



Specification for the:

*WISHBONE System-on-Chip (SoC)
Interconnection Architecture
for Portable IP Cores*

Revision: B.32, Released: ~~To be determined~~ ~~October 10, 2001~~

Preliminary

Steward's Change Copy

Wade D. Peterson – Revised: November 30, 2001

This document shows changes to the WISHBONE specification that are recommended or under review by the Steward. The purpose of this document is to convey all known issues that need to be addressed during the next revision of the specification. This is a working document that is subject to change without notice. All changes shown herein are NOT APPROVED, and are presented for informational purposes only.

Silicore Corporation

6310 Butterworth Lane – Corcoran, MN 55340
TEL: (763) 478-3567 FAX: (763) 478-3568
www.silicore.net



This Page is Intentionally Blank

Stewardship

Stewardship for this specification is maintained by Silicore Corporation. Questions, comments and suggestions about this document are welcome and should be directed to:

Wade D. Peterson, Silicore Corporation
6310 Butterworth Lane – Corcoran, MN USA 55340
TEL: (763) 478-3567; FAX: (763) 478-3568
E-MAIL: wadep@silicore.net URL: www.silicore.net

Silicore Corporation maintains this document as a service to its customers and to the IP core industry as a whole. The intent of this specification is to improve the quality of Silicore products, as well as to foster cooperation among the users and suppliers of IP cores.

Copyright & Trademark Release / Royalty Release / Patent Notice

Notice is hereby given that this document is not copyrighted, and has been placed into the public domain. It may be freely copied and distributed by any means.

The name ‘WISHBONE’ and the ‘WISHBONE COMPATIBLE’ rubber stamp logo are hereby placed into the public domain (within the scope of System-on-Chip design, System-on-Chip fabrication and related areas of commercial use). The WISHBONE logo may be freely used under the compatibility conditions stated elsewhere in this document.

This specification may be used for the design and production of System-on-Chip (SoC) components without royalty or other financial obligation to Silicore Corporation.

The author(s) of this specification are not aware that the information contained herein, nor of products designed to the specification, cause infringement on the patent, copyright, trademark or trade secret rights of others. However, there is a possibility that such infringement may exist without their knowledge. The user of this document assumes all responsibility for determining if products designed to this specification infringe on the intellectual property rights of others.

Disclaimers

In no event shall Silicore Corporation be liable for incidental, consequential, indirect, or special damages resulting from the use of this specification. By adopting this specification, the user assumes all responsibility for its use.

This is a preliminary document, and is subject to change.

Silicore® is a registered service mark and trademark of Silicore Corporation.
Verilog® is a registered trademark of Cadence Design Systems, Inc.

Document Format, Binding and Covers

This document is formatted for printing on double sided, 8½” x 11” white paper stock. It is designed to be bound within a standard cover. The preferred binding method is a black coil binding with outside diameter of 9/16” (14.5 mm). The preferred cover stock is Paper Direct part number KVR09D (forest green) and is available on-line at: www.paperdirect.com. Binding can be performed at most full-service copy centers such as Kinkos (www.kinkos.com).

Acknowledgements

Like any great technical project, the WISHBONE specification could not have been completed without the help of many people. The Steward wishes to thank the following for their ideas, suggestions and contributions:

Ray Alderman
Yair Amitay
Danny Cohan
Marc Delvaux
Richard Herveille*
Volker Hetzer
Magnus Homann
Brian Hurt
Linus Kirk
Damjan Lampret
Barry Rice
John Rynearson
Rudolf Usselmann

(*). Special thanks to Richard Herveille for his many valuable comments and suggestions.

Revision History

The various revisions of the WISHBONE specification, along with their changes and revision history, can be found at www.silicore.net/wishbone.htm.

Table of Contents

CHAPTER 1 - INTRODUCTION	7
1.1 WISHBONE FEATURES.....	8
1.2 WISHBONE OBJECTIVES.....	10
1.3 SPECIFICATION TERMINOLOGY.....	13 12
1.4 USE OF TIMING DIAGRAMS.....	14 13
1.5 SIGNAL NAMING CONVENTIONS.....	16 15
1.6 WISHBONE LOGO.....	17 16
1.7 GLOSSARY OF TERMS.....	17 16
1.8 REFERENCES.....	29 25
CHAPTER 2 – INTERFACE SPECIFICATION	30 26
2.1 REQUIRED DOCUMENTATION FOR IP CORES.....	30 26
2.1.1 <i>General Requirements for the WISHBONE DATASHEET</i>	30 26
2.1.2 <i>Signal Naming</i>	31 27
2.1.3 <i>Logic Levels</i>	32 28
2.2 WISHBONE SIGNAL DESCRIPTION.....	33 29
2.2.1 <i>SYSCON Signals</i>	33 29
2.2.2 <i>Signals Common to MASTER and SLAVE Interfaces</i>	34 29
2.2.3 <i>MASTER Signals</i>	35 30
2.2.4 <i>SLAVE Signals</i>	37 31
CHAPTER 3 – BUS CYCLES	40 33
3.1 GENERAL OPERATION.....	40 33
3.1.1 <i>Reset Operation</i>	40 33
3.1.2 <i>Handshaking Protocol</i>	42 35
3.1.3 <i>Use of [STB_O]</i>	45 38
3.1.4 <i>Use of [ACK_O], [ERR_O] and [RTY_O]</i>	45 38
3.1.5 <i>Use of [TAGN_I] and [TAGN_O] Signals</i>	45 38
3.2 SINGLE READ / WRITE CYCLES.....	47 39
3.2.1 <i>SINGLE READ Cycle</i>	48 40
3.2.2 <i>SINGLE WRITE Cycle</i>	50 42
3.3 BLOCK READ / WRITE CYCLES.....	52 44
3.3.1 <i>BLOCK READ Cycle</i>	53 45
3.3.2 <i>BLOCK WRITE Cycle</i>	56 48
3.4 RMW CYCLE.....	59 51
3.5 DATA ORGANIZATION.....	62 54
3.5.1 <i>Nomenclature</i>	62 54
3.5.2 <i>Transfer Sequencing</i>	65 57
3.5.3 <i>Data Organization for 64-bit Ports</i>	65 57
3.5.4 <i>Data Organization for 32-bit Ports</i>	67 59
3.5.5 <i>Data Organization for 16-bit Ports</i>	68 60
3.5.6 <i>Data Organization for 8-bit Ports</i>	69 61
3.6 REFERENCES.....	69 61

CHAPTER 4 – TIMING SPECIFICATION	<u>7062</u>
REFERENCES	<u>7264</u>
APPENDIX A – WISHBONE TUTORIAL.....	<u>7665</u>
A.1 AN INTRODUCTION TO WISHBONE.....	<u>7665</u>
A.2 TYPES OF WISHBONE INTERCONNECTION.....	<u>7968</u>
A.2.1 <i>Point-to-point Interconnection</i>	<u>7968</u>
A.2.2 <i>Data Flow Interconnection</i>	<u>7968</u>
A.2.3 <i>Shared Bus Interconnection</i>	<u>8069</u>
A.2.4 <i>Crossbar Switch Interconnection</i>	<u>8170</u>
A.3 THE WISHBONE INTERFACE SIGNALS.....	<u>8372</u>
A.4 THE WISHBONE BUS CYCLES.....	<u>8574</u>
A.4.1 <i>SINGLE READ/WRITE Cycle</i>	<u>8574</u>
A.4.2 <i>BLOCK READ/WRITE Cycle</i>	<u>8776</u>
A.4.3 <i>READ-MODIFY-WRITE (RMW) Cycle</i>	<u>8776</u>
A.5 ENDIAN.....	<u>8877</u>
A.6 SLAVE I/O PORT EXAMPLES.....	<u>9079</u>
A.6.1 <i>Simple 8-bit SLAVE Output Port</i>	<u>9079</u>
A.6.2 <i>Simple 16-bit SLAVE Output Port With 16-bit Granularity</i>	<u>9382</u>
A.6.3 <i>Simple 16-bit SLAVE Output Port With 8-bit Granularity</i>	<u>9483</u>
A.7 WISHBONE MEMORY INTERFACING.....	<u>9786</u>
A.7.1 <i>FASM Synchronous RAM and ROM Model</i>	<u>9786</u>
A.7.2 <i>Simple 16 x 8-bit SLAVE Memory Interface</i>	<u>9988</u>
A.7.3 <i>Memory Primitives and the [ACK_O] Signal</i>	<u>10190</u>
A.8 POINT-TO-POINT INTERCONNECTION EXAMPLE.....	<u>10391</u>
A.9 SHARED BUS EXAMPLE.....	<u>10794</u>
A.9.1 <i>Choosing Between Multiplexed and Non-multiplexed Bus Topology</i>	<u>10794</u>
A.9.2 <i>Choosing Between Three-State and Multiplexor Interconnection Logic</i>	<u>10895</u>
A.9.3 <i>Creating the Interconnection Topology</i>	<u>11198</u>
A.9.4 <i>Full vs. Partial Address Decoding</i>	<u>114101</u>
A.9.5 <i>The System Arbiter</i>	<u>117104</u>
A.9.6 <i>Creating and Benchmarking the System</i>	<u>117104</u>
A.10 REFERENCES.....	<u>120107</u>
INDEX.....	<u>121108</u>

Chapter 1 - Introduction

The WISHBONE¹ System-on-Chip (SoC) Interconnection Architecture for Portable IP Cores is a flexible design methodology for use with semiconductor IP cores. Its purpose is to foster design reuse by alleviating System-on-Chip integration problems. This is accomplished by creating a common interface between IP cores. This improves the portability and reliability of the system, and results in faster time-to-market for the end user.

Previously, IP cores used non-standard interconnection schemes that made them difficult to integrate. This required the creation of custom glue logic to connect each of the cores together. By adopting a standard interconnection scheme, the cores can be integrated more quickly and easily by the end user.

This specification can be used for soft core, firm core or hard core IP. Since firm and hard cores are generally conceived as soft cores, the specification is written from that standpoint.

This specification does not require the use of specific development tools or target hardware. Furthermore, it is fully compliant with virtually all logic synthesis tools. However, the examples presented in the specification do use the VHDL hardware description language. These are presented only as a convenience to the reader, and should be readily understood by users of other hardware description languages (such as Verilog®). Schematic based tools can also be used.

The WISHBONE interconnect is intended as a general purpose interface. As such, it defines the standard data exchange between IP core modules. It does not attempt to regulate the application-specific functions of the IP core.

The WISHBONE architects were strongly influenced by three factors. First, there was a need for a good, reliable System-on-Chip integration solution. Second, there was a need for a common interface specification to facilitate structured design methodologies on large project teams. Third, they were impressed by the traditional system integration solutions afforded by microcomputer buses such as PCI bus and VMEbus.

In fact, the WISHBONE architecture is analogous to a microcomputer bus in that they both: (a) offer a flexible integration solution that can be easily tailored to a specific application; (b) offer a variety of bus cycles and data path widths to solve various system problems; and (c) allow products to be designed by a variety of suppliers (thereby driving down price while improving performance and quality).

¹ Webster's dictionary defines a WISHBONE as "the forked clavicle in front of the breastbone of most birds." The term 'WISHBONE interconnect' was coined by Wade Peterson of Silicore Corporation. During the initial definition of the scheme he was attempting to find a name that was descriptive of a bi-directional data bus that used either multiplexors or three-state logic. This was solved by forming an interface with separate input and output paths. When these paths are connected to three-state logic it forms a 'Y' shaped configuration that resembles a wishbone. The actual name was conceived during a Thanksgiving Day dinner that included roast turkey. Thanksgiving Day is a national holiday in the United States, and is observed on the third Thursday in November. It is generally celebrated with a traditional turkey dinner.

However, traditional microcomputer buses are fundamentally handicapped for use as a System-on-Chip interconnection. That's because they are designed to drive long signal traces and connector systems which are highly inductive and capacitive. In this regard, System-on-Chip is much simpler and faster. Furthermore, the System-on-Chip solutions have a rich set of interconnection resources. These do not exist in microcomputer buses because they are limited by IC packaging and mechanical connectors.

The WISHBONE architects have attempted to create a specification that is robust enough to insure complete compatibility between IP cores. However, it has not been over specified so as to unduly constrain the creativity of the core developer or the end user. It is believed that these two goals have been accomplished with the publication of this document.

1.1 WISHBONE Features

The WISHBONE interconnection makes System-on-Chip and design reuse easy by creating a standard data exchange protocol. Features of this technology include:

- Simple, compact, logical IP core hardware interfaces that require very few logic gates.
- Supports structured design methodologies used by large project teams.
- Full set of popular data transfer bus protocols including:
 - READ/WRITE cycle
 - BLOCK transfer cycle
 - RMW cycle
- Data bus widths² and operand sizes up to 64-bits.
- Supports both BIG ENDIAN and LITTLE ENDIAN data ordering.
- Variable core interconnection methods support point-to-point, shared bus, crossbar switch, and switched fabric interconnections.
- Handshaking protocol allows each IP core to throttle its data transfer speed.
- Supports single clock data transfers.
- Supports normal cycle termination, retry termination and termination due to error.
- Address widths³ up to 64-bits.

² Specifications are given for data port and operand sizes up to 64-bits. However, the basic architecture can theoretically support any data width (e.g. 128-bit, 256-bit etc.). ~~Also, zero bit data bus accesses are permissible (generally used in FIFO interfaces).~~

- Partial address decoding scheme for SLAVES. This facilitates high speed address decoding, uses less redundant logic and supports variable address sizing and interconnection means.
- User-defined tags ~~support~~. ~~These is are~~ useful for applying information to an address bus, a data bus or a bus cycle. They are especially helpful when modifying a bus cycle to identify information such as: ~~identifying transfers such as:~~
 - Data transfers
 - Parity or error correction bits
 - Interrupt vectors
 - Cache control operations
- MASTER / SLAVE architecture for very flexible system designs.
- Multiprocessing (multi-MASTER) capabilities. This allows for a wide variety of System-on-Chip configurations.
- Arbitration methodology is defined by the end user (priority arbiter, round-robin arbiter, etc.).
- Supports various IP core interconnection means, including:
 - Point-to-point
 - Shared bus
 - Crossbar switch
 - Data flow interconnection
 - Unidirectional bus
 - Bi-directional bus
 - Multiplexor based interconnections
 - Three state based interconnections
 - ~~Off chip I/O~~
- Synchronous design assures portability, simplicity and ease of use.
- Very simple, variable timing specification.
- Documentation standards simplify IP core reference manuals. ~~requirements allow the end user to quickly evaluate interface needs.~~
- Independent of hardware technology (FPGA, ASIC, etc.).

³ Specifications are given for address widths between zero and 64-bits. However, the basic architecture can theoretically support any address width. Also, zero bit address buses (used for FIFO interfaces) are allowed.

- Independent of delivery method (soft, firm or hard core).
- Independent of synthesis tool, router and layout tool technology.
- Independent of FPGA and ASIC test methodologies.
- Seamless design progression between FPGA prototypes and ASIC production chips.

1.2 WISHBONE Objectives

The main objective of the specification is to create a flexible interconnection means for use with semiconductor IP cores. This allows various IP core modules to be connected together to form a System-on-Chip.

A further objective of the specification is to enforce good compatibility between IP core modules. This enhances design reuse.

A further objective of the specification is to create a robust standard, but one that does not unduly constrain the creativity of the core developer or the end user.

A further objective of the specification is to make it easy to understand by both the core developer and the end user.

A further objective of the specification is to facilitate structured design methodologies on large project teams. With structured design, individual team members can build and test small parts of the design. Each member of the design team can interface to the common, well-defined WISHBONE specification. When all of the sub-assemblies have been completed, the full system can be integrated.

A further objective of the specification is create a portable interface that is independent of the underlying semiconductor technology. For example, WISHBONE interconnections can be made that support both FPGA and ASIC target devices.~~the interconnect must be capable of working with both FPGA and ASIC hardware target devices.~~

A further objective of the specification is to make WISHBONE~~the~~ interfaces independent of logic signaling levels.

A further objective of the specification is to create a flexible interconnection scheme that is independent of the IP core delivery method. For example, it may be used with ‘soft core’, ‘firm core’ or ‘hard core’ delivery methods.

A further objective of the specification is to be independent of the underlying hardware description. For example, soft cores may be written and synthesized in VHDL, Verilog® or some other hardware description language. Schematic entry may also be used.

A further objective of the specification is to require a minimum standard for documentation. This takes the form of the WISHBONE DATASHEET, and allows IP core users to quickly evaluate and integrate new cores.

A further objective of the specification is to eliminate extensive interface documentation on the part of the IP core developer. In most cases, this specification along with the WISHBONE DATASHEET is sufficient to completely document an IP core data interface.

A further objective of the specification is to allow users to create SoC components without infringing on the patent rights of others. While the use of WISHBONE technology does not necessarily prevent patent infringement, it does provide a reasonable safe haven where users can design around the patent claims of others. The specification also provides cited patent references, which describes the field of search used by the WISHBONE architects.

A further objective of the specification is to identify critical System-on-Chip interconnection technologies, and to place them into the public domain at the earliest possible date. This makes it more difficult for individuals and organizations to create proprietary SoC technologies through the use of patent, trademark, copyright and trade secret protection mechanisms. ~~This objective applies only to the interconnection of IP cores, but not to the IP cores themselves.~~

A further objective of the specification is to support a business model whereby IP Core suppliers can cooperate at a technical standards level, but can also compete in the commercial marketplace. This improves the overall quality and value of products through market forces such as price, service, delivery, performance and time-to-market. This business model also allows open source IP cores to be offered as well.

A further objective is to create an architecture that has a smooth transition path to support new technologies. This increases the longevity of the specification as it can adapt to new, and as yet un-thought-of, requirements.

A further objective is to create an architecture that allows various interconnection means between IP core modules. This insures that the end user can tailor the System-on-Chip to his/her own needs. For example, the entire interconnection system (which is analogous to a backplane on a standard microcomputer bus like VMEbus or cPCI) can be created by the system integrator. This allows the interconnection to be tailored to the final target device.

A further objective is to create an architecture that requires a minimum of glue logic. In some cases the System-on-Chip needs no glue logic whatsoever. However, in other cases the end user may choose to use a more sophisticated interconnection method (for example with FIFO memories or crossbar switches) that requires additional glue logic.

A further objective is to create an architecture with variable address and data path widths to meet a wide variety of system requirements.

A further objective is to create an architecture that fully supports the automatic generation of interconnection and IP Core systems. This allows components~~the WISHBONE interconnection~~ to be generated with parametric core generators.

A further objective is to create an architecture that supports both BIG ENDIAN and LITTLE ENDIAN data transfer organizations.

A further objective is to create an architecture that supports one data transfer per clock cycle.

A further objective is to create a flexible architecture that allows address, data and bus cycles to be tagged, to be tagged. TagsAGs are user defined signals that allow user~~each IP core~~ to modify a bus cycle with additional information, communicate with the rest of the system. They are especially useful when novel or unusual control signals (such as parity, cache control or interrupt acknowledge) are needed on an interface.

A further objective is to create an architecture with a MASTER/SLAVE topology. Furthermore, the system must be capable of supporting multiple MASTERS and multiple SLAVES with an efficient arbitration mechanism.

A further objective is to create an architecture that supports point-to-point interconnections between IP cores.

A further objective is to create an architecture that supports shared bus interconnections between IP cores.

A further objective is to create an architecture that supports crossbar switches between IP cores.

A further objective is to create an architecture that supports switched fabrics.

A further objective is to create a synchronous protocol to insure ease of use, good reliability and easy testing. Furthermore, all transactions can be coordinated by a single clock.

A further objective is to create a synchronous protocol that works over a wide range of interface clock speeds. The effects of this are: (a) that the WISHBONE interface can work synchronously with all attached IP cores, (b) that the interface can be used on a large range of target devices, (c) that the timing specification is much simpler and (d) that the resulting semiconductor device is much more testable.

A further objective is to create a variable timing mechanism whereby the system clock frequency can be adjusted so as to control the power consumption of the integrated circuit.

A further objective is to create a synchronous protocol that provides a simple timing specification. This makes the interface very easy to integrate.

A further objective is to create a synchronous protocol where each MASTER and SLAVE can throttle the data transfer rate with a handshaking mechanism.

A further objective is to create a synchronous protocol that is optimized for System-on-Chip, but that is also suitable for off-chip I/O routing. Generally, the off-chip WISHBONE interconnect will operate at slower speeds.

1.3 Specification Terminology

To avoid confusion, and to clarify the requirements for compliance, this specification ~~makes use of~~ five keywords. ~~They are: to define the operation of the WISHBONE interconnect. The key words are:~~

- **RULE**
- **RECOMMENDATION**
- **SUGGESTION**
- **PERMISSION**
- **OBSERVATION**

Any text not labeled with one of these keywords describes the operation in a narrative style. The keywords are defined as follows:

RULE

Rules form the basic framework of the specification. They are sometimes expressed in text form and sometimes in the form of figures, tables or drawings. All rules **MUST** be followed to ensure compatibility between interfaces. Rules are characterized by an imperative style. The upper-case words **MUST** and **MUST NOT** are reserved exclusively for stating rules in this document, and are not used for any other purpose.

RECOMMENDATION

Whenever a recommendation appears, designers would be wise to take the advice given. Doing otherwise might result in some awkward problems or poor performance. While this specification has been designed to support high performance systems, it is possible to create an interconnection that complies with all the rules, but has very poor performance. In many cases a designer needs a certain level of experience with the system architecture in order to design interfaces that deliver top performance. Recommendations found in this document are based on this kind of experience and are provided as guidance for the user.

SUGGESTION

A suggestion contains advice which is helpful but not vital. The reader is encouraged to consider the advice before discarding it. Some design decisions are difficult until experience has been gained. Suggestions help a designer who has not yet gained this experience. Some suggestions have to do with designing compatible interconnections, or with making system integration easier.

PERMISSION

In some cases a rule does not specifically prohibit a certain design approach, but the reader might be left wondering whether that approach might violate the spirit of the rule, or whether it might

lead to some subtle problem. Permissions reassure the reader that a certain approach is acceptable and will not cause problems. The upper-case word MAY is reserved exclusively for stating a permission and is not used for any other purpose.

OBSERVATION

Observations do not offer any specific advice. They usually clarify what has just been discussed. They spell out the implications of certain rules and bring attention to things that might otherwise be overlooked. They also give the rationale behind certain rules, so that the reader understands why the rule must be followed.

1.4 Use of Timing Diagrams

Figure 1-1 shows some of the key features of the timing diagrams in this specification. Unless otherwise noted, the MASTER signal names are referenced in the timing diagrams. In some cases the MASTER and SLAVE signal names are different. For example, in the ~~single MASTER / single SLAVE configuration~~, point-to-point interconnections the [ADR_O] and [ADR_I] signals are connected together. Furthermore, the actual waveforms at the SLAVE may vary from those at the MASTER. That's because the MASTER and SLAVE interfaces can be connected together in different ways. Unless otherwise noted, the timing diagrams refer to the connection diagram shown in Figure 1-2.

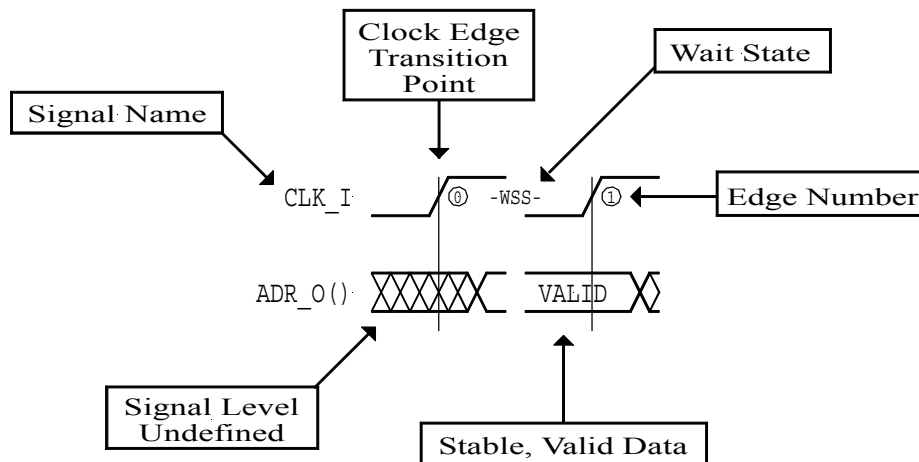


Figure 1-1. Use of timing diagrams.




Figure 1-2. Standard connection for timing diagrams.

~~Some individual~~ signals may or may not be present on an specific interface. That's because many of the signals are optional.

Two symbols are also presented in relation to the [CLK_I] signal. These include the positive going clock edge transition point and the clock edge number. In most diagrams a vertical guideline is shown at the positive-going edge of each [CLK_I] transition. This represents the theoretical transition point at which flip-flops register their input value, and transfer it to their output. The exact level of this transition point varies depending upon the technology used in the target device. The clock edge number is included as a convenience so that specific points in the timing diagram may be referenced in the text. The clock edge number in one timing diagram is not related to the clock edge number in another diagram.

Gaps in the timing waveforms may be shown. These indicate either: (a) a wait state or (b) a portion of the waveform that is not of interest (in the context of the diagram). When the gap indicates a wait state, the symbols '-WSM-' or '-WSS-' are placed in the gap along the [CLK_I] waveform. These correspond to wait states inserted by the MASTER or SLAVE interfaces respectively. They also indicate that the signals (with the exception of clock transitions and hatched regions) will remain in a steady state during that time.

Undefined signal levels are indicated by a hatched region. This region indicates that the signal level is undefined, and may take any state. It also indicates that the current state is undefined, and should not be relied upon. When signal arrays are used, stable and predictable signal levels are indicated with the word 'VALID'.

1.5 Signal Naming Conventions

All signal names used in this specification have the '_I' or '_O' characters attached to them. These indicate if the signals are an input (to the core) or an output (from the core). For example, [ACK_I] is an input and [ACK_O] is an output. This convention is used to clearly identify the direction of each signal.

~~In some cases, the input and output characters 'I' and 'O' may be omitted and replaced by an 'X'. For example: [TAG3_X]. This is not an actual signal name, but rather a shorthand form to indicate both the [TAG3_I] and [TAG3_O] signal.~~

Signal arrays are identified by a name followed by ~~the array boundaries in a set of~~ parenthesis. For example, [DAT_I(63..0)] is a signal array. Array limits may also be shown within the parenthesis. In this case the first number of the array limit indicates the most significant bit, and the second number indicates the least significant bit. For example, [DAT_I(63..0)] is a signal array with upper array boundary number sixty-three (the most significant bit), and lower array boundary number zero (the least significant bit). Furthermore, the array boundaries indicate the full range of the permissible array size. The array size on any particular core may vary. In many cases the array boundaries are omitted if they are irrelevant to the context of the description.

Special user defined signals, called *tags*, can also be used. Tags are assigned a *tag type* that indicates the exact timing to which the signal must adhere. For example, if a parity bit such as [PAR_O] is added to a data bus, it would probably be assigned a tag type of TAG TYPE: TGD_O(). This indicates that the signal will adhere to the timing diagrams shown for [TGD_O()], which are shown in the timing diagrams for each bus cycle. Also note that, while all tag types are specified as arrays (with parenthesis ‘()’), the actual tag does not have to be a signal array. It can also be non-arrayed signal.

When used as part of a sentence, signal names are enclosed in brackets ‘[]’. This helps to discriminate signal names from the words in the sentence.

1.6 WISHBONE Logo

The WISHBONE logo can be affixed to SoC documents that are compatible with this standard. Figure 1-3 shows the logo.



Figure 1-3. WISHBONE logo.

PERMISSION 1.00

Documents describing a WISHBONE compatible SoC component that are 100% compliant with this specification standard, MAY use the WISHBONE logo.

1.7 Glossary of Terms

0x (numerical prefix)

The ‘0x’ prefix indicates a hexadecimal number. It is the same nomenclature as commonly used in the ‘C’ programming language.

Active High Logic State

A logic state that is ‘true’ when the logic level is a binary ‘1’ (high state). The high state is at a higher voltage than the low state.

Active Low Logic State

A logic state that is ‘true’ when the logic level is a binary ‘0’ (low state). The low state is at a lower voltage than the high state.

Address Tag

One or more user defined signals that modify a WISHBONE address. For example, they can be used create a parity bit on an address bus, to indicate an address width (16-bit, 24-bit etc.) or can be used by memory management hardware to indicate a protected address space. All address tags must be assigned a tag type of [TGA_I0] or [TGA_O0]. Also see *tag, tag type, data tag* and *cycle tag*.

ASIC

Acronym for: Application Specific Integrated Circuit. AgGeneral term which describes a generic array of logic gates or analog building blocks which are programmed by a metalization layer at a silicon foundry. High level circuit descriptions are impressed upon the logic gates or analog building blocks in the form of metal interconnects.

Asserted

(1) A verb indicating that a logic state has switched from the inactive to the active state. When active high logic is used it means that a signal has switched from a logic low level to a logic high level. (2) *Assert*: to cause a signal line to make a transition from its logically false (inactive) state to its logically true (active) state. Opposite of *negated*.

Bus

(1) A common group of signals. (2) A signal line or a set of lines used by a data transfer system to connect a number of devices.

Bus Interface

An electronic circuit that drives or receives data or power from a bus.

Bus Cycle

The process whereby digital signals effect the transfer of data across a bus by means of an interlocked sequence of control signals. Also see: *Phase (bus cycle)*.

Crossbar Interconnection (Crossbar Switch)

Crossbar switches are mechanisms that allow modules to connect and communicate. Each connection channel can be operated in parallel to other connection channels. This increases the data transfer rate of the entire system by employing parallelism. Stated another way, two 100 Mbyte/second channels can operate in parallel, thereby providing a 200 Mbyte/second transfer rate. This makes the crossbar switches inherently faster than traditional bus schemes. Crossbar routing mechanisms generally support dynamic configuration. This creates a configurable and reliable network system. Most crossbar architectures are also scalable, meaning that families of crossbars can be added as the needs arise. A crossbar interconnection is shown in Figure 1-4.

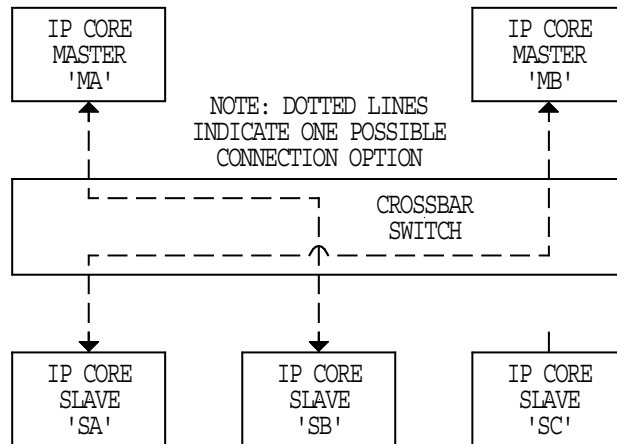


Figure 1-4. Crossbar (switch) interconnection.

Cycle Tag

One or more user defined signals that modify a WISHBONE bus cycle. For example, they can be used to discriminate between WISHBONE SINGLE, BLOCK and RMW cycles. All cycle tags must be assigned a tag type of [TGC_I()] or [TGC_O()]. Also see tag type, address tag and data tag.

Data Flow Interconnection

An interconnection where data flows through a prearranged set of IP cores in a sequential order. Data flow architectures often have the advantage of parallelism, whereby two or more functions are executed at the same time. Figure 1-5 shows a data flow interconnection between IP cores.




Figure 1-5. Data flow interconnection.

Data Organization

The ordering of data during a transfer. Generally, 8-bit (byte) data can be stored with the most significant byte of a multi-byte transfer at the higher or the lower address. These two methods are generally called BIG ENDIAN and LITTLE ENDIAN, respectively. In general, BIG ENDIAN refers to byte lane ordering where the most significant byte is stored at the lower address. LITTLE ENDIAN refers to byte lane ordering where the most significant byte is stored at the higher address. The terms BIG ENDIAN and LITTLE ENDIAN for data organization was coined by Danny Cohen of the Information Sciences Institute, and was derived from the book Gulliver's Travels by Jonathan Swift.

Data Tag

One or more user defined signals that modify a WISHBONE data transfer. For example, they can be used carry parity information, error correction codes or time stamps. All data tags must be assigned a tag type of [TGD_I()] or [TGD_O()]. Also see *tag type*, *address tag* and *cycle tag*.

DMA Unit

Acronym for Direct Memory Access Unit. (1) A device that transfers data from one location in memory to another location in memory. (2) A device for transferring data between a device and memory without interrupting program flow. (3) A device that does not use low-level instructions and is ~~intended~~ for transferring data between memory or I/O locations.

ENDIAN

See the definition under 'Data Organization'.

FIFO

Acronym for: First In First Out. A type of memory used to transfer data between ports on two devices. In FIFO memories, data is removed in the same order that they were added. The FIFO memory is very useful for interconnecting cores of differing speeds.

Firm Core

An IP Core that is delivered in a way that allows conversion into an integrated circuit design, but does not allow the design to be easily reverse engineered. It is analogous to a binary or object file in the field of computer software design.

Fixed Interconnection

An interconnection system that is fixed, and *cannot* be changed without causing incompatibilities between bus modules (or SoC/IP cores). Also called a *static interconnection*. Examples of fixed interconnection buses include PCI, cPCI and VMEbus. Also see: *variable interconnection*.

Fixed Timing Specification

A timing specification that is based upon a fixed set of rules. Generally used in traditional microcomputer buses like PCI and VMEbus. Each bus module must conform to the ridged set of timing specifications. Also see: *variable timing specification*.

Foundry

See silicon foundry.

FPGA

Acronym for: Field Programmable Gate Array. Describes a generic array of logical gates and interconnect paths which are programmed by the end user. High level logic descriptions are impressed upon the gates and interconnect paths, often in the form of IP Cores.

Full Address Decoding

A method of address decoding where each SLAVE decodes all of the available address space. For example, if a 32-bit address bus is used, then each SLAVE decodes all thirty-two address bits. This technique is used on standard microcomputer buses like PCI and VMEbus. Also see: *partial address decoding*.

Gated Clock

~~A type of SYSCON interface. A clock that can be stopped and restarted, where clock signal [CLK_O] can be stopped and restarted.~~ In WISHBONE, a gated clock generator allows [CLK_O] to be stopped. The signal is always stopped in its low state. This technique is often used to reduce the power consumption of an integrated circuit. Under WISHBONE, the gated clock generator is optional. Also see: *variable clock generator*.

Glue Logic

(1) Logic gates and interconnections required to connect IP cores together. The requirements for glue logic vary greatly depending upon the interface requirements of the IP cores. (2) A family of logic circuits consisting of various gates and simple logic elements, each of which serve as an interface between various parts of a computer system.

Granularity

The smallest unit of data transfer that a port is capable of transferring. For example, a 32-bit port can be broken up into four 8-bit BYTE segments. In this case, the granularity of the interface is 8-bits. Also see: *port size* and *operand size*.

Hard Core

An IP Core that is delivered in the form of a mask set (i.e. a graphical description of the features and connections in an integrated circuit).

Hardware Description Language (HDL)

(1) Acronym for: Hardware Description Language. Examples include VHDL and Verilog®. (2) A general-purpose language used for the design of digital electronic systems.

Interface

A combination of signals and data-ports on a module which are capable of either generating or; receiving; ~~controlling or interconnecting IP cores bus cycles~~. WISHBONE ~~defines these as MASTER, SLAVE, INTERCON and SYSCON interfaces respectively~~ defines these as MASTER and SLAVE interfaces respectively. Also see: MASTER; and SLAVE interfaces; INTERCON and SYSCON.

INTERCON

A WISHBONE interfacemodule that interconnects MASTER; and SLAVE ~~and SYSCON~~ interfaces.

IP Core

Acronym for: Intellectual Property Core. Also see: *soft core, firm core and hard core*.

Mask Set

A graphical description of the features and connections in an integrated circuit.

MASTER

A WISHBONE interface that is capable of generating bus cycles. All systems based on the WISHBONE interconnect must have at least one MASTER interface. Also see: SLAVE; SYSCON and INTERCON.

Memory Mapped Addressing

An architecture that allows data to be stored and recalled in memory at individual, binary addresses.

Minimization (Logic Minimization)

A process by which HDL synthesis, router or other software development tools remove unused logic. This is important in WISHBONE because there are optional signals defined on many of the interfaces. If a signal is unused, then the logic minimization tools will remove these signals and their associated logic, thereby making a faster and more efficient design.

Module

In the context of this specification, it's another name for an IP core.

Multiplexor Interconnection

An interconnection that uses multiplexors to route address, data and control signals. Often used for System-on-Chip (SoC) applications. Also see: *three-state bus interconnection*.

Negated

A verb indicating that a logic state has switched from the active to the inactive state. When active high logic is used it means that a signal has switched from a logic high level to a logic low level. Also see: *asserted*.

Off-Chip Interconnection

An off-chip interconnection is used when a WISHBONE interface extends off-chip. See Figure 1-6.




Figure 1-6. Off-chip interconnection.

Operand Size

The operand size is the largest single unit of data that is moved through an interface. For example, a 32-bit DWORD operand can be moved through an 8-bit port with four data transfers. Also see: *granularity* and *port size*.

Parametric Core Generator

A software tool used for the generation of IP cores based on input parameters. One example of a parametric core generator is a DSP filter generator. These are programs that create lowpass, bandpass and highpass DSP filters. The parameters for the filter are provided by the user, which causes the program to produce the digital filter as a VHDL or Verilog® hardware description. Parametric core generators can also be used create WISHBONE interconnections.

Partial Address Decoding

A method of address decoding where each SLAVE decodes only the range of addresses that it requires. For example, if the module needs only four addresses, then it decodes only the two least significant address bits. The remaining address bits are decoded by the interconnection system. This technique is used on SoC buses, and has the advantages of: less redundant logic in the system, it supports variable address buses, it supports variable interconnection buses and is relatively fast. Also see: *full address decoding*.

PCI

Acronym for: Peripheral Component Interconnect. Generally used as an interconnection scheme between integrated circuits. It also exists as a board level interconnection known as Compact PCI (or cPCI). While this specification is very flexible, it isn't practical for SoC applications.

Phase (Bus Cycle)

A periodic portion of a bus cycle. For example, a WISHBONE BLOCK READ cycle could contain ten phases, with each phase transferring a single 32-bit word of data. Collectively, the ten phases form the BLOCK READ cycle.

Point-to-point Interconnection

(1) An interconnection system that supports a single WISHBONE MASTER and a single WISHBONE SLAVE interface. It is the simplest way to connect two cores. See Figure 1-7. (2) A connection with only two endpoints.




Figure 1-7. Point to point interconnection.

Port Size

The width of the WISHBONE data ports in bits. Also see: *granularity* and *operand size*.

Router

A software tool that physically routes interconnection paths between logic gates. Applies to both FPGA and ASIC devices.

RTL

(1) Register-transfer logic. A design methodology that moves data between registers. Data is latched in the registers at one or more stages along the path of signal propagation. The WISHBONE specification uses a synchronous RTL design methodology where all registers use a common clock that requires that each register be clocked with a common clock. (2) Register-transfer level. A description of computer operations where data transfers from register to register, latch to latch and through logic gates. (3) A level of description of a digital design in which the clocked behavior of the design is expressly described in terms of data transfers between storage elements, (which may be implied); and combinatorial logic; (which may represent any computing logic or arithmetic-logic-unit logic). RTL modeling allows design hierarchy that represents a structural description of other RTL models.

Shared Bus Interconnection

The shared bus interconnection is a system where a MASTER initiates addressable bus cycles to a target SLAVE. Traditional buses such as VMEbus and PCI bus use this type of interconnection. As a consequence of this architecture, only one MASTER at a time can use the interconnection resource (i.e. bus). Figure 1-8 shows an example of a WISHBONE shared bus interconnection.




Figure 1-8. Shared bus interconnection.

Silicon Foundry

A factory that produces integrated circuits.

SLAVE

A WISHBONE interface that is capable of receiving bus cycles. All systems based on the WISHBONE interconnect must have at least one SLAVE. Also see: ~~MASTER, SYSCON and INTERCON.~~

Soft Core

An IP Core that is delivered in the form of a hardware description language or schematic diagram.

SoC

Acronym for System-on-Chip. Also see: *System-on-Chip*.

Structured Design

(1) A popular method for managing complex projects. Often used with large project teams. When structured design practices are used, individual team members build and test small parts of the design with a common set of tools. Each sub-assembly is designed to a common standard. When all of the sub-assemblies have been completed, the full system can be integrated and tested. This approach makes it much easier to manage the design process. (2) Any disciplined approach to design that adheres to specified rules based on principles such as modularity and top-down design.

Switched Fabric Interconnection

A type of interconnection that uses large numbers of crossbar switches. These are organized into arrays that resemble the threads in a fabric. The resulting system is a network of redundant interconnections.

SYSCON

A WISHBONE module ~~and interface.~~ ~~The SYSCON module is the only module in the design which may contain the SYSCON interface.~~ ~~The SYSCON interface is the only interface in the design which may drive the~~ that drives the system clock [CLK_O] and reset [RST_O] signals.

System-on-Chip (SoC)

A method by which whole systems are created on a single integrated circuit chip. In many cases, this requires the use of IP cores which have been designed by multiple IP core providers. System-on-Chip is similar to traditional microcomputer bus systems whereby the individual components are designed, tested and built separately. The components are then integrated to form a finished system.

Tag

One or more characters or signals associated with a set of data, containing information about the set. Also see: tag type.

Tag Type~~AG TYPE~~

A special class of signals called *tag types* are also used. Tag types themselves are not signals used by the WISHBONE interface. Rather, they represent classes of signals. When a user defined signal is specified, it is assigned a tag type that indicates the precise timing to which the signal must conform. This simplifies the creation of new signals. There are three basic tag types. These include address tags, data tags and cycle tags. These allow additional information to be attached to an address transfer, a data transfer or a bus cycle (respectively). The uppercase form TAG TYPE is used when specifying a tag type in the WISHBONE DATASHEET. For example, TAG TYPE: TGA O() indicates an address tag. Also see: *address tag*, *data tag* and *cycle tag*.

Target Device

The semiconductor type (or technology) onto which the IP core design is impressed. Typical examples include FPGA and ASIC target devices.

Three-State Bus Interconnection

A microcomputer bus interconnection that relies upon three-state bus drivers. Often used to reduce the number of interconnecting signal paths through connector and IC pins. Three state buffers can assume a logic low state ('0' or 'L'), a logic high state ('1' or 'H') or a high impedance state ('Z'). Three-state buffers are sometimes called Tri-State® buffers. Tri-State® is a registered trademark of National Semiconductor Corporation. Also see: *multiplexor interconnection*.

Variable Clock Generator

A type of SYSCON ~~module~~ interface where the frequency of [CLK_O] can be changed dynamically. The frequency can be changed by way of a programmable phase-lock-loop (PLL) circuit or other control mechanism. Among other things, this technique is used to reduce the power consumption of the circuit. In WISHBONE, the variable clock generator capability is optional. Also see: *gated clock generator* and *variable timing specification*.

Variable Interconnection

A microcomputer bus interconnection that *can* be changed without causing incompatibilities between bus modules (or SoC/IP cores). Also called a dynamic interconnection. An example of a variable interconnection bus is the WISHBONE SoC architecture. Also see: *fixed interconnection*.

Variable Timing Specification

A timing specification that is not fixed. In WISHBONE, variable timing can be achieved in a number of ways. For example, the system integrator can select the ~~SYSCON~~ frequency ~~rate~~ of [CLK_O] by enforcing a timing specification during the circuit design. Variable timing can also be achieved during circuit operation with a variable clock generator. Also see: *gated clock generator* and *variable clock generator*.

Verilog®

A textual based hardware description language (HDL) intended for use in circuit design. The Verilog® language is both a synthesis and a simulation tool. Verilog® was originally a proprietary language first conceived in 1983 at Gateway Design Automation (Acton, MA), and was later refined by Cadence Corporation. It has since been greatly expanded and refined, and much of it has been placed into the public domain. Complete descriptions of the language can be found in the IEEE 1364 specification.

VHDL

Acronym for: VHSIC Hardware Description Language. [VHSIC: Very High Speed Integrated Circuit]. A textual based computer language intended for use in circuit design. The VHDL language is both a synthesis and a simulation tool. Early forms of the language emerged from US Dept. of Defense ARPA projects in the 1960's, and have since been greatly expanded and refined. Complete descriptions of the language can be found in the IEEE 1076, IEEE 1073.3, IEEE 1164 specifications.

VMEbus

Acronym for: Versa Module Eurocard bus. A popular microcomputer (board) bus. While this specification is very flexible, it isn't practical for SoC applications.

WISHBONE DATASHEET

A type of documentation required for WISHBONE compatible IP cores. This helps the end user understand the detailed operation of the core, and how to connect it to other cores. The WISHBONE DATASHEET can be included as part of an IP core technical reference manual, or as part of the IP core hardware description.

WISHBONE Signal

A signal that is defined as part of the WISHBONE specification. Non-WISHBONE signals can also be used on the IP core, but are not defined as part of this specification. For example, [ACK_O] is a WISHBONE signal, but [CLK100_I] is not.

WISHBONE Logo

A logo that, when affixed to a document, indicates that the associated SoC component is compatible with the WISHBONE standard.

Wrapper

A circuit element that converts a non-WISHBONE IP Core into a WISHBONE compatible IP Core. For example, consider a 16-byte synchronous memory primitive that is provided by an IC vendor. The memory primitive can be made into a WISHBONE compatible SLAVE by layering a circuit over the memory primitive, thereby creating a WISHBONE compatible SLAVE. A wrapper is analogous to a technique used to convert software written in 'C' to that written in 'C++'.

1.8 References

IEEE 100: The Authoritative Dictionary of IEEE Standards Terms, Seventh Edition. IEEE Press 2000.

Chapter 2 – Interface Specification

This chapter describes the signaling method between the MASTER ~~interface~~, SLAVE ~~interface~~, and SYSCON ~~module and INTERCON interfaces~~. This includes numerous options which may or may not be present on a particular interface. Furthermore, it describes a minimum level of required documentation that must be ~~must be~~ created for each IP core.

2.1 Required Documentation for IP Cores

WISHBONE compatible IP cores ~~must~~ include documentation that describes the interface. This helps the end user understand the operation of the core, and how to connect it to other cores. This documentation takes the form of a WISHBONE DATASHEET. It can be included as part of the IP core technical reference manual, it can be embedded in source code or it can take other forms as well.

2.1.1 General Requirements for the WISHBONE DATASHEET

RULE 2.00

Each WISHBONE compatible IP core MUST include a WISHBONE DATASHEET as part of the IP core documentation.

~~RULE 2.05~~

~~The WISHBONE DATASHEET MUST include the revision level of the WISHBONE specification to which it was designed.~~

~~RULE 2.10~~

~~The WISHBONE DATASHEET MUST indicate whether it is a MASTER, SLAVE, SYSCON or INTERCON interface. Furthermore, it MUST indicate the types of bus cycles it supports.~~

RULE 2.15

The WISHBONE DATASHEET for MASTER and, SLAVE ~~and INTERCON~~ interfaces MUST include the following information:

- (1) The revision level of the WISHBONE specification to which is was designed.
- (2) The type of interface: MASTER or SLAVE.
- (3) The signal names that are defined for the WISHBONE SoC interface. If a signal name is different than that defined in this specification, then it MUST be cross-referenced to the corresponding signal name which is used in this specification.

- (14) If a MASTER supports the optional [ERR_I] signal, then the WISHBONE DATASHEET MUST describe how it reacts in response to the signal. If a SLAVE supports the optional [ERR_O] signal, then the WISHBONE DATASHEET MUST describe the conditions under which the signal is generated.
- (52) If a MASTER supports the optional [RTY_I] signal, then the WISHBONE DATASHEET MUST describe how it reacts in response to the signal. If a SLAVE supports the optional [RTY_O] signal, then the WISHBONE DATASHEET MUST describe the conditions under which the signal is generated.
- (63) All interfaces that support ~~the tag signal(s) [TAGN_I] or [TAGN_O] signal(s)~~ MUST describe ~~their name, TAG TYPE and operation of the tag use~~ in the WISHBONE DATASHEET.
- (74) The WISHBONE DATASHEET MUST indicate the port size. The port size MUST be indicated as: 8-bit, 16-bit, 32-bit or 64-bit.
- (85) The WISHBONE DATASHEET MUST indicate the port granularity. The granularity MUST be indicated as: 8-bit, 16-bit, 32-bit or 64-bit.
- (96) The WISHBONE DATASHEET MUST indicate the maximum operand size. The maximum operand size MUST be indicated as: 8-bit, 16-bit, 32-bit or 64-bit. If the maximum operand size is unknown, then the maximum operand size shall be the same as the granularity.
- (107) The WISHBONE DATASHEET MUST indicate the data transfer ordering. The ordering MUST be indicated as BIG ENDIAN or LITTLE ENDIAN. When the port size equals the granularity, then the interface shall be specified as BIG/LITTLE ENDIAN ~~and/or LITTLE ENDIAN.~~ [When the port size equals the granularity, then BIG ENDIAN and LITTLE ENDIAN transfers are identical].
- (118) The WISHBONE DATASHEET MUST indicate the sequence of data transfer through the port. If the sequence of data transfer is not known, then the datasheet MUST indicate that it asis UNDEFINED ~~undefined~~.
- (129) The WISHBONE DATASHEET MUST indicate if there are any constraints on the [CLK_I] signal. These constraints include (but are not limited to) clock frequency, application specific timing constraints, the use of gated clocks or the use of variable clock generators.

2.1.2 Signal Naming

RULE 2.20

Signal names MUST adhere to the rules of the native tool in which the IP core is designed.

PERMISSION 2.00

Any signal name MAY be used to describe the WISHBONE signals.

OBSERVATION 2.00

Most hardware description languages (such as VHDL or Verilog®) have naming conventions. For example, the VHDL hardware description language defines the alphanumeric symbols which may be used. Furthermore, it states that UPPERCASE and LOWERCASE characters may be used in a signal name.

RECOMENDATION 2.00

It is recommended that the interface use the signal names defined in this document.

OBSERVATION 2.05

Core integration is simplified if the signal names match those given in this specification. However, in some cases (such as IP cores with multiple WISHBONE interconnects) they cannot be used. The use of non-standard signal names will not result in any serious integration problems since all hardware description tools allow signals to be renamed.

RULE 2.25

~~The WISHBONE DATASHEET MUST include the signal names that are defined for a WISHBONE SoC interface. If a signal name is different than defined in this specification then it MUST be cross-referenced to the corresponding signal name which is used in this specification.~~

PERMISSION 2.05

Non-WISHBONE signals MAY be used with IP core interfaces.

OBSERVATION 2.15

Most IP cores will include non-WISHBONE signals. These are outside the scope of this specification, and no attempt is made to govern them. For example, a disk controller IP core could have a WISHBONE interface on one end and a disk interface on the other. In this case the specification does not dictate any technical requirements for the disk interface signals.

2.1.3 Logic Levels

RULE 2.30

All WISHBONE interface signals MUST use active high logic.

OBSERVATION 2.10

In general, the use of active low signals does not present a problem. However, RULE 2.30 is included because some tools (especially schematic entry tools) do not have a standard way of indicating an active low signal. For example, a reset signal could be named [#RST_I], [/RST_I] or [N_RST_I]. This was found to cause confusion among users and incompatibility between modules. This constraint should not create any undue difficulties, as the system integrator can invert any signals before use by the WISHBONE interface.

PERMISSION 2.05

~~Non-WISHBONE signals MAY be used with IP core interfaces.~~

OBSERVATION 2.15

~~Most IP cores will include non-WISHBONE signals. These are outside the scope of this specification, and no attempt is made to govern them. For example, a disk controller IP core could have a WISHBONE interface on one end and a disk interface on the other. In this case the specification does not dictate any technical requirements for the disk interface signals.~~

OBSERVATION 2.20

~~[TAGN_I] and [TAGN_O] are user defined signals that must adhere to the timing specifications given in this document.~~

2.2 WISHBONE Signal Description

This section describes the signals used in the WISHBONE interconnect. Some of these signals are optional, and may or may not be present on a specific interface.

2.2.1 SYSCON Module Signals

CLK_O

The system clock output [CLK_O] is generated by the SYSCON ~~module~~interface. It coordinates all activities for the internal logic within the WISHBONE interconnect. The INTERCON ~~mod-~~ule connects the [CLK_O] output to the [CLK_I] input on MASTER and SLAVE interfaces.

RST_O

The reset output [RST_O] is generated by the SYSCON ~~module~~interface. It forces all WISHBONE interfaces to restart. All internal self-starting state machines are forced into an initial state. The INTERCON connects the [RST_O] output to the [RST_I] input on MASTER and SLAVE interfaces.

2.2.2 Signals Common to MASTER and SLAVE Interfaces

CLK_I

The clock input [CLK_I] coordinates all activities for the internal logic within the WISHBONE interconnect. All WISHBONE output signals are registered at the rising edge of [CLK_I]. All WISHBONE input signals must be stable before the rising edge of [CLK_I].

DAT_I()

The data input array [DAT_I()] is used to pass binary data. The array boundaries are determined by the port size, with a maximum port size of 64-bits (e.g. [DAT_I(63..0)]). Also see the [DAT_O()] and [SEL_O()] signal descriptions.

DAT_O()

The data output array [DAT_O()] is used to pass binary data. The array boundaries are determined by the port size, with a maximum port size of 64-bits (e.g. [DAT_I(63..0)]). Also see the [DAT_I()] and [SEL_O()] signal descriptions.

RST_I

The reset input [RST_I] forces the WISHBONE interface to restart. Furthermore, all internal self-starting state machines will be forced into an initial state. This signal only resets the WISHBONE interface. It is not required to reset other parts of an IP core (although it may be used that way).

TGD_I()

Data tag type [TGD_I()] is used on MASTER and SLAVE interfaces. It contains information that is associated with a data lines [DAT_I()], and is qualified by signal [STB_I]. For example, parity protection, error correction and time stamp information can be attached to the data bus. These tag bits simplify the task of defining new signals because their timing (in relation to every bus cycle) is pre-defined by this specification. The name and operation of a data tag must be defined in the WISHBONE DATASHEET.

TGD_O()

Data tag type [TGD_O()] is used on MASTER and SLAVE interfaces. It contains information that is associated with a data lines [DAT_O()], and is qualified by signal [STB_O]. For example, parity protection, error correction and time stamp information can be attached to the data bus. These tag bits simplify the task of defining new signals because their timing (in relation to every bus cycle) is pre-defined by this specification. The name and operation of a data tag must be defined in the WISHBONE DATASHEET.

TAGN_I

The tag input(s) [TAGN_I] are user defined, and are used with either MASTER or SLAVE interfaces. 'N' in this signal name refers to a tag number because multiple tags may be used (e.g. [TAG3_I]). Tag inputs are used whenever an IP core needs specific information from the interconnection. For example, a MASTER can be designed to monitor the state of a FIFO.

TAGN_O

The tag output(s) [TAGN_O] are user defined, and are used with either MASTER or SLAVE interfaces. For example, the tag output(s) can be used to indicate the type of data transfer in progress. Furthermore, 'N' in this signal name refers to a tag number because multiple tags may be used. For example, [TAG1_O] may indicate a valid data transfer cycle, [TAG2_O] may indicate an interrupt acknowledge cycle and so on. The exact meaning of each tag is defined by the IP core provider in the WISHBONE DATASHEET.

2.2.3 MASTER Signals

ACK_I

The acknowledge input [ACK_I], when asserted, indicates the termination of a normal bus cycle. Also see the [ERR_I] and [RTY_I] signal descriptions.

ADR_O(63..0)

The address output array [ADR_O(63..0)] is used to pass a binary address, with the most significant address bit at the higher numbered end of the signal array. The maximum size of the array is specified as [ADR_O(63..0)]. However, the higher array boundary is specific to the address width of the core, and the lower array boundary is determined by the lower array boundary is specific to the data port size (e.g. the maximum array size on a 32-bit data port is [ADR_O(63..2)]. The higher array boundary is core specific. In some cases (such as FIFO interfaces) the array may not be present on the interface.

CYC_O

The cycle output [CYC_O], when asserted, indicates that a valid bus cycle is in progress. The signal is asserted for the duration of all bus cycles. For example, during a BLOCK transfer cycle there can be multiple data transfers. The [CYC_O] signal is asserted during the first data transfer, and remains asserted until the last data transfer. The [CYC_O] signal is useful for interfaces with multi-port interfaces (such as dual port memories). In these cases, the [CYC_O] signal requests use of a common bus from an arbiter. Once the arbiter grants the bus to the MASTER, it is held until [CYC_O] is negated.

DAT_I(63..0)

The data input array [DAT_I(63..0)] is used to pass binary data. The array boundaries are determined by the port size. Also see the [DAT_O(63..0)] and [SEL_O(7..0)] signal descriptions.

DAT_O(63..0)

~~The data output array [DAT_O(63..0)] is used to pass binary data. The array boundaries are determined by the port size. Also see the [DAT_I(63..0)] and [SEL_O(7..0)] signal descriptions.~~

ERR_I

The error input [ERR_I] indicates an abnormal cycle termination. The source of the error, and the response generated by the MASTER is defined by the IP core supplier. Also see the [ACK_I] and [RTY_I] signal descriptions.

RTY_I

The retry input [RTY_I] indicates that the interface is not ready to accept or send data, and that the cycle should be retried. When and how the cycle is retried is defined by the IP core supplier. Also see the [ERR_I] and [RTY_I] signal descriptions.

SEL_O(7..0)

The select output array [SEL_O(7..0)] indicates where valid data is expected on the [DAT_I(63..0)] signal array during READ cycles, and where it is placed on the [DAT_O(63..0)] signal array during WRITE cycles. The array boundaries are determined by the granularity of a port. For example, if 8-bit granularity is used on a 64-bit port, then there would be an array of eight select signals with boundaries of [SEL_O(7..0)]. Each individual select signal correlates to one of eight active bytes on the 64-bit data port. For more information about [SEL_O()], please refer to the data organization section in Chapter 3 of this specification. Also see the [DAT_I(63..0)], [DAT_O(63..0)] and [STB_O] signal descriptions.

STB_O

The strobe output [STB_O] indicates a valid data transfer cycle. It is used to qualify various other signals on the interface such as [SEL_O(7..0)]. The SLAVE ~~must~~ asserts either the [ACK_I], [ERR_I] or [RTY_I] signals in response to every assertion of the [STB_O] signal.

TGA_O()

Address tag type [TGA_O()] contains information associated with address lines [ADR_O()], and is qualified by signal [STB_O]. For example, address size (24-bit, 32-bit etc.) and memory management (protected vs. unprotected) information can be attached to an address. These tag bits simplify the task of defining new signals because their timing (in relation to every bus cycle) is defined by this specification. The name and operation of an address tag must be defined in the WISHBONE DATASHEET.

TGC_O()

Cycle tag type [TGC_O()] contains information associated with bus cycles, and is qualified by signal [CYC_O]. For example, data transfer, interrupt acknowledge and cache control cycles can be uniquely identified with the cycle tag. They can also be used to discriminate between WISHBONE SINGLE, BLOCK and RMW cycles. These tag bits simplify the task of defining new signals because their timing (in relation to every bus cycle) is defined by this specification. The name and operation of a cycle tag must be defined in the WISHBONE DATASHEET.

WE_O

The write enable output [WE_O] indicates whether the current local bus cycle is a READ or WRITE cycle. The signal is negated during READ cycles, and is asserted during WRITE cycles.

2.2.4 SLAVE Signals

ACK_O

The acknowledge output [ACK_O], when asserted, indicates the termination of a normal bus cycle. Also see the [ERR_O] and [RTY_O] signal descriptions.

ADR_I(63..0)

~~The address input array [ADR_I(63..0)] is used to pass a binary address, with the most significant address bit at the higher numbered end of the signal array. The maximum size of the array is specified as [ADR_I(63..0)]. However, the higher array boundary is specific to the address width of the core, and the lower array boundary is determined by the data port size (e.g. the maximum array size on a 32-bit data port is [ADR_O(63..2)]. The lower array boundary is specific to the data port size. The higher array boundary is core specific.~~ In some cases (such as FIFO interfaces) the array may not be present on the interface.

CYC_I

The cycle input [CYC_I], when asserted, indicates that a valid bus cycle is in progress. The signal is asserted for the duration of all bus cycles. For example, during a BLOCK transfer cycle there can be multiple data transfers. The [CYC_I] signal is asserted during the first data transfer, and remains asserted until the last data transfer.

DAT_I(63..0)

~~The data input array [DAT_I(63..0)] is used to pass binary data. The array boundaries are determined by the port size. Also see the [DAT_O(63..0)] and [SEL_O(7..0)] signal descriptions.~~

DAT_O(63..0)

~~The data output array [DAT_O(63..0)] is used to pass binary data. The array boundaries are determined by the port size. Also see the [DAT_I(63..0)] and [SEL_O(7..0)] signal descriptions.~~

ERR_O

The error output [ERR_O] indicates an abnormal cycle termination. The source of the error, and the response generated by the MASTER is defined by the IP core supplier. Also see the [ACK_O] and [RTY_O] signal descriptions.

RTY_O

The retry output [RTY_O] indicates that the interface is not ready to accept or send data, and that the cycle should be retried. When and how the cycle is retried is defined by the IP core supplier. Also see the [ERR_O] and [RTY_O] signal descriptions.

SEL_I(7..0)

The select input array [SEL_I(7..0)] indicates where valid data is placed on the [DAT_I(63..0)] signal array during WRITE cycles, and where it should be present on the [DAT_O(63..0)] signal array during READ cycles. The array boundaries are determined by the granularity of a port. For example, if 8-bit granularity is used on a 64-bit port, then there would be an array of eight select signals with boundaries of [SEL_I(7..0)]. Each individual select signal correlates to one of eight active bytes on the 64-bit data port. For more information about [SEL_I()], please refer to the data organization section in Chapter 3 of this specification. –Also see the [DAT_I(63..0)], [DAT_O(63..0)] and [STB_I] signal descriptions.

STB_I

The strobe input [STB_I], when asserted, indicates that the SLAVE is selected. A SLAVE shall respond to other WISHBONE signals only when this [STB_I] is asserted, ~~—~~(except for the [RST_I] signal which should always be responded to). The SLAVE ~~must~~ asserts either the [ACK_O], [ERR_O] or [RTY_O] signals in response to every assertion of the [STB_I] signal.

TGA_I

Address tag type [TGA_I()] contains information associated with address lines [ADR_I()], and is qualified by signal [STB_I]. For example, address size (24-bit, 32-bit etc.) and memory management (protected vs. unprotected) information can be attached to an address. These tag bits simplify the task of defining new signals because their timing (in relation to every bus cycle) is pre-defined by this specification. The name and operation of an address tag must be defined in the WISHBONE DATASHEET.

TGC_I()

Cycle tag type [TGC_I()] contains information associated with bus cycles, and is qualified by signal [CYC_I]. For example, data transfer, interrupt acknowledge and cache control cycles can be uniquely identified with the cycle tag. They can also be used to discriminate between WISHBONE SINGLE, BLOCK and RMW cycles. These tag bits simplify the task of defining new signals because their timing (in relation to every bus cycle) is pre-defined by this specification. The name and operation of a cycle tag must be defined in the WISHBONE DATASHEET.

WE_I

The write enable input [WE_I] indicates whether the current local bus cycle is a READ or WRITE cycle. The signal is negated during READ cycles, and is asserted during WRITE cycles.

Chapter 3 – Bus Cycles

WISHBONE bus cycles are described in terms of their general operation, reset operation, handshaking protocol and the data organization during transfers. Additional requirements for bus cycles (especially those relating to the common clock) can be found in the timing specifications in Chapter 4.

3.1 General Operation

Each MASTER and SLAVE are interconnected with a set of signals that permit them to exchange data. For descriptive purposes these signals are cumulatively known as a *bus*, and are contained within a functional module called the INTERCON. Address, data and other information is impressed upon this bus in the form of *bus cycles*.

3.1.1 Reset Operation

All hardware interfaces **are must be** initialized to a pre-defined state. This is accomplished with the reset signal [RST_O], which can be asserted at any time. It is also used for test simulation purposes by initializing all self-starting state machines and counters which may be used in the design. The reset signal [RST_O] is driven by the SYSCON **module interface**. It is connected to the [RST_I] signal on all MASTER and SLAVE interface. Figure 3-1 shows the reset cycle.

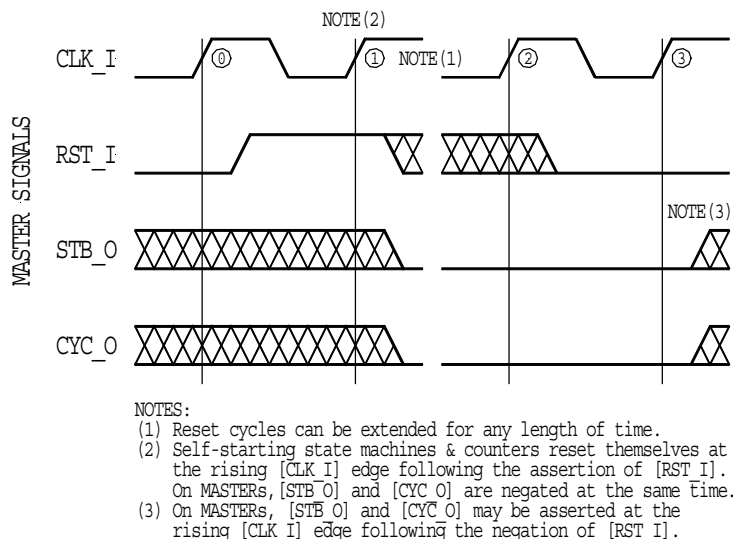


Figure 3-1. Reset cycle.

RULE 3.00

All WISHBONE interfaces MUST initialize themselves at the rising [CLK_I \ominus] edge following the assertion of [RST_I \ominus]. They MUST stay in the initialized state until the rising [CLK_I \ominus] edge that follows the negation of [RST_I \ominus].

RULE 3.05

[RST_I] MUST be asserted for at least one complete clock cycle on all WISHBONE interfaces.

PERMISSION 3.00

[RST_I \ominus] MAY be asserted for more than one clock cycle, and MAY be asserted indefinitely.

RULE 3.10

All WISHBONE interfaces MUST be capable of reacting to [RST_I] at any time.

RULE 3.15

All self-starting state machines and counters in WISHBONE interfaces MUST initialize themselves at the rising [CLK_I] edge following the assertion of [RST_I]. They MUST stay in the initialized state until the rising [CLK_I] edge that follows the negation of [RST_I].

OBSERVATION 3.00

In general, self-starting state machines do not need to be initialized. However, this may cause problems because some simulators may not be sophisticated enough to find an initial starting point for the state machine. Furthermore, self-starting state machines can go through an indeterminate number of initialization cycles before finding their starting state, thereby making it difficult to predict their behavior at start-up time. The initialization rule prevents both problems by forcing all state machines to a pre-defined state in response to the assertion of [RST_I].

RULE 3.20

The following MASTER signals MUST be negated at the rising [CLK_I] edge following the assertion of [RST_I], and MUST stay in the negated state until the rising [CLK_I] edge that follows the negation of [RST_I]: [STB_O], [CYC_O]. The state of all other MASTER signals are undefined in response to a reset cycle.

OBSERVATION 3.05

On MASTER interfaces [STB_O] and [CYC_O] may be asserted beginning at the rising [CLK_I] edge following the negation of [RST_I].

OBSERVATION 3.10

SLAVE interfaces automatically negate [ACK_O], [ERR_O] and [RTY_O] when their [STB_I] is negated.

RECOMENDATION 3.00

Design SYSCON ~~modules~~~~circuits~~ so that they assert [RST_O] during a power-up condition. [RST_O] should remain asserted until all voltage levels and clock frequencies in the system are stabilized. When negating [RST_O], do so in a synchronous manner that conforms to this specification.

OBSERVATION 3.15

If a gated clock generator is used, and if the clock is stopped, then the WISHBONE interface is not capable of responding to its [RST_I] signal.

SUGGESTION 3.00

Some circuits require an *asynchronous* reset capability. If an IP core or other SoC component requires an asynchronous reset, then define it as a non-WISHBONE signal. This prevents confusion with the WISHBONE reset [RST_I \ominus] signal, which uses a purely synchronous protocol, and need be applied only to the WISHBONE interface.

OBSERVATION 3.20

All WISHBONE *interfaces* ~~must~~ respond to the reset signal. However, the IP Core connected to a WISHBONE interface does not necessarily need to respond to the reset signal.

3.1.2 Handshaking Protocol

All bus cycles use a handshaking protocol between the MASTER and SLAVE interfaces. As shown in Figure 3-2, the MASTER asserts [STB_O] when it is ready to transfer data. [STB_O] remains asserted until the SLAVE asserts one of the cycle terminating signals [ACK_I], [ERR_I] or [RTY_I]. At every rising edge of [CLK_I] the terminating signal is sampled. If it is asserted, then [STB_O] is negated. This gives both the MASTER and SLAVE interfaces the possibility to control the rate at which data is transferred.

If the SLAVE guarantees it can keep pace with all MASTER interfaces, and if the [ERR_I] and [RTY_I] signals are not used, then the [ACK_I] signal may be tied to the SLAVE's [STB_I] input. The interface will function normally under these circumstances.




Figure 3-2. Local bus handshaking protocol.

Most of the examples in this specification describe the use of [ACK_I] to terminate a local bus cycle. However, the SLAVE can optionally terminate the cycle with an error [ERR_O], or request that the cycle be retried [RTY_O].

All interfaces include the [ACK_I] terminator signal. Asserting this signal during a bus cycle causes it to terminate normally.

Asserting the [ERR_I] signal during a bus cycle will terminate the cycle. It also serves to notify the MASTER that an error occurred during the cycle. This signal is generally used if an error was detected by SLAVE logic circuitry. For example, if the SLAVE is a parity-protected memory, then the [ERR_I] signal can be asserted if a parity fault is detected. This specification does not dictate what the MASTER will do in response to [ERR_I].

Asserting the optional [RTY_I] signal during a bus cycle will terminate the cycle. It also serves to notify the MASTER that the current cycle should be aborted, and retried at a later time. This signal is generally used for shared memory and bus bridges. In these cases SLAVE circuitry ~~would~~ asserts [RTY_I] if the local resource is busy. This specification does not dictate when or how the MASTER will respond to [RTY_I].

RULE 3.25

As a minimum, the MASTER interface MUST include the following signals: [ACK_I], [CLK_I], [CYC_O], [RST_I] and [STB_O]. As a minimum, the SLAVE interface MUST include the following signals: [ACK_O], [CLK_I] and [RST_I]. All other signals are optional.

PERMISSION 3.05

MASTER and SLAVE interfaces MAY be designed to support the [ERR_I] and [ERR_O] signals. In these cases, the SLAVE asserts [ERR_O] to indicate that an error has occurred during the bus cycle. This specification does not dictate what the MASTER does in response to [ERR_I].

PERMISSION 3.10

MASTER and SLAVE interfaces MAY be designed to support the [RTY_I] and [RTY_O] signals. In these cases, the SLAVE asserts [RTY_O] to indicate that the interface is busy, and that the bus cycle should be retried at a later time. This specification does not dictate what the MASTER will do in response to [RTY_I].

RULE 3.30

If a SLAVE supports the [ERR_O] or [RTY_O] signals, then the SLAVE MUST NOT assert more than one of the following signals at any time: [ACK_O], [ERR_O] or [RTY_O].

RECOMMENDATION 3.00

Design WISHBONE MASTER interfaces so that there are no intermediate logic gates between a registered flip-flop and the signal outputs on [STB_O] and [CYC_O]. Delay timing for [STB_O] and [CYC_O] are very often the most critical paths in the system. This prevents sloppy design practices from slowing down the interconnect because of added delays on these two signals.

RULE 3.35

SLAVE interfaces MUST be designed so that the [ACK_O], [ERR_O] and [RTY_O] signals are asserted and negated in response to the assertion and negation of [STB_I]. Furthermore, this activity MUST occur asynchronous to the [CLK_I] signal (i.e. there is a combinatorial logic path between [STB_I] and [ACK_O], etc.).

OBSERVATION 3.25

The asynchronous logic requirement assures that the interface can accomplish one data transfer per clock cycle. Furthermore, it simplifies the design of arbiters in multi-MASTER applications.

RECOMMENDATION 3.05

Design interconnection logic to prevent deadlock conditions when MASTER accesses are made to unused address locations. One solution to this problem is to include a watchdog timer function that monitors the MASTER's [STB_O] signal, and asserts [ERR_I] or [RTY_I] if the cycle exceeds some pre-defined time limit.

PERMISSION 3.15

Under certain circumstances SLAVE interfaces MAY be designed to hold [ACK_O] in the asserted state. This situation occurs on point-to-point interfaces where there is a single SLAVE on the interface, and that SLAVE always operates without wait states.

RULE 3.40

MASTER interfaces MUST be designed to operate normally when SLAVE interface holds [ACK_I] in the asserted state.

3.1.3 Use of [STB_O]

RULE 3.45

MASTER interfaces MUST qualify the following signals with [STB_O]: [ADR_O], [DAT_O()], [SEL_O()], [WE_O], [SEL_O] and [TAGN_O].

RULE 3.50

MASTER interfaces MUST assert [CYC_O] for the duration of SINGLE READ / WRITE, BLOCK and RMW cycles. [CYC_O] MUST be asserted no later than the rising [CLK_I] edge that qualifies the assertion of [STB_O]. [CYC_O] MUST be negated no earlier than the rising [CLK_I] edge that qualifies the negation of [STB_O].

3.1.4 Use of [ACK_O], [ERR_O] and [RTY_O]

RULE 3.55

SLAVE interfaces MUST qualify the following signals with [ACK_O], [ERR_O] or [RTY_O]: [DAT_O()].

3.1.5 Use of TAG TYPES ~~[TAGN_I]~~ and ~~[TAGN_O]~~ Signals

The WISHBONE interface can be modified with user defined signals. This is done with a technique known as tagging. Tags are a well known concept in the microcomputer bus industry. They allow user defined information to be associated with an address, a data word or a bus cycle.

All tag signals must conform to set of guidelines known as TAG TYPES. Table 3-1 lists all of the defined TAG TYPES along with their associated data set and signal waveform. When a tag is added to an interface it is assigned a TAG TYPE from the table. This explicitly defines how the tag operates. This information must also be included in the WISHBONE DATASHEET.

~~The TAG signals [TAGN_I] and [TAGN_O] are used by both the MASTER and SLAVE interfaces. They are used for three purposes: (a) to tag data with information such as parity or time stamps, (b) to identify specialty bus cycles (like interrupts or cache control operations) and (c) to communicate with the bus interconnection. These signals are user defined.~~

~~For example, the designer of a MASTER may wish to add parity check bits to its bus cycle. In this case a [TAGN_O] signal is defined by the IP core designer, and logic would be created to generate the bit. Furthermore, the signal would be described in the WISHBONE DATASHEET.~~

~~In another example, the designer of a SLAVE interface may wish to notify the bus interconnection logic with the size of its data interface. In this case a [TAGN_O] signal is defined by the IP core designer, and logic would be created to reflect the bus size. The signal would also be described in the WISHBONE DATASHEET.~~

<u>Table 3-1. TAG TYPES.</u>				
	<u>MASTER</u>		<u>SLAVE</u>	
<u>Description</u>	<u>TAG TYPE</u>	<u>Associated With</u>	<u>TAG TYPE</u>	<u>Associated With</u>
<u>Address tag</u>	<u>TGA_O()</u>	<u>ADR_O()</u>	<u>TGA_I()</u>	<u>ADR_I()</u>
<u>Data tag, input</u>	<u>TGD_I()</u>	<u>DAT_I()</u>	<u>TGD_I()</u>	<u>DAT_I()</u>
<u>Data tag, output</u>	<u>TGD_O()</u>	<u>DAT_O()</u>	<u>TGD_O()</u>	<u>DAT_O()</u>
<u>Cycle tag</u>	<u>TGC_O()</u>	<u>Bus Cycle</u>	<u>TGC_I()</u>	<u>Bus Cycle</u>

~~For example, consider a MASTER interface where a parity protection bit named [PAR_O] is generated from an output data word on [DAT_O(15..0)]. It's an 'even' parity bit, meaning that it's asserted whenever there are an even number of '1's in the data word. If this signal were added to the interface, then the following information (in the WISHBONE DATASHEET) would be sufficient to completely define the timing of [PAR_O]:~~

~~SIGNAL NAME: PAR_O
DESCRIPTION: Even parity bit
MASTER TAG TYPE: TGD_O()~~

RULE 3.60

~~The [TAGN_I] and [TAGN_O] signals MUST adhere to the timing specifications given in this document.~~

RULE 3.57

~~All user defined tags MUST be assigned a TAG TYPE. Furthermore, they MUST adhere to the timing specifications given in this document for the indicated TAG TYPE.~~

PERMISSION 3.17

~~While all TAG TYPES are specified as arrays (with parenthesis '(')'), the actual tag MAY be a non-arrayed signal.~~

RECOMMENDATION 3.XX

~~If a MASTER interface supports more than one defined bus cycle over a common set of signal lines, then include a cycle tag to identify each type of bus cycle. This allows INTERCON and SLAVE interface circuits to discriminate between these bus cycles (if needed). Define the signals as TAG TYPE: [TGC_O()], using signal names of [SGL_O], [BLK_O] and [RMW_O] when identifying SINGLE, BLOCK and RMW cycles respectively.~~

3.2 SINGLE READ / WRITE Cycles

The SINGLE READ / WRITE cycles perform one data transfer at a time. These are the basic cycles used to perform data transfers on the WISHBONE interconnect.

RULE 3.65

All MASTER and SLAVE interfaces that support SINGLE READ or SINGLE WRITE cycles MUST conform to the timing requirements given in sections 3.2.1 and 3.2.2.

PERMISSION 3.20

MASTER and SLAVE interfaces MAY be designed so that they do not support the SINGLE READ or SINGLE WRITE cycles.

3.2.1 SINGLE READ Cycle

Figure 3-3 shows a SINGLE READ cycle. The bus protocol works as follows:

CLOCK EDGE 0: MASTER presents a valid address on [ADR_O()] and ~~[TGA_O()TAGN_O]~~.
MASTER negates [WE_O] to indicate a READ cycle.
MASTER presents bank select [SEL_O()] to indicate where it expects data.
MASTER asserts [CYC_O] and [TGC_O()] to indicate the start of the cycle.
MASTER asserts [STB_O] to indicate the start of the phase. ~~to qualify [ADR_O()], [SEL_O()] and [WE_O].~~

SETUP, EDGE 1: SLAVE decodes inputs, and responding SLAVE responds by asserting as-
serts [ACK_I].
SLAVE presents valid data on [DAT_I()] and [TGD_I()].
SLAVE asserts [ACK_I] in response to [STB_O] to indicate valid data.
~~SLAVE presents [TAGN_O].~~
~~MASTER monitors [TAGN_I].~~
MASTER monitors [ACK_I], and prepares to latch data on [DAT_I()] and
[TGD_I()].

Note: SLAVE may insert wait states (-WSS-) before asserting [ACK_I], thereby allowing it to throttle the cycle speed. Any number of wait states may be added.

CLOCK EDGE 1: MASTER latches data on [DAT_I()] and [TGD_I()].
~~MASTER latches [TAGN_I].~~
MASTER negates [STB_O] and [CYC_O] to indicate the end of the cycle.
~~SLAVE negates [ACK_I] in response to negated [STB_O].~~

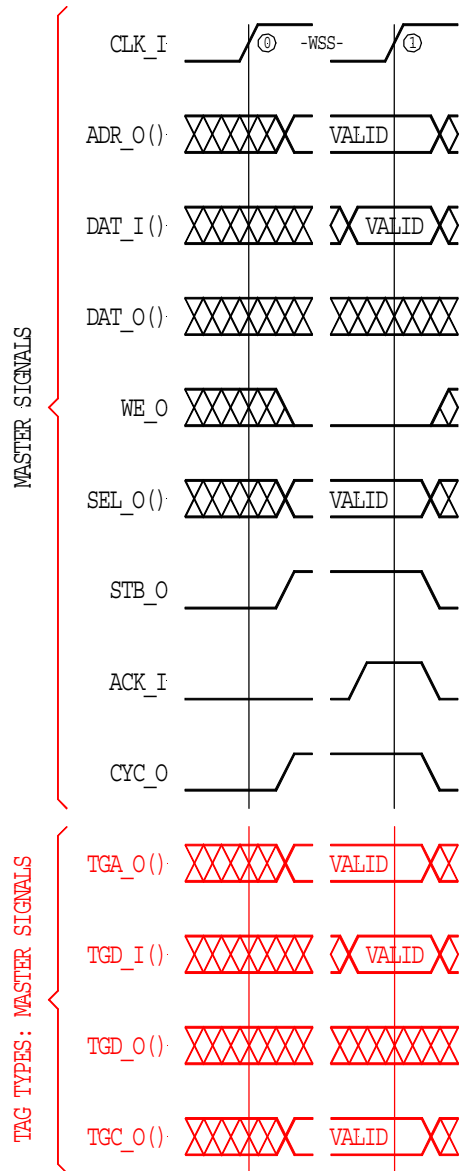


Figure 3-3. SINGLE READ cycle.

3.2.2 SINGLE WRITE Cycle

Figure 3-4 shows a SINGLE WRITE cycle. The bus protocol works as follows:

CLOCK EDGE 0: MASTER presents a valid address on [ADR_O()] and [TGA_O()TAGN_O].
MASTER presents valid data on [DAT_O()] and [TGD_O()].
MASTER asserts [WE_O] to indicate a WRITE cycle.
MASTER presents bank select [SEL_O()] to indicate where it sends data.
MASTER asserts [CYC_O] and [TGC_O()] to indicate the start of the cycle.
MASTER asserts [STB_O] to indicate the start of the phase to qualify [ADR_O()], [SEL_O()] and [WE_O].

SETUP, EDGE 1: SLAVE ~~decodes~~ inputs, and responding SLAVE asserts responds by asserting [ACK_I].
SLAVE ~~presents~~ prepares to latch data on [DAT_O()] and [TGD_O()].
SLAVE asserts [ACK_I] in response to [STB_O] to indicate latched data.
SLAVE presents [TAGN_O].
MASTER monitors [TAGN_I].
MASTER monitors [ACK_I], and prepares to terminate the cycle.

Note: SLAVE may insert wait states (-WSS-) before asserting [ACK_I], thereby allowing it to throttle the cycle speed. Any number of wait states may be added.

CLOCK EDGE 1: SLAVE latches data on [DAT_O()] and [TGD_O()].
MASTER latches [TAGN_I].
MASTER negates [STB_O] and [CYC_O] to indicate the end of the cycle.
SLAVE negates [ACK_I] in response to negated [STB_O].

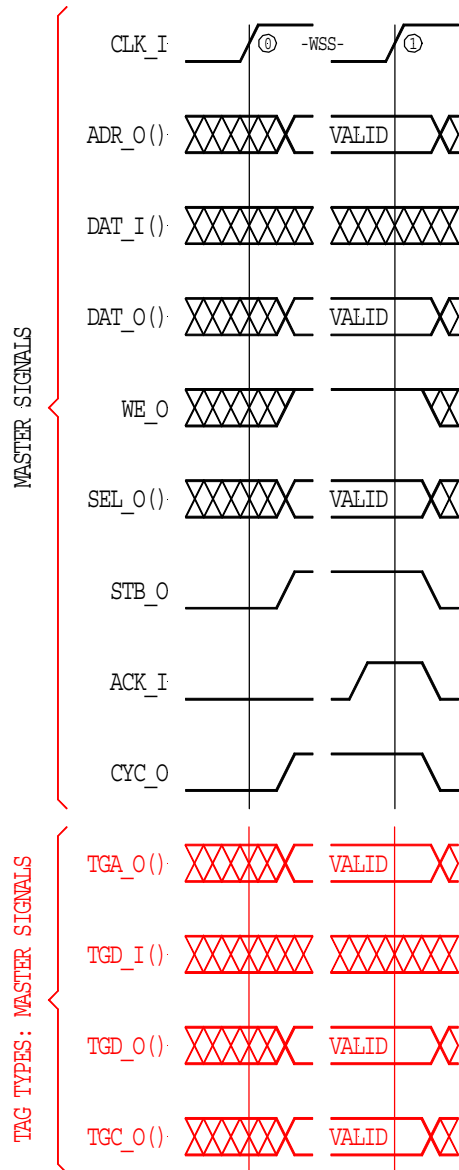


Figure 3-4. SINGLE WRITE cycle.

3.3 BLOCK READ / WRITE Cycles

The BLOCK transfer cycles perform multiple data transfers. They are very similar to single READ and WRITE cycles, but have a few special modifications to support multiple transfers.

During BLOCK cycles, the interface basically performs SINGLE READ/WRITE cycles as described above. However, the BLOCK cycles are modified somewhat so that these individual cycles (*called phases*) are combined together to form a single BLOCK cycle. This function is most useful when multiple MASTERS are used on the interconnect. For example, if the SLAVE is a shared (dual port) memory, then an arbiter for that memory can determine when one MASTER is done with it so that another can gain access to the memory.

As shown in Figure 3-5, the [CYC_O] signal is asserted for the duration of a BLOCK cycle. This signal can be used to request permission to access from a shared resource from a local arbiter, and hold the access until the end of the current cycle. During each of the data transfers *phases* (within the block transfer), the normal handshaking protocol between [STB_O] and [ACK_I] is maintained.




Figure 3-5. Use of [CYC_O] signal during BLOCK cycles.

It should be noted that the [CYC_O] signal does not necessarily rise and fall at the same time as [STB_O]. [CYC_O] may be asserted at the same time as [STB_O], or one or more [CLK_I] edges before [STB_O]. Similarly, [CYC_O] may be negated at the same time as [STB_O], or after an indeterminate number of [CLK_I] cycles.

RULE 3.70

All MASTER and SLAVE interfaces that support BLOCK cycles MUST conform to the timing requirements given in sections 3.3.1 and 3.3.2.

PERMISSION 3.25

MASTER and SLAVE interfaces MAY be designed so that they do not support the BLOCK cycles.

3.3.1 BLOCK READ Cycle

Figure 3-6 shows a BLOCK READ cycle. The BLOCK cycle is capable of a data transfer on every clock cycle. However, this example also shows how the MASTER and the SLAVE interfaces can both throttle the bus transfer rate by inserting wait states. A total of five transfers (phases) are shown. After the second transfer the MASTER inserts a wait state. After the fourth transfer the SLAVE inserts a wait state. The cycle is terminated after the fifth transfer. The protocol for this transfer works as follows:

CLOCK EDGE 0: MASTER presents a valid address on [ADR_O()] and [TGA_O()]~~[TAGN_O]~~.
MASTER negates [WE_O] to indicate a READ cycle.
MASTER presents bank select [SEL_O()] to indicate where it expects data.
MASTER asserts [CYC_O] and [TGC_O()] to indicate the start of the cycle.
MASTER asserts [STB_O] to indicate the start of the first phase.

Note: the MASTER ~~must~~ asserts [CYC_O] and/or [TGC_O()]~~[TAGN_O]~~ at, or anytime before, clock edge 1. ~~The use of [TAGN_O] is optional.~~

SETUP, EDGE 1: SLAVE decodes inputs, and responding SLAVE asserts responds by asserting [ACK_I].
SLAVE presents valid data on [DAT_I()] and [TGD_I()].
~~SLAVE presents [TAGN_O].~~
~~MASTER monitors [TAGN_I].~~
MASTER monitors [ACK_I], and prepares to latch ~~data on~~ [DAT_I()] and [TGD_I()].

CLOCK EDGE 1: MASTER latches data on [DAT_I()] and [TGD_I()].
~~MASTER latches [TAGN_I].~~
MASTER presents new [ADR_O()] and [TGA_O()]~~[TAGN_O]~~.
MASTER presents new bank select [SEL_O()] to indicate where data is.

SETUP, EDGE 2: SLAVE decodes inputs, and responds by asserting [ACK_I].
SLAVE presents valid data on [DAT_I()] and [TGD_I()].
~~SLAVE presents [TAGN_O].~~
~~MASTER monitors [TAGN_I].~~
MASTER monitors [ACK_I], and prepares to latch ~~data on~~ [DAT_I()] and [TGD_I()].

CLOCK EDGE 2: MASTER latches data on [DAT_I()] and [TGD_I()].
~~MASTER latches [TAGN_I].~~
MASTER negates [STB_O] to introduce a wait state (-WSM-).

SETUP, EDGE 3: SLAVE negates [ACK_I] in response to [STB_O].

Note: any number of wait states can be inserted by the MASTER.

CLOCK EDGE 3: MASTER presents new [ADR_O()] and [~~TGA_O()~~~~TAGN_O~~].
MASTER presents new bank select [SEL_O()].
MASTER asserts [STB_O].

SETUP, EDGE 4: SLAVE decodes inputs, and responds by asserting [ACK_I].
SLAVE presents valid data on [DAT_I()] and [TGD_I()].
~~SLAVE presents [TAGN_O].~~
~~MASTER monitors [TAGN_I].~~
MASTER monitors [ACK_I], and prepares to latch ~~data on~~ [DAT_I()] and [TGD_I()].

CLOCK EDGE 4: MASTER latches data on [DAT_I()] and [TGD_I()].
MASTER presents [ADR_O()] and [~~TGA_O()~~~~TAGN_O~~].
~~MASTER latches [TAGN_I].~~
MASTER presents new bank select [SEL_O()] to indicate where it expects data.

SETUP, EDGE 5: SLAVE decodes inputs, and responds by asserting [ACK_I].
SLAVE presents valid data on [DAT_I()] and [TGD_I()].
~~SLAVE presents [TAGN_O].~~
~~MASTER monitors [TAGN_I].~~
MASTER monitors [ACK_I], and prepares to latch ~~data on~~ [DAT_I()] and [TGD_I()].

CLOCK EDGE 5: MASTER latches data on [DAT_I()] and [TGD_I()].
~~MASTER latches [TAGN_I].~~
SLAVE negates [ACK_I] to introduce a wait state.

Note: any number of wait states can be inserted by the SLAVE at this point.

SETUP, EDGE 6: SLAVE decodes inputs, and responds by asserting [ACK_I].
SLAVE presents valid data on [DAT_I()] and [TGD_I()].
MASTER monitors [ACK_I], and prepares to latch ~~data on~~ [DAT_I()] and [TGD_I()].

CLOCK EDGE 6: MASTER latches data on [DAT_I()] and [TGD_I()].
MASTER terminates cycle by negating [STB_O] and [CYC_O].




Figure 3-6. BLOCK READ cycle.

3.3.2 BLOCK WRITE Cycle

Figure 3-7 shows a BLOCK WRITE cycle. The BLOCK cycle is capable of a data transfer on every clock cycle. However, this example also shows how the MASTER and the SLAVE interfaces can both throttle the bus transfer rate by inserting wait states. A total of five transfers are shown. After the second transfer the MASTER inserts a wait state. After the fourth transfer the SLAVE inserts a wait state. The cycle is terminated after the fifth transfer. The protocol for this transfer works as follows:

CLOCK EDGE 0: MASTER presents [ADR_O()] and [~~TGA_O()~~TAGN_O].
MASTER asserts [WE_O] to indicate a WRITE cycle.
MASTER presents bank select [SEL_O()] to indicate where it expects data.
MASTER asserts [CYC_O] and [~~TGC_O()~~TAGN_O] to indicate cycle start.
MASTER asserts [STB_O].

Note: the MASTER ~~must~~ asserts [CYC_O] and/or [~~TGC_O()~~TAGN_O] at, or anytime before, clock edge 1. ~~The use of [TAGN_O] is optional.~~

SETUP, EDGE 1: SLAVE decodes inputs, and responds by asserting [ACK_I].
SLAVE prepares to latch data on [DAT_O()] and [TGD_O()].
~~SLAVE presents [TAGN_O].~~
~~MASTER monitors [TAGN_I].~~
MASTER monitors [ACK_I], and prepares to terminate current data phase.

CLOCK EDGE 1: SLAVE latches data on [DAT_O()] and [TGD_O()].
~~MASTER latches [TAGN_I].~~
MASTER presents [ADR_O()] and [~~TGA_O()~~TAGN_O].
MASTER presents new bank select [SEL_O()].

SETUP, EDGE 2: SLAVE decodes inputs, and responds by asserting [ACK_I].
SLAVE prepares to latch data on [DAT_O()] and [TGD_O()].
~~SLAVE presents [TAGN_O].~~
~~MASTER monitors [TAGN_I].~~
MASTER monitors [ACK_I], and prepares to terminate current data phase.

CLOCK EDGE 2: SLAVE latches data on [DAT_O()] and [TGD_O()].
~~MASTER latches [TAGN_I].~~
MASTER negates [STB_O] to introduce a wait state (-WSM-).

SETUP, EDGE 3: SLAVE negates [ACK_I] in response to [STB_O].

Note: any number of wait states can be inserted by the MASTER at this point.

CLOCK EDGE 3: MASTER presents [ADR_O()] and [~~TGA_O()~~TAGN_O].
MASTER presents bank select [SEL_O()] to indicate where it expects data.

MASTER asserts [STB_O].

SETUP, EDGE 4: SLAVE decodes inputs, and responds by asserting [ACK_I].
SLAVE prepares to latch data on [DAT_O()] and [TGD_O()].
~~SLAVE presents [TAGN_O].~~
~~MASTER monitors [TAGN_I].~~
MASTER monitors [ACK_I], and prepares to terminate data phase.

CLOCK EDGE 4: SLAVE latches data on [DAT_O()] and [TGD_O()].
~~MASTER latches [TAGN_I].~~
MASTER presents [ADR_O()] and [TGA_O()]~~[TAGN_O]~~.
MASTER presents new bank select [SEL_O()] to indicate where it expects data.

SETUP, EDGE 5: SLAVE decodes inputs, and responds by asserting [ACK_I].
SLAVE prepares to latch data on [DAT_O()] and [TGD_O()].
~~SLAVE presents [TAGN_O].~~
~~MASTER monitors [TAGN_I].~~
MASTER monitors [ACK_I], and prepares to terminate data phase.

CLOCK EDGE 5: SLAVE latches data on [DAT_O()] and [TGD_O()].
SLAVE negates [ACK_I] to introduce a wait state.
~~MASTER latches [TAGN_I].~~

Note: any number of wait states can be inserted by the SLAVE at this point.

SETUP, EDGE 6: SLAVE decodes inputs, and responds by asserting [ACK_I].
SLAVE prepares to latch data on [DAT_O()] and [TGD_O()].
MASTER monitors [ACK_I], and prepares to terminate data phase.

CLOCK EDGE 6: SLAVE latches data on [DAT_O()] and [TGD_O()].
MASTER terminates cycle by negating [STB_O] and [CYC_O].




Figure 3-7. BLOCK WRITE cycle.

3.4 RMW Cycle

The RMW (read-modify-write) cycle is used for indivisible semaphore operations. During the first half of the cycle a single read data transfer is performed. During the second half of the cycle a write data transfer is performed. The [CYC_O] signal remains asserted during both halves of the cycle.

RULE 3.75

All MASTER and SLAVE interfaces that support RMW cycles MUST conform to the timing requirements given in section 3.4.

PERMISSION 3.30

MASTER and SLAVE interfaces MAY be designed so that they do not support the RMW cycles.

Figure 3-8 shows a read-modify-write (RMW) cycle. The RMW cycle is capable of a data transfer on every clock cycle. However, this example also shows how the MASTER and the SLAVE interfaces can both throttle the bus transfer rate by inserting wait states. Two transfers are shown. After the first (read) transfer, the MASTER inserts a wait state. During the second transfer the SLAVE inserts a wait state. The protocol for this transfer works as follows:

CLOCK EDGE 0: MASTER presents [ADR_O()] and [~~TGA_O()~~TAGN_O].
MASTER negates [WE_O] to indicate a READ cycle.
MASTER presents bank select [SEL_O()] to indicate where it expects data.
MASTER asserts [CYC_O] and [~~TGC_O()~~TAGN_O] to indicate the start of cycle.
MASTER asserts [STB_O].

Note: the MASTER ~~must~~ asserts [CYC_O] and/or [~~TGC_O()~~TAGN_O] at, or anytime before, clock edge 1. The use of [TAGN_O] is optional.

SETUP, EDGE 1: SLAVE decodes inputs, and responds by asserting [ACK_I].
SLAVE presents valid data on [DAT_I()] and [TGD_I()].
~~SLAVE presents [TAGN_O].~~
~~MASTER monitors [TAGN_I].~~
MASTER monitors [ACK_I], and prepares to latch ~~data on~~ [DAT_I()] and [TGD_I()].

CLOCK EDGE 1: MASTER latches data on [DAT_I()] and [TGD_I()].
~~MASTER latches [TAGN_I].~~
MASTER negates [STB_O] to introduce a wait state (-WSM-).

SETUP, EDGE 2: SLAVE negates [ACK_I] in response to [STB_O].
MASTER asserts [WE_O] to indicate a WRITE cycle.

Note: any number of wait states can be inserted by the MASTER at this point.

CLOCK EDGE 2: ~~MASTER presents the same [ADR_O()] and [TAGN_O] as was on clock 1.~~
MASTER presents WRITE data on [DAT_O()] and [TGD_O()].
MASTER presents new bank select [SEL_O()].
MASTER asserts [STB_O].

SETUP, EDGE 3: SLAVE decodes inputs, and responds by asserting [ACK_I] (when ready).
SLAVE presents valid data on [DAT_I()] and [TGD_I()].
~~SLAVE presents [TAGN_O].~~
~~MASTER monitors [TAGN_I].~~
MASTER monitors [ACK_I], and prepares to latch ~~data on~~ [DAT_I()] and [TGD_I()].

Note: any number of wait states can be inserted by the SLAVE at this point.

CLOCK EDGE 3: SLAVE latches data on [DAT_O()] and [TGD_O()].
~~MASTER latches [TAGN_I].~~
MASTER negates [STB_O] and [CYC_O] indicating the end of the cycle.
SLAVE negates [ACK_I] in response to negated [STB_O].

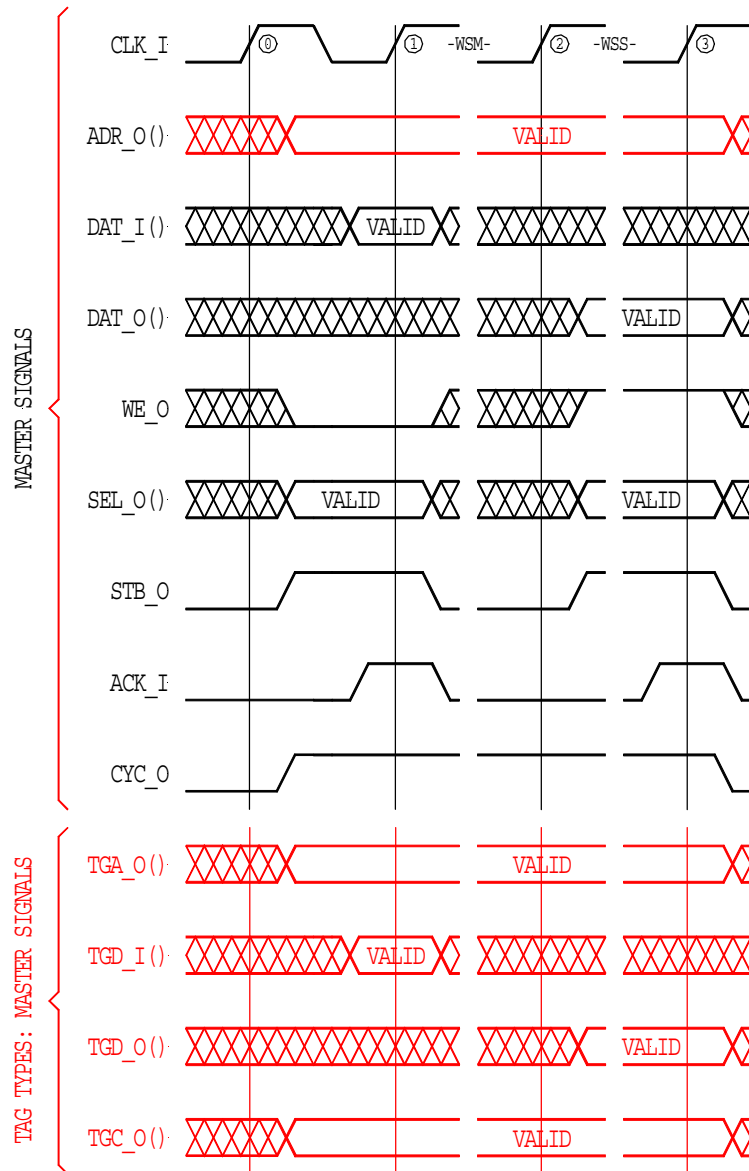


Figure 3-8. RMW cycle.

3.5 Data Organization

Data organization refers to the ordering of data during transfers. There are two general types of ordering. ~~These are -which-are~~ called BIG ENDIAN and LITTLE ENDIAN. BIG ENDIAN refers to data ordering where the most significant portion of an operand is stored at the lower address. LITTLE ENDIAN refers to data ordering where the most significant portion of an operand is stored at the higher address. The WISHBONE architecture supports both methods of data ordering.

3.5.1 Nomenclature

A BYTE(N), WORD(N), DWORD(N) and QWORD(N) nomenclature is used to define data ordering. These terms are defined in Table 3-1. Figure 3-9 shows the operand locations for input and output data ports.


Table 3-1. Data Transfer Nomenclature		
Nomenclature	Granularity	Description
BYTE(N)	8-bit	An 8-bit BYTE transfer at address 'N'.
WORD(N)	16-bit	A 16-bit WORD transfer at address 'N'.
DWORD(N)	32-bit	A 32-bit Double WORD transfer at address 'N'.
QWORD(N)	64-bit	A 64-bit Quadruple WORD transfer at address 'N'.

The table also defines the granularity of the interface. This indicates the minimum unit of data transfer that is supported by the interface. For example, the smallest operand that can be passed through a port with 16-bit granularity is a 16-bit WORD. In this case, an 8-bit operand cannot be transferred.


Figure 3-10 shows an example of how the 64-bit value of 0x0123456789ABCDEF is transferred through BYTE, WORD, DWORD and QWORD ports using BIG ENDIAN data organization. Through the 64-bit QWORD port the number is directly transferred with the most significant bit at DAT_I(63) / DAT_O(63). The least significant bit is at DAT_I(0) / DAT_O(0). However, when the same operand is transferred through a 32-bit DWORD port, it is split into two bus cycles. The two bus cycles are each 32-bits in length, with the most significant DWORD transferred at the lower address, and the least significant DWORD transferred at the upper address. A similar situation applies to the WORD and BYTE cases.

Figure 3-11 shows an example of how the 64-bit value of 0x0123456789ABC is transferred through BYTE, WORD, DWORD and QWORD ports using LITTLE ENDIAN data organization. Through the 64-bit QWORD port the number is directly transferred with the most significant bit at DAT_I(63) / DAT_O(63). The least significant bit is at DAT_I(0) / DAT_O(0). However, when the same operand is transferred through a 32-bit DWORD port, it is split into two bus cycles. The two bus cycles are each 32-bits in length, with the least significant

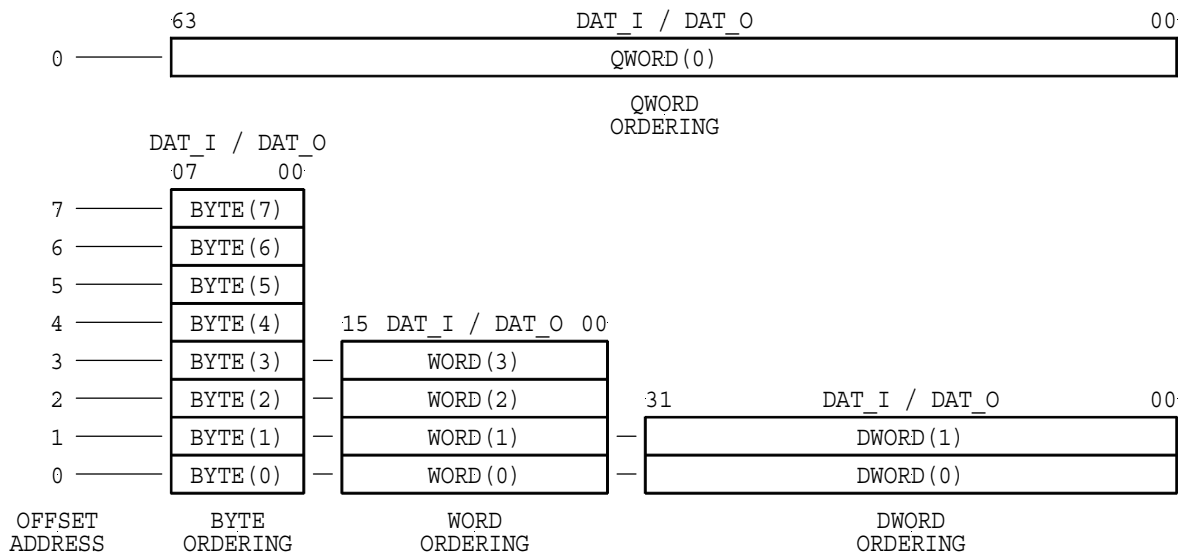
DWORD transferred at the lower address, and the most significant DWORD transferred at the upper address. A similar situation applies to the WORD and BYTE cases.



(a) BIG ENDIAN BYTE, WORD, DWORD and QWORD positioning in a 64-bit operand.



(b) LITTLE ENDIAN BYTE, WORD, DWORD and QWORD positioning in a 64-bit operand.



(c) Address nomenclature.

Figure 3-9. Operand locations for input and output data ports.

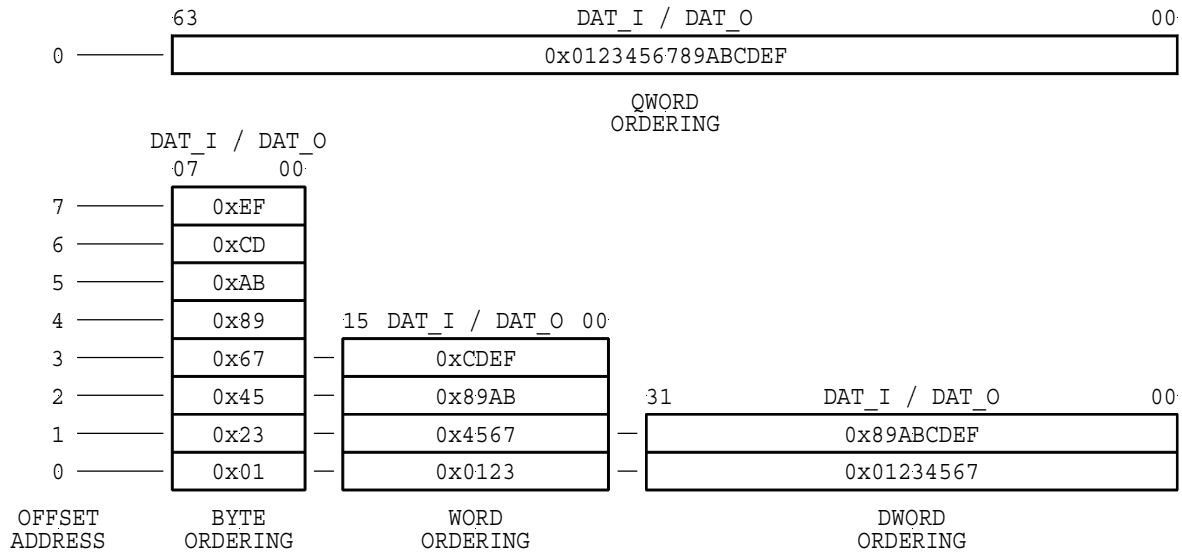


Figure 3-10. Example showing a variety of BIG ENDIAN transfers over various port sizes.

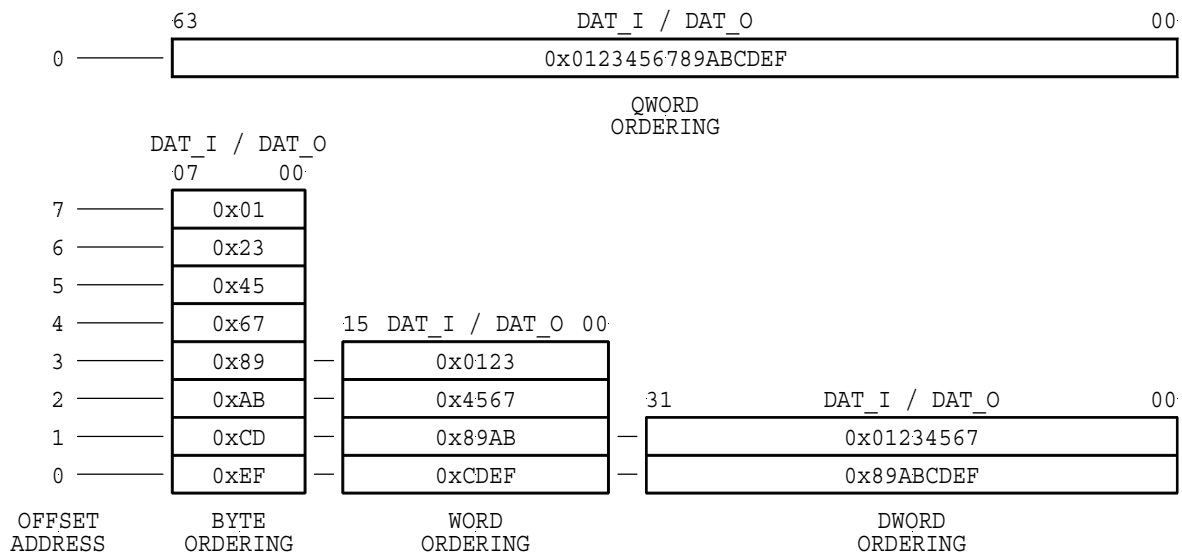


Figure 3-11. Example showing a variety of LITTLE ENDIAN transfers over various port sizes.

RULE 3.80

Data organization MUST conform to the ordering indicated in Figure 3-9.

3.5.2 Transfer Sequencing

The sequence in which data is transferred through a port is not regulated by this specification. For example, a 64-bit operand through a 32-bit port will take two bus cycles. However, the specification does not require that the lower or upper DWORD be transferred first.

RECOMMENDATION 3.10

Design interfaces so that data is transferred sequentially from lower addresses to a higher addresses.

OBSERVATION 3.30

The sequence in which an operand is transferred through a data port is not highly regulated by the specification. That is because different IP cores may produce the data in different ways. The sequence is therefore application-specific.

3.5.3 Data Organization for 64-bit Ports**RULE 3.85**

Data organization on 64-bit ports MUST conform to Figure 3-12.

64-bit Data Bus With 8-bit (BYTE) Granularity									
	Address Range:	Active Portion of Data Bus							
	ADR_I ADR_O (63..73)	DAT_I DAT_O (63..56)	DAT_I DAT_O (55..48)	DAT_I DAT_O (47..40)	DAT_I DAT_O (39..32)	DAT_I DAT_O (31..24)	DAT_I DAT_O (23..16)	DAT_I DAT_O (15..08)	DAT_I DAT_O (07..00)
	Active Select Line	SEL_I(7) SEL_O(7)	SEL_I(6) SEL_O(6)	SEL_I(5) SEL_O(5)	SEL_I(4) SEL_O(4)	SEL_I(3) SEL_O(3)	SEL_I(2) SEL_O(2)	SEL_I(1) SEL_O(1)	SEL_I(0) SEL_O(0)
BYTE Ordering	BIG ENDIAN	BYTE(0)	BYTE(1)	BYTE(2)	BYTE(3)	BYTE(4)	BYTE(5)	BYTE(6)	BYTE(7)
	LITTLE ENDIAN	BYTE(7)	BYTE(6)	BYTE(5)	BYTE(4)	BYTE(3)	BYTE(2)	BYTE(1)	BYTE(0)

64-bit Data Bus With 16-bit (WORD) Granularity					
	Address Range	Active Portion of Data Bus			
	ADR_I ADR_O (63..72)	DAT_I DAT_O (63..48)	DAT_I DAT_O (47..32)	DAT_I DAT_O (31..16)	DAT_I DAT_O (15..00)
	Active Select Line	SEL_I(3) SEL_O(3)	SEL_I(2) SEL_O(2)	SEL_I(1) SEL_O(1)	SEL_I(0) SEL_O(0)
WORD Ordering	BIG ENDIAN	WORD(0)	WORD(1)	WORD(2)	WORD(3)
	LITTLE ENDIAN	WORD(3)	WORD(2)	WORD(1)	WORD(0)

64-bit Data Bus With 32-bit (DWORD) Granularity			
	Address Range	Active Portion of Data Bus	
	ADR_I ADR_O (63..71)	DAT_I DAT_O (63..32)	DAT_I DAT_O (31..00)
	Active Select Line	SEL_I(1) SEL_O(1)	SEL_I(0) SEL_O(0)
DWORD Ordering	BIG ENDIAN	DWORD(0)	DWORD(1)
	LITTLE ENDIAN	DWORD(1)	DWORD(0)

64-bit Data Bus With 64-bit (QWORD) Granularity		
	Address Range	Active Portion of Data Bus
	ADR_I ADR_O (63..70)	DAT_I DAT_O (63..00)
	Active Select Line	SEL_I(0) SEL_O(0)
QWORD Ordering	BIG ENDIAN	QWORD(0)
	LITTLE ENDIAN	QWORD(0)

Figure 3-12. Data organization for 64-bit ports.

3.5.4 Data Organization for 32-bit Ports

RULE 3.90

Data organization on 32-bit ports MUST conform to Figure 3-13.

32-bit Data Bus With 8-bit (BYTE) Granularity					
	Address Range	Active Portion of Data Bus			
	ADR_I ADR_O (63..02)	DAT_I DAT_O (31..24)	DAT_I DAT_O (23..16)	DAT_I DAT_O (15..08)	DAT_I DAT_O (07..00)
	Active Select Line	SEL_I(3) SEL_O(3)	SEL_I(2) SEL_O(2)	SEL_I(1) SEL_O(1)	SEL_I(0) SEL_O(0)
BYTE Ordering	BIG ENDIAN	BYTE(0) BYTE(4)	BYTE(1) BYTE(5)	BYTE(2) BYTE(6)	BYTE(3) BYTE(7)
	LITTLE ENDIAN	BYTE(3) BYTE(7)	BYTE(2) BYTE(6)	BYTE(1) BYTE(5)	BYTE(0) BYTE(4)

32-bit Data Bus With 16-bit (WORD) Granularity			
	Address Range	Active Portion of Data Bus	
	ADR_I ADR_O (63..01)	DAT_I DAT_O (31..16)	DAT_I DAT_O (15..00)
	Active Select Line	SEL_I(1) SEL_O(1)	SEL_I(0) SEL_O(0)
WORD Ordering	BIG ENDIAN	WORD(0) WORD(2)	WORD(1) WORD(3)
	LITTLE ENDIAN	WORD(1) WORD(3)	WORD(0) WORD(2)

32-bit Data Bus With 32-bit (DWORD) Granularity		
	Address Range	Active Portion of Data Bus
	ADR_I ADR_O (63..00)	DAT_I DAT_O (31..00)
	Active Select Line	SEL_I(0) SEL_O(0)
DWORD Ordering	BIG ENDIAN	DWORD(0) DWORD(1)
	LITTLE ENDIAN	DWORD(0) DWORD(1)

Figure 3-13. Data organization for 32-bit ports.

3.5.5 Data Organization for 16-bit Ports

RULE 3.95

Data organization on 16-bit ports MUST conform to Figure 3-14.

16-bit Data Bus With 8-bit (BYTE) Granularity			
	Address Range	Active Portion of Data Bus	
		DAT_I DAT_O (15..08)	DAT_I DAT_O (07..00)
	Active Select Line	SEL_I(1) SEL_O(1)	SEL_I(0) SEL_O(0)
BYTE Ordering	BIG ENDIAN	BYTE(0) BYTE(2) BYTE(4) BYTE(6)	BYTE(1) BYTE(3) BYTE(5) BYTE(7)
	LITTLE ENDIAN	BYTE(1) BYTE(3) BYTE(5) BYTE(7)	BYTE(0) BYTE(2) BYTE(4) BYTE(6)

16-bit Data Bus With 16-bit (WORD) Granularity		
	Address Range	Active Portion of Data Bus
		DAT_I DAT_O (15..00)
	Active Select Line	SEL_I(0) SEL_O(0)
WORD Ordering	BIG ENDIAN	WORD(0) WORD(1) WORD(2) WORD(3)
	LITTLE ENDIAN	WORD(0) WORD(1) WORD(2) WORD(3)

Figure 3-14. Data organization for 16-bit ports.

3.5.6 Data Organization for 8-bit Ports

RULE 3.100

Data organization on 8-bit ports MUST conform to Figure 3-15.

8-bit Data Bus With 8-bit (BYTE) Granularity		
	Address Range	Active Portion of Data Bus
	ADR_I ADR_O (63..00)	DAT_I DAT_O (07..00)
	Active Select Line	SEL_I (0) SEL_O (0)
BYTE Ordering	BIG ENDIAN	BYTE (0) BYTE (1) BYTE (2) BYTE (3) BYTE (4) BYTE (5) BYTE (6) BYTE (7)
	LITTLE ENDIAN	BYTE (0) BYTE (1) BYTE (2) BYTE (3) BYTE (4) BYTE (5) BYTE (6) BYTE (7)

Figure 3-15. Data organization for 8-bit ports.

3.6 References

Cohen, Danny. *On Holy Wars and a Plea for Peace*. IEEE Computer Magazine, October 1981. Pages 49-54. [Description of BIG ENDIAN and LITTLE ENDIAN.]

Chapter 4 – Timing Specification

The WISHBONE specification is designed to provide the end user with very simple timing constraints. Although the application specific circuit(s) will vary in this regard, the interface itself is designed to work without the need for detailed timing specifications. In all cases, the only timing information that is needed by the end user is the maximum clock frequency (for [CLK_I]) that is passed to a place & route tool. The maximum clock frequency is dictated by the time delay between a positive clock edge on [CLK_I] to the setup on a stage further down the logical signal path. This delay is shown graphically in Figure 4-1, and is defined as $T_{pd,clk-su}$.

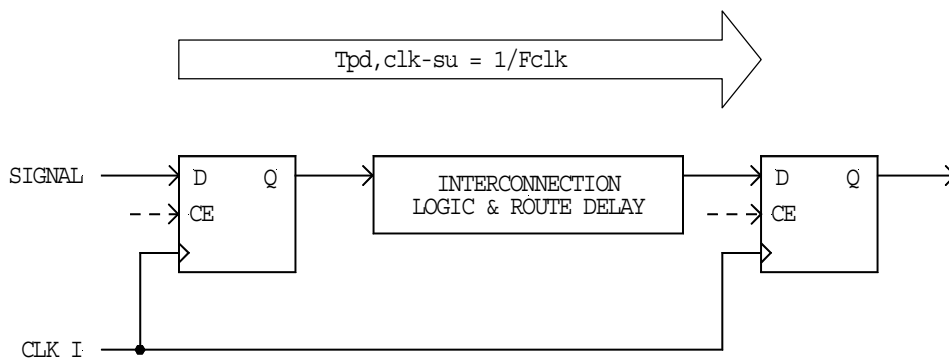


Figure 4-1. Definition for $T_{pd,clk-su}$.

RULE 4.00

The clock input [CLK_I] to each IP core MUST coordinate all activities for the internal logic within the WISHBONE interface. All WISHBONE output signals are registered at the rising edge of [CLK_I]. All WISHBONE input signals **MUST** be stable before the rising edge of [CLK_I].

PERMISSION 4.00

The user's place and route tool MAY be used to enforce **RULE 4.00. this rule.**

OBSERVATION 4.00

Most place and route tools can be easily configured to enforce **RULE 4.00 this rule.** Generally, it only requires a single timing specification for $T_{pd,clk-su}$.

RULE 4.05

The WISHBONE interface MUST use synchronous, RTL design methodologies that, given nearly infinitely fast gate delays, will operate over a nearly infinite range of clock frequencies on [CLK_I].

OBSERVATION 4.05

Realistically, the WISHBONE interface will never be expected to operate over a nearly infinite frequency range. However this requirement eliminates the need for non-portable timing constraints (that may work only on certain target devices).

OBSERVATION 4.10

The WISHBONE interface logic assumes that a low-skew clock distribution scheme is used on the target device, and that the clock-skew shall be low enough to permit reliable operation over the environmental conditions.

PERMISSION 4.05

The IP core connected to a WISHBONE interface MAY include application specific timing requirements.

RULE 4.10

The clock input [CLK_I] MUST have a duty cycle that is no less than 40%, and no greater than 60%.

PERMISSION 4.10

The SYSCON ~~moduleinterface~~ MAY use a variable clock generator. In these cases the clock frequency can be changed by the SYSCON ~~moduleinterface~~ so long as the clock edges remain clean and monotonic, and if the clock does not violate the duty cycle requirements.

PERMISSION 4.15

The SYSCON ~~moduleinterface~~ MAY use a gated clock generator. In these cases the clock shall be stopped in the low logic state. When the gated clock is stopped and started the clock edges ~~are required to~~**must** remain clean and monotonic.

SUGGESTION 4.00

When using a gated clock generator, turn the clock off when the WISHBONE interconnection is not busy. One way of doing this is to create a MASTER interface whose sole purpose is to acquire the WISHBONE interconnection and turn the clock off. This assures that the WISHBONE interconnection is not busy when gating the clock off. When the clock signal is restored the MASTER then releases the WISHBONE interconnection.

OBSERVATION 4.15

This specification does not attempt to govern the design of gated or variable clock generators.

SUGGESTION 4.10

Design an IP core so that all of the circuits (including the WISHBONE interconnect) follow the aforementioned RULEs, as this will make the core portable across a wide range of target devices and technologies.

References

~~Orton et al. CHANGING CLOCK FREQUENCY US Patent No. 6,118,306~~

Chapter 5 – Cited Patent References

This chapter contains a partial list of the patents that have been reviewed by the WISHBONE steward and others. In the opinion of the steward, the WISHBONE System-on-Chip (SoC) Interconnection Architecture for Portable IP Cores does not infringe on any of these patents. However, the possibility exists that an integrated circuit device designed to the WISHBONE specification could infringe on the intellectual property rights of others. The user assumes all responsibility for determining if their WISHBONE design infringes on the rights of others. All of these documents contain information relevant to SoC design and integration. Noteworthy patents are marked with a '(*)'.

This list is maintained for three reasons. First, a public domain specification can't depend on patented ideas (unless permission to use the patent is obtained). This is the list of documents that have been reviewed to see if WISHBONE infringes on any known patents. Second, it provides a starting point for IC designers and other researchers who need to know if a specific SoC design infringes on the rights of others. Third, the patent database is a wonderful place to learn how other people have solved similar SoC problems. Patent documents are very good in this regard, as they must fully describe how to reproduce an invention.

5.1 General Methods Relating to SoC

Hartmann, Alfred C. - US Patent No. 6,096,091

DYNAMICALLY RECONFIGURABLE LOGIC NETWORKS INTERCONNECTED BY FALL-THROUGH FIFOS FOR FLEXIBLE PIPELINE PROCESSING IN A SYSTEM-ON-A-CHIP

Luk et al. - US Patent No. 5,790,839

SYSTEM INTEGRATION OF DRAM MACROS AND LOGIC CORES IN A SINGLE CHIP ARCHITECTURE

Luk et al. - US Patent No. 5,883,814

SYSTEM-ON-CHIP LAYOUT COMPILATION

Wingard et al. - US Patent No. 5,948,089

FULLY-PIPELINED FIXED-LATENCY COMMUNICATIONS SYSTEM WITH A REAL TIME DYNAMIC BANDWIDTH ALLOCATION.

Wingard et al. - US Patent No. 6,182,183 B1

COMMUNICATION SYSTEM AND METHOD WITH MULTILEVEL CONNECTION IDENTIFICATION

5.2 Methods Relating to SoC Testability

Edwards, et al.- US Patent No. 6,298,394 B1 (*)

SYSTEM AND METHOD FOR CAPTURING INFORMATION ON AN INTERCONNECT IN AN INTEGRATED CIRCUIT

Flynn, David W. - US Patent No. 5,525,971
INTEGRATED CIRCUIT

5.4 Methods Relating to Variable Clock Frequency

Gandhi et al. - US Patent No. 6,185,691 B1
CLOCK GENERATION

Kardach et al. - US Patent No. 5,473,767
METHOD AND APPARATUS FOR ASYNCHRONOUSLY STOPPING THE CLOCK ON A PROCESSOR.

Kardach et al. - US Patent No. 5,918,043
METHOD AND APPARATUS FOR ASYNCHRONOUSLY STOPPING THE CLOCK ON A PROCESSOR.

Maitra, Amit K. - US Patent No. 5,623,647
APPLICATION SPECIFIC CLOCK THROTTLING

Orton et al. - US Patent No. 6,118,306 (*)
CHANGING CLOCK FREQUENCY

Poplinger et al. - US Patent No. 6,173,379 B1
MEMORY DEVICE FOR A MICROPROCESSOR REGISTER FILE HAVING A POWER MANAGEMENT SCHEME AND METHOD FOR COPYING INFORMATION BETWEEN MEMORY SUB-CELLS IN A SINGLE CLOCK CYCLE

Stinson et al. - US Patent No. 6,127,858
METHOD AND APPARATUS FOR VARYING A CLOCK FREQUENCY ON A PHASE BY PHASE BASIS

Thomas, Thomas P. - US Patent No. 6,140,883
TUNABLE, ENERGY EFFICIENT CLOCKING SCHEME

Wong et al. - US Patent No. 5,586,307
METHOD AND APPARATUS SUPPLYING SYNCHRONOUS CLOCK SIGNALS TO CIRCUIT COMPONENTS

Young, Bruce - US Patent No. 6,079,022
METHOD AND APPARATUS FOR DYNAMICALLY ADJUSTING THE CLOCK SPEED OF A BUS DEPENDING ON BUS ACTIVITY

5.5 Methods Relating to Selection of IP Cores

Lee et al. - US Patent No. 6,102,961

METHOD AND APPARATUS FOR SELECTING IP BLOCKS

Methods Relating to Data Flow Architectures

Cismas, Sorin C. - US Patent No. 6,145,073

DATA FLOW INTEGRATED CIRCUIT ARCHITECTURE

5.6 Methods Relating to Crossbar Switch Architectures

Brewer et al. - US Patent No. 5,577,204

PARALLEL PROCESSING COMPUTER SYSTEM INTERCONNECTIONS UTILIZING UNIDIRECTIONAL COMMUNICATION LINKS WITH SEPARATE REQUEST AND RESPONSE LINES FOR DIRECT COMMUNICATION OR USING A CROSSBAR SWITCHING DEVICE

Nelson et al. - US Patent No. 6,138,185

HIGH PERFORMANCE CROSSBAR SWITCH

Van Krevelen et al. - US Patent No. 6,230,229 B1

METHOD AND SYSTEM FOR ARBITRATING PATH CONTENTION IN A CROSSBAR INTERCONNECT NETWORK

Appendix A – WISHBONE Tutorial⁴

By: Wade D. Peterson, Silicore Corporation

The WISHBONE System-on-Chip (SoC) interconnection is a method for connecting digital circuits together to form an integrated circuit ‘chip’. This tutorial provides an introduction to the WISHBONE design philosophy and its practical applications.

The WISHBONE architecture solves a very basic problem in integrated circuit design. That is, how to connect circuit functions together in a way that is simple, flexible and portable. The circuit functions are generally provided as ‘IP Cores’ (Intellectual Property Cores), which system integrators can purchase or make themselves.

Under this topology, IP Cores are the functional building blocks in the system. They are available in a wide variety of functions such as microprocessors, serial ports, disk interfaces, network controllers and so forth. Generally, the IP cores are developed independently from each other and are tied together and tested by a third party system integrator. WISHBONE aides the system integrator by standardizing the IP Core interfaces. This makes it much easier to connect the cores, and therefore much easier to create a custom System-on-Chip.

A.1 An Introduction to WISHBONE

WISHBONE uses a MASTER/SLAVE architecture. That means that functional modules with MASTER interfaces initiate data transactions to participating SLAVE interfaces. As shown in Figure A-1, the MASTERS and SLAVES communicate through an interconnection interface called the INTERCON. The INTERCON is best thought of as a ‘cloud’ that contains circuits. These circuits allow MASTERS to communicate with SLAVES.

The term ‘cloud’ is borrowed from the telecom community. Oftentimes, telephone systems are modeled as clouds that represent a system of telephone switches and transmission devices. Telephone handsets are connected to the cloud, and are used to make phone calls. The cloud itself represents a network that carries a telephone call from one location to another. The activity going on inside the cloud depends upon where the call is made and where it is going. For example, if a call is made to another office down the hall, then the cloud may represent a local telephone switch located in the same building. However, if the call is made across an ocean, then the cloud may represent a combination of copper, fiber optic and satellite transmission systems.

⁴ This tutorial is not part of the WISHBONE specification.

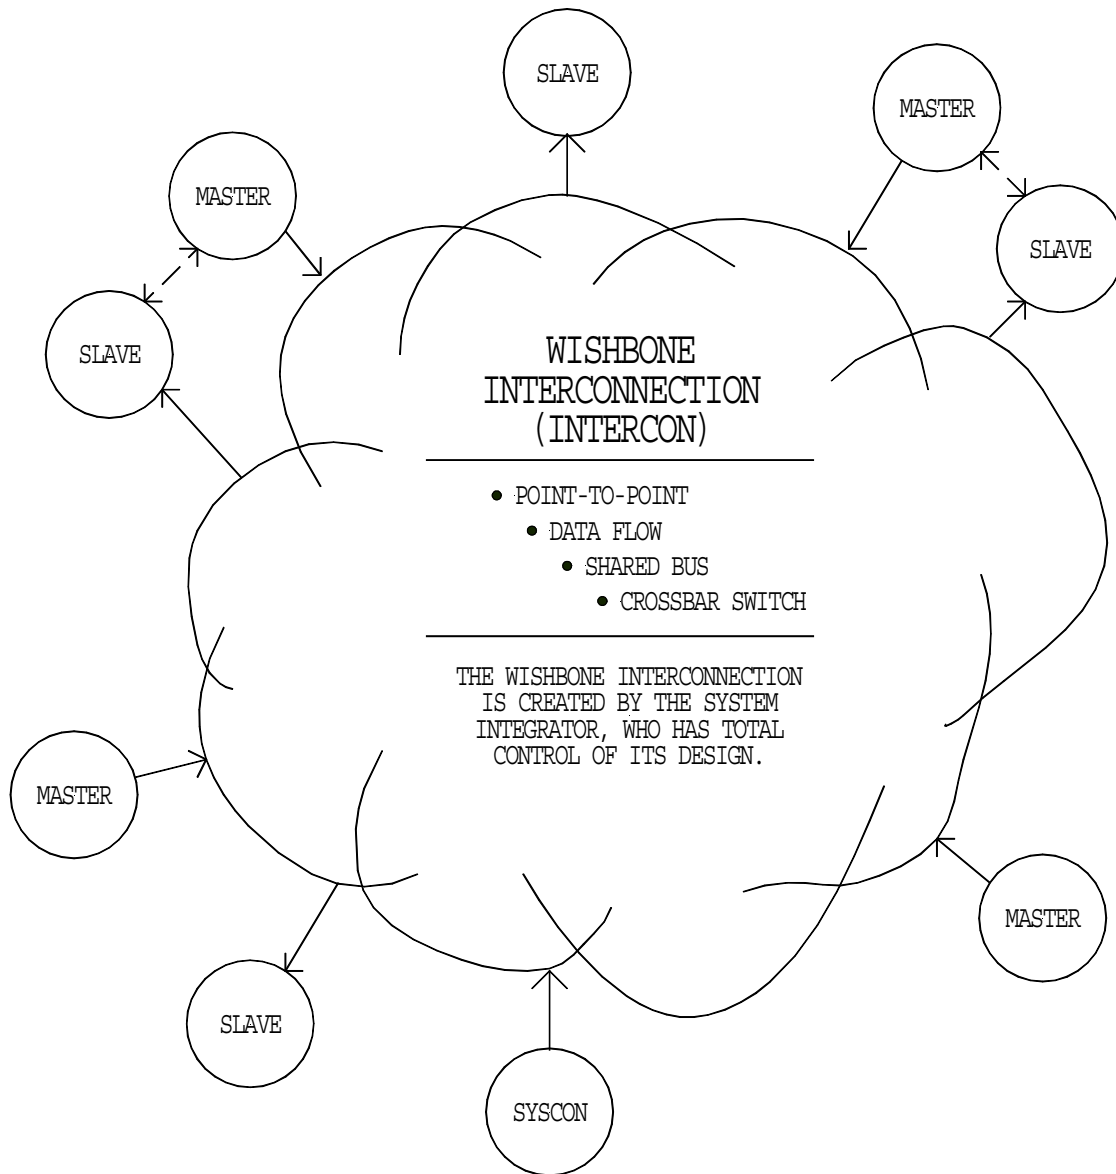


Figure A-1. The WISHBONE interconnection.

The cloud analogy is used because WISHBONE can be modeled in a similar way. MASTER and SLAVE interfaces (which are analogous to the telephones) communicate through an interconnection (which is analogous to the telephone network ‘cloud’). The WISHBONE interconnection network can be changed by the system integrator to suit his or her own needs. In WISHBONE terminology this is called a *variable interconnection*.

Variable interconnection allows the system integrator to change the way in which the MASTER and SLAVE interfaces communicate with each other. For example, a pair of MASTER and SLAVE interfaces can communicate with point-to-point, data flow, shared bus or crossbar switch topologies.

The variable interconnection scheme is very different from that used in traditional microcomputer buses such as PCI, cPCI, VMEbus and ISA bus. Those systems use printed circuit backplanes with hardwired connectors. The interfaces on those buses can’t be changed, which severely limits how microcomputer boards communicate with each other. WISHBONE overcomes this limitation by allowing the system integrator to change the system interconnection.

This is possible because integrated circuit chips have interconnection paths that can be adjusted. These are very flexible, and take the form of logic gates and routing paths. These can be ‘programmed’ into the chip using a variety of tools. For example, if the interconnection is described with a hardware description language like VHDL or Verilog®, then the system integrator has the ability to define and adjust the interconnection. Interconnection libraries can also be formed and shared.

The WISHBONE interconnection itself is nothing more than a large, synchronous circuit. It must be designed to *logically* operate over a nearly infinite frequency range. However, every integrated circuit has physical characteristics that limit the maximum frequency of the circuit. In WISHBONE terminology this is called a *variable timing specification*. This means that a WISHBONE compatible circuit will theoretically function normally at any speed, but that its maximum speed will always be limited by the process technology of the integrated circuit.

At Silicore Corporation we generally define our WISHBONE interconnections using the VHDL hardware description language. These take the form of an ASCII file that contains the VHDL language instructions. This allows us to fully define our interconnections in a way that best fits the application. However, it also allows us to share our interconnections with others, who can adjust them to meet their own needs. It’s important to note, though, that there’s nothing magic about the interconnection itself. It’s just a file, written with off-the-shelf tools, that fully describes the hardware in the interconnection.

A.2 Types of WISHBONE Interconnection

The WISHBONE variable interconnection allows the system integrator to change the way that IP cores connect to each other. There are four defined types of WISHBONE interconnection, and include:

- Point-to-point
- Data flow
- Shared bus
- Crossbar switch

A fifth possible type is the off-chip interconnection. However, off-chip implementations generally fit one of the other four basic types. For example, WISHBONE interfaces on two different integrated circuits can be connected using a point-to-point interconnection.

The WISHBONE specification does not dictate how any of these are implemented by the system integrator. That's because the interconnection itself is a type of IP Core interface called the INTERCON. The system integrator can use or modify an off-the-shelf INTERCON, or create their own.

A.2.1 Point-to-point Interconnection

The point-to-point interconnection is the simplest way to connect two WISHBONE IP cores together. As shown in Figure A-2, the point-to-point interconnection allows a single MASTER interface to connect to a single SLAVE interface. For example, the MASTER interface could be on a microprocessor IP core, and the SLAVE interface could be on a serial I/O port.




Figure A-2. The point-to-point interconnection.

A.2.2 Data Flow Interconnection

The data flow interconnection is used when data is processed in a sequential manner. As shown in Figure A-3, each IP core in the data flow architecture has both a MASTER and a SLAVE interface. Data flows from core-to-core. Sometimes this process is called *pipelining*.




Figure A-3. The data flow interconnection.

The data flow architecture exploits parallelism, thereby speeding up execution time. For example, if each of the IP cores in the Figure represents a floating point processor, then the system has three times the number crunching potential of a single unit. This assumes, of course, that each IP Core takes an equal amount of time to solve its problem, and that the problem can be solved in a sequential manner. In actual practice this may or may not be true, but it does illustrate how the data flow architecture can provide a high degree of parallelism when solving problems.

A.2.3 Shared Bus Interconnection

The shared bus interconnection is useful for connecting two or more MASTERS with one or more SLAVES. A block diagram is shown in Figure A-4. In this topology a MASTER initiates a bus cycle to a target SLAVE. The target SLAVE then participates in one or more bus cycles with the MASTER.

An arbiter (not shown in the Figure) determines when a MASTER may gain access to the shared bus. The arbiter acts like a ‘traffic cop’ to determine when and how each MASTER accesses the shared resource. Also, the type of arbiter is completely defined by the system integrator. For example, the shared bus can use a priority or a round robin arbiter. These grant the shared bus on a priority or equal basis, respectively.

The main advantage to this technique is that shared interconnection systems are relatively compact. Generally, it requires fewer logic gates and routing resources than other configurations, especially the crossbar switch. Its main disadvantage is that MASTERS may have to wait before gaining access to the interconnection. This degrades the overall speed at which a MASTER may transfer data.




Figure A-4. Shared bus interconnection.

The WISHBONE specification does not dictate how the shared bus is implemented. Later on, we'll see that it can be made either with multiplexer or three-state buses. This gives the system integrator additional flexibility, as some logic chips work better with multiplexor logic, and some work better with three-state buses.

The shared bus is typically found in standard buses like PCI and VMEbus. There, a MASTER interface arbitrates for the common shared bus, and then communicates with a SLAVE. In both cases this is done with three-state buses.

A.2.4 Crossbar Switch Interconnection

The crossbar switch interconnection is used when connecting two or more WISHBONE MASTERS together so that each can access two or more SLAVES. A block diagram is shown in Figure A-5. In the crossbar interconnection, a MASTER initiates an addressable bus cycle to a target SLAVE. An arbiter (not shown in the diagram) determines when each MASTER may gain access to the indicated SLAVE. Unlike the shared bus interconnection, the crossbar switch allows more than one MASTER to use the interconnection (as long as two MASTERS don't access the same SLAVE at the same time).




Figure A-5. Crossbar switch interconnection.

Under this method, each master arbitrates for a ‘channel’ on the switch. Once this is established, data is transferred between the MASTER and the SLAVE over a private communication link. The Figure shows two possible channels that may appear on the switch. The first connects MASTER ‘MA’ to SLAVE ‘SB’. The second connects MASTER ‘MB’ to SLAVE ‘SA’.

The overall data transfer rate of the crossbar switch is higher than shared bus mechanisms. For example, the figure shows two MASTER/SLAVE pairs interconnected at the same time. If each communication channel supports a data rate of 100 Mbyte/sec, then the two data pairs would operate in parallel at 200 Mbyte/sec. This scheme can be expanded to support extremely high data transfer rates.

One disadvantage of the crossbar switch is that it requires more interconnection logic and routing resources than shared bus systems. As a rule-of-thumb, a crossbar switch with two MASTERS and two SLAVES takes twice as much interconnection logic as a similar shared bus system (with two MASTERS and two SLAVES).

The crossbar interconnection is a typical configuration that one might find in microcomputer buses like⁵ RACEway, SKY Channel or Myrinet.


⁵ Raceway: ANSI/VITA 5-1994. SKYchannel: ANSI/VITA 10-1995. Myrinet: ANSI/VITA 26-1998. For more information about these standards see www.vita.com.

A.3 The WISHBONE Interface Signals


WISHBONE MASTER and SLAVE interfaces can be connected together in a number of ways. This requires that WISHBONE interface signals and bus cycles be designed in a very flexible and reusable manner. The signals were defined with the following requirements:

- The signals allow MASTER and SLAVE interfaces to support point-to-point, data flow, shared bus and crossbar switch interconnections.
- The signals allow three basic types of bus cycle. These include SINGLE READ/WRITE, BLOCK READ/WRITE and RMW (read-modify-write) bus cycles. The operation of these bus cycles are described below.
- A handshaking mechanism is used so that either the MASTER or the participating SLAVE interface can adjust the data transfer rate during a bus cycle. This allows the speed of each bus cycle (or phase) to be adjusted by either the MASTER or the SLAVE interface. This means that all WISHBONE bus cycles run at the speed of the slowest interface.
- The handshaking mechanism allows a participating SLAVE to accept a data transfer, reject a data transfer with an error or ask the MASTER to retry a bus cycle. The SLAVE does this by generating the [ACK_O], [ERR_O] or [RTY_O] signals respectively. Every interface must support the [ACK_O] signal, but the error and retry acknowledgement signals are optional.
- All signals on MASTER and SLAVE interfaces are either inputs or outputs, but are never bi-directional (i.e. three-state). This is because some FPGA and ASIC devices do not support bi-directional signals. However, it is permissible (and sometimes advantageous) to use bi-directional signals in the interconnection logic if the target device supports it.
- Address and data bus widths can be altered to fit the application. 8, 16, 32 and 64-bit data buses, and 0-64 bit address buses are defined.
- As shown in Figure A-6, all signals are arranged so that MASTER ~~and~~ SLAVE ~~and SYSCON~~ interfaces can be connected directly together to form a simple point-to-point interface. This allows very compact and efficient WISHBONE interfaces to be built. For example, WISHBONE could be used as the external system bus on a microprocessor IP Core. However, it's efficient enough so that it can be used for internal buses inside of the microprocessor.
- User defined signals in the form of 'tags' are allowed. This allows the system integrator to add special purpose signals to each WISHBONE interface. For example, the system integrator could add a parity bit to the address or data buses.

A comprehensive list of the WISHBONE signals and their descriptions is given in the specification.



(A) FORMING A POINT-TO-POINT INTERCONNECTION.



(B) POINT-TO-POINT INTERCONNECTION.

Figure A-6. The WISHBONE signals are selected to permit a MASTER, and SLAVE and SYSCON interfaces to be directly connected, thereby forming a simple point-to-point interface.

A.4 The WISHBONE Bus Cycles

There are three types of defined WISHBONE bus cycles. They include:

- SINGLE READ/WRITE
- BLOCK READ/WRITE
- READ MODIFY WRITE (RMW)

A.4.1 SINGLE READ/WRITE Cycle

The SINGLE READ/WRITE is the most basic WISHBONE bus cycle. As the name implies, it is used to transfer a single data operand. Figure A-7 shows a typical SINGLE READ cycle.

The WISHBONE specification shows all bus cycle timing diagrams as if the MASTER and SLAVE interfaces were connected in a point-to-point configuration. They also show all of the signals on the MASTER side of the interface. This provides a standard way of describing the interface without having to describe the whole system. For example, the Figure shows a signal called [ACK_I], which is an input to a MASTER interface. In this configuration it is directly connected to [ACK_O], which is driven by the SLAVE. If the timing diagram were shown from the perspective of the SLAVE, then the [ACK_O] signal would have been shown. The SINGLE READ cycle operates thusly:

1. In response to clock edge 0, the MASTER interface asserts [ADR_O()], [WE_O], [SEL_O], [STB_O] and [CYC_O].
2. The SLAVE decodes the bus cycle by monitoring its [STB_I] and address inputs, and presents valid data on its [DAT_O()] lines. Because the system is in a point-to-point configuration, the SLAVE [DAT_O()] signals are connected to the MASTER [DAT_I()] signals.
3. The SLAVE indicates that it has placed valid data on the data bus by asserting the MASTER's [ACK_I] acknowledge signal. Also note that the SLAVE may delay its response by inserting one or more wait states. In this case, the SLAVE does not assert the acknowledge line. The possibility of a wait state in the timing diagrams is indicated by '-WSS-'.
-WSS-
4. The MASTER monitors the state of its [ACK_I] line, and determines that the SLAVE has acknowledged the transfer at clock edge 1.
5. The MASTER latches [DAT_I()] and negates its [STB_O] signal in response to [ACK_I].

The SINGLE WRITE cycle operates in a similar manner, except that the MASTER asserts [WE_O] and places data on [DAT_O]. In this case the SLAVE asserts [ACK_O] when it has latched the data.




Figure A-7. SINGLE READ cycle.

A.4.2 BLOCK READ/WRITE Cycle

The BLOCK READ/WRITE cycles are very similar to the SINGLE READ/WRITE cycles. The BLOCK cycles can be thought of as two or more back-to-back SINGLE cycles strung together. In WISHBONE terminology the term *cycle* refers to the whole BLOCK cycle. The small, individual data transfers that make up the BLOCK cycle are called *phases*.

The starting and stopping point of the BLOCK cycles are identified by the assertion and negation of the MASTER [CYC_O] signal (respectively). The [CYC_O] signal is also used in shared bus and crossbar interconnections because it informs system logic that the MASTER wishes to use the bus. It also informs the interconnection when it is through with its bus cycle.

A.4.3 READ-MODIFY-WRITE (RMW) Cycle

The READ-MODIFY-WRITE cycle is used in multiprocessor and multitasking systems. This special cycle allows multiple software processes to share common resources by using semaphores. This is commonly done on interfaces for disk controllers, serial ports and memory. As the name implies, the READ-MODIFY-WRITE cycle reads and writes data to a memory location in a single bus cycle. It prevents the allocation of a common resource to two or more processes. The READ-MODIFY-WRITE cycle can also be thought of as an *indivisible cycle*.

The read portion of the cycle is called the *read phase*, and the write portion is called *the write phase*. When looking at the timing diagram of this bus cycle, it can be thought of as a two phase BLOCK cycle where the first phase is a read and the second phase is a write.

The READ-MODIFY-WRITE cycle is also known as an indivisible cycle because it is designed for multiprocessor systems. WISHBONE shared bus interconnections must be designed so that once an arbiter grants the bus to a first MASTER, it will not grant the bus to a second MASTER until the first MASTER gives up the bus. This allows a single MASTER (such as a microprocessor) to read some data, modify it and then write it back...all in a single, contiguous bus cycle. If the arbiter were allowed to change MASTERS in the middle of the cycle, then the two processors could incorrectly interpret the semaphore. The arbiter does this by monitoring the [CYC_O] cycle from each MASTER on the interconnection. The problem is averted because the [CYC_O] signal is always asserted for the duration of the RMW cycle.

To illustrate this point, consider a two processor system with a single disk controller. In this case each processor has a MASTER interface, and the disk controller has a SLAVE interface. Oftentimes, these systems require that only one processor access the disk at any given time⁶. To satisfy this requirement, a semaphore bit somewhere in memory is assigned to act as a ‘traffic cop’ between the two processors. If the bit is cleared, then the disk is available for use. If it’s set, then the disk controller is busy.

⁶ This is a common requirement to prevent one form of disk ‘thrashing’. In this case, if both processors were allowed to access the disk during the same time interval, then one processor could request data from one sector of the disk while the other requested data from another sector. This could cause a situation where the disk head is constantly moved between the two locations, thereby degrading its performance or causing it to fail altogether.

Now consider a system where the two processors both need to use the disk. We'll call them processor #0 and processor #1. In order for processor #0 to acquire the disk it first reads and stores the state of the semaphore bit, and then sets the bit by writing back to memory. The reading and setting of the bit takes place inside of a single RMW cycle.

Once the processor is done with the semaphore operation, it checks the state of the bit it read during the first phase of the RMW cycle. If the bit is clear it goes ahead and uses the disk controller. If the other processor attempts to use the disk controller at this time, it reads a '1' from the semaphore, thereby preventing it from accessing the disk controller. When the first processor (#0) is done with the disk controller, it simply clears the semaphore bit by writing a '0' to it. This allows the other processor to gain access to the controller the next time it checks the semaphore.

Now consider the same situation, except where the semaphore is set and cleared using a SINGLE READ cycle followed by a SINGLE WRITE cycle. In this case it is possible for both processors to gain access to the disk controller at the same time...a situation that would crash the system. That's because the arbiter can grant the bus in the following order:

- Processor #0 reads '0' from the semaphore bit.
- Processor #1 reads '0' from the semaphore bit.
- Processor #0 writes '1' to the semaphore bit.
- Processor #1 writes '1' to the semaphore bit.

This leads to a system crash because both processors read a '0' from the semaphore bit, thereby causing both to access the disk controller.

It is important to note that a processor (or other device connected to the MASTER interface) must support the RMW cycle in order for this to be effective. This is generally done with special instructions that force a RMW bus cycle. Not all processors do this. A good example is the 680XX family of microprocessors. These use special compare-and-set (CAS) and test-and-set (TAS) instructions to generate RMW cycles, and to do the semaphore operations.

A.5 Endian

The WISHBONE specification regulates the ordering of data. This is because data can be presented in two different ways. In the first way, the most significant byte of an operand is placed at the higher (bigger) address. In the second way, the most significant byte of an operand can be placed at the lower (smaller) address. These are called BIG ENDIAN and LITTLE ENDIAN data operands, respectively. WISHBONE supports both types.

ENDIAN becomes an issue when the granularity of a WISHBONE port is smaller than the operand size. For example, a 32-bit port can have an 8-bit (BYTE wide) granularity. This results in a fairly ambiguous situation where the most significant byte of the 32-bit operand could be

placed at the higher or lower byte address of the port. However, ENDIAN is not an issue when the granularity and port size are the same.

The system integrator may wish to connect a BIG ENDIAN interface to a LITTLE ENDIAN interface. In many cases the conversion is quite straightforward, and does not require any exotic conversion logic. Furthermore, the conversion does not create any speed degradation in the interface. In general, the ENDIAN conversion takes place by renaming the data and select I/O signals at a MASTER or SLAVE interface.

Figure A-8 shows a simple example where a 32-bit BIG ENDIAN MASTER output (CORE 'A') is connected to a 32-bit LITTLE ENDIAN SLAVE input (CORE 'B'). Both interfaces have 32-bit operand sizes and 8-bit granularities. As can be seen in the diagram, the ENDIAN conversion is accomplished by cross coupling the data and select signal arrays. This is quite simple since the conversion is accomplished at the interconnection level, or using a wrapper. This is especially simple in soft IP cores using VHDL or Verilog® hardware description languages, as it only requires the renaming of signals.

In some cases the address lines may also need to be modified between the two cores. For example, if 64-bit operands are transferred between two cores with 8-bit port sizes, then the address lines may need to be modified as well.

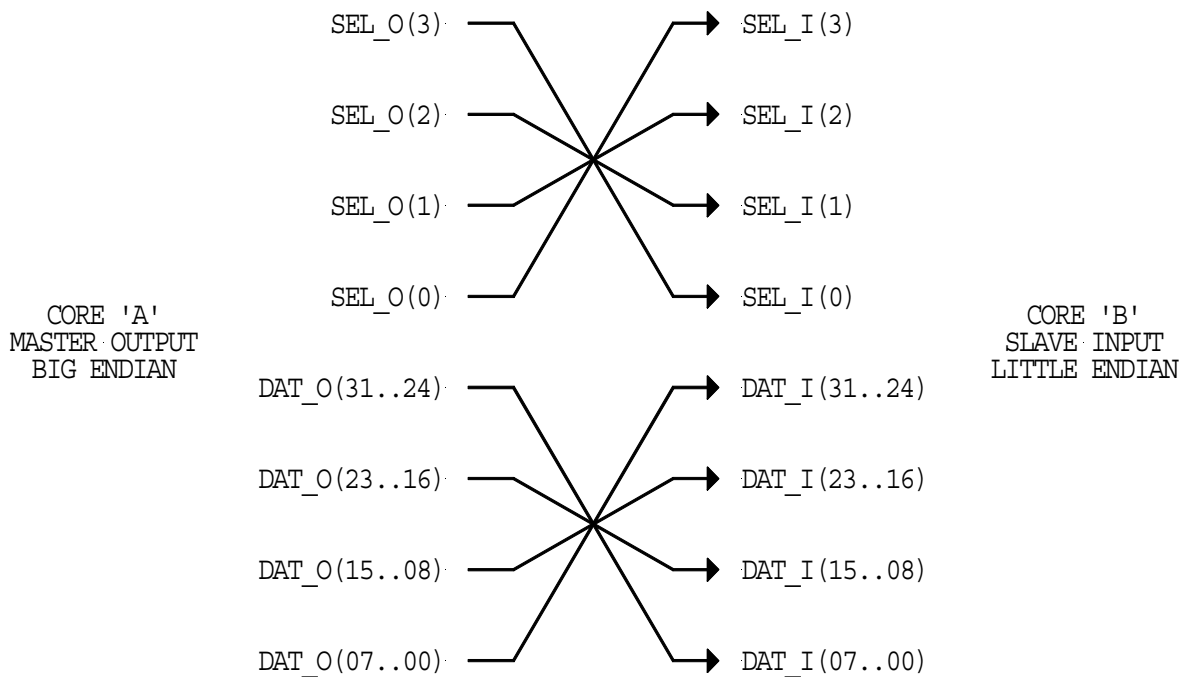


Figure A-8. BIG ENDIAN to LITTLE ENDIAN conversion example.

A.6 SLAVE I/O Port Examples

In this section we'll investigate several examples of WISHBONE interface for SLAVE I/O ports. Our purpose is to:

- Show some simple examples of how the WISHBONE interface operates.
- Demonstrate how simple interfaces work in conjunction with standard logic primitives on FPGA and ASIC devices. This also means that very little logic is needed to implement the WISHBONE interface.
- Demonstrate the concept of *granularity*.
- Provide some portable design examples.
- Give examples of the WISHBONE DATASHEET.
- Show VHDL implementation examples.

A.6.1 Simple 8-bit SLAVE Output Port

Figure A-9 shows a simple 8-bit WISHBONE SLAVE output port. The entire interface is implemented with a standard 8-bit 'D-type' flip-flop register (with synchronous reset) and a single AND gate. During write cycles, data is presented at the data input bus [DAT_I(7..0)], and is latched at the rising edge of [CLK_I] when [STB_I] and [WE_I] are both asserted.




Figure A-9. Simple 8-bit WISHBONE SLAVE output port.

The state of the output port can be monitored by a MASTER by routing the output data lines back to [DAT_O(7..0)]. During read cycles the AND gate prevents erroneous data from being latched into the register.

This circuit is highly portable, as all FPGA and ASIC target devices support D-type flip-flops with clock enable and synchronous reset inputs.

The circuit also demonstrates how the WISHBONE interface requires little or no logic overhead. In this case, the WISHBONE interface does not require any extra logic gates whatsoever. This is because WISHBONE is designed to work in conjunction with standard, synchronous and combinatorial logic primitives that are available on most FPGA and ASIC devices.

The WISHBONE specification requires that the interface be documented. This is done in the form of the WISHBONE DATASHEET. The standard does not specify the form of the datasheet. For example, it can be part of a comment field in a VHDL or Verilog® source file or part of a technical reference manual for the IP core. Table A-1 shows one form of the WISHBONE DATASHEET for the 8-bit output port circuit.

The purpose of the WISHBONE DATASHEET is to promote design reuse. It forces the originator of the IP core to document how the interface operates. This makes it easier for another person to re-use the core.

Table A-1. WISHBONE DATASHEET for the 8-bit output port example.																	
Description	Specification																
General description:	8-bit SLAVE output port.																
Supported cycles:	SLAVE, READ/WRITE SLAVE, BLOCK READ/WRITE SLAVE, RMW																
Data port, size:	8-bit																
Data port, granularity:	8-bit																
Data port, maximum operand size:	8-bit																
Data transfer ordering:	Big endian and/or little endian																
Data transfer sequencing:	Undefined																
Supported signal list and cross reference to equivalent WISHBONE signals:	<table border="0"> <thead> <tr> <th><u>Signal Name</u></th> <th><u>WISHBONE Equiv.</u></th> </tr> </thead> <tbody> <tr> <td>ACK_O</td> <td>ACK_O</td> </tr> <tr> <td>CLK_I</td> <td>CLK_I</td> </tr> <tr> <td>DAT_I(7..0)</td> <td>DAT_I()</td> </tr> <tr> <td>DAT_O(7..0)</td> <td>DAT_O()</td> </tr> <tr> <td>RST_I</td> <td>RST_I</td> </tr> <tr> <td>STB_I</td> <td>STB_I</td> </tr> <tr> <td>WE_I</td> <td>WE_I</td> </tr> </tbody> </table>	<u>Signal Name</u>	<u>WISHBONE Equiv.</u>	ACK_O	ACK_O	CLK_I	CLK_I	DAT_I(7..0)	DAT_I()	DAT_O(7..0)	DAT_O()	RST_I	RST_I	STB_I	STB_I	WE_I	WE_I
<u>Signal Name</u>	<u>WISHBONE Equiv.</u>																
ACK_O	ACK_O																
CLK_I	CLK_I																
DAT_I(7..0)	DAT_I()																
DAT_O(7..0)	DAT_O()																
RST_I	RST_I																
STB_I	STB_I																
WE_I	WE_I																

Figure A-10 shows a VHDL implementation of same circuit. The WBOPRT08 entity implements the all of the functions shown in the schematic diagram of Figure A-9.

```

library ieee;
use ieee.std_logic_1164.all;

entity WBOPRT08 is
port(
    -- WISHBONE SLAVE interface:

    ACK_O:      out   std_logic;
    CLK_I:      in    std_logic;
    DAT_I:      in    std_logic_vector( 7 downto 0 );
    DAT_O:      out   std_logic_vector( 7 downto 0 );
    RST_I:      in    std_logic;
    STB_I:      in    std_logic;
    WE_I:      in    std_logic;

    -- Output port (non-WISHBONE signals):

    PRT_O:      out   std_logic_vector( 7 downto 0 )
);
end entity WBOPRT08;

architecture WBOPRT081 of WBOPRT08 is

    signal Q: std_logic_vector( 7 downto 0 );

begin

    REG: process( CLK_I )
    begin

        if( rising_edge( CLK_I ) ) then

            if( RST_I = '1' ) then
                Q <= B"00000000";
            elsif( (STB_I and WE_I) = '1' ) then
                Q <= DAT_I( 7 downto 0 );
            else
                Q <= Q;
            end if;

        end if;

    end process REG;

    ACK_O <= STB_I;
    DAT_O <= Q;
    PRT_O <= Q;

end architecture WBOPRT081;

```

Figure A-10. VHDL implementation of the 8-bit output port interface.

A.6.2 Simple 16-bit SLAVE Output Port With 16-bit Granularity

Figure A-11 shows a simple 16-bit WISHBONE SLAVE output port. Table A-2 shows the WISHBONE DATASHEET for this design. It is identical to the 8-bit port shown earlier, except that the data bus is wider. Also, this port has 16-bit granularity. In the next section, it will be compared to a 16-bit port with 8-bit granularity.




Figure A-11. Simple 16-bit WISHBONE SLAVE output port with 16-bit granularity

Table A-2. WISHBONE DATASHEET for the 16-bit output port with 16-bit granularity.																	
Description	Specification																
General description:	16-bit SLAVE output port.																
Supported cycles:	SLAVE, READ/WRITE SLAVE, BLOCK READ/WRITE SLAVE, RMW																
Data port, size:	16-bit																
Data port, granularity:	16-bit																
Data port, maximum operand size:	16-bit																
Data transfer ordering:	Big endian and/or little endian																
Data transfer sequencing:	Undefined																
Supported signal list and cross reference to equivalent WISHBONE signals:	<table border="1"> <thead> <tr> <th>Signal Name</th> <th>WISHBONE Equiv.</th> </tr> </thead> <tbody> <tr> <td>ACK_O</td> <td>ACK_O</td> </tr> <tr> <td>CLK_I</td> <td>CLK_I</td> </tr> <tr> <td>DAT_I(15..0)</td> <td>DAT_I()</td> </tr> <tr> <td>DAT_O(15..0)</td> <td>DAT_O()</td> </tr> <tr> <td>RST_I</td> <td>RST_I</td> </tr> <tr> <td>STB_I</td> <td>STB_I</td> </tr> <tr> <td>WE_I</td> <td>WE_I</td> </tr> </tbody> </table>	Signal Name	WISHBONE Equiv.	ACK_O	ACK_O	CLK_I	CLK_I	DAT_I(15..0)	DAT_I()	DAT_O(15..0)	DAT_O()	RST_I	RST_I	STB_I	STB_I	WE_I	WE_I
Signal Name	WISHBONE Equiv.																
ACK_O	ACK_O																
CLK_I	CLK_I																
DAT_I(15..0)	DAT_I()																
DAT_O(15..0)	DAT_O()																
RST_I	RST_I																
STB_I	STB_I																
WE_I	WE_I																

A.6.3 Simple 16-bit SLAVE Output Port With 8-bit Granularity

Figure A-12 shows a simple 16-bit WISHBONE SLAVE output port. This port has 8-bit granularity, which means that data can be transferred 8 or 16-bits at a time.




Figure A-12. Simple 16-bit WISHBONE SLAVE output port with 8-bit granularity.

This circuit differs from the aforementioned 16-bit port because it has 8-bit granularity. This means that the 16-bit register can be accessed with either 8 or 16-bit bus cycles. This is accomplished by selecting the high or low byte of data with the select lines [SEL_I(1..0)]. When [SEL_I(0)] is asserted, the low byte is accessed. When [SEL_I(1)] is asserted, the high byte is accessed. When both are asserted, the entire 16-bit word is accessed.

The circuit also demonstrates the proper use of the [STB_I] and [SEL_I(0)] lines for slave devices. The [STB_I] signal operates much like a chip select signal, where the interface is selected when [STB_I] is asserted. The [SEL_I(0)] lines are only used to determine where data is placed by the MASTER or SLAVE during read and write cycles.

In general, the [SEL_I(0)] signals should never be used by the SLAVE to determine when the IP core is being accessed by a master. They should only be used to determine where data is placed on the data input and output buses. Stated another way, the MASTER will assert the select signals [SEL_O(0)] during every bus cycle, but a particular slave is only accessed when it monitors

that its [STB_I] input is asserted. Stated another way, the [STB_I] signal is generated by address decode logic within the WISHBONE interconnect, but the [SEL_I()] signals may be broadcasted to all SLAVE devices.

Table A-3 shows the WISHBONE DATASHEET for this IP core. This is very similar to the 16-bit data port with 16-bit granularity, except that the granularity has been changed to 8-bits.

It should also be noted that the datasheet specifies that the circuit will work with READ/WRITE, BLOCK READ/WRITE and RMW cycles. This means that the circuit will operate normally when presented with these cycles. It is left as an exercise for the user to verify that the circuit will indeed work with all three of these cycles.

Table A-3. WISHBONE DATASHEET for the 16-bit output port with 8-bit granularity.																	
Description	Specification																
General description:	16-bit SLAVE output port with 8-bit granularity.																
Supported cycles:	SLAVE, READ/WRITE SLAVE, BLOCK READ/WRITE SLAVE, RMW																
Data port, size:	16-bit																
Data port, granularity:	8-bit																
Data port, maximum operand size:	16-bit																
Data transfer ordering:	Big endian and/or little endian																
Data transfer sequencing:	Undefined																
Supported signal list and cross reference to equivalent WISHBONE signals:	<table style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left; border-bottom: 1px solid black;"><u>Signal Name</u></th> <th style="text-align: left; border-bottom: 1px solid black;"><u>WISHBONE Equiv.</u></th> </tr> </thead> <tbody> <tr> <td>ACK_O</td> <td>ACK_O</td> </tr> <tr> <td>CLK_I</td> <td>CLK_I</td> </tr> <tr> <td>DAT_I(15..0)</td> <td>DAT_I()</td> </tr> <tr> <td>DAT_O(15..0)</td> <td>DAT_O()</td> </tr> <tr> <td>RST_I</td> <td>RST_I</td> </tr> <tr> <td>STB_I</td> <td>STB_I</td> </tr> <tr> <td>WE_I</td> <td>WE_I</td> </tr> </tbody> </table>	<u>Signal Name</u>	<u>WISHBONE Equiv.</u>	ACK_O	ACK_O	CLK_I	CLK_I	DAT_I(15..0)	DAT_I()	DAT_O(15..0)	DAT_O()	RST_I	RST_I	STB_I	STB_I	WE_I	WE_I
<u>Signal Name</u>	<u>WISHBONE Equiv.</u>																
ACK_O	ACK_O																
CLK_I	CLK_I																
DAT_I(15..0)	DAT_I()																
DAT_O(15..0)	DAT_O()																
RST_I	RST_I																
STB_I	STB_I																
WE_I	WE_I																

Figure A-13 shows a VHDL implementation of same circuit. The WBOPRT16 entity implements the all of the functions shown in the schematic diagram of Figure A-12.

```

entity WBOPRT16 is
port (
    -- WISHBONE SLAVE interface:
    ACK_O:      out   std_logic;
    CLK_I:      in    std_logic;
    DAT_I:      in    std_logic_vector( 15 downto 0 );
    DAT_O:      out   std_logic_vector( 15 downto 0 );
    RST_I:      in    std_logic;
    SEL_I:      in    std_logic_vector( 1 downto 0 );
    STB_I:      in    std_logic;
    WE_I:       in    std_logic;

    -- Output port (non-WISHBONE signals):
    PRT_O:      out   std_logic_vector( 15 downto 0 )
);
end entity WBOPRT16;

architecture WBOPRT161 of WBOPRT16 is
    signal QH: std_logic_vector( 7 downto 0 );
    signal QL: std_logic_vector( 7 downto 0 );
begin

    REG: process( CLK_I )
    begin
        if( rising_edge( CLK_I ) ) then
            if( RST_I = '1' ) then
                QH <= B"00000000";
            elsif( (STB_I and WE_I and SEL_I(1)) = '1' ) then
                QH <= DAT_I( 15 downto 8 );
            else
                QH <= QH;
            end if;
        end if;

        if( rising_edge( CLK_I ) ) then
            if( RST_I = '1' ) then
                QL <= B"00000000";
            elsif( (STB_I and WE_I and SEL_I(0)) = '1' ) then
                QL <= DAT_I( 7 downto 0 );
            else
                QL <= QL;
            end if;
        end if;

    end process REG;

    ACK_O <= STB_I;
    DAT_O( 15 downto 8 ) <= QH;
    DAT_O( 7 downto 0 ) <= QL;
    PRT_O( 15 downto 8 ) <= QH;
    PRT_O( 7 downto 0 ) <= QL;

end architecture WBOPRT161;

```

Figure A-13. VHDL implementation of the 16-bit output port with 8-bit granularity.

A.7 WISHBONE Memory Interfacing

In this section we'll examine WISHBONE memory interfacing and present some examples. The purpose of this section is to:

- Introduce the FASM synchronous RAM and ROM models.
- Show a simple example of how the WISHBONE interface operates.
- Demonstrate how simple interfaces work in conjunction with standard logic primitives on FPGA and ASIC devices. This also means that very little logic (if any) is needed to implement the WISHBONE interface.
- Present a WISHBONE DATASHEET example for a memory element.
- Describe portability issues with regard to FPGA and ASIC memory elements.

A.7.1 FASM Synchronous RAM and ROM Model

The WISHBONE interface can be connected to any type of RAM or ROM memory. However, some types will be faster and more efficient than others. If the memory interface closely resembles the WISHBONE interface, then everything will run very fast. If the memory is significantly different than WISHBONE, then everything will slow down. This is such a fundamental and important issue that both the WISHBONE interface and its bus cycles were designed around the most common memory interface that could be found.

This was very problematic in the original WISHBONE concept. That's because there are very few portable RAM and ROM types used across all both FPGA and ASIC devices. In fact, the most common memory type that could be found are what we call 'FASM', or the FPGA and ASIC Subset Model⁷.

The FASM synchronous RAM model conforms to the connection and timing diagram shown in Figure A-14. The WISHBONE bus cycles all are designed to interface directly to this type of RAM. During write cycles, FASM Synchronous RAM stores input data at the indicated address whenever: (a) the write enable (WE) input is asserted, and (b) there is a rising clock edge.

During read cycles, FASM Synchronous RAM works like an asynchronous ROM. Data is fetched from the address indicated by the ADR() inputs, and is presented at the data output (DOUT). The clock input is ignored. However, during write cycles, the output data is updated immediately during a write cycle.

A good exercise for the user is to compare the FASM Synchronous RAM cycles to the WISHBONE SINGLE READ/WRITE, BLOCK READ/WRITE and READ-MODIFY-WRITE cycles. You will find that these three bus cycles operate in an identical fashion to the FASM Synchronous RAM model. They are so close, in fact, that FASM RAMs can be interfaced to WISHBONE with as little as one NAND gate.

⁷ The original FASM model actually encompasses many type of devices, but in this tutorial we'll focus only on the FASM synchronous RAM and ROM models.

While most FPGA and ASIC devices will provide RAM that follows the FASM guidelines, you will probably find that most devices also support other types of memories as well. For example, in some brands of FPGA you will find block memories that use a different interface. Some of these will still interface very smoothly to WISHBONE, while others will introduce a wait-state. In all cases that we found, all FPGAs and most ASICs did support at least one style of FASM memory.




Figure A-14. Generic FASM synchronous RAM connection and timing diagram.

The FASM ROM closely resembles the FASM RAM during its read cycle. FASM ROM conforms to the connection and timing diagram shown in Figure A-15.

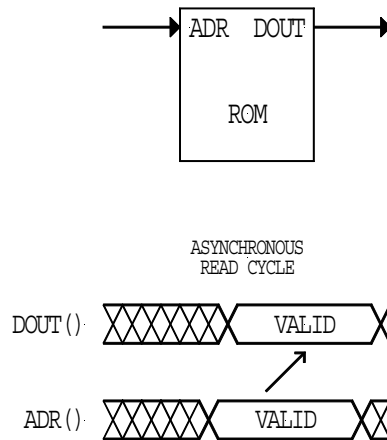


Figure A-15. FASM asynchronous ROM connection and timing diagram.

A.7.2 Simple 16 x 8-bit SLAVE Memory Interface

Figure A-16 shows a simple 8-bit WISHBONE memory. The 16 x 8-bit memory is formed from two 16 x 4-bit FASM compatible synchronous memories. Besides the memory elements, the entire interface is implemented with a single AND gate. During write cycles, data is presented at the data input bus [DAT_I(7..0)], and is latched at the rising edge of [CLK_I] when [STB_I] and [WE_I] are both asserted. During read cycles, the memory output data (DO) is made available at the data output port [DAT_O(7..0)].




Figure A-16. Simple 16 x 8-bit SLAVE memory.

The memory circuit does not have a reset input. That's because most RAM memories do not have a reset capability.

The circuit also demonstrates how the WISHBONE interface requires little or no logic overhead. In this case, the WISHBONE interface requires a single AND gate. This is because WISHBONE is designed to work in conjunction with standard, synchronous and combinatorial logic primitives that are available on most FPGA and ASIC devices.

The WISHBONE specification requires that the interface be documented. This is done in the form of the WISHBONE DATASHEET. The standard does not specify the form of the datasheet. For example, it can be part of a comment field in a VHDL or Verilog® source file or part of a technical reference manual for the IP core. Table A-4 shows one form of the WISHBONE DATASHEET for the 16 x 8-bit memory IP core.

The purpose of the WISHBONE DATASHEET is to promote design reuse. It forces the originator of the IP core to document how the interface operates. This makes it easier for another person to re-use the core.

Table A-4. WISHBONE DATASHEET for the 16 x 8-bit SLAVE memory.																	
Description	Specification																
General description:	16 x 8-bit memory IP core.																
Supported cycles:	SLAVE, READ/WRITE SLAVE, BLOCK READ/WRITE SLAVE, RMW																
Data port, size:	8-bit																
Data port, granularity:	8-bit																
Data port, maximum operand size:	8-bit																
Data transfer ordering:	Big endian and/or little endian																
Data transfer sequencing:	Undefined																
Clock frequency constraints:	NONE (determined by memory primitive)																
Supported signal list and cross reference to equivalent WISHBONE signals:	<table border="1"> <thead> <tr> <th>Signal Name</th> <th>WISHBONE Equiv.</th> </tr> </thead> <tbody> <tr> <td>ACK_O</td> <td>ACK_O</td> </tr> <tr> <td>ADR_I(3..0)</td> <td>ADR_I()</td> </tr> <tr> <td>CLK_I</td> <td>CLK_I</td> </tr> <tr> <td>DAT_I(7..0)</td> <td>DAT_I()</td> </tr> <tr> <td>DAT_O(7..0)</td> <td>DAT_O()</td> </tr> <tr> <td>STB_I</td> <td>STB_I</td> </tr> <tr> <td>WE_I</td> <td>WE_I</td> </tr> </tbody> </table>	Signal Name	WISHBONE Equiv.	ACK_O	ACK_O	ADR_I(3..0)	ADR_I()	CLK_I	CLK_I	DAT_I(7..0)	DAT_I()	DAT_O(7..0)	DAT_O()	STB_I	STB_I	WE_I	WE_I
Signal Name	WISHBONE Equiv.																
ACK_O	ACK_O																
ADR_I(3..0)	ADR_I()																
CLK_I	CLK_I																
DAT_I(7..0)	DAT_I()																
DAT_O(7..0)	DAT_O()																
STB_I	STB_I																
WE_I	WE_I																
Special requirements:	Circuit assumes the use of synchronous RAM primitives with asynchronous read capability.																

A.7.3 Memory Primitives and the [ACK_O] Signal

Memory primitives, specific to the FPGA or ASIC target device, are usually used for the RAM storage elements. That's because most high-level languages (such as VHDL and Verilog®) don't synthesize these very efficiently. For this reason, the user should verify that the memory primitives are available for the target device.

The memory circuits shown above are highly portable, but do assume that FASM compatible memories are available. During *write* cycles, most synchronous RAM primitives latch the input data when at the rising clock edge when the write enable input is asserted. However, during *read* cycles the RAM primitives may behave in different ways.

There are two types of RAM primitives that are generally found on FPGA and ASIC devices: (a) those that synchronously present data at the output after the rising edge of the clock input, and (b) those that asynchronously present data at the output after the address is presented to the RAM element.

The circuit assumes that the RAM primitive is the FASM, asynchronous read type. That's because during read cycles the WISHBONE interface assumes that output data is valid at the rising

[CLK_I] edge following the assertion of the [ACK_O] output. Since the circuit ties the [STB_I] signal back to the [ACK_O] signal, the asynchronous read RAM is needed on the circuit shown here.

For this reason, if *synchronous* read type RAM primitives are used, then the circuit must be modified to insert a single wait-state during all read cycles. This is quite simple to do, and only requires an additional flip-flop and gate in the [ACK_O] circuit.

Furthermore, it can be seen that the circuit operates faster if the asynchronous read type RAM primitives are used. That's because the [ACK_O] signal can be asserted immediately after the assertion of [STB_I]. If the synchronous read types are used, then a single-clock wait-state must be added.

A.8 Customization with Tags and TAG TYPEs

One fundamental problem with traditional microcomputer buses is that they can't be customized very easily. That's because they rely on fixed printed circuit boards and connectors. Customization of these components are costly (in terms of time and money) and very often result in system components that aren't compatible.

SoC interconnections like WISHBONE are much easier to customize. That's because their interconnections are programmable in FPGA and ASIC target devices. Furthermore, they are supported by a rich set of development tools such as the VHDL and Verilog® hardware description languages, as well as a wide variety of routing tools. These make it possible to quickly change the interface. However, this can also lead to incompatible interfaces.

WISHBONE allows MASTER and SLAVE interfaces to be customized, while still retaining a highly compatible interface. This is accomplished with a technique known as a *tagged architecture*. This technique is not new to the microcomputer bus industry. The IEEE Authoritative Dictionary defines it as: a computer architecture in which each word is 'tagged' as either an instruction or a unit of data. A similar concept is used by WISHBONE.

Custom tags can be added to the WISHBONE interface as long as they conform to something called a TAG TYPE. There are three general TAG TYPEs defined by WISHBONE. They include an address tag, a set of data tags and a cycle tag. When a custom signal is added to a WISHBONE interface it is assigned a TAG TYPE. This indicates the exact timing that the signal must adhere to. Table A-5 shows some examples for each of the TAG TYPEs.

Table A-5. TAG TYPE examples.			
<u>General Type of tag (MASTER):</u>	<u>Address Tag</u>	<u>Data Tag</u>	<u>Cycle Tag</u>
<u>Assigned TAG TYPE:</u>	<u>TGA_O()</u>	<u>TGD_I() & TGD_O()</u>	<u>TGC_O()</u>
<u>Examples:</u>	<ul style="list-style-type: none"> • <u>Address width (e.g. 16-bit, 24-bit etc.)</u> • <u>Memory management (e.g. user vs. protected address)</u> 	<ul style="list-style-type: none"> • <u>Parity bits</u> • <u>Error correction codes (ECC)</u> • <u>Time stamps</u> • <u>Endian type</u> 	<ul style="list-style-type: none"> • <u>Cycle type (e.g. SINGLE, BLOCK and RMW cycles)</u> • <u>Cache control</u> • <u>Interrupt acknowledge cycles</u> • <u>Instruction vs. data cycles</u>

For example, consider a MASTER interface that is adapted to generate a parity bit called [PAR_O]. Since this signal modifies the output data bus, it would be assigned a data TAG TYPE of: TGD_O(). If the MASTER's input bus were similarly modified, it would be assigned a TAG TYPE of: TGD_O(). Assignment of the TAG TYPE does two things: (a) it links the parity bit [PAR_O] to the data bus and (b) it defines the exact timing during SINGLE, BLOCK and RMW bus cycles. However, the creator of the MASTER interface would still need to define the operation of [PAR_O] in the WISHBONE DATASHEET.

Also note that the [TGD_O()] TAG TYPE is an array (i.e. it includes a parenthesis), while the [PAR_O] tag is not arrayed. This is allowed under the WISHBONE specification.

A.98 Point-to-point Interconnection Example

Now that we've reviewed some of the WISHBONE basics, it's time to try them out with a simple example. In this section we'll create a complete WISHBONE system with a point-to-point interconnection. The system includes a 32-bit MASTER interface to a DMA⁸ unit, and a 32-bit SLAVE interface to a memory. In this example the DMA transfers data to and from the memory using block transfer cycles.

The purpose of this system is to create a portable benchmarking device. Although the system is very simple, it does allow the user to determine the typical maximum speeds and minimum sizes on any given FPGA or ASIC target device⁹.

⁸ DMA: Direct Memory Access.

⁹ Benchmarking can be a difficult thing to do. On this system the philosophy was to create a 'real-world' SoC that estimates 'typical maximum' speeds and 'typical minimum' size. It's akin to the 'flight envelope' of an airplane. A

Source code for this example can be found in the [WISHBONE Public Domain Library for VHDL](#) (under 'EXAMPLE1' in the EXAMPLES folder). The library also has detailed documentation for the library modules, including detailed circuit descriptions and timing diagrams for the INTERCON, SYSCON, DMA and memory [modules/interfaces](#). The reader is encouraged to review and experiment with all of these files.

Figure A-17 shows the system. The WISHBONE interconnection system (INTERCON) can be found in file ICN0001a. That system connects a simple DMA MASTER (DMA0001a) to an 8 x 32-bit register based memory SLAVE (MEM0002a). The reset and clock signals are generated by system controller SYSCON (SYC0001a).

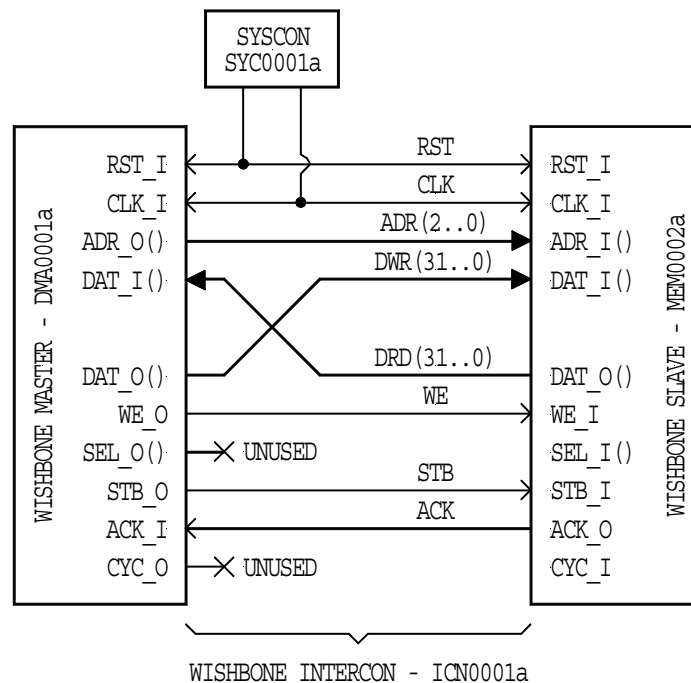


Figure A-17. Point-to-point interconnection example.

This system was synthesized and routed on two styles of Xilinx¹⁰ FPGA: the Spartan 2 and the Virtex 2. For benchmarking purposes the memories were altered so that they used Xilinx dis-

flight envelope is a graph that shows the altitude vs. the speed of the aircraft. It's one 'benchmark' for the airplane. While the graph may show that the airplane is capable of flying at MACH 2.3 at an altitude of 28,000 meters, it may never actually fly in that situation during its lifetime. The graph is simply a tool for quickly understanding the engineering limits of the design. The same is true for the WISHBONE benchmarks given in this tutorial. However, having said this it is important to remember that the benchmarks are real systems, and do include all of the logic and routing resources needed to implement the design.

¹⁰ Xilinx is a registered trademark of Xilinx, Inc.

tributed RAMs instead of the register RAMs in MEM0002a. A memory interface for the Xilinx RAMs can be found in MEM0001a, which is substituted for MEM0002a.

It should be noted that the Xilinx distributed RAMs are quite efficient on the WISHBONE interface. As can be seen in the source code, only a single ‘AND’ gate was needed to interface the RAM to WISHBONE.

The system for the Xilinx Spartan 2 was synthesized and operated on a Silicore evaluation board. This was a ‘reality check’ that verified that things actually routed and worked as expected. Some of the common signals were brought out to test points on the evaluation board. These were monitored with an HP54620a logic analyzer to verify the operation. Figure A-18 shows an example trace from the logic analyzer. Address lines, data write lines and several control signals were viewed. That Figure shows a write cycle with eight phases followed by a read cycle with eight phases. The data lines always show 0x67 because that’s the data transferred by the DMA in this example.

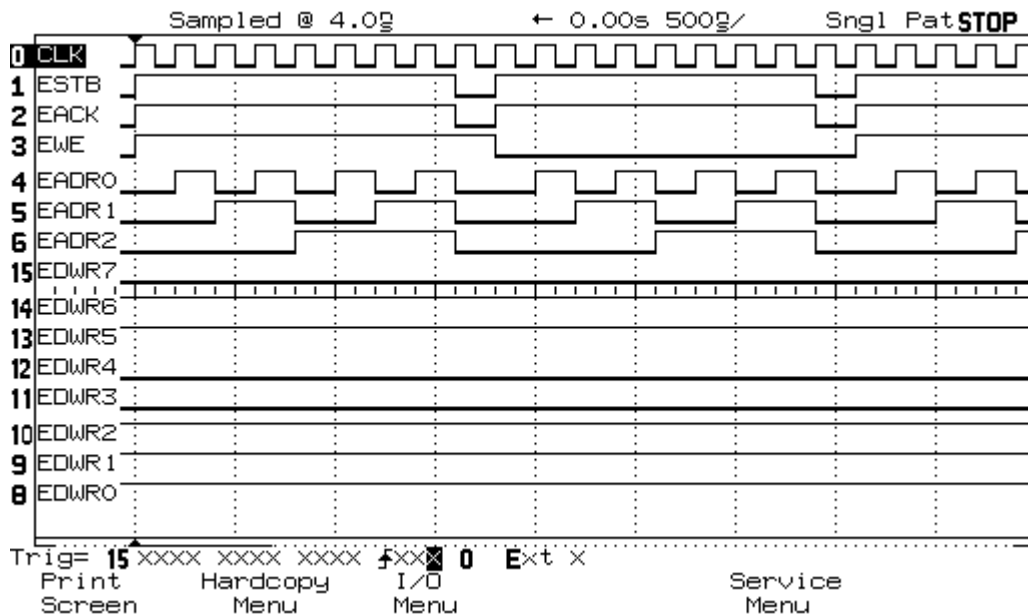


Figure A-18. Logic analyzer trace on the Spartan 2 evaluation board¹¹.

Table A-65 shows the speed of the system after synthesis and routing. The Spartan 2 benchmarked at about 428 Mbyte/sec, and was tested in hardware (HW TEST). The Virtex 2 part was synthesized and routed, but was only tested under software (SW TEST).

¹¹ The logic analyzer samples at 500 Mhz, so the SoC was slowed down to make the traces look better. This trace was taken with a SoC clock speed of 5 MHz. Slowing the clock down is also a good way to verify that the speed of the WISHBONE interface can be ‘throttled’ up and down.

Table A-65. 32-bit Point-to-point Interconnection Benchmark Results							
MFG & Type	Part Number	HW TEST	SW TEST	Size	Timing Constraint (MIN)	Actual Speed (MAX)	Data Transfer Rate (MAX)
Xilinx Spartan 2 (FPGA)	XC2S50-5-PQ208C	√		53 SLICE	100 MHz	107 MHz	428 Mbyte/sec
Xilinx Virtex 2 (FPGA)	XC2V40-4-FG256C		√	53 SLICE	200 MHz	203 MHz	812 Mbyte/sec

Notes:

VHDL synthesis tool: Altium Accolade PeakFPGA 5.30a

Router: Xilinx Alliance 3.3.06I_V2_SE2

Hardware evaluation board: Silicore 170101-00 Rev 1.0

Listed size was reported by the router.

Spartan 2 test used '-5' speed grade part (slower than the faster '-6' part).

A.109 Shared Bus Example

Now that we've built a WISHBONE point-to-point interconnection, it's time to look at a more complex SoC design. In this example, we'll create a 32-bit shared bus system with four MASTERS and four SLAVES. Furthermore, we'll re-use the same DMA, memory and SYSCON [modules/interfaces](#) that we used in the point-to-point interconnection example above. This will demonstrate how WISHBONE interfaces can be adapted to many different system topologies.

This example will require the introduction of some new concepts. As the system integrator, we'll need to make some decisions about how we want our system to work. After that, we'll need to create the various parts of the design in order to finish the job. Some of the decisions and tasks include:

- Choosing between multiplexed and non-multiplexed bus topology.
- Choosing between three-state and multiplexor based interconnection logic.
- Creating the interconnection topology.
- Creating an address map (using variable address decoding).
- Creating a four level round-robin arbiter.
- Creating and benchmarking the system.

A.109.1 Choosing Between Multiplexed and Non-multiplexed Bus Topology

The first step in designing a shared bus is to determine how we'll move our data around the system. In this section we'll explore multiplexed and non-multiplexed buses, and explore some of the trade-offs between them.

The big advantage of multiplexed buses is that they reduce the number of interconnections by routing different types of data over the same set of signal lines. The most common form of multiplexed bus is where address and data lines share a common set of signals. A multiplexed bus is shown in Figure A-19. For example, a 32-bit address bus and 32-bit data bus can be combined to form a 32-bit common address/data bus. This reduces the number of routed signals from 64 to 32.

The major disadvantage of the multiplexed bus is that it takes twice as long to move the information. In this case a non-multiplexed, synchronous bus can generally move address and data information in as little as one clock cycle. Multiplexed address and data buses require at least two clock cycles to move the same information.

Since we're creating a benchmarking system that is optimized for speed, we'll use the non-multiplexed scheme for this example. This will allow us to move one data operand during every clock cycle.

It should be noted that multiplexed buses have long been used in the electronics industry. In semiconductor chips the technique is used to reduced the number of pins on a chip. In the mi-

microcomputer bus industry the technique is often used to reduce the number of backplane connector pins.

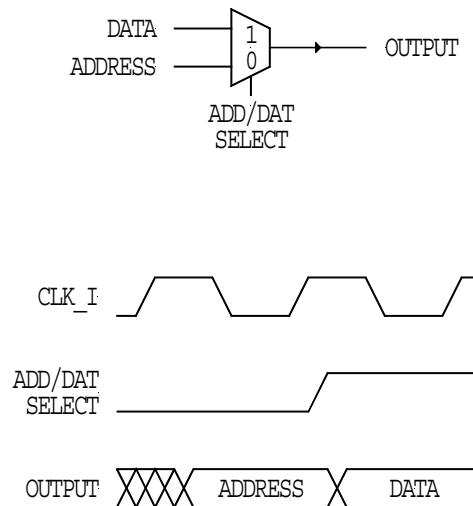


Figure A-19. Circuit and timing diagram for a multiplexed address/data bus.

A.109.2 Choosing Between Three-State and Multiplexor Interconnection Logic

WISHBONE interconnections can use three-state¹² or multiplexor logic to move data around a SoC. The choice depends on what makes sense in the application, and what's available on the integrated circuit.

Three-state I/O buffers have long been used in the semiconductor and microcomputer bus industries. These reduce the number of signal pins on an interface. In microcomputer buses with master-slave architectures, the master that 'owns' the bus turns its buffers 'on', while the other master(s) turn their buffers 'off'¹³. This prevents more than one bus master from driving any signal line at any given time. A similar situation also occurs at the slave end. There, a slave that participates in a bus cycle enables its output buffers during read cycles.

In WISHBONE, all IP core interfaces have 'in' and 'out' signals on the address, data and other internal buses. This allows the interface to be adapted to point-to-point, multiplexed and three-state I/O interconnections. Figure A-20 shows how the 'in' and 'out' signals can be connected to a three-state I/O bus¹⁴.

¹² Three-state buffers are sometimes called Tri-State® buffers. Tri-State® is a registered trademark of National Semiconductor Corporation.

¹³ Here, 'on' and 'off' refer to the three-state and non three-state conditions, respectively.

¹⁴ Also note that the resistor/current source shown in the figure can also be a pull-down resistor or a bus terminator, or eliminated altogether.

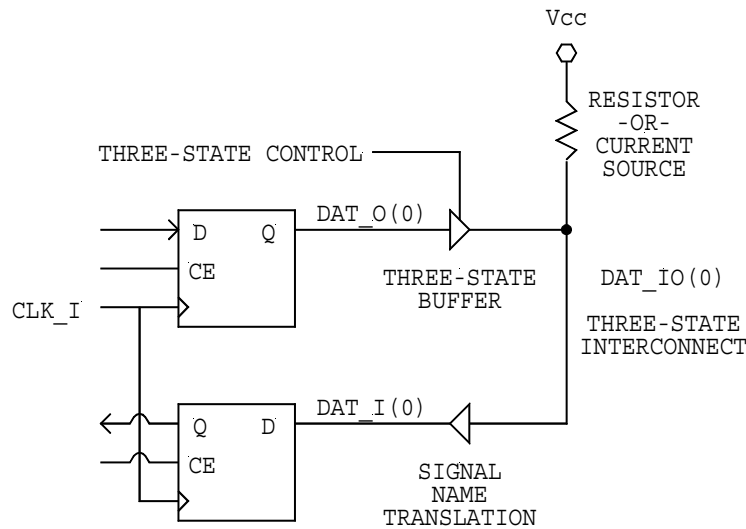


Figure A-20. Connection of a data bus bit to a three-state interconnection.

A simple SoC interconnection that uses three-state I/O buffers is shown in the block diagram of Figure A-21(a). There, the data buses on two master and two slave modules are interconnected with three-state logic. However, this approach has two major drawbacks. First, it is inherently slower than direct interconnections. That's because there are always minimum timing parameters that must be met to turn buffers on-and-off. Second, many IC devices do not have any internal three-state routing resources available to them, or they are very restrictive in terms of location or quantity of these interconnects.

As shown in Figure A-21(b), the SoC bus can be adapted to use multiplexor logic to achieve the same goal. The main advantage of this approach is that it does not use the three-state routing resources which are not available on many FPGA and ASIC devices.


The main disadvantage of the multiplexor logic interconnection is that it requires a larger number of routed interconnects and logic gates (which are not required with the three-state bus approach).

However, there is also a growing body of evidence that suggests that this type of interconnection is easier to route in both FPGA and ASIC devices. Although this is very difficult to quantify, the author has found that the multiplexor logic interconnection is quite easily handled by standard FPGA and ASIC routers. This is because:


- Three-state interconnections force router software to organize the SoC around the fixed three-state bus locations. In many cases, this constraint results in poorly optimized and/or slow circuits.

- Very often, ‘bit locations’ within a design are grouped together. In many applications, the multiplexor logic interconnection is easier to handle for place & route tools.
- Pre-defined, external I/O pin locations are easier to achieve with multiplexor logic interconnections. This is especially true with FPGA devices.

For the shared bus example we will use multiplexor logic for the interconnection. That’s because multiplexor logic interconnections are more portable than three-state logic designs. The shared bus design in this example is intended to be used on many brands of FPGA and ASIC devices.



(A) THREE-STATE BUS INTERCONNECTION



(B) MULTIPLEXOR LOGIC INTERCONNECTION

Figure A-21. Three-state bus interconnection vs. multiplexor logic interconnection.

A.109.3 Creating the Interconnection Topology

In the previous two sections it was decided to use multiplexor interconnections with non-multiplexed address and data buses. In this section we'll refine those concepts into a broad interconnection topology for our system. However, we'll save the details for later. For now, we're just interested in looking at some of the general issues.

In WISHBONE nomenclature, the interconnection is also called the INTERCON. The INTERCON is an IP Core that connects all of the MASTER ~~and~~; SLAVE ~~interfaces and SYSCON cores~~ together.

Figure A-22 shows the generic topology of an INTERCON that supports multiplexor interconnections with multiplexed address and data buses. By 'generic', we mean that there are 'N' MASTERS and SLAVEs shown in the diagram. The actual number of MASTER and SLAVE interfaces can be adjusted up or down, depending upon what's needed in the system. In the shared bus example we'll use four MASTERS and four SLAVEs. However, for now we'll think in more general terms.

A ~~module~~ interface called the SYSCON provides the system with a stable clock [CLK_O] and reset signal [RST_O]. For now, we'll assume that the clock comes from off-chip, and that the reset signal is synchronized from some global system reset.

After the initial system reset, one or more MASTERS request the interconnection, which we'll call a 'bus' for now. MASTERS do this by asserting their [CYC_O] signal. An arbiter, which we'll discuss shortly, determines which MASTER can use the bus. One clock edge after the assertion of a [CYC_O] signal the arbiter grants the bus to one of the MASTERS that requested it. It grants the bus by asserting grant lines GNT0 – GNTN and GNT(N..0). This informs both the INTERCON as to which MASTER can own the bus.

Once the bus is arbitrated, the output signals from the selected MASTER are routed, via multiplexors, onto the shared buses. For example, if MASTER #0 obtains the bus, then the address lines [ADR_O()] from MASTER #0 are routed to shared bus [ADR()]. The same thing happens to the data out [DAT_O()], select out [SEL_O()], write enable [WE_O] and strobe [STB_O] signals. The shared bus output signals are routed to the inputs on the SLAVE interfaces.

The arbiter grant lines are also used to enable the terminating signals [ACK_I], [RTY_I] and [ERR_I]. These are enabled at the MASTER that acquired the bus. For example, if MASTER #0 is granted the bus by the arbiter, then the [ACK_I], [RTY_I] and [ERR_I] are enabled at MASTER #0. Other MASTERS, who may also be requesting the bus, never receive a terminating signal, and therefore will wait until they are granted the bus.

During this interval the common address bus [ADR()] is driven with the address lines from the MASTER. The address lines are decoded by the address comparator, which splits the address space into 'N' sections. The decoded output from the comparator is used to select the slave by way of its strobe input [STB_I]. A SLAVE may only respond to a bus cycle when its [STB_I] is

asserted. This is also a wonderful illustration of the partial address decoding technique used by WISHBONE, which we'll discuss in depth below.

For example, consider a system with an addressing range of sixteen bits. If the addressing range were evenly split between all of the SLAVEs, then each SLAVE would be allocated 16 Kbytes of address space. This is shown in the address map of Figure A-23. In this case, the address comparator would decode bits [ADR(15..14)]. In actual practice the system integrator can alter the address map at his or her discretion.

Once a SLAVE is selected, it participates in the current bus cycle generated by the MASTER. In response to the cycle, the SLAVE must assert either its [ACK_O], [RTY_O] or [ERR_O] output. These signals are collected with an 'or' gate, and routed to the current MASTER.

Also note that during read cycles, the SLAVE places data on its [DAT_O()] bus. These are routed from the participating SLAVE to the current MASTER by way of a multiplexor. In this case, the multiplexor source is selected by some address lines which have been appropriately selected to switch the multiplexor.

Once the MASTER owning the bus has received an asserted terminating signal, it terminates the bus cycle by negating its strobe output [STB_O]. If the MASTER is in the middle of a block transfer cycle, it will generate another phase of the block transfer. If it's performing a SINGLE cycle, or if its at the end of a BLOCK cycle, the MASTER terminates the cycle by negating its [CYC_O] signal. This informs the MASTER that it's done with the bus, and that it can re-arbitrate it.




Figure A-22. WISHBONE shared bus with multiplexor interconnections.

0xFFFF	SLAVE #3
0xC000 0xBFFF	SLAVE #2
0x8000 0x7FFF	SLAVE #1
0x4000 0x3FFF	SLAVE #0
0x0000	

Figure A-23. Address map example.

A.109.4 Full vs. Partial Address Decoding

The address comparator in our INTERCON example is a good way to explain the concept of WISHBONE partial address decoding.

Many systems, including standard microcomputer buses like PCI and VMEbus, use *full address decoding*. Under that method, each slave module decodes the full address bus. For example, if a 32-bit address bus is used, then each slave decodes all thirty-two address bits (or at least a large portion of them).

SoC buses like WISHBONE use *partial address decoding* on slave modules. Under this method, each slave decodes only the range of addresses that it uses. For example, if the slave has only four registers, then the WISHBONE interface uses only two address bits. This technique has the following advantages:

- It facilitates high speed address decoders.
- It uses less redundant address decoding logic (i.e. fewer gates).
- It supports variable address sizing (between zero and 64-bits).
- It supports the variable interconnection scheme.
- It gives the system integrator a lot of flexibility in defining the address map.

For example, consider the serial I/O port (IP core) with the internal register set shown in Figure A-24(a). If *full address decoding* is used, then the IP core must include an address decoder to select the module. In this case, the decoder requires: 32 bits – 2 bits = 30 bits. In addition, the IP core would also contain logic to decode the lower two bits which are used to determine which I/O registers are selected.

If *partial address decoding* is used, then the IP core need only decode the two lower address bits ($2^2 = 4$). The upper thirty bits are decoded by logic outside of the IP core. In this case the decoder logic is shown in Figure A-24(b).

Standard microcomputer buses always use the full address decoding technique. That's because the interconnection method does not allow the creation of any new signals on the interface. However, in WISHBONE this limitation does not exist. WISHBONE allows the system integrator to modify the interconnection logic and signal paths.


One advantage of the partial address decoding technique is that the size of the address decoder (on the IP core) is minimized. This speeds up the interface, as decoder logic can be relatively slow. For example, a 30-bit full address decoder often requires at least 30 XOR gates, and a 30-input AND gate.

Another advantage of the partial address decoding technique is that less decoder logic is required. In many cases, only one 'coarse' address decoder is required. If full address decoding is used, then each IP core must include a redundant set of address decoders.

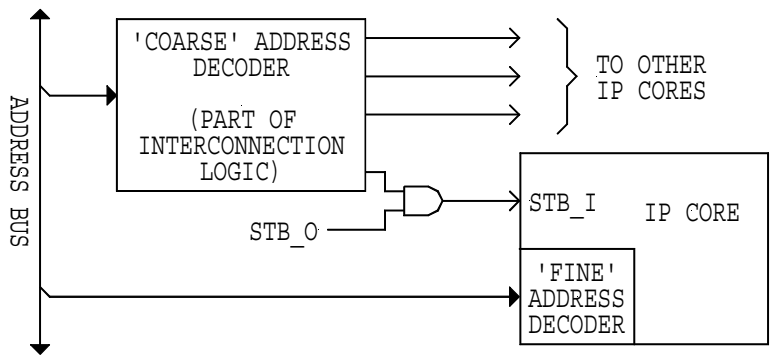
Another advantage of the partial address decoding technique is that it supports variable address sizing. For example, on WISHBONE the address path can be any size between zero and 64-bits. Slave modules are designed to utilize only the block of addresses that are required. In this case, the full address decoding technique cannot be used because the IP core designer is unaware of the size of the system address path.

Another advantage of the partial address decoding technique is that it supports the variable interconnection scheme. There, the type of interconnection logic is unknown to the IP core designer. The interconnection scheme must adapt to the types of slave IP cores that are used.

The major disadvantage of the partial address decoding technique is that the SoC integrator must define part of the address decoder logic for each IP core. This increases the effort to integrate the IP cores into the final SoC.



(A) SAMPLE IP CORE REGISTER SET



(B) IP CORE ADDRESS DECODING

Figure A-24. WISHBONE partial address decoding technique.

A.109.5 The System Arbiter

The system arbiter determines which MASTER can use the shared bus. The WISHBONE specification allows a variety of arbiters to be used. However, in this example a four level round-robin arbiter is used.

Round-robin arbiters give equal priority to all of the MASTERS. These arbiters grant the bus on a rotating basis much like the four position rotary switch shown in Figure A-25. When a MASTER relinquishes the bus (by negating its [CYC_O] signal), the switch is turned to the next position, and the bus is granted to the MASTER on that level. If a request is not pending on a certain level, the arbiter skips over that level and continues onto the next one. In this way all of the MASTERS are granted the bus on an equal basis.




Figure A-25. Round-robin arbiters grant the bus on a rotating basis much like a rotary switch.

Round-robin arbiters are popular in data acquisition systems where data is collected and placed into shared memory. Often these peripherals must off-load data to memory on an equal basis. Priority arbiters (where each MASTER is assigned a higher or lower level of priority) do not work well in these applications because some peripherals would receive more bus bandwidth than others, thereby causing data ‘gridlock’.

The arbiter used in this example can be found in the WISHBONE Public Domain Library for VHDL. ARB0001a is used for the example.

A.109.6 Creating and Benchmarking the System

The final task in our shared bus system example is to create and benchmark the entire system. The INTERCON in our example system is based on the generic shared bus topology that was described above. However, that system is fine tuned to give the exact features that we will need.

The final system supports four DMA0001a MASTERS, four MEM0002a memories (SLAVES), a 32-bit data bus, a five bit address bus, a single SYC0001a system controller and a ARB0001a

four level round-robin arbiter. The resulting VHDL file can be found under ICN0002a in the [WISHBONE Public Domain Library for VHDL](#).

In this application, the round-arbiter was chosen because all of the MASTERS are DMA controllers. That means that all four MASTERS continuously vie for the bus. If a priority arbiter were used, then only the one or two highest priority MASTERS would ever get the bus.

As we'll see shortly, the error and retry signals [ERR_I] and [RTY_I] are not supported on the MASTER and SLAVE interfaces on our example system. That's perfectly okay because these signals are optional in the WISHBONE specification. We could have added these signals in there, but they would have been removed by synthesis and router logic minimization tools.

Since all of the MASTERS and SLAVES on this system have identical port sizes and granularities, the select [SEL] interconnection has been omitted. This could have been added, but it wasn't needed.

The INTERCON system includes a partial address decoder for the SLAVES. This decoder creates the system address space shown in Figure A-26. The final address map is shown in Table A-76.

0x1F	SLAVE 3
0x18 0x17	
0x10 0x0F	SLAVE 2
0x08 0x07	SLAVE 1
0x00	SLAVE 0

Figure A-26. Address map used by the INTERCON example.

Table A-76. Address spaces used by INTERCON.			
DMA Master:	DMA's To:	At Addresses	Using Cycles
MASTER #0	SLAVE #0	0x00 – 0x07	BLOCK READ/WRITE
MASTER #1	SLAVE #1	0x08 – 0x0F	BLOCK READ/WRITE
MASTER #2	SLAVE #2	0x10 – 0x17	BLOCK READ/WRITE
MASTER #3	SLAVE #3	0x18 – 0x1F	SINGLE READ/WRITE

Source code for this example can be found in the [WISHBONE Public Domain Library for VHDL](#) (in the EXAMPLES folder). The library also has detailed documentation for the library

modules, including detailed circuit descriptions and timing diagrams. The reader is encouraged to review and experiment with all of these files.

This system was synthesized and routed on two styles of Xilinx¹⁵ FPGA: the Spartan 2 and the Virtex 2. For benchmarking purposes the memories were altered so that they used Xilinx distributed RAMs instead of the register RAMs in MEM0002a. A memory interface for the Xilinx RAMs can be found in MEM0001a, which is substituted for MEM0002a.

It should be noted that the Xilinx distributed RAMs are quite efficient on the WISHBONE interface. As can be seen in the source code, only a single ‘AND’ gate was needed to interface the RAM to WISHBONE.

The system for the Xilinx Spartan 2 was synthesized and operated on a Silicore evaluation board. In order to verify that the system actually does run correctly, an HP54620a logic analyzer was connected to test pins on the board, and some of the signals were viewed. Figure A-27 shows the trace. Address lines, data write lines and several control signals are shown.

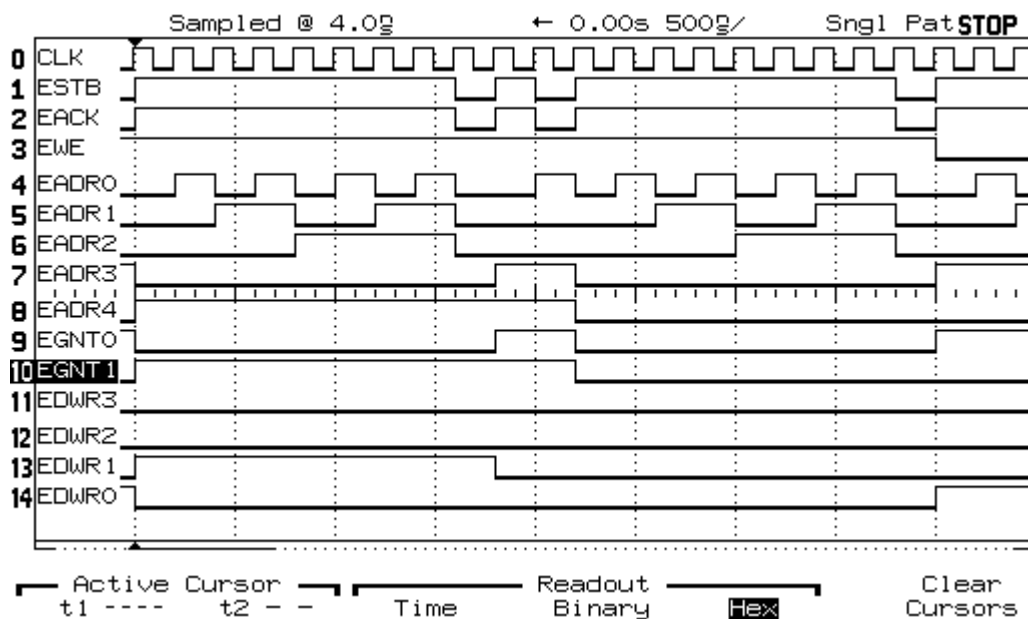


Figure A-27. Logic analyzer trace on the Spartan 2 evaluation board¹⁶.

Table A-87 shows the speed of the system after synthesis and routing. The Spartan 2 benchmarked at about 220 Mbyte/sec, and was tested in hardware (HW TEST). The Virtex 2 part was only synthesized and routed, and showed a maximum speed of about 404 Mbyte/sec (SW TEST).

¹⁵ Xilinx is a registered trademark of Xilinx, Inc.

¹⁶ The logic analyzer samples at 500 Mhz, so the SoC was slowed down to make the traces look better. This trace was taken with a SoC clock speed of 5 MHz.

Table A- 87 . 32-bit Shared Bus Interconnection Benchmark Results							
MFG & Type	Part Number	HW TEST	SW TEST	Size	Timing Constraint (MIN)	Actual Speed (MAX)	Data Transfer Rate (MAX)
Xilinx Spartan 2 (FPGA)	XC2S50-5-PQ208C	√		356 SLICE	55 MHz	55MHz	220 Mbyte/sec
Xilinx Virtex 2 (FPGA)	XC2V250-5-FG256C		√	355 SLICE	99 MHz	101 MHz	404 Mbyte/sec

Notes:

VHDL synthesis tool: Altium Accolade PeakFPGA 5.30a

Router: Xilinx Alliance 3.3.06I_V2_SE2

Hardware evaluation board: Silicore 170101-00 Rev 1.0

Listed size was reported by the router.

Spartan 2 test used '-5' speed grade part (slower than the faster '-6' part).

A.10.7 Other Benchmarks

Other benchmarks have also been run on the shared bus interconnection. At present, the fastest implementation of the 4 MASTER / 4 SLAVE shared bus system has been with 64-bit interfaces. This has resulted in data transfer rates in excess of 800 Mbyte/sec.

A.110 References

Di Giacomo, Joseph. Digital Bus Handbook. McGraw-Hill 1990. ISBN 0-07-016923-3.

Peterson, Wade D. The VMEbus Handbook, 4th Edition. VITA 1997. ISBN 1-885731-08-6

INDEX

- 0x (prefix), 17
- ACK_I signal, 36, 44
- ACK_O signal, 38, 46
- active high logic state, 17
- active low logic state, 18
- address tag. *See* tags
- addressing
 - memory mapped, 23
 - partial, 24
- ADR_I() signal array, 38
- ADR_O() signal array, 36
- arbiter, 118
- ASIC, 18
- asserted signal, 18
- big endian. *See* endian
- BLOCK cycle, 53, 88
- BLOCK READ cycle, 54
- BLOCK WRITE cycle, 57
- bus cycle, defined, 18
- bus cycles
 - BLOCK, 53, 88
 - BLOCK READ, 54
 - BLOCK WRITE, 57
 - READ/WRITE, 48
 - RMW, 60, 88
 - SINGLE, 86
 - SINGLE READ, 49
 - SINGLE WRITE, 51
- bus interface, 18
 - general operation, 41
 - handshaking protocol, 43
 - reset operation, 41
- bus, defined, 18
- BYTE(N), 63
- CLK_I signal, 35
- CLK_O signal, 34
- clock edge number, 16
- clock generator
 - gated, 72
 - variable, 29, 72
- clock transition, 16
- crossbar switch, 12, 18, 82
- CYC_I signal, 38
- CYC_O signal, 36, 46
- cycle tag. *See* tags
- DAT_I() signal array, 35, 37, 39
- DAT_O() signal array, 35, 37, 39
- data flow interconnection, 19, 80
- data organization, 20, 63
- data tag. *See* tags
- DMA, 20
- documentation standard, 11, 31
- DWORD(N), 63
- endian, 12, 20, 63
- endian conversion, 89
- ERR_I signal, 37, 44
- ERR_O signal, 39, 46
- FASM
 - ROM, 99
 - synchronous RAM, 98
- FIFO, 20
- firm core, 20
- fixed interconnection, 21
- fixed timing specification, 21
- foundry. *See* silicon foundry
- FPGA, 21
- full address decoding, 21
- gated clock, 21
- gated clock generator, 72
- glue logic, 11, 21
- granularity, 21, 63
- handshaking protocol, 43
- handshaking, protocol, 84
- hard core, 21
- Hardware Description Language (HDL), 23
- I/O routing, 13
- INTERCON, 80
- interconnection
 - crossbar switch, 82
 - data flow, 80
 - data-flow, 19
 - fixed, 21
 - multiplexor, 23
 - off-chip, 13, 24
 - point-to-point, 26, 80
 - shared bus, 26, 81
 - switched fabric, 12, 27
 - three-state, 29
 - variable, 29
- interface, defined, 23
- IP core, 23, 77
- little endian. *See* endian
- logic minimization. *See* minimization
- logic state
 - active high, 17
 - active low, 18
 - signal name, 34
- logo. *See* WISHBONE logo
- mask set, 23
- MASTER, 23
- MASTER signals, 36
- MASTER/SLAVE topology, 12, 77

- memory mapped, 23
- minimization, 23
- module, 23
- multiplexed bus, 108
- multiplexor interconnection, 23
- negated signal, 24
- OBSERVATION, 14
- off-chip I/O routing, 13
- off-chip interconnection, 24
- operand size, 24
- parametric core generator, 24
- partial address decoding, 24, 115
- PCI, 24
- PERMISSION, 14
- phase (bus cycle), 25, 88
- pipelining, 80
- point-to-point interconnection, 26, 80
- port size, 26
- portable interface, 10
- ports
 - 16-bit, 69
 - 32-bit, 68
 - 64-bit, 66
 - 8-bit, 70
- power-up reset, 43
- QWORD(N), 63
- RAM
 - FASM compatible, 98
- RECOMMENDATION, 13
- reset
 - asynchronous generation, 43
 - operation, 41
- .
- RMW cycle, 60, 88
- router, 26
- RST_I signal, 35
- RST_O signal, 35
- RTL design methodology, 26
- RTY_I signal, 37, 44
- RTY_O signal, 39, 46
- RULE, 13
- SEL_I(7..0) signal array, 39
- SEL_O(7..0) signal array, 37
- shared bus example, 108
- shared bus interconnection, 26, 81
- signal
 - asserted state, 18
 - description, 34, 84
 - input, 16, 17
 - logic levels, 10
 - name. *See* signal name
 - naming conventions, 16, 17, 33
 - negated, 24
 - non-WISHBONE, 33, 34
 - output, 16, 17
- silicon foundry, 27
- SINGLE READ Cycle, 49
- SINGLE WRITE cycle, 51
- SLAVE, 27
- SoC, 27
- soft core, 27
- STB_I signal, 39
- STB_O signal, 37, 46
- structured design, 10, 27
- SUGGESTION, 14
- switched fabric interconnection, 12, 27
- synchronous protocol, 12
- SYSCON, 21, 27, 34
- system controller. *See* SYSCON
- System-on-Chip (SoC), 28
- TAG TYPE. *See* tags
 - , 34, 36, 46
 - , 34, 36, 46
- tags
 - address tag, 18, 38, 40
 - cycle tag, 19, 38, 40
 - data tag, 20, 35
 - defined, 28
 - examples of, 103
 - rules for assigning, 47
 - TAG TYPE, 17, 28, 32, 46, 103
 - TGA_I() TAG TYPE, 40
 - TGA_O() TAG TYPE, 38
 - TGC_I() TAG TYPE, 40
 - TGC_O() TAG TYPE, 38
 - TGD_I() TAG TYPE, 35
 - TGD_O() TAG TYPE, 35
- target device, 28
- three-state bus interconnection, 29, 109
- timing delay, 45
- timing diagrams, 14
- timing specification, 29, 71, 74
- transition path, 11
- variable address path, 11
- variable clock generator, 29, 72
- variable data path, 11
- variable interconnection, 29, 79, 80
- variable timing specification, 29, 79
- Verilog, 29
- VHDL, 29
- VMEbus, 30
- WE_I signal, 40
- WE_O signal, 38
- WISHBONE
 - copyright release, 3
 - DATASHEET, 30, 31, 32, 33, 46
 - disclaimer, 3
 - documentation standard, 31
 - logo, 3, 17, 30
 - revision history, 4
 - , 31
 - signal, 30

steward, 3
WORD(N), 63
wrapper, 30

WSM (wait state MASTER), 16
WSS (wait state SLAVE), 16

Ω