EK-TCAAB-SP-005

TURBOchannel

System Parameters

On-line version.

Previous versions of this document are obsolete and should be discarded. This document supersedes all previous versions.

digital™

September 1991

Digital's TRI/ADD Program provides technical and marketing support worldwide to third-party vendors using the SCSI, TURBOchannel, VME, and Futurebus+ interconnects to develop add-on products for open systems.

To receive free technical support and notice of new TURBOchannel documentation, contact Digital's TRI/ADD Program about free membership at the numbers below.

Digital Equipment Corporation
TRI/ADD Program
100 Hamilton Avenue
Palo Alto, CA, U.S.A. 94301

FAX 1.415.853.0155
Internet address: triadd@decwrl.dec.com

U.S.A/Canada	**Australia**	**France**	**Italy**
1.800.678.OPEN	0014.800.125.388	05.90.2874	1678.19087
Japan	**U.K.**	**Germany**	
0031.12.2363	0800.89.2610	0130.81.1974	

TURBOchannel Technology Transfer Agreement

Grant of Right to Use TURBOchannel Technology

In exchange for your agreeing to the warranty disclaimer and liability limitation stipulated in this Technology Transfer Agreement, Digital Equipment Corporation (Digital) grants at no cost to you a royalty-free nonexclusive license to use TURBOchannel technology (as specified in the *TURBOchannel Specifications*) to design and develop any kind of option board, computer system, or application-specific integrated circuit (ASIC). This Agreement does not grant you any other rights in Digital's patents, copyrights, trade secrets, trademarks, or licenses to TURBOchannel technology. The purchase cost of the TURBOchannel kit is basically the cost to reproduce the materials.

Warranty Disclaimer

The TURBOchannel technology is transferred "as is." Digital expressly disclaims all implied warranties including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Digital does not warrant, guarantee, or make any representations regarding the use of or the results of the use of the specification and related documents in terms of correctness, accuracy, or reliability. Digital believes the documentation is accurate; however, you must assume the risk as to the results and performance of any product you develop that is based on the TURBOchannel technology.

Limits of Liability

You agree that Digital shall not be liable to you under this Agreement for any damages, including without limitation any lost profits or lost savings, or any consequential, incidental, or punitive damages arising out of the use or inability to use the TURBOchannel Hardware Specification and related documents, or for any claim by another party. Your exclusive remedy under this Agreement shall be the furnishing by Digital of the technical support provided herein. You agree to hold Digital harmless for all claims and damages arising from any third party as a result of their use of or inability to use any product you develop based on TURBOchannel technology.

Important Changes in TURBOchannel System Parameters

This version of TURBOchannel System Parameters includes all changes made since the original document was released. Revision bars mark technical changes added since Version 2B of the specification.

Changes in Version 2C

- **Implementation Notes**: This section has been added.
- **DECstation 5000 Model 100 System Parameters**: This chapter has been added.

Changes in Version 2B

- None.

Changes in Version 2A

- **I/O Transaction**: This section was added to explain operation of the processor write buffer implementation of read-over-write priority retiring of outstanding processor transactions. Option device drivers must wait for the write buffer to drain via the kernel **wbFlush( )** routine.

Changes in Version 2.0

- **Transaction Parameters**: The system type has been added.

1

DECstation/DECsystem 5000 Model 200 System Parameters

Transaction Parameters

System type	0x82020220
Clock frequency	25 MHz
Slot size	4 Mbytes
Number of integral slots	3
Number of expansion slots	3
I/O timeout	255 ~*sel* asserted cycles
DMA arbitration priority	Fixed (2 > 1 > 0)
DMA address range	480 Mbytes
Maximum DMA burst	128 words
Parity	Not supported

All other parameters are defined in the *TURBOchannel Hardware Specification*.

Physical Address Space

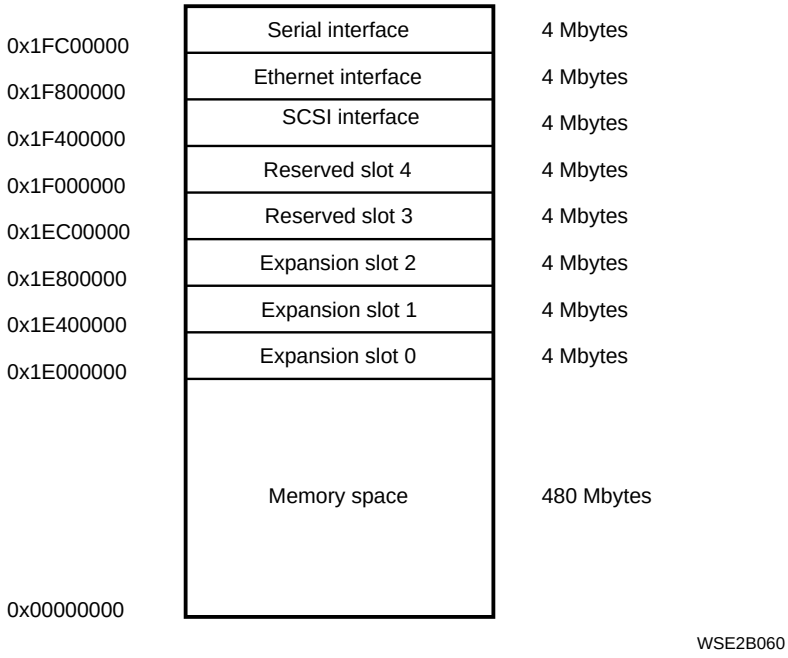

DMA transactions can access only the 480-megabyte memory space region. When viewed from the rear of the system enclosure, expansion slot 0 is on the left, expansion slot 1 is in the middle, and expansion slot 2 is on the right.

Transaction Performance

I/O read minimum R3000 cycles	8 (load stall/fixup)
I/O write minimum R3000 cycles	3 (sustained)
DMA arbitration minimum cycles	0
DMA read data latency cycles	9
DMA write data trailer cycles	6
Memory refresh overhead	5/195 cycles

The processor subsystem implements six-stage write-buffer logic that accepts R3000 store instructions without stall cycles when the buffer is not full. However, the TURBOchannel can sustain only I/O writes at a 3-cycle rate. The DMA write trailer cycles refers to the number of cycles that the memory system remains busy after DMA writes complete on the TURBOchannel. Peak DMA bandwidth (100 megabytes per second) is time multiplexed between expansion slots 2, 1, and 0 in fixed priority order.

I/O Transactions

The processor subsystem write buffer implements read-over-write priority retiring of outstanding processor transactions. That is, read transactions bypass buffered write transactions. Hence, processor I/O reads may issue to options before previously executed store instructions. An options that requires strict ordering of I/O transactions must have its device driver explicitly wait for the write buffer to drain via the kernel **wbFlush()** routine.

The write buffer contains address conflict logic that automatically retries bypassed read transactions if it detects a buffered write to the same eight-word address block. If this occurs on an I/O read transaction, options are issued two I/O read transactions for one processor load instruction. Options that cannot tolerate such behavior must have their device driver explicitly wait for the write buffer to drain via the kernel **wbFlush()** routine.

Because of the processor subsystem write buffer, options that terminate I/O transactions with the *~conflict* signal may be issued other I/O transactions before the conflicted I/O transaction is reissued.

The processor subsystem indefinitely reissues I/O transactions that terminate with the *~conflict* signal asserted. Hence, options should give a normal termination to I/O transactions as soon as possible.

TURBOchannel Interrupts

Each slot interrupt signal is independently enabled through a system control register, and is independently visible through a system control register regardless of the enable logic state. All enabled slot interrupts are merged into a single R3000 hardware interrupt level.

TURBOchannel Resets

Each slot reset signal follows the state of the system reset logic. The system reset logic is activated by the system power supply during power-up and power-down events, or from an operator switch.

Implementation Notes

In regard to DMA transactions, two cases have been identified in which the DECstation/DECsystem 5000 Model 200 platform does not comply with the TURBOchannel Hardware Specification.

1. If a TURBOchannel option asserts *~rReq* within the eight cycles trailing the last assertion of *~wReq* from that same option, the DECstation/DECsystem 5000 Model 200 will use an incorrect address for the DMA read. The DECstation/DECsystem 5000 Model 200 will also return one more word than originally requested. And, if a CPU read request or memory refresh request occurs simultaneously, the system may hang until reset.
2. If a TURBOchannel option asserts *~rReq* in the same cycle in which *~ack* is de-asserted for a previous DMA read from that same option, the DECstation/DECsystem 5000 Model 200 will use an incorrect address for the DMA read. The DECstation/DECsystem 5000 Model 200 will also return one more word than originally requested.

An ECO to a base system dmaCTL PAL on the DECstation/DECsystem 5000 Model 200 is being investigated in an effort to correct number 1 above. A decision is pending on how to deal with number 2.

2

DECstation 5000 Model 100 System Parameters

Transaction Parameters

System type	0x801204B8
Clock frequency	12.5 MHz
CPU frequency	Variable (20 MHz, 25 MHz, and up)
Slot size	64 Mbytes
Number of integral slots	1
Number of expansion slots	3
I/O timeout	127 ~*sel* asserted cycles
DMA arbitration priority	Fixed (2 > 1 > 0)
DMA address range	128 Mbytes
Maximum DMA burst	128 words
Parity	Not supported

All other parameters are as defined in the *TURBOchannel Hardware Specification*.

Physical Address Space

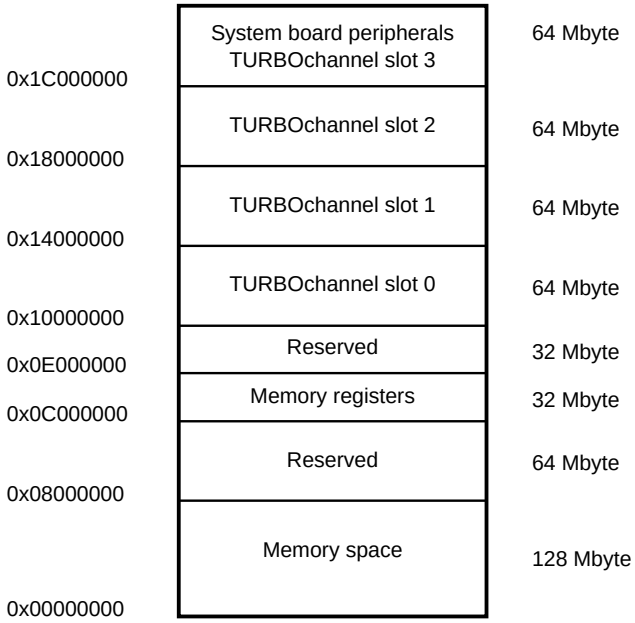

DMA transactions can access only the 128-megabyte memory space region. When viewed from the rear of the system enclosure slot 0 is left most, slot 1 is in the middle, and slot 2 is on the right nearest to the power supply.

Transaction Performance

Since DECstation 5000/1XX CPU daughter cards utilize CPU chips at varying clock rates, it is necessary to evaluate the I/O performance of a system in terms of these clock rates. An outline of how the calculations are made follows.

- I/O reads and writes are placed in a single queue and processed in order.
- The I/O write bandwidth is 12.5 megabytes per second for options that assert *~rdy* concurrent with the data cycle.
- I/O read bandwidth: Assume one TURBOchannel wait state on read.

$$\frac{4000}{(400\ \text{ns}) + 7\ \text{cpu\_cycle time (ns)}} = \text{I/O read bandwidth (Mbytes/sec)}$$

DECstation 5000/120 - 20 MHz(tcyc=50 ns):
4000/(400 ns + 350 ns) = 5.33 Mbytes/sec

DECstation 5000/125 - 25 MHz(tcyc=40 ns):
4000/(400 ns + 280 ns) = 5.88 Mbytes/sec

Note: The actual number of cycles required by the CPU may change with each new daughter card implementation.

DMA Performance

DMA arbitration minimum	1
DMA address cycle	1
DMA read data latency cycles	2
DMA write data trailer cycles	1
Memory refresh overhead	5/193 cycles
DMA peak bandwidth	50 Mbytes/sec

The DECstation 5000 Model 100 is designed around a 12.5 MHz TURBOchannel implemented as a bus. There are six devices or drops on this bus: three option slots (0, 1, 2), the memory subsystem, the I/O subsystem (implemented as TURBOchannel slot 3), and the CPU. The CPU is implemented as a daughter card that runs asynchronous to the rest of the system, which allows upgrades to faster technology.

I/O Transactions

The processor subsystem write buffer may or may not implement a read-over-write priority of outstanding processor transactions. That is, read transactions may bypass buffered write transactions. However, to be safe, the option designer should provide a driver that expects this read bypass to occur. Since processor I/O reads may issue to options before previously executed I/O write instructions, options that require strict ordering of I/O transactions must have their device driver explicitly wait for the write buffer to drain via the kernel **wbFlush()** routine.

The write buffer contains address conflict logic that automatically retries bypassed read transactions if it detects a buffered write to the same eight-word address block. If this occurs on an I/O read transaction, options are issued two I/O read transactions for one processor load instruction. Options that cannot tolerate such behavior must have their device driver explicitly wait for the write buffer to drain via the kernel **wbFlush()** routine.

Because of the processor subsystem write buffer, options that terminate I/O transactions with the *~conflict* signal may be issued other I/O transactions before the conflicting I/O transaction is reissued.

The processor subsystem indefinitely reissues I/O transactions that terminate with the *~conflict* signal asserted. Hence, options should give a normal termination to I/O transactions as soon as possible.

DMA Transactions

As defined in the *TURBOchannel Hardware Specification*, addresses used for DMA transactions are physical addresses. For the DECstation 5000 Model 100, the valid physical address space for system memory is from 0 to 128 megabytes. Thus, DMA byte address bits 27 to 33, which correspond to *ad* bits 30 to 31 and 0 to 4, must be set to 0. If these bits are nonzero, successful completion of the DMA transaction is not guaranteed.

TURBOchannel Interrupts

Each TURBOchannel interrupt is mapped to a specific processor interrupt. Slot 0 interrupts at 0, slot 1 interrupts at 1, and so on. The interrupts are visible in the R3000 cause register.

Error Information

If a TURBOchannel error occurs during an option DMA, the system saves some information about that error. The following table shows bits defined in the memory error register (MER) at 0x0C400000.

[16]	Set when any DMA crosses a 2-Kbyte boundary
[15]	Set when any DMA exceeds 128 words
[11]	Set when a memory parity error occurs on byte 3
[10]	Set when a memory parity error occurs on byte 2
[9]	Set when a memory parity error occurs on byte 1
[8]	Set when a memory parity error occurs on byte 0

Bits not shown above are to be considered reserved and should be ignored. The MER is included here as a debug aid; the MER should never be written.

System Reset

Each slot reset signal follows the state of the system reset logic. The system reset logic is activated by the system power supply during power-up and power-down events.

System Halt

The push button in the back of the system unit generates an interrupt at level 4 that causes a halt. This allows the user to halt the system without losing its state.