

Mark, Alex, Hamid,

Attached (Larry's are below) you will find the design review comments from Larry Doolittle, Chris Ziomek and myself. The overall consensus is to proceed with the hardware design while paying close attention to our individual recommendations and suggestions. A theme that ran through all three reviewer's comments (and a recurring discussion of the review board) is the hardware specifications. I think Larry said it best, "There is a project-wide need to finish the flow-down from accelerator specs (amplitude and phase errors) to hardware specifications." We highly recommend that this be completed by the end of January.

Curt

LLRF "Review" December 17, 2002 Larry Doolittle

General comments:

- * Much progress has been made. There's still a lot of work to be done.
- * I see too much emphasis on telling other people what to do, and not enough buckling down and doing the work.
- * I have some concern about over-reliance on complex black-box software. If you can't judge whether the results of a computation make physical sense, using your brain and the back of an envelope, you haven't solved the problem.
- * As I said in the meeting, as a clearly dissenting opinion, the team should standardize on standards, and buy products because they implement the standards well. Standardizing on products is short-sighted, closed-minded, and (apparently) needlessly expensive.
- * There is a project-wide need to finish the flow-down from accelerator specs (amplitude and phase errors) to hardware specifications. I would almost ignore the digital control contribution at this point, other than giving it a non-zero budget. Key items are thermal drift (cables, mixers, amplifiers, ADCs), plus mixer distortion.

Re: SNS LLRF Platform Specifications

1. VXI bus is clearly acceptable, politically expedient, and the choice that has been made. Calling it a "requirement" is out of line.
2. Please clarify "max" and "nominal". How does it relate to full scale

and damage threshold? Are these levels adjustable with a screwdriver, soldering iron, or software?

3. Please separate the specs for the 50 MHz and 402.5/805 MHz input channels. The needs and implementation are too distinct to share.

3a: Bandwidth is wrong -- should be 41 to 59 MHz

3e: -10 dBc output of mixer? -70 dBc output of filter?

4. (Output) overspecified, and I/Q terminology is obsolete. The output stage of a feedback loop is not the place for tight specifications.

5. (Clock) <1ms jitter -> noise on the ADC
>16ms drift -> can be corrected in software

8. This wish-list of memory allocation overconstrains hardware design. Can shoehorn functionality into whatever is available, above some soft minimum in the 20 kByte realm. The 6 kByte in the current LBNL design is arguably inadequate for the long term.

- 8l. If you want a loop delay specification, list it for pure proportional feedback. Fancier application firmware (filters) can slow this down, but that tradeoff is adjustable on-line. Separately, someone needs to investigate other filter possibilities (besides linear phase FIR bandpass) that could do the job with lower group delay.

SNS LLRF Analog Front End Specifications (John Power)

7. (drift) Distinguish between absolute and differential phase drift
Mention amplitude sensitivity of phase measurement (AM/PM conversion)

SNS LLRF Digital Front End Specifications (Matt Stettler)

(was there a paper handout for this, other than the schematic and pin list?)

- Say on paper that the FPGA program comes in via JTAG from the host. The XC2V1500 needs 5,166,240 configuration bits pumped over this link. At approximately 1 bit/usec (using the existing VXI interface), that's 5 seconds of clock time at boot. The LBNL system, which uses bit operations local to the CPU chip, takes approximately 1 second to load 1,040,128 configuration bits.

- Triple check that the mechanical spec on the connectors "allows" multiple connectors on a single daughter card. (Probably OK since they have individual pins.)

- Highly recommend serial number on both DFE and AFE, basically any "FRU" (field replaceable unit) with analog components.
- Power supply: recommend reducing output current requirement, for the first rev at least. Focus on the needs of an XC2V1500.
Can a switching regulator be found that has a synchronization input?
Plan for filtering beyond that provided within the module.
- Abandon support for the XC2V250 - it's smaller than the XC2S150 I use in the current LBNL board (3072 logic cells vs. 3456).
- Missing "system ADC". I see two crucial analog measurements:
VCO control voltage, FPGA core supply current.
- Continue revisit of DAC choice. DAC902 is available in a smaller package than what the LBNL board uses. High resolution is not needed in the output stage of a feedback loop.
- Double check SSO (simultaneous switched outputs) restrictions on FPGA; package and drive strength dependent

LLRF Control System Review 12-17-2002 – C. Ziomek

1. **Specifications** The design specifications need to be top-down where the error budget is allocated throughout the system and the individual designs are specified according to that error budget.
 - a. I recommend separating the specifications for short-term jitter and transient errors from the specifications for long-term drift.
 - b. The comprehensive list of error sources should be evaluated thoroughly. At present, I believe that many error sources have not been adequately specified.
 - c. All designs should be thoroughly validated according to the required specifications before being released for manufacturing or revision.
2. **Analog Front End** In general, the specs look fine. I suggest separating the non-critical signal specs (FWD, REF) from the critical signal specs (IF, CAV). For the critical signals, the following issues must be considered:
 - a. Tight specification of SINAD over signal band of interest – do not over-specify for out-of-band signals
 - b. Mixer IP3 and sensitivity to AM-to-PM
 - c. IF filter return loss at IF and image frequencies
 - d. IF input return loss and interaction with IF filter output return loss (or whatever is on far end of cable)
3. **RF Distribution** The RF distribution is a key part of the LLRF system and must not be ignored because it is not strictly the responsibility of LANL LLRF.
4. **Digital Front End** The design concept looks good and seems to be proceeding. Areas to be careful with include:
 - a. Proper ground, power and signal isolation for low-noise design
 - b. We prefer buried stripline routing surrounded by ground planes and via fences for low-noise analog signals
 - c. We prefer linear regulators for ADC supplies to reduce switching noise mixing into IF signal
 - d. We use differential signaling wherever possible for low-noise high-speed designs
 - e. Be careful to select proper amplifiers to not degrade SINAD
5. **Digital Processing** The FPGA development is the critical path item and could possibly significantly delay the delivery of the system. Also, I believe that the overall system performance will ultimately be limited by the FIR filter. I suggest the following:
 - a. Leverage the work done at BNL as much as possible.

- b. Use a standard set of implementation tools, language, and style guidelines to ensure compatibility and code reuse amongst the many code developers.
- c. Look at alternatives to the FIR design. The 1.6 microsecond group delay is not acceptable. My experience is that a 50 kHz bandwidth is not sufficient. For example, if your klystron ripple is at 20 kHz, you will have very little gain and noise rejection at this frequency. At a minimum, a non-symmetric, non-linear phase FIR filter will cut the delay in half. Having a mode filter in the feedback path does not eliminate the possible excitation of that mode by transients or feedforward. I suggest that a notch or comb filter in the forward path is the preferred solution.

SNS LLRF Hardware Review
December 17, 2002
Curt Hovater

The hardware conceptual design is fairly straight forward and on track. It appears more than adequate for implementing all of the necessary control algorithms especially if the team follows the strategy used in the LBL system. At this stage I would say go for it and build the hardware. Any problems can be worked out on the next rev (it's more important to get some hardware and commissioning under the teams belt).

A concern I (and others) have is the system specifications. We don't have an integrated specification especially as it flows from the golden grail requirements of 1% - 1°. The system/board specifications seem to me good educated guesses for some items and loosely tied back to the overall field control specifications. While it is my belief that the system will be capable of meeting its overall specifications this has not been demonstrated.

Lab interplay also seems to be still rough. The system is being built across laboratory lines and comments such as "that is their problem not ours" still exist. I was surprised that the cavity down converter mixer was not immediately known or included in this review. The reference and distribution were also excluded and seeing that the jitter and drift plays a huge role in this it can't be ignored nor is it fair to the people designing the hardware not to have a known jitter budget (see previous paragraph). The review was almost exclusively dominated by a "LANL view" with out much input from LBL or ORNL. I thought the intention from the get go was to grow from the LBL experience. It is obvious that the hardware design (block down converter, single FPGA) has been guided by this strategy, but it is not clear if the firmware is proceeding along the same path.

One thing that was immediately apparent is that there is really no over all system engineer leading the effort. The three laboratory team leaders Mark, Hamid and Alex do not have the experience or expertise to make the decisions as it applies to the hardware or firmware. This in my mind has hindered the speed of the development especially as it applies to certain design choices. One example where a system engineer could have helped is the FIR filter. Questions have repeatedly been raised about this and better ways suggested, but know one has sat down and looked at other options. Another area is the model simulation and firmware, it is not clear to me that these efforts will converge in the FPGA. I don't think you can implement the model into the system with out some how tying these two together. In addition questions during the review with vague or incomplete answers made it obvious that a lead system engineer is desperately needed.

Overall I have good confidence the hardware team will deliver a system. The schedule does seem optimistic especially since the team (Power and Stettler) are also preoccupied with diagnostics projects. What you do with this hardware is another matter. I am concerned that with out a lead system engineer who understands signal processing and

control systems that the hardware will not be implemented correctly or efficiently. Given that I would start with the LBL firmware (that's what was sold to the review board) and go from there. Finally I recommend a lead system engineer be chosen (assigned) for all hardware and firmware decisions.

Specific Comments

1. Analog Front End: Although I have not seen the board specification (what you send to Bergoz) what was presented Tuesday was some what confusing. Ranges need to be added for SNR (over what dynamic range etc.). Isolations and filter specifications are awkward. I would add an IP3 specification for the various channels (mixers). Some of the specifications may be redundant/conflicting especially as they concern distortion and isolation. An acceptance test plan as well as test procedures for some of the more subtle measurements needs to be made. Are the BPM switches needed? If you have a loop back feature great. It has been my experience that switches lower board isolation. John Power should continue his System View models looking at the various channels.
2. Remote Mixer: While this was not presented it is one of the most important aspects of this design. I would like to see its specification like the AFE. There is still confusion as to where this lives!
3. Analog Output: I suppose that this is an up converter and a switch now??
4. Clock: As has been recommended in the past I would come up with slow/ fast drift numbers and put bandwidths onto them. I think this was done but it was not obvious. I would also put every thing into ps. With as many frequencies as you have using degrees can be confusing. Using ps will normalize everything.
5. Master Reference and Distribution: While this was not explicitly presented it is critical that this be included in the LLRF system. You are only as good as the master reference and some one should look at the jitter/drift as it emanates from the master oscillator out to the LO's at each mixer.
6. Closed Loop Bandwidth: This was only loosely discussed in reference with the system latency. I would like to see a Bode Plot showing the gain and bandwidth needed for field control for the various cavities (nc and sc) under a typical environment and operational parameters. This bandwidth (with contingency) should drive what system latency you must have and therefore influencing the firmware and hardware options.