

Testing LHC electronics

J. Christiansen, CERN, Geneva, Switzerland (email: jorgen.christiansen@cern.ch)

Abstract

The special testing requirements of LHC electronics are analysed and compared with standard testing techniques used in industry. General testing problems at the IC, MCM, board and system level are analysed and related to the construction of the large and complicated electronics systems required in an experiment.

1. INTRODUCTION

Testing aspects of the construction of the large and complex electronics systems, required for LHC experiments, receives more and more attention as the LHC electronics community is moving from a conceptual and design phase into a real production phase. During the conceptual and early design phase of the electronics systems, testing problems have to a large extent been ignored. The main focus has been assigned to the construction of small demonstration systems, required to prove the correct function of a proposed architecture and its implementation. The construction of these demonstration systems, and their related components, has in many cases been performed in a competitive environment and under a significant time pressure. Testing aspects have received little attention, as it was not a major problem in the construction of small demonstration systems.

Many LHC sub-detectors are now in a situation where it has been proven that electronic systems can be built with the required physics performance. It is now required to prove that the final systems, many orders of magnitude larger and more complicated, can be built from the components used in the demonstration systems. Many integrated circuits have been designed and shown to work in small systems. These integrated circuits have though in many cases not been optimised for being produced and tested in large quantities. Some components used in demonstration systems even have serious design flaws, which could be handled in the demonstration systems, but can not be tolerated in the final system.

In the electronics industry it is known from the start of a project that no profit can be gained from a product before the design has been qualified to be of sufficient quality and that it can be produced in large quantities. The market time window of commercial products in the electronics domain is so short that a few additional months, used to solve qualification and testing problems, can reduce the final profit significantly. It is therefore common that testing and qualification aspects are dealt with from the start of a project. A part of a design team is made responsible for insuring that the

final design can be transferred to a production phase as quickly as possible. A whole set of design approaches, design tools and internal design reviews are used to minimise the problems of transferring the design into production.

2. MOTIVATION FOR TESTING

The main motivation to spend significant amounts of resources and time on testing is to be capable of making reliable systems at minimum total cost. One of the main testing philosophies is sketched in figure 1. Testing is divided into two significantly different domains. Design verification consists of qualifying a design before it is released for production. Production testing consists of testing each individual component to remove devices with failures coming from the production process.

Level	Failure mechanism	Cost	
Specification	Functionality Performance, Testability, Reliability, Inter-operability, . . .	10\$	Design verification (per design)
Design	Circuit architecture, Speed, Performance, Stability, . . .	1.000\$	
Prototype	Verification, Qualification, Production margins, . . .	100.000\$	
Wafer	Yield, Speed, Noise, Gain, . . .	1\$	Production (per chip)
Chip	Cutting, Bonding, . . .	10\$	
Module	Soldering, Substrate, ESD, . . .	100\$	
System	Cables, Connectors, . . .	1.000\$	
At customer (in experiment)	Reliability, Temperature, Vibrations, Corrosion, Radiation, High voltage, . . .	10.000\$	

Figure 1. Cost of finding and repairing a failing design/chip.

The design process has been divided into a specification, a design and finally a prototype verification phase. A missing feature is very expensive to discover when the first prototype chip is tested (~100K\$ for a commercial prototype run) but could have been added at much reduce cost at the design or specification phase. The critical part of design verification testing is if an imperfect design is released for production. In the best case a complete production lot is lost. In the worst case scenario the chips will be used in the final system, which will not function correctly or will encounter frequent breakdowns.

Production testing of integrated circuits can also be performed at different levels as shown in the figure. The higher level a failing device is detected, the higher the cost. The cost of detecting a failing chip at wafer level testing could be of the order of 1\$ per chip. At the next level, when the chip has been packaged, the detection of a chip failure means that the additional costs of packaging will also be lost (if packaging is cheap and the yield is high, it may though be cheaper to skip wafer testing and only perform testing of packaged

ICs). If a failing component makes it all the way to the final system installation, it may become very expensive to diagnose the cause of a system failure and perform the required repair. In LHC experiments a system repair may be further complicated by the fact that the electronics may be inaccessible for extended periods of time.

3. ELECTRONICS FOR LHC EXPERIMENTS

Electronics for LHC experiments are often technically characterised by the fact that they must be capable of handling very small detector signals. These signals must be buffered, during one or several levels of triggers in the front-end electronics, before finally being transferred to the data acquisition system in a digital form. It is for the front-end systems that a large set of different types of electronics has been developed. This spans from analogue low noise amplifiers to special purpose digital processors (trigger systems) over highly integrated and complex mixed signal devices. For the DAQ systems a high level of standardisation is possible, and commercial modules are used to a large extent.

A large set of different integrated circuits is required to deal with the different kinds of signals from the detector technologies used in each sub-detector. An estimated number of one hundred different ASICs (Application Specific Integrated Circuit) has been developed for the four major LHC experiments. The total production volume of ASICs for the same four experiments is estimated to be of the order of 1 to 2 million. Hundreds of new modules and MCMs (Multi Chip Modules) must be developed and a total volume of hundreds of thousands must be tested after production.

A large majority of these ASICs and modules are developed within the university environment by a large number of small design groups spread over the large world-wide High-Energy Physics (HEP) community. These kinds of groups have in most cases shown themselves capable of designing the required integrated circuits, but they only have limited experience in producing large quantities of high reliability circuits. High reliability is in many cases required from the front-end electronics, as a large part will be mounted inside the detectors where it can only be serviced on a yearly (or a few years) basis. In addition, it is required that the components maintain high reliability in a very hostile environment with high levels of radiation, magnetic fields and high voltages.

A significant part of the front-end systems must be built using MCM technology, to comply with the limited space available inside the detectors. An important question concerning MCMs is the possibility

of performing repairs if one of its components is found to be faulty. If MCMs can not be repaired it requires a very high quality of its components to get sufficient yields. A typical front-end MCM with 12 integrated circuits will only have 50% chance to be fully working if each integrated circuit has a 5% risk of being malfunctioning ($0.95^{12} = 50\%$ yield). This is even under the assumption that additional failure mechanisms (MCM substrate faults, bonding faults, etc.) are ignored. It is therefore critical for these applications to perform very exhaustive tests at the component level. Testing of the MCM itself is also problematic, as normal testing schemes used for PCBs can not easily be used for MCM testing.

For a significant part of the front-end systems special radiation hard or tolerant IC technologies must be used. These kinds of technologies are only produced in low quantities and may therefore suffer from significant lower yields than mainstream commercial technologies.

4. DESIGN VERIFICATION TESTING

Design verification testing consists of proving that a designed circuit behaves in a way that is compatible with its required role in the final system. A design is normally started from a written textual specification. From this a behavioural model using Verilog or VHDL for the digital functions can (should) be built to optimise and verify the performance of a given architecture. Analogue behavioural modelling can potentially be performed to verify the architecture of the analogue functions. Based on this, the design is mapped into a chosen technology using a given design methodology (full custom, standard cell, gate array). Detailed gate level or transistor level simulations are then used to verify the correct function of the designed circuits. It is here important to take into account all parameter variations which may result from variations in the fabrication process and in the environment (temperature, supply voltage, radiation, etc) the component has to work in. When finally a prototype chip is produced it must be verified if it complies with the original specification.

Statistics from ASIC designs in industry show that 50% of new designs are found to be working correctly during the design verification testing. The circuits are then plugged into the system, where they have to be used, and it has been found that only 25 % of the designs are found to work correctly within the system. The remaining 25 % fail in the system because the original design specification did not cover in sufficient details the functions required. This gives a clear indication of the importance of system level simulations, to verify the correct function of the behavioural model of the integrated circuit, before detailed design at the gate/transistor level is started. In HEP it is though often seen that bugs in IC designs are

circumvented, by running the system in a restricted manner.

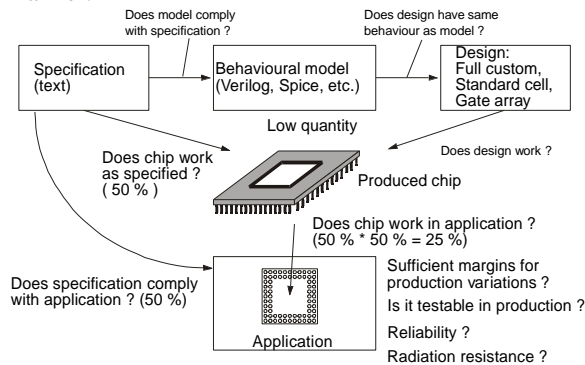


Figure 2. Design verification testing.

When a prototype has been shown to work correctly it must be verified if it is ready to be transferred into production. Before this can be done it must be insured that the design has sufficient margins to variations in process parameters. It is also important that the design can be sufficiently tested within an acceptable testing time. It must in many cases also be verified that the circuit has sufficient resistance to radiation.

For design verification testing it is important to have access to a complete set of flexible test equipment, where the circuit can be exercised with a large set of tests. To be capable of tracing the cause of possible malfunctions it is important that the designers are actively involved in the design verification testing. The total test time per circuit is not an important issue for this kind of test. The flexibility and ease of use of the test system must be considered first priority.

Design verification will in most cases be a significant part of the design time and development costs of an integrated circuit. For complicated mixed signal integrated circuits the design verification testing can be up to half of the total development costs.

5. PRODUCTION TESTING

Production testing consists of testing each produced unit to insure its correct function, before it is used in the final application. To reach a sufficient quality level the produced circuits must pass a whole set of tests. A typical production test of a digital IC consists of the following steps: functional test, internal speed test, external speed test, and finally test of IO signal levels. Mixed signal ICs must in addition pass a set of extensive analogue tests. All these tests must be performed with sufficient margins, to take into account the precision of the test equipment used and the environment in which the IC must be guaranteed to work. Testing at worst case temperature poses a particular practical problem, as the ICs must be preheated before testing is performed.

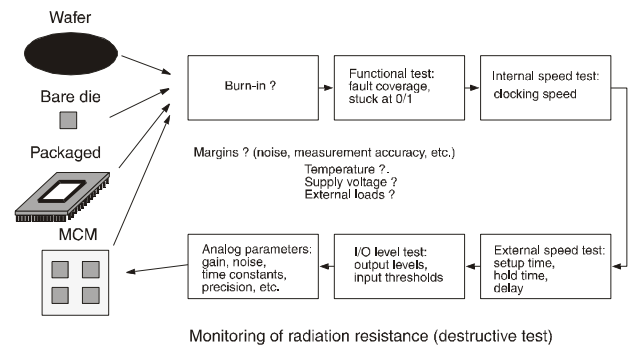


Figure 3. Production testing.

For production test systems it is important that a sufficient throughput (tested ICs/boards per time unit) can be obtained. For large scale productions the time needed to test each circuit must be minimised through a set of optimisations, still keeping a sufficient level of fault coverage. The ease of generating tests is less important as the time invested is amortised over the complete production volume.

As previously mentioned, IC's can be tested at different levels: wafer level, bare die or packaged. For large-scale productions it is often cost effective to perform test both at the wafer level and when packaged. Performing testing at the wafer level can not be used to skip the testing of packaged devices as new failure mechanisms are introduced during bonding and packaging.

Obtaining sufficiently tested circuits at the bare die level poses a specific problem, as they are very difficult to handle mechanically. Wafer level testing is no guarantee that the devices have not been damaged during cutting and the related handling. As previously mentioned, sufficient quality of components are vital to the final yield of MCMs. In addition it may be required to perform burn-in to ensure sufficient reliability (requiring an additional test after burn-in).

The development of efficient production test procedures can be a significant part of the total development budget. The cost of testing each component can be up to 50% of the final component cost in case of complex mixed signal circuits. An additional complication in the testing of HEP circuits is the monitoring of radiation resistance, which requires destructive tests to be performed on a representative sample of the production lot.

6. TEST OF INTEGRATED CIRCUITS

The resources needed to perform effective tests of integrated circuits are often largely underestimated (not only in HEP). The major driving force behind the required high level of testing of integrated circuits is the problem of yield, related to the critical and very sensitive processing steps needed for the production of modern IC's. The expected yield of a given chip area,

assuming a constant defect density, can be seen in figure 4. High volume commercial processes, used for components where yield optimisations of the design has been performed, has significantly better yield than low volume technologies used for specialised ASICs (radiation hard technologies). Some types of components can significantly improve production yield by having redundant sub-circuits (e.g. memories). In HEP circuits a failing front-end channel can in some cases be accepted, thereby significantly reduce the number of chips to reject. It must though be kept in mind, that accepting chips with certain failures may have an influence on their long-term reliability. A good overview of manufacturing yield and reliability of semiconductors can be found in [1].




Figure 4. Yield of different IC technologies.

Reliability

Reliability of integrated circuits is a special worry for applications where repairs are difficult to perform. It is known that integrated circuits have a rather high failure rate during their first few months up to a full year. After this time period it has been found that ICs have very low failure rates for tens of years. The circuits failing during the initial time interval are normally termed infant mortalities and can in some cases be of the order of one to a few percents.




Figure 5. Reliability of integrated circuits.

A one percent failure rate of the integrated circuits of the previously mentioned MCM with 12 ICs translates into a MCM failure rate of 12 %, which in most cases is unacceptable. These weak components can be screened by means of a burn-in procedure, where the circuits are heated to 100 – 125 degrees, resulting in an

acceleration factor of the order of 30 – 40. A few days at this temperature is the equivalent of several months at normal working conditions. For this burn-in scheme to be efficient to sort out the weak population it is also required to power the devices during this period (static burn-in) and if possible continuously stimulate them to keep their internal logic working (dynamic burn-in).

The reliability of components can in certain cases be reduced significantly. Badly designed components may have problems with electromigration, if the power distribution network on-chip has not been sufficiently sized. Circuits working at elevated temperatures because of insufficient cooling will also have reduced lifetimes. Contamination problems related to improper packaging or passivation may be a problem in certain working environments. Careless handling of CMOS circuits have also been seen to provoke small ESD (Electro Static Discharge) damages that may not be seen immediately. The mounting of bare die ICs on a mechanical substrate may introduce stress-based failures to occur if the thermal expansion coefficients of the IC and the substrate are incompatible. Failures from this mechanical stress are a specific problem for direct flip-chip mounting. Radiation effects are one of the major worries when it comes to the reliability and lifetime of the electronics located inside the detectors or in the caverns of LHC experiments.

Basic IC testing problems

To be capable of making effective tests of integrated circuits it is important to understand some of the basic problems in IC testing. A simple combinatorial circuit with N inputs requires 2^N test patterns to perform an exhaustive functional test. If the circuit is of sequential nature, containing M storage elements, an exhaustive test will require $2^{(N+M)}$ test patterns. It is evident that modern integrated circuits with many thousand storage elements can not be tested with this kind of brute force approach. The topology of the circuit must be used to reduce the number of test patterns to a level that can be generated by available test equipment.

The topology of the circuit can be looked at from different abstraction levels: Transistor level (layout), Gate level (netlist) or functional level of macros. A defined set of fault mechanisms must also be taken into account to limit the length of the test: shorts to ground, shorts to Vdd, broken lines, bridges between lines, etc.. Faults in simple basic components (gates) used in all digital designs poses surprisingly large problems in testing. Two examples have been chosen to illustrate this.

A simple inverter as shown in figure 6 is normally tested by asserting a logic one and a zero at the input, and then verify that the inverse values are present at the output. In case the PMOS transistor of a CMOS inverter is constantly conducting (stuck on) the failing

circuit will resemble an old fashioned NMOS logic inverter with a pull up. If the NMOS transistor used in the inverter is sufficiently strong, it will still be capable of driving the output voltage below the threshold voltage of the following gates in a logic circuit. In this case it is impossible to detect the PMOS transistor being stuck on with a simple functional test. The failure may though have serious consequences for the correct function of the circuit. The noise margins of the generated “digital” signal is seriously deteriorated and small levels of noise may result in functional failures. The propagation delay through the inverter will also be significantly changed in the described failure mode.

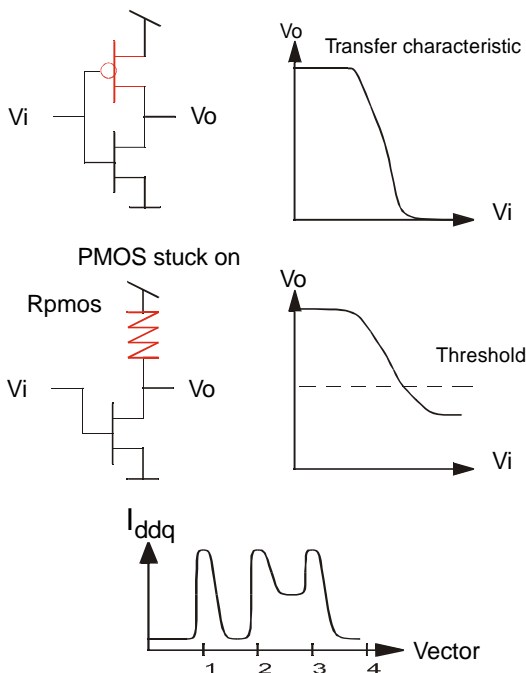


Figure 6. Failing inverter can not be detected.

The power consumption of the gate will be significantly increased when both transistors are in a conductive state. For a large IC this will just give a slight increase in power consumption, but may at long-term overload the internal power supply distribution network locally and cause electromigration effects. The increase in power consumption can actually be used to detect this kind of failure by measuring the steady state power supply current for a given set of test patterns. When a state is reached where both transistors are conducting a significant increase in steady state current consumption can be seen in CMOS logic circuits (called Iddq testing).

An even more worrisome problem in CMOS circuits is the fact that simple logic gates can start to function as sequential elements if one of their transistors is stuck open. This is illustrated in figure 7 where one on the PMOS transistors in a two input nand gate is stuck open. When the output is supposed to be driven by the faulty transistor the parasitic capacitance on the output

node will “remember” the previous output value and thereby appear as a storage element. A set of basic test patterns for a nand gate is shown and it can be seen that in a give sequence the fault will be detected, but an alternative sequence of the same patterns will leave the fault unnoticed.

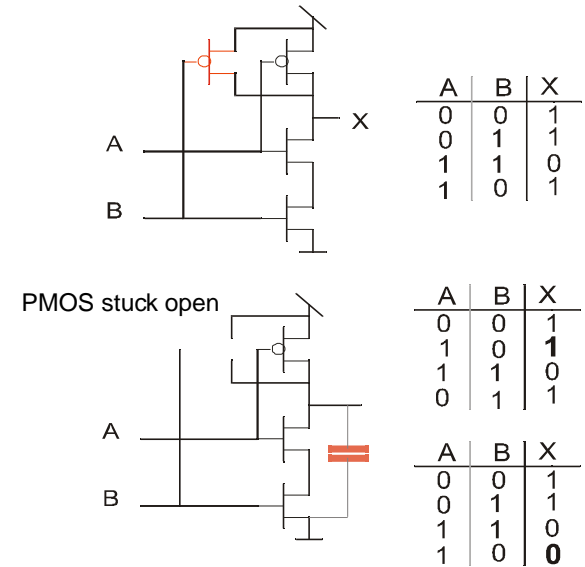


Figure 7. Failing nand gate becomes sequential.

Fault coverage

Fault coverage is normally used as a measure of the efficiency of a certain set of test patterns. Fault coverage is obtained by a fault simulation of the given design, to determine the faults detectable by a given test. Fault simulation requires large computing resources to determine if all faults are detectable by a given set of patterns. It is necessary to limit the number of fault types taken into account to limit the computing resources necessary. The most commonly used fault model for large digital circuits is the “stuck at zero”/ “stuck at one” model. This only considers the effects of any node in the gate netlist being tied to logic one or logic zero. Bridging faults and open faults are in this case simply ignored. Based on the fault simulation the fault coverage is calculated as the ratio of detected faults to the total number of possible single stuck at faults.

As previously demonstrated the stuck at zero/one fault model does not take into account even simple failure mechanisms in CMOS logic at the transistor level. The conclusion from this is that a chip, which has passed a test with a 100% fault coverage, can in fact not be guaranteed to be fully functional !. The percentage of failing ICs passing such a test is dependent on the specific implementation and is very hard to estimate.

Testability

To arrive at designs that can be efficiently tested during both design verification testing and production testing, it is very important that the design has been made with testability in mind. A design made with testability features can obtain very good fault coverage with a limited number of test vectors as illustrated in figure 8. Designs made without any support for testing may require order of magnitude longer test patterns, if it is possible at all to reach the required quality level. What is seen in normal designs is that the first part of a set of test patterns obtains a quick increase in fault coverage. It is the coverage of potential faults in “hidden” parts of the design which makes it very hard to obtain the final few percent of fault coverage (even assuming the simple stuck at 0/1 fault model).




Figure 8. Testability of different designs.

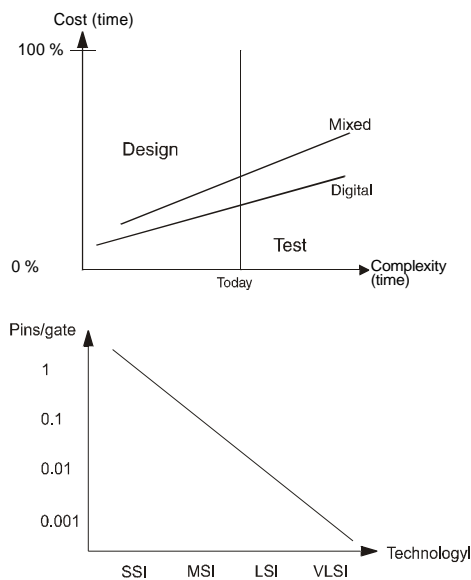


Figure 9. Decreasing testability with increased integration.

Testability in integrated circuits has continuously decreased because the number of gates per pin typically increases an order of magnitude for each new technology generation as shown in figure 9. This has meant that resources needed, to obtain sufficiently testable designs, has been steadily increasing. For mixed signal ICs this tendency has been even more

pronounced. A whole set of design methodologies and design tools has been developed in the CAE industry, to optimise and automate the process of obtaining sufficient testability.

Use of scan-path and JTAG

The use of scan paths in digital designs is one of the main schemes used to obtain sufficient testability. If all storage nodes can be accessed (read and write), in a special test mode, the testing can be performed efficiently using these virtual signal pins. The testing can in fact be so simplified that test patterns with 100% fault coverage (assuming stuck at 0/1 fault model) can be generated with Automatic Test Pattern Generation tools (ATPG). IBM has for decades obliged all in-house designs to have complete scan paths, using a scheme called level sensitive scan design (LSSD).

The IEEE 1149.1 standard has been defined to enable internal test features in integrated circuits to be accessed in a standardised manner using a minimum of pins (4). This standard, also known under the name of JTAG (Joint Test Action Group), has a large set of features improving testability at the component and board level. A simple serial protocol allows direct access to internal scan paths and Built In Self Test (BIST) features. An additional scan path gives direct access to all physical pins of the device to enable efficient tests of the connections between chips at the board level. Most commercial ICs, above a certain complexity level, supports JTAG boundary scan for board testing. These chips have in most (all) cases also extensive internal test features to insure the required fault coverage during production testing. These internal test features are though never publicly documented, as they are of no practical use to the normal user.

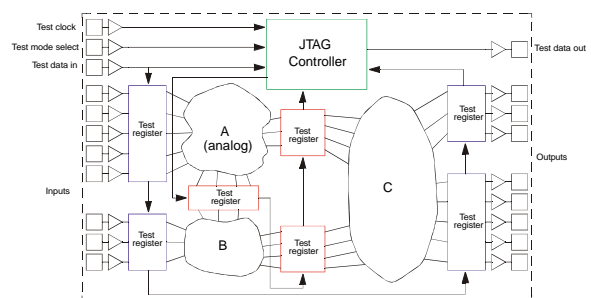


Figure 10. JTAG scan path architecture.

These schemes to improve testability have so far seen little use in integrated circuits for high-energy physics. A lot of reasons (excuses) can be found for this: not in specification, too complicated, takes too much power and silicon area, etc., etc. The main reason in most cases must though be considered to be the fact that the question of testability has received very little

attention during the specification and design phase of the projects.

7. IC TEST EQUIPMENT

Test equipment for high performance integrated circuits is very expensive with a price range from 500K\$ up to 10M\$. VLSI testers are very complicated machines with very stringent requirements. They must be capable of testing ICs with millions of test patterns at several hundred MHz on hundreds of channels (pins), with very accurate time resolution (tens of Pico seconds). Production testers must in addition have a very high throughput to be cost effective.



Figure 11. Commercial high-end VLSI tester.

Digital testers can be bought as standard commercial systems. Mixed signal test systems are though not available as standardised systems as each individual mixed signal IC has special testing requirements. Mixed signal test systems for certain types of mixed signal ICs are though now appearing (DAC/ADC, Telecom, etc.).

CERN IC test installation

The CERN microelectronics group, consisting of the order of 10 IC designers finalising several new IC designs per year, has during several years had a significant testing problem. ICs have been tested with custom-made test set-ups for each new design. The design of these dedicated test systems requires a significant effort. In many cases more time was spent debugging the test system than the time used on testing the integrated circuit itself. This kind of home-made test systems only has limited flexibility and testing performance. A significant set of parameters is limited by the system (test frequency, timing on signals, logic voltage levels, etc.). There is neither any kind of calibration available to guarantee the quality of the tested components.

It was considered to use testing facilities available in industry (manufactures or specialised testing houses). This kind of service is though very hard to use for design verification of complicated mixed signal

integrated circuits. In many cases the required mixed signal testing features are not available in these facilities. Detailed design verification of mixed signal ICs requires access to test equipment for several months and needs a close interaction with the designers. External industrial test facilities are appropriate for production testing of large quantities (high throughput, Automated handling equipment), when well defined tests are available, but were not found appropriate for detailed design verification.

Because of financial constraints it was not possible to buy a flexible high-end IC tester with sufficient mixed signal capabilities. It was therefore investigated what kind of test equipment could be purchased for a total value below a million Swiss francs. The total test system should have digital and mixed signal capabilities and also include a wafer prober in a clean room. The clocking speed of the digital part should be 100 MHz or more, to cover the 40MHz LHC bunch clocking rate with sufficient margins, and also cover ICs using double sampling. Timing resolution of the order of 100ps was considered acceptable. For mixed signal tests high speed and high-resolution arbitrary waveform generators and digitizers were required.




Figure 12. CERN mixed signal IC test system.

A system based on a "low-cost" digital design verification tester and a VXI system with the required analogue instruments was found to be an appropriate solution with a high level of flexibility. The microelectronics group was though not in a position to assemble the system and write all the required software (this part is always underestimated). A commercial company specialised in this kind of test equipment was identified to do the system integration and deliver a set of their software tools to obtain a fully integrated system (from point of view of hardware and software). This system has now been used over a period of two years and a large set of digital and mixed signal ICs has successfully been tested. This test system is available to the LHC electronics community to the extent that sufficient tester time is available. Detailed information

about this test system can be found on the web page of the microelectronics group:

<http://pcvlsi5.cern.ch/MicDig/tester/tester.htm>

During the use of this test system experience in IC testing has been gained and set of lessons learnt:

- Testing is always underestimated
- Testing requirements are often badly defined in specifications.
- Test developments must be performed by designers or in close collaboration with designers.
- You can never put to many testing facilities in your chip.
- IC testing in some cases gets delayed because of urgent beam tests. These beam tests will though often be of limited use, as function of ICs not fully understood.
- Mixed signal testing can be quite slow:
A: Lacking test facilities in IC.
B: Slow transfer and processing of large amounts of acquired data per IC.
- Synchronisation between instruments important.
- Do not have write only registers.
- Scan path and BIST test facilities extremely useful.
- Redundancy or self-checking features can be hard (impossible) to test.
- ICs with PLLs requires special attention to initialisation and synchronisation.

8. RADIATION TESTING

Verification of sufficient radiation hardness of integrated circuits is an additional complication of the testing problems in HEP. Radiation testing has no equivalence in the normal commercial electronics industry. Only highly specialised domains like space and military applications have similar problems. In these domains, only a very limited quantity of devices is needed and they can therefore better accept the high costs related to this (the launch of a satellite is already very expensive). Space industry can not accept any major component failures as repair is excluded. In LHC experiments repairs can in principle be performed, but only at infrequent intervals.

Radiation testing requires significant time and is expensive as many different effects must be investigated:

- Total dose effect
- Dose rate effects (bipolar)
- Single event latch-up
- Single event upsets
- Gate rupture (high power MOS devices).

In addition the radiation environment in the experiments is not known with a high certitude. For the radiation tests it is basically impossible to generate a realistic environment comparable to the final application. A whole set of tests with different kinds of particles with different energies must be performed to

get some kind of confidence that the components can survive. Certain technologies are “guaranteed” radiation hard or radiation tolerant, but this must still be verified both in the design verification phase and also during production.

The use of Commercial Off The Shelf (COTS) components for radiation tolerant applications has for obvious reasons received a lot of attention. This is though by no means trivial as commercial IC’s often have significant variations in their tolerance to radiation. This can be explained by the fact that the manufacturer may have several different process lines producing the same component. The detailed processing steps may also be changed by the manufacturer, without any notice to the customers. It has even been seen that a specific component type from a manufacturer in some cases comes from a processing line of an other manufacturer (hidden second sourcing). Even in the case of a special agreement with a manufacturer, that chips will come from the same process line with exactly the same processing, there is still no guarantee that the produced chips will be as radiation resistant as the chips previously tested.

For components and systems located in places with low radiation dose levels it is very hard to determine a safe limit that a standard electronics module may accept. Normal (especially modern sub-micron) CMOS IC technologies in most cases works correctly above a total dose of 10 Krad. Some very sensitive components may though start to fail at a dose level below 1Krad. Single event upsets can also be a serious problem for complicated electronic modules (with processors, memories, FPGA’s, etc.) even in a low dose rate environment.

9. MCM AND BOARD TESTING.

When performing MCM or board testing, it is normally assumed that the individual components are correctly working. It must though be taken into account that components may have been damaged during the mounting process (Soldering, ESD or incorrect handling of bare dies). MCM and board testing is to a large extent concentrated on identifying missing or wrong components and verifying the correct connections between components. It is not only important to detect if the module works or not. It must also be identified why it does not work to enable quick and effective repair.

Traditional board testing is performed using in-circuit tests and/or functional tests. In-circuit testing probes all nets on a board via a bed of nails fixture. This enables all connections on the board to be tested and allows all components to be verified with a simple set of tests. Failing or missing components can be identified directly insuring an easy repair procedure. In-circuit test has been a very popular test procedure in

industry for many years. It is unfortunately encountering significant problems on modern high-density modules using surface mount technology. It is not any more possible to probe directly all nets on the board.



Figure 13. Industrial in-circuit tester.

Functional testing only connects to the module via its normal external interface. The correct function of the module can be verified but it is very difficult to identify the cause of the failure.

MCM testing is a particular difficult case. In-circuit testing can obviously not be used because of the directly bonded (or flip-chip) mounting of the components. This only leaves functional tests that can not directly identify the failure. If the MCM is not made to be repairable, then this does not pose a problem. It is though in some cases required to be capable of repairing MCMs to get sufficient production yield.

The boundary scan feature of JTAG was included to solve these test problems at the board level. When all I/O pins of all devices on the board can be directly accessed via the boundary scan path, it is possible to test all board connections. The hardware equipment needed for this is also extremely simple, as only four pins of the module need to be connected to a computer. Specialised software products are now available from several suppliers, which automates the whole process of making a full test. A netlist of the board plus the description of the boundary scan paths of the components are enough to automatically generate a test that can pin-point the exact cause of failures. This testing approach unfortunately encounters problems for analogue circuits and when a significant amount of the digital components do not have boundary scan.

10. SYSTEM TESTING

Efficient system level testing features are extremely important to be capable of making large and complicated systems that can be made to work reliably.

The electronic systems for HEP experiments are very large and complicated systems which in addition must work in a hostile environment. Test procedures must be available to test all parts of the system in-situ. To insure this level of testing capabilities, the required functions must be in the specifications of the sub-systems, boards and components of the total system. Many electronic sub-systems, in the front-ends of experiments, are very hard to access. It must therefore be possible to identify the exact cause of a failure to perform repairs as effectively as possible.

Electronic systems for HEP experiments use large quantities of data links to transport data. These links must also have testing procedures to identify if they are the cause of a system failure.

To perform efficient system testing it must be possible to access the different parts of the system even in case of a major system failure. This data path will in HEP experiments normally consist of the DCS control path (in fact not considered a part of DCS in some LHC experiments) to the different sub-systems. In case only one combined data path (combined readout and control) is available it becomes very difficult (impossible) to identify why the system has failed. It should also be insured that all registers in the electronics can both be written and read, to have a means to verify that data has actually arrived at the intended destination.

As previously mentioned the systems have to work in a hostile environment. System failures should be expected to happen, caused by single event upsets, glitches and alike. This makes it important that the system is to a large extent self-testing during normal running, to identify if errors have occurred.

11. CONCLUSIONS

It must be considered a significant challenge to produce the required electronics for LHC experiments with sufficient quality. The size and complexity of the electronics needed are an order of magnitude larger and more complicated than what have been used in previous generations of high-energy physics experiments. To be capable of building these systems it is of vital importance that testing and qualification procedures are defined for all levels of the complete system: components, modules and sub-systems. The heavy use of custom-made integrated circuits poses a new challenge for the community. Integrated circuits can not be "repaired" once produced. In previous generations of HEP experiments, based on electronics with standard commercial components, small bugs could often be repaired at the module level. With the introduction of new technologies (IC, MCM) this will not any more be possible. It is therefore of outmost importance that all designs have been properly qualified and that all components have been extensively tested after

production. The problems of accessibility and radiation damage to large parts of the electronics is an additional complication, which has never before been faced at this scale within high energy physics (nor anywhere else).

References

- [1] Way Kuo and Taeho Kim, "An overview of manufacturing yield and reliability of semiconductor products". IEEE Proceedings, August 1999, pp. 1329 – 1344.