PCI-Bus 3 Use and Data Buffering

25 July 2000

1. Introduction

PCI-bus 3 is heavily used during normal data taking. This may lead to contention and requires that the elements accessing PCI-3 implement some sort of buffering for their data. PCI-3 is used for two main purposes: communication between the Daughterboards (DB) and Buffer Controllers (BC) and VME access to the boards. Figure 1 shows a simplified block diagram of the accesses to PCI-3. Table 1 lists the data transfers that occur on PCI-3. Marked in yellow are those that occur often enough to pose potential bus conflict problems.

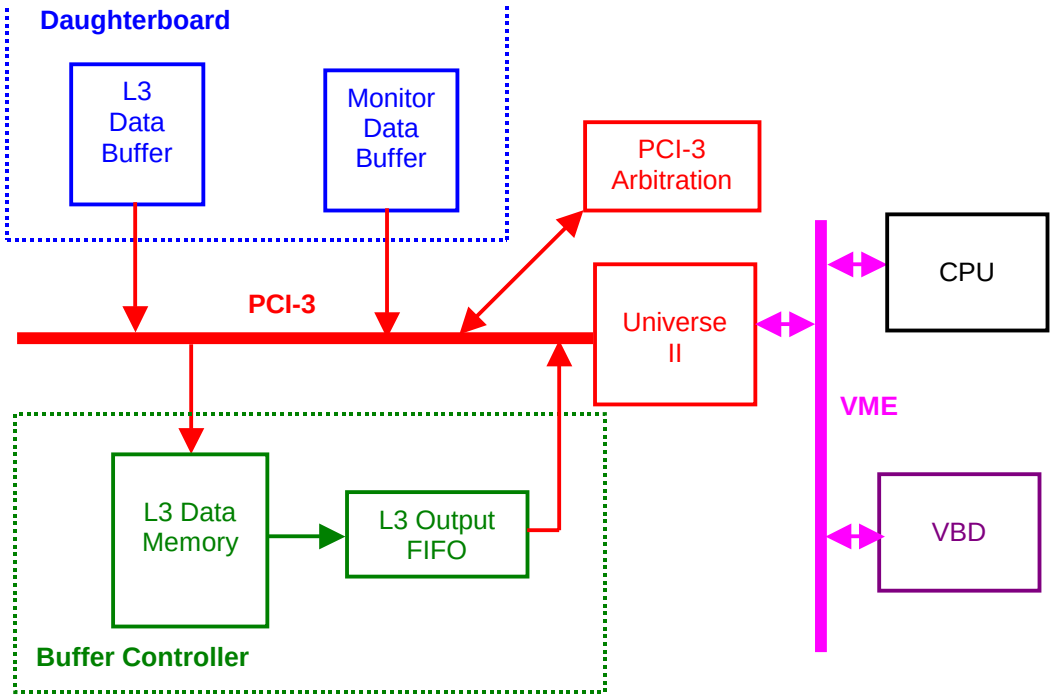

Figure 1: A simplified block diagram of elements accessing PCI-3.

Data on PCI-3		VME	Use	When
Requester	Source			
BC	DB	no	Data to L3 Buffer	L1 Accept
BM	CPU	irq	SCL_INIT interrupt request	SCL_INIT
BM	CPU	irq	Monitor interrupt request	L1Acc & Coll Stat
CPU	DB	yes	Collect Mon Data from DBs	L1Acc & Coll Stat
CPU	DB	yes	Read SCL_INIT status	SCL_INIT
VBD	BC	yes	Read L3 Data from Output Fifo	L2 Accept
CPU	DB	yes	Download / Test	out of run

Table 1: Uses of PCI-Bus 3.

2. PCI Transfer Times

The time for PCI block transfers is composed of two parts[1]: setup time and data transfer time. Setup time is around 8 clock cycles. Data words are then transferred at the rate of one (32-bit) word per clock assuming that the data can be supplied at this rate. The total transfer time is then given by:

$$t_{blk} = (8 + N_{data}) / PCI\text{-}freq = [\ (8 + N_{data}) / 32\ ]\ \mu s.$$

3. Event Timing

Times between various events are given in Table 2. These are taken from the Geographic Section document[2] and Ulrich's Trigger Timing slide[3]. Since there is some delay between the event (*e.g.* L2 Accept) and the action (*e.g.* the VBD reads the BC output fifo) the times in Table 2 only set the scale of things. The actual time between PCI-3 requests can be almost arbitrarily small.

		Min time [μs]		
Time between	Ave [μs]	Limit	STT	Comments
SCL	0.132			
L1 Accept	100	2.6	~9.0	set by SMT readout time
L2 Accept	1000	1.6	1.6	L3 falls behind
Monitor	5 s			comes with L1 Accept
L1 – L2		0.132	0.132	

Table 2: Minimum and average times between various events in the STT system. Limit refers to the minimum enforced by the trigger framework, while STT is a more realistic number during running.

4. Consequences

Since the transfer time for data on PCI-3 will occasionally be longer than the time between actions that will attempt to load a specific PCI-3 data source, each source must

buffer its data and must also hold off new data if this buffer becomes full. The depth of any buffer accessible by PCI-3 is set by four parameters:

1) Fluctuations in the amount of data in an event stored in the buffer.
2) The time between actions that fill the buffer.
3) The longest time required for a PCI-3 transfer. This will probably correspond to the transfer of unbiased events from the output FIFO on the BC to the VBD for L3 readout.
4) Depth of other buffers in the rest of the trigger system.

4.2 Buffer Depths in Events

Depth of buffers in events for the three time-critical PCI-3 data sources is given in Table 4.

The depth of the **Monitor Data Buffer** only needs to be one event (of the maxiumum possible size) because of the relative infrequency at which *collect status* requests are issued.

The depth of the **L3 Output FIFOs** on the BCs that are read by the VBD is set by DØ-wide requirements to be 8 events deep.

The depth of the **L3 Data Buffers** on the DBs is a more complicated issue. By default, these would be set to the L2-standard size for buffers containing L1 data – 16 events. However, this would result in buffers of the same size as the memory in the BCs being required on each DB. This is probably not necessary for efficient operation since the only parameter controlling when these buffers are read is the availability of PCI-3 on the card containing the DB. This is much less restrictive than the parameters controlling reading from and writing to the BC memory, which are affected by the state of the entire STT crate. The depth of the DB L3 Data buffers is therefore set by the maximum amount of time that PCI-3 can be busy on another action (readout of the card's data to VBD or readout of the card's monitoring data). Note that the relevant time here is that for an action on a single card – since we are only interested in the occupancy of that card's PCI-3.

The largest data volume transferred over PCI-3 (for the STCs and TFCs) will be for the L3 data for *unbiased events.* Both the data going from the DB to the BC and the data going from the BC to the VBD will be of approximately the same size here (give or take headers and trailers). However, the transfer time for the DB to BC transfer is set by the PCI frequency (32 MHz), while the BC to VBD transfer is set by VME (16 ??? MHz). The longest time will therefore occur when the VBD requests readout of the card's unbiased L3 data from the BC Output FIFO, over PCI-3 to the Universe II chip and then to VME.

Since the volume of data produced by the FRC is small (54 words or less), the largest data block transmitted on the FRC PCI-3 will come from the monitoring data. We have not finalized the specifications for FRC monitoring, but estimates are less than 100 words.

Data sizes and the corresponding readout times are given in Table 3. Numbers for data sizes are taken from Ulrich Heintz' note *Data Output from L2STT to Level 3*[4]. The relevant times for transfers from the DBs to the BC are given in the *PCI* row, since this is the only bus used for the DB to BC transfer. For data collected by the VBD, the times are given by the *VME* bus frequency since the Universe II chip only requests data from PCI at the VME frequency (16 ??? MHz).

L3 data needs to be buffered in the DBs for at least as many events as occur during the readout of the maximum size unbiased or monitoring event. Using the average time between L1 accepts (100 μs) to set the scale for when events arrive at the DB L3 data buffer, we find the minimum depth of the buffer in events required to cover any possible PCI-3 action. These numbers are given in Table 4.

4.2 Buffer Depths in Words

The depths of the various buffers (in 32-bit words) used for PCI-3 transfers are given in Table 4.

The depth of the **Monitor Data Buffers** is set by the amount of data to be monitored.

The depth of the **L3 Output FIFOs** on the BCs is simply the size of a single buffer in the BC memory (32k) multiplied by 8.

The **L3 Data Buffers** on the DBs should be large enough to fulfill either of the conditions below.

a) One unbiased event of maximum size plus one normal event of maximum size.
b) *N* normal events of maximum size.

For the FRC and TFC, case a) is the relevant one given the size of the maximum normal and unbiased events. The STC size is set by case b).

This assumes that there will be at most one unbiased event per 16 normal events. If we run with a higher rate of unbiased events, the STT will not be able to keep up! Aside from that the buffer depth requirements are conservative in that normal events of the maximum size will not occur at the average L1 accept rate. If this does happen for some reasonably short period of time the STT will react by asserting L1_BUSY and causing readout deadtime.

4.3 Buffer Full Actions (Busy)

In the event that a buffer fills up, it must hold off the appropriate triggers until it can drain itself. This is done by asserting L1/L2_BUSY to the Hub. Each buffer must be capable of setting the appropriate status line on J3. The OR of these over all the cards in the system is collected by the BM, which generates the requests to the SCL mezzanine card.

This means that each DB must use the L1_BUSY line in its direct connection to J3, while the BC must use the L2_BUSY line to J3.

		Ave	Max-Normal	Max-Unbiased
FRC	L3 Data Size [32-bit words]	0	0	54
	PCI Transfer Time [μs]	0	0	1.7
	VME Transfer Time [μs]	0	0	2.7
STC	L3 Data Size [32-bit words]	40	2304	5847
	PCI Transfer Time [μs]	1.3	72.0	182.7
	VME Transfer Time [μs]	2.0	115.2	292.4
TFC	L3 Data Size [32-bit words]	6	94	6173
	PCI Transfer Time [μs]	0.2	2.9	192.9
	VME Transfer Time [μs]	0.3	4.7	308.7

Table 3: L3 data sizes and transfer times for DB to BC transfers (PCI) and BC to VBD transfers (VME).

PCI-3 Source Buffer	Location	Depth [events]	Minimum Depth [words]	Action when Full
L3 Data	DB-FRC	2	154	L1_BUSY
	DB-STC	4	9216	
	DB-TFC	4	6267	
L3 Output fifo	BC	8	256k	L2_BUSY
Monitor Data	DB	1	?	L1_BUSY

Table 4: Minimum PCI-3 data source buffer depths and actions to take when buffers are full.

References

[1] PCI reference manual.
[2] revision 22-Oct-99 is linked at http://www.nevis.columbia.edu/~evans/stt/docs/GS_spec1.pdf
[3] U. Heintz, *DØ Trigger Timing*, (30-Jun-00), http://physics.bu.edu/~heintz/STT/trigger_timing.pdf
[4] U. Heintz, *Data Output from L2STT to Level 3*, (4-Feb-00) http://physics.bu.edu/~heintz/STT/L3_readout.pdf