

# **SVT Upgrade Review Committee Report**

Review of 29 June 2004

## **Run IIb SVT Upgrade Review Committee**

Bill Ashmanskas, Matthew Jones, Jonathan Lewis (Chair), Steve Nahn, Burkart Reisert

### **Summary**

The committee was impressed with the overall plan and commends the SVT group for its diligent work. The plan that was presented is well matched to the physics requirements for the Run 2b era. We strongly endorse the plan to use Pulsar boards with new mezzanine cards to perform the functions of the Associative Memory Sequencer, Hit Buffer, and Track Fitter in the upgraded system. We believe that the PULSAR boards and required mezzanine cards can be delivered quickly enough to meet the SVT upgrade project's needs, provided that mezzanine specification converges quickly and the necessary funds are made available promptly. There is a tremendous advantage to working from an existing and well-tested PCB design, for which an extensive body of CDF expertise exists. We note that in the past year a set of twelve PULSAR boards was successfully assembled, programmed, tested, and commissioned as "Road Warrior" boards in SVT, reducing SVT's execution time by 3 microseconds. This project also familiarized SVT hardware experts with the Pulsar design, and PULSAR experts with SVT issues.

We find there is minimal technical risk in the project. The Pulsar can perform the necessary functions with the addition of mezzanine cards that are will be of a straightforward design.

Our principal concerns are:

- 1) Detailed technical specifications are required for each board
- 2) The project relies on the completion of the challenging AM++ on time
- 3) An complete installation plan is needed
- 4) Maintaining operational efficiency from the present through 2009

### **Technical Specifications**

The bulk of the effort for the project under review consists of writing firmware for FPGAs. Firmware coding, like software coding, proceeds much more rapidly once a complete design is in place. The design should be sufficiently detailed to allow realistic input and output test vectors to be generated for each significant block of the data processing (board, FPGA, RAM, perhaps even individual FPGA subdesign). Typically, once the data flow has been described in detail and shown to be functionally correct, implementation and verification are straightforward (given sufficient hardware resources).

While other approaches may be viable, we believe that the most effective route to design and test the necessary firmware and mezzanine cards is to first implement the new system in simulation. Much of the work on the firmware is planned to be done by people who either have minimal experience writing firmware, are new to SVT or both. The first step of this design effort should be

to implement the emulation software, data structures, and configuration files needed to calculate bit-for-bit the correct flow of data at the input and output of each SVT board for a given CDF event (including Run II data already on tape). Producing this bit-level specification will require the experts at an early stage to discuss and to agree upon the most important details of the lookup tables and algorithms used within each board. This process also ensures that the system as a whole will correctly find tracks, as long as each board and its emulator agree upon the data flow, and that the configuration constants will exist to make the system work on day one. The firmware specification should be reviewed by individuals experienced in writing Pulsar firmware.

The specification of the system contained in the simulation should include the detailed function of the mezzanine cards. Once the details have been worked out, it will be possible to determine if a common mezzanine card can be used or multiple designs will be required. A detailed plan for the new mezzanine cards must be completed, including components, pin allocations, and bit allocations. The schedule for the mezzanine cards should reflect the simulation required and include the possibility of needing multiple designs.

With the board-level behavior agreed upon and coded into the existing Svtsim library, the board experts can focus on programming their boards to meet the specification. Specifically, we recommend that all of the red and yellow boxes on Alex Cerri's schedule graphic be fully implemented by October 2004.

## **Schedule Risk**

We are concerned about possible schedule slip in the event that a second (four-month) fabrication cycle of the AM++ standard-cell chip is needed after a first set of chips is tested in December. How would such a slip affect the ability to commission the Pulsar-based subsystems in a timely way?

We recommend that the SVT group evaluates the extent to which the testing and commissioning of the various Pulsar boards can be done independently. This may allow work to be completed as early as possible, integrating boards (and their experts) into operations.

The Track Fitter appears to be a drop-in replacement for the current device. Because the new AM format is different from the old, backward compatibility is an additional complexity. We encourage the SVT group to evaluate whether it is worthwhile to include backward compatibility in the AMS and/or HB without leading to significant delays in the design of firmware or mezzanine cards. The evaluation should be completed as part of the bit-wise simulation described above.

## **Installation**

It is necessary to develop a detailed installation and commissioning plan early in the project. The 2005 shutdown will not be a long one. Therefore, unless all components are available and tested by the end of July 2005, the upgrade cannot be installed and tested with fake hits and debugged during a shutdown. There are at least two ways to proceed. When all components are ready, it would be possible to split the output of the mergers and send the hits to a crate of new components operating parasitically in order to bring up the new system one crate at a time. Alternatively, during the shutdown, the crates could be reconfigured to allocate two slots each for AMS, HB and TF. Then

the pulsar-based boards could be integrated as a group for all crates. The commissioning plan should include two scenarios: that the AM++ is ready by 2005 shutdown or that it significantly delayed. The decision of which way to proceed should be made no later than May 2005 in order to have the correct hardware and software infrastructure prepared. The plan for installation during Tevatron running should include estimates of downtime both for CDF and SVT needed at each stage.

Because each Pulsar board occupies the space previously occupied by two boards, crate space may be an important issue. How will the final set of boards be physically organized, and how will the existing physical arrangement of boards, cables, and backplane connections evolve into a completed upgrade with minimal downtime? The SVT upgrade group should provide a list of crate and slot allocations for each stage of development and installation.

While we note that test stands exist in Pisa and Chicago for development and board check-out, we are concerned about the adequacy of test-stand resources in B0 during the commissioning period. Problems encountered during test runs in the real system will need to be understood quickly offline in the B0 test area. A realistic plan for crate space, work space, and computing stations (for running tests, diagnosing failures, updating software, compiling firmware, downloading firmware) should account for the number of experts and the number of board variants that will be involved and their anticipated commissioning schedules. It is important to check now whether enough SVT backplanes will be available for test stands and spares.

## **Operations**

The upgrade will require an impressive amount of firmware coding, software coding, testing, installation, and commissioning to occur in the coming year. This effort may be in competition with ongoing SVT operations, for which there is currently one FTE at FNAL, while the stated need (in Alberto's talk) for the existing system is "one or two" FTE at FNAL. If possible, the upgrade work should be used as a means to attract new experts to the SVT project, with a commitment to spend some time learning about and maintaining the current SVT and its upgrade. In order to ensure that the SVT can be maintained through 2009, it is crucial that all hardware, firmware, and software be properly documented and archived. The SVT upgrade group should describe the long-term maintenance needs including manpower and test stands. Those needs should be included in updated MOUs.

The 28-board purchase is probably insufficient. We recommend having at least two available spares of each firmware type in addition to those used for test stands at Fermilab, Pisa and Chicago. If spares are to be common with Level 2, an integrated spares plan should be worked out before the new purchase. The number of spares would have to be a little larger if the AMS uses the P3 modification for backward compatibility. Also, the installation plan should include when the existing RW Pulsars are available for the upgrade. That will factor into the number of boards that need to be purchased.

## **Appendix A: Individual concerns from committee members, comments on documentation, etc.**

### **Bill Ashmanskas**

I have a minor comment about Alberto's comments on the timing measurements. He stated that "the timing measurement is underestimated due to pipelining." The SVT processing time for a given event clearly ends when that event's end-event word is sent to L2. But when does it begin? If one begins at the corresponding L1A, then the SVT time will be overestimated on events for which a previous event's SVX readout (or its SVT processing) is still active when L1A occurs. To avoid this overestimate, the SVT processing time measurement starts at the given event's L1A time or the time at which SVT has finished processing the previous event--whichever comes later. It would be trivial to add a bit to the SVT timing measurement to indicate whether or not SVT was idle at the time the given event's L1A was issued. Requiring SVT to have been idle will provide an unambiguous set of timing measurements.

The physics case for the "hadronic B" trigger is already compelling: three of CDF-II's first four published papers required this trigger. The stated case for the high-pT b jet triggers would benefit from studies of Run II data e.g. quantifying new-physics searches' sensitivities with and without SVT or at least showing that these triggers enhance CDF's ability to record b-tagged jets in conjunction with other trigger objects. Likewise, the case for preserving muon trigger efficiency would benefit from an illustrative graph based on a study of Run II data.

### **Steve Nahn**

I concur with what I took as general consensus that the hardware procurement end of the bargain can be done within the time constraints, although I would not be surprised if the installation were not complete on time. Given that the new HB and AMS/RW require the AM++ bitmap addressing scheme, the actual testing of the real system may be held up by the availability of the AM board. I am skeptical about the utility of putting a Pulsar in for the HB with firmware using the old addressing scheme, since I think the real work here is making the firmware work, so you won't learn too much about the final system by taking this intermediate step. Given that the AM Chips can be modeled by FPGAs just not with so many patterns, could there be a "Pulsar as AM" solution to provide testing for the HB and AMS/RW which runs at speed but doesn't give as many patterns?

Regarding performance it looks like it will do the job. I would still like to see the timing evaluations decoupled from SVXII time and not based on  $BESTATE == 0$  measurements, but that's just me. I'm not totally convinced that 4 L1A in a row can be factorized by extrapolation, but I think what they showed is good enough. When you go to 3E32, the peak of the timing distribution shifts considerably, presumably just due to the higher occupancy in SVXII, but it is unclear how much the peak plays a role in deadtime versus the tails (which are really what the upgrade addresses, according to the plots).

### **Burkard Reiser**

On all sketches of Implementations on Pulsar boards (Fig 10, Fig. 12 and fig 14) I find bidirectional connections between Control and DataIO FPGAs. Only in one case (Fig 12) a bandwidth of 43 bits

is specified, I tried counting the connections I come up with 42 (but I maybe miss some lines which could be reused in a clever way?). If they should be bidirectional how is the traffic on the board handled? As the hardware already exists, and we advertise that no changes are needed we need to be very careful here to not getting sidetracked on some firmware development which then later on turns out can't be implemented and needs major rework to make it fit.

I think there have to be reviews of the firmware implementation specifications by people how have experience on the board (design, implementation of firmware, board checkout). For example I am confused by the TF implementation as described on page 22. The SVT input is available to all 3 FPGAs, but from how the implementation is described I got the impression the Control FPGA receives the SVT inputs and sends it to the Fit-1(2) FPGAs.

## **Appendix B: Charge to the Committee**

### **Charge to Run 2B SVT Review Committee**

Review date June 29, 2004

Version 1.0

#### **Introduction**

A significant component of the original baseline for the CDF Run 2B DAQ and Trigger upgrades was a project upgrade the Silicon Vertex Tracker (SVT). This project was required to handle changes in the geometry of the Run 2B silicon detector. With the cancellation of the Silicon upgrade, the primary motivation of the SVT upgrade ended.

Recognizing that improvements in the SVT processing time could be critical to improving the L1 bandwidth, an SVT project was retained in the DAQ and Trigger upgrades but no scope was defined. The SVT group has now defined the scope of the SVT upgrade, which is described in CDF note 7064. Key components to the upgrade include new Associative Memory (AM) chips, Associative Memory boards and AM plug boards. This new AM system will provide many more patterns than the current design. The AM upgrades are in the prototyping phase with full funding already provided by INFN. The upgrade proposed in CDF7064 adds replacement Hit Buffer (HB), Associative Memory Sequencer (AMS) and Track Fitter (TF) boards to the AM upgrade. All of these boards would be implemented using Pulsar boards designed for the L2 decision crate. The bulk of the SVT upgrade effort will be in designing a new mezzanine card for the Pulsar and writing firmware for each of the board types. The schedule for this work is very tight since it must fit within the requirements of the Run 2B project. Delivery of the hardware to CDF must be (production and checkout completed) by the end of FY2005 (Sept 30, 2005).

The primary mission of the SVT upgrade is the same as the L2 Decision crate upgrade, maintain if not improve upon the current L1 Trigger bandwidth ( $\sim 20\text{kHz}$ ) as the luminosity increases to  $3 \times 10^{32} \text{ cm}^{-2}\text{-sec}^{-1}$  while maintaining deadtime below 10%. All of this is to be achieved while maintaining L2 rejection power. The SVT upgrade will also provide additional handles for tracking in both the Central and Forward regions that may be negatively impacted by high occupancies in the inner COT layers and/or changes in XF T configuration.

#### **Charge**

The charge for this committee is to evaluate the proposed SVT upgrade design and schedule to determine if the system will provide the necessary improvements in performance within the time constraints of the Run 2B project. The goal of the review is to determine the actual scope SVT upgrade project. It is crucial that there is a high confidence in the ability of the SVT upgrade group to design, build, and test upgrades to the system within the constraints of the Run 2B project and with minimal interruption to ongoing CDF operations.

The primary scope of the review is to evaluate the portions of the system to be funded by Run 2B DOE project funds. These include:

1. Production of Pulsar boards for use as Hit Buffer, Associative Memory Sequencer and Track Fitter boards
2. Design of associated firmware

### 3. Design, fabrication and assembly of new Pulsar mezzanine boards

Although it is not covered by Run 2B funds, the committee should also evaluate the plans for the new AM boards and how that upgrade impacts those parts covered by the project.

In particular, the committee should consider the following specific points:

1. Evaluate the projected performance of the baseline design against the performance benchmark of maintaining or improving the current L1 bandwidth at luminosities of  $2-3 \times 10^{32}$ . Will the proposed design perform maintain or improve the quality of the tracks provided to the L2 Decision crate?
  - a. If there is not sufficient information to determine the performance of the upgraded system, are there further tests and/or studies that should be performed to provide that information?
2. Is the scope of the proposed design appropriate for the needs of the experiment?
  - a. What requirements, if any are placed on the Hit Buffer, Associative Memory Sequencer and Track Fitter by the AM upgrade?
  - b. Are there additional components of the system that should be considered for inclusion in the upgrade?
  - c. To what extent can the individual components of the upgrade be built, tested and installed independently? What incremental improvements can be achieved through partial or staged implementation? For example, could an upgraded Track Fitter be successfully utilized prior to completion of the AM upgrade if it is ready first?
3. Is the resource loaded schedule for design and construction of the SVT upgrade credible?
  - a. Does the plan have a complete list of tasks needed to build, test and commission the system? This should include hardware, firmware and software.
  - b. Are there sufficient labor resources within the existing groups to complete the project on schedule?
  - c. Identify any specific areas where additional manpower is necessary or highly desirable. What type of manpower is needed (physicist, engineer, technician...)?
4. What are the risks to ongoing CDF operations from testing, installation and commissioning efforts for the upgraded SVT system?
  - a. Does the plan include consideration of needs for testing with the detector and the disruption to operations that will result from commissioning?
  - b. How do these compare to the risks of continuing to operate with the existing system or a less ambitious upgrade?
  - c. Will the manpower needed to operate and maintain the existing system be diverted to building the upgrade?
  - d. Will the upgrade train new people to operate the SVT for the remainder of Run 2?

Given the tight schedule for the upcoming Directors Review, it is requested that the committee provide a written executive summary of their findings and recommendations to the Run 2B project management within five days of the review. If committee members have more detailed recommendations these should be reported within 2-3 weeks of the review.

## Appendix C: Review Agenda

- Introduction - Pat Lukens
- System overview and current performance - Alberto Annovi
- Physics motivation for upgrade - Mel Shochet
- Description of AM upgrade (background, not reviewed) - Paola Gianetti
- Associated Memory Sequencer/Road Warrior - Franco Spinella
- Hit Buffer - Taka Maruyama
- Track Fitter - Mel Shochet
- Software - Roberto Carosi
- Timing studies - Ivan Furic
- Overall Project Plan - Mel Shochet
- Questions and Discussion
- Executive Session

Copies of all talks are available through the archived WebTalks page:

[http://fcdfwww.fnal.gov/internal/WebTalks/Archive/0406/040629\\_run\\_iib\\_svt\\_upgrade\\_review/](http://fcdfwww.fnal.gov/internal/WebTalks/Archive/0406/040629_run_iib_svt_upgrade_review/)