



PCI:SBUS Comparison

Solaris™ 2.5.1 Hardware: 8/97

Sun Microsystems
2550 Garcia Avenue
Mountain View, CA 94043
U.S.A. 415-960-1300

Part No. 805-1628-10
August 1997, Revision A

Copyright 1997 Sun Microsystems, Inc., 2550 Garcia Avenue, Mountain View, California 94043-1100 U.S.A.
All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Portions of this product may be derived from the UNIX[®] system, licensed from Novell, Inc., and from the Berkeley 4.3 BSD system, licensed from the University of California. UNIX is a registered trademark in the United States and other countries and is exclusively licensed by X/Open Company Ltd. Third-party software, including font technology in this product, is protected by copyright and licensed from Sun's suppliers.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

Sun, Sun Microsystems, the Sun logo, Solaris, NFS, SunATM, SunFastEthernet, RSM, Online Backup, Online: DiskSuite, Solstice DiskSuite, X11/NeWS, JumpStart, Netra, SunATM, SunFDDI, Ultra, Enterprise, RSM Array 2000, and Sun-4 are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. SunDoc is a service mark of Sun Microsystems, Inc. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the United States and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK[®] and Sun[™] Graphical User Interfaces were developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

Prestoserve is a trademark of Legato Systems, Inc

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.



Please
Recycle



Adobe PostScript

Copyright 1997 Sun Microsystems Inc., 2550 Garcia Avenue, Mountain View, Californie 94043-1100, U.S.A. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou de sa documentation associée ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Des parties de ce produit pourront être dérivées du système UNIX® licencié par Novell, Inc. et du système Berkeley 4.3 BSD licencié par l'Université de Californie. UNIX est une marque enregistrée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company Ltd. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Sun, Sun Microsystems, le logo Sun, Solaris, NFS, SunFDDI, Sun ATM, SunFastEthernet, RSM, Online Backup, Online: DiskSuite, Solstice DiskSuite, X11/NeWS, JumpStart, Netra, Ultra, Enterprise, Sun RSM Array 2000, SunATM, et Sun-4 sont des marques déposées ou enregistrées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC, utilisées sous licence, sont des marques déposées ou enregistrées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

Les interfaces d'utilisation graphique OPEN LOOK® et Sun™ ont été développées par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant aussi les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

Prestoserve est une marque de Legato Systems, Inc.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" SANS GARANTIE D'AUCUNE SORTE, NI EXPRESSE NI IMPLICITE, Y COMPRIS, ET SANS QUE CETTE LISTE NE SOIT LIMITATIVE, DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DES PRODUITS A RÉPONDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ILS NE SOIENT PAS CONTREFAISANTS DE PRODUITS DE TIERS.

Contents

Preface ix

Related Documents ix

Ordering Sun Documents x

Sun Welcomes Your Comments x

1. Feature Comparison 1

Features Tables 1

Bus Transaction Participants 4

 SBus Controller 5

 PCI Arbiter 5

 SBus Master 5

 PCI Master 5

 SBus Slave/Target 6

 PCI Target 6

Protocol 6

 SBus Basic Cycles 6

 PCI Basic Cycles 7

 Bus Arbitration 7

Latency 8

PCI Bridge Information	9
Improving Performance with Cache Line Size	9
PCI-PCI Bridge Device Supported	9
Index	13

Tables

TABLE 1	Electrical Features Comparison	1
TABLE 2	Firmware Features Comparison	1
TABLE 3	Hardware Features Comparison	2
TABLE 4	Software Features Comparison	3
TABLE 5	Normal Transaction Cycle	4
TABLE 6	PCI Bus Access Latency Components	9
TABLE 7	Sun Host/PCI Bridge Characteristics	10

Preface

The PCI:SBus Comparison document describes feature differences between the PCI interconnect bus and the SBus interconnect bus. This document is intended to help developers who are in transition from developing for SBus to developing for PCI bus. In addition, this document is intended to help field engineers who are assisting these developers.

Related Documents

The following documentation contains the topics related to the information discussed in this *PCI:SBus Comparison* manual.

Table P-1 Related Documents

Subject	Title	Part Number
PCI	<i>PCI Local Bus Specification, Revision 2.1</i>	802-2387-02
	<i>PCI System Architecture</i> , by Tom Shanley and Don Anderson, MindShare, Inc.	
	<i>PCI--A Developer Overview</i> , December, 1996 (technical brief)	
	Sun Microsystems™ PCI migration Web page: http://shasta.corp.sun.com/Rte66/	
	PCI Developer Kit	
	IEEE 1275-1994 Standard for Boot Firmware	
	PCI Bus Binding to IEEE Std1275-1994 Standard for Boot (Initialization Configuration) Firmware, Rev 2..0 at: fhhttp://playground.sun.com/pub/1275/bindings/pci	

Table P-1 Related Documents

Subject	Title	Part Number
PCI and SBus	<i>SBus and PCI Bus: a Comparison</i> (white paper)	
SBus	<i>1496-1993 IEEE Standard for a Chip Module Interconnect Bus: SBus</i> , IEEE Computer Society	SH16659-NYF
	<i>1275.2-1994 IEEE Standard for Boot (Initialization Configuration) Firmware:Bus Supplement for IEEE 1496</i>	SH94236-NYF
	<i>SBus Specification B.0</i> , by Edward H. Frank and Jim Lyle, Sun Microsystems, Inc., December, 1990	800-5922-10
	<i>SBus:Information, Applications, and Experience</i> , James D. Lyle, Springer-Verlag	
	<i>SBus Handbook</i> , by Susan A. Mason, SunSoft Press/Prentice-Hall	

Ordering Sun Documents

The SunDocsSM program provides more than 250 manuals from Sun Microsystems, Inc. If you live in the United States, Canada, Europe, or Japan, you can purchase documentation sets or individual manuals using this program.

For a list of documents and how to order them, see the catalog section of the SunExpressTM Internet site at <http://www.sun.com/sunexpress>.

Sun Welcomes Your Comments

Please use the *Reader Comment Card* that accompanies this document. We are interested in improving our documentation and welcome your comments and suggestions.

If a card is not available, you can email or fax your comments to us. Please include the part number of your document in the subject line of your email or fax message.

- Email: smcc-docs@sun.com
- Fax: SMCC Document Feedback
1-415-786-6443

Feature Comparison

This document contains information and tables to show the comparison of the major PCI features to those of SBus, as well as, some PCI *only* features.

Features Tables

Tables 1 through 5 compare the SBus and PCI electrical, firmware, hardware, software, and normal transaction cycle features.

TABLE 1 Electrical Features Comparison

Feature	SBus	PCI	Difference
Power consumption	P=VI	25W (5V x 5A)	None
Power supply	5V +/- .25V 2A max per connector 12V +/- .75V 30mA max per connector	5V +/- 5% 5A max per connector 3.3V +/- .3V & .6A max per connector 12V +/- 5% 500ma per connector -12V +/- 10% 100ma per connector	None

TABLE 2 Firmware Features Comparison

Feature	SBus	PCI	Difference
FCode required	All devices	Boot and console devices	None

TABLE 3 Hardware Features Comparison

Feature	SBus	PCI	Difference
Address allocation	Static mapping. Slot has known start address and length	Dynamic mapping for I/O and memory address spaces. Slot distinguished only by address of configuration space header.	None
Addressing mode	Use virtual addressing. Require MMU in SBus controller and address translation. 32-bit virtual address for masters, 28-bit physical address space per slave.	Normal address mode	None
Address spaces	Standard chunk of memory-mapped space (28 bits/slot)	Three physical address spaces: memory, I/O, configuration.	None
Auto-configuration			
Burst: Transfer mode size	Available. Up to 64 bytes: size declared in advance in SIZ{2:0}.	Yes Variable size, no limit; determined by PCI device and PCI bridge.	None
Bus bandwidth: (data transfer rate)	25MHz/32-bit, max. 100MBytes/sec for 32-bit. 200 MBytes/sec for 64-bit	33MHz (rev 2.0) or 66 MHz for PCI/66, 32/64bit, maximum rate from 132MByte/second (33MHz/32-bit) to 528MByte/second (66 MHz/64-bit).	None
Bus parking	No	Yes. An arbiter may grant the busses to a master when the bus is idle and masters are not generating requests for the bus. If the master that the bus is parked on subsequently issues a request from the bus, it can immediately access it.	None
Bus width			x 32-to 64-bit
Clock frequency/ bus speed	Max CLK 25MHz CLK to out: 22nS Input Set-up: 15nS	33MHz max CLK CLK to out:22ns Input Set-up: 7 ns	Synchronous; signals referenced to rising clock edge
Connectors	One type of connector	Either 5V or 3.3V 32-and 64-bit connectors; the 64-bit being an extension of the 32-bit	None
Data path	32-bit, except 64-bit for extended transfers	32/64-bit	None

TABLE 3 Hardware Features Comparison *(Continued)*

Feature	SBus	PCI	Difference
Dynamic Bus Sizing	Yes. Enables slave to control data width it accepts during non-burst transfers	No, but same function is performed by Byte Enable setting during Data Phase	None
Form factors	Single form factor per slot (double and triple-wide can share slots)	Three form factors: short, long, and variable-height short	None
Interrupts	Seven levels	Four levels (INTA-D)	None
Max devices per bus	Usually limited by electrical loading	32 slots	None
Max functions per device	1	8	None
Max masters per bus	8	32	None
Pin count (connector)	96	94 pins for both 32-bit and 64-bit connectors	None

TABLE 4 Software Features Comparison

Feature	SBus	PCI	Difference
Parity	Optional on data and virtual address transfers if parity generation and checking is implemented on controller and installed masters and slaves.	Default during address and data phases; must be performed by all PCI-compliant devices.	None

TABLE 5 Normal Transaction Cycle

Address decode	Slave selected by decode from controller	Each target performs full 32-bit decoding and drives DEVSEL# if selected
Burst size	32 words	No limit
Bus driving and turnaround	None	A turnaround cycle is necessary for signals driven by more than one agent, to enable contention-avoidance when bus driving agents transfer a signal.
Byte ordering and placement	Byte-lane swapping	DWORD swapping puts bytes in correct lane base on byte address
Byte ordering	Big-endian	Little-endian
Cycle participants	Controller, Master, Slave	PCI Agents: Master/Initiator, Target
Cycle terminology and composition	Transfer = 1. Arbitration Phase 2. Translation Phase 3. Extended Transfer Information Phase 4. Transfer Phase (also called Slave Cycle)	Transfer = 1 or more clock cycles Phase = 1 or more read or write transfers, including (hidden) arbitration while current initiator is performing a data transfer Read or write transaction = 1 address phase + 1 or more data phases

Bus Transaction Participants

The bus transaction participants include the SBus controller, PCI arbiter, SBus and PCI masters, and the SBus and PCI targets.

SBus Controller

The SBus controller arbitrates contention between bus masters during the arbitration phase.

The SBus controller also performs the following:

- Address strobe
- Bus arbitration
- Bus time-outs
- Data transfer count
- SBus system clock
- Slave selects
- Virtual address translation and page size restrictions

PCI Arbiter

The PCI arbiter has no true controller but performs functions equivalent to the controller to arbitrate between bus masters. Bus masters can terminate transactions on completion or time-out; targets can also terminate transactions. The arbiter is typically integrated into the host/PCI or the PCI/expansion bus bridge chip.

SBus Master

The SBus master controls operations that produce error-free data read and write tasks between itself and an SBus slave.

The master-initiated transaction types are:

- Arbitration
- Translation
- Default transfer
- Extended transfer

PCI Master

The PCI master becomes an initiator when it has arbitrated for and gained access to the PCI bus. The initiator starts transfers but can also abort, terminate, and time out.

The master also does the following:

- Starts the Address Phase
- Inserts wait states during data transfer
- Terminates transactions

SBus Slave/Target

The SBus slave/target performs as follows:

- Monitors the SBus to determine if a master is requesting a data transfer
- Provides the data requested to the SBus master during the transfer phase
- Does not perform address decoding; the master does this
- Participates in burst transfers, dynamic bus-sizing, extended transfer phases
- Controls the data transfer rate by controlling the data acknowledgment rate
- Terminates transactions

PCI Target

The PCI target performs the following functions:

- Determines that it is the target of a transaction
- Receives a data object from the initiator
- Decodes addresses
- Participates in special cycles
- Inserts wait states during data transfer (controls data transfer rate)
- Terminates transactions

Protocol

Protocol contains the SBus and PCI basic transaction cycles for each bus, and bus arbitration.

SBus Basic Cycles

The following are definitions of the SBus basic transaction cycle for each bus.

Arbitration Phase—During this phase, masters request bus access. When there is contention between masters, the controller determines which master performs the next transfer. After arbitration, the controller is responsible for monitoring the transfer.

Translation Phase—The master and controller participate in conversion of the virtual addresses to physical address and selection signals that are used by the master and the slaves.

Extended Transfer Information Phase—This phase is used only for the cycles requested by the master that have a SIZ[2:0] value of Extended Transfer.

Transfer Phase—During the Transfer Phase (slave cycle), data is moved to or from the slave.

Dynamic Bus-sizing—This feature enables a master to communicate more easily with slaves of varying widths.

Burst Transfers—SBus burst transfer protocol is the same as that for SBus basic transactions, except that multiple words are transferred. Neither dynamic bus-sizing or varying-width slaves are allowed.

Extended Transfer Mode—These 64-bit transfers enable increased performance. Up to twice the bandwidth is possible, as the data path is twice as wide.

PCI Basic Cycles

The following are definitions of the PCI basic transaction cycle for each bus.

Address Phase—Every PCI transaction begins with this phase, which includes concurrent hidden arbitration.

The initiator identifies the target device and transaction type.

Data Phase—At the end of the address phase, the address/data bus transfers data in one or more data phases. The clock immediately following the Address Phase begins the data phase.

During the data phase, a data object is transferred between initiator and target during every rising edge of the PCI bus clock.

Bus Idle State – When the last data transfer has finished, the initiator returns the bus to the idle state, which is the condition of having no transaction in progress on the bus.

Bus Arbitration

Bus arbitration by the SBus enables concurrent arbitration and transfer execution. The PCI arbitration is access-based, not time-slot-based, to minimize access latency.

SBus

When more than one SBus master requests bus access, the controller grants access to one of the requesters. Enabling arbitration concurrent while the master makes a data transfer referred to as hidden arbitration.

PCI

In PCI arbitration, a bus master must arbitrate for each bus access.

The PCI specification does not define the PCI bus arbitration scheme. The arbiter may use any scheme, but the 2.1 specification requires that the arbiter implement a fairness algorithm to avoid deadlocks.

Latency

SBus and PCI bus latency characteristics are similar. Both are low-latency, high-throughput buses. The number of wait states that targets and masters can add to a transaction is limited. Also, masters have programmable timers that limit their times on the bus during heavy-traffic periods. The limits, plus bus arbitration order, ensure that bus acquisition latencies can accurately be predicted for any bus master.

PCI and SBus latencies are functions of:

- The number of bus masters
- The arbitration method and its overhead time
- The length of an SBus translation phase or a PCI data phase
- The time the slave or target takes to finish the transfer
- The occurrence of retries and errors.

Bus access latency is the elapsed time from the moment that a bus master requests bus access until it finishes the first data transfer of the transaction. Table 6 describes the bus access latency components. PCI Bridge Information

TABLE 6 PCI Bus Access Latency Components

Component	Description
Bus access latency	The elapsed time from the moment a master requests bus access until it finishes the transactions's first data transfer
Arbitration latency	The time that the master issues a request to the time when the arbiter asserts the master's grant
Bus acquisition latency	The elapsed time that the requesting master receives the grant until the current master surrenders the bus
Target latency	The elapsed time that transaction starts until the currently-addressed target is ready to finish the transactions's first data transfer

PCI Bridge Information

PCI bridge device PCI characteristics information is contained in this section. See Table 7 for bus type A and bus type B specific breakdowns

Improving Performance with Cache Line Size

The cacheline size for Sun™ SPARC™ platforms is 64 bytes. PCI devices for SPARC platforms should use 64 bytes for best performance.

PCI-PCI Bridge Device Supported

For PCI-PCI bridge device support, refer to "PCI Developer's Frequently-Asked Questions" at <http://shasta.corp.sun.com/Rte66/>.

TABLE 7 Sun Host/PCI Bridge Characteristics

Characteristics	Bus A	Bus B
Data Transfer Width	64-bit	64-bit
Clock Frequency	33/66 MHz capable	33 MHz capable
Burst Size	64 bytes	64 bytes
Number of Read Buffers	One 64-byte for DMA	One 64-byte for DMA
Number of Write Buffers	Two 64-byte for DMA One 64-byte for PIO	Two 64-byte for DM One 64-byte for PIO
Dual Address Cycles	Bypass DMA only	Bypass DMA only
Fast back-to-back device	In target mode only	In target mode only
Byte Swapping	Yes for DMA	Yes for DMA
Interrupt Latency	6 Cycles inside the IDU	6 Cycles inside the IDU
Cache Line Size	64 bytes	64 bytes
Number of cache lines	16	16
Disconnected on Cache Line	Target mode (DMA) only	Target mode (DMA) only
Configuration Mechanism (per Section 3.7.4 of PCI 2.1 Specification)	Configuration Mechanism #2	Configuration Mechanism #2
Configuration Space	256 byte; starting at physical address 1FE.0101.0000	256 bytes; starting at physical address 1FE.0100.0000
I/O Space	8K; at physical address 1FE.0200.0000	8K; at physical address 1FE.0201.0000
Memory Space	2G; at physical address 1FF.0000.0000	2G; at physical address 1FF.8000.0000
Configuration Cycles	Master mode only	Master mode only
Special Cycle	Master mode only	Master mode only
Arbitrary byte enables	Consistent DMA only	Consistent DMA only
Peer-to-peer DMA	On a single segment	On a single segment
Interrupt	Four interrupt lines shared among PCI devices	Four interrupt lines shared among PCI devices
IOMMU page size	8K and 64K.; only 8K page size is used in the STC.	8K and 64K.; only 8K page size is used in the STC.

TABLE 7 Sun Host/PCI Bridge Characteristics *(Continued)*

Characteristics	Bus A	Bus B
DVMA addressing space (set by pci nexus driver)	64 Mbyte in Solaris 2.5.1; may be changed in later release of Solaris	64 Mbyte in Solaris 2.5.1; may be changed in later Solaris releases
PIO read size	1, 2, 4, 8, 16, 64 bytes for memory cycles 1,2, and 4 bytes for I/O or configuration cycles	1, 2, 4, 8, 16, 64bytes for memory cycles 1,2, and 4 bytes for I/O or configuration cycles
PIO write size	0-16 arbitrary byte enable and 64-byte aligned for memory cycles 0-4 arbitrary byte enables for I/O or configuration cycles	0-16 arbitrary byte enable and 64-byte aligned for memory cycles 0-4 arbitrary byte enables for I/O or configuration cycles
Cache-line wrap addressing mode	Not supported	Not supported
Local (on-PCI) cache	Not supported	Not supported
Exclusive access to main memory	LOCK# signal not connected	LOCK# signal not connected
Address/data stepping	Not supported	Not supported
DOS compatibility hole	Not supported	Not supported
External arbiter	Not supported	Not supported
Subtractive decode	Not supported	Not supported

Index

A

address
 phase, 7
 strobe, 5
arbitration, 5, 7, 8
arbitration fairness algorithm (PCI), 8
arbitration phase (SBus), 6

B

basic transaction cycle, 6, 7
burst transfer
 protocol, 7
burst transfers, 7
bus time-outs, 5
bus transaction participants, 1

C

controller (SBus), 5

D

data
 acknowledgment rate, 6
 phase, 7
data transfer
 count, 5
 rate, 6

default transfer, 5
dynamic bus-sizing, 7

E

extended transfer, 5
 information phase, 7
 mode, 7

I

initiator, 5, 7

L

latency, 8

M

master (bus), 5
master (PCI), 5
master-initiated transaction types, 5

P

PCI arbiter, 5
power consumption, 1, 2, 3
power dissipation, 1, 2

protocol, 6

S

SBus master, 5

SBus system clock, 5

SIZ[2:0], 7

slave, 6

 cycle, 7

 selects, 5

T

target, 6

target (PCI), 6

transaction termination (PCI), 5

transfer phase, 7

transfers, 64-bit, 7

translation, 5

translation phase, 6

V

virtual address translation, 5