30
12-14-87
WB.

(3)

I - 12743

OR 0350-6

SLAC - PUB - 4453
October 1987
(E/I)

ENGINEERING TRADEOFFS IN MINIATURIZATION OF ELECTRONICS FOR VERY LARGE DETECTORS*

R. S. LARSEN

Stanford Linear Accelerator Center
Stanford University, Stanford, California 94305

SLAC-PUB--4453

DE88 003607

1. Introduction

In a recent issue of Proceedings of the IEEE, special section on Application-Specific IC's, the editors state:

"The field of integrated circuits has been undergoing a metamorphosis. The change is away from chips designed by great experts and manufactured in high volume, to unique chips designed by systems engineers with little or no expertise in semiconductor technology and manufactured in comparatively small quantities. Typically... the systems engineers reside in close proximity to where the systems that use the IC chips are made. This metamorphosis is being accomplished through the cooperation and changing roles of the expert IC designers and systems engineers and the increased availability of CAD tools. The primary impetus for these changes is the need to bring new and unique products to market in the shortest possible time while incorporating the advantages integration offers, such as reduced weight and volume, higher reliability, improved performance, and low power consumption [1]."

The editorial goes on to document the spectacular growth of ASIC's and projects that over half of all IC's will soon be of the application-specific variety.

The advent of the ASIC together with several related developments is being felt in the field of very large scale detectors in a unique way. Indeed, these developments are coming at a time when yet another increase in the magnitude of detectors for the Superconducting Super Collider (SSC) is being considered. Undoubtedly the timely advent of ASIC's is doing its share in fueling development of systems considered impractical a few short years ago.

Research, by its nature, demands innovation. However, very large scale detectors require a level of highly disciplined engineering if they are to be built within time and budgetary constraints, and an excessive level of innovation in the design can be a hindrance rather than an advantage toward the chief end. How do we judge the level to which new electronics technology should be introduced into a new detector design? What are the costs versus the benefits of doing so? Drawing from experience associated with the on-going design of the large SLAC detector known as SLD, slated for completion by 1989, this paper will explore these questions.

2. Engineering Goals for Electronics in Detector Design

Depending upon the size and scope of a project, one may set different goals for the electronics design. For example, if the project is small and consists of a relatively modest investment in electronics, the best engineering solution may be to purchase off-the-shelf standard modules and exercise innovation only as necessary to integrate the various components together. Even at this level, one must be very careful not to "over-engineer" the system with a clever approach or new design, unless there is a much longer-range and cost-justified view in mind. Innovation at this level introduces more problems than benefits. At the opposite extreme, namely the very largest physics detectors being built around the world today, the situation is much more difficult to assess; and the potential benefits of innovation are much more seductive. Appropriately, the pitfalls are considerably deeper and more numerous.

The electronics engineering groups in many of our National Laboratories are typically not entities in which vigorous R&D is practiced without a very specific application in mind. Therefore, the New Detector is not only an exciting task for the physicists, but it may well be the only vehicle by which engineers may practice their creative instincts as well. This can be a happy symbiosis of need and desire, or it can be a recipe for disaster. A detector which comes on-line a year late because of problems associated with some innovative subsystem would be considered at least a modest disaster even if it ultimately worked very well. However, there is a potential for disaster with a detector of even conventional design, just by the nature of its magnitude and complexity. It should be remembered that the detectors of today rival in size, cost and complexity the accelerators of a decade or two ago; so even without innovation, they are considerable feats of engineering.

The major goal, therefore, is to build a high-performance detector, in which the electronics serves the end goals of performing to state-of-the-art specifications, and of being built in time and within budget. That is, if compromises must be made between innovation and sound engineering, the latter must predominate.

What do we mean by sound engineering? Besides timeliness and planning within budget constraints, sound engineering also means designing for performance, testability, reliability, and ease of operation and maintenance. None of today's detectors are "easy" to operate, but engineering should address the problems of the best of its ability with these goals in mind.

The use of ASIC's and other modern technologies, therefore, must support this overall goal. For the system designer, this translates into a requirement that design, prototyping and testing tools for the new technologies be mastered. A practical interpretation of this requirement is that design tools and fast prototype turn around be made routine in the laboratory setting, by a combination of in-house and supporting local industry capabilities. For example, design at all levels may be accomplished in the laboratory, with prototyping support from either in-house or local vendors. The main objective is to arrive at a final, tested design as quickly and efficiently as possible, with production-quality documentation which can then be confidently submitted to vendors for competitive bid.

3. The People Problem

A major problem faced by laboratories embarking on the design of a large new system, whether it be a detector or an entire accelerator or upgrade, is the problem of acquiring and maintaining a team of people with the necessary mix of skills. Since such projects are not the steady-state diet of many of our laboratories, many of the necessary skills may not be present, or may have fallen into disuse, and may have to be developed anew. The skills being referred to are those of large-scale system conceptual design, systems engineering, test engineering, and project engineering. In addition, the laboratory may be poorly equipped at the level of support services, such as mechanical and electronic designers, prototype shops, production shops, and system coordination. If all of these elements are not present, or at least latent, the overall project execution will be poor.

Naturally, the laboratories abound with creative people, both physicists and engineers, who are eager to participate in the conceptual design of a project. This is the level at which one has an opportunity to exercise one's most interesting and constructive ideas, and eventually see them become reality. But this level lasts a very short time on a given project, after which a team of engineers, physicists and support groups of all kinds must systematically grind through a series of very

*Work supported by the Department of Energy, contract DE-AC03-76SF00515.

[1] IEEE Proceedings, Vol. 75, No. 6, June 1987.

MASTER

80

difficult problems of detailed design, space wars, management conflicts, personnel difficulties, laboratory priorities, funding difficulties, shop problems, inter-laboratory collaboration conflicts, etc., etc., for an intensive period which typically lasts four or five years.

In other words, ten percent innovation must be followed by ninety percent dedication and very hard, sometimes boring, work. It is a difficult management feat under the best of conditions to hold together a team for such a long period; it is doubly difficult if the job is not technically interesting for others besides the physicists who will ultimately garner their Nobel Prizes using the instrument.

Therefore, one advantage of an approach involving some interesting level of innovation is that the engineering and perhaps also the support staff (non-R&D staff) will be more highly motivated, and the team will function better. The danger, of course, is that the level of innovation must be manageable by the engineering team and support, and not overwhelming; otherwise, schedules will slip, and team morale and effectiveness will suffer. Regardless of whether it is a high or a low level of innovation, the team will function best and morale will remain high when goals are set and consistently met. The higher the level of innovation, the more difficult this is to achieve.

4. Technical Risks versus Benefits

At the outset of this paper, the current technology revolution spurred by the advent of application specific IC's was mentioned, in the context of powerful new tools for the systems engineer engaged in bringing products to the marketplace. In no less of a sense, the task of engineering in the large physics environment is to bring products to the "market" of a successful completed design. The usual questions need to be answered. Is the proposed design going to have a lifetime beyond this particular project? If so, additional development time and money may be warranted. Is the design taking maximum advantage of the technology at hand, as well as developments which are "just around the corner" and will become available in time to be used in this project? This is a dangerous area: many a project has been scuttled because the engineer designed it with the latest "available" chip, only to find the product later abandoned by the manufacturer, or worse, left in limbo, being neither available nor unavailable. Interestingly, ASIC's have the potential to overcome this problem which is common among manufacturers who must have a singularly large customer base to justify mass production of a less specific (and therefore less attractive, less cost effective) product. This at least will be true as long as silicon foundries continue to offer attractive pricing and deliveries for custom IC production.

Another question is, how risky is the technology being proposed? Examples of high risk are true custom IC's which have about a one-year design cycle, and for which more than one round of rework would be intolerably costly in terms of time and funds. Less risky would be the implementation of a commercial Gate Array or Linear Array where the cells are very well characterized and can be fully computer simulated *a priori*, or a semicustom IC design using standard cells and manufacturer's design tools. High risk technology can be undertaken at the very beginning of a five-year project, but if it fails, the consequences may be too severe for complete recovery. It would be better to have a lead time for truly innovative developments so that by the time a major project was to start, basic building blocks would be on hand, and the less risky tasks of forming these parts into complete subsystems could be undertaken with confidence.

The major part of any project, after the design of first prototypes, is to achieve a full set of true production models from which to construct the complete system. This is the most engineering-intensive part of the project, and probably the one where most mistakes are made, and most over-runs experienced. First of all, the danger of over-design is ever present. This may happen in the first place because of a desire to build a great deal of "smarts" or sophistication into the device or product. This is always an attraction of a new design effort—the desire to set right all the foibles of past history in a single, elegant design. The additional design load, however, coupled

with the difficulty of getting agreement on scaling down of such goals, will often lead to a syndrome of endless cycles of minor redesign, rework, and improvement; the consequent "improvements" are invariably more complex, thus compounding the implementation problem. Neither the physicists nor the engineers wish to compromise what they perceive at the time as important performance goals, with the result that the project manager finds that his production goals have slipped by months and designs are still the same incremental distance from completion. After such commitments are made, it is very hard, or impossible perhaps, to find a more modest path.

5. The Design Tools Problem

Along with the appropriation of new design architectures and higher levels of systems integration comes a requirement for new design tools and methods, specifically CAD/CAE (computer aided design, computer aided engineering) tools. The lower level tools are those which allow design and drafting of circuits and printed circuit boards, and the higher level comprises similar tools for circuit simulation and analysis, and semicustom or full custom IC design. For a large and complex project, all of these tools are needed at some level, either for use directly by the laboratory staff, or accessed via a vendor; in the latter case, the work may be done using a vendor's standard tools, or the design job may be contracted to the vendor entirely once a specification is agreed upon.

Theoretically, with the powerful new CAE tools now available, all new analog, digital and system designs should be modeled in software and exercised by computer simulation before a single wire is soldered to the circuit board. From these design simulations, detailed designs can be undertaken, further simulated, and eventually the design files passed directly to the automated PC board router to produce the first, undoubtedly perfect, prototype which will need only one small wire moved, or a minor modification in the mechanics, before being confidently committed to mass production. Simultaneously, the design file from CAD is passed through to a universal tester which can "easily" be software configured to test any analog or digital parameter on any shape or size of board, hybrid, or even a custom integrated circuit. The production tester whizzes through all such designs in microseconds, and even helps the troubled field and maintenance technician pinpoint knotty problems which the human mind is too frail to perceive. In operation, of course, the devices themselves are self-diagnosing and will call up the machine or experiment operator, before they have a malfunction, to say "I think I'm going to be sick." Because the entire design was accomplished on an integrated set of CAD tools, documentation is perfect, and the field technician or user can walk to the nearest available graphics terminal to find out anything that could possibly be known about the device in question.

This scenario is approximately one light-year from reality, galactically speaking. Although some fragment of the necessary tools exist, they fall far short of the manufacturer's "hype" in performance. No single tool does even a major fraction of the necessary tasks, which leaves us with the rather formidable job of integrating such systems ourselves in our own laboratories—or, conceivably, deciding not to do so on the basis that it is an impossible task. The too-rapid introduction of too many of these tools with their exaggerated claims can cause a rapid decrease, rather than the desired increase, in productivity and quality of the designs.

Interestingly, all of the impressive design tools which are emerging and which are extremely useful, do not solve the problem of improving engineering quality and overall productivity, which should be the chief goal. This goal remains to be achieved by a combination of sound management coupled with sound engineering judgment, and not by the tools themselves. The tools, in fact, because of the need to invest so heavily in training and structural changes within the organization, will at least temporarily cause disruption and a decrease in productivity. A slow designer, whether engineer or technician, may become, and REMAIN even slower when he or she is handed an unfamiliar tool. For some of the staff, of course, adaptability will be rapid and even spectacular; but some other fraction probably will become ineffective, which causes a real disruption and a management burden which is not easily discharged.

In the introduction of new design tools and technologies, two major problems present themselves. The first is the technical learning curve for the engineering staff, which is generally the easier to accomplish; but more significant is the management and support component, which affects the entire technical/managing support structure and ALL personnel who must eventually cope with the new technology. Generally speaking, the engineering staff are more adaptable than the technical and managerial support staff, especially in a laboratory where the staff is mature and has not been used to adapting rapidly to change; this results in another insidious problem, namely that the engineering staff finds itself doing more and more of the work which one would really prefer to have done by the support staff, in particular the design and finishing of highly complex circuits, hybrids, and printed circuit board designs; but also including the more mundane aspects such as the contracting of vendor services of various kinds and the specifying and ordering of custom IC's, hybrids and components. This situation is forced upon the innovative engineers even more strongly if the overall volume of such designs exceeds the capacity of that small portion of the technical/managing support staff which is prepared to cope with it. The condition persists until some later time when the support staff comes up to speed.

A related problem in the design support area is that once people are trained in a new technology utilizing sophisticated CAD design tools, they become a scarce commodity which is in high demand by the industry as a whole; attrition of such people, followed by the arduous problem of recruiting and retraining, can be a deadly combination. Thus we may see technological risks emerging midway through a long project in the form of a loss of key support or engineering personnel.

Acquisition of the needed design tools and capabilities therefore must be done with care toward not only which are the most cost-effective or most powerful tools, but also toward a consideration of the larger support structure needed to sustain those tools, in particular the key personnel who will utilize and maintain them. In terms of costs, the least cost is for the tools themselves, and the highest cost is the investment needed to retrain staff and develop a long term stable support structure. Since decisions of this nature have an impact on the laboratory as a whole, it is important to give strong consideration to building tools which not only operate well in isolation, but which fit well into an overall integrated system of design tools and support services.

6. The Testing Problem

It is difficult to reconcile, but the fact remains that in modern microelectronics design in large data acquisition systems, AT LEAST as much engineering effort must be applied to the design and implementation of testers and testing, as to the design itself. As subsystems become more complex, larger, and higher density, this problem becomes exacerbated. Moreover, the larger and more complex pieces are inevitably more intricately coupled, and since the standard interfaces one is used to dealing with are removed farther and farther from the bulk of the circuitry, the problems are difficult to separate, and tend to fall more heavily onto a single design engineer. We have noticed this trend some years ago with the FASTBUS board, which equals roughly 3 CAMAC boards; and we observe the same trend more recently with highly integrated front-end systems, where in fact the amount and complexity of circuitry being placed on a single front-end assembly equals or exceeds the complexity found in an entire crate of electronics a few short years ago. Obviously, we need to revise our thinking on the proper way to manage the engineering of such projects.

The problem of designing and testing a large, very high density FASTBUS board is an interesting example. With today's electronics components, we can place enormous power on such a piece of real estate. The question is, where should we draw the line? What is the tradeoff between complexity and total real estate? Consider the SLD Waveform Sampling Module (WSM), the most densely packed and sophisticated module we have attempted to date. This single-width module can receive eight parallel data streams into its auxiliary port which represent the data from eight analog fiber optic channels, all

being clocked at 1.5 μ sec per word. Each of the eight channels are corrected for linearity and offset through a full custom IC, sparsified, and the corrected reduced data stored in local data memory. Calibration memory is on-board so that each channel in the system operates with full parallelism for maximum throughput. The density of memory requires special packaging techniques using the densest available chips (two-sided daughter-boards plugged into the FASTBUS board).

This design is a challenge, both in its function, which has gone through several iterations of choice of on-board processors and software, as well as in packaging. Once a basic channel design is achieved, it is worth packaging as many channels as possible onto the board, power and cooling permitting, in order to minimize the number of modules and crates in the system, since there is a relatively fixed, and rather expensive, overhead associated with each FASTBUS slot. Therefore the density goals are clearly desirable.

What about the costs? The module has taken more than twice the design effort originally envisaged, partly because of the number of iterations, which translates into twice the design time and twice the design cost. Turnaround of even a partially instrumented unit in the shop is extremely time consuming, even with our powerful CAD tools. One reason is, that there are no shortcuts in design, and ALL complex pieces of the design must proceed in concert and/or sequentially. Parallel efforts were applied to certain subassemblies, such as the custom IC, the memories, the processor, and the FASTBUS interface; however, in the end, all of these pieces must be combined into a single integrated board, and some small number of engineers must meticulously sort out the interfaces. In some cases, the pieces designed by different people do not fit, and more work is needed—sometimes a very large amount of work—to make them fit. One would hope that very careful specification of all the subassemblies would avoid such problems, but this is difficult even with a fully documented standard interface such as CAMAC or FASTBUS, let alone with an artificial boundary placed in the middle of a design.

The same problems which plague design are carried over into the prototyping and testing phase. One can, and does, design and test the separate pieces, but the integrated unit must eventually be tackled. Presumably, if one made a deliberate choice to parcel the design into smaller functional units each in a FASTBUS module, more engineering horsepower could be brought to bear; however some of the important performance criteria would also be compromised. In research physics, above all we are trying to build very high performance tools, and compromises in the electronics are not looked upon kindly.

Finally, one must face the testing problems. This first of all demands another intensive engineering effort, since some level of adaptation of an existing tester is at least required, which involves both hardware and software effort. This problem will tend to fall again on the design engineer, or in the case of a large complex assembly such as described above, onto several design engineers and programmers. But in addition, the problems of production testing must be dealt with, which involves additional technical support manpower and training, and probably additional test equipment.

In modern, high density, high performance electronics, testing becomes a major engineering and technical support effort. For example, since the cost per unit area of the electronics is on the increase, even though the overall cost of the finished product is hopefully on the decrease; and since certain packaging techniques such as hybrid and board surface mount make it more difficult to mount and remove components by the usual methods; and since our electronics is becoming mounted in more inaccessible places within the detector assembly and therefore must operate with considerably higher reliability than otherwise; it is imperative that a much higher level of testing be performed during assembly than has been acceptable in the past. The various levels can be categorized as follows:

1. Custom IC Testing
2. Custom Hybrid Testing
3. Printed Circuit Board Testing

4. Subassembly Testing
5. Subsystem Testing
6. Full System Testing

Each of the above categories offers a different set of testing requirements and challenges.

Custom IC testing consists of three parts. The first prototypes will generally be manually probed and the critical pieces of the circuitry checked for functionality using standard test equipment jury-rigged on a probe station. To test production die, a test program must be defined and implemented on a standard programmable wafer tester; this often involves a vendor other than the foundry which makes the chips, since the latter are not set up for production testing and their testing costs are prohibitive. The third level of die testing is a full functional test of packaged die at the customer's site, since the commercial test machines in general cannot perform tests at the speeds required for most physical applications, a condition which is likely to persist. The latter typically involves designing a module which can accept some samples of the packaged die and test them by program control under all the conditions expected in practice.

Custom hybrid testing can be simple or complex, depending upon the complexity of the device. Since devices used in high energy physics front ends tend to be complex, the more complex test problem will be described. The most difficult problems we have encountered are with large high density multilayer hybrids with double substrates; examples are the SLD Drift and Liquid Argon preamplifiers. Building double substrate hybrids is not recommended, but in these two cases the packaging constraints were severe, and the design goals ambitious. For example, both units are equipped with a full randomly addressable and programmable calibration system, on-hybrid pulse power control to reduce power dissipation, and in the case of the Drift unit, a fully implemented trigger detection section consisting of discriminator-latch-serial readout for each channel, the digital section of which is instrumented on a custom gate array.

As many as six different testers are required for such a unit: a separate substrate tester for each of the two bare substrates; a functional tester for each loaded substrate; a tester for the assembled hybrid; and a test box for use in conjunction with laser trimming of the calibration section. Most of the requirements can be met by a single design with variations of test heads (probe cards) and software for each of the tests to be performed. In practice, it may be necessary to provide a vendor with all six testers as stand-alone units, depending upon the production sequence envisaged. Obviously it is necessary to standardize and simplify the testing approach as much as possible.

7. ASIC and Hybrid Procurement Problems

Besides the engineering problems, all of the new technologies require a higher level of vendor support than the more common designs using standard parts on printed circuit boards. In addition to fabrication of the special devices, special packaging and interconnect designs are required, in particular, surface mount technology (SMT). Industry is busily looking toward very extensive use of SMT, which will soon become the dominant technology for sophisticated electronics.

It is vital for the laboratories to approach procurement of the new technologies on a competitive bid basis, which in some cases is difficult to achieve. ASIC suppliers glibly promise second sourcing of designs, but in reality this cannot happen unless the designer *a priori* creates a "portable" design;

this in turn requires a very detailed review of the design rules and tooling specifications for the "target" vendors, and the rendering of a design which can accommodate multiple vendor's rules. This desire almost dictates that the design be contracted independently of a given vendor, after which the prototype fabrication and production can be let to bid. The chief danger to this approach is that some glitch may develop in the process, or the vendor may become uncooperative if all work is not done under his direct control.

Mask-making can be procured competitively, but again, some foundries will insist that this be done by themselves, or they refuse to process the job. Their mask costs may be double those of a third-party supplier—even though the foundry may ultimately use the same supplier!

In a competitive bid, care must be taken to structure the bid process to achieve a guaranteed result, which may result in some compromises. A test program must be produced in cooperation with a vendor, which in general is a subcontractor to the foundry. This phase is easily left out of the foundry bid and is better handled separately. It is also more economical to purchase wafers rather than tested die from a foundry.

Hybrid vendors are more numerous but there appears to be a very wide spread in cost quotations for the same part. Qualification of such vendors is an important precursor to award, since "low-balling" is not an uncommon practice, and, especially for difficult multilayer hybrids, great care must be taken in qualifying the vendor and some small fraction of the production run, prior to entering full production.

All such procurements will proceed much more smoothly if full production design packages are produced in-house, and if the cost-estimating expertise exists in-house so that unreasonably high bids can be analyzed and weeded out. A reasonably broad vendor base is necessary to support prototyping of small quantities as well as production of the large runs.

Quite often fiscal year constraints will necessitate delivery options in the bid package which will not be exercised until some later time. It is obviously best to build in options for the full production run at the initiation of the package, in order to obtain the best possible pricing.

8. Conclusions

The trend toward ASIC's and similar systems-on-a-chip technologies is fueling a new wave of innovation in detector electronics, just in time to address some of the problems being introduced by detectors which will approach a million channels of electronics. The cost-effectiveness of these technologies can be easily demonstrated, and the trend of the past twenty years of achieving more powerful electronics at a lower per-channel cost should receive a major impetus.

The investment required in the new technologies will reshape the work force of most laboratories, by providing more and better tools, and by requiring training or retraining of significant numbers of personnel. The need for new instrumentation standards will arise at new levels in the detectors of the future.

The laboratories must also invest heavily in integrating various CAE/CAD/CAM tools into a smoothly functioning system. They must also establish a new and different kind of working relationships with vendors and suppliers of both basic devices as well as standard packaged products.

The evolution of ASIC's poses many challenges but also promises potentially more benefits to the user community than any development since the introduction of the first microprocessors over ten years ago.

DISCLAIMER

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.