

DEVELOPING AND COMMISSIONING BABAR ELECTRONICS

A.J. Lankford, Department of Physics and Astronomy, University of California,
Irvine, CA 92697-4575 (email: lankford@lankford.ps.uci.edu)

for the *BABAR* Collaboration

Abstract

Many aspects of the architecture and performance requirements of the electronics system for the *BABAR* Experiment are similar to those for the much larger LHC experiments. We briefly describe the requirements and architecture of *BABAR* Electronics, focusing on aspects that are similar to the LHC. We then discuss the experience of developing the system from design to operation, including such topics as prototype and system tests, manufacturing, grounding and shielding, integration, and initial performance. We focus primarily on front-end electronics, including IC development; however, we also discuss data acquisition.

1. INTRODUCTION

This rather informal paper is a recollection of the experience of developing and commissioning the *BABAR* Electronics System. Hopefully, this account can serve as a reminder to LHC electronics developers of some of the issues, concerns, and pitfalls to remember during the development process. This paper is not intended as a description of the *BABAR* Electronics System. A brief introduction will point out how the *BABAR* Electronics System resembles electronics systems under construction for the LHC experiments, despite the fact that it differs in scale. However, neither the common aspects of the system nor the detector-specific electronics will be described. Summary descriptions of these aspects can be found in the paper that I presented at the 1998 LHC Electronics Board Workshop [1] and in the papers to which it refers, and in references [2-5].


Furthermore, this paper is by no means a complete guide to the "Do's" and "Don'ts" of electronics development. It is anecdotal, relating personal impressions and the comments of people involved in the *BABAR* project. It recounts some of the problems and setbacks encountered during development of *BABAR* Electronics; however, it does not recount all. It also includes lessons learned by the individuals during prior projects that were felt to have been successfully applied to *BABAR*.

Several of the other leaders of the *BABAR* electronics development contributed to the preparation of this paper. I have frequently included their comments verbatim in the text. In these cases, I have chosen to include their

comments as quotations without identifying the specific contributor since these comments were often informal. The contributors of the quotations are identified in the Acknowledgements.

2. THE BABAR EXPERIMENT

The *BABAR* Experiment will study CP violation at the SLAC B-Factory. Its detector consists of five major systems: a silicon vertex tracker [6], a cylindrical drift chamber with dE/dx capability [7], a particle identification system (DIRC) based on imaging of Cerenkov rings [8,9], a cesium-iodide crystal calorimeter [10], and a muon identification system (IFR) based on resistive plate chambers [11]. The specialised requirements of each detector system are addressed by front-end electronics customised to the detector technology but integrated into a uniform data acquisition architecture.



3. BABAR ELECTRONICS OVERVIEW

In order to address the requirements of its detector and operating environment, *BABAR* has designed an electronics, trigger, and data acquisition architecture that is quite similar to architectures being designed for the LHC. For instance, the *BABAR* architecture is multilevel, pipelined, and nearly deadtime free. It employs detector-specific custom ICs to realise the full performance of detector systems. Front-end electronics is detector-mounted. It simultaneously digitises analogue signals, writes data to buffers, and is read out to the data acquisition system. Although analogue front-ends and

digital readout circuits are located in close proximity, full analogue performance has been realised without digital noise. The trigger system has two levels, with the option for an additional intermediate level. The first level trigger is based on pipelined hardware processors and utilises tracking and calorimeter data. The higher level trigger is based on an online event-processing farm of commercial processors and is integrated with a network-based event builder into the data acquisition system. The online system also includes a prompt reconstruction phase that performs full “offline” processing of the data.

The *BABAR* design incorporates an unusual level of standardisation across detector-specific subsystems. Front-end electronics is customised to detector-specific needs; however, front-end buffer architecture and front-end links and protocols, for both readout and control, are standard. All off-detector electronics is standardised. The *BABAR* Readout Module provides standard interfaces to timing, data acquisition, and detector controls systems. It also provides a standard platform and code framework for detector-specific pre-processing and calibration code. Data acquisition crates, fast control and timing modules, and much detector controls hardware are also standard across *BABAR* systems.

All front-end electronics subsystems of *BABAR* share a common architecture. Each front-end chain consists of an amplifier, digitisation (discriminator or flash ADC), a circular buffer (to store data during level 1 latency), and a derandomizing buffer (to store data between the level 1 accept and transfer to a Readout Module). Analogue signal processing, digitisation, and data readout occur simultaneously. All front-end subsystems except the IFR provide data sparsification. All level 1 latency buffers are digital; hence, it is possible to store data longer than analogue buffers would allow. The buffers of all front-end systems are managed by a common protocol. All level 1 latency buffers function as pipelines of fixed length, and all derandomizing buffers function as FIFOs which are capable of storing a fixed number of events, regardless of the actual implementation of the buffers. For instance, in the readout IC of the vertex detector [7, 8], the derandomizing buffer holds four events in a series of a registers, two buffers, and the sparsification logic; whereas, in the drift chamber digitizer IC [9, 10], the derandomizing buffer is an SRAM that stores waveforms for four events.

Each front-end subsystem also shares standard *BABAR* interfaces to the detector-spanning, or common, electronics subsystems. In all cases, the front-end electronics is mounted directly on the detector for performance reasons. This solution also substantially reduces required cable plant. The design of each detector-specific subsystem balances its individual and common requirements in order to achieve a cost-effective, robust, and easy to maintain implementation. Custom integrated circuit solutions have been adopted for

four of the five detector subsystems. Seven custom integrated circuits were developed.

3.1 Catalogue of *BABAR* custom ICs

For the silicon vertex detector, a 128-channel amplifier/shaper/discriminator IC (ATOM) [2,12-15] was developed by Pavia, LBNL, and Santa Cruz. It features a programmable shaping time and a 4-bit, time-over-threshold pulse height measurement. It contains a level 1 latency buffer, sparsification, derandomizing buffer, readout control logic, and a 60Mbps serial output. It was fabricated in Honeywell rad-hard 0.8 μ m CMOS.

For the drift chamber electronics [3], a four-channel amplifier/shaper IC (DCAC) with analogue and discriminated outputs was developed by Santa Cruz in Maxim CPi semi-custom bipolar. In addition, an 8-channel drift chamber digitizer chip (ELEFANT) [16-18] was developed by LBNL. It features a TDC with 1ns bins and a 15MHz, bilinear, 6-bit flash ADC for each channel. It also includes a level 1 latency buffer, sparsification, and derandomizing buffer. It was fabricated in Hewlett Packard triple-metal 0.8 μ m CMOS.

For the DIRC electronics [4], an eight-channel, low-noise, amplifier for phototube signals with zero-crossing discriminator and a multiplexed analogue output [19] was developed by Orsay in AMS 1.2 μ m CMOS. A 16-channel DIRC Digital Chip [20] containing TDCs with better than 250ps linearity was developed by LPNHE Paris 6-7. It also contains a programmable level 1 latency buffer, sparsification, and derandomizing buffer. It was fabricated in Atmel-ES2 double-metal 0.8 μ m CMOS.

For the calorimeter, a single-channel low-noise amplifier chip was developed by SLAC. It contains only one channel for reliability/redundancy reasons. It also provides a split-range output and a cable driver. A 4-channel Calorimeter Auto-Ranging & Encoding chip (CARE) [5] was also developed by SLAC. It provides a sample & hold on four amplitude scales and encodes the index of the scale of interest. Both calorimeter chips were fabricated in AMS 1.2 μ m CMOS.

4. ORGANIZATION

4.1 Organisation of the Electronics System

The *BABAR* Electronics System incorporated all electronics related activities into a single comprehensive electronics working group. The Electronics System included front-end electronics, data acquisition, trigger, including the level 3 trigger, electronics for detector controls, and electronics infrastructure. The only electronics components that were not provided by this overall working group were the detector-mounted components (readout IC and hybrid) of the vertex detector, although the overall electronics organisation was involved in reviews of these components. In general,

the same working group also provided all software directly associated with electronics. For instance, the software of the data acquisition system, of the level 3 triggers, and for electronics initialisation, diagnostics, and maintenance was provided by the electronics working group. Of course, the overall working group included subgroups working on subprojects, such as drift chamber electronics, and these subgroups further split into smaller subgroups as appropriate.

The purpose of organising all the electronics in a single detector system working group was to facilitate co-ordination and to foster common solutions. As commented by a leader of one of the electronics subsystems, "The grouping at the design and realisation phase of all electronics in one electronics group proved very effective to assure commonality and uniformity. The caveat here is to maintain enough relationship with the subdetector group, but this is not very hard to achieve." We found that the bond between individual electronics subsystems and their respective detector systems was sufficiently strong that it was not necessary to explicitly place the electronics subsystem group within the detector system organisation. We found that our organisation bound together, in a closer way than would otherwise have been possible, all the electronics subsystems, including the trigger and data acquisition systems. This organisation worked quite well throughout the construction project. It contributed to a coherent overall system design, including shared design solutions among electronics subsystems. It contributed to the constructive attitude of design reviews performed by peers. It resulted in subsystems helping each other during design and during implementation. It also afforded the project leaders flexibility to redirect resources and effort when necessary.

4.2 Working with multiple institutions

When a working group consists of a large number of institutions, there is need for good communications and frequent contact. The *BABAR* Electronics working group included about 28 institutions from five nations. One of the electronics subsystem leaders commented, "Developing electronics with the involvement of several labs on two continents requires very good communication between teams. Thanks to the Web and videoconferences, it is now easy to exchange quickly any kind of documents. However, at the very beginning and at the end obviously, nothing can replace meetings and work *in-situ*."

Even with the Web and video conferencing, significant travel is needed as part of the communications. Participants must travel to meetings regularly, although not all participants from each institution need to travel to each meeting. During the design phase, *BABAR* held meetings of the overall working group about six times per

year, at four collaboration meetings and at two additional electronics workshops. These meetings were held both in the U.S. and in Europe, with at least two per year in Europe.

The electronics co-ordinators, particularly at the level of the "system manager" and "system engineer", must also perform site visits to demonstrate their interest, to foster close co-operation, and to meet the participants at institutions that they might not meet through meetings. Site visits of course also offer the opportunity to see first hand the facilities at participating institutions and to see the work being performed.

In multi-institutional projects, design reviews are even more important than they otherwise would be. This importance is because they help compensate for difficult communications. For instance, they provide a critical opportunity to carefully examine interface definitions and implementations of interfaces.

By paying attention to the need to communicate, that is by arranging frequent meetings and providing mechanisms for communications between meetings, it is possible to draw upon the strengths of many institutions without compromising the integrity of the electronics design. It is important to have at least one very qualified person at each institution in order to make certain that work is on track and compatible with the overall system design. The importance of communicating between meetings should be stressed. The timely completion of the project requires that questions and open issues not be left unanswered or unresolved until the next meeting. Email, telephone, or special meetings should be used to resolve issues promptly as they are identified and as resolution is needed.

5. REQUIREMENTS

Defining requirements at the outset of the project is essential. Requirements provide the yardstick against which performance of prototype and production systems is measured. Without this metric, the collaboration risks that the system under development will not meet the performance requirements of the experiment. There are also the risks that the design effort will not converge as the design team strives for unnecessary performance or that the design solution will be unnecessarily complex, causing undue technical risk.

Requirements should be realistic. Overly stringent requirements can cause the same risks relating to unnecessary performance and complexity mentioned above. Proper discourse between physicists and engineers during the definition of requirements, and ultimately throughout the design and prototyping phase, can help tremendously in achieving a set of requirements that are both adequate and practical.

Achieving appropriate design solutions requires a good knowledge of the problem being addressed, including an understanding of the context of the problem. This

knowledge and understanding can be more difficult to achieve in the multi-institutional efforts of large experiments. As observed by one of our electronics subsystem leaders, "With the growing size of the experiments, it is more and more important that people work with a sufficient knowledge of the context in the experiment, where a piece of circuitry is sitting, how it is intended to be used, which issues can be expected if some requirements are not met." An appropriate set of requirements properly defines the problem to be addressed. Documentation explaining the requirements can provide an understanding of the context.

When defining requirements, one should keep in mind that "Simple is beautiful!" Again, collaborative projects exacerbate the tendency to define requirements or design solutions that are unnecessarily complex. As observed by one of our electronics subsystem leaders, "Everyone can add features that delay and sometimes compromise the main goals. Unnecessary complexity is often the result of too many people interacting. This is a drawback of easy communication." The tendency towards complexity should be properly checked. Why take undue technical risk? That risk leads ultimately to cost risk, schedule risk, or both.

Requirements should be documented. Documentation is necessary in order that the requirements are not forgotten during the design process. It also makes the requirements available to newcomers on the project. In addition, the motivation or justification for each requirement should be documented. Documenting the justification enables designers to understand and recall the origin of any requirement, frequently enabling a designer to question a requirement that is proving difficult to satisfy. Documenting the justification for requirements also provides a method of checking that requirements are still valid as time passes and a project evolves, in case the underlying assumptions of some key requirement change.

6. SYSTEM DESIGN

A coherent system design is the central component of an adequate technical solution. The system design should be developed first, before developing detailed solutions to aspects of the overall problem. The system design should be developed from the "top down" from the requirements that should be addressed and from a clear vision of a solution that addresses the entire scope of the task. "Bottoms up" designs, which derive from a solution to only a part of the task or that originate in a desire to employ a particular technical solution, should be avoided. Although it seems obvious to say that the first step in design should be a top-down system design, in practice, design often does not happen that way.

For *BABAR*, we felt that it was important to develop a system design with unified solutions for the mechanisms that control the front-end electronics, for data acquisition

for all front-end electronics subsystems, for monitoring of electronics, and for slow control and monitoring of detector systems. Such common solutions simplify the overall design and greatly ease integration.

When establishing common solutions, it is important to gain acceptance of all electronics subsystems and to give subsystems flexibility in the detailed implementation of the solution. Acceptance is of course important to the completeness and success in the implementation of common solutions. In *BABAR* we gained acceptance by involving the entire community, through the overall electronics working group, in the development of *BABAR* protocols and standards. Flexibility in the detailed implementation of common solutions allows subsystems to tailor protocols and standards to the details of their system. Allowing subsystem-specific details in the implementation also helps subsystem designers to take "ownership" of the design of their subsystems.

Each design team should understand its requirements, and it should depart from convention when a departure simplifies system design or allows an important optimisation. For instance, although *BABAR* adopted a "standard" deadtimeless, multilevel architecture, certain characteristics of *BABAR* allow its architecture to incorporate some unusual features. For instance, all level 1 latency buffers in *BABAR* are digital, *i.e.* there is no analogue storage during the level 1 latency. This characteristic allowed *BABAR* to adopt an unusually long level 1 latency (12.5 μ s) without impacting the cost of front-end electronics. The benefit of this long latency was a level 1 decision time significantly longer than usual values of two to five microseconds. The longer decision time simplifies trigger requirements and reduced the cost of level 1 trigger. Similarly, modest hit occupancies in *BABAR* detector systems result in bandwidth requirements that fit comfortably within the capabilities of existing optical link technologies. The resulting headroom in bandwidth facilitated adoption of a "data-pull" architecture between the buffers in the front-end electronics and the data acquisition system, as opposed to a "data-driven" or "push" architecture, without a consequent increase in deadtime. This architecture avoids data loss, and it greatly simplifies buffer management and event synchronisation.

In order to result in a high performance system that is easy to bring into operation, the system design must include attention to power supplies and power distribution, grounding and shielding, cabling, and cooling. These are demanding areas that require attention during the design phase in order to avoid awkward remedies to problems discovered during commissioning. In *BABAR*, we were fairly successful in achieving attention to detail; however, problems arose in some areas where we were not sufficiently attentive. For instance, one of our front-end subsystems required several retrofits

to the location and mounting of electronics on the detector.

Because system design demands insight into requirements, a broad view of possible technical solutions, and attention to detail, an experienced electronics engineer can contribute invaluable to the success of the system design.

6.1 Connectors and connections

With respect to reliability and signal integrity, connectors and connections are often the weakest link of the electronics chain. Furthermore, connection problems can be very difficult to diagnose. For instance, in *BABAR*, we had a problem with some of the backplanes in our data acquisition crates. These were custom backplanes that we fabricated separately from the crates. They were somewhat thicker than a standard backplane, and the pins did not always contact the pathways in the backplane. Contact would sometimes be lost when cards were plugged into the crate, resulting in unpredictable behaviour of the backplane that could not be diagnosed by interchanging modules and spares.

Because problems with connections can be difficult to diagnose, they should be avoided in advance by careful consideration during the design phase. Testing, both during the design phase and during production, can help avoid problems. As commented by the leader of one of our electronics subsystems, "Pay special attention to cables and connectors. Get the best quality you can afford. Cheap connectors in particular can lead to lots of headaches. Make sure the vendor does a thorough test of each cable (not just a random sampling of connections). Specify what resistance constitutes an open or a short. Get an early sample and give it a workout to see if any shorts or opens develop. We thought we had done this, but when we actually started working with the cables we wished we had spent even more on connectors."

Problems with connections often arise when the electrical connection also provides mechanical support. To avoid such problems, always provide mechanical support that is independent of the electrical connection. A simple example is to strain relieve cables where they connect to printed circuit boards. Electrical connections can also be broken due to mechanical stress in the connections between motherboards and daughterboards if adequate mechanical support is not provided. For instance, in *BABAR* we experienced problems between two detector-mounted boards on the calorimeter. Our calorimeter Input/Output Boards (IOBs) directly plugged into right-angle connectors on ADC Boards (ADBs). The electrical connection between the connector and the board on the ADBs was frequently broken during mounting or remounting of the IOBs. This type of problem is particularly insidious because a new problem (a broken connection on an ADB) can easily be created when trying to remedy an existing problem (a failure on an IOB).

6.2 Grounding and shielding

Details of grounding and shielding are critical to the performance of analogue front-end electronics. Each electronics subsystem needs to invest careful design in its grounding and shielding plan. In fact, one subsystem with poorly designed grounding or shielding can adversely affect other subsystems.

In *BABAR*, we had a fairly successful experience with grounding and shielding in our final installation. Detailed grounding and shielding plans were part of design and were included in design reviews. In summary, detector subsystems in *BABAR* are isolated from one another. Each has a single point ground. Extensive work was done on grounding and shielding during prototype and system tests. Consequently, relatively little further work was needed during final installation. The *in situ* performance exceeds requirements, and is as good as on the bench. The one exception is the calorimeter power supply problem discussed in Section 12, the whose remedy is presently being implemented.

7. PLANNING

Planning a project such that enough time is allocated to each stage of the project, for instance to design and prototyping, procurement and fabrication, and installation and commissioning, can be important to the timely success of the project. It is particularly important to budget adequate time for design. Adequate time on design can avoid later problems, as pointed out by one of the leaders of our electronics project, "Spending sufficient time to make sure the design (overall and detail) is right and that it communicates properly with other parts is very important, since making it later work by testing and debugging the hardware is much more time consuming and expensive."

In order to maximise the amount of time available for design of *BABAR* electronics, while budgeting adequate time for subsequent project phases, we drew up our project schedule by planning backwards from the earliest date that electronics could be used by each detector system. In reverse order, we conservatively scheduled time for commissioning, installation, testing, production, and prototyping. Then we added some schedule contingency to our estimates for those phases. Finally, we allocated all the remaining time for design. Fortunately, we performed this planning exercise sufficiently early that this approach led to ample allocations of design time.

7.1 Design iterations

For all board-level components, we scheduled three iterations: one full-functionality prototype, one preproduction model, and one production run. In some cases, the full-functionality prototype was preceded by partial prototypes. The preproduction model was intended to be identical to the production version and was intended to

validate the final design, although in some cases we allowed further small design changes between the preproduction model and the production version. In such cases, we generally assembled and tested first articles of the production design before assembling the entire production run. Likewise, for high volume items, we also generally assembled and tested first articles before completing the full system.

For custom integrated circuits, we planned for more prototyping rounds than were budgeted for board-level components. Our experience was that our custom ICs required between four and eight iterations in total from prototype through production. Our Drift Chamber Amplifier Chip (DCAC), of which we needed approximately two thousand chips manufactured in a Maxim CPi semi-custom bipolar process, required about five iterations (1 prototype run, 2 preproduction runs, and 2 production runs). This number of preproduction and production runs included mistakes made by the manufacturer which, although not really design iterations, nevertheless cost time in our schedule. Our Drift Chamber Digitizer Chip (ELEFANT) [16-18], of which we needed about one thousand chips in a Hewlett Packard triple-metal 0.8 μ m CMOS process, required about four iterations (preprototypes of key functional blocks, 1 full-functionality prototype run, 1 preproduction run, and 1 production run). Our DIRC Analogue Chip [19], of which we needed about 1500 chips in an AMS 1.2 μ m CMOS process, required about six iterations (3 prototype runs, 2 preproduction runs, and 1 production run). Our DIRC Digital Chip [20], of which we needed approximately 750 chips in an Atmel-ES2 double-metal 0.8 μ m CMOS process, required approximately five iterations (3 prototype runs, 1 preproduction run, and 1 production run). Our Calorimeter Amplifier Chip, of which we needed approximately thirteen thousand chips in an AMS 1.2 μ m BiCMOS process, required approximately eight iterations (1 semicustom preprototype run, 3 prototype runs, 1 preproduction run, and 3 production runs). In this case, the number of production runs included issues beyond our control at the manufacturer. Our Calorimeter Auto-Ranging & Encoding Chip (CARE) [5], of which we needed approximately two thousand chips in an AMS 1.2 μ m BiCMOS process, required approximately seven iterations (1 semicustom preprototype run, 3 prototype runs, 1 preproduction run, and 2 production runs). Again in this case, the number of production runs was increased by incidents beyond our control. Note that the above statistics are not exact, and do not count multiple prototyping and production runs in the same way from chip to chip.

7.2 Manpower

In *BABAR*, one of our biggest problems was the loss of engineers during the project. This problem was particularly significant at U.S. and U.K. laboratories. We

suffered many delays from this problem. For example, on the vertex detector readout IC, as described by one of the project leaders, "We basically had 100% turnover of the U.S. engineers, which was a real nightmare. Luckily, the Pavia group stayed on the project and provided some continuity ...". Loss of engineers during the project was primarily caused by job opportunities in industry.

Engineers are still needed after the design and prototyping phase. They are needed through the debugging and commissioning phase as well. As observed by the leader of one of the *BABAR* electronics subsystems, "With the long time frame [of the project], we couldn't keep any engineer on board, and had no engineer available for debugging. This really hurt." Engineering work is not complete until the readout system is fully commissioned and operational.

The possibility of losing an engineer during a project, means that having a single engineer working on a complex, critical project is dangerous. As stated by the leader of one of *BABAR*'s electronics subsystems, "I think the main lesson is to have few or no one-engineer sub-projects, since they become zero or minus one engineer projects with Poisson statistics." For instance, at one point, we lost the one engineer working on our Readout Module that serves all detector systems. Although we recovered in time to complete Readout Module production on schedule, this debacle delayed our data acquisition software development by many months, making it much more difficult to prepare detector-specific software for commissioning. The loss of this engineer was such a significant setback for *BABAR* that it is perhaps worthwhile to describe how we recovered, and, in this account, to sketch the design of the Readout Module as it was completed.

7.3 The *BABAR* Readout Module

When the lead engineer for the Readout Module (ROM) design left the project unexpectedly in late summer 1996, a prototype ROM had already been built, and the design of the preproduction model of the ROM was largely complete. Nevertheless, in order to complete the ROM project as rapidly as possible, we decided to start a new design that was quite different from the original. The new ROM conceptual design addressed the same requirements as the original; however, it had the added goal of not requiring the same high level of engineering, *i.e.* to require a level of engineering less sophisticated than the original design. The new design did benefit from some of the experience gained during the initial design. Being modular, it facilitated multiple engineers to participate in the design, and it also somewhat decreased the coupling between ROM hardware design and data acquisition software design. The production cost of the new solution was more costly than would have been the case with the original design, because the new design uses embedded commercial CPU

boards. Nonetheless, the additional production cost was offset by reductions in the number of ROMs required and by reduced engineering costs. The new design and first prototype of this complex module was completed in just over one year.

The implementation of the *BABAR* ROM consists of four printed circuit cards that assemble into a single-width 9U VME module. A commercial single board VME computer (Motorola MVME2306) running a real-time operating system (Wind River *VxWorks*) buffers and processes data and interfaces to VME. A custom controller board provides the interface to the fast control and timing system and manages front-end buffers and the transfer of data. A custom personality module interfaces to the controller card via a private bus and interfaces to the front-end electronics through *BABAR*-standard control and data links. There are two types of personality module. The first type is used by all detector systems for control and by all systems except the calorimeter for data transfer. The second type is used only by the calorimeter for data transfer. Finally, a custom PCI mezzanine card (PMC) interfaces between the CPU and controller boards and provides the DMA engine (Intel i960) for data transfer from event data buffers on the personality card to CPU memory via a PCI/i960-bus bridge.

8. REVIEWS

Reviews played an important role in the development of *BABAR* Electronics. We instituted a series of three design reviews for each component or subsystem. The general scheme of our reviews followed a plan proposed by R. Jared of LBNL for the SDC Experiment. Our system of reviews was fairly similar to those of the quality assurance plan of the Rutherford Laboratory. Each electronics subsystem and each electronics component, either IC or board, underwent a series of three reviews. The three reviews were:

PDRR+CDR: This review was the Preliminary Design Requirements Review and Conceptual Design Review. It was held when requirements definition and conceptual design were complete and before detailed design began. It reviewed the appropriateness of the requirements that had been defined, and it reviewed whether the proposed conceptual design addressed requirements.

PDR: This review was the Preliminary Design Review. It was held following completion of detailed design and before fabricating the first full-functionality prototype. It was a detailed review of interface specifications and schematics. It also reviewed results of partial prototypes.

FDR: This review was the Final Design Review. It was held following completion of prototyping and of system tests and before production fabrication was started. It focused on completeness of the detailed design and on results of prototype and systems tests, where test results were measured against the requirements that had

been defined at the first review. The Final Design Review also reviewed plans for acceptance testing of production units.

For integrated circuits, there were often additional design reviews before each submission.

Note that *BABAR* electronics reviews were tied to milestones in each subproject. Consequently, each review was held when there was a definite purpose to the review. *BABAR* did not perform periodic reviews of the electronics. We felt that reviews tied to milestones were more effective than periodic reviews. Because the development of *BABAR* Electronics was rather rapid, the period between reviews for any subsystem or component was never longer than about nine months. Note that in addition to these reviews, each subsystem reported progress at collaboration meetings and electronics workshops.

In *BABAR*, most reviews were tied to the submission of fabrication runs. This timing avoids the unnecessary expense of fabricating something that will not work because it is designed to the wrong functional or interface specification. A leader of one of *BABAR*'s subsystems suggests, "Have a sign-off procedure so that no component can be submitted for production (prototype or final) until the people responsible for the components connected to it in any way (mechanical or electrical) have reviewed and signed off on the design. System manager should also sign. ... [This procedure] caught many problems before it was too late."

Our review procedure was well accepted within the electronics leadership and community. As commented by one of the subsystem leaders, "The concept of regular reviews at fixed intervals in the project was a very good thing. Our own experience ... was the first [PDRR+CDR] and final [FDR] one were very useful." Another comment was "Do reviews early on."

A qualified and appropriate review committee is essential to an effective review. In *BABAR*, each review committee included a physicist from the appropriate detector system who was familiar with its requirements and performance. The largest part of the review committee consisted of engineers and physicists from other electronics subsystems. These members were familiar with the architecture of *BABAR* Electronics and with the issues of concern in developing electronics. Each committee also included engineers and physicists from systems that interface to the system under review and who were familiar with interface requirements. Finally, each committee included one or two outside reviewers, who provided their experience and wisdom to the review.

Each *BABAR* electronics review required formal documentation. Electronics management provided the list of required documentation. Most of the required documentation was documentation that should be developed during the normal course of design; consequently, very little additional documentation was

required. For instance, the following documentation was required for the Final Design Review:

Updated materials from the PDR:

- Requirements Document
- Hierarchical set of block diagrams
- System Description (at level of block diagrams)
- Interface Specifications Document
- Grounding and Shielding Plan
- List of deliverables
- Cost estimate
- Production schedule

New materials for the FDR:

- Summary of design changes since PDR
- Detailed schematics of production components
- Mechanical drawings of production components
- Reliability analysis of components and of system
- Results of tests of individual preproduction units
- Results of system tests of preproduction units
- Quality Assurance Plan, including:
 - Production plans
 - Burn-in plans
 - Acceptance test plan
 - Maintenance plan
- Cooling and Access Plan
- Cable specifications (including safety ratings)

Documentation was circulated, usually by posting on the Web, approximately one week in advance of each review. We found that the review was generally not effective when documentation was not available well in advance.

For each review, the committee provided a written report. Before finalisation, draft committee reports were sent to the group being reviewed in order that factual errors in the report could be corrected. These reports served as recommendations to the Electronics System leaders. Written responses to committee reports were not requested. After consultation between the leaders of the Electronics System and the group under review, it was agreed which recommendations would be implemented. Generally all recommendations were implemented; however, there were cases when it was agreed that some particular committee recommendation should not be implemented, either because the recommendation was not practical or not in the experiment's best overall interests.

In order to allow the fullest participation of the group whose work was under review, and to foster the collaborative aspects of the review process, design reviews were usually held at the home institution of one of collaborating design teams.

In order to perform an adequate review, a design review must be sufficiently long to allow the committee to understand the design and to allow ample time for questions and discussion. In *BABAR*, we felt that it was particularly important that the review committee be left with no outstanding questions or misunderstandings during their deliberations. Consequently, we also

scheduled discussion following committee deliberations. Our electronics reviews were generally two to three days in length, although reviews of individual components were frequently performed in a single day.

Peer reviews can have a positive impact on the collaborative spirit and effectiveness of electronics development. As commented by one of the participants, "Reviews offer opportunities to exchange ideas, stay on (or leave a bad) track, get out of your computer screen."

9. MANAGEMENT

One purpose of project management is to help avoid setbacks and mistakes. The following comments are related to this purpose.

- Everyone involved in electronics development and management should appreciate that "If it can go wrong, it will go wrong." Fortunately, this statement is not completely true, but plan for setbacks.
- Problems don't go away on their own. If left unattended, problems only grow worse. Be proactive, focus on priorities, keep everything moving forward.
- There is no problem that cannot be solved, in a system that is well-designed and well-planned. Nonetheless, solutions often require time and/or money.
- Good system engineering helps avoid many mistakes and setbacks.

10. COST CONSIDERATIONS

Clearly, cost must be a consideration during design. Nevertheless, one must avoid undue technical risk when trying to reduce costs. In particular, systems should be designed to meet requirements in their first production iteration, even if more expensive. The cost and schedule impact of failing usually overshadows any small cost savings. As pointed out by one of the leaders of *BABAR* Electronics, "Cutting cost in order to save money will cost a lot more later to fix the problem."

Cost considerations during parts selection can often lead to costly retrofits or long term problems. For example, *BABAR* experienced a significant problem with reliability of the optoelectronics used to transmit data from the calorimeter to its Readout Modules. This subsystem used approximately 300 G-links and was experiencing about one failure of the optotransmitters per week until all transmitters were replaced. Failures caused loss of data, and required frequent resynchronisation of the troubled links. The problematic transmitters were not the same unit used during prototyping. Prototypes and other *BABAR* subsystems used a transmitter from Finisar. The problematic transmitters were purchased because of lower cost. The units purchased were manufactured by Methode under license from Finisar. At the time of

purchase, we were not aware of the fact that the Methode parts we purchased used CD lasers; whereas, the Finisar units available at that time used VCSELs. Short lifetime of the CD lasers caused the problematic failure rate. In fact, Methode no longer fabricates the model that *BABAR* used, and Methode Gigabit Ethernet transmitters use VCSELs. Until replacement, these transmitters were our largest reliability problem.

11. SYSTEM TESTS

System tests can avoid integration problems and problems of scale. A system test is a test of the complete readout chain, from front-end to data acquisition. It includes an actual or prototype detector, enough channels to detect large scale system problems, and actual power supplies, grounding, shielding, *etc.* The purpose of a system test is to validate that system performance is consistent with requirements. In *BABAR*, we required a system test for each subsystem before full scale production. Although our tests did not reveal any serious problems, they did help us refine grounding and shielding plans in some of our subsystems. They also gave us the confidence to enter production without worry of significant setbacks that would be costly in terms of money or valuable schedule time. In the case of one *BABAR* electronics subsystem, if the system test had involved a larger number of channels, two problems found later might have been avoided. As that subsystem leader commented regarding the "Importance of pre-testing a reasonably large sample of pre-production parts before going into production. We had two failures, which really were components. One was in manufacturing (solder-on surface-mount connectors), and the other was the bad transmitter problem [described above]. If we had made a larger preproduction and had a longer time to break it in, we might have spotted it."

12. VENDOR QUALITY

Inexperienced vendors will often take longer to complete projects within specifications than experienced vendors. Moreover, there is some risk that an inexperienced vendor will never be able to manufacture production units that consistently meet specifications. *BABAR* experienced two important problems with manufacturers of electronics components, one in the calorimeter subsystem and one in the vertex detector subsystem.

The problem in the calorimeter electronics subsystem was with power supplies for the front-end electronics. As described by one of the subsystem leaders, "In open European bidding, we got much lower bids from companies with no experience in building the sort of product we wanted. Our power supplies would no doubt have been well made by" experienced companies "but the company that made ours learned on the job." The

consequence was that the supplies were delivered at the last moment and were out of specification. They contributed common-mode coherent noise at about two times design value until additional external filters were installed on all supplies.

The problem in the vertex detector subsystem involved hybrids for the detector-mounted electronics, and ultimately led to changing vendors. The experience led one of the subsystem leaders to comment, "The key was to identify a good vendor, and not to waste a lot of time working with vendors who could not quite meet specifications. Once we found the right vendor it turned out very well." The original vendor, who produced toys and consumer electronics, was not familiar with hybrids of the quality needed for the vertex detector. The consequence was a huge delay in our ability to test the vertex detector readout chip in large numbers with detectors. Such key components "should be on a fast-track with a very experienced and responsible group in charge." It is often a delicate issue to change vendors; however, a timely change can often save time or even rescue a project.

13. CUSTOM INTEGRATED CIRCUITS

Custom integrated circuits have profoundly changed the way front-end electronics systems are designed. Indeed, there is no other way to build most present high energy physics experiments. However, the use of custom integrated circuits implies years of development, and the development period is almost always underestimated. As pointed out by one of the *BABAR* designers of integrated circuits, one of the reasons for underestimating development time, among others, is that there is essentially no room for even the smallest error during the multi-step development process. "If simulated, a chip may work; if not, it never works." The same is true for IC evaluation testing. Furthermore, throughout the development process, "One relies strongly on industry, academic design centres, and students. One single failure is able to compromise months of efforts, and due to the size of projects, it is impossible to cross-check everything." The demanding task of developing a full-performance IC can be further complicated by unnecessary complexity. Consequently, as pointed out by one of the subsystem leaders, "You can never start custom IC design too early in the project, it is usually the pacing item. Always budget time and money for additional prototype runs to fix the mistakes you didn't know you would make." And as remarked by an IC designer, "This is the price to pay for getting both quantity and quality."

13.1 Unexpected IC setbacks

Development of custom integrated circuits can also be delayed by unforeseen circumstances beyond the control of the design team. *BABAR* experienced several such

delays. Because an engineering change order twice was not transmitted properly between foundry and packaging house, two consecutive fabrication runs of our Drift Chamber Amplifier Chip (DCAC) were incorrectly bonded and were useless. In another incident, the complete order of our Calorimeter Auto-Ranging & Encoding (CARE) chip disappeared in U.S. Customs. In the case of our Calorimeter Amplifier Chip, we ended up with three slightly different shaping times that required amplifiers to be sorted and grouped for trigger reasons. The second shaping time arose because the foundry fully processed only a portion of the production run until receiving our evaluation of its performance, and then processed the remainder of the run. The third shaping time arose because the manufacturer overestimated the yield, and did not produce enough working chips in the first production run. One can draw separate lessons from each incident; however, the general conclusion is that one should plan for setbacks and delays. In *BABAR*, we experienced delays beyond our control on three of our seven custom ICs.

13.2 Radiation-hard ICs

Radiation-hard integrated circuits pose additional challenges to IC development. In the case of the *BABAR* vertex detector readout chip (ATOM), one of the subsystem leaders drew the following conclusions, some of which are general and some of which are particularly relevant to development of a radiation-hard IC. In fact, the ATOM development followed most of these suggestions.

- "If the project is distributed across more than one institution, try to use a common set of design and layout tools (seems obvious, but we had two different sets.)"
- "If you are prototyping in a rad-soft process for a rad-hard design, adopt a set of layout rules combining the most conservative rules of both processes so that the layout does not need to be modified in going from prototype in the rad-soft process to pre-production in the rad-hard process. We prototyped in HP 0.8 μ m CMOS and fabricated in Honeywell 0.8 μ m CMOS (both processes had 3 metal layers)."
- "Submit any analogue circuit elements to the rad-hard process as early as possible in the design process."
- "If the rad-hard process is not well-characterised by the vendor, plan to invest in the tools and the effort to characterise it yourself and to understand what process variations you can expect."
- "Develop your own test structures to be inserted in the reticule along with the one used by the vendor to qualify the run so you can figure out why a run failed when the vendor says it passed QA and wants payment. We did this after problems with

the first production run. The test structure we submitted was a full channel with a lot of probe points, and it was very useful for debugging the prototype. (Unfortunately we didn't include it in the pre-production!)"

13.3 Vendor IC testing

Production testing of several of the custom integrated circuits developed for *BABAR* was performed by the manufacturers. We were definitely pleased with the results of doing the testing this way. It was fast, effective at eliminating bad chips, and cost-effective. The calorimeter amplifier IC was tested at AMS. *BABAR* supplied the test jig. After some initial difficulties of getting the test jig to work properly, testing proceeded very quickly. We did not retest the chips until after they were mounted in pairs on the amplifier boards. The Calorimeter Auto-Ranging & Encoding (CARE) IC was also tested at AMS. For this chip, we supplied the test sequence and the limits of acceptance. The test vectors and the circuitry were developed at the test house. We did not retest the chips prior to assembly on the ADC boards. Less than about 0.5% of the chips failed upon retest after assembly. The vertex readout IC (ATOM) was tested as die at Honeywell. In this case, we supplied the probe card and test vectors. Honeywell supplied the tester. The test set-up was difficult to debug at the foundry, but it gave a very rapid turn-around once the wafers were out of fabrication. The DIRC Digital IC was tested by Atmel-ES2. The test vectors used at the factory covered about 60% of the chip. The delivered yield was very good.

14. PRODUCTION TESTING

Acceptance testing of production units before installation in the experiment is an important step in the development process. Adequate production testing can be a very time consuming process. Developing a test stand suitable for production testing can also be very time consuming. Acceptance testing should be planned in advance of receiving the production run. In order to ensure that acceptance test procedures were adequate in *BABAR*, we reviewed Acceptance Test Plans for each component before production as part of the Final Design Review. One of the *BABAR* subsystem leaders suggests, "Agree on a set of testing specifications and procedures for every component, to be followed by the responsible institution(s) and documented (on the Web if possible). In this way the history of every part is available on the Web and problems later on can be tracked down. Also you can be sure that each part really meets the specs."

Easy-to-operate test fixtures are often needed for acceptance testing components in large numbers. Such fixtures often require extensive software suites for tests. *BABAR* found that multiple test set-ups and stations were sometimes needed. More set-ups can be needed than expected or planned. A detector system leader suggests,

"Don't forget to build enough prototypes to provide a complete electronics readout chain to every institution that will be building detector modules or testing front-end components. ... Budget enough spares to provide a similar final test set-up wherever it will be needed to carry out any maintenance or repairs."

14.1 Industrial board production and testing

Our French colleagues working on the *BABAR* DIRC subsystem had a very satisfactory experience manufacturing and testing their principal component, the DIRC Front-end Board (DFB), in industry. As described by the subsystem leader, "We have been lucky enough to interest a big company [Thomson CSF] for which cost was not so much an issue but that wanted to learn our techniques. Besides the (usual I guess) very high requirements on quality and control for the fabrication process, the key action was to implement OUR test bench in the factory to fully validate the fabrication and to train THEM to use it." The tests at the factory included complete point-to-point coverage of the card, burn-in, and testing with a test bench provided by *BABAR*. Fifty percent of the DFBs passed the tests the first time at the factory. When the cards were received in Orsay, further tests were made. One hundred percent of the cards passed the same tests in Orsay as they did in the factory. Five percent did not pass a set of more stringent final tests used in Orsay. The screening provided by these tests was quite successful. Since installation, we have had only one or two minor problems on 168 cards. As an additional benefit of manufacture in industry, the manufacturer as part of the bid process computed the MTBF of the DFBs (for free). Note that *BABAR* required MTBF calculations before production of all major components as part of the Final Design Review.

15. INSTALLATION AND CHECKOUT

BABAR generally installed and checked out each front-end electronics subsystem on its respective detector system before the detector system was installed in the experiment. Installation and checkout in this way greatly facilitates and accelerates final commissioning *in situ*. It also provides a better, more peaceful environment for commissioning. For instance, for the particle identification system (DIRC), front-end electronics was assembled and tested on DIRC sectors in France before shipment to SLAC. They were then remounted and retested on the DIRC sectors in a staging area at SLAC before installation of the DIRC. Meanwhile, off-detector electronics was installed in the counting house prior to DIRC installation. Only final cabling and checkout occurred after installation

We also found that system tests are a useful prelude to installation and checkout. Initial checkout is also greatly facilitated by existence of a subsystem of the final data acquisition system or, in the absence of final data

acquisition components, by suitable test modules that could perform the essential data acquisition functionality. We were not as successful in *BABAR* at having full subsystems of the final data acquisition system available for early checkout as we would have liked, primarily due to our problem with turnover of engineers. This situation made the leader of one of the electronics subsystems remark that "The lack of a readout platform before production was complete really made debugging system problems extremely difficult." We did however have test modules and prototype data acquisition modules available for early checkout, and partial subsystems of data acquisition for final installation. The full data acquisition framework was only available at the last moment. Finally, we found that cosmic ray tests can provide an excellent opportunity to perform final checkout with the detector.

Installation and checkout are a manpower intensive effort, and teams of ample strength should be assembled. We did not have enough effort available during this phase in some of our subsystems. These were larger subsystems that had less capability of being checked out outside the experiment prior to installation. The consequence was a protracted commissioning period for these subsystems, and lower quality initial performance.

Installation and checkout require a team of physicists that is larger than needed during design and production phases. However, it is often difficult to identify teams to commission a subsystem that did not have a role in design and production of some component of the system. Consequently, the teams that will eventually do the commissioning should be included in the project near the beginning of the development process.

Commissioning requires appropriate software suites of diagnostics. These can be the basis of final *in situ* diagnostic programs, if they are based upon (or compatible with) the APIs of the data acquisition system. *In situ* diagnostics should be capable of validating proper operation of the entire electronics chain, including connections to the detector and of localising failures to the replaceable component (board, cable, *etc.*). More refined diagnostics for board repair can be reserved for the test bench. Developing the needed diagnostic suites is another major software effort associated with electronics development.

16. DATA ACQUISITION

16.1 Partitioning

Partitioning refers to the simultaneous and autonomous operation of portions of the data acquisition system with full functionality. For example, the drift chamber and particle ID system can use the cosmic trigger, while the vertex detector does threshold scans, while the level 1 electromagnetic energy trigger uses the calorimeter for checkout, while the calorimeter diagnoses failures in one crate, *etc.* Partitioning is important for efficient

integration of detectors into the central data acquisition system, for efficient checkout and commissioning of detector subsystems, and for efficient calibration and maintenance of detector subsystems. Partitioning is pervasive. It must be implemented in all systems that control and initialise electronics subsystems, that distribute fast timing and trigger signals, that acquire data from electronics subsystems, and that control the detector (*e.g.* high voltage).

BABAR implemented a system that allows partitions with any combination of crates, utilising any combination of triggers. It also allows trigger crates to participate in partitions with detector-specific crates. This architecture was facilitated by our compact data acquisition system (~25 crates), which allows fully flexible interconnections between "partition masters" and data acquisition crates.

16.2 Synchronisation

Because the flow of data is largely asynchronous from each front-end element in pipelined readout architectures, synchronisation of the flow of data into the data acquisition system is of central importance. This issue is often further complicated when flow control does not extend all the way to the source of data. That is, there is often no mechanism to throttle triggers when front-end buffers fill. Data flow synchronisation is a central design issue. The system should be robust against lack or loss of synchronisation; *i.e.* the system should not crash. Otherwise, it can be very difficult to commission the system or to debug synchronisation problems.

BABAR has approximately one thousand independent front-end sections. Although its readout architecture is pipelined, *BABAR* implemented a "data-pull" operation between the front-end electronics and the Readout Modules in order to simplify the synchronisation problem. In *BABAR*, the difference in deadtime for a data-pull connection *vs.* a data-driven connection between front ends and Readout Modules was very small, even at a 10KHz level 1 trigger rate *vs.* the nominal 2KHz rate. Deadtime is minimised by the multiple event buffer in the front-end preceding the readout link (D-link), and by the ample bandwidth of the D-links. Data-pull simplifies synchronisation by eliminating potential data loss between the front-end electronics and the Readout Modules. This is accomplished by allowing backpressure at the level of the Readout Module to throttle the level 1 trigger, and by creating added event coherency in the Readout Module.

In *BABAR*, despite our best intentions, thorough documentation, and design meetings, some subsystems implemented slightly different interpretations of our standard protocol, thus causing synchronisation problems. Fortunately, all of the differences were capable of remedy by reprogramming logic.

Once running in synchronisation, the *BABAR* system remains synchronised in the absence of link problems that

"damage" data. The system was more difficult to get running in the first place, than to keep running. Integrating each front-end subsystem into the data acquisition system required a very big investment of effort by the data acquisition group. Standard buffers, protocol, and readout module facilitated integration.

The most complex synchronisation issues during *BABAR* commissioning were in subsystems that require synchronisation between their front-end electronics and the trigger as well as between their front-end electronics and data acquisition. We experienced some sequencing problems in these systems.

16.3. Software

During the planning phase, software effort is frequently, or even usually, underestimated. Software is needed throughout the development cycle. It is needed to test and evaluate prototypes, to perform system tests, and to perform production acceptance testing. It is needed to check out and commission the system. It is needed to diagnose and monitor the system during operation and to maintain and repair failed components. As pointed out by the system engineer of one of electronics subsystems, "Last but not least, hardware and software have never been so closely mixed in electronics design. The time it takes to get a chip working, optimise the implementation of an algorithm, is 90% writing code: C, VHDL/Verilog, assembly, tests vectors. Even pure analogue design requires programming for the hardware test set-up."

17. CONCLUSION

A long and tortuous path, with many potential obstacles, winds between conceptual design and operation. Nevertheless, it is a path that can be navigated successfully. The development of the *BABAR* Electronics System was completed in about four years. The *BABAR* Experiment was approved in spring 1995. At that time, the conceptual design of the Electronics System was not yet complete. Electronics installation on detector systems started in Early 1998. Detector assembly, without the vertex detector, was completed in October 1998. A cosmic ray run was performed from November 1998 to January 1999. The detector, with the vertex detector, was rolled on beam line in spring 1999. Physics running started in May 1999. *BABAR* was able to record physics quality data within weeks, without an engineering run. Although some wrinkles in the Electronics System are still being ironed out, with few exceptions, the electronics has performed as designed since the start of data taking.

18. ACKNOWLEDGEMENTS

BABAR Electronics is the product of a large group of talented scientists and engineers, from laboratories and universities in Canada (*Montreal*), France (*LAL Orsay*, *LPNHE Paris 6-7*, *X-PNHE Palaiseau*), Italy (*Frascati*,

Genova, Milano, Napoli, Padova, Pavia, Pisa, Torino), the United Kingdom (Bristol, Edinburgh, Imperial College, Royal Holloway, RAL), and the United States (UC Irvine, UC Santa Cruz, Caltech, Colorado, Colorado State, Iowa, Iowa State, LBNL, Maryland, Pennsylvania, Princeton, SLAC, Stanford). This group deserves credit for the outstanding work described here. Dr. G. Haller of SLAC led the technical development.

The author would especially like to acknowledge the following individuals who contributed to the content of this paper: D. Freytag, J.F. Genat, G. Haller, M. Huffer, J.F. Kral, N. Roe, and G. Wormser.

The author's work is supported in part by U.S. Department of Energy Grant DE FG03 91ER40679.

19. REFERENCES

1. A.J. Lankford, Architecture of the *BABAR* Electronics System, Proceedings of the Fourth Workshop on Electronics for LHC Experiments, Rome, Italy (1998) 383.
2. V. Re, The Radhard Readout System of *BABAR* Silicon Vertex Tracker, Nucl. Instr. and Meth. A 409 (1998) 354.
3. J. Albert, *et al.*, Electronics for the *BABAR* Central Drift Chamber, Presented at the 1998 Nuclear Science Symposium, Toronto, to be published in IEEE Trans. Nucl. Sci.
4. P. Bailly, *et al.*, *BABAR* DIRC Electronics Front End Chain, Proceedings of the Third Workshop on Electronics for LHC Experiments, London, England, (1997) 227.
5. G. Haller and D. Freytag, Analog Floating Point BiCMOS Sampling Chip and Architecture of the *BABAR* CsI Calorimeter Front-end Electronics System at the SLAC B-Factory, IEEE Trans. Nucl. Sci. NS-43 (1996) 1610.
6. C. Bozzi, *et al.*, The *BABAR* Silicon Vertex Tracker, Nucl. Instr. and Meth. A 435 (1999) 25.
7. F. Ferroni, *et al.*, The *BABAR* Drift Chamber Project, Nucl. Instr. and Meth. A 409 (1998) 46.
8. I. Adam, *et al.*, DIRC, The Internally Reflecting Ring Imaging Cerenkov Detector for *BABAR*, IEEE Trans. Nucl. Sci. 45 (1998) 657.
9. L. Del Buono, Test DIRC: Beam Test Results and Performances, Nucl. Instr. and Meth. A 409 (1998) 382.
10. A. Stahl, *et al.*, The *BABAR* Electromagnetic Calorimeter, Nucl. Instr. and Meth. A 409 (1998) 615.
11. F. Anulli, *et al.*, The Muon and Neutral Hadron Detector for *BABAR*, Nucl. Instr. and Meth. A 409 (1998) 542.
12. I. Kipnis, *et al.*, A Time-over-Threshold Machine: the Readout Integrated Circuit for the *BABAR* Silicon Vertex Tracker, IEEE Trans. Nucl. Sci. 44 (1997) 289.
13. P.F. Manfredi, *et al.*, The Analog Front-end Section of the *BABAR* Silicon Vertex Tracker Readout IC, Nucl. Phys. B 61B (1998) 532.
14. R. Becker, *et al.*, Signal Processing in the Front-end Electronics of *BABAR* Vertex Detector, Nucl. Instr. Meth. A 377 (1996) 459.
15. P.F. Manfredi, *et al.*, Functional Characteristics and Radiation Tolerance of ATOM, the front-end chip of *BABAR* vertex detector, Presented at the 1998 Nuclear Science Symposium, Toronto, to be published in IEEE Trans. Nucl. Sci.
16. S.F. Dow, *et al.*, Design and Performance of the ELEFANT Digitizer IC for the *BABAR* Drift Chamber, IEEE Trans. Nucl. Sci., 46 (1999) 785.
17. A. Chan, *et al.*, A Multi-Channel Time-to-Digital Converter Chip for Drift Chamber Readout, IEEE Trans. Nucl. Sci. 43 (1996).
18. D. Santos, *et al.*, A CMOS Delay Locked Loop and Sub-Nanosecond Time-to-Digital Converter Chip, IEEE Trans. Nucl. Sci. 43 (1996).
19. J. Ardelean, *et al.*, A Discriminator and Shaper Circuit Realised in CMOS Technology for *BABAR*, Nucl. Instr. and Meth. A 409 (1998) 322.
20. P. Bailly, *et al.*, A 16-Channel Digital TDC Chip with Internal Buffering and Selective Readout for the DIRC Cerenkov Counter of the *BABAR* Experiment, Nucl. Instr. and Meth. A 432 (1999) 157.