Booster damper logbook part 3

Progress notes by Bill Ashmanskas

2008-07-16

The big goal for today is to make the damper diagnostics more usable. Right now, it takes far too long to see what the damper is doing for each pulse. Now that the damper is basically working, I don't think it is necessary to look bunch-by-bunch very much of the time. I think that the most frequent mode of readout will be to look at a small number of summary words per turn (or every nth turn). Of course I want to preserve the ability to read out data bunch-by-bunch for every turn, for debugging or for e.g. looking at the frequency spectrum of coupled-bunch modes.

Examples of summary words would be
- turn ID, RF frequency, milliseconds since reset
- Quadrature and InPhase signals for x,y planes, summed over bunches
- abs(filter) and abs(kick) signals for x,y planes, summed over bunches

Other possibilities include
- RF frequency (though this can be tabulated)
- milliseconds since reset (though this can be tabulated)
- values for selected bunch instead of summed over bunches
- maximum absolute value seen over all bunches

The first thing I did was to run the 'dofifo' script without having changed anything, just to reproduce where I am. The damper is off, so I collected data for a $1D (miniboone) cycle.

Next, I added a pipeline stage in front of the existing path for the FIFO data and write_enable. This is the easiest way for me to insert a parallel logic path without screwing up this path: I can multiplex between the old and new readout schemes at the added pipeline stage. So I am recompiling the program that contains no change except to add the pipeline stage, to make sure I didn't break anything.

OK, step zero succeeds! On to step one.

Step one (multiplexing between the existing write code and new turn-by-turn code) at least didn't screw up the old writes. But I did screw up the gating of FIFO writes with the FIFO_control logic, so I didn't manage to capture data for a meaningful time interval. Try again....

Success! Now I have my little one-per-turn FIFO ntuple. One irritation is that it gets confused when the bunch counter is reset to zero when my notch finder does its thing around turn 2000. Perhaps I'll fix that. But a higher priority is to add the filter and kick variables. I also see that the horizontal ADC is saturating at transition -- something I should probably fix.

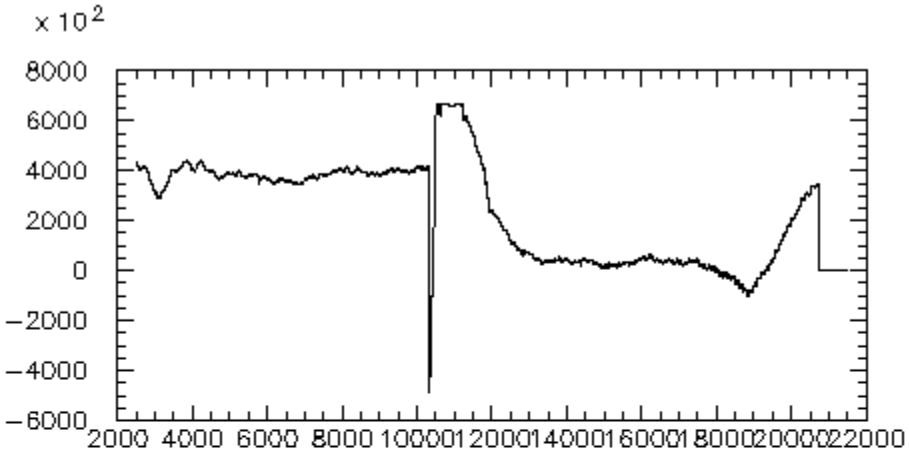

xq VS. turn

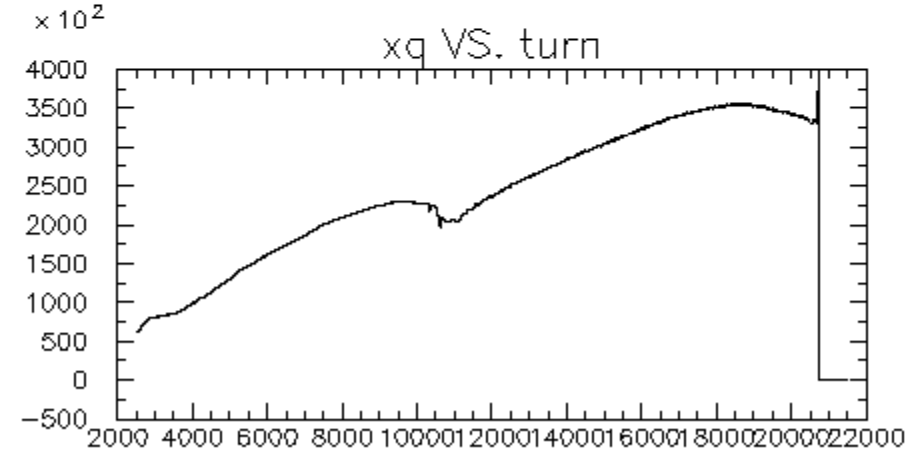

yq VS. turn

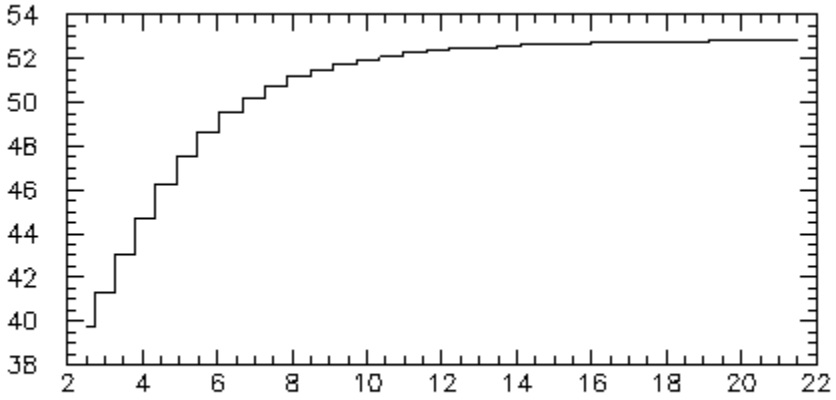

frf/1000 VS. turn/1000

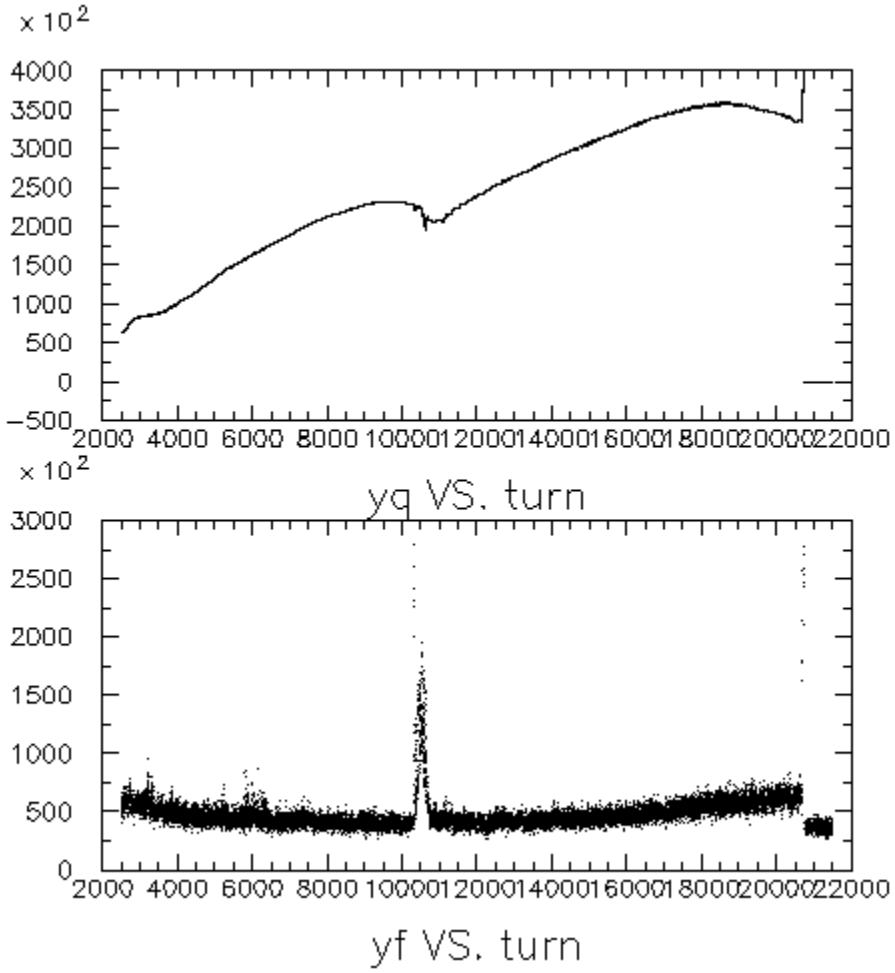

I got the filter signals into the FIFO, too. Here are the quadrature and filter signal for the vertical damper. Quadrature is signed quantity summed over buckets for one turn. Filter is absolute value summed over buckets for one turn. It looks as if there is activity worth damping immediately after transition!

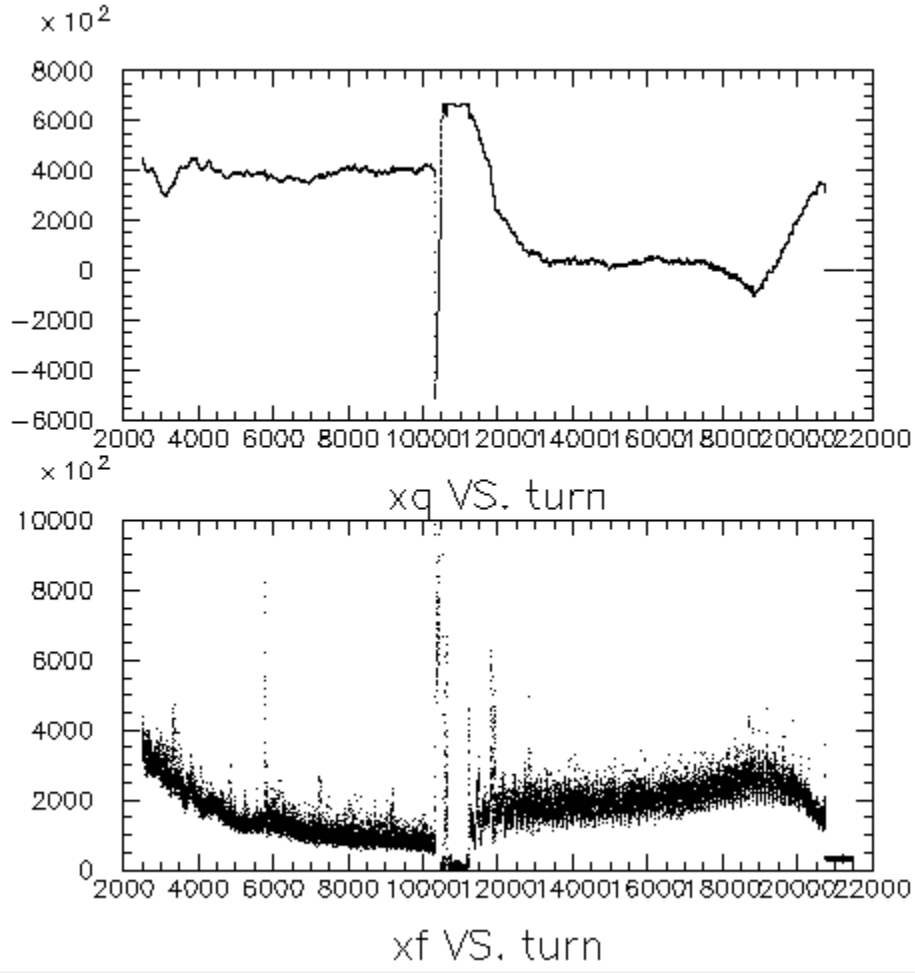

Here is analogous graph for X damper. There is some saturation just after transition, so a bit more attenuation is needed.

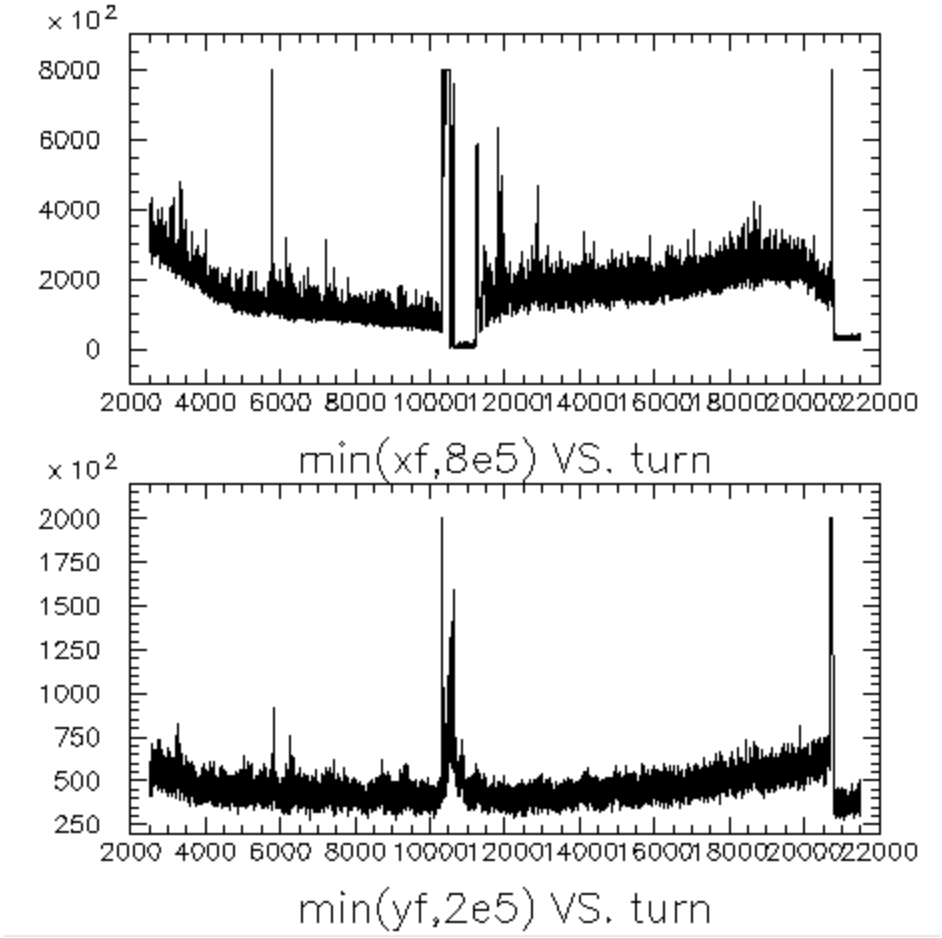

Here are filter signals (with very large values truncated) for both X and Y, for another pulse. By the way, damper is off for all of today's graphs.

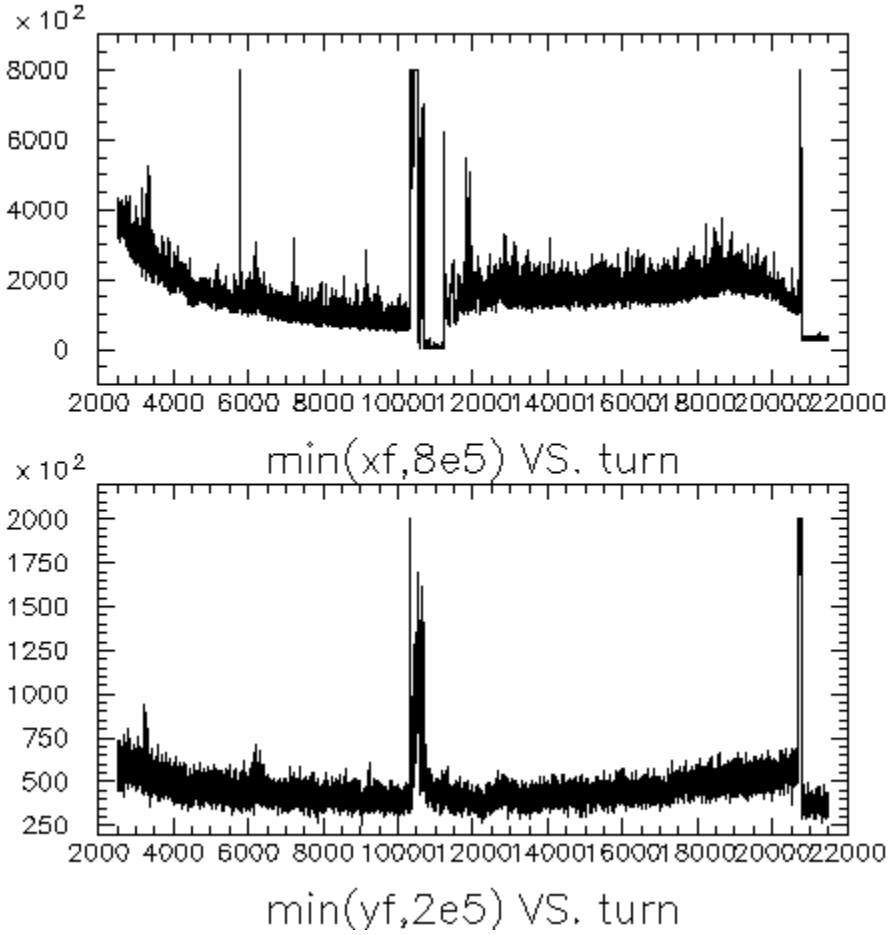

Same graph for another pulse; X and Y filters. X channel still saturates just above transition.

Just before heading home, I replaced 3dB pads with 6dB pads on horizontal channel.

2008-07-17

Let's see whether 3dB of attenuation helped the horizontal saturation.

Rats. Access day. If I had taken the time to read the accelerator plan, I could have worked late into the night and slept all day.

2008-07-18

Let's finally see whether 3dB of attenuation helped the horizontal saturation.

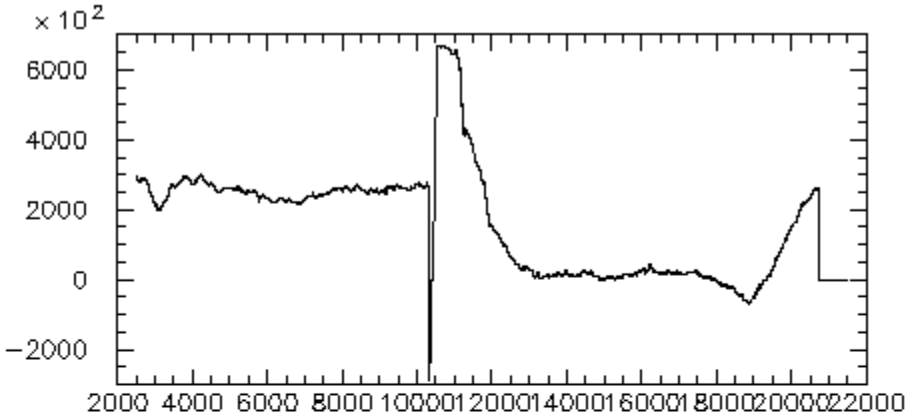

xq VS. turn

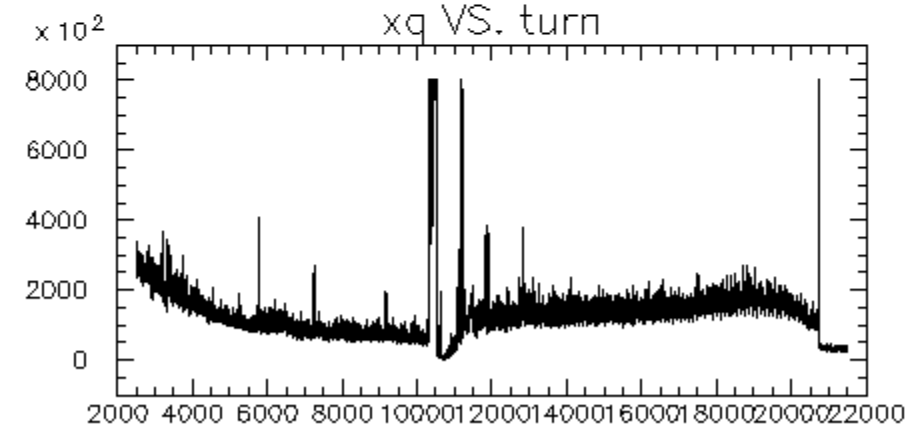

min(xf,8e5) VS. turn

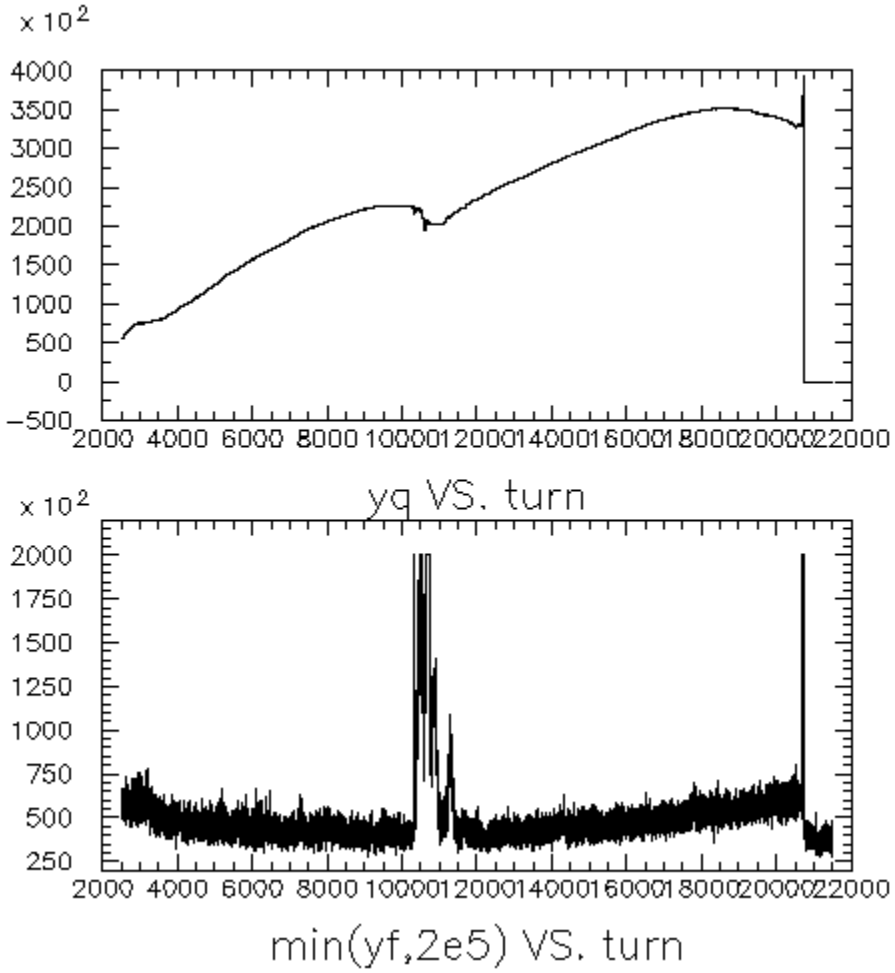

Looks as if saturation is fixed.

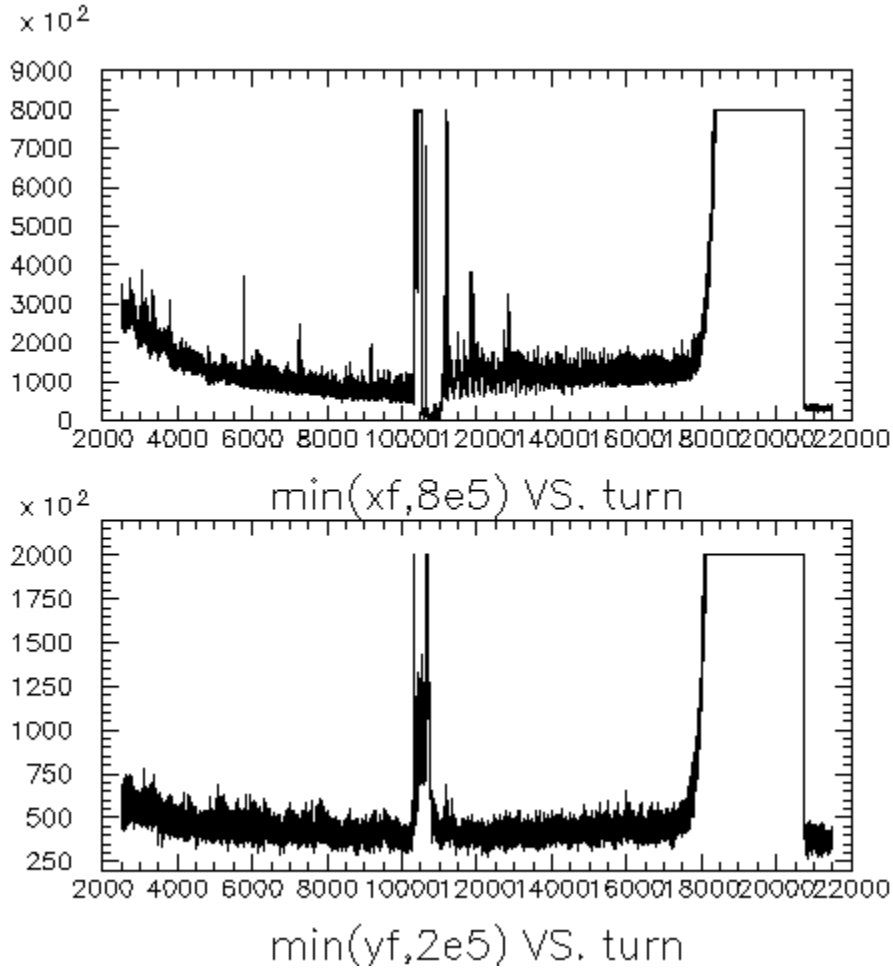

Here is a $17 cycle (damper off). You can clearly see the blow-up toward the end of the cycle.

Toward the end of the day today, I got the kick timed in for turn ~ 6000 (i.e. pretty early in the cycle), and thus (by making the bucket offset linear in RF frequency and leaving it as is for the end of the cycle) got it timed in for the entire cycle. I think this worked, though damping early in the cycle could use more testing.

Tomorrow, I want (a) to read the kick out in the FIFO and (b) to make the gain vary throughout the cycle, as a lookup table on high bits of turn number. I also want to make the read-out filter value be before the gain factor.

2008-07-19

I now read out the summed filter values before the shifting by DxxGAI or the scaling by DxxBBG, which makes it much easier to see what's going on when gain is being tuned. Also, I am writing to the FIFO typically every 3rd turn now, and typically skipping the fills of the big lookup tables for DAC delays vs. turn.

OK, now I'm writing the summed kick values to FIFO, too, 16 bits each.

OK, now I have gain ramp, for now loaded with values proportional to momentum, by taking beta=f/(53.1 MHz), p ~ gamma*beta. I think, however, that 6 bits of gain will not be adequate. Also, it would be good to have an overall gain scale factor by which momentum table is scaled.

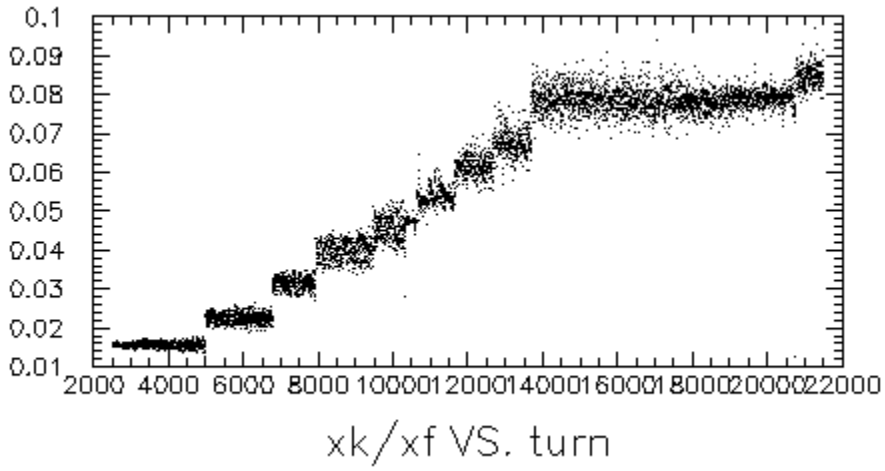

Other questions:

Am I using enough of the DAC bits?
What is the range of DAC voltages output by board? Does it reach +/- 1V?
What is scale factor between kick FIFO readout and DAC voltage?

2008-07-21

I spoke with Valeri Lebedev about how we ought to put the damper to use. He sees two uses:

reduce losses for NuMI slip-stacking
increase proton brightness for tevatron shots

In the second case, one wants to reduce emittances in all three planes. In the first case, one wants to run the longitudinal quadrupole damper more aggressively, to decrease longitudinal emittance above transition. He suggests using Main Injector diagnostics, particularly for measuring longitudinal emittance.

2008-07-23

Arrgh, a somewhat frustrating couple of days. I have been developing/running the damper readout/analysis software on balrog, my behind-the-firewall Linux box, which lives out in the AP50 counting room. Balrog's fan has been making awful noises for several weeks. Yesterday, I took the time to replace it, so that the machine doesn't overheat and fail catastrophically. Success. But I should really migrate over to a centrally supported machine. Toward that end, this morning I started running the damper readout on adclx (clx55), a console linux node maintained by the controls department.

I had to make numpy work on adclx. It would be nice eventually to run the most up-to-date numpy/scipy, but for now I copied the balrog:~ashmansk/code/numpy-0.9.5 directory tree over to clx55:~ashmansk/code/numpy and did 'python setup.py install --prefix=$HOME --exec-prefix=$HOME', which put its output files in ~ashmansk/bin and ~ashmansk/lib. Then I had to add $HOME/lib/python2.3/site-packages to PYTHONPATH and $HOME/lib to LD_LIBRARY_PATH in .bashrc. The second was so that I could copy liblapack.so.3 and libblas.so.3 from balrog's /usr/lib to clx55:~ashmansk/lib. I also copied libcern.so to ~/lib -- this is my shared-object version of the entire CERNLIB (hbook, higz, etc.), including my C interfaces to the parts that I use. I need it so that I can run my 'convntup1' program that reads a flat binary file and writes a PAW ntuple.

The up side of this is that I will no longer be working on a private machine. The down side is that all of this stuff will still only work out of my own directory (though Phil could trivially run it from his area by adding 2-3 lines to his .bashrc). Once the dust starts to settle, I need to put the control/readout into standard ACNET form. I'm running now out of clx55:~ashmansk/proj/damper/

OK, now that I am running on clx55, the next thing I want to do is to reload the damper firmware and to test the changes I compiled on Monday. This reminds me that it is completely shocking and scandalous that I do not have a second damper board programmed as a tester, or at least some way to inject test signals into the damper board and to record its output, without using the beam for this purpose. What would Ted Liu say?!

Anyway, Monday's recompile seems not to have broken anything. I succeeded at antidamping then quickly re-damping (vertically) starting at turn 14000, 13000, 12000, 11000.

Then Bill Pellico came over and pointed out to me that the main thing stopping us from trying out the damper on real accelerator cycles is that there aren't very good diagnostics for the Booster damper. (The most useful non-expert diagnostics for the MI damper are 720 Hz Fast Time Plots. But 720 Hz isn't really adequate for the Booster.) So I spent much of the day considering all kinds of options for getting suitable data out of the damper. Probably the easiest option, suggested by Dennis Nicklaus, would be to make some array devices that store waveforms, which one could read out shortly after the cycle of interest. Snapshot plots in the OAC don't seem to be an option. I guess snapshot plots will also not be trivial in the VME front end. This issue deserves more thought. For now, I will probably send the filter and kick signals to an oscilloscope.

Nathan and I also set up one of the VME damper boards in the VME crate that lives by the damper. Nathan set up pbar-scope-03 to look at the state of the PLL phase lock on the VME damper board. I had a big adventure getting the TDS7104 scope (node name jorkins) to display its traces on a remote computer in the MCR. It required a firmware update to provide the 'display:showremote on' command.

After hours, I started trying to put the new scope to use, displaying the fitler signal throughout the cycle. But nothing I did seemed to give consistent results. I suspected that I had introduced a bug into the Altera program and tried backing out. I should have become more suspicious when the FIFO data path started to misbehave. It turns out that some combination of VME damper board and TDS3000 oscilloscope (located just beneath the non-VME damper board) caused the damper's clock to glitch. I finally started to suspect

this when I got junk data out of the FIFO even after backing out to a version of the damper firmware that had worked yesterday morning. So the VME damper crate and the TDS3000 scope are off for now. Also, I moved the scope over to the next rack. Later in the day, Nathan and I tracked down the source of the clock interference to the NIM fan in/out board that we've been using for the RF clock. If we use a second quadrant of the module to buffer the clock before it drives the VME damper board, the problem disappears. Hmmm.

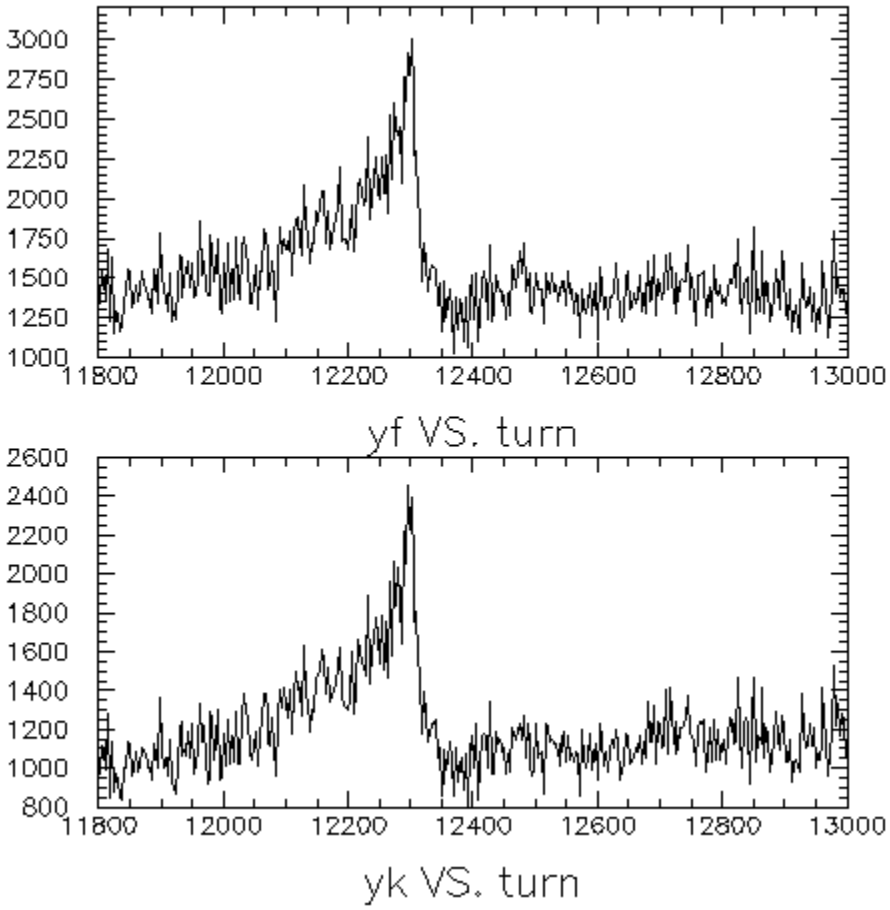

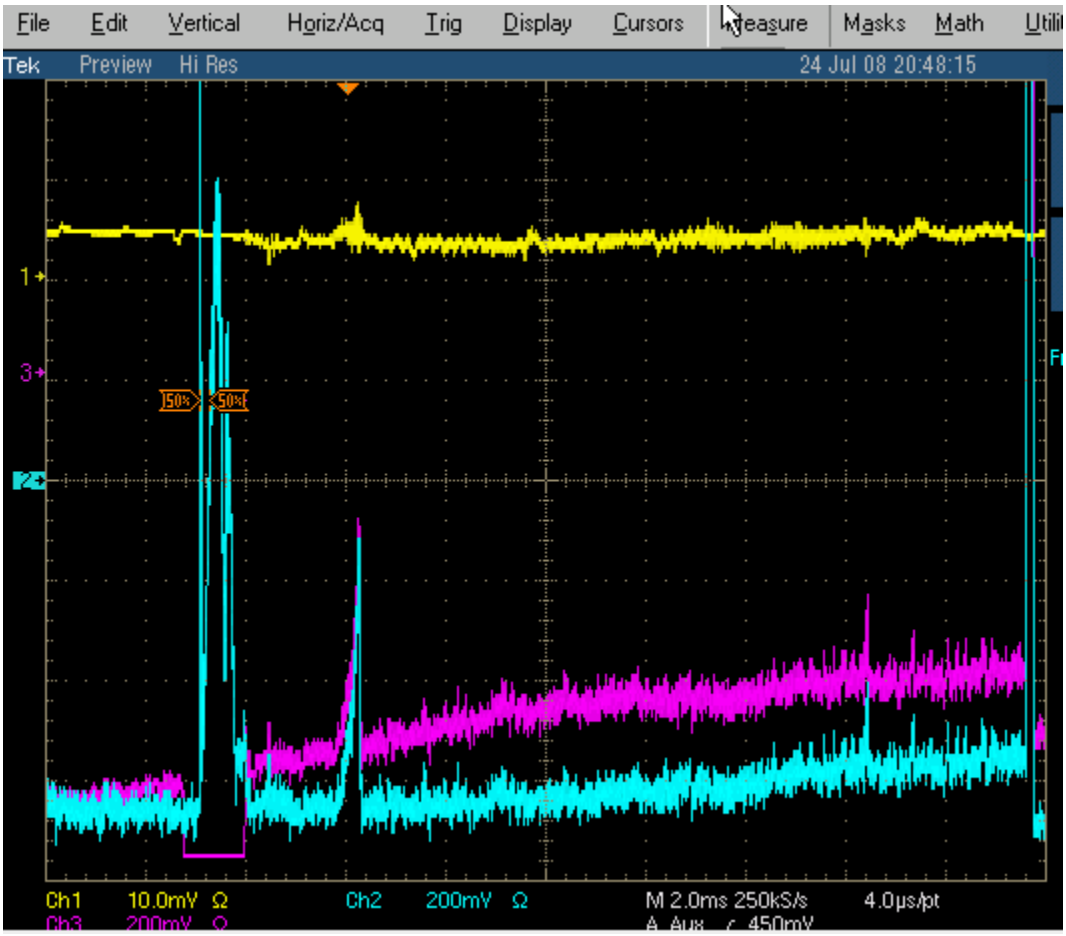
File Edit Vertical Horiz/Acq Trig Display Cursors Measure Masks Math Utili
Tek Preview Hi Res
24 Jul 08 20:48:15
1
3
50%
50%
Ch1 10.0mV Ω
Ch2 200mV Ω
M 2.0ms 250kS/s
4.0µs/pt

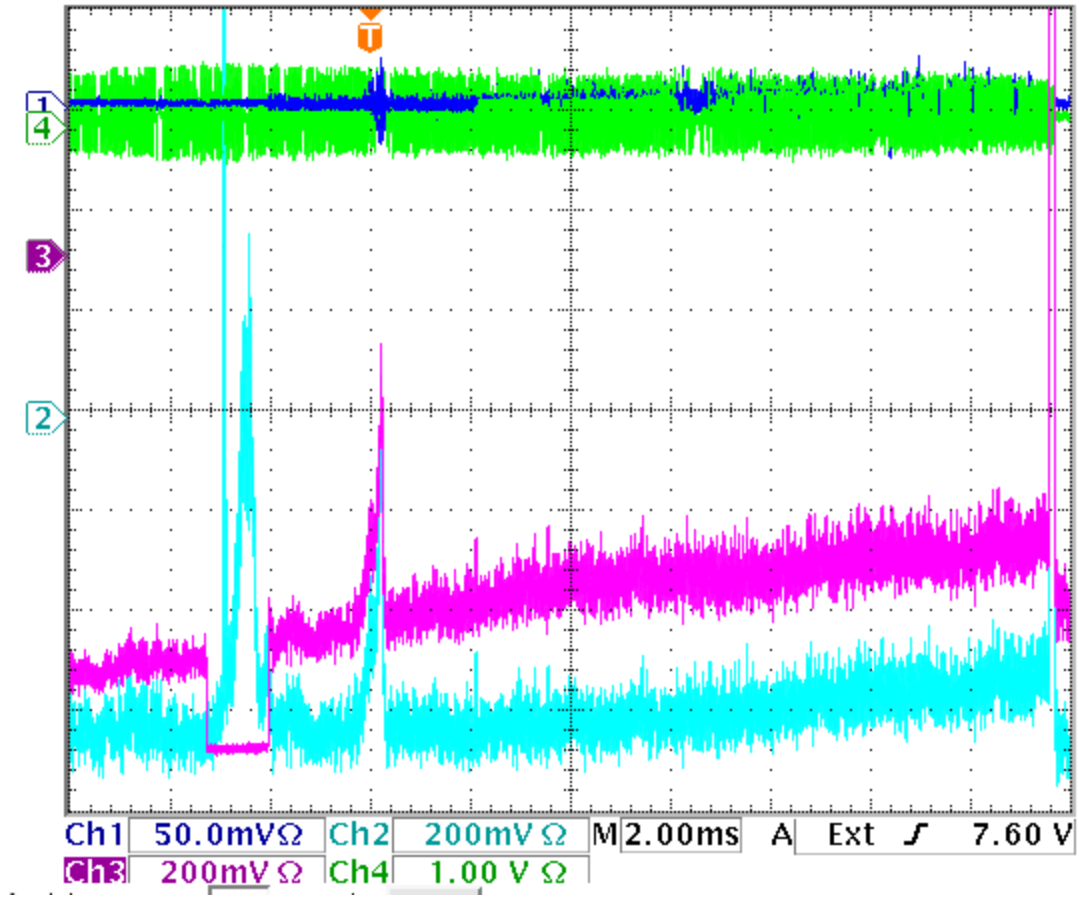

2008-07-25

We had a big meeting this morning over by the damper racks with Nathan Eddy, Jim Crisp, John Seraphin, and Bill Pellico, to decide what would go where. John has a nice plan to clean up the racks, so that the cables aren't such a rat's nest. We've already made good progress on cleaning up old cables, modules, etc. Also, one of the two CAMAC crates is now gone, and the other will be removed next week. Linda set up a rack-mount PC for working near the damper. Steve Conlon pointed out how we can use the rollable consoles for ACNET, by plugging in two racks over. Steve will pull an additional four ethernet cables next week, so that we can keep everything online at once.

As shown above, last night I got two comfort display traces onto an oscilloscope, using the two spare DAC outputs of the damper board. But I'd really like at least four comfort display traces, to fill an oscilloscope's four channels. Also, one can put four traces onto a Fast Time Plot, etc. Most of the time, I think one would look at Xfilt, Yfilt, Xkick, and Ykick, summed

over buckets for a turn. The disgusting (but ingenious!) idea that we decided to use for this is to send a ~ 25 Mbps bit stream from a TTL output of the damper board into a TTL input of a spare AP2 BPM board (NIM module), whose FPGA will decode the bit stream and drive four 16-bit serial DACs, which will update at something like 10-100 kHz and drive both an oscilloscope and a set of MADC channels. Sten and Terry are kludging up a spare AP2 BPM board, to send its gain/offset DACs instead to front-panel LEMO outputs.

Bill Pellico wants remote control/readback of amplifier power supply status, ASAP.

2008-07-31

BOOSTER DAMPER EQUIPMENT RACK LAYOUT

	G11-RR4-1		G11-RR4-2	G11-RR4-3	
1u	BLANK		BLANK	BLANK	1u
2u	CABLE		CABLE	CABLE	2u
3u	EASEMENT		EASEMENT	EASEMENT	3u
4u					4u
5u			MISC	MISC	5u
6u	NIM		EQUIPT	EQUIPT	6u
7u			SHELF	SHELF	7u
8u	CRATE				8u
9u			BLANK	BLANK	9u
10u	BLANK		MISC		10u
11u			EQUIPT		11u
12u	VME		SHELF		12u
13u				PC	13u
14u	CRATE		BLANK		14u
15u				and	15u
16u			MISC		16u
17u	BLANK		EQUIPMENT	MONITOR	17u
18u	T CLK		SHELF		18u
19u					19u
20u			BLANK		20u
21u			DAMPER		21u
22u			CHASSIS		22u
23u			BLANK		23u
24u			BLANK		24u
25u			4		25u
26u			POWER		26u
27u			SUPPLY		27u
28u				DELAY	28u
29u			3		29u
30u	BLANK		POWER	LINE	30u
31u			SUPPLY		31u
32u				CABLE	32u
33u			2		33u
34u			POWER		34u
35u			SUPPLY		35u
36u					36u
37u			1		37u
38u			POWER		38u
39u			SUPPLY		39u
40u					40u

NOTE : 1U = 1.7

John Seraphin made up a nice rack layout and is helping to move toward this configuration, vs. the rat's nest that was there before. Rupe Crouch et al. took away two unneeded CAMAC crates, gave us a new TCLK fanout box, and got us 16 new MADC channels in the adjacent rack. Sten kludged an AP2 BPM board to send four of its gain/offset DACs to front-panel LEMOs, which which I can drive oscilloscope/MADC channels for diagnostic output. I wrote the FPGA firmware to receive a 25 Mbps serial bit stream containing four channels' 10-bit ADC data and to copy the bits over to the four DACs.

Here's an illustration of oscilloscope display (which can be seen on a Booster CATV channel too) and MADC readout (which is capable of high-speed snapshot plotting).

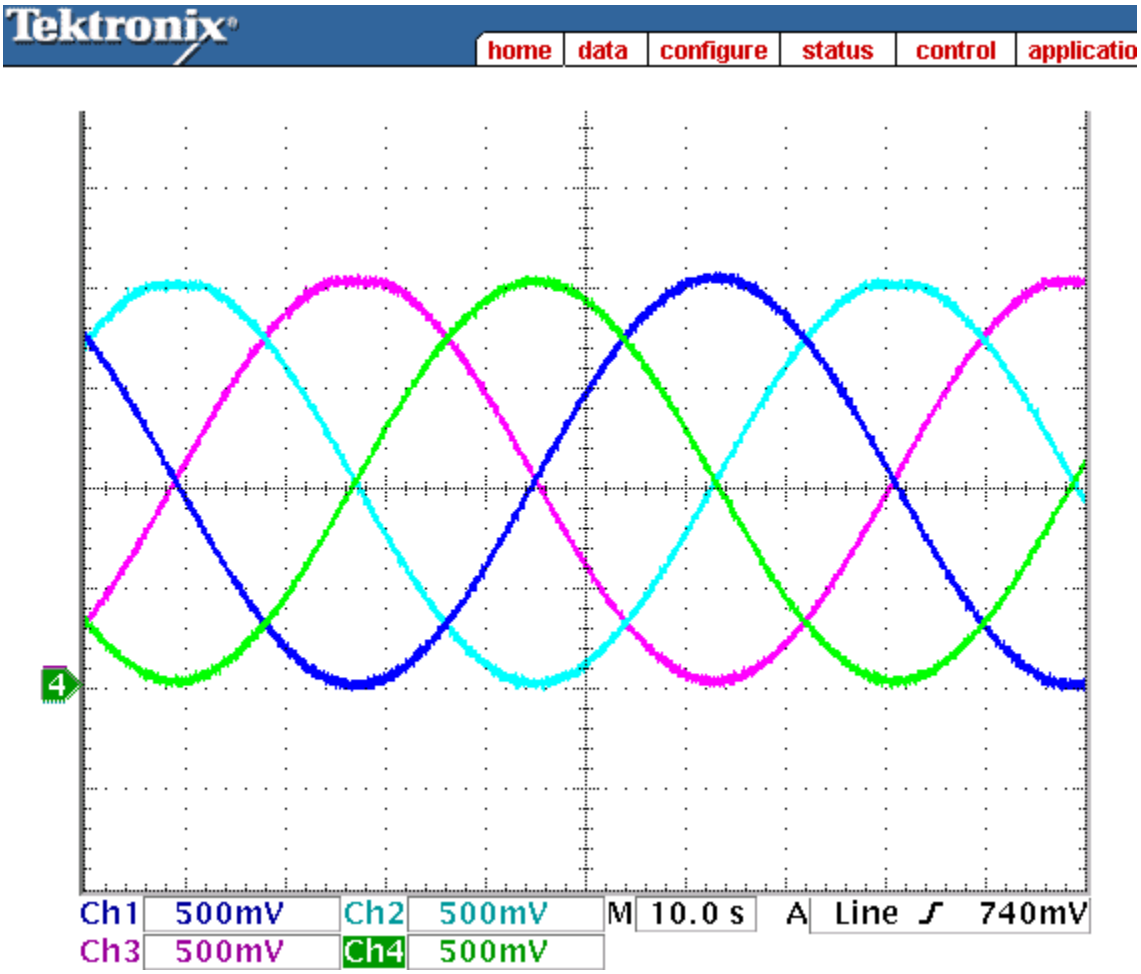
Tektronix
home
data
configure
status
control
applicatio
4
Ch1 500mV
Ch2 500mV
M 10.0 s
A Line 740mV
Ch3 500mV
Ch4 500mV

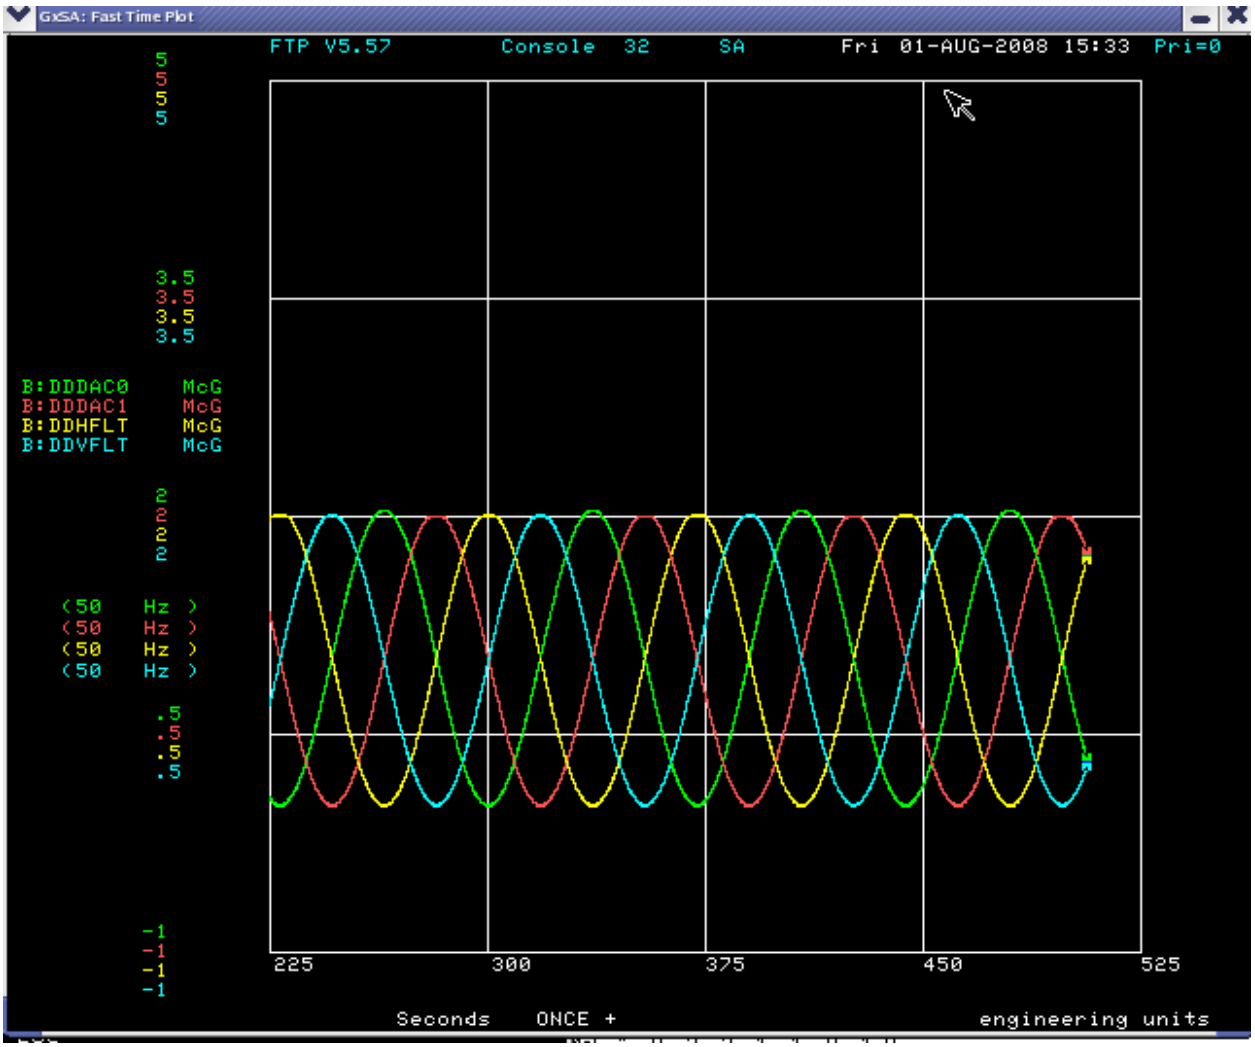

2008-08-05

Haven't written for several days, but there has been considerable progress.

The damper has several spare TTL outputs but doesn't really have analog outputs suitable for monitoring at the rate of a snapshot plot. So after I programmed Sten's AP2 BPM board to receive 40-bit frames of ADC data at 25 Mbps, I programmed the damper board to transmit the data, sending the X and Y filter and X and Y kick signals (summing absolute values over all buckets, for ten turns) at 10-turn intervals for display on oscilloscope and MADC channels. It works! We now have a convenient diagnostic display for tuning the damper in the MCR. Eventually I will create some ACNET devices that allow 2 of the 4

channels to be multiplexed between several sources, e.g. I and Q values for checking cable lengths.

Phil Adamson informed me that Charlie Briegel had set up the software for remote GPIB control (on/off at least) of the 500W MI transverse damper amplifiers, but that it had never been fully installed because the MI damper is always operational. Since we're using the two spares now for the Booster (horizontal plane), he thought it would be a good chance (a) to let Charlie test his work so that it could more easily be used to control the MI amplifiers and (b) to get us the remote control that Bill Pellico wants. Charlie thinks he can get that set up tomorrow. Meanwhile, Nathan tells me that on/off control for the 100W Recycler amplifier (whose spare we borrowed for the Booster vertical plane) can be a simple TTL level, which in the Recycler case is driven by the damper board.

Yesterday afternoon, Bill Pellico showed me how to turn the horizontal chromaticity (or at least the sextupole currents -- not sure what value cancels the natural chromaticity) way down on $17 cycles. The next two plots show (upper) damper on and (lower) damper off, with the SEXTL currents close to zero. (Should add a graph of those currents.) DDHFLT and DDVFLT are the horizontal and vertical filter signals from the damper. CHG0 is the beam current. DDDAC0 is the horizontal kick signal. The filter and kick signals are absolute values accumulated over 10 turns. Display uses snapshot plot from MADC channel, as described above.

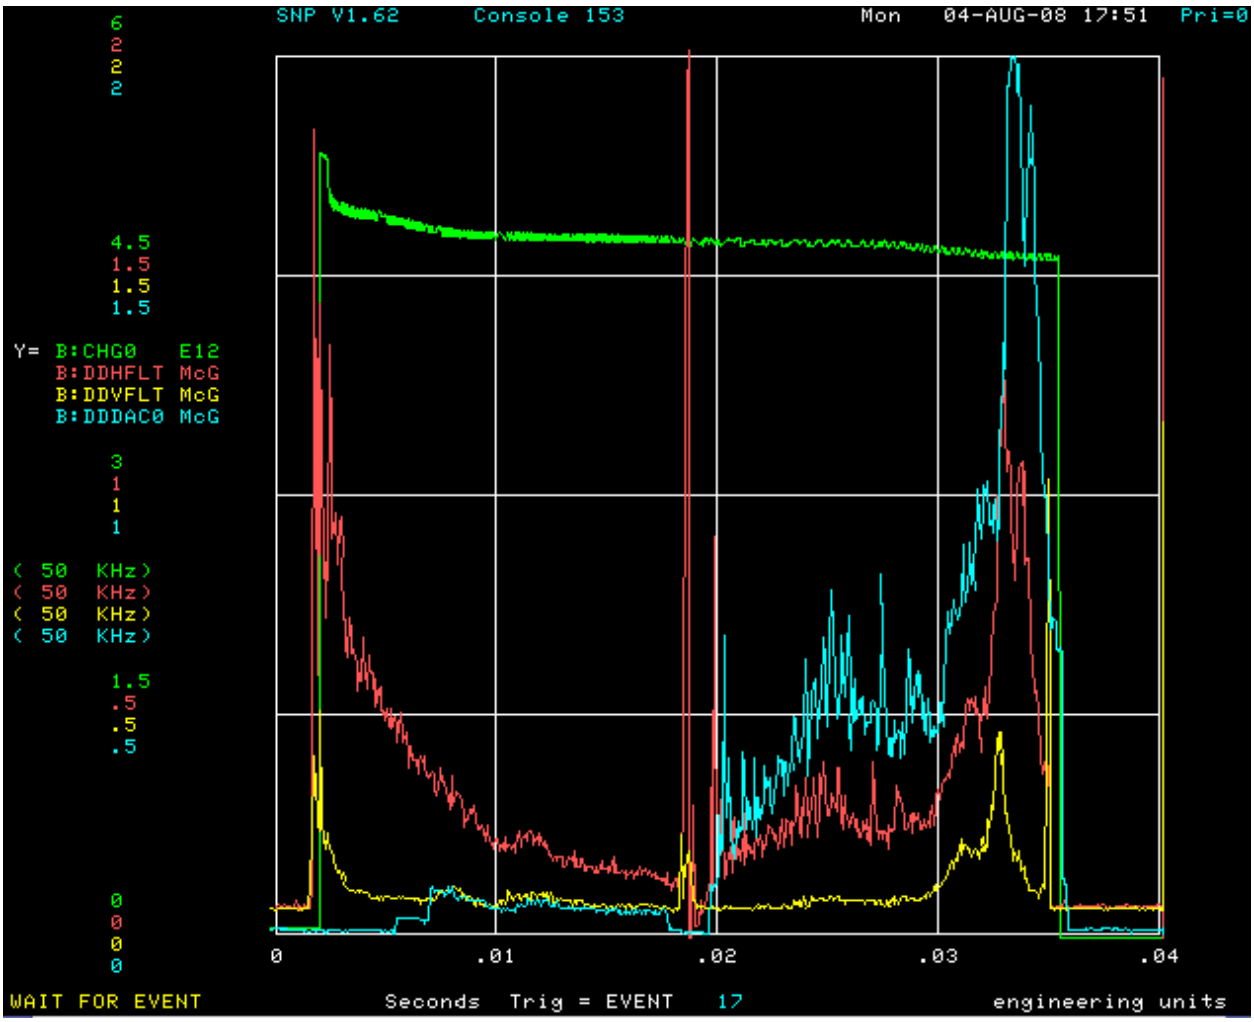
SNP V1.62
Console 153
Mon 04-AUG-08 17:51 Pri=0
6
2
2
2
4.5
1.5
1.5
1.5
Y= B:CHG0 E12
B:DDHFLT McG
B:DDVFLT McG
B:DDDAC0 McG
3
1
1
1
(50 KHz)
(50 KHz)
(50 KHz)
(50 KHz)
1.5
.5
.5
.5
0
0
0
0
0
.01
.02
.03
.04
WAIT FOR EVENT
Seconds Trig = EVENT 17
engineering units

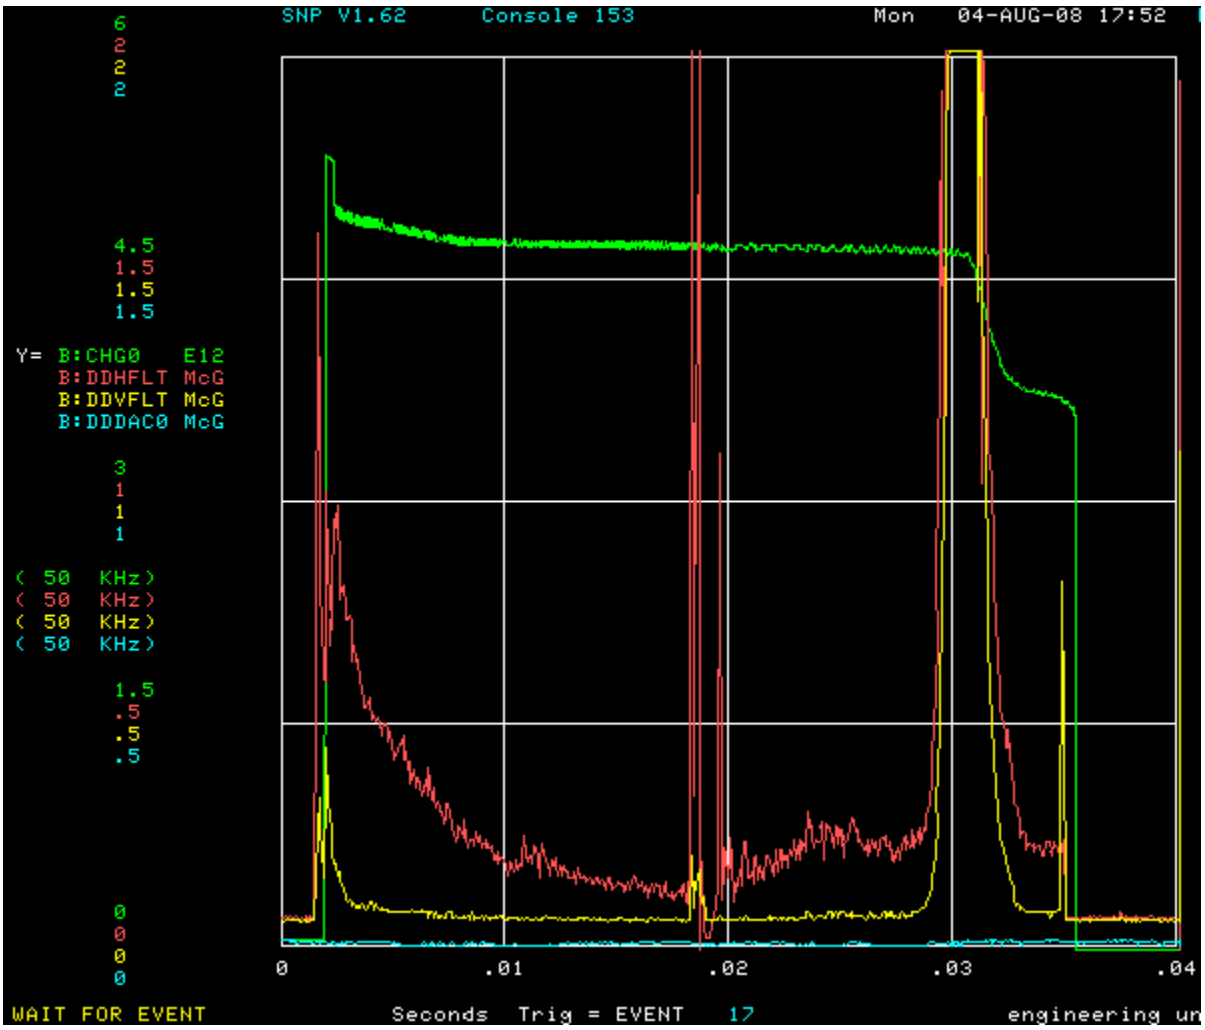

Incidentally, I'm going through and deleting from my files old graphs that I don't think I need anymore, and I found this one. I am still curious what causes the structure in the lower graph (below). I think it is a chromatic effect, with recoherence occurring once per synchrotron period. The horizontal axis is revolution number and the vertical axis is Y[turn n] minus Y[turn n-2], which is a good measure of betatron amplitude.

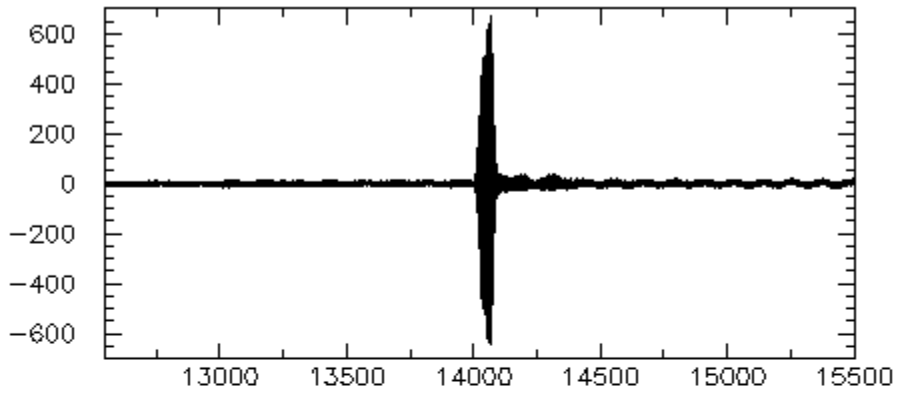

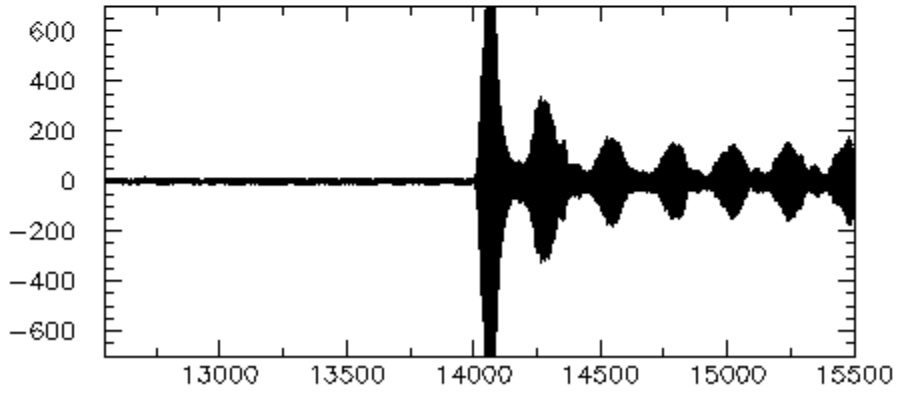

To summarize (and to add a few things I didn't write above), here is what has happened in the past couple of weeks.

- Display/control from MCR:
    - Diagnostic display via CATV-enabled scope and MADC/snapshot exists now and looks quite nice
    -
    - Remote power control for amplifiers should come soon
- Maintainability:
    - Instrumentation department people are rearranging the relay racks to make things neater and more maintainable
    - Jim Crisp has ordered 4 (plus 1 spare) 100W amplifiers, same model as Recycler uses, so that we won't depend on the Main Injector and Recycler spares
    - VME board
        - Jim and Nathan have a summer student who will work on porting the project over to the new VME board

        - We have a VME crate, CPU, etc. all set up (thanks to Dennis Nicklaus) for the VME board
- Other:
    - It looks as if the damper will allow us to run at very low chromaticity!
    -
    - A more up-to-date deep-memory Tek 7000-series oscilloscope is in the pipeline (used/refurbished)
    - We now have a rack-mounted PC over by the dampers, which makes it much easier to work out there

Meanwhile, here are some things I'd like to do before too long

- big picture things
    - Read out FIFO data for all buckets in parts of cycle that consistently show large values of damper error signal, and do fourier analysis to see what the source may be
    - Is it possible to make damper immune from rapid closed-orbit motion?
        - For instance, replace X[N] - X[N-2] with X[N] - 2*X[N-2] + X[N-4] (in fractional tune=0.75 limit) ?
    - Try looking at beam profile upon extraction, for study cycles on which chromaticity is reduced and damper is on
    - Try damper on NuMI cycles
    - Make damper usable for entire Booster cycle, not just above transition
- small picture things
    - Put back multiplicative gain factor (DDXBBG), in series with the gain ramp table that I recently added
    - Put sum signal into third ADC channel, so that in principle we can reconstruct beam position from damper FIFO data, and can factor out the various longitudinal effects that affect the amplitudes we see on pickups
        - This also makes it possible lay the groundwork for eventual longitudinal bunch-by-bunch damper!
    -
    - Multiplex even more diagnostic signals for slow DAC output
    - Replace the sum/difference hybrids I borrowed from Dave Peterson with Brian Fellenz's sum/difference amplifiers
    - Remove excess cable length from ADC and clock inputs
    - Play out a fake betatron oscillation signal on third high-speed DAC channel
        - Makes it possible to test firmware updates without (ugh!) using the beam as a pulser
        - Makes it possible to ensure that digital circuit is doing what we think it is doing, under controlled circumstances

            - Enables open-loop "beam transfer function" measurement, to verify negative feedback near tune lines
        - Think through, finalize, and eventually document the various damper control knobs
            - The knobs that will stay should move from I: to B: device names
- things needed before initial try at operational use
    - Bill Pellico would like to ask Alex Waller to write a damper console application so that the operators have one-stop-shopping for damper knobs and diagnostics
    - Follow through on remote control of amplifiers' on/off state