April 2, 1999

ROD Review

March 25-26, 1999
Santa Cruz Institute of Particle Physics
University of California Santa Cruz

Ad Hoc ROD Review Committee: J. Beringer, G. Darbo, K. Einsweiler, A. Gara, A. Grillo, P.T. Keener, M. Levi, J. Parsons, L. Premisler, L. Saphoznikov, A. Seiden, R. Wastie

An Ad Hoc group, with membership listed above, was convened on March 25th and March 26th to provide a recommendation on the choice of architecture for the SCT and Pixel ROD. The group heard presentations on the status of two alternative architectures: one based entirely on DSP's and the other on FPGA's in the data stream and DSP's for the other tasks such as monitoring and calibration. In both cases an impressive amount of work had been accomplished with much still to do to provide a real design. The committee commends the system leader, Andy Lankford, and the design teams on their excellent work. It is clear that the full set of physicists, engineers, technicians, and students now engaged in the project will be necessary to produce a successful ROD which serves the needs of the SCT and Pixel detectors of ATLAS and is completed in a timely way.

The Ad Hoc committee believes that both architectures could likely be made to work. However, in making a final recommendation, which was unanimous, the committee has chosen the architecture which was felt to entail the smallest risk. We believe this is the second of the two which uses FPGA's in the data stream and DSP's for the other tasks. This scheme should allow a natural sharing of the design tasks since the two groups have primarily focused on FPGA's and DSP's, respectively, during their recent design efforts. Much creative work is left to complete the design and test it with detectors and a real DAQ system. The data rates which can be handled are uncomfortably close to the edge and enough safety has to be ensured to deal with unusual conditions combined with the very high rates. As part of the design a significant amount of software will also be required, and its development will be an ongoing task as the ROD is used and its monitoring functions defined in real conditions. We discuss below some of the specific technical issues that underlie our recommendation.

Technical Comments

The functions of the ROD can be divided into four different tasks: data flow, FE electronics configuration, calibration, and monitoring. In the FPGA based design the data flow task is decoupled from the other three tasks, while in the DSP design it is not. In the case of the DSP ROD the monitoring task will suffer in case of increased occupancy, and each change of the monitoring program may have an impact on the data acquisition performance. The DSP based ROD is faster for calibration, while the FPGA based ROD is better suited for the monitoring task.

The committee feels that both architectures will be able to fulfill the basic requirements of monitoring and calibration. Both designs have their advantages and disadvantages. Because of the decoupling of the data flow tasks from the other tasks, we consider the FPGA design as the safer one. In the proposed FPGA architecture, the DSP functionality now added for calibration, histogramming, etc. is very desirable. While the FPGA development is quite advanced, the development of the DSP part seems, by comparison, at a very early stage. There are many non-

trivial engineering details to work out in order to integrate the DSPs into the ROD properly. In this area, the expertise of UCI is critically needed.

We believe that both proposed architectures can be made to meet the present ROD data flow requirements, but there is not an enormous amount of headroom in either. The data path in the proposed FPGA architecture operates at full input link bandwidth and does not mix the data transfer and monitoring functions as in the proposed DSP architecture. We believe this offers less risk for the FPGA proposal to handle unexpected conditions such as an extra noisy link. The DSP architecture operating at 100 kHz trigger rate does not have much "free time", which reduces the headroom to deal with problems by adding more code.

The DSP proposal assumes a 250 MHz DSP in order to meet the performance marks presented. While this is technically feasible with the newly available DSP, it requires the design of a 125 MHz memory bus on the PCB. This represents a difficult design challenge and may require more than one iteration to achieve. As for technical risk of the FPGA solution, the utilization of the particular FPGA proposed for the Decoder appears to be quite tight. We recommend that the design move now to the larger FPGA in order to provide more flexibility to cope with unforeseen problems.

The committee feels that it is important that the ROD designers carefully consider the current and possible future need to access the various pieces of information available in the ROD, e.g. S-Link status, buffer occupancy information, front end status, etc. As much of this information as possible needs to be made available to the ROD controller. In addition, connections should be made between unused pins on the FPGAs to mitigate the risk of unforeseen requirements. These provisions are important both to handle unforeseen circumstances and to ease the debugging of the board.

Software Engineering

While the committee recognizes that much effort has gone into trying to carefully design the software for the ROD, it feels strongly that the software should be engineered as fully and as carefully as the hardware, given its importance to the overall ROD design. By software, we mean both the DSP program running in the ROD as well as the interfaces and low level primitives used to access the ROD from the ROD Crate Controller.

This software needs to be carefully designed and implemented and needs to be reviewed at each phase. The software needs to be managed and versioned not only for released code, but also during development. Documentation must also be maintained throughout the process. This will help complete the ROD on schedule and will ease maintenance.

Some effort must be put into understanding production versioning control to mitigate the issues involved in understanding which code is being used at any given time. These versioning control issues are also applicable to the FPGA programming bit streams.

Costs/Schedule

The committee reviewed costs of both the DSP and FPGA approaches. The material costs of the FPGA solution (without spares) is $842 K, the materials cost of the DSP solution, without spares, is $926 K. Both are significantly below the $ 1.5 M baseline materials cost. Consequently, it was

not felt by the committee that cost could be used as a selector between the two approaches, and it had no bearing on the final decision. Neither design presented figures for EDIA or labor and therefore this was not reviewed. Given the potential for technical risk in the DSP approach there would likely be some upside cost risk to this approach. Of greater concern to the committee was the significant time lost from the schedule in coming to conclusion for the choice of architecture.

The first version of the ROD was to be delivered by March 1999, clearly this is no longer feasible. There still exists sufficient time to deliver the first articles of the production ROD, for the important SCT production barrel tests by September 2001 with prototypes made available to SCT module test by February 2000.

However the design must now proceed very expeditiously towards these important deadlines. Furthermore, the creation of separate versions of ROD99's and prototypes may not be practical and the group must work with ATLAS management and the subsystems on a revised schedule. The lack of time remaining highlights the need to settle on an architecture that minimizes the overall technical and design risk. The ROD engineering team will need to work closely with the Atlas SCT, Pixel, and DAQ communities to finalize the requirements for the ROD now and should not view ROD99 or the prototype as a means to discover new architectural requirements.

Resources

In accordance with the presented milestone chart, the UW and UCI groups identified seven full time equivalents (FTE) to complete the remaining R & D, detailed design/architecture tasks, up to and including Module "0." The available resources that can be committed are:

	UW (FTE)	**UCI (FTE)**
Physicists	1	2.15
EE	2.35	1.9

***Note**: Technicians and students are not considered and can always be added when necessary.

Only together, can the groups meet the stated manpower required. However, the review committee suggests that the manpower estimate needed for completion of the work is understated and should be reviewed again. It is essential for the success of the experiment that the groups continue to cooperate and complete the tasks through commissioning of ATLAS.

For the chosen approach, the committee suggests that DSP architecture/design tasks be performed by UCI and the FPGA tasks be undertaken by UW because of their respective expertise in these areas.