

The CLEO-III Trigger: Design Methodology Case Study¹

C. Bebek³, T.J. Bergfeld², J.A. Ernst², G.D. Gollin², M.J. Haney², R.M. Hans², E.E. Johnson²,
C.L. Plager², C. Sedlack², M.A. Selen², C.R. Strohmman³, and J. Williams²

²Department of Physics, University of Illinois at Urbana-Champaign, 1110 West Green Street, Urbana, Illinois 61801

³Wilson Laboratory, Cornell University, Ithaca, NY 14853-8001

Abstract

The design of the CLEO-III Trigger produced 24 different multilayer printed circuit board designs, and over 100 FPGA configurations. Only one board design was intentionally prototyped and respun; only one board design was respun due to design/layout error; only one board design had more than 6 ECO wires. A detailed case study of the CLEO-III Trigger design methodology is presented. Scheduling, budgeting, manpower allocation, and architectural refinements are reviewed. In each area, comparisons are provided between the original projected needs and the actual final consumption. These quantitative data are offered as a sample reference for use in project planning and estimation.

I. INTRODUCTION

The design of the CLEO-III Trigger produced 24 different multilayer printed circuit board designs, and over 100 FPGA configurations. Only one board design was intentionally prototyped and respun; only one was respun due to design/layout error; only one had more than 6 correction wires.

This project represents a significant departure from past approaches used by the High Energy Physics Group of the University of Illinois, in that the majority of design and engineering work was undertaken by physicists and students overseen by a single engineer, rather than executed by engineers. Extensive use of design templates, sophisticated simulation tools, field (re)programmable gate arrays, and conservative yet not overly confining design rules allowed a relatively large and sophisticated system to be developed with very few problems.

A detailed case study of the CLEO-III Trigger design methodology is presented. Scheduling, budgeting, and personnel allocation are reviewed. In each area, comparisons are provided between the original projected needs and the actual final consumption. These quantitative data are offered as a sample reference for use in project planning and estimation. It is hoped that this information can be of value both to the planning and execution of future experiments.

II. ENVIRONMENT AND RESOURCES

The CLEO-III Trigger hardware collaboration of the High

Energy Physics Group of the University of Illinois included two physics faculty members, two post docs, five graduate students, four undergraduate students, and the assistance of one engineer from a staff of three engineers, and a technician.

The engineering group has had many years of experience in designing and implementing data acquisition and control electronics using sophisticated computer aided engineering (CAE) tools. The principle tools included the VeriBest (ne' Daisy/Dazix/Intergraph) suite of design capture and printed circuit design tools, as well as Verilog-based modeling and simulation. At the onset of the design, these tools were running on WindowsNT platforms, but only available to the engineers.

One member of the physics faculty had experience with designing Altera field programmable gate array devices. Both physics faculty members had significant hardware design experience, in digital and analog systems. One of the post docs had significant hardware design experience as well. The students had varying levels of design experience and basic course work.

III. SCOPE OF PROJECT

In the Spring of 1994, the University of Illinois undertook to design the CLEO-III trigger. The details of this trigger system can be found in three other papers in these proceedings [1, 2, and 3]. Briefly, the features of the trigger include the following:

i) 7,800 doped CsI crystals feed 504 mixer shaper modules, which produce 3 μ s rise time (180 μ s fall time) pulses. These signals are examined in geographically overlapping groups to improve the trigger efficiency in cases where the energy deposition of a shower is spread over multiple crystal groups.

ii) Calorimetry clusters are counted and organized by topology.

iii) 1700 axial drift chamber wire signals are examined directly to find tracks with transverse momentum in excess of 200MeV/c.

iv) 8100 stereo drift chamber wires, treated as 4x4 blocks, are examined in the U and V planes separately to find tracks in excess of 250 MeV/C.

v) Axial and stereo tracks are correlated by topology to find long tracks.

vi) A reconfigurable, free running pipelined trigger with access to all of the above, makes trigger decisions every 42 ns, with a total trigger system latency (pipe depth) of 2.5 μ s.

vii) Trigger decisions are conditioned by the availability of the data acquisition (DAQ) system, with corresponding bookkeeping for all suppressed triggers.

¹Presented as NSS579: Nuclear Science Symposium and Medical Imaging Conference, Lyon, France, October 15-20, 2000. Work supported, in part, by Department of Energy contract DE-FG02-91ER40677.

IV. INITIAL PROJECT PLANNING

Preliminary trigger design concepts were developed, and presented in the Spring of 1995. The follow sections outline the initial rules used to estimate the time, cost, and personnel.

A. Scheduling

The preliminary architecture identified 9 "projects" to be completed. The initial scheduling metric limited two "projects" per person per year. Multi-board sets were treated as single projects. Consideration was given to the limited amount of technician support for board assembly and testing.

Table 1 lists the individual printed circuit boards, with their corresponding schedule.

Table 1
Printed circuit board projects and schedules

	Schedule Dec-95	Schedule Jan-97	Schedule Dec-97	Actual (bare board)
Pulser		Mar-97		10-Feb-97
QPRTST		Nov-96	done	Nov-96
P_Tile		Apr-97	dropped	dropped
P_AXTR		Feb-97	done	Spring-97
QVME		Feb-98		18-Aug-99
tile control		Jan-98	bought	bought
TILE	Jan-97	Sep-97	Sep-98	Dec-98
TILED				24-Apr-98
TILED_D				28-Apr-99
TILEP3		Dec-97		98
ASUM	May-97	Nov-97	dropped	dropped
TPRO	May-97	Nov-97	Oct-98	Jan-99
TLP5P6			Oct-98	98
SURF				Feb-99
AXTR	Jan-97	Jun-97	Jul-98	1-Sep-98
AXX	Jan-97	Jun-97	May-98	26-May-98
AXP5P6		Aug-97		4-Jun-98
AXPR			Oct-98	3-Nov-98
TRCR	Sep-97	Aug-97		9-Sep-99
L1TR	Sep-97	Feb-98	Oct-98	16-Dec-98
LUMI				7-Jan-99
L1P5P6				4-Jan-99
CCGL			Oct-98	Mar-99
STTR	May-97		Dec-98	19-Nov-99
STX			dropped	dropped
STP5P6				25-Jun-99
GCAL	Jan-98	Jan-98	Aug-98	25-Feb-99
DFC	Sep-97	Mar-98	Jun-98	16-Mar-99
UPATCH				29-Sep-99
VPATCH				14-Oct-99

Architectural refinements in 1996 led to a significant increase in the number of recognized circuit boards required for the system, as seen in Table 1. More formal scheduling metrics were introduced. Table 2 lists these metrics.

Table 2

Printed circuit board project scheduling metrics (Jan 1997)

preliminary design	4 weeks
final design	4 weeks
layout	3 weeks
fabrication	4 weeks
load+test	4 weeks

A commercial project planning software package was used to produce the schedules indicated as "Jan-97" in Table 1. Again, consideration was given to the personnel resources available to assemble the electronics.

Throughout 1998 and 1999, efforts focussed away from planning and concentrated on design, assembly, testing, and installation. No scheduling data are available from that period, other than the actual completion dates.

B. Budgeting

Table 3 indicates the cost metrics used over the scope of the development. All printed circuit board (pcb) projects were assumed to have a minimum of two power planes distinct from any signal routing layers. Set up charges are more sensitive to photoplotting expenses than to tooling, and were estimated on a per layer basis. Although initially treated as part of the setup, electrical test was later called out as a separate cost item scaled by board area and insensitive to layer count. Component cost was based on board area. Assembly was initially assumed to be done in-house.

Table 3

Printed circuit board (pcb) costs for fabrication and assembly

	Jul-96	Jul-97	Apr-98
pcb setup	317/layer	150/layer	150/layer
pcb fab	0.6/squinlay	0.6/squinlay	0.75/squinlay
pcb test	included in setup	1000 per 9Ux400mm	1000 per 9Ux400mm
silicon	2000 per 9Ux400mm	2000 per 9Ux400mm	1600 per 9Ux400mm
assembly	in house	1000 per 9Ux400mm	1000 per 9Ux400mm

For several board projects, electrical test was forgone. At the time, netlist based testing was prohibitively expensive, and small quantity (less than 5) board projects would offer such poor statistics in board-board comparison that the additional cost of electrical test was considered unjustified.

Table 3 shows the evolution of the costing metrics. The increase in fabrication cost from 0.6 to 0.75 per square-inch-layer (squinlay) reflected actual feedback in board costs based on smaller projects completed to date. The decrease in

component cost estimation was based more on a combination of optimism and budget pressure than on actual data, as none of the 9U boards had been completed by that time.

Throughout 1998 and 1999 efforts were focussed on final design implementation and installation. No budgeting data are available from that period beyond the actual costs (below).

C. Personnel

It was immediately determined that this project far exceeded the resources of the engineering staff. The traditional model for project development was for one or more engineers to undertake all engineering responsibilities, with guidance and support from the faculty physicists. As two engineers were committed to other experiments, it was unrealistic for one engineer to complete the entire trigger in the required time.

Since the two physics faculty members and one of the post docs had substantial experience in electronics, a substantially different approach was adopted. With minimal exceptions, all engineering would be undertaken by the physicists and students, from preliminary design through layout, assembly, and installation, with one engineer as a full time consultant to all projects. All of the physicists and students would be given easy access to the full suite of CAE tools used by the engineers, and trained as needed by the consulting engineer.

Initial scheduling estimates (above) limited two projects per person per year. With the preliminary architecture (Spring 95), the 9 projects could be undertaken by 5 people in one year; two years were allocated to allow for sequential dependencies between projects, and bottlenecks created by limited technician support for assembly.

Since one key design criterion was to be able to reproduce the CLEO-II capability, it was determined that the stereo trigger represented pure value added. As such, it was initially allocated to another collaborator. When they became unable to undertake the task due to changing personnel, the stereo trigger became the "last thing to worry about" task.

V. ARCHITECTURAL REFINEMENTS

An important theme in the trigger design was that it should be reconfigurable to accommodate changing trigger objectives, but capable of reproducing the trigger functions of the CLEO-II system [4]. However, as always, the best laid plans are often derailed by both reality and opportunity. Architectural changes were made in the initial design, because they were needed, or might be:

i) Input buffer memory was added wherever possible. The memories would serve the dual purpose of recording input patterns (and hence the outputs from the preceding module), as well as provide a self-diagnostic resource for exercising modules by loading test patterns to be used as surrogate inputs.

ii) To separate analog and digital electronics in the calorimetry trigger, analog and digital subracks were specified. This caused the development of the "quiet VME" (QVME) interface as well as the purely digital tile processor (TPRO).

iii) As the tasks of the TPRO were refined, it became clear that a support (SURF) board, was required to combine the results from multiple TPRO boards.

iv) The analog sum (ASUM) module was removed from the architecture, and listed as a future enhancement.

v) Previously, the tile and endcap analog calorimetry circuits had been specified for separate subracks. It became clear that this was an unnecessary distinction.

vi) To combine the results of the axial trackers (AXTR), a separate axial processor (AXPR) was introduced.

vii) To simplify cabling between the digital calorimetry (SURF) boards and the level 1 trigger (L1TR), a simple cable connection board (CCGL) was introduced.

viii) It was determined that the stereo transition modules (STX) were not needed, as the necessary connections could be made through the front panel of the stereo trigger (STTR). However, two passive patch panels (UPATCH and VPATCH) were introduced to rearrange signals into advantageous patterns for the STTR.

VI. ACTUAL DESIGN

The following is a brief summary of the actual times, costs, and personnel required to complete the project.

A. Scheduling

Table 1 indicates the actual fabrication dates of each printed circuit board. While these dates do not reflect the operational date for each board (assembled, tested, functioning), the fabrication dates are objective and well documented, whereas the operational dates are not.

B. Budgeting

Table 4 lists actual board and silicon (component) costs, as computed from the bills of materials for the boards, compared to the last budget estimate from Table 3. Actual costs, reflecting minimum order expenses extremely difficult to obtain, and are not included in the average. Space precludes project by project cost; these are available on the web [5].

Table 4
Actual (average) costs

	Apr-98	Actual (average)	Std Dev
pcb setup	150/layer	76.25/layer	43
pcb fab	0.75/squinlay	0.44/squinlay	0.41
pcb test	1000 per 9Ux400mm	939 per 9Ux400m	673
silicon	1600 per 9Ux400mm	1390 per 9Ux400mm	854
assembly	1000 per 9Ux400mm	not generally applicable	

Although estimated for the 1998 budget, contract assembly of boards was limited to one project. Of the boards assembled

by the University of Illinois, only the 615 TILE daughterboards justified the expense of contract assembly.

It must be noted that several board projects designed by the University of Illinois were fabricated, and some contract assembled, by Cornell University due to budgetary constraints. Cost data from these assemblies are not available.

As can be seen from the standard deviations shown in Table 4, the variability in actual costs is quite large. Several different printed circuit board vendors were employed. In some cases competitive quotes for the same project differed by nearly a factor of 2. However, on several occasions, selection of the lowest bidder resulted in substantial delays and questionable quality.

C. Manpower

For most of 1999, little or no technician support was available. This made component purchasing and board assembly difficult. Although student labor was readily available, the consulting engineer spent large amounts of time with clerical and supervisory tasks instead of engineering.

For both personnel and financial reasons, several printed circuit board projects were "farmed out" to Cornell University, which had available resources. In most cases, the design and layout was completed by the University of Illinois, while fabrication and assembly was coordinated (and paid for) by Cornell. In the case of the TILE motherboard, preliminary schematics were provided to Cornell for their implementation.

VII. REVIEW OF FAILURES

Below is an abbreviated list of the difficulties experienced with the trigger design organized by their proximal causes.

A. Inadequate Information

Although extensive analog and digital simulation was used to verify designs prior to fabrication, very few accurate device models were available for this task. Incorrect or incomplete understanding of devices led to incorrect or incomplete models, which led to design errors. These included the lack of pull-up resistors on a power supervisor chip (first board only), the incorrect use of a PROM power-down pin as a chip select (first board only), unused SRAM address lines allowed to float (seven boards), and missing pull-up resistors for open collector lines shared between boards (one board only).

B. Improper use of Tools

Since the design entry and layout were performed by individuals with no formal training in the tools, problems were introduced as a result of improper tool use. The CAE toolset employed for the trigger allowed significant freedom of representation. However, that freedom was not supported between tools, so freedoms elected in one tool readily produced conflicts as a result of limitations imposed by the next. Time and effort was lost in correcting such conflicts.

Excess faith in the quality of the autorouter resulted in one board nearly being sent out with an overly long DTACK trace. Since the component placement could not be corrected, a

single-gate surface mount buffer was introduced near the backplane connector to isolate the trace.

C. Lack of Experience

A substantial amount of quality design work was performed by physicists and students with only one consulting engineer, most often telling them what they could not do. Several simple mistakes were made simply based on the lack of experience of the designers (*e.g.* misrouting of sensitive analog signals too close to noisy digital signals), or the lack of consulting skills of the engineer (*e.g.* not specifying the DTACK requirements as described above).

D. Lack of Resources

The loss of technician support during the development resulted in substantial delays, for lack of personnel, and in running out of parts because purchasing was not coordinated.

Both budget and schedule issues gave rise to boards being fabricated outside of the control of the group. Largely due to the difficulty in communicating detailed technical issues to third parties, the boards completed outside of the group tended to be more problematic.

E. Premature adoption of Standards

The most frustrating difficulties associated with the trigger arose as a result of adopting standard solutions were relatively new to the industry. The most significant of these was the use of 2mm pitch "hard metric" connectors. While these connectors had been available and in use for Compact PCI (CPCI) systems, their application to 9Ux400mm boards was new. Perceiving the need for pins, most of the trigger boards were designed with 5x57 type B (unkeyed) connectors.

Although the 5x25 type B connector was readily available, the 5x22 counterpart was not. Hence all of the trigger boards (and backplanes) were constructed using dual 5x25 connectors, with 3 rows depopulated from above P5. This violates the mechanical space allocated for the (unused) P4 connector.

It was decided to adopt the then new IEEE 1101.10 specification for subrack mechanics, at a time when only one manufacturer was producing prototypes for evaluation. However, to accommodate the CLEO power supply assembly, the subracks were custom extended by the manufacturer from 19" to 24". Newly designed card injector/ejector handles, engaging relatively thin punched aluminum extrusion, extended 5" beyond its original design, were asked to provide in excess of 150 lbs of insertion force produced by dual 5-row DIN connectors and 5x57 pin hard metric connectors. Even minor subrack misalignments caused the insertion forces to increase rapidly. Mechanical distortions, erosions of the in/ejector handles, and crushed hard metric backplane pins were typical. Only the last trigger board was spared these hardships, as it employed dual type A (keyed) connectors. The keyway provided substantial vertical and horizontal alignment prior to pin mating.

Low voltage differential (LVDS) signalling was used for all interboard connections. But due to improper grounding and

return for LVDS, the "failsafe" aspect of broken cables asserting Busy/Error and sending no triggers was forfeit.

Extensive use was made of dual-density blade+beam connectors on twist+flat ribbon cable. Due to anisotropies in the cable (thicker vinyl on one side) and geometry/topology issues that forced specific sidedness, many of the cables were excessively difficult to assemble and tended to break apart.

VIII. INTERPRETATION OF SUCCESSES

Of the over two dozen printed circuit boards designed by the High Energy Physics Group, eight had 2 correction wires (seven boards had the SRAM pull-up error described above; one board was changed to add functionality), one board had 3 wire (temperature sense error), and one had over 6 (TILE). All other boards are fully functional without corrections wires. With the exception of the TILE daughterboards, none of the designs were respun due to design errors.

It is typically easier to recognize failure than success. That which works is taken for granted; that which does not stands out. This section attempts to identify specific successes in the design process, and to evaluate why they succeeded.

A. *Quality Evolution*

The first design started from a carefully crafted design template created by the consulting engineer. Each subsequent design leveraged off of the previous, correcting the errors, and improving. In spite of substantial differences between designs, each shared 20-40% of the same design infrastructure (VME interface, event buffer memory, input history memory) which were convey by cut+paste to each new design.

As repetitive steps were identified, Perl scripts were developed to relieve the workload and enforce procedure.

Also, as problems were discovered, a checklist of rules and processes was refined. The first board was implemented almost entirely by a post doc talking to the consulting engineer. The last board was implemented largely by a grad student following a well crafted list.

B. *Simulation Tools*

Extensive use of simulation was employed at all levels. Analog simulation of the daughterboard, digital simulation of all FPGA designs, board level simulations (in Verilog), and in some cases board-board simulation (also in Verilog).

Verilog was the simulation language of choice, dictated by the tool set; Altera's AHDL was the synthesis language of choice again dictated by the tools. In hindsight this was an extremely advantageous mix, since the more rigorous AHDL enforced simpler design which synthesized easily and were more likely to work faster than expected.

C. *Field Programmable Gate Arrays*

To maximize flexibility, the trigger was developed as a predominantly FPGA+buffer architecture. The majority of logic was incorporated in reprogrammable devices with simple buffers and latches providing board-board communication. Later designs benefited from the earlier in this regard.

The reprogrammable nature of the design, combined with in-system (JTAG) programming, allowed several design errors to be corrected after the fact with no correction wires.

D. *Conservative Design Rules*

To simplify design decisions and component purchasing, only a small library of components was supported. The so-called "DigiKey" metric was employed; if it can not be purchased overnight, then dispensation must be sought for why that device must be allowed.

It was imperative that all of the projects either work, or be made to work. Surface mount technology was readily encouraged to provide component density, but all devices with a pin-pin pitch less than 0.05" inches employed through-hole sockets. This would allow for drill-out and rewiring if needed.

A detailed checklist was created, enumerating every step and rule of the CAE design process. As a final check, all 9U board artwork was given a final review by the consulting engineer before fabrication.

E. *Ever improving shop methods*

In spite of the loss of technician support during the development, the electronics shop continued to improve in process and performance, largely due to the continuity of our extremely conscientious undergraduate hourly employees.

One extremely valuable development was the recognition that reflowing bismuth (or other metals) into existing solder joints substantially reduces the metaling point of the junction. For large surface mount devices, each pin can be "treated" separately with a conventional soldering iron (200 C), after which the entire device can be removed at a lower temperature (100 C) provided by underboard and overboard hot air with no discernable damage to the board.

VII. SUMMARY

A case study of the CLEO-III trigger has been presented. Scheduling, budgeting, and personnel issues have been examined, as well as the failures and success associated with the development.

Detailed spreadsheets containing schedule dates, costs, and the categorization of failure case and effect are available on the web [5].

VIII. REFERENCES

- [1] see also "The CLEO-III Trigger: Analog and Digital Calorimetry" in these proceedings.
- [2] see also "The CLEO-III Trigger: Axial and Stereo Tracking" in these proceedings.
- [3] see also "The CLEO-III Trigger: Decision and Gating" in these proceedings.
- [4] Y. Kubota, et al., "The CLEO-II Detector" *Nuclear Instruments and Methods in Physics Research* **A320**, 66 (1992).
- [5] <http://web.hep.uiuc.edu/Engin/CLEO3/>