

Recycler BPM - Bug #12915

Missing first injection and other odd behavior with 6+6 slip stacking state

06/13/2016 10:31 AM - John Diamond

Status:	Assigned	Start date:	06/13/2016
Priority:	High	Due date:	
Assignee:	John Diamond	% Done:	50%
Category:		Estimated time:	0.00 hour
Target version:		Spent time:	20.25 hours
Description			
MI folks began noticing strange behavior when they tried moving to a 6+6 slip stacking state:			
On some cycles -			
<ul style="list-style-type: none">• the first injection flash measurement was not being triggered• the state machine will stay in closed orbit or flash mode after the first injection			

History

#1 - 06/13/2016 10:36 AM - John Diamond

Peter and I were able to setup the RRBPM development test stand and reproduce the problem. Peter was able to diagnose a problem in the timing board firmware that was causing duplicate BES triggers. After flashing the new firmware we were able to determine that the system was behaving as expected with 12 flash measurements but adding a thirteenth (12 injections and 1 extraction) would cause the problem to occur. We believe the problem has to be the amount of data being transferred after the end-of-beam. The proposed solution is to make the number of data points captured on a flash measurement variable (right now it is hard-coded to be 512).

#2 - 06/13/2016 12:12 PM - John Diamond

- % Done changed from 0 to 90

Reduced the size of RRBPM_FLASH_TBT_BUFFER_SIZE from 512 to 448. Also, a couple of spots in RRBPMEchoTekPool::safeWriteTo(..) had 512 hard-coded as a magic number. Replaced these occurrences with RRBPM_FLASH_TBT_BUFFER_SIZE. Compiled, deployed dev and tested upstairs. Noticing no bad cycles with bpmShowModes but we did notice at least one bad cycle upstairs on Peter's scope but for the most part it seems that reducing the number of turns in the flash turn buffer has fixed the problem.

However, this may not be the best solution. Changing the buffer size probably will mean that R43, R39 and other programs need to be modified. It's up to MI group if this is preferable to removing a single flash measurement from the command list.

#3 - 06/13/2016 05:05 PM - John Diamond

- % Done changed from 90 to 50

Phil says that 13 flash measurements with 512 turns is a necessity. We promised to look for opportunities to optimize the end-of-beam data transfer and position calculations.

#4 - 06/23/2016 12:47 PM - John Diamond

6+6 and Muon Timing

- The 6+6 slip-stacking state (RR state 6) has 13 flash measurements (12 injections and 1 extraction) and uses the TCLK \$E3 as a cycle reset
- The Muon campus g-2 state (RR state 20) has only 2 flash measurements and uses the TCLK \$E9 as a cycle reset
- Flash measurements are collected and readout after the end-of-beam event TCLK \$E6 in one DMA transfer per board, one after another
- After readout, positions are calculated for each flash measurement buffer. Each buffer contains 512 turns of I/Q data for all BPM channels (~80 channels per node).
- After positions are calculated the first or last turn of beam is found in each flash buffer using a linear search

A typical supercycle begins like this:

Time	Event	
0.000	\$E3	Beginning of first 6+6 NuMI cycle
0.865	\$E6	End of beam
1.000	\$E9	Beginning of Muon cycle, time between cycles = 135ms
1.26	\$E6	End of beam

1.333	\$E3	2nd NuMI Cycle, time between cycles = 73ms
2.198	\$E6	End of beam
2.666	\$E3	3rd NuMI Cycle, time between cycles = 468ms
3.351	\$E6	End of Beam
... Repeats until slow spill ...

The time between NuMI cycles is 468ms while the time between the NuMI cycle and the Muon cycle is much shorter - 135ms. The time between the Muon cycle and the next NuMI cycle is even shorter - 73ms.

Using the system clock I timed the following tasks to be completed after the end-of-beam event -

Task	Time	
Echotek Readout	151ms	Readout is always performed for 12000 samples per channel to support 20 flash measurements buffers
Calculate positions	228ms	For 13 flash buffers or ~17.5ms per flash measurement
Find First/Last Turn	47ms	For 13 empty flash buffers or ~3.6ms per buffer but would be much less with actual beam signals
Total	426ms	

If we reduce the number of samples collected from 12000 to 6656 (13 * 512) then the Echotek Readout task improves to 91ms, total **366ms** but still more than twice (2.7x) as much time as we're given between the end of state 6 and the beginning of state 20.

Conceivably, we could use an MVME7100 in place of the MVME5500. This board has a 1.3GHz dual-core CPU. The faster clock frequency should get us a 30% improvement in the processing steps, ~82.5ms. Re-engineering the code to split the position calculations between both cores should yield another 142ms for a total time between cycle requirement of 142ms. Much better but still short by 7ms. Charlie points out that we would have to upgrade to VxWorks 6.9 for SMP support.

However, even with the increased processing power the time to readout the Echoteks is still 91ms, more time than we're given between the end of the Muon cycle and the beginning of the next NuMI cycle. We would have to re-engineer the readout routine to only readout 1024 samples in this case ~14ms, leaving only 59ms for processing.

#5 - 06/23/2016 02:56 PM - John Diamond

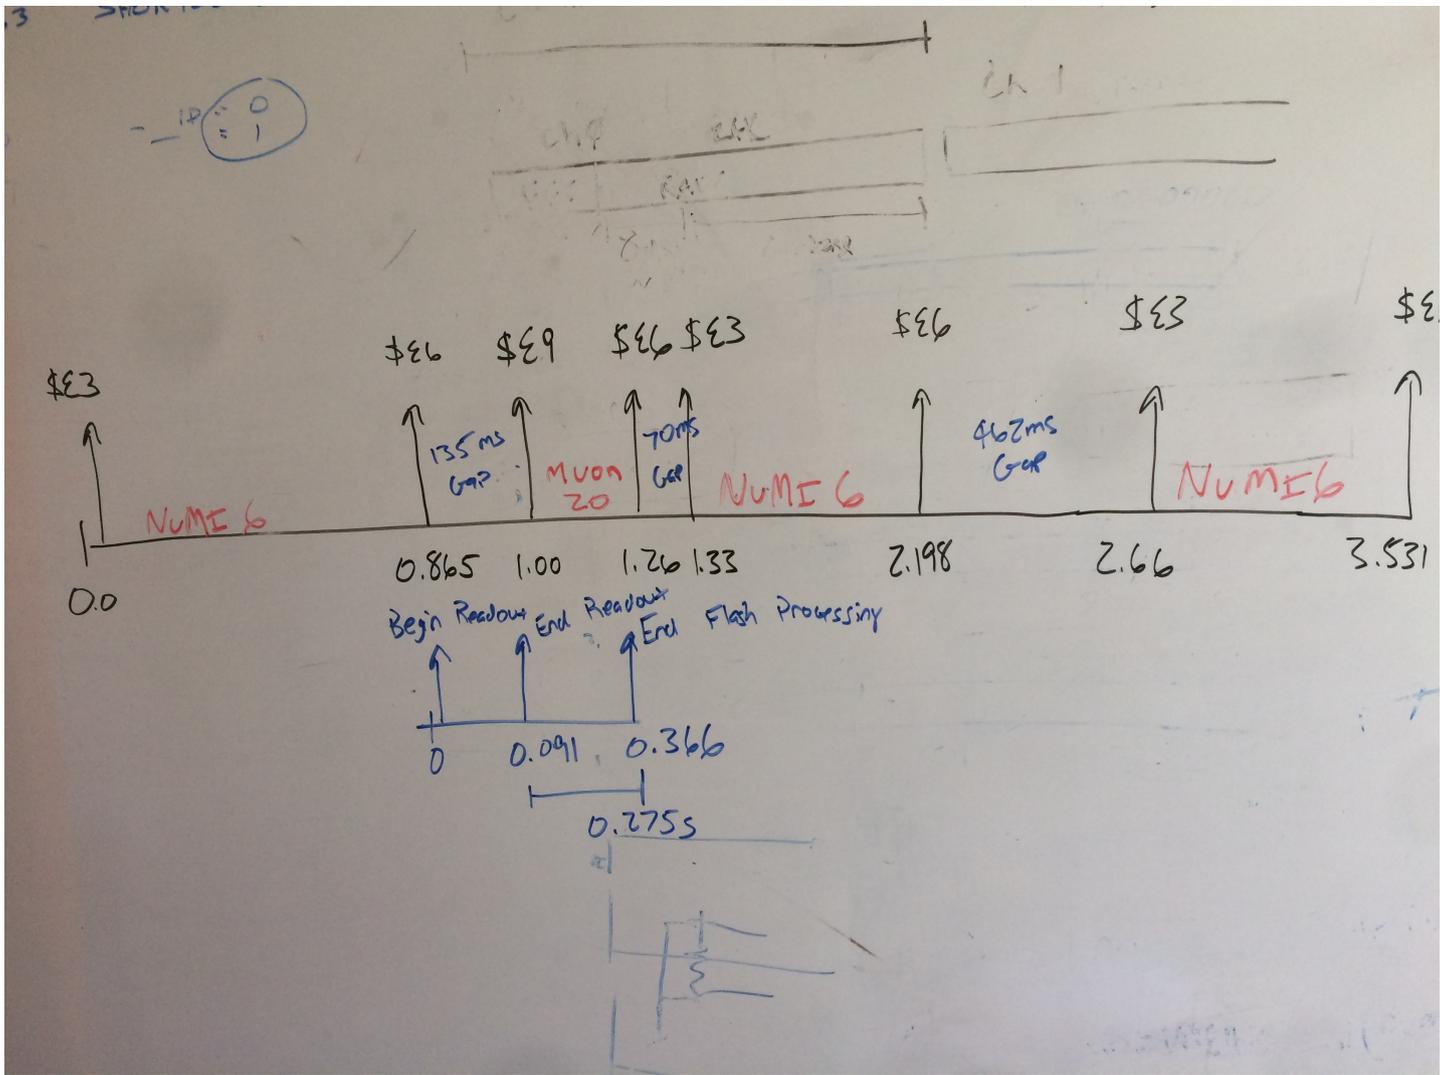
Modified the ecdr814fc_mibpm readout function to dispatch all DMAs at once rather than one at a time. This resulted in about a 0.5 - 1ms reduction in the readout time, not really that big of a gain so I'm going to roll it back.

FWIW - I calculated that for a 40MB/s bus the readout of 6656 samples for all channels on 10 Echotek's should take 53ms and we're at 91ms.

#6 - 06/27/2016 02:20 PM - John Diamond

- File RR BPM Timeline.JPG added

Attaching a timing diagram for the table in comment [#4](#).



#7 - 06/27/2016 02:21 PM - John Diamond

- File deleted (RR BPM Timeline.JPG)

#8 - 06/27/2016 02:23 PM - John Diamond

- File RR_BPM_Timeline.jpg added

Files

RR_BPM_Timeline.jpg	1.08 MB	06/27/2016	John Diamond
---------------------	---------	------------	--------------