M U L T I C L U S T E R Series

Hardware Documetation
&
Software Documentation

Author:

Winfried Mularski

M S C

Mass Storage Controler

Busless Transputer Module
with SCSI and floppy interface

Technical Documentation

Version 1.3 , May 1989

PART I Hardware

1. Block Diagram and General Description

1.1 Introduction

The MSC board is part of the MULTICLUSTER and SUPERCLUSTER series. The MSC board serves as a mass storage controller, which interfaces via SCSI bus and floppy bus to SCSI and loppy disk drive devices. It can operate as a host or as a fileserver/mass storage subsystem in a transputer network.

Multiple MSC boards can be used to drastically increase the I/O bandwidth of a transputer network. This is achieved by connecting to every MSC it's own mass storage devices (mainly winchesters). A MSC can do SCSI bus transfers in parallel with link data transfers over all 4 links without significant performance degradation.

The big memory size of 4 MByte enables implementation of sophisticated buffer algorithms.

The MEGAFRAME modular concept also supports fail save systems: For example two or more MSC's, each with its own winchester of the same type, perform in parallel read and write operations with the same data (n-fold redundancy).

Furtheron in case of a program crash the MSC board can be reset by sending to it a link reset signal over any of the 4 MEGAFRAME links.

Features:

- 32 bit Transputer T800 (opt. T414)
- 4 MByte memory with parity checking (16 MByte, when 4 MBit DRAMs are avaiable)
- 4 RS-422 driven MULTICLUSTER links
- directly interfaces to SCSI bus (ANSI SCSI X3T9.2)
- directly interfaces to floppy bus
- average asynchronous SCSI transfer rate: 1.1 MByte/s
- average synchronous SCSI transfer rate: 1.8 MByte/s
- bidirectional data transfer over all 4 links in parallel with SCSI bus data transfer
- additional onboard 64 pin connector with transputer interface for extension boards
- software package available with complete set of medium level communication procedures (clear, load, unload, read, write etc.); dynamical multisector blocks
- 5 Volt only, low power
- small board size: extended euro card

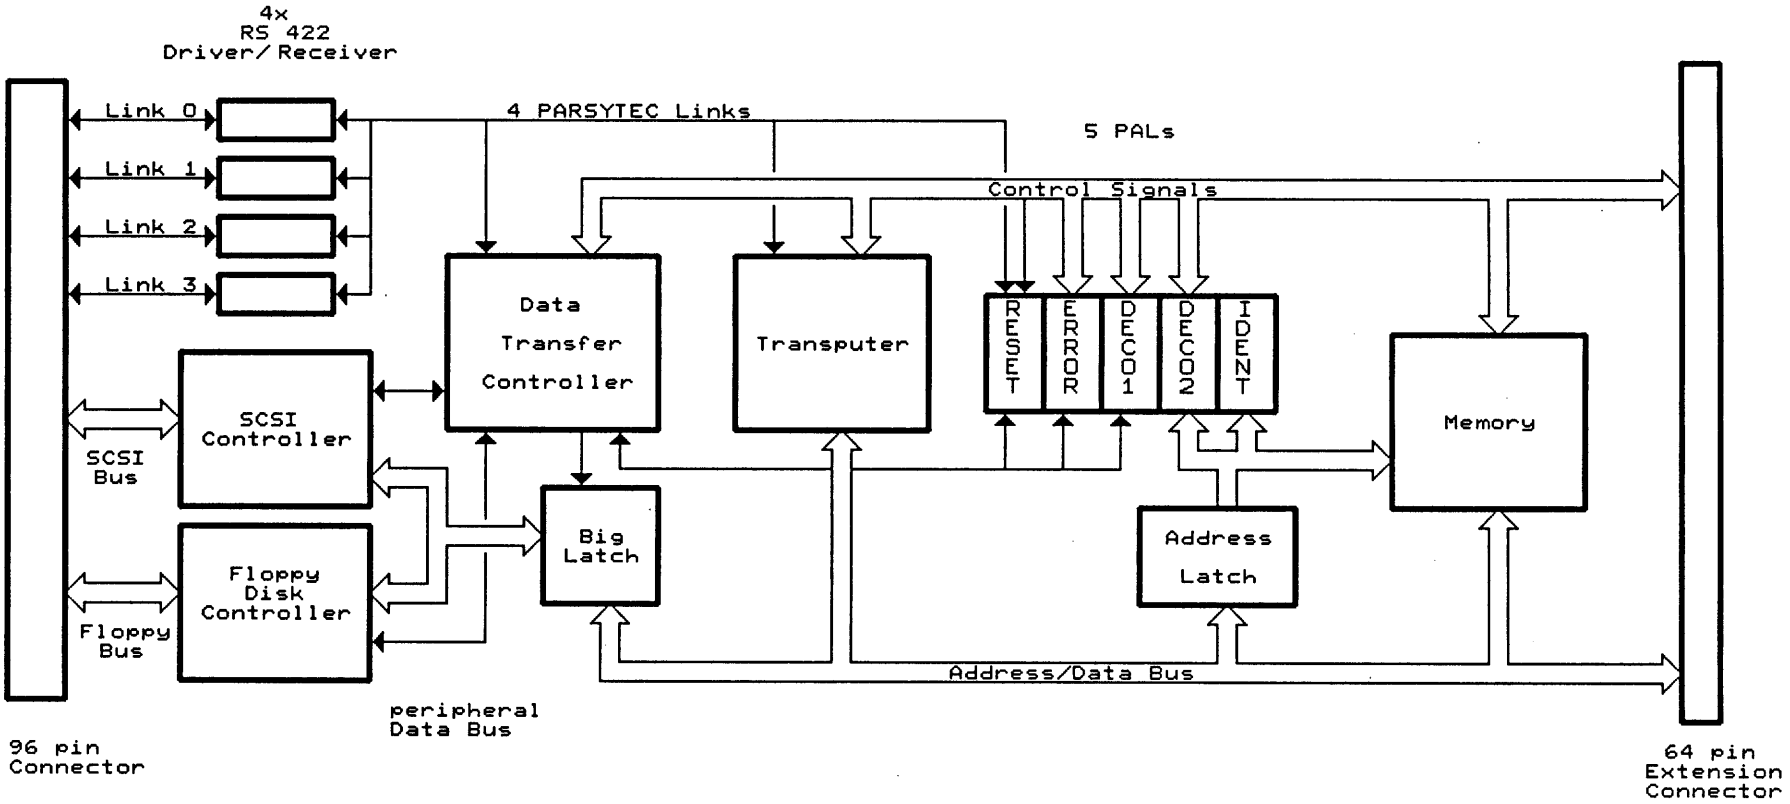

PARSYTEC GmbH		
Size A	Document Number MSC schematic	REV 1.1
Date: February 19, 1988	Sheet of	

1.2 Processor

The MSC board runs with a T800 or T414 Transputer. The T800 is a 32 bit processor with 4 K bytes of on-chip static RAM. It runs at 20 MHz to perform an instruction throughput of 10 MIPS. Four high speed serial links (10 or 20 Mbit/sec) support the communication with other transputers in a network.

Furthermore the T800 has a floating point unit on chip to perform 1.5 MFLOPS/sec.

1.3 Clock

All transputers derive their processor clock from an internal oscillator, which is synchronized by an internal PLL to an external 5 MHz oscillator. The transputer clock speed is defined by jumpers. See Jumper Allocation.

1.4 Links

The four bidirectional serial links of the transputer operate independently of the processing element when transfering data from or to memory by using fast DMA. So the use of the links does only lightly degrade processor performance. The links have a default speed of 10 Mbit/sec and can also operate at 5 and 20 Mbit/sec. The link speeds can be selected by jumpers.

The MSC board has four MULTICLUSTER links. A MULTICLUSTER link consists of four signals: Two signals transfer the serial link data, one for both directions, and two reset signals, also one for both directions. Using these reset signals the transputer can generate a reset for any of its four neighbour transputers. Further on such a link reset does not effect a normal transputer reset but a so called Analyse/Reset. The Analyse/Reset preserves the internal status of the resetted transputer section. This internal status can be analysed by a user written procedure, which must be downloaded into the transputer (see "Error and Analyse").

Link and link reset signals, which leave the board, are driven by RS422 drivers. Link and link reset signals which enter the board are conditioned by RS-422 receivers. This provides a considerably higher noise immunity and longer distances for data transmission (up to 10 meters at 20 MBits/sec, up to 30 meters at 10 or 5 MBits/sec)).

1.5 Memory

The memory is organized as 1M x 32 bit, i.e. 4 MByte, and consists of dynamic RAM's. The 4 bytes in the 32 bit word are parity checked by additional 1M x 4 bit memory. The MSC board can be populated with 16 MByte memory when the next generation dynamic RAM's (4 Mbit DRAMs) are available.

1.6 SCSI/Floppy Controller

The WD33C93 SCSI Controller provides the interface to the SCSI Bus. To speed up SCSI data transfer a multiplexer-buffer-latch, called **BigLatch** (see "The BigLatch"), connects the 8 bit controller data bus to the 32 bit transputer data bus.

The WD37C65 Floppy Controller interfaces to the floppy bus. Both controllers and the BigLatch are under control of the Data Transfer Controller DTC.

1.7 Booting

After a transputer reset there are two ways to boot the transputer. They are selected by a jumper. The first is that the transputer executes the boot code of an EPROM, which resides on an optional extension board. In this case physical memory location #7FFFFFFE is interpreted as an instruction and executed. The second way is to load the boot code over a transputer link: The first incoming data from one of the links will be interpreted as boot code.

2. The SCSI Section

2.1 The SCSI Controller

The MSC board contains the WD33C93 SCSI-bus interface controller from Western Digital Corp.. The internal registers are accessible by the transputer (see "Overview of the OCCAM Adress Space").

The controller works in non multiplexed bus mode. That means, to access an internal register of the controler the transputer must first write the address of the desired register into the controlers address register and then accesses the register.

There are three ways for the transputer to notice an interrupt of the SCSI Controller.

- The interrupt bit of the Auxiliary Status Register of the SCSI Controller can be polled.
- The interrupt line can be polled (see "BigLatch Status Register")
- The interrupt line can activate the EVENTREQ input of the transputer (see "Events" and "Jumper J9").

The reset input is activated at power on, external reset (see "Reset Signals") or by clearing the flip flop PRES (see "Flip flop PRES").

The DBA mode (Direct Buffer Access) of the WD33C93 is implemented, because it delivers the maximum SCSI performance. In this mode the controller actively performs read/write cycles during SCSI data phases.

2.2 The SCSI Interface

The WD33C93 SCSI controller is fully compatible with ANSI SCSI X3T9.2 specifications. The MSC board achieves an average data transfer rate of 1.1 MByte/sec (asynchronous) and 1.8 MByte/sec (synchronous). The chip includes 48 mA drivers for direct connection to the single ended SCSI bus, therefore the SCSI bus can be up to 6 meters in length.

All SCSI signals are terminated with 220/330 ohm on board. The SCSI RST signal is connected to the 96 pin bus connector and cannot be activated by the transputer.

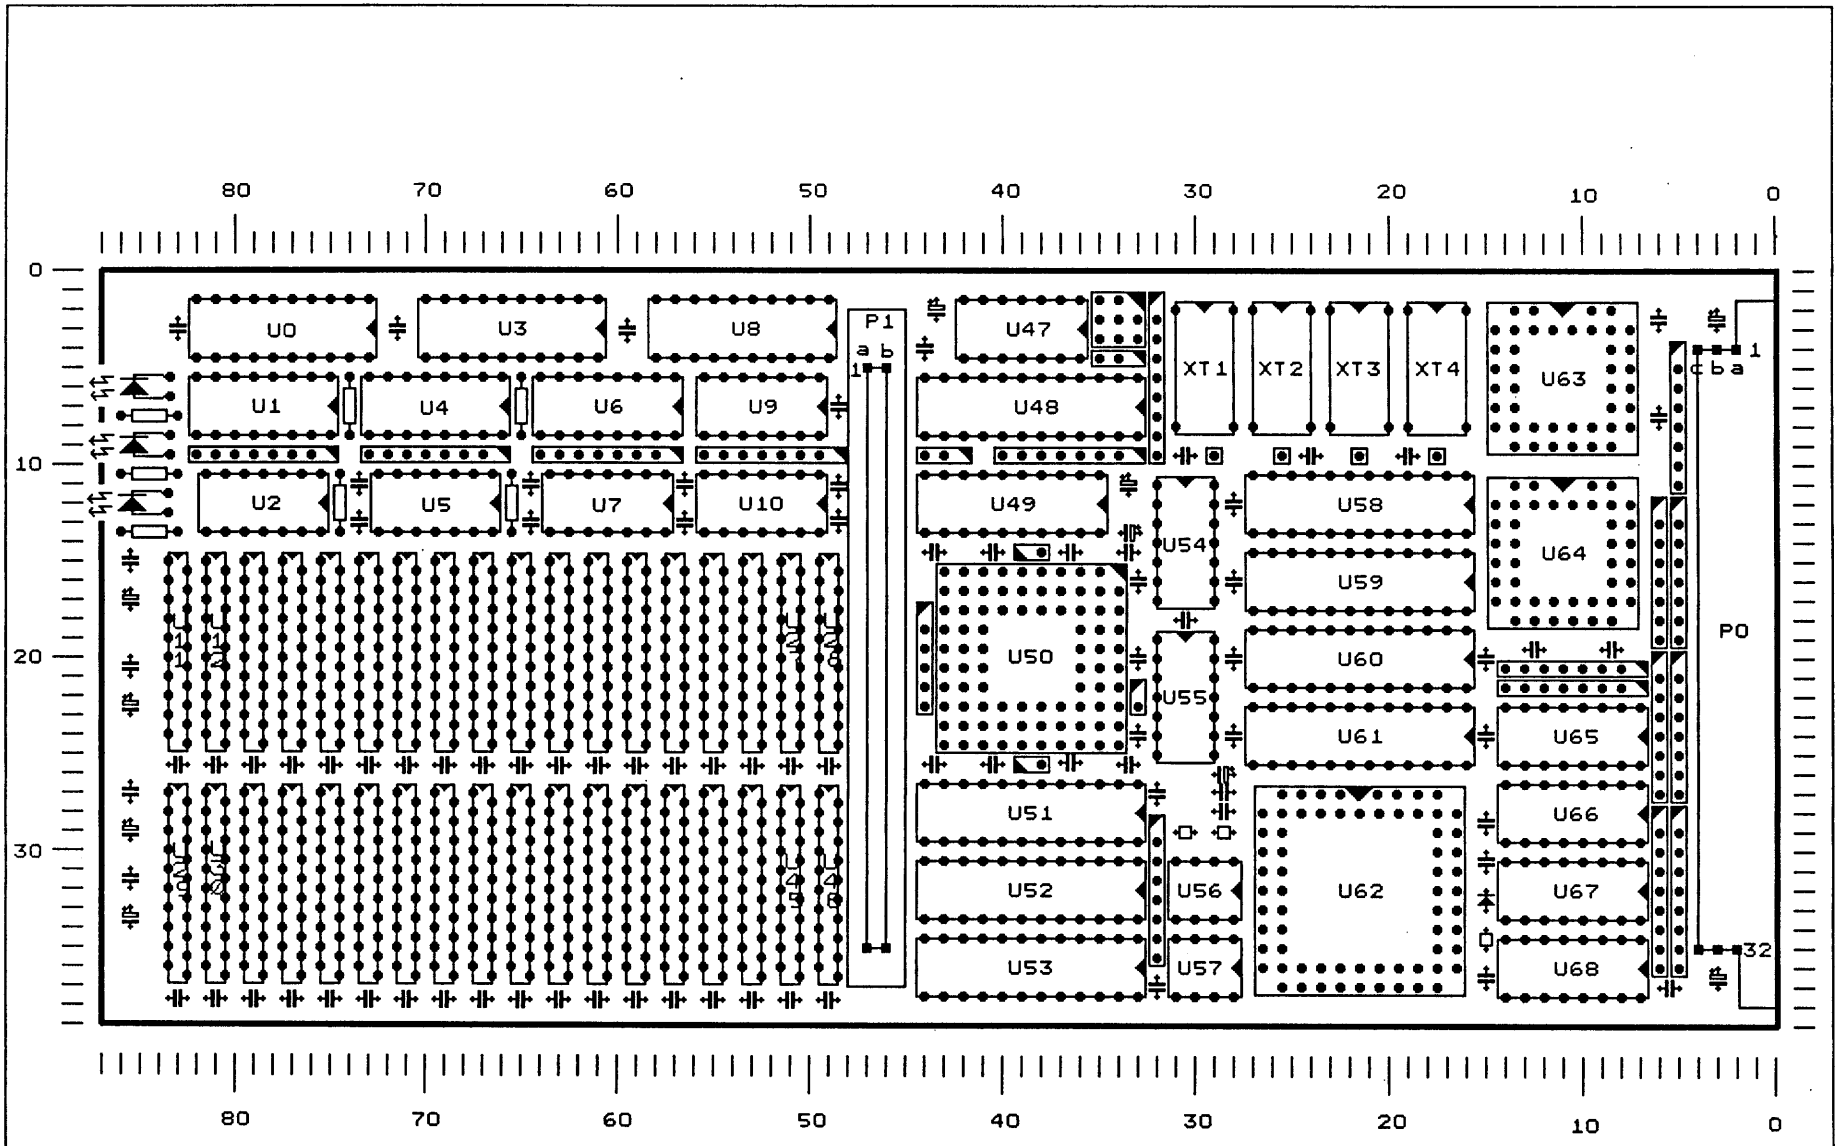

PARSYTEC GmbH Juelicher Str. 338 D - 5100 Aachen		
Size A	Document Number MSC Board Layout	REV 1.1
Date: February 10, 1988	Sheet 7 of 8	

2.3 The speed of the SCSI Interface

The MSC achieves a peek rate of 3.0 - 3.5 MBytes/sec in asynchronous mode. This high data rate can not be sustained by most winchesters, so the limiting factor in SCSI bandwidth is nearly always the winchester.

See sheet "MSC SCSI & Link transfer rate". Due to overhead of the SCSI protocol and internal operation in the winchester, the continuous data rate will be lower than 3.0 MBytes/Sec.

MSC SCSI and Link Transfer Rate

(All values measured with stretched WREN IV)

- SCSI transfer, process switching and all 4 links (8 DMA engines) can operate in parallel.
- SCSI peak transfer rate is 3.5 MByte/sec during a sector transfer. (independant of number of running links)
- Average transfer rate during long transfers is more than 1 MByte /sec.

SCSI Bus activity (measured with stretched WREN IV)

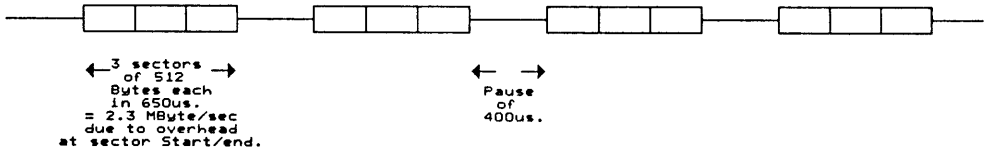

- Test environment:

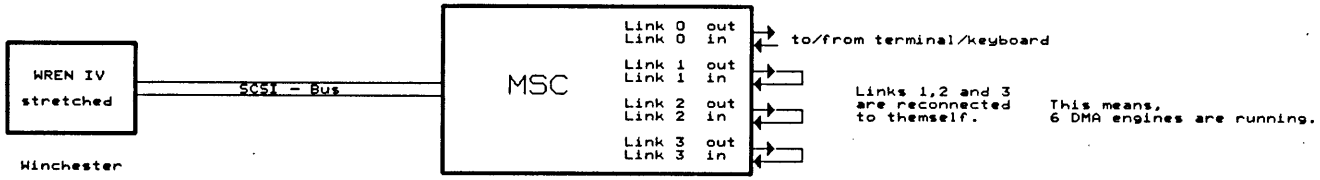

- Link transfer rate during SCSI sector transfer:
  (During SCSI pause the links can reach their maximum speed.)

T8: 1.1 MByte/sec on any link, even if all 6 DMA engines are running.

T4: 570 KByte/sec on any link, even if all 6 DMA engines are running.

PARSYTEC GmbH		
Title: MSC SCSI & Link Transfer Rate		
Size: B	Document Number: MSCRATE.ORC	REV: 1.1
Date: October 21, 1988	Sheet of	

3. The Floppy Section

3.1 The Floppy Controller

The MSC board contains the WD37C65 Floppy Disk Subsystem Controller from Western Digital Corporation. The features are: IBM PC/AT compatible format, dual speed spindle drive support, direct floppy disk drive interface, drives up to 4 floppy or micro floppy disk drives, data rates of 125, 250, 300, and 500 kbit/sec. Input XT1 (26) is driven by a 16 MHz clock.

The internal registers are accessible by the transputer (see "Overview of the OCCAM Address Space").

At the end of a multisector data transfer the Terminal Count input of the floppy controller must be activated to signal the transfer of the last data byte to the controller. See "Flip flop Read" in "Overview of the OCCAM address space".

Floppy data transfers are intended to be done without DMA but by polling. So this flip flop combination was used for this special purpose.

There are three ways for the transputer to notice an interrupt of the floppy controler.

- Reading the Status Register of the floppy controller.
- The interrupt line can be polled (see "BigLatch Status Register")
- The interrupt line can activate the EVENTREQ input of the transputer (see "Events" and "Jumper J9").

The reset input is activated at power on, external reset (see "Reset Signals") or by clearing the flip flop PRES (see "Flip flop PRES").

3.2 The Floppy Interface

Outgoing signals are open collector. Incoming signals have pull up resistors of 1 k ohm. See "Pin out of 34-way Floppy Connector".

The msc.driver software package uses the WD37C65 in IBM-AT mode. This means, Motor Select 1 and Drive Select 1 are both active during accesses to Floppy Drive 1. Floppy Drive 2 is accessed, when Motor Select 2 and Drive Select 2 are both active. The floppy drive 1 must be connected at the end of the 34-wire floppy cable. The wires 10 - 16 comprise a subcable, which must be twistet half a turn before connecting to drive 1. See wiring diagramm of "MSC Backplane".

4. Transputer to Controllers Interface

4.1 The BigLatch

The BigLatch is a multifunctional symmetrical bidirectional driver/latch, which connects an 8 bit bus with a 32 bit bus. Its purpose is to buffer data during SCSI data transfers. The bytewise incoming high speed SCSI data is assembled in the BigLatch to produce a stream of 32 bit words at lower speed for the transputer and vice versa.

For convenience lets call the two busses the P (peripheral) data bus and the T (transputer) data bus. At both bus interfaces there are 4 8-bit latches. The BigLatch connects the 8 bit wide peripheral data bus P to the 32 bit wide transputer data bus T (see sheet "MSC schematic").

The BigLatch can transfer the data on the T bus to the P bus and vice versa (transparent mode).

The BigLatch can latch the data of one of the busses and later send this data to the other bus (latch mode).

There are the following functions:

If the transputer reads a register in the SCSI or floppy controler, the BigLatch works like a transparent bus driver. The 8 bit output produced by the accessed controler is transmitted fourfold to the four bytes of the transputer bus.

If the transputer writes to a register in the SCSI or floppy controler, the BigLatch works like a transparent bus driver in the other direction. The contents of the bits 0-7 of the transputer data bus are transmitted to the accessed controler. Bits 8-31 are don't-care bits.

The latch facility of the BigLatch is used only during SCSI data transfers, when the SCSI controler works in DBA mode: The data transfer controler converts the read/write cycles of the SCSI controler to direct the data bytes into/out of the BigLatch. The data transfer controler supervises the SCSI data transfer which takes place between the SCSI controler and the BigLatch and the transputer. The transputer on the other side of the BigLatch gets/supplies the SCSI data by normal memory read/writes accesses.

During SCSI DBA (direct buffer access of WD33C93) read the SCSI controler supplies data bytewise. Up to four bytes can be buffered in the BigLatch. The DBA transfer has to wait (under control of the data transfer controler) until the transputer reads out the four stored data bytes in the Biglatch by a single word read access.

During SCSI DBA write the DBA transfer has to wait until the transputer writes a word (i.e. four bytes) into the BigLatch. Then the DBA proceeds: The data transfer controler directs the BigLatch to send the four stored bytes one after the other to the requesting SCSI controler.

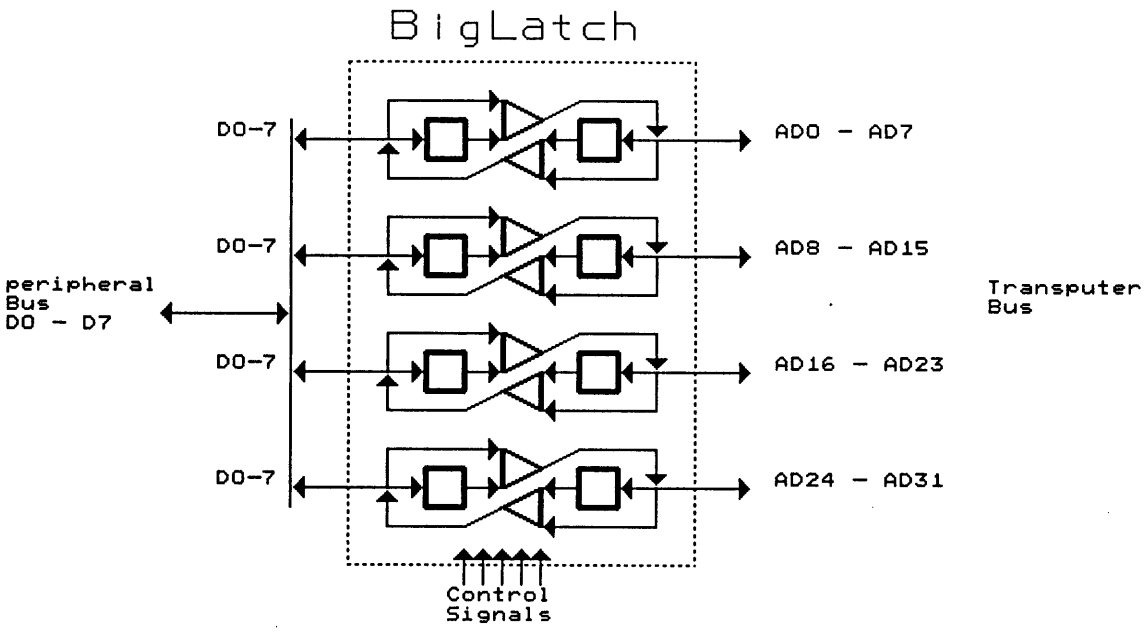

8 Bit Latch

Multiplexer 2x8 Bit -> 8 Bit

PARSYTEC GmbH		
Size A	Document Number MSC schematic details	REV 1.1
Date: February 18, 1988	Sheet of	

4.2 The data transfer controler

The main purpose of the data transfer controler is to establish highspeed SCSI data transfer:

- By controling the BigLatch.
- The transputer and the SCSI controler are tightly coupled during high speed data transfer under the control of the data transfer controler.

The data transfer controler operates internally synchronous with respect to the clock of the transputer and the SCSI controler. It operates asynchronous to the floppy controler.

The data transfer controler contains three registers, one state machine, i.e. an automaton, a watch dog timer and 6 user settable flip flops.

The automaton is called **Byte Counter**, which counts how many bytes are written into (during DBA read) or read out (during DBA write) of the BigLatch. It halts DBA transfer or transputer accesses when the BigLatch needs service from the transputer or from the SCSI controller, because it is full or empty.

The automaton is invisible to the programmer and is triggered by transputer and DBA accesses. But the knowledge about it is somewhat helpful in understanding the data transfer controler.

The six flip flops are called PRES, SCSI, READ, DBA, EXTO and DUMMY. For a description see "Overview of OCCAM Address Space".

The state **BigLatch is full/empty** of the Byte Counter Automaton can be polled by the transputer: Bit 5 of the BigLatch Status Register. This enables SCSI data transfer completely under software control in contrast to high speed transfer, where the transputer makes a block move operation between the BigLatch and the memory.

5. Programming the MSC Board

5.1 Overview of the OCCAM Address Space

The T800/T414 transputer can address 1 Giga-words of 32 bit. In OCCAM the address space ranges from #00000000 to #3FFFFFFF (adresses in present OCCAM-2-Implementation PLACEment as word address):

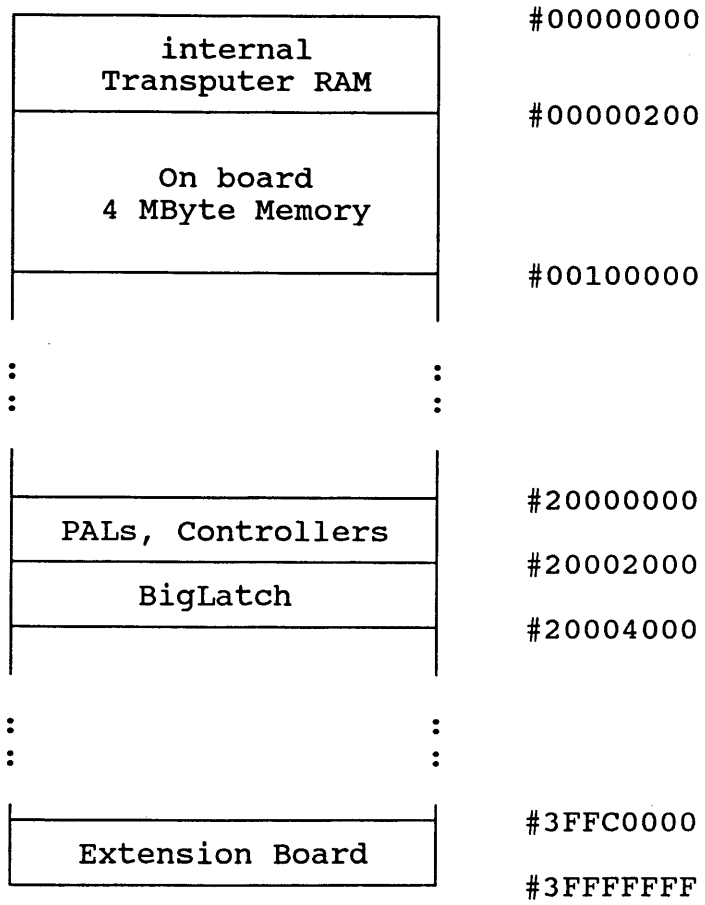

The addresses of all peripherals except BigLatch are located between #20000000 and #20000050. Each peripheral device (controller, PAL) is accessible through the lowest byte of a word (bits 0-7).

Transputer accesses into unused address areas result in an address error (see bit 1 of "Error Register" and "Error and Analyse").

Overview of peripheral addresses

Word address	Bits	Function
#20000000	0-7	Ident bytes from Ident PAL (read only)
#20000020	0-5	Error register of Error PAL (read only)
#20000030	0-3	Link reset out register of Reset PAL (write only)
#20000046 controler	4-7	Status register of data transfer (read only)
#20000046 controler	4-7	Control Register 1 of data transfer (write only)
#20000047 controler	4-7	Control register 2 of data transfer (write only)
#20002000 - #20003FFF	0-31	BigLatch (read/write)

The SCSI and floppy controler share some addresses. They are distinguished by the state of the flip flop SCSI.

———— case flip flop SCSI = 1 ————

#20000040	0-7	Auxiliary Status Register of SCSI controler (read only)
#20000040	0-7	Address Register of SCSI controler (write only)
#20000041	0-7	Register File of SCSI controler (read/write)

———— case flip flop SCSI = 0 ————

#2000040 controler	0-7	Master Status Register of floppy (read only)
#20000041	0-7	Data Register of floppy controler (read/write)
#20000042	0-7	Control Register of floppy controler (write only)
#20000043	0-7	Operations Register of floppy controler (write only)

A read access to any of the floppy or SCSI controler registers, all of which are bytewide, results in a word containing 4 identical bytes. E.g. a read of a status register, which contains #13, returns the word #13131313.

Ident PAL at #20000000

The Ident PAL works like a bytewide PROM and is read only. It can be programmed to hold several bytes for identification or other purposes. There is one default value of #5A.

Error Register of Error PAL at #20000020

This Register latches the following error conditions:

Bit	Error Condition
0	Cleared if ERROR output of transputer was activated.
1	Cleared if transputer accesses an unused address.
2	Cleared if parity error in byte 0 (i.e. bits 0-7).
3	Cleared if parity error in byte 1 (i.e. bits 8-15).
4	Cleared if parity error in byte 2 (i.e. bits 16-23).
5	Cleared if parity error in byte 3 (i.e. bits 24-31).

After the register is initialized (see below) it monitors the 6 error conditions mentioned above. The register latches the first occuring error condition. At the same time the register will be locked, i.e. error conditions which will come up afterwards will not change the status of the error register.