>. If not <code>UDI_OK</code>, the status code may already contain a correlate value.</p> <p>result is a result code set by a <code>UDI_PIO_END</code> transaction.</p>
DESCRIPTION	<p><code>udi_pio_trans</code> performs a series of one or more PIO transactions to the area of a device indicated by the PIO handle, pio_handle. The list of (possibly repeated) transactions to perform is given by the trans_list list which was associated with the PIO handle at the time of <code>udi_pio_map</code>. See <code>udi_pio_trans_t</code> on page 3-13 for a detailed description of PIO transactions.</p>

In order to prevent race conditions between multiple pending `udi_pio_trans` calls (possibly initiated from multiple regions), the environment will serialize concurrent `udi_pio_trans` operations to the same register set of the same device no matter which driver or region invoked `udi_pio_trans`. Each transaction list passed to `udi_pio_trans` is guaranteed to be processed without interruption, but separate lists may be processed in any order, regardless of the order of the calls to `udi_pio_trans`, except for calls made from the same region, which will be processed in order.

If the PIO attributes associated with `pio_handle` indicate that the endianness of the device's data structures are different from the driver's endianness, `udi_pio_trans` will automatically perform necessary byte-swapping. The driver must not do the byte-swapping itself, as `udi_pio_trans` may be able to take advantage of special hardware support for byte-swapped PIO transactions.

Pacing delays and other attributes encoded in `pio_handle`, such as ordering and atomicity requirements, will be applied to each individual I/O operation.

Once the last transaction has been executed, including any pacing delay, the callback routine will be called. An implicit `UDI_PIO_BARRIER` will be executed at the end of the transaction list, but if `UDI_PIO_SYNC` or `UDI_PIO_SYNC_OUT` are required they must be explicitly included.

WARNINGS

Use of the `mem_ptr` parameter must conform to the rules described in "Using Memory Pointers with Asynchronous Service Calls" on page 5-2 of the UDI Core Specification.

STATUS VALUES

A UDI status code indicating the success or failure of the PIO transactions. If a hardware error, such as a bus timeout or parity error, is detected in conjunction with a PIO transaction, execution of the PIO transaction list is terminated and status is set to `UDI_STAT_HW_PROBLEM`. Transactions following the one that caused the error may or may not be executed.

Environments are not required to detect errors during `udi_pio_trans`; in such environments, the effect of a PIO hardware error is indeterminate and may include driver or system termination. If a driver expects that the PIO might reasonably fail (e.g. because the device is not present) it must use `udi_pio_probe` instead of `udi_pio_trans`.

REFERENCES

`udi_pio_map`, `udi_pio_trans_t`, `udi_pio_probe`

NAME	udi_pio_probe	<i>Probe a PIO device that might not be present</i>
SYNOPSIS	<pre> #include <udi.h> #include <udi_physio.h> void udi_pio_probe (udi_pio_probe_call_t *callback, udi_cb_t *gcb, udi_pio_handle_t pio_handle, void *mem_ptr, udi_ubit32_t pio_offset, udi_ubit8_t tran_size, udi_ubit8_t direction); typedef udi_pio_probe_call_t (udi_cb_t *gcb, udi_status_t status); /* Values for direction */ #define UDI_PIO_IN 0x00 #define UDI_PIO_OUT 0x20 </pre>	
ARGUMENTS	<p>callback, gcb are standard arguments described in the “Asynchronous Service Calls” section of “Standard Calling Sequences” in the UDI Core Specification.</p> <p>pio_handle is a PIO handle previously allocated via <code>udi_pio_map</code>.</p> <p>mem_ptr is a pointer to a memory location which is the source or destination of the PIO transaction.</p> <p>pio_offset is the offset into the mapped device at which to perform the transaction. There are no alignment requirements on <code>pio_offset</code>.</p> <p>tran_size is the power-of-2 size of the PIO transaction. The actual size is $2^{\text{tran_size}}$ bytes. The tran_size value must be from 0 to 5, corresponding to a transaction size of 1 to 32 bytes (256 bits). The mnemonic constants <code>UDI_PIO_1BYTE</code> through <code>UDI_PIO_32BYTE</code> are available for use in setting tran_size.</p> <p>direction is exactly one of <code>UDI_PIO_IN</code> or <code>UDI_PIO_OUT</code>, to indicate the direction of transfer.</p> <p>status indicates the success or failure of the PIO probe attempt.</p>	
DESCRIPTION	<p>This service call attempts to perform a PIO access to a device that may or may not be present. It will do this in such a way that there will be no adverse side effects (such as driver or system aborts) if there is in fact no device present. If a different device is present at the same or overlapping location then it might</p>	

be accessed instead of the intended device, causing an indeterminate effect on the state of that device; environments are not required to prevent this from happening.

The trans list associated with **pio_handle** is ignored; instead, a single basic PIO transaction of 2^{tran_size} bytes is performed, in the direction indicated by **direction**.

If a driver's device might not have been reliably enumerated and might not be actually present, `udi_pio_probe` must be used to determine if the device is present before using it for normal operation. Otherwise, `udi_pio_probe` should not be used. It is likely to be slower than `udi_pio_trans`, possibly affecting overall system performance as well. In some environments, use of `udi_pio_probe` may force the driver to be run only during certain stages of system initialization; such an environment may require system re-initialization in order to use a driver that uses `udi_pio_probe`.

In order for a driver to use `udi_pio_probe`, its region's "pio_probe" region attribute must be set to "yes". Some environments may refuse to use drivers with "pio_probe" set to "yes". See Section 1.4, "Extensions to Static Driver Properties," on page 1-2 for details.

Transactions using `udi_pio_probe` are not guaranteed to be atomic.

If a driver determines that its device is not present, it must log an error using `udi_log_write` and respond to the Management Agent with a status of `UDI_STAT_NOT_RESPONDING` via `udi_bind_to_parent_ack`.

WARNINGS

Use of the **mem_ptr** parameter must conform to the rules described in "Using Memory Pointers with Asynchronous Service Calls" on page 5-2 of the UDI Core Specification.

STATUS VALUES

`UDI_OK`

`UDI_STAT_HW_PROBLEM` – the device failed to respond to the PIO access. Some environments may return "garbage" data in such cases, rather than indicating an error with `UDI_STAT_HW_PROBLEM`.

REFERENCES

`udi_pio_trans`

3.4 PIO Handle Transferability

Since PIO handles are transferable between regions, they need to be represented in layout specifiers for use by metalanguages and by drivers that specify inline layouts. This section defines an extension to the data layout specifier defined in the UDI Core Specification.

NAME	udi_layout_t (PIO)	<i>Data layout specifier for PIO</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> typedef udi_ubit8_t udi_layout_t; /* PIO Handle Layout Element Type Code */ #define UDI_DL_PIO_HANDLE_T 200</pre>	
DESCRIPTION	<p>This page lists additional layout specifier codes that can be used with the <code>udi_layout_t</code> type, defined in the UDI Core Specification, to specify PIO-related data layouts.</p> <p>A data layout specifier consists of an array of one or more <code>udi_layout_t</code> layout elements. Each element contains a type code indicating one of the UDI data types that can be passed into a channel operation, either as a field in the control block or as an additional parameter. Each successive element of the array represents successive offsets within the described structure, with padding automatically inserted for alignment purposes as if the specified data types had appeared in a C <code>struct</code> declaration.</p> <p>A <code>UDI_DL_PIO_HANDLE_T</code> layout element represents a PIO handle, of type <code>udi_pio_handle_t</code>, which may be <code>UDI_NULL_PIO_HANDLE</code>.</p>	
REFERENCES	<code>udi_layout_t</code>	



UDI Physical I/O Specification

Section 2: Bus Bridge Metalanguage



Bus Bridge Metalanguage

4

4.1 Overview

This chapter details the channel operations and their parameters for the Bus Bridge Metalanguage, which allow a driver for a direct-attached hardware device to communicate with the driver for its parent bus bridge. The device may be a leaf device or another bridge. For purposes of discussion we always refer to the child (furthest from the processor) as the device and the parent as the bridge, and we'll refer to the device driver and the bridge driver.

Each subsection defines the channel operation calls, control block type declarations, the rationale for the operation's existence, constraints and guidelines for the use of each operation, and error conditions that can occur. (The common errors described earlier are not repeated.)

The Bus Bridge Metalanguage operations are grouped into four roles. For each role there's a corresponding channel ops vector, which is registered via an associated `udi_ops_init_t` structure. The Bus Bridge Metalanguage bridge role supports binding/unbinding and interrupt registration operations invoked on a parent bridge object by a child device driver. The Bus Bridge Metalanguage device role supports binding/unbinding and interrupt registration operations invoked on a child device object by a parent bridge driver. The Bus Bridge Metalanguage interrupt dispatcher role supports interrupt acknowledgment operations invoked on a parent bridge object by a child device driver. The Bus Bridge Metalanguage interrupt handler role supports interrupt event operations invoked on a child device object by a parent bridge driver.

In the Bus Bridge Metalanguage, since each control block type is designed to be used across a group of related operations, a separate control block group is defined per individual type of control block. See **`udi_cb_init_t`** on page 10-11 of the UDI Core Specification for additional information.

Bus Bridge Metalanguage operations are not abortable with `udi_channel_op_abort`.

4.1.1 Versioning

All functions and structures defined in this chapter are part of the “`udi_bridge`” interface, currently at version “0x100”. A driver that conforms to and uses the Bus Bridge Metalanguage of the UDI Physical I/O Specification, Version 1.0, must include the following declaration in its `udiprops.txt` file (see Chapter 31, “*Static Driver Properties*”, of the UDI Core Specification):

```
requires udi_bridge 0x100
```

Compile-time versioning and header files for the Bus Bridge Metalanguage are covered by the general requirements for the UDI Physical I/O Specification defined in Section 1.2, “General Requirements”.

A portable implementation of the Bus Bridge Metalanguage must include a corresponding “provides” declaration in its `udiprops.txt` file and conform to the same compile-time versioning and header file requirements as for drivers.

4.2 Binding/Unbinding Operations

These operations are used during the binding and unbinding of a device driver to a bridge driver.

NAME	udi_bus_device_ops_t <i>Device driver entry point ops vector</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> typedef struct { udi_channel_event_ind_op_t *channel_event_ind_op; udi_bus_bind_ack_op_t *bus_bind_ack_op; udi_bus_unbind_ack_op_t *bus_unbind_ack_op; udi_intr_attach_ack_op_t *intr_attach_ack_op; udi_intr_detach_ack_op_t *intr_detach_ack_op; } udi_bus_device_ops_t; /* Bus Device Ops Vector Number */ #define UDI_BUS_DEVICE_OPS_NUM 1</pre>
DESCRIPTION	A driver using the Physical I/O services will specify the <code>udi_bus_device_ops_t</code> as part of its <code>udi_init_info</code> in order to register its entry points for managing bus bindings and handling interrupts (as opposed to dispatching interrupts).
REFERENCES	<code>udi_init_info</code> , <code>udi_ops_init_t</code> , <code>udi_bus_bridge_ops_t</code>
EXAMPLE	<p>The driver's initialization structure definitions might include the following:</p> <pre>#define MY_CHILD_OPS 10 /* Ops for my child role */ #define MY_BRIDGE_OPS 11 /* Ops for my bridge role */ #define MY_BUS_META 1 /* Meta index for Bus Bridge Metalanguage */ static const udi_bus_device_ops_t ddd_bus_device_ops = { ddd_bus_device_channel_event_ind, ddd_bus_bind_ack, ddd_bus_unbind_ack, ddd_intr_attach_ack, ddd_intr_detach_ack }; ... static const udi_ops_init_t ddd_ops_init_list[] = { { MY_DEVICE_OPS, MY_BUS_META, UDI_BUS_DEVICE_OPS_NUM, 0, /* chan_context_size */ (udi_ops_vector_t *)&ddd_bus_device_ops }, { 0 } };</pre>

NAME	udi_bus_bridge_ops_t <i>Bridge driver entry point ops vector</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> typedef struct { udi_channel_event_ind_op_t *channel_event_ind_op; udi_bus_bind_req_op_t *bus_bind_req_op; udi_bus_unbind_req_op_t *bus_unbind_req_op; udi_intr_attach_req_op_t *intr_attach_req_op; udi_intr_detach_req_op_t *intr_detach_req_op; } udi_bus_bridge_ops_t; /* Bus Bridge Ops Vector Number */ #define UDI_BUS_BRIDGE_OPS_NUM 2</pre>
DESCRIPTION	A driver using the Physical I/O services and which acts as an “interrupt dispatcher” (as opposed to an “interrupt handler”) will specify the <code>udi_bus_bridge_ops_t</code> structure as part of its <code>udi_init_info</code> to register its entry points for managing bus bindings and interrupts attachments/detachments.
REFERENCES	<code>udi_init_info</code> , <code>udi_ops_init_t</code> , <code>udi_bus_device_ops_t</code>
EXAMPLE	<p>The driver’s initialization structure definitions might include the following:</p> <pre>#define MY_DEVICE_OPS 10 /* Ops for my device role */ #define MY_BRIDGE_OPS 11 /* Ops for my bridge role */ #define MY_BUS_META 1 /* Meta index for Bus Bridge Metalanguage */ static const udi_bus_bridge_ops_t ddd_bus_bridge_ops = { ddd_bus_bridge_channel_event_ind, ddd_bus_bind_req, ddd_bus_unbind_req, ddd_intr_attach_req, ddd_intr_detach_req }; ... static const udi_ops_init_t ddd_ops_init_list[] = { { MY_DEVICE_OPS, MY_BUS_META, UDI_BUS_DEVICE_OPS_NUM, 0, /* chan_context_size */ (udi_ops_vector_t *)&ddd_bus_device_ops }, { MY_BRIDGE_OPS, MY_BUS_META, UDI_BUS_BRIDGE_OPS_NUM, 0, /* chan_context_size */ (udi_ops_vector_t *)&ddd_bus_bridge_ops }, ... };</pre>

NAME	udi_bus_bind_cb_t	<i>Control block for bus bridge binding operations</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> typedef struct { udi_cb_t <i>gcb</i>; } udi_bus_bind_cb_t; /* Bus Bind Control Block Group Number */ #define UDI_BUS_BIND_CB_NUM 1</pre>	
MEMBERS	<i>gcb</i>	is a generic control block header, which includes a pointer to the scratch space associated with this control block. The driver may use the scratch space while it owns the control block, but the values are not guaranteed to persist across channel operations.
DESCRIPTION	<p>The bus bind control block is used between the bus device driver and the bus bridge driver to complete binding and unbinding over the bind channel.</p> <p>This control block must be declared by specifying the control block index value UDI_BUS_BIND_CB_NUM in a udi_cb_init_t in the driver's udi_init_info.</p> <p>The bus device driver obtains the udi_bus_bind_cb_t structure to use with the udi_bus_bind_req or udi_bus_unbind_req operation by calling udi_cb_alloc with a <i>cb_idx</i> that has been defined for the UDI_BUS_BIND_CB_NUM control block.</p>	
REFERENCES	udi_bus_bind_cb_t, udi_cb_alloc, udi_bus_bind_req, udi_bus_bind_ack, udi_bus_unbind_req, udi_bus_unbind_ack	

NAME	udi_bus_bind_req	<i>Request a binding to a bridge driver</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_bus_bind_req (udi_bus_bind_cb_t *cb);</pre>	
ARGUMENTS	cb	is a pointer to a bus bind control block.
TARGET CHANNEL	The target channel for this operation is the bind channel connecting a device driver with its parent bus bridge driver.	
DESCRIPTION	<p>A device driver uses this operation to bind to its parent bus bridge driver.</p> <p>The device driver must prepare for the <code>udi_bus_bind_req</code> operation by allocating a bus bind control block (calling <code>udi_cb_alloc</code> with a <code>cb_idx</code> corresponding to the <code>UDI_BUS_BIND_CB_NUM</code> control block type).</p> <p>Next, the device driver sends the bus bind control block to the bridge driver with a <code>udi_bus_bind_req</code> operation.</p>	
REFERENCES	<code>udi_bus_bind_cb_t</code> , <code>udi_bus_bind_ack</code> , <code>udi_bus_unbind_req</code>	

NAME	udi_bus_bind_ack	<i>Acknowledge a bus bridge binding</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_bus_bind_ack (udi_bus_bind_cb_t *cb, udi_ubit8_t preferred_endianness, udi_status_t status); /* Values for preferred_endianness */ #define UDI_DMA_BIG_ENDIAN (1U<<5) #define UDI_DMA_LITTLE_ENDIAN (1U<<6) #define UDI_DMA_ANY_ENDIAN (1U<<0)</pre>	
ARGUMENTS	<p>cb is a pointer to a bus bind control block.</p> <p>preferred_endianness indicates the device endianness which works most effectively with the bridges in this path. It may be set to one of the following values:</p> <p style="margin-left: 40px;">UDI_DMA_LITTLE_ENDIAN UDI_DMA_BIG_ENDIAN UDI_DMA_ANY_ENDIAN</p> <p>status indicates whether or not the binding was successful.</p>	
TARGET CHANNEL	The target channel for this operation is the bind channel connecting a bus bridge driver with one of its child device drivers.	
DESCRIPTION	<p>The <code>udi_bus_bind_ack</code> operation is used by a bridge driver to acknowledge binding with a child device driver (or failure to do so, as indicated by <code>status</code>), as requested by a <code>udi_bus_bind_req</code> operation.</p> <p>Some devices are bi-endian; that is, they can be placed in either a little-endian mode or a big-endian mode. preferred_endianness provides a hint to drivers for such devices, as to which endianness is likely to be most efficient. If this is set to <code>UDI_DMA_ANY_ENDIAN</code>, at least one interposed bridge is bi-endian, so either endianness can be supported without significant additional cost (i.e. without software byte swapping).</p> <p>Drivers for fixed-endianness devices can ignore preferred_endianness.</p>	
STATUS VALUES	UDI_STAT_CANNOT_BIND	
WARNINGS	The control block must be the same control block as passed to the driver in the corresponding <code>udi_bus_bind_req</code> operation.	
REFERENCES	<code>udi_bus_bind_cb_t</code> , <code>udi_bus_bind_req</code> , <code>udi_channel_close</code>	

NAME	udi_bus_unbind_req	<i>Request a bridge driver unbinding (child to bridge)</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_bus_unbind_req (udi_bus_bind_cb_t *cb);</pre>	
ARGUMENTS	cb	is a pointer to a bus bind control block.
TARGET CHANNEL	The target channel for this operation is the bind channel connecting a device driver with its parent bus bridge driver.	
DESCRIPTION	<p>A device driver uses this operation to unbind from its parent bus bridge driver.</p> <p>The device driver must prepare for the <code>udi_bus_unbind_req</code> operation by allocating a bus bind control block (calling <code>udi_cb_alloc</code> with a cb_idx that indicates a <code>udi_bus_bind_cb_t</code>).</p> <p>Next, the device driver sends the bus unbind control block to the bridge driver with a <code>udi_bus_unbind_req</code> operation.</p>	
REFERENCES	<code>udi_bus_bind_cb_t</code> , <code>udi_bus_unbind_ack</code>	

NAME	udi_bus_unbind_ack	<i>Acknowledge a bus bridge unbinding</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_bus_bind_ack (udi_bus_bind_cb_t *cb);</pre>	
ARGUMENTS	cb is a pointer to a bus bind control block.	
TARGET CHANNEL	The target channel for this operation is the bind channel connecting a device driver with its parent bus bridge driver.	
DESCRIPTION	<p>The <code>udi_bus_unbind_ack</code> operation is used by a bridge driver to acknowledge an unbinding with a child device driver as requested by a <code>udi_bus_unbind_req</code> operation.</p> <p>There is no status parameter associated with this operation; the bridge driver is expected to always be able to handle the unbind request and respond appropriately. If, for example, the bridge driver were to receive an unbind from a child without having first received a bind (or two unbinds in a row from the child), the bridge driver may log this condition but must always respond with this acknowledgment.</p>	
WARNINGS	The control block must be the same control block as passed to the driver in the corresponding <code>udi_bus_unbind_req</code> operation.	
REFERENCES	<code>udi_bus_bind_cb_t</code> , <code>udi_bus_unbind_req</code> , <code>udi_channel_close</code>	

4.3 Interrupt Registration Operations

These operations are used to register and de-register interrupt handlers for interrupt sources routed through the bus bridge.

NAME	udi_intr_attach_cb_t	<i>Control block for interrupt registration operations</i>
SYNOPSIS	<pre> #include <udi.h> #include <udi_physio.h> typedef struct { udi_cb_t <i>gcb</i>; udi_index_t <i>interrupt_idx</i>; udi_ubit8_t <i>max_event_pend</i>; udi_ubit16_t <i>event_buf_size</i>; udi_pio_handle_t <i>preprocessing_handle</i>; } udi_intr_attach_cb_t; /* Bridge Attach Control Block Group Number */ #define UDI_BUS_INTR_ATTACH_CB_NUM 2 </pre>	
MEMBERS	<p><i>gcb</i> is a generic control block header, which includes a pointer to the scratch space associated with this control block. The driver may use the scratch space while it owns the control block, but the values are not guaranteed to persist across channel operations.</p> <p><i>interrupt_idx</i> is used to select one of possibly several interrupt sources from the interrupt handler's device that are managed by a particular interrupt dispatcher. Zero indicates the first interrupt source for the device, one indicates the second source, etc. The dispatcher driver will have previously associated the handler's device properties with its channel context during the initial binding sequence.</p> <p>The definition of <i>interrupt_idx</i> is bus and device dependent, and must be determined by the driver via the <i>bus_type</i> instance attribute (see Chapter 16, "<i>Instance Attribute Management</i>" in the UDI Core Specification) in conjunction with the corresponding UDI bus bindings (see "Section 3: Bus Bindings"). For example, if <i>bus_type</i> is "pci" then the settings would be based on the definitions in the <i>UDI PCI Bus Binding Specification</i>.</p> <p><i>max_event_pend</i> is the maximum number of interrupt events the device and driver can have pending at one time. This is a hint to the dispatcher, so it can determine how many event control blocks to allocate for maximum concurrency.</p> <p><i>event_buf_size</i> is the valid data size, in bytes, of buffers to be allocated by the dispatcher and sent to the handler with each event, via <i>udi_intr_event_ind</i>. If <i>event_buf_size</i> is zero, no buffers shall be allocated.</p>	

DESCRIPTION

preprocessing_handle is an optional PIO handle to be used to preprocess interrupts. If not null, the interrupt dispatcher will execute the associated PIO transaction list before delivering the interrupt to the interrupt handler. If null, the handler must run in an interrupt region.

The interrupt attach control block is used between the interrupt handler and the interrupt dispatcher to attach interrupt handler channels (i.e., register an interrupt handler for a particular interrupt source).

This control block must be declared by specifying the control block index value `UDI_BUS_INTR_ATTACH_CB_NUM` in a `udi_cb_init_t` in the driver's `udi_init_info`.

The bus device driver obtains the `udi_intr_attach_cb_t` structure to use with the `udi_intr_attach_req` by calling `udi_cb_alloc` with a ***cb_idx*** that has been associated with `UDI_BUS_INTR_ATTACH_CB_NUM`.

REFERENCES

`udi_intr_attach_req`, `udi_intr_attach_ack`,
`udi_init_info`, `udi_cb_init_t`, `udi_cb_alloc`

NAME	udi_intr_attach_req	<i>Request an interrupt attachment</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_intr_attach_req (udi_intr_attach_cb_t *intr_attach_cb);</pre>	
ARGUMENTS	<i>intr_attach_cb</i> is a pointer to an interrupt attach control block.	
TARGET CHANNEL	The target channel for this operation is the bound bus bridge channel connecting a device driver with its parent bus bridge driver, established during the initial binding between the child device and its bus bridge.	
DESCRIPTION	<p>An interrupt handler driver uses this operation to request an interrupt dispatcher driver to begin accepting interrupts for the handler driver's device and delivering them to the handler. If a device has multiple interrupt sources, the interrupt handler indicates which one to attach by using <i>interrupt_idx</i> in the control block. The handler driver can issue an <code>udi_intr_attach_req</code> for an <i>interrupt_idx</i> that is already attached (i.e., uses the same <i>interrupt_idx</i> as an existing attachment), in which case the <code>udi_intr_attach_req</code> simply changes the interrupt pre-processing for that interrupt source and ignores all other parameters.</p> <p>The handler driver must prepare for the <code>udi_intr_attach_req</code> operation by allocating an interrupt attach control block, filling in the relevant members of the control block, and spawning its end of a new interrupt event channel.</p> <p>The new channel will be used to send interrupt event operations from the interrupt dispatcher to the handler. The handler driver spawns its end by calling <code>udi_channel_spawn</code>, passing in its end of the bus bridge channel and setting spawn index equal to the <i>interrupt_idx</i> for this attachment. Using <i>interrupt_idx</i> for the spawn index ensures that it is unique with respect to any other spawn operations in progress on the same bus bridge channel, allowing the environment use this index to match up the two parts of the spawn operation.</p> <p>Next, the interrupt handler sends the interrupt control block to the dispatcher driver with a <code>udi_intr_attach_req</code> operation.</p> <p>When the interrupt dispatcher driver receives this request, it spawns its end of the new interrupt event channel, calling <code>udi_channel_spawn</code>, passing in its end of the bus bridge channel and using <i>interrupt_idx</i> for the spawn index. The dispatcher then sends a <code>udi_intr_attach_ack</code> operation back on the original channel to the handler driver, which may now reuse the spawn index.</p> <p>The dispatcher driver may begin sending interrupt events over the new channel as soon as it has completed the spawn operation. Thus, the handler may receive interrupt event operations even before it receives the <code>udi_intr_attach_ack</code>.</p>	

If **preprocessing_handle** is not null (use UDI_HANDLE_IS_NULL to test it), events on this channel will be preprocessed. Otherwise, the handler driver must ensure that the handler's end of the new channel is anchored in an interrupt region (i.e. a region with the "interrupt" attribute set in the driver's static driver properties; see Section 1.4, "Extensions to Static Driver Properties," on page 1-2).

Some of the advantages of interrupt preprocessing are that the interrupt handler does not need to be in an interrupt region and that the udi_intr_event_ind handler code is not restricted in the operations it may perform before responding with udi_intr_event_res, unlike in the non-preprocessed case.

Any interrupts for the selected interrupt source that are processed by the dispatcher after receiving a udi_intr_attach_req with a non-null **preprocessing_handle** will be preprocessed. This means that the dispatcher will use udi_pio_trans to execute the specified trans list (associated with **preprocessing_handle**). This trans list must include all operations necessary to de-assert its interrupt (if it was asserted) in all cases.

When the dispatcher calls udi_pio_trans it will set the **buf** argument to refer to the buffer that will be sent to the handler driver along with the udi_intr_event_ind operation. This buffer will have exactly **event_buf_size** bytes of valid data (see **udi_intr_attach_cb_t** on page 4-11), but the data values are unspecified. The trans list can use this buffer to pass data to the handler driver's upper-level handler.

If the bus type supports event info, the dispatcher driver must set the **mem_ptr** argument to point to a memory block containing the event info for this interrupt. Otherwise, **mem_ptr** must be set to NULL.

The preprocessing trans list must not use UDI_PIO_SCRATCH, as the size and contents of the scratch area for interrupt preprocessing are unspecified.

The **result** from the udi_pio_trans call (set with UDI_PIO_END) is used to determine the disposition of the interrupt, as described in **udi_intr_event_cb_t** on page 4-22 and **udi_intr_event_ind** on page 4-24.

The bridge driver must free the **preprocessing_handle** after it has finished with it (as a result of either another udi_intr_attach_req with a new handle or a udi_intr_detach_req), by using udi_pio_unmap.

WARNINGS

The **mem_ptr** value must follow the rules for memory object usage described in Section 5.2.1.1, "Using Memory Pointers with Asynchronous Service Calls," on page 5-2 of the UDI Core Specification.

REFERENCES

udi_intr_attach_cb_t, udi_channel_spawn,
udi_intr_attach_ack, udi_intr_event_ind,
udi_pio_trans, udi_pio_map, udi_pio_unmap

NAME	udi_intr_attach_ack	<i>Acknowledge an interrupt attachment</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_intr_attach_ack (udi_intr_attach_cb_t *intr_attach_cb, udi_status_t status);</pre>	
ARGUMENTS	<p><i>intr_attach_cb</i> is a pointer to an interrupt attach control block.</p> <p><i>status</i> indicates whether or not the attachment was successful.</p>	
TARGET CHANNEL	The target channel for this operation is the bound bus bridge channel connecting a bus bridge driver with one of its child device drivers, established during the initial binding between the child device and its bus bridge.	
PROXIES	udi_intr_attach_ack_unused	<i>Proxy for udi_intr_attach_ack</i>
	<pre>udi_intr_attach_ack_op_t udi_intr_attach_ack_unused;</pre> <p>udi_intr_attach_ack_unused may be used as a device driver's udi_intr_attach_ack entry point if the driver never calls udi_intr_attach_req (and therefore expects to receive no acknowledgements).</p>	
DESCRIPTION	<p>The udi_intr_attach_ack operation is used by an interrupt dispatcher driver to acknowledge attachment of an interrupt handler (or failure to do so, as indicated by status), as requested by a udi_intr_attach_req operation.</p> <p>Before sending the acknowledgment, the interrupt dispatcher driver must spawn and anchor its end of the new interrupt event channel, setting the spawn index equal to the <i>interrupt_idx</i> for this attachment.</p> <p><i>interrupt_idx</i> in the returned control block must be the same as that passed to udi_intr_attach_req.</p> <p>The interrupt handler may free the control block using udi_cb_free or save it to pass to udi_intr_detach_req.</p> <p>Upon failure indication, the handler driver must be sure to close the half-spawned channel, by calling udi_channel_close.</p>	
STATUS VALUES	UDI_STAT_MISTAKEN_IDENTITY – This device has no interrupt source corresponding to the <i>interrupt_idx</i> value in the control block.	
WARNINGS	The control block must be the same control block as passed to the driver in the corresponding udi_intr_attach_req operation.	
REFERENCES	udi_intr_attach_cb_t, udi_intr_attach_req, udi_intr_detach_req, udi_channel_close	

NAME	udi_intr_detach_cb_t	<i>Control block for interrupt detachment operations</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> typedef struct { udi_cb_t <i>gcb</i>; udi_index_t <i>interrupt_idx</i>; } udi_intr_detach_cb_t; /* Bridge Detach Control Block Group Number */ #define UDI_BUS_INTR_DETACH_CB_NUM 3</pre>	
MEMBERS	<p><i>gcb</i> is a generic control block header, which includes a pointer to the scratch space associated with this control block. The driver may use the scratch space while it owns the control block, but the values are not guaranteed to persist across channel operations.</p> <p><i>interrupt_idx</i> is used to select one of possibly several interrupt sources from the interrupt handler's device that are managed by a particular interrupt dispatcher. Zero indicates the first interrupt source for the device, one indicates the second source, etc. The dispatcher driver will have previously associated the handler's device properties with its channel context during the initial binding sequence.</p>	
DESCRIPTION	<p>The interrupt detach control block is used between the interrupt handler and the interrupt dispatcher to detach interrupt handler channels (i.e., deregister an interrupt handler for a particular interrupt source).</p> <p>This control block must be declared by specifying the control block index value UDI_BUS_INTR_DETACH_CB_NUM in a udi_cb_init_t in the driver's udi_init_info.</p> <p>The bus device driver obtains the udi_intr_detach_cb_t structure to use with the udi_intr_detach_req by calling udi_cb_alloc with a <i>cb_idx</i> that has been associated with UDI_BUS_INTR_DETACH_CB_NUM.</p>	
REFERENCES	<pre>udi_intr_detach_req, udi_intr_detach_ack, udi_init_info, udi_cb_init_t, udi_cb_alloc</pre>	

NAME	udi_intr_detach_req	<i>Request an interrupt detachment</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_intr_detach_req (udi_intr_detach_cb_t *intr_detach_cb);</pre>	
ARGUMENTS	<i>intr_detach_cb</i> is a pointer to an interrupt detach control block.	
TARGET CHANNEL	The target channel for this operation is the bound bus bridge channel connecting a device driver with its parent bus bridge driver, established during the initial binding between the child device and its bus bridge.	
DESCRIPTION	<p>When the interrupt handler driver wishes to detach its handler from the interrupt dispatcher, it invokes the <code>udi_intr_detach_req</code> channel operation. A prior attachment must have been completed, as indicated to the interrupt handler driver by receipt of the attach acknowledgment.</p> <p><i>interrupt_idx</i> in the control block must be the same as that passed to <code>udi_intr_attach_req</code> (and returned with <code>udi_intr_attach_ack</code>).</p>	
REFERENCES	<code>udi_intr_detach_cb_t</code> , <code>udi_intr_detach_ack</code>	

NAME	udi_intr_detach_ack	<i>Acknowledge an interrupt detachment</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_intr_detach_ack (udi_intr_detach_cb_t *intr_detach_cb);</pre>	
ARGUMENTS	<i>intr_detach_cb</i> is a pointer to an interrupt detach control block.	
TARGET CHANNEL	The target channel for this operation is the bound bus bridge channel connecting a bus bridge driver with one of its child device drivers, established during the initial binding between the child device and its bus bridge.	
PROXIES	udi_intr_detach_ack_unused	<i>Proxy for udi_intr_detach_ack</i>
	<pre>udi_intr_detach_ack_op_t udi_intr_detach_ack_unused;</pre> <p><i>udi_intr_detach_ack_unused</i> may be used as a device driver's <i>udi_intr_detach_ack</i> entry point if the driver never calls <i>udi_intr_detach_req</i> (and therefore expects to receive no acknowledgements).</p>	
DESCRIPTION	<p>After detaching the interrupt handler and closing the local end of the corresponding interrupt event channel, the interrupt dispatcher driver sends the control block back to the interrupt handler using the channel operation, <i>udi_intr_detach_ack</i>.</p> <p><i>interrupt_idx</i> in the returned control block must be the same as that passed to <i>udi_intr_detach_req</i>.</p> <p>Upon receipt of the <i>udi_intr_detach_ack</i>, the interrupt handler must close its end of the interrupt event channel (if it did not already do so as a result of a <i>udi_channel_event_ind</i>).</p> <p>The interrupt handler may free the control block using <i>udi_cb_free</i> or reuse it for another detachment.</p>	
WARNINGS	The control block must be the same control block as passed to the driver in the corresponding <i>udi_intr_detach_req</i> operation.	
REFERENCES	<i>udi_intr_detach_cb_t</i> , <i>udi_intr_detach_req</i>	

4.4 Interrupt Event Operations

These operations are used to deliver and respond to interrupt events between interrupt dispatchers and interrupt handlers.

Interrupt events may be handled in one of two ways: with interrupt preprocessing, or by using interrupt regions. See **udi_intr_attach_req** on page 4-13 for details on interrupt preprocessing. Any region using `udi_intr_event_ind` or `udi_intr_event_res` without interrupt preprocessing enabled must be an *interrupt region*; i.e. it must have its `type` region attribute set to `interrupt` (see Table 1-1, “Physical I/O Region Attributes,” on page 1-2). The primary region of a driver instance cannot be an interrupt region.

All environment services are available to interrupt handlers, but since interrupt regions can require critical resources to be held while executing, drivers should minimize the amount of time spent executing in interrupt regions.

There is, however, a restriction on the *order* in which service calls and channel operations can be invoked from an interrupt handler, when interrupt preprocessing is not used. For details, see `udi_intr_event_ind` and `udi_intr_event_res`.

NAME	udi_intr_handler_ops_t <i>Interrupt handler ops vector</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> typedef struct { udi_channel_event_ind_op_t *channel_event_ind_op; udi_intr_event_ind_op_t *intr_event_ind_op; } udi_intr_handler_ops_t; /* Interrupt Handler Ops Vector Number */ #define UDI_BUS_INTR_HANDLER_OPS_NUM 3</pre>
DESCRIPTION	A driver which wishes to register for handling interrupts (as opposed to “dispatching” interrupts declares the <code>udi_intr_handler_ops_t</code> structure to define its entry point for receiving interrupt events.
REFERENCES	<code>udi_init_info</code> , <code>udi_ops_init_t</code> , <code>udi_intr_dispatcher_ops_t</code>
EXAMPLE	<p>The driver’s initialization structure definitions might include the following:</p> <pre>#define MY_INTR_HANDLER_OPS 2 /* My interrupt handler ops */ #define MY_BUS_META 1 /* Meta index for the Bus Bridge Metalanguage */ static const udi_intr_handler_ops_t ddd_intr_handler_ops = { ddd_intr_handler_channel_event_ind, ddd_intr_event_ind }; ... static const udi_ops_init_t ddd_ops_init_list[] = { { MY_INTR_HANDLER_OPS, MY_BUS_META, UDI_BUS_INTR_HANDLER_OPS_NUM, 0, /* chan_context_size */ (udi_ops_vector_t *)&ddd_intr_handler_ops }, { 0 } };</pre>

NAME	udi_intr_dispatcher_ops_t <i>Interrupt dispatcher ops vector</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> typedef struct { udi_channel_event_ind_op_t *channel_event_ind_op; udi_intr_event_res_op_t *intr_event_res_op; } udi_intr_dispatcher_ops_t; /* Interrupt Dispatcher Ops Vector Number */ #define UDI_BUS_INTR_DISPATCH_OPS_NUM 4</pre>
DESCRIPTION	A bus driver which delivers interrupt indications uses the <code>udi_intr_dispatcher_ops_t</code> to declare the interface operations for receiving interrupt acknowledgements from the interrupt handler.
REFERENCES	<code>udi_init_info</code> , <code>udi_ops_init_t</code> , <code>udi_intr_handler_ops_t</code>
EXAMPLE	<p>The driver's initialization structure definitions might include the following:</p> <pre>#define MY_INTR_DISP_OPS 2 /* My interrupt dispatcher ops */ #define MY_BUS_META 1 /* Meta index for the Bus Bridge Metalanguage */ static const udi_intr_dispatcher_ops_t ddd_intr_dispatcher_ops = { ddd_intr_channel_event_ind, ddd_intr_event_res }; ... static const udi_ops_init_t ddd_ops_init_list[] = { { MY_INTR_DISP_OPS, MY_BUS_META, UDI_BUS_INTR_DISPATCH_OPS_NUM, 0, /* chan_context_size */ (udi_ops_vector_t *)&ddd_intr_dispatcher_ops }, { 0 } };</pre>

NAME	udi_intr_event_cb_t	<i>Control block for interrupt event ops</i>
SYNOPSIS	<pre> #include <udi.h> #include <udi_physio.h> typedef struct { udi_cb_t <i>gcb</i>; udi_buf_t *<i>event_buf</i>; udi_ubit16_t <i>intr_result</i>; } udi_intr_event_cb_t; /* Flag values for intr_result */ #define UDI_INTR_UNCLAIMED (1U<<15) #define UDI_INTR_NO_EVENT (1U<<14) /* Bus Interrupt Event Control Block Group Number */ #define UDI_BUS_INTR_EVENT_CB_NUM 4 </pre>	
MEMBERS	<p><i>gcb</i> is a generic control block header, which includes a pointer to the scratch space associated with this control block. The driver may use the scratch space while it owns the control block, but the values are not guaranteed to persist across channel operations.</p> <p><i>event_buf</i> is a buffer containing bus-specific or driver-specific piece of extra information delivered from the interrupting device hardware or first-level handler along with the interrupt. This buffer must contain <i>event_buf_size</i> bytes of valid data, as specified during <code>udi_intr_attach_req</code>. If <i>event_buf_size</i> was zero, <i>event_buf</i> must be NULL.</p> <p><i>intr_result</i> is, in the interrupt preprocessing case, set by the dispatcher to the <i>result</i> value from the <code>udi_pio_trans</code> call used to preprocess the interrupt. In the non-preprocessing case, <i>intr_result</i> is set by the interrupt handler to indicate whether or not its device asserted the interrupt.</p> <p>Whether indicated as the <i>result</i> value from <code>udi_pio_trans</code> in the preprocessing case, or passed in <i>intr_result</i> with <code>udi_intr_event_res</code> in the non-preprocessing case, the UDI_INTR_UNCLAIMED flag, if set, indicates that interrupt was not generated by this handler's device, and that the dispatcher should proceed to process any other shared interrupts for this interrupt line.</p> <p>Drivers that cannot tell whether or not their device actually asserted an interrupt, or whose device's interrupts are for some other reason non-sharable, must include a "nonsharable_interrupt" declaration in their static driver properties file (see Section 1.4, "Extensions to Static Driver Properties," on page 1-2), and must not use UDI_INTR_UNCLAIMED.</p>	

DESCRIPTION

The interrupt event control block is used between the interrupt handler and the interrupt dispatcher to deliver and acknowledge interrupt events.

If the interrupt event was pre-processed (as indicated by the initial value of *intr_result* passed to *udi_intr_event_ind*), the *event_buf* will contain any data filled in during the first-level handler by using the device driver's pre-registered PIO transaction list. (See *preprocessing_handle* in *udi_intr_attach_cb_t*.)

Otherwise, *event_buf* will contain bus-type specific event info, as defined in the bus binding specification for the type of bus to which the device is attached. If the size of the event info is greater than *event_buf_size*, only the first *event_buf_size* bytes will be placed in the buffer. If the event info size is less than *event_buf_size*, the remaining bytes will be part of the buffer's valid data range, but their values are unspecified.

If a driver supports a device which might make use of interrupt event info, it must first determine what type of bus the device is plugged into, by looking at its driver instance attributes. It may then interpret the event info using the corresponding bus binding. See Chapter 5, "Introduction to Bus Bindings", for information on the format of interrupt event info.

REFERENCES

udi_intr_attach_cb_t

NAME	udi_intr_event_ind <i>Interrupt event indication</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_intr_event_ind (udi_intr_event_cb_t *intr_event_cb, udi_ubit8_t flags); /* Values for flags */ #define UDI_INTR_MASKING_NOT_REQUIRED (1U<<0)</pre>
ARGUMENTS	<p><i>intr_event_cb</i> is a pointer to an interrupt event control block.</p> <p><i>flags</i> specifies optional flags, described below.</p>
TARGET CHANNEL	<p>The target channel for this operation is the interrupt event channel for this attachment, jointly spawned by the interrupt handler and interrupt dispatcher during interrupt attachment (see <i>udi_intr_attach_req</i>, <i>udi_intr_attach_ack</i>). The interrupt handler driver can use the channel context pointer for this channel to distinguish this event from events for other attached interrupts.</p>
DESCRIPTION	<p>Upon receipt of an interrupt condition that may have been generated by a particular interrupt source, an interrupt dispatcher will prepare to deliver the interrupt to the appropriate handler. If the handler has requested interrupt preprocessing, it will execute the pre-registered PIO trans list (see <i>udi_intr_attach_req</i> on page 4-13) before delivering the interrupt to the handler driver.</p> <p>Since some interrupt handlers change between non-preprocessing and preprocessing modes, the value of <i>intr_result</i> in <i>intr_event_cb</i> indicates whether or not this particular interrupt was pre-processed. The dispatcher must set <i>intr_result</i> to UDI_INTR_UNCLAIMED if the interrupt was not preprocessed; otherwise this bit will never be set.</p> <p>The following applies only to the preprocessing case:</p> <p style="padding-left: 40px;">The <i>result</i> value from the <i>udi_pio_trans</i> call used to preprocess the interrupt is used by the dispatcher to set <i>intr_result</i> in the interrupt event control block. If either the UDI_INTR_UNCLAIMED or the UDI_INTR_NO_EVENT flags are set, no <i>udi_intr_event_ind</i> operation will be sent to the handler driver; otherwise, the dispatcher must set <i>event_buf</i> to the buffer used with the <i>udi_pio_trans</i> call, set <i>intr_result</i> to the <i>result</i> value, and deliver the event to the dispatcher using <i>udi_intr_event_ind</i>.</p>

The following applies only to the non-preprocessing case:

The dispatcher driver must fill in the **event_buf** buffer with the bus type-specific event info for the interrupt, if any, set **intr_result** to UDI_INTR_UNCLAIMED, and deliver the event to the appropriate interrupt handler using the channel operation, `udi_intr_event_ind`.

When the interrupt handler receives an interrupt event indication it must either pass the event on to another level of interrupt handler using another `udi_intr_event_ind` operation or respond back to the dispatcher (after de-asserting the device condition) using `udi_intr_event_res`.

If the interrupt was not preprocessed, the `udi_intr_event_ind` operation must execute in an interrupt region, and before invoking one of the above interrupt event operations, the driver must not invoke any other channel operation or any asynchronous service calls except `udi_pio_trans`, and must not depend on any additional callbacks or channel ops entry points, except the `udi_pio_trans` callback(s), in order to complete the sequence needed to invoke the required interrupt event operation.

If the interrupt was preprocessed, there are no restrictions on the service calls and channel operations that `udi_intr_event_ind` may invoke before invoking `udi_intr_event_ind` or `udi_intr_event_res`, since the interrupt condition was already dismissed in the first-level handler.

The **flags** values are interpreted as follows:

UDI_INTR_MASKING_NOT_REQUIRED - this flag indicates that an interrupt handler that is also an interrupt dispatcher does not need to mask off the interrupt before passing the event on to the next level (because a higher level interrupt dispatcher interprets interrupts as one-shot events—rather than continuous assertion—and will not pass continuous assertion through).

REFERENCES

`udi_intr_event_cb_t`, `udi_intr_attach_req`,
`udi_intr_event_res`

NAME	udi_intr_event_res <i>Acknowledge an interrupt event</i>
SYNOPSIS	<pre>#include <udi.h> #include <udi_physio.h> void udi_intr_event_res (udi_intr_event_cb_t *intr_event_cb);</pre>
ARGUMENTS	<i>intr_event_cb</i> is a pointer to an interrupt event control block.
TARGET_CHANNEL	The target channel for this operation is the interrupt event channel for this attachment, jointly spawned by the interrupt handler and interrupt dispatcher during interrupt attachment (see <i>udi_intr_attach_req</i> , <i>udi_intr_attach_ack</i>). The interrupt dispatcher driver can use the channel context pointer for this channel to distinguish this response from responses for other attached interrupts, or it can key off of the <i>initiator_context</i> in the control block.
DESCRIPTION	<p>An interrupt handler driver uses <i>udi_intr_event_res</i> to send an interrupt event control block back to the dispatcher driver. It should do this as soon as possible after it ensuring that the interrupt is no longer asserted.</p> <p>For a leaf interrupt source, this often means accessing some device-specific register which causes the device to stop asserting the interrupt. Leaf drivers must call <i>udi_intr_event_res</i> before returning from <i>udi_intr_event_ind</i>. Bridge drivers may complete the interrupt handling by passing the <i>udi_intr_event_res</i> on to their parent bridge, or may invoke another interrupt handler with <i>udi_intr_event_ind</i>.</p> <p>For an interrupt handler which is also a dispatcher, <i>udi_intr_event_res</i> should be called immediately if the <i>UDI_INTR_MASKING_NOT_REQUIRED</i> flag is set, or after masking the interrupt if possible. If the interrupt could not be masked, then <i>udi_intr_event_res</i> should be called upon receipt of a <i>udi_intr_event_res</i> operation from the child handler driver.</p> <p>In either case, if not using interrupt preprocessing, the handler driver must set <i>intr_result</i> in the control block to zero if its device asserted the interrupt (or the interrupt is non-sharable as indicated by the “nonsharable_interrupt” declaration) or to <i>UDI_INTR_UNCLAIMED</i> if the device was not the source of the interrupt. When using interrupt preprocessing, <i>intr_result</i> is ignored for <i>udi_intr_event_ind</i>.</p> <p>The interrupt dispatcher may free the control block using <i>udi_cb_free</i> or (more likely) save it to reuse for another event.</p>
WARNINGS	The control block must be the same control block as passed to the driver in the corresponding <i>udi_intr_event_ind</i> operation.
REFERENCES	<i>udi_intr_event_cb_t</i> , <i>udi_intr_attach_ack</i> , <i>udi_intr_event_ind</i>

4.5 Instance Attribute Bindings

One enumeration attribute is defined for all uses of the bus bridge metalanguage: `bus_type`. This attribute, of type `UDI_ATTR_STRING`, is set to the name of the I/O bus type to which the child adapter is connected, as defined by the relevant Bus Binding.

Additional enumeration attributes are specified in each Bus Binding to apply to each bus type.

4.6 Bus Bridge Trace Events

The Bus Bridge Metalanguage defines the use of `UDI_TREVENT_META_SPECIFIC_1` for bus bridge drivers when they are about to call `udi_pio_trans` to execute a preprocessing transaction list.

The `UDI_TREVENT_IO_SCHEDULED` and `UDI_TREVENT_IO_COMPLETED` events are not defined for use with the Bus Bridge Metalanguage.

4.7 Bus Bridge Metalanguage States

The following events change the state of a device driver with respect to the Bus Bridge Metalanguage:

1. Binding initiated
2. Binding complete
3. Interrupt attachment initiated
4. Interrupt attachment complete
5. Interrupt detachment initiated
6. Interrupt detachment complete
7. Unbinding initiated
8. Unbinding complete

The following events change the state of a bus bridge driver with respect to the Bus Bridge Metalanguage:

1. Binding to a new child
2. New interrupt attached
3. Interrupt detached
4. Child unbound

4.8 Bus Bridge Status Codes

No metalanguage-specific status codes are defined for the Bus Bridge Metalanguage.



UDI Physical I/O Specification

Section 3: Bus Bindings



Introduction to Bus Bindings

5

Some usage details of UDI Physical I/O services and metalanguages vary depending on the physical I/O bus used, requiring a Bus Binding definition to provide these details. This chapter defines the general requirements on Bus Binding definitions, and lists the items that must be specified in a UDI Bus Binding specification.

Note – There are no aspects of DMA services that can vary from one bus binding to another.

5.1 Normative References

Some bus bindings reference non-UDI standards that, through reference in the bus binding specification, constitute provisions of the UDI bus binding. Each such bus binding shall list all normative references.

5.2 Header Files

Each bus binding shall define a bus binding header file that the driver must include to obtain bus binding declarations for a particular bus type. This header must be included by drivers that support devices on this bus type. It must be included after `udi.h` and `udi_physio.h`.

5.3 PIO Bindings

5.3.1 udi_pio_map

The definition of the `regset_idx` parameter to `udi_pio_map` is bus and device dependent. Each bus binding defines legal values for `regset_idx` and the semantics thereof. Drivers must set `regset_idx` to a value appropriate for the bus type of their supported device.

5.4 Interrupt Bindings

5.4.1 Interrupt Index Values

Interrupt index values for `udi_intr_attach_req` are bus-specific. Each bus binding must define the range of legal `interrupt_idx` values and their semantics.

5.4.2 Event Info

Interrupt event info is a bus-specific piece of extra information delivered from the interrupting device hardware along with the interrupt.

The size and structure of event info associated with an interrupt is bus type dependent. Each bus bindings defines the size and semantics of event info for their bus type.

5.5 Instance Attribute Bindings

A number of instance attributes are bus-specific in nature and are thus defined in bus bindings. These attributes shall include a set of enumeration attributes that are required to be set by any bridge driver (or mapper) when it enumerates the children attached to a bus. (See **udi_enumerate_ack** in the UDI Core Specification.)

5.5.1 Enumeration Attributes

In all cases, the “bus_type” attribute is considered an enumeration attribute, and the value used to identify each bus type must be defined in the corresponding bus binding. Each bus binding shall also specify additional enumeration attributes. Any enumerator for a particular bus type must support the specified enumeration attributes for the corresponding bus binding.

5.5.2 Filter Attributes

A subset of enumeration attributes must be specified as filter attributes. Filter attributes may be used with **udi_enumerate_req** to specify enumeration filters. Any enumerator for a particular bus type must support the specified filter attributes for the corresponding bus binding.

5.5.3 Generic Attributes

Each bus binding shall specify the form of string value for each of the generic attributes, “identifier”, “address_locator”, “physical_locator” and “physical_slot”, for the corresponding bus type. The locator attribute is used to distinguish one instance of the same type of device from another. Locator attribute values must be unique for all children of a particular parent bus bridge.

5.5.4 Parent-Visible Attributes

Some bus bindings may also specify parent-visible attributes. Any bus bridge driver for a particular bus type must support the specified parent-visible attributes for the corresponding bus binding.



UDI Physical I/O Specification

Section 4: Appendices

-
- atomic transaction** an I/O bus transaction which is indivisible with respect to other transactions. Other concurrent transactions will either see the state of the referenced memory (or device register) as it was before the atomic transaction or after it, but not an intermediate value. Further, the adapter being accessed will see the access as exactly one transaction. For example, a WRITE4 transaction (4-byte write) is atomic if no observer of the same address can see a mixture of old and new byte values, and two WRITE4's for the same address are atomic with respect to each other if one overwrites the other without mixing byte values.
- bridge** see *bus bridge*.
- bridge driver** see *bus bridge driver*.
- bridge metalanguage** see *bus bridge metalanguage*.
- bus address** see *card-relative address*.
- bus bridge** hardware which connects one kind of bus to another. For example, the bridge(s) between the PMI and primary I/O bus(es), or between a primary I/O bus and an attached (secondary) I/O bus. Bus bridges may be transparent to the software or may require more active participation of a bus bridge driver.
- bus bridge driver** device driver software responsible for managing a bus bridge.
- bus bridge metalanguage** a UDI metalanguage which covers registration and de-registration of interrupt handlers plus the delivery of bus-related events. It is (currently) subdivided into four interface sets, one for each of four roles: the bridge, the device, the interrupt handler and the interrupt dispatcher.
- card-relative address** a physical memory address from the point of view of an adapter card's DMA engine. Such an address is used by the device to access main system memory, but may go through a translation process (in hardware) before reaching the memory.
- DMA** Direct Memory Access by an I/O card to system memory without host processor intervention.
- interrupt dispatcher** a driver which is responsible for managing interrupt control and routing hardware.
- interrupt handler** a driver that is responsible for the interrupt-related, device-specific, software control and management of interrupt-associated hardware.
- interrupt router** the entity, transparent to the interrupt handler driver, that steers an interrupt along a specific signal path between the interrupting device and an interrupt slot on a particular interrupt controller.

interrupt slot	the index of an interrupt request source (starting from zero) relative to a particular interrupt controller. In UDI, the slot number is found in the device's instance attributes.
interrupt source	a specific interrupt request line (or source ID, for messaged based interrupts) through which a device signals its asynchronous events. Depending on the interrupt routing, one or more interrupt sources may share the same interrupt slot. Multiple interrupt sources sharing the same interrupt slot is known as <i>shared interrupts</i> .
ISA	1) Instruction Set Architecture. Defines the binary machine language syntax and semantics for a particular type of processor or processor family. 2) Industry Standard Architecture. An I/O bus type originally designed for the IBM AT and used in many PCs.
PCI	Peripheral Control Interconnect. A specific type of I/O bus.
PIO	Programmed-I/O; i.e. transfer of data between the external device and processor memory under direct program control of the CPU. See DMA for contrast.
PMI	Processor-Memory Interconnect: the bus or interconnect that links the host CPU with system memory and primary I/O bus bridges.
scatter-gather (scgth) element	. a single address/length pair associated with a scatter/gather structure, that identifies one piece of data buffer memory which is contiguous with respect to card-relative addresses.
scatter-gather (scgth) segment	an array of one or more scatter/gather elements. Each scatter/gather structure references one or more scatter/gather segments, as necessary to address the entire data buffer. The memory containing a scatter/gather segment is itself contiguous with respect to card-relative addresses.
scatter-gather (scgth) structure	a driver-visible UDI structure that reflects the card-relative addresses of a possibly physically discontinuous but logically contiguous data buffer, which the driver uses to set up DMA transfers.
shared interrupts	see <i>interrupt source</i> .
sync	an environment service that synchronizes the hosts view of memory with an I/O card's view. This includes cache flushing and purging, flushing buffers in bus bridges, etc.



Index

B

Bus Bridge Metalanguage

roles

- binding & interrupt registration 4-1

- interrupt acknowledgements 4-1

- interrupt events 4-1

D

DMA

- address-length pairs 2-1

- constraints 2-15

- definition 2-1

handle

- allocation 2-24

- definition 2-1

- transferability 2-6

- memory allocation limits 2-7

- memory ordering barriers 2-35

- scatter-gather list 2-1

- synchronization 2-1, 2-32, 2-34

I

Instance attributes

- !bus_type 3-4, 4-11

- interrupt region 4-19

P

- physical device 1-1

PIO

- attributes 3-1

- data alignment 3-6

- data ordering 3-5

- data translation 3-5

- definition 3-1

handle

- allocation 3-2, 3-4

- deallocation 3-2, 3-8

- definition 3-1

- transferability 3-3, 3-28

- multiple mappings 3-3, 3-7

- pacing 3-6

R

Region attributes

- pio_probe 1-2, 3-27

- type 1-2

- Region kill 3-10

T

Transferability

- of DMA handles 2-6

- of PIO handles 3-3

U

- UDI_PHYSIO_VERSION 1-1

