DD

Thesis
RAL-TH-2001-007

Radiation Effects in Electronics for the CMS Tracking Detector

J R Fulcher

Thesis-2001-Fulcher

CERN LIBRARIES, GENEVA

CM-P00040988

July 2001

#2316655

COUNCIL FOR THE CENTRAL LABORATORY OF THE RESEARCH COUNCILS

© Council for the Central Laboratory of the Research Councils 2001

Enquiries about copyright, reproduction and requests for additional copies of this report should be addressed to:

The Central Laboratory of the Research Councils
Library and Information Services
Rutherford Appleton Laboratory
Chilton
Didcot
Oxfordshire
OX11 0QX
Tel: 01235 445384 Fax: 01235 446403
E-mail library@rl.ac.uk

ISSN 1362-0215

Neither the Council nor the Laboratory accept any responsibility for loss or damage arising from the use of information contained in any of their reports or in any communication about their tests or investigations.

Radiation Effects in Electronics for the CMS Tracking Detector

Jonathan Richard Fulcher

A thesis submitted for the degree of Doctor of Philosophy of the University of London and Membership of the Diploma of Imperial College

January 2001

Imperial College
London

Abstract

This thesis presents a study into the CMS tracker analogue front-end amplifier readout chip (APV), which during the period of the study was fabricated in three different VLSI technologies. The early versions were fabricated in a total dose radiation hardened Harris 1.2μm process. Later it was transferred to a DMILL 0.8μm process and the latest version is in a 0.25μm technology. Part of this thesis describes a test system which was designed to thoroughly test APV chips on the silicon wafer and produce a comprehensive data set for each chip to enable confident selection of good chips. The main study is on the effects that large dose radiation environments can cause in the individual parts of the chip. With the chips fabricated in different technologies it was possible to make some comparisons of the magnitude of the effects between the Harris and the 0.25μm technologies, but most of the work was aimed towards understanding the effects within the 0.25μm technology. Single Event Upset (SEU) was the main consideration behind the experimental and simulation work. The study had two main goals: the first was to investigate how SEU would affect the operation of the CMS detector in the expected high radiation environment of the Large Hadron Collider (LHC). The second goal was to look at SEU from a more academic viewpoint, enabling a full understanding of how it is caused and what factors affect its magnitude. Simulations were performed in order to reconstruct the conditions brought about by highly ionising particles striking certain parts of the sensitive circuits, along with careful consideration of the mechanisms behind the effect such as: ionised charge collection within the semiconductor parts of the chip, how this charge deposition affects the circuit and how the effects manifest themselves within larger devices. A good set of results was collected from specially designed experiments, from which a confirmation of the theoretical effect was produced.

Acknowledgements

During the three years which I have spent in the High Energy Physics group at Imperial College, I have received much appreciated help and guidance in all aspects of my work. I wish to thank Mark Raymond for his unrelenting support, tolerance and especially sense of humour, not to mention all his hard work building the electronics for all the experiments. I would also like to thank Geoff Hall for his guidance throughout and PPARC and Rutherford Appleton Laboratory for funding my research over the three years.

Thanks to Federico Faccio and Mika Huhtinen for their work on SEU prediction in CMS, Irving Din Doyal for his contribution to the transistor measurements, Etam Noah for his contribution to transistor and APV irradiation, charge collection simulation and work on SEGR, and also to Lih-King Lim for his contribution to the APV test data analysis.

Thanks go to Greg, Barry, Jan, Etam, Emlyn and Rob for the great atmosphere at work and occasional moments of inspiration over a coffee or down the local. I would also like to thank my best friends: Maria, Rupert and Paul for supporting me and keeping me sane during the writing of my thesis. Thanks also to everyone in the clan, you are all great friends. I would like to dedicate this thesis to my mum, dad and three sisters without any of whom this would never have been possible. Also love goes to both of my grandmothers for staying with us long enough to see me through my education.

It has been a great three years and a real pleasure to be a part of such a lively and exciting group. Finally I would like to thank Peter Dornan for allowing me to work with the group.

Contents

List of Figures

List of Tables

Introduction

The subject of physics has spread throughout the realms of nature and now virtually encompasses the foundations of all aspects of science. One complex realm of the modern world is that of electronic systems and their interaction with their surroundings. In fact if it were not for our understanding of electronics we would be unable to probe deep into the subatomic heart of the natural world. Experimental physics now relies heavily upon our ability to create and understand complicated electronic systems that are capable of high-speed interactions with the carefully manipulated world surrounding them.

The area of physics that describes the interactions between the fundamental particles and forces of nature has undergone a great deal of evolution during the past 100 years. The study of fundamental particles and their interactions is the underlying driving force behind high-energy physics experiments. It is presently believed that the basic building blocks of matter are contained within two families of particles (the leptons and quarks), and four fundamental forces (the electromagnetic, weak nuclear, strong nuclear and gravitational). From the point of view of high-energy physics, gravity is ignored since it is many orders of magnitude weaker than the other forces. The Standard Model is a marriage of theories, based on gauge group symmetries, which describe the interactions of particles via the electromagnetic, weak nuclear and strong nuclear forces. Unification of the electromagnetic and the weak nuclear forces has been achieved within the standard model, but gravitation has not yet been encompassed.

The lepton family of particles contains six members, which are grouped into three pairs of particles with charges -1 and 0. The uncharged members, called neutrinos, until recently were believed to either have no mass at all or extremely small masses. Recent findings suggest that they do have a mass but this is still a point for conjecture. The charged members have well defined masses. Leptons interact by means of the electromagnetic and weak nuclear forces.

There are also six quarks that can be grouped similarly into three generations of pairs; their existence has been inferred from the behaviour of quark-antiquark pairs (mesons) and quark triplets (baryons). Particles composed of quarks carry integer charge and are collectively known as hadrons. Quarks interact by means of the weak nuclear, electromagnetic and strong nuclear forces; their coupling via the strong force distinguishes them from leptons.

Quantum Electrodynamics (QED) is the quantum theory of electromagnetism. It is through the application of this theory that one can predict the interaction of charged particles. QED postulates that the electromagnetic force is carried by a massless particle with unit spin called the photon. The main attribute of the equations of QED is their invariance under certain transformations of the electromagnetic field known as gauge transformations[1.1]. The unification of the electromagnetic and weak nuclear forces is achieved within Electroweak theory and requires the existence of three more force carriers (denoted the W^+ W^- and Z^0 bosons). The weak nuclear force has a short range, this fact gives rise to the requirement for the bosons to be massive; experiments have shown that their masses are large (80.2 GeV/c^2 for the W bosons and 91.2 GeV/c^2 for the Z boson). However, for the theory to exhibit gauge symmetry the force carriers must be massless. This apparent contradiction can be resolved by postulating the existence of the so-called Higgs field, mediated by the Higgs boson. The interaction of the Higgs field with the weak nuclear bosons allows them to acquire mass at low energies through the mechanism of spontaneous symmetry breaking. The search for this Higgs boson is the main purpose of the CMS detector and to gather new experimental evidence in the TeV energy regime, which has been, as yet, un-investigated.

Quantum Chromodynamics (QCD) is the theory of the strong force. Like QED it is a gauge invariant quantum field theory and contains eight mediating unit spin particles called gluons. Each quark carries one of six varieties of "colour" charge (red, green, blue, anti-red, anti-green and anti-blue), which the gluons couple to. The gluons also carry colour and thus interact among themselves, restricting the range of the strong force to sub-nuclear distances[1.1].

The Road to Electronics in High Energy Physics

The role of electronics in High Energy Physics is vastly important. It is quite clear that without the powerful tools that we can create using modern electronics, we would stand no chance of making any of the experimental measurements we need in order to probe deeper into the nature of particles and forces. Today we rely so heavily on electronic detectors and computer systems to perform the particle detection and data analysis that it is easy to forget just how recently the development of electronics and microelectronics took place.

In 1904 JJ Thomson was struggling to explain how the atom might be made up of electrons embedded in a positively charged sphere, the nature of which was a mystery. During these times it was the role of chemistry to give insight into the nature of elements and the structure of matter, but it was still mainly guesswork. At this time detecting individual charged particles was achieved by means of special screens made of substances such as ZnS, which produced a flash of light when struck by a charged particle.

One of the first real breakthroughs on the road to particle physics and understanding the structure of the atom was made in 1909 by Ernest Rutherford, when he discovered the nucleus by firing alpha particles at a gold foil and measuring the scattering angle. Expecting a distribution of small angles, he was surprised to find that very occasionally an alpha-particle would be scattered through more than 90^0. He postulated that this implied the existence of a very small extremely dense object within the atom, the nucleus. For this experiment Rutherford used no electronic equipment, employing a ZnP screen and visually detecting the flashes of incident particles. Tools like this were important in the search for an understanding of the structure of matter, but the possibility of making an experiment that could track the path of individual particles during collision experiments had not even been considered. It was years until the first electronic equipment appeared in the particle detection world.

The first effective experimental apparatus that was capable of detecting individual charged particles and tracking their path was the bubble chamber. This worked without electronics and was extremely effective at tracking particles during collision experiments. High energy physics experiments were dominated by the use of bubble chambers between 1950 and 1970. They were typically made from a large vessel, a meter or two in diameter, filled with superheated liquid hydrogen. Kinetic energy from a passing charged particle would be transferred to the hydrogen causing it to boil and produce tracks of bubbles. After a short time to allow the bubbles to expand, ~ 10 ms, the chamber was photographed from several directions enabling a 3-D reconstruction of the tracks. Unfortunately this system was slow to reset. After one event it would take about 1s for the bubbles to be removed by increasing the pressure. So it took a long time to gather enough statistics of one type of collision/decay event before it was possible to make a good measurement of the cross-section. To add to this was the problem of scanning the photographic films to look for interesting events. This was a very time consuming task, with specially employed scanners to do the job.

At this time the first computers were just being developed. They were not very powerful in comparison to modern day computers, and the size of a large room, which made them quite impractical to begin with. Over the next few years, electronic systems began to appear within the detectors themselves. Systems were developed in which film was used to record the data, and coincidence monitors (normally two scintilators combined in a logical AND that could be used to trigger spark chambers or the photography of bubble chambers) were used to trigger on interesting particles. Following this, position sensitive ionising chambers were developed by placing many anode wires between two cathode plates. Electrons were released by the ionising particle and then drifted towards the anode wire where they avalanched in the high field region near the wire, amplifying the signal. The signal, in the form of a transient current, could be detected and read out from each wire, by very simple electronic circuits, which would measure the current, in order to detect the position of the particle. Heavier ions that were also produced during the process drifted towards the cathodes, producing a larger signal over a longer time period. Both of these signals were used in order to reconstruct the tracks with a resolution as precise as 200μm.

In 1974 the J/Ψ meson was discovered. This was one of the discoveries that began to transform high energy physics. This meson did not contain any of the previously

discovered quarks (up, down and strange), but a quark-antiquark pair of a new heavier quark, charm. Over the next few years the community began to search for other mesons, which were expected after the discovery of the charm quark. These other mesons were likely to have shorter lifetimes, of the order of 10^{-12}s. This meant that the spatial resolution of the detectors had to be even better than 100μm. At this time resolutions of this order were possible in bubble chambers and emulsions, but because of the low data taking rate and extremely slow analysis method, it was very impractical to make many measurements of these rare events. At this point in the seventies, high energy physicists began to research new ways of achieving better special resolution and faster readout and analysis methods. Semiconductors had long since been discovered and were already being used as high precision energy measuring devices in other areas of physics. Soon after research and development had started on finding new ways of tracking particles the silicon microstrip detector was born. The modern versions typically consist of many long parallel p-n junction diodes spaced at a pitch of about 50μm, on a large area silicon wafer. By reverse biasing the diodes, a strong electric field is produced within the device, which quickly sweeps away electron hole pairs that are produced by an ionising particle. The problem with silicon microstrips is that, unlike ionising gas chambers, they do not intrinsically amplify their own signal, therefore it is necessary to design dense amplifier circuits to amplify the signals and transmit them for analysis. Later these circuits were reduced in size and became application specific integrated circuits (ASICs). Such ASICs must be placed right next to the silicon microstrips in order to preserve as much signal as possible. By now computers were beginning to become more practical, although still very cumbersome by today's standards, and very quickly the high energy community adopted them to perform analysis of the data from the new breed of experiments that were emerging.

With readout electronics now inside the experiment it has become increasingly important to keep their size to a minimum in order to reduce the material inside the detector thus improving the detection efficiency. This is important, as the modern experiments are getting larger.

At current e^+e^- colliders such as LEP the environment in which the tracker electronics must operate is moderately hostile, with a total radiation dose, over the lifetime of the experiment, of a few tens of krads. It does not pose a threat to the operation of the readout electronics, which have been fabricated in technologies known

to be tolerant to this magnitude of integrated radiation. However, at the Large Hadron Collider (LHC), the p-p collider under construction at CERN, the radiation environment is far more hostile. One of the proposed experiments, to be built at LHC, is the Compact Muon Solenoid (CMS)[1.2]. In the case of CMS the total dose of ionising radiation that the inner layers of the tracker will be exposed to is as high as 10 Mrads, and a neutron fluence of ~ $3x10^{14}$ cm^{-2}, over its 10 year lifetime. Not only does this raise concerns about total dose radiation tolerance of the proposed silicon microstrips and readout electronics, but the high particle flux, up to 10^7 $cm^{-2}s^{-1}$ in the inner layers, has introduced questions about transient radiation effects or so-called Single Event Effects (SEEs).

Since the introduction of CMOS (Complementary Metal Oxide Semiconductor) technology, electronics components have been continuously reducing in size. Although this has enabled more sophisticated ASICs, with advanced on board signal processing, and faster readout data rates, it has also introduced a potential new susceptibility to SEEs. The main focus of this thesis is a new generation of ASICs, the APV front-end readout chip, which has been designed specifically for the CMS tracker. Chapter 1 describes the design and operation of the APV and its constituent transistors. Research into the vulnerability of the APV to SEEs, in particular Single Event upset (SEU), formed a large part of the work for this thesis. Chapter 3 describes in detail the radiation effects which occur in microelectronics, explaining the origin of SEU and in particular the effects on the APV chip. Chapter 4 shows results from experiments that were performed to establish the sensitivity of the APV radiation, both total dose and SEEs, and investigates how this will affect the operation of CMS.

The CMS tracker has been designed to contain ~ 10 million channels. Each APV chip will read out 128 channels, hence ~ 100,000 APV chips will be required in the final system. Before these chips can be assembled, along with the microstrips, into the detector, they must be qualified. This means every chip must undergo a thorough test procedure, following which only chips passing the screening will be used in the detector. In order to test such a large number of chips, it is necessary to test them while they are still on the uncut wafer, each wafer holding approximately 400 APVs. This requires a sophisticated testing apparatus, capable of probing chips on the wafer and running an exhaustive set of tests on each. This system has been under development at Imperial College since 1996, and it is described in detail in chapter 2 of this thesis.

The work which I have contributed consists of the design and implementation of the APV testing and corresponding data analysis, design of experiments for the SEU testing, analysis of SEU data, simulation of SEU in the APV25, theory of SEU upset modes and heavy ion cross-section curve, and some of the transistor measurements. Other work contributed by colleagues in the group was some analysis of APV test data by Lih-King Lim, total dose radiation testing of transistors by Mark Raymond, Etam Noah and Irving Din Doyal.

Although there is a lot of research continuing into new particle detectors and front-end readout, these systems have only been developed as a vehicle to improve our understanding of the fundamentals of particle properties and interactions. Chapter 1 begins with a brief description of CMS and the specific way it has been designed in order to optimise its potential for discovering new and exciting physics.

Chapter 1

The LHC & CMS

The next generation of High Energy Physics experiments will begin with the completion of the new Large Hadron Collider (LHC) to be built at CERN, Geneva, Switzerland. The LHC has been designed to enable us to probe even deeper into the unknown, with an increased centre of mass energy for proton-proton collisions of up to 15.4 TeV. Two counter rotating bunches of protons will be accelerated around a 27 km ring at CERN with a bunch crossing frequency of 40 MHz. At this frequency the proton bunches will cross at 25 nanosecond intervals, achieving a peak luminosity of 1.7×10^{34} $cm^{-2}s^{-1}$. Completion of the LHC is due in the year 2005 when it will commence operation at a lower luminosity of about 10^{33} $cm^{-2}s^{-1}$, building up to full luminosity over a period of a few years. This period will be utilized for low luminosity physics such as B-physics and CP-violation studies among others[1.3]. The Compact Muon Solenoid (CMS) is one of the principal LHC experiments which have been designed to take advantage of the new physics possibilities offered by the LHC. Its design has been optimised to find and characterise the Higgs Boson in the mass range 90 GeV to 1 TeV. At the required 40 MHz operating frequency and during high luminosity running the central tracking regions of CMS will be operating in an extremely severe radiation environment. Over a ten-year running period the deposited energy will be around 10 Mrads. This provides the challenge of designing detectors and readout electronics that are not only fast enough to operate at this speed but capable of withstanding this sort of damaging radiation.

The CMS (Compact Muon Solenoid) detector has been designed to search for the Higgs boson in the mass range 100 GeV to 1 TeV by identifying and measuring the momentum and energies of its decay products. The mass range up to 100 GeV and just beyond has been covered by existing experiments at CERN.

The Standard Model predicts that some of the principal decay products of the Higgs will be electrons, muons and photons; thus the CMS design has been optimised for the detection and measurement of these particles over a large energy range at high luminosity. The design is also well suited to studying CP violation effects in the B-meson system and for studying the physics of the top quark. In addition to being operated as a proton-proton collider, the LHC will be used to collide heavy ions at a centre of mass energy of 5.5 TeV per nucleon pair, thus enabling CMS to be used to study the physics of quark gluon plasmas[1.1].

CMS is made up of many equally important sections[1.4], not all of which there is need to discuss in this thesis. The main focus of attention is the Tracker, which will consist of a central silicon pixel section surrounded by a silicon microstrip section. The high radiation environment of the tracker, coupled with requirements for greater resolution in the tracking of the collision products, has resulted in extremely demanding design requirements for both the individual detectors and the readout electronics. To achieve a sufficient resolution, the granularity of the tracker has to be such that the total number of channels is approximately 12 million. Coping with such a vast amount of data at high speed is no mean feat and the required radiation tolerance compounds the problem further.

The LHC will be situated in the LEP tunnel at CERN, and the detectors will sit in custom designed experimental halls at various points around the ring. Protons will be collided with a nominal centre of mass energy of $\sqrt{s}$ = 14 TeV. Even though the total cross section for p-p interaction increases with energy, the cross section for an individual process reduces as $1/E^2$, and therefore high luminosity is required for high statistics of interesting interactions. An integrated luminosity of up to 10^5 pb^{-1} per year is anticipated[1.5].

1.1 The LHC Radiation Environment

In an experiment such as CMS it is essential to consider the hostile radiation environment produced under operation at high luminosity. Extensive simulations show that the dose due to ionizing particles and photons is almost independent of detector geometry and detector material, but scales with rapidity and distance from the interaction point. Low energy charged particles curling inside the magnetic field

increase the dose in the inner tracker, as do secondary interactions in the tracker material and the beam pipe. It has been shown [1.6][1.7] that close to the vertex, the particle flux is dominated by low energy pions. Figure 1.1[1.8] shows the total dose, and neutron and charged hadron fluxes within the tracker. It can be seen that the ionizing radiation dose varies from 9 Mrad per year at r = 7.5 cm to 0.1 Mrad per year at r = 100 cm, with a dose of about 3 Mrad per year at r = 20 cm. The neutron flux varies by about an order of magnitude inside the tracking cavity, which extends to a radius of 120cm, with a typical value of 10^{13} cm^{-2} per year, and the charged hadron flux decreases from 2.5 x 10^{14} cm^{-2} per year at r = 7.5 cm to 10^{12} cm^{-2} per year at r = 1 m. To a first approximation, the particle flux varies with the inverse square of the distance from the interaction point, as does the dose rate.

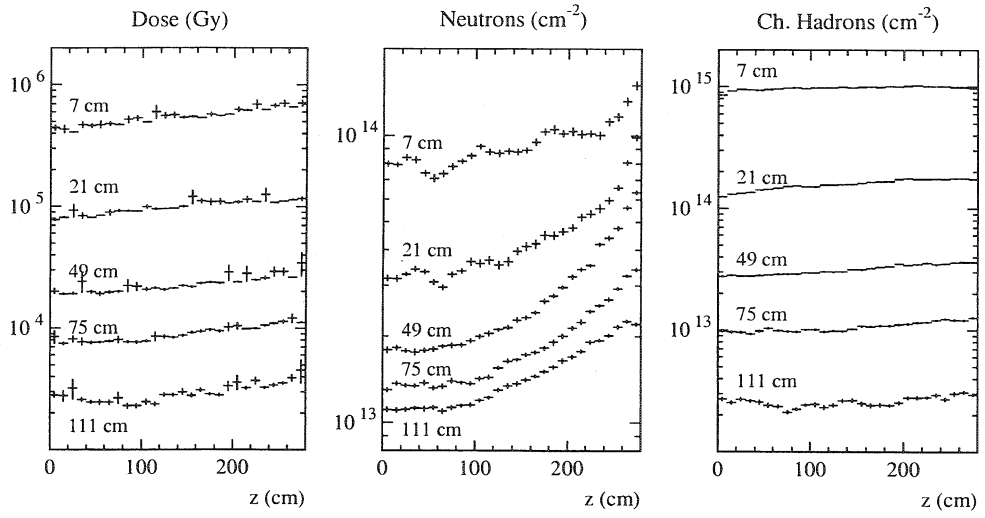

Figure 1.1 Radiation levels at selected radii in the CMS Tracker region. All values correspond to an integrated luminosity of $5x10^5$ pb^{-1}. The error bars indicate only the statistics of the simulations. The neutron fluences include only the part of the spectrum above 100 keV.

1.2 The CMS Experiment

The CMS collaboration have decided to build a detector based on a solenoid magnet with a compact design, therefore a high magnetic field is necessary. This is good because the field is parallel to the beam, therefore the bending of the muon tracks is in the transverse plane and, since the vertex will be known to within 20 μm, triggers can be based on pointing to the vertex. The momentum measurement in a solenoid starts at R=0, therefore there is little energy loss of the interaction products prior to their

momentum measurement and the design is compact. However, this means that the high precision measurements must be made in the area of the detector with the highest occupancy, which requires high granularity and puts tough demands on the detector design. High precision measurements must also be made within the muon chambers, unlike with a torroidal magnet, which also has the advantage that the momentum measurements are made at a larger radius and hence lower occupancy, putting less stringent demands on the tracking detector design. On the other hand a wide range of experience exists in the high energy physics community of working with solenoids.

The high field will give good momentum resolution for muons of high momentum (~ 1 TeV) up to $\eta = 2.5$ (Note that the polar angle of any point in the detector from the interaction point, which can be represented in angular units of radians or degrees, can also be represented by the quantity pseudo-rapidity, η, which is defined by $\eta = -\ln(\tan(\theta/2))$).

As has been shown, the CMS tracker must have a very good and redundant muon system and the best possible ECAL system, without significant compromise of the muon system. A high quality inner tracker is also desirable to supplement these systems. Figure 1.2 shows the current design for the CMS detector [1.3].

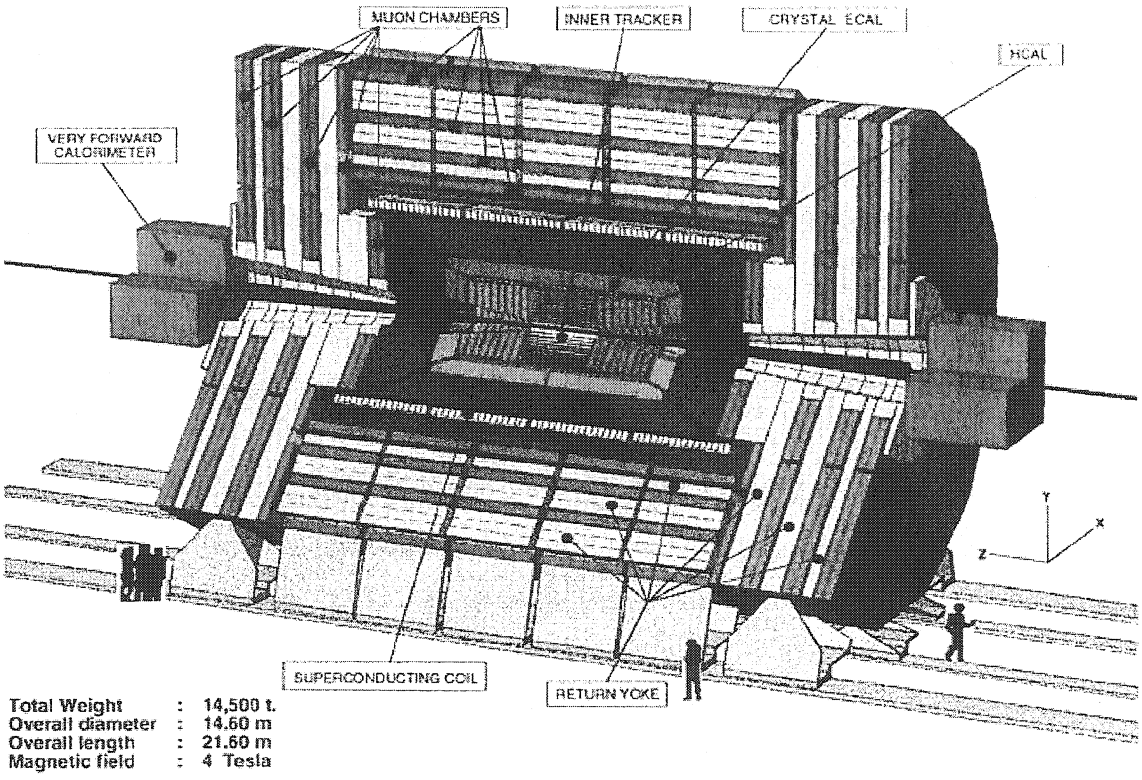

Figure 1.2 A three-dimensional view of the CMS detector.

1.2.1 The Tracking System

The Tracker is mainly concerned with reconstructing the particle tracks to enable measurement of their momentum. It is required to cover a large degree of rapidity and for high quality muon resolution it is necessary to be able to reconstruct high p_T tracks with an efficiency of greater than 95%, and high p_T tracks within jets with an efficiency of greater than 90% up to $\eta = 2.6$[1.3]. It must also be able to separate electrons and photons by precise energy-momentum matching, and to isolate calorimeter showers.

The very centre of the tracker contains two layers of silicon pixels, which are placed close to the beam pipe to optimise the pattern recognition capability. The outside of this region is made up of silicon microstrip detectors. The overall requirement for high resolution and a radial coverage of 1.3 m leads to a very large number of individual readout channels (~10^7). It is the purpose of the front-end system to capture the data from these channels and transmit any data, which may be of interest, to the receiver module.

1.2.2 The Front End System

Careful consideration of the requirements of the system as a whole resulted in the choice of an analogue, as opposed to an apparently much simpler binary system. It has been shown that an analogue system can improve the resolution because of its ability to detect charge sharing between adjacent strips and its immunity to common mode noise[1.8], whereas a binary system has a considerably smaller volume of data to read out; in a system with 1% occupancy, approximately 1 strip per 128 channel chip will be hit, and typically 1 or 2 neighbouring strips would also be read out. A full address and time stamp for neighbouring hits is not required, since that information can be inferred from the central strip. Therefore only a flag of 4 bits needs to be sent, which means neighbour hits contribute very little to the total data rate. The flag along with the address of the central strip, which must be of 7-10 bits, and a timestamp of at least 10 bits leads to ~20 bits for each event. The analogue system however, reads out each channel at the same speed as one binary bit (40 MHz), plus a 12-bit header, a total of 140 clock cycles, hence there is a big difference in the number of links for an analogue and binary system. Although there is an increase in the amount data to transmit, there is also a considerable gain in the amount of information about the event, making the analogue system more luxurious but more expensive.

In the central region of the detector there will be up to 10 Mrads (10^5 Gy) of deposited energy over a period of ten years of high luminosity running. With this in mind, it is important to monitor regularly the extent of the damage. Analogue data provides a much better indicator of the state of the front-end electronics and detectors , enabling tuning of the system throughout its lifetime.

A schematic of the CMS tracker readout scheme is shown in Figure 1.3[1.8]. The first stage of the system is the APV front end amplifier chip, the purpose of which is to sample, amplify and store, for up to 4 μs, the signals from their associated detectors. Upon external triggering the APV transmits the corresponding data via an optical analogue link to the receiver module.

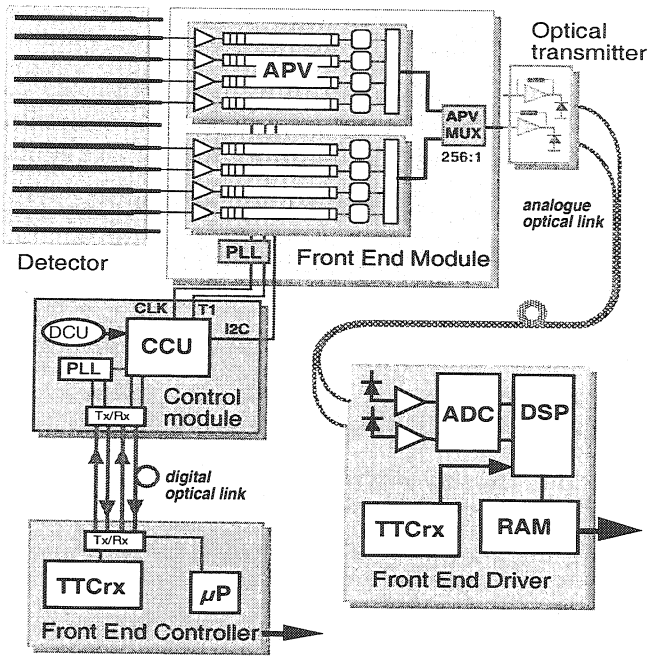

Figure 1.3 Diagram of the front-end readout

By transmitting the analogue data to the receiver module, one eliminates the need for ADCs inside the tracker volume. External ADC's can then be used which do not have to be radiation hard, and can be chosen from a large number of commercially available circuits making the system more cost effective.

The identification of particles crossing the tracking layers must be performed with good timing precision to enable the reconstruction of charged particle tracks, therefore it is essential that charge pulse data can be linked to the exact bunch crossing from which they originated. The signal from a minimum ionizing particle passing through a 300 μm silicon microstrip detector is approximately 25,000 electrons, over a period of 10 ns.

This signal must be stored for up to 4 μs until the triggering system decides whether the data were of interest. Therefore the front-end chip must be capable of amplifying a 25,000 electrons signal, operating at the bunch crossing frequency of 40 MHz, storing the analogue signal for up to 4 μs and confining the signal to one bunch crossing. In order to measure efficiently the low energy region of the energy deposition spectrum, it must also not have a noise greater than 2000 electrons and be able to minimize the loss of data due to many consecutive events. The overall power requirements, which are set by a compromise between lower noise, higher costs and reduced cooling requirements, mean that the chip must not draw more than 2 mW per channel. The simplest way to confine the signal to one bunch crossing with low noise is with a very high-speed amplifier, but this requires a large amount of power. An alternative, developed by the RD20 collaboration [1.9], involves low noise slow amplification (a charge sensitive preamplifier and shaper with 50 ns shaping time) followed by an analogue pipeline to hold the data until a triggering decision has been made, with read out through an analogue deconvolution filter[1.10] to restore the signal to one bunch crossing. The product of this research is the APV (Analogue Pipeline Voltage mode) front-end readout chip[1.11][1.12].

1.3 The APV readout chip

Research and development has been underway on the front-end readout chip since 1993, the culmination of this work has lead to the CMS collaboration adopting a readout system based upon the APV chip series. The APV has gone through many iterations leading up to the production of three 128 channel versions; beginning with the APV6 in 1997, fabricated in the Harris AVLSI-RA Bulk CMOS process [1.13], followed by the APVM fabricated in the same process, but adapted to readout MSGCs (Micro-strip gas chambers). Then at a similar time the APVD, fabricated in the DMILL (Durci Mixte su Isolant Logico-Linéaire) process and finally the APV25 in 2000, fabricated in a commercial 0.25μm process[1.12]. The APV6, APVM and APVD have feature sizes of approximately 1 μm, with the APV25 much smaller at 0.25 microns. The APV employs MOSFET transistors (see section 1.4), which are more intrinsically radiation hard than other types of transistor to bulk silicon damage because they are fabricated in a thin layer on the surface of the chip. However, their degree of radiation tolerance is still dependent on many factors (see section 3.2), which were carefully considered in the

choice of technologies listed above. The work documented in this thesis is confined to studies of the APV6 and the APV25.

All three versions of the APV are essentially identical in operation, with the addition of some improvements to the APV25. The APV25 consists of 128 channels, each made up of a pulse amplifier and shaper, which feeds a 192 deep analogue pipeline capable of storing input pulses for up to 4 μs. The pipeline samples the output from the amplifier at 25ns intervals, storing the values until they are either overwritten on the next pass, or a level 1 trigger is received, when the data are marked for readout. On reading out the marked pipeline locations, the data are retrieved and output via a 128:1 multiplexer, in one of two modes: peak or deconvolution. Peak mode takes the value of the pipeline location at the peak of the pulse; thus only pulses that are separated by many beam-crossing intervals can be resolved efficiently. However, in deconvolution mode three consecutive temporal samples are required by the deconvolution filter, which performs a weighted sum on the data, enabling resolution of consecutive pulses. The pipeline can have up to 32 events tagged for readout at any one time. The addresses of the data to be read out are stored in a FIFO, following which the deconvolution filter, or analogue pulse shape processor (APSP), reads the data from the pipeline as soon as the previous event has been processed. Data taking and readout occur simultaneously.

Control of the various chip operation modes and bias settings is achieved via a standard I^2C[1.14] serial bus link. The output stream consists of a digital header, containing information about the status of the chip and also the pipeline location from which the data were retrieved. Following this the analogue levels are retrieved from all 128 channels.

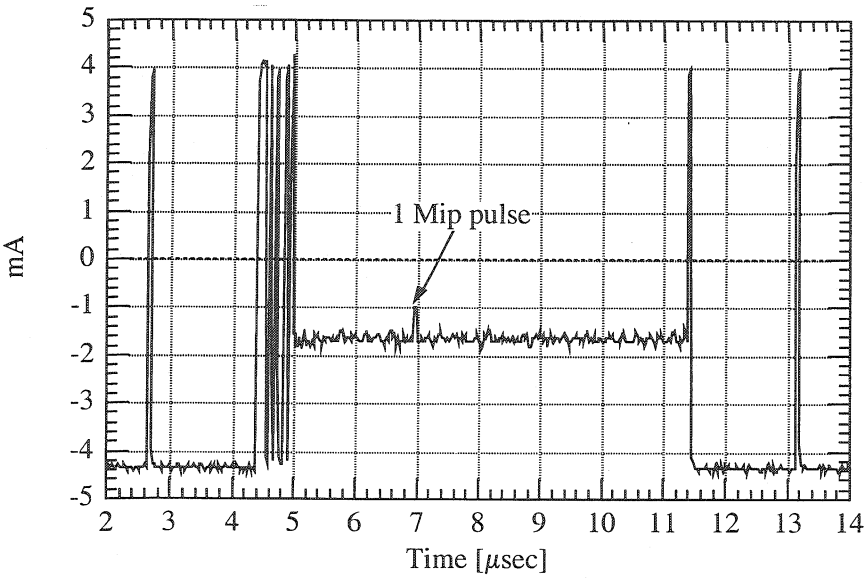

Figure 1.4 Example data frame from APV25, a signal of 1 mip can be seen

On the above scale one can see the digital header at about 4.4 μs, preceding the analogue data. Before and after the frame one can see the tick marks, which are output every 1.75 μs in time with the readout cycle of the APSP. Each APV has an address that can be set, which enables the I^2C bus to communicate separately with individual chips. Another important feature is the internal calibration system. This creates an input signal internally which can be stored in the pipeline and analysed in the output stream. All 128 channels can be driven to enable testing of amplifier gain and pulse shaping uniformity. This feature in particular is essential for the testing procedures shown in Chapter 2.

Figure 1.5 APV25-S1 the latest version of the APV

Figure 1.5 shows that APV-S1. The chip area is dominated by the analogue circuits: amplifier, shaper, pipeline, APSP and multiplexer. The four digital control circuit blocks are situated along the bottom edge of the chip, as shown in Figure 1.5. From left to right these are the bias registers, pipeline pointers, FIFO readout buffer and the control logic.

1.4 MOSFET Transistors

In order to understand how the APV will behave under the conditions of CMS we must understand how the component parts, the transistors, will behave. In order to do this we must first understand how they operate.

The type of transistors used in both the analogue and digital circuits are Metal Oxide Semiconductor Field Effect Transistors (MOSFET usually shortened to MOS). MOS transistors are three terminal devices, which utilize a conducting channel between two terminals, the *source* and *drain.* The third terminal, the *gate,* is used to modulate the conductance of the channel. MOS transistors are classified by the type of carriers that flow in the channel; an n-channel device is one in which the conducting carriers are electrons and a p-channel device, holes. There is also an additional terminal, the *bulk,* which connects to the substrate on which the transistor is fabricated. Varying the voltage on this contact may affect the threshold behaviour. The drain, source and bulk

contacts of the MOS transistor are not ohmic contacts since they are semiconductor to metal junctions(usually Schottky diodes).

Figure 1.6 shows the cross-section of a simple NMOS and PMOS transistor.

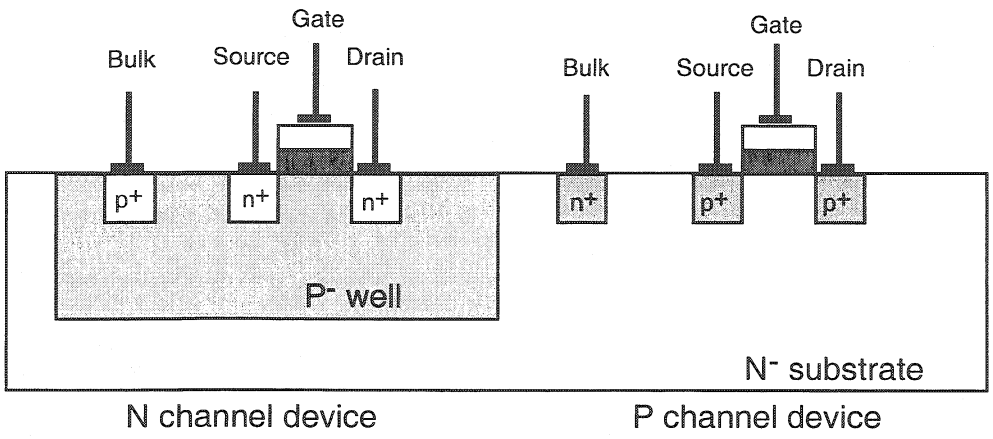

Figure 1.6 Cross-section through an ideal NMOS and PMOS transistor.

This is typically how they would be constructed on an n^- substrate, with the PMOS implanted directly into the substrate and the NMOS implanted into a p^- well, these are usually called p-well processes. For more modern technologies such as the deep sub-micron process, both N and PMOS are constructed on wells, which are implanted into the substrate. These are called dual-well processes.

The fabrication process of such devices is long and complicated; detailed explanations can be found in[1.15]. After all the doping, oxide growth and etching, which goes into such a process, in reality the devices do not look quite so well defined as in Figure 1.6. A scanning electron microscope picture of a cut-away cross-section of an NMOS transistor is shown in Figure 1.7[1.16]. The n^+ diffusions are the light-grey areas in the dark substrate. In the centre of the picture is the gate and to the left is a metal contact to the n^+ diffusion.

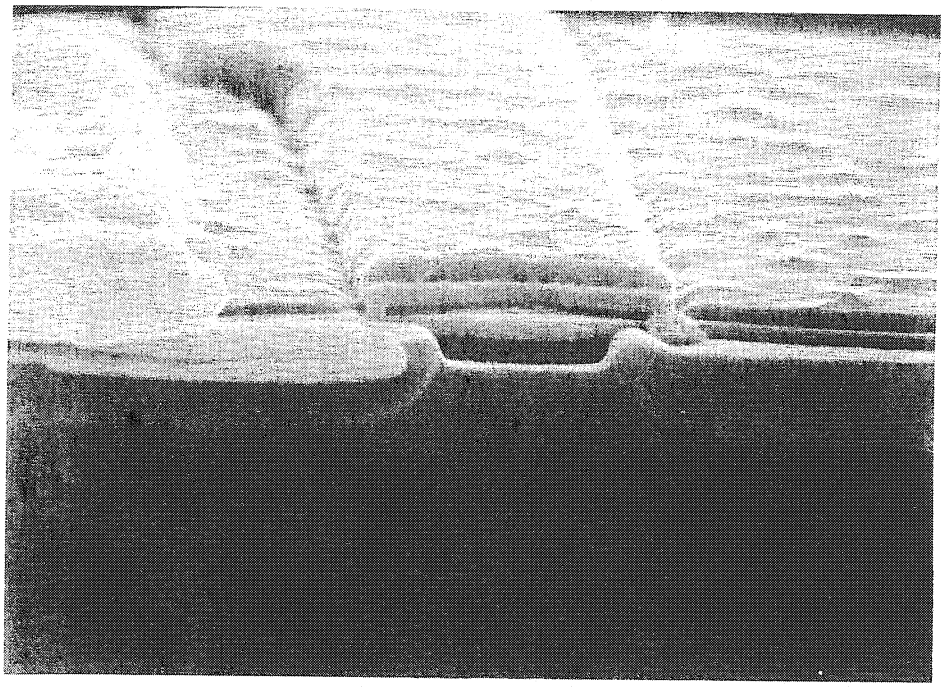

Figure 1.7 Scanning electron microscope photograph of the cross section of an NMOS transistor.

1.4.1 MOSFET operation

The basic operation of a NMOS device can be described as follows. Begin by considering a small voltage drop across the gate and the bulk (Figure 1.8). The gate oxide is an insulator and allows the gate to act as a capacitor producing an electric field between the gate and the bulk. This field drives holes at the surface of the bulk away and gives rise to a depletion region at the surface. This combines with the natural depletion regions of the source and drain implants.

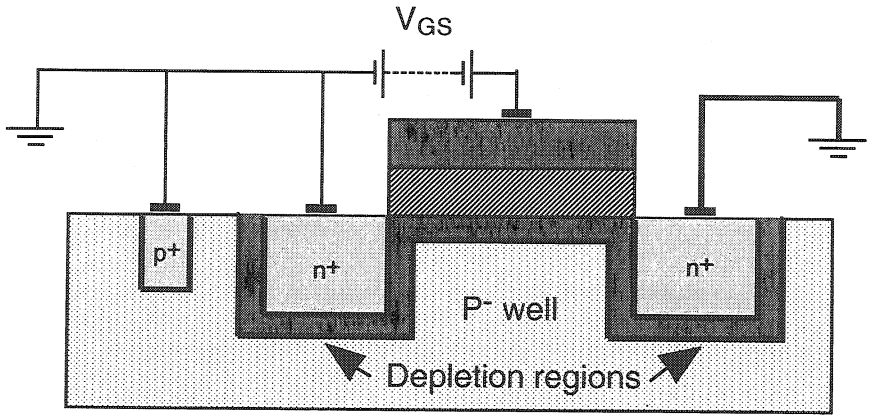

Figure 1.8 NMOS transistor under initial bias conditions.

The depth of the depletion region in the channel is given by equation [1.1].

$$x_d = \sqrt{\left(\frac{2\varepsilon_{si}|\Phi_s - \Phi_F|}{q_e N_A}\right)} \qquad \left(|\Phi_S - \Phi_F| \geq 0\right) \qquad [1.1]$$

where:

ε_{Si} = dielectric constant in silicon = $11.7\varepsilon_0$.

q_e = electronic charge.

N_A = acceptor concentration.

Φ_F = semiconductor equilibrium electrostatic potential

Φ_S = surface electrostatic potential

Following on from this one can show that the total immobile charge due to acceptor ions that have been stripped of their holes is given by:

$$Q_B \cong -\sqrt{2q_e N_A \varepsilon_{si}|\Phi_s - \Phi_F|} \qquad [1.2]$$

As the voltage at the gate is increased more holes are driven away and electrons are attracted into the surface channel from the source and drain n^+ implants. In a somewhat simplified picture, as this continues eventually more electrons than holes are present at the surface. This is opposite to what is expected in a p-type semiconductor and so the surface is said to be inverted (Figure 1.9). Now there is a channel of conduction electrons between the source and the drain allowing a current to flow when there is an applied voltage between them. The charges in this inverted surface, Q_I, are all within a very small distance of the surface, and are therefore considered to be a charge sheet. The threshold voltage of the transistor is the gate voltage at which current just starts to flow in the channel, the more accurate definition of which is given later.

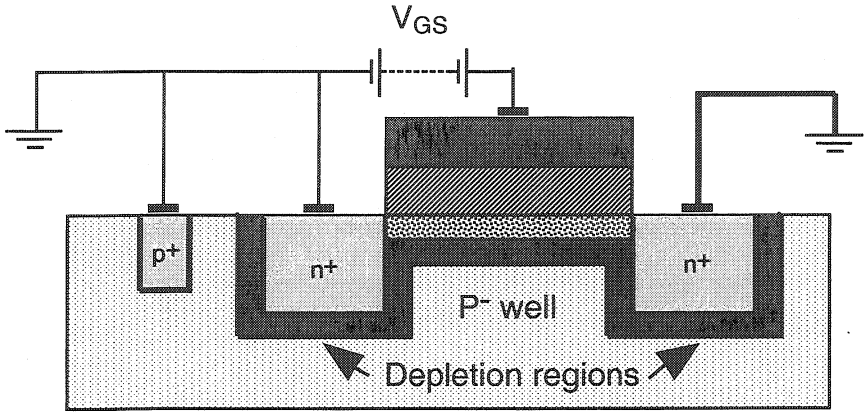

Figure 1.9 NMOS transistor during inversion.

As the gate voltage causes the surface potential to increase from $-\Phi_F$ to zero, only the junction leakage current flows. When the gate voltage increases, Q_B increases steadily, and Q_I is still very small. However, due to diffusion of this charge a small current will flow. When the diffusion current is larger than the leakage current, the channel is in a state of *weak inversion.* A further increase in the gate voltage causes the surface potential to increase until, at approximately $\Phi_S = +\Phi_F$, Q_I increases more rapidly with increasing gate voltage, and the current is dominated by charge drift. At this point the channel is in a state of *strong inversion.* The gate voltage required to achieve this is defined as the threshold voltage, V_T which can be shown to be given by:

$$V_T = (\Phi_{GB}) + \left(-2\Phi_F - \frac{Q_B}{C_{ox}} \right) + \left(-\frac{Q_{ss}}{C_{ox}} \right) \qquad [1.3]$$

Φ_{GB} = gate to bulk contact potential

C_{ox} = gate to channel capacitance

Q_{ss} = silicon surface trapped charge

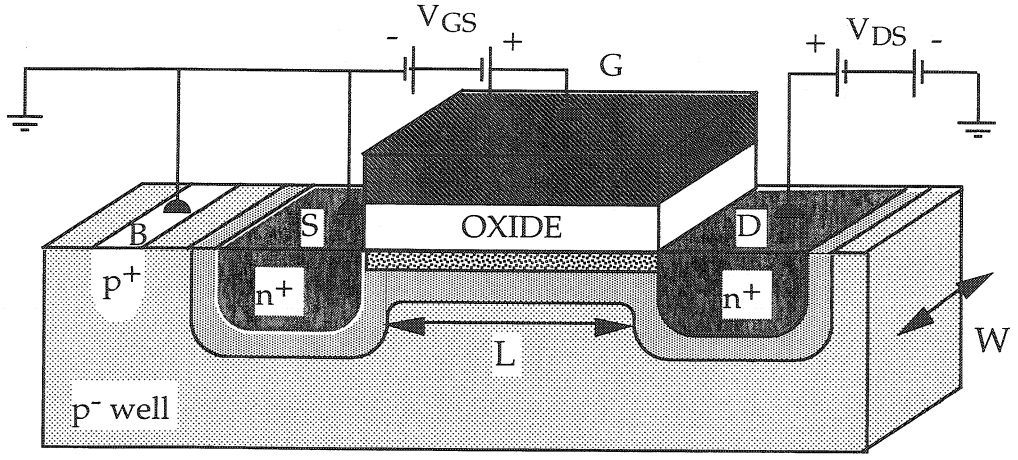

Figure 1.10 An n-type transistor biased in strong inversion.

Figure 1.10 shows an NMOS transistor biased in strong inversion. When $V_{GS} > V_T$ a current flows through the channel under the applied drain-source voltage V_{DS}. However, the channel does not behave like a simple ohmic conductor with the current proportional to the applied voltage. Instead there are two regions of operation; non-saturation - when the channel extends fully between the drain and source and I_{DS} increases rapidly under increasing V_{DS} ; and saturation - when V_{DS} is high enough so that it causes the depletion region surrounding the drain to extend to the surface, and the

channel is *pinched off,* quenching the current increase almost to zero[1.17]. The current in these two regions is given by equations [1.4] and [1.5].

$$I_{DS} = \frac{\mu C_{OX} W}{L}\left[(V_{GS} - V_T) - \frac{V_{DS}}{2}\right]V_{DS} \qquad [1.4]$$

$$I_{DS} = \frac{\mu C_{ox} W}{2L}(V_{GS} - V_T)^2 \qquad [1.5]$$

where μ is electron mobility and all other terms are defined in Figure 1.10. Equation [1.5] shows no dependence on V_{DS} and so we would expect the current to have totally saturated. However, in reality I_{DS} is not quite constant with respect to V_{DS} after saturation and so we must modify the equation to account for the change in the channel length due to the pinching off of the end of the channel. This is achieved by including a channel length modulation parameter, λ.

$$I_{DS} = \frac{\mu C_{ox} W}{2L}(V_{GS} - V_T)^2(1 + \lambda V_{DS}) \qquad [1.6]$$

where:

$$L_{mod} = L(1 - \lambda V_{DS}) \qquad [1.7]$$

Figure 1.11 shows the two regions of operation, including the channel length modulation parameter gives rise to the increasing current in the saturation region.

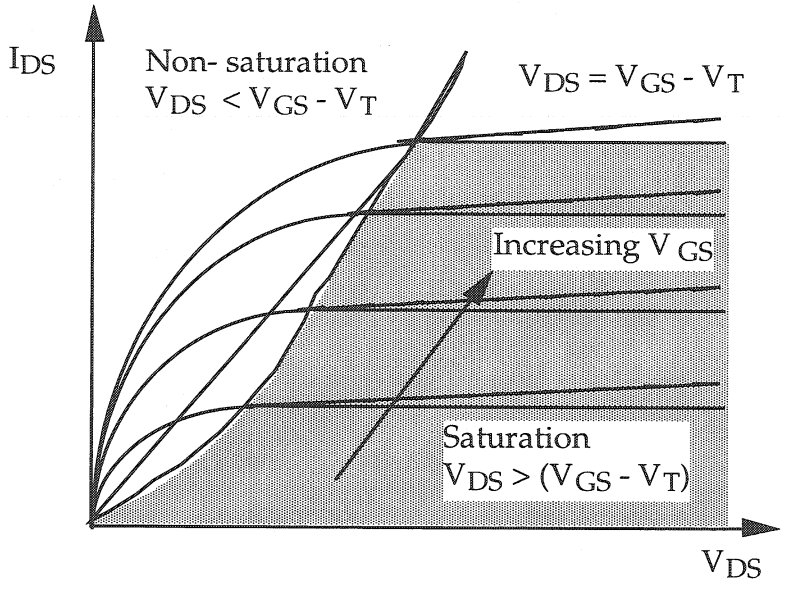

Figure 1.11 Operation of a transistor in the saturation region. The constant plateau represents λ=0, rising plateau represents λ>0.

The transfer characteristics are commonly represented by a quantity called the transconductance which is given by equation[1.8] for non saturation and equation[1.9] for saturation.

$$g_m = \frac{\partial I_D}{\partial V_{GS}} = \frac{I_{DS}}{[(V_{GS} - V_T) - V_{DS}/2]} \qquad [1.8]$$

$$g_m = \sqrt{\frac{2\mu C_{OX} W}{L} I_D} \qquad [1.9]$$

1.4.1.1 Subthreshold characteristics

The equations above assume that the transistor is in a state of strong inversion. By sweeping V_{GS} one moves though different states of the channel, from weak inversion to strong inversion. If a constant value of V_{GS} is applied during this process, one can see how the current is modulated by the state of the channel. The current in a real Harris transistor, measured as described in chapter 4 is shown in Figure 1.12. Above a certain voltage, V_{ON}, the transistor is in a state of strong inversion. Below the threshold voltage (-0.98 V in this case), the transistor is in a state of *weak inversion*. Between these two regions is a less well defined region, called the moderate inversion region.

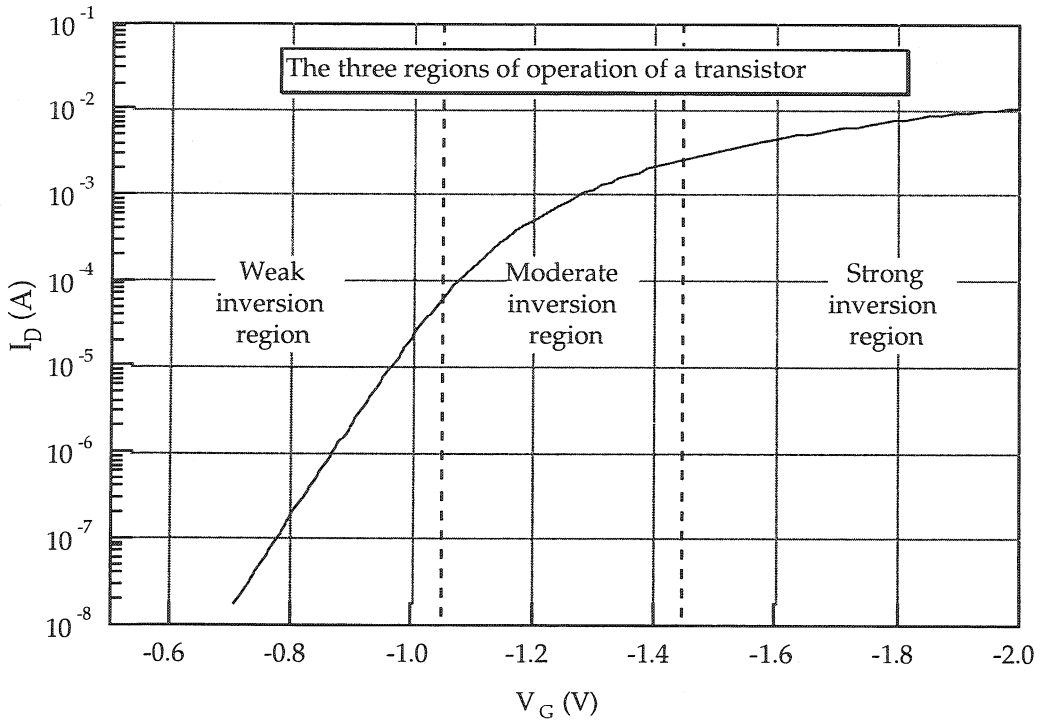

Figure 1.12 The three regions of operation on an MOS transistor. Data is from a 2000 μm × 1.4 μm PMOS Harris transistor, measured as described in chapter 4.

The drain current of the transistor in weak inversion can be described by equation [1.10].

$$I_{DS} \approx \frac{W}{L} I_{D0} \exp\left(\frac{q_e V_{GS}}{nkT}\right) \qquad [1.10]$$

and the transconductance can be calculated as:

$$g_m = I_D \frac{q_e}{nkT} \qquad [1.11]$$

In most analogue circuits the transistors are operated in the strong inversion region and in digital circuits the transistors act as switches, either fully on or off, therefore the exact behaviour in the weak inversion region is not of such importance.

1.4.2 Noise in MOS transistors

Noise in transistors is generally expressed as fluctuations in the current. There are three basic noise sources in an MOS transistor:

- Thermal noise
- Shot noise
- 1/f noise, or flicker noise.

They are uncorrelated[1.18], therefore the total mean square noise is the sum of the individual mean square noise values.

1.4.2.1 Thermal noise

Thermal noise is caused by thermal excitation of carriers in the channel in thermal equilibrium. The thermal noise current spectral density in a resistor is given by:

$$i_d^2 = \frac{4kT}{R} \qquad [1.12]$$

Where k is the Boltzman constant, T is the absolute temperature and R is the resistance. However, the channel of a transistor is not as simple as a resistor and must be considered carefully in order to derive the noise[1.18]. The dominant thermal noise in a MOSFET is from the channel and is given by:

$$i_d^2 = 4kT\gamma g_m \qquad [1.13]$$

where γ is a very complex function based on the transistor parameters and bias conditions. For the technologies in question, γ is situated between 0.6 and 1[1.18].

1.4.2.2 Shot Noise

Shot noise is due to the fact that charge carriers are discrete and behave in a random manner. Since large numbers of carriers make up the steady state current this current is simply the mean of the distribution of individual carriers, while the standard deviation is shot noise. The mean square current fluctuations may be written as;

$$i_d^2 = \gamma 2 q_e I \qquad [1.14]$$

where γ is equal to 1 for a purely random process, however in a real transistor, carriers are not emitted individually, but in groups, and γ will be less than 1. In most MOS transistors shot noise is much smaller than the associated thermal noise and can be ignored when calculating the expected noise response.

1.4.2.3 1/f Noise

1/f noise dominates at low frequencies and is named so because the power of the noise varies with the inverse of the frequency. The dominant model in MOS transistors suggest that the main source is the fluctuation of carrier charge density due to exchange of carriers between the channel and interface traps and also charge fluctuation between surface states. The form of 1/f noise is given in equation [1.15].

$$i_f^2 = \frac{K_f}{(C_{ox})^2 W L f^{\alpha}} \qquad [1.15]$$

where K_f is a parameter that depends on technology and device construction and C_{ox} is the gate capacitance per unit area.

In general at low frequencies 1/f noise tend to dominate and at higher frequencies thermal noise takes over. The general noise spectrum from an MOS transistor is shown in Figure 1.13.

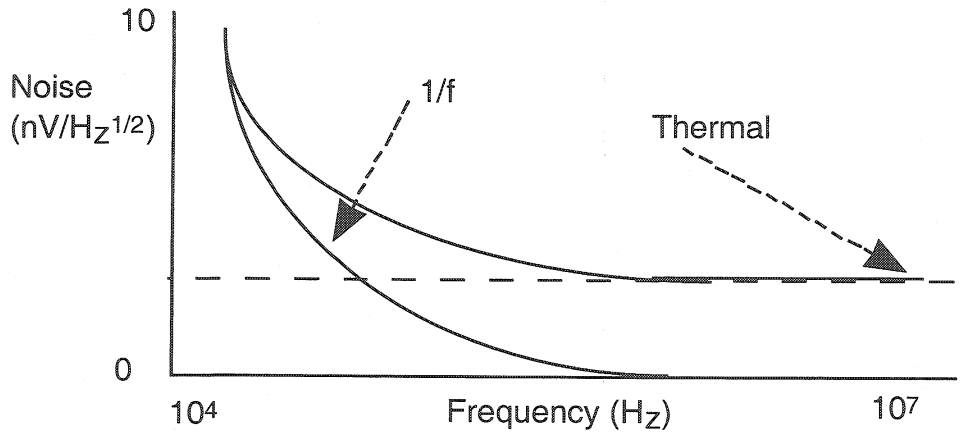

Figure 1.13 General noise spectrum for an MOS transistor.

The point where the two regions overlap is called the corner frequency. Since PMOS transistors have a much lower corner frequency than NMOS of a similar size, the input transistor at the pre-amplifier stage in the APV25 chip was chosen to be a PMOS transistor.

1.5 Conclusions

When designing front-end systems such as the one for the CMS tracker, one must pay attention to all of the requirements such as; low noise, low power, fast signals, radiation tolerance and other general functionality, and attempt to design a system which is a compromise between all of them. This chapter has covered only a small part of a wide background of information in an attempt to give the reader a flavour of the technical considerations involved in employing such technologies.

Chapter 2

Testing The APV Chip

2.1 Motivations

After fabrication, the APV chips will be received on eight-inch diameter silicon wafers, each containing ~400 individual chips. Before these chips can be sawn from the wafer and distributed for assembly into the detector modules, one must establish whether they are fully functional, partly functional or useless. The percentage of chips per wafer that is fully functional is called the yield and it is important that the yield is high to save on fabrication and testing costs. Therefore, it is also important that wafer screening is highly reliable and provides assurance that no possibly useful chips are failed unnecessarily. Conversely, it is important that no bad chips are passed and then found to be faulty after distribution.

It is possible to perform rudimentary tests upon wafers by hand using a manual probe station. In this case the tester must manually control the individual probes, making sure that they make good contact with the chip contact pads at each site. Once connections have been made the tester can then perform a number of tests such as checking currents drawn by the power rails or looking for bad channels. Manual testing can cover a wide range of tests and would be reliable. The tester knows the exact conditions of each input and output at any time and can form a good overview of the chip performance. There are obvious disadvantages with manual testing:

- It is very time consuming; when the time comes for mass screening of the final product there will be roughly 100,000 chips to test. This would not only be highly tedious but also impractical.
- It is not sufficient only to test the chips, but it is also important that information which the tester gathers is stored for future reference. Certain chips may contain various non-fatal errors such as bad channels or pipeline locations.

- To make the best possible use of these chips it will be necessary to provide detailed information of chip errors. Without the use of a computer controller, it would be extremely difficult and time consuming to store all the necessary statistics and information.

These are the main driving forces behind the development of an automated wafer probing setup. An automatic probe station can be programmed to enable it to scan through a whole wafer without the need for intervention from an operator. This alone speeds up the screening process considerably. To further increase efficiency, a controlling computer can implement the full set of tests for each chip. As the setup performs each test on the chip it stores all the required statistical information in a file on the computer. This file can then be retrieved for later reference.

2.2 Development of The System

The system has been under development since 1996. Previous iterations of the APV were developed in the Harris AVLSI-RA CMOS process, which used 4 inch silicon wafers. The APV6 was the first generation to be tested by this system, which was originally capable of screening a whole 4 inch wafer, containing 61 chips, with a minimal set of the following tests:

- A check that the power rails are drawing the correct currents.
- An I^2C[1.14] bus read and write test.
- A check for a good analogue output stream.
- A check of the APV addressing.
- A check that the analogue level adjustment can be set correctly.
- A check that the calibrate pulse heights are good.

This system had been used to test 8 APV6 wafers from one batch (D40208), a total of 488 chips. It was at this time that I began to develop the system. Before any more chips were tested about 8 months work went into redesigning and rebuilding the system. A great deal of the time spent was on becoming completely familiar with the system hardware and testing methods. Once a full understanding had been reached it was possible to diagnose errors in the existing software. The remaining time was spent correcting these errors, incorporating more important tests and developing a

comprehensive format in which the data could be stored for later statistical analysis and summary. The system was semi-automated, which allowed testing procedures to be monitored for the first few wafers. This provided confidence in the system before the process was fully automated.

By December 1998 the first set of improvements had been completed and the system was used to test 25 wafers from 2 batches, a total of 1,525 APV6 chips. Following this the system was modified in order to test the new APVM, which meant slight changes in the interfacing to the chip. In May 1999 the first batch (DAP4P) of APVMs were tested and a second batch (DAP9D) in July 1999. In total 17 wafers, containing 1037 chips, were tested.

At this point work had commenced on the APV25 and there were no untested wafers. By the beginning of 2000 the APV25-S0 had been completed and work commenced to convert the setup to test the new chip. Improvements were made and more tests added. Testing of the APV25-S0 began in July 2000 and since the chips had already been cut into individual die they had to be mounted, aligned and tested one by one. Testing of 501 chips was completed in a total of 30 days with an average time of 5 minutes per chip. Most of this time was spent in mounting and aligning the chips.

The analysis of the data from all of the chips tested is presented in this chapter.

2.3 Equipment and Test Criteria

The system is described as it was for the testing of the APV6 and APVM, for which the changes made between the two iterations were extremely small and can be ignored. The improvements made for the testing of the APV25 are discussed separately.

2.3.1 APV6 and APVM Equipment

Figure 2.1 is a schematic representation of the test setup. Control and testing is performed by a Macintosh using LabVIEW[2.1] software via VME interface. The clock and trigger signals to the APV are provided by a SEQSI sequencer. They are buffered by optical relays on the card interface board. The analogue output from the APV is taken directly to the ADC, and may also go to a scope. The OUTE signal is taken to a discriminator to achieve a fast pulse (~5ns) representing the leading edge. This signal is used to start the logic that triggers the ADC.

Data cannot be read from the APV6 into the computer at the full rate of triggers (typically 100 kHz), therefore a veto is required. This is achieved by setting and un-setting bit 6 of the SEQSI when the software is ready to accept data. The resulting pulse is used to set a latch, which then only allows the subsequent OUTE to set the ADC gate.

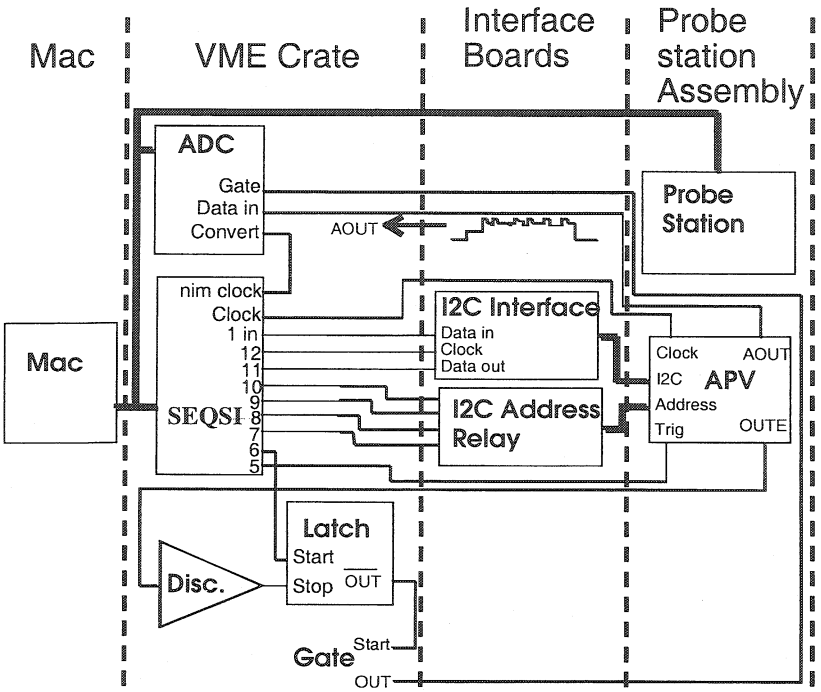

Figure 2.1 Schematic of the probe-station test setup.

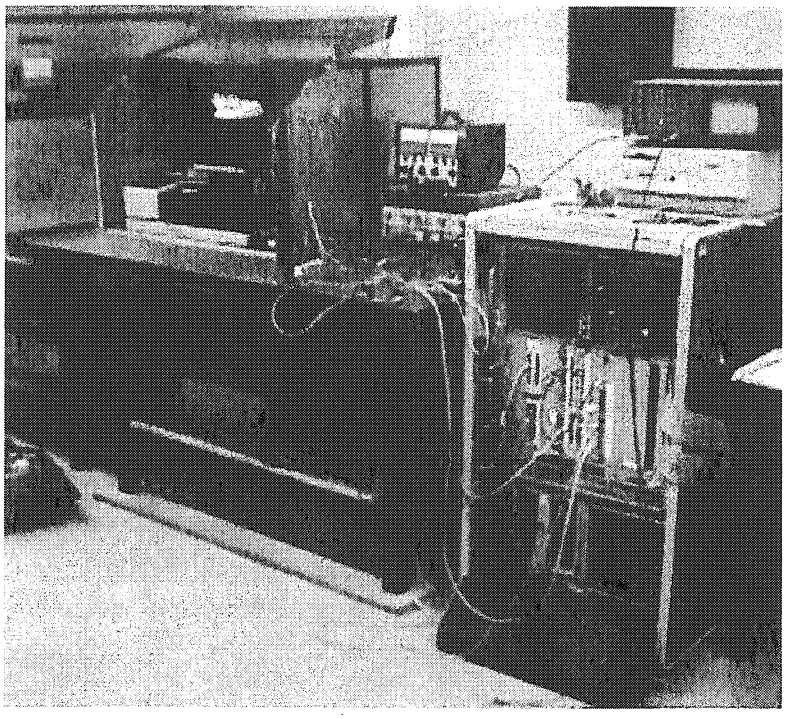

Figure 2.2 Probe station and VME interface.

The I^2C signals and addressing are provided by the SEQSI, which can run the I^2C interface at a bit rate of a few kHz. However, the APV6 can accept data at about 400 kHz. Significant improvements in speed could be achieved by incorporating a VI^2C interface board. All I^2C transactions are software generated.

Figure 2.3 shows the custom designed probe card, which is used to directly interface with the APV.

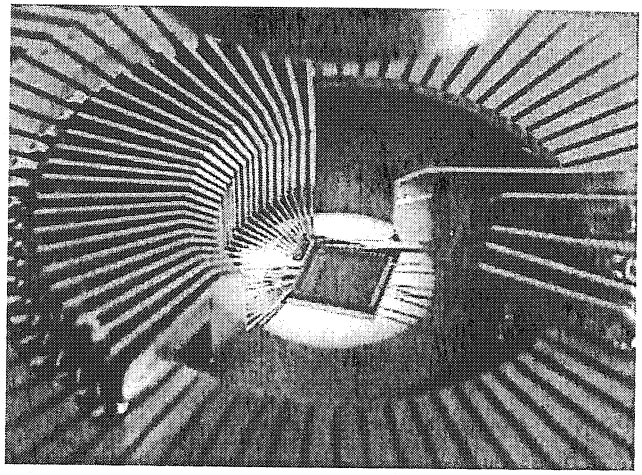

Figure 2.3 Probe card interfacing an APV25.

2.3.2 Recent Modifications

The modifications and improvements that were made in order to test the APV25 were as follows:

- A new interface card was designed in order to convert to the new geometry of the APV25.
- Various modifications were made to the control code to adapt the testing to the specifics of the new chip. A new set of bias register addresses and data had to be used.
- Improvements were made to the saved data format, making it more comprehensive.

2.3.3 Preparing the APV for testing

Before any tests are performed on a chip, a looped trigger sequence is set up, which is sent to the trigger input of the chip. It is a simple sequence that resets the chip and then reads out the data from one pipeline location. This causes the chip to produce a continuous output stream of all 128 channels, containing the normal levels when there

have been no input pulses. These so-called channel pedestals are subtracted from the channel values during the retrieval of pulse data.

Following this the chip is programmed with a standard set of register values. Table 2.1 shows the register settings for the APV6 and APVM, and Table 2.2 the APV25 taken from the user manuals[2.2][2.3].

Register	Value	
	APV6	APVM
Mode	-	-
Latency	119	119
IPRE	111	111
ISHA	88	88
IPSP	83	83
ISFB	43	43
VPRE	150	150
VSHA	50	50
VADJ	-	-
VCAS	0	0
CLVL	152	152
CSKW	-	-
CDRV	-	-
VCMP	-	150
IREF	-	255

Table 2.1 Programmed register values for APV6 and APVM.

Register	Value
Mode	-
Latency	119
MUXGAIN	-
IPRE	110
IPCASC	60
IPSF	50
ISHA	50
ISSF	50
IPSP	80
IMUXIN	50
ICAL	100
VFP	0
VFS	50
VPSP	-
CSEL	-
CDRV	-

Table 2.2 Programmed register values for APV25.

The Mode value depends on the mode of operation; all chips were tested in both Peak and Deconvolution mode. VADJ for the APV6 and APVM is set to produce an analogue baseline of 180μA, the equivalent register in the APV25 called VPSP is set to produce a baseline of 2 mA (this is ¼ of the way up from the lowest possible value). Once this has been achieved the value of VADJ is stored in the output data file. Test 11 makes sure it is possible to set the desired VADJ or VPSP value. CLVL is set to 152, which corresponds to a calibrate pulse amplitude of 9.12 fC, ~ 2mips, and ICAL, the equivalent register in the APV25, is set to 100, which corresponds to ~ 10 fC. CSKW (CSEL for the APV25) and CDRV are used to create the calibrate pulse. The values are cycled in order to map out a pulse. VCMP and IREF are registers that were introduced in the APVM and so are unused for the APV6. All other registers set the bias voltages and currents.

2.3.4 Test Criteria

The total number of individual tests is 18. A list of the test names and brief description of each follows:

1. No Probe Contact

Built into the probing mechanism is one pin whose job is simply to sense when contact is made with the chip. This enables the probing mechanism to know when the pins are in contact with the chip, at which point the pins are lowered a further set distance (called the over-travel) to ensure good contact has been made by all the pins. In the event that no contact is sensed, error 1 is produced and the system moves onto the next site. In the cases where this occurs it is necessary to go back and lower the probes manually enabling the site to be tested. Experience suggests that this is not a common error.

2. I^2C Write Test

To ensure that the I^2C communication bus will operate correctly, it is tested to see if the chip is receiving the intended data. During normal operation of an I^2C write transaction the chip sends an acknowledgement when it receives information from the bus. In the event that there is no acknowledgement from the chip, error 2 is produced.

3. I^2C read

This is identical to test 2 except that an acknowledgement is awaited after an I^2C read.

4. No OUTE

One of the outputs from the chip is the out enable (OUTE), which produces a pulse in time with the start of the analogue data stream. This provides a trigger for an analogue to digital conversion of the data. If this output is either missing or out of sync, then error 4 is produced.

5. No AOUT

If the analogue output stream is missing then the ADC will return a value of 0 for every channel. In this case error 5 is produced.

6. Stuck bits in I^2C registers

It is possible that some bits in the I^2C registers may be stuck either high or low. This may not be picked up in an I^2C read test. By writing the value 10101010 to the registers, comparing the read value with the written value and then repeating this with the value 01010101, it is possible to determine whether there are any stuck bits.

7. I^2C read write compare

Data written to the chip in an I^2C write transaction is compared with data received from the chip in the I^2C read transaction. If these data do not match then error 7 is produced. This test is included mainly because, at this point in the procedure, the chip is programmed with a set of default data, which is required for the chip to run correctly during the remaining tests. If this test is failed then there could be a reason for the other tests to fail also. Experience shows that if the chip passes tests 2, 3 and 6, it will also pass this test. However, if it failed test 6 then it is possible that this test may fail.

8. Bad currents

The APV has been designed to draw particular currents during normal operation. These currents are checked and if they lie outside a predetermined range then error 8 is produced.

9. Failed some addresses

See 10.

10. Failed all addresses

The address for the chip is scanned through all possibilities. During each step the I^2C bus addresses the chip and if there is no response, the chip has failed that address. A note is made of how many addresses fail; if more than 0 fail then error 9 is produced, if they all fail then error 10 is produced.

11. Can't set VADJ

The default analogue pedestal heights can be adjusted by the register called VADJ to optimise the performance of the chip. This register is ramped up from 0 until the analogue output stream reaches 180μA. If this cannot be achieved then error 11 is produced.

12. Some bad pedestals

During continual triggering of the chip the analogue data stream is retrieved by an ADC. While there are no input pulses on all 128 channels, the data stream is recorded 1000 times. An average is taken for each channel giving good values for the channel pedestals, which are recorded for later use in pedestal subtraction. If any of the channels have a pedestal value outside a predetermined range then error 12 is produced. Channels 1-4 and 128 are discarded since they are known to have lower than average pedestals, which is due to an unserious flaw in the design. Also channels 111-114 are discarded from the tests because the inputs to these channels were probed and thus their noise was much large than normal. Occasionally one finds channels that have somehow been shorted to ground or power.

13. Noisy analogue output

Again the analogue data stream is retrieved, but this time from throughout the pipeline. Each pipeline location is read out 50 times and averaged. A two-dimensional array containing the pedestal values from all of the 160 pipeline locations for each of the 128 channels is compiled. The average and standard deviation of the pipeline pedestals from each channel is then calculated. If the sum of the standard deviations from all 128 channels is greater than a pre-determined value (roughly 1,000,000, but to be refined as required), then error 13 is produced.

14. Some bad channels

This is the same as test 12 except it tests all the channels in all pipeline locations. If more than a certain number of pipeline locations in a channel fail the test then that channel is failed.

15. Some bad pipeline locations

Again the analogue data stream is retrieved, but this time the data are retrieved from throughout the pipeline. Each pipeline location is read out 50 times and averaged. A two-dimensional array containing the pedestal values from all 160 pipeline locations for each of the 128 channels is compiled. The average and standard deviation of the

pipeline pedestals from each channel is then calculated. If there are any pipeline locations which deviate from the average by more than 5 standard deviations (this value will be refined as required), then error 15 is produced and a note is made of the channel and pipeline location.

16. Channels with low gain

The chip is run in calibrate mode and each channel driven with a test pulse in both peak and deconvolution modes. The pulse heights in each mode are recorded for all channels. The average and standard deviation of the pulse heights is calculated. If any channels have pulse heights more than 5 standard deviations (this value will probably be refined as required) away from the mean then error 16 is produced.

17. Some bad pipeline addresses

Each pipeline address is read from the digital header and compared to the expected code. If any are incorrect then this error is flagged.

18. Failed all pipeline addresses

If the procedure in 17 reveals more than an acceptable number of discrepancies. (Note that no errors are acceptable so this test is redundant except for information purposes.)

2.4 Test Data Storage and Retrieval

After each chip is tested, data are stored in arrays and subsequently saved to disk. These data are also compiled into a single sheet (Figure 2.4), which can be printed and sent with the chip to the user. This sheet includes seven plots, which are valuable information for anyone using the chip. If a situation arises where the data sheet does not clearly show the required information, the raw data can be retrieved and scrutinised. Analysis of data from all the tested chips has been performed and is presented in this chapter.

2.4.1 Example APV6 Data Sheet

An example data sheet taken from one APV6 chip from wafer 14 of run D40208A is shown in Figure 2.4. All values for pedestals and pulse heights are in ADC counts. 1 ADC count = 2.42 μA The Pedestals are all stored after common mode subtraction. So the average of all the channels is 0. An ADC count of 40 corresponds to ~ 95 μA, ~ 2 MIPs.

The data sheet contains four plots: "Pedestals" shows the average channel pedestals. "Channel Cross-section" contains 160 plots showing the channel output values after pedestal subtraction, one for each pipeline location. "Pipeline Cross-section" contains 128 plots of the pipeline pedestals, one for each channel. The combination of these two provides a simple way of looking at the two-dimensional array, which contained these data. In the ideal case both of these plots would exhibit only values of zero since these data are taken after pedestal subtraction. "Pulse Heights Peak Mode" and "Pulse Heights Deconvolution Mode" contain one plot showing the calibration pulse heights in the two modes of operation. "Pulse Shapes Peak Mode" and "Pulse Shapes Deconvolution Mode" show the shape of the maximum, minimum and average pulses taken from every channel in both modes; the time scale is in nsec. From this one can inspect the peaking time and width of the pulse. The inversion of the pulse shape between the peak and deconvolution is just a feature of the software that was never amended, which has no adverse effects on the test procedure. The large undershoot, which is visible in the deconvolution pulse, is a feature of the APV6 and has been improved vastly in the APV25.

The data sheet is reasonably representative of a fully operational chip with normal channel pedestals. One feature to note is the apparent regular pattern throughout the channels, this is due to the reordering of the channels by the multiplexer at the output stage. Since the pedestals have smoothly varying heights before the multiplexer reorders them, the process of chopping up and interleaving gives rise to the visible pattern. The spread of the pipeline pedestals can be seen clearly on the plots titled 'Pipeline Pedestals Cross-section' and 'Channel Cross-section'; i.e. each channel bin contains all 160 pipeline locations and the full spread is less than 10 ADC counts. Comparing this to the size of an average pulse height, roughly 110 ADC counts, one can see that there is only a maximum deviation of ~5%. There are three channels of interest: 111,112 and 113. These channels appear to have a less uniform pipeline. The effect that this would have on data must be investigated and some threshold must be set which puts an upper limit on the allowable spread throughout the pipeline. The plot of the pipeline shows that there are no outstanding bad locations, but one is able to see that the non-uniformity of channels 111,112 and 113 is spread throughout the pipeline. The spread of the pulse heights is acceptable, again roughly 10%. The pulse shapes are also within acceptable limits.

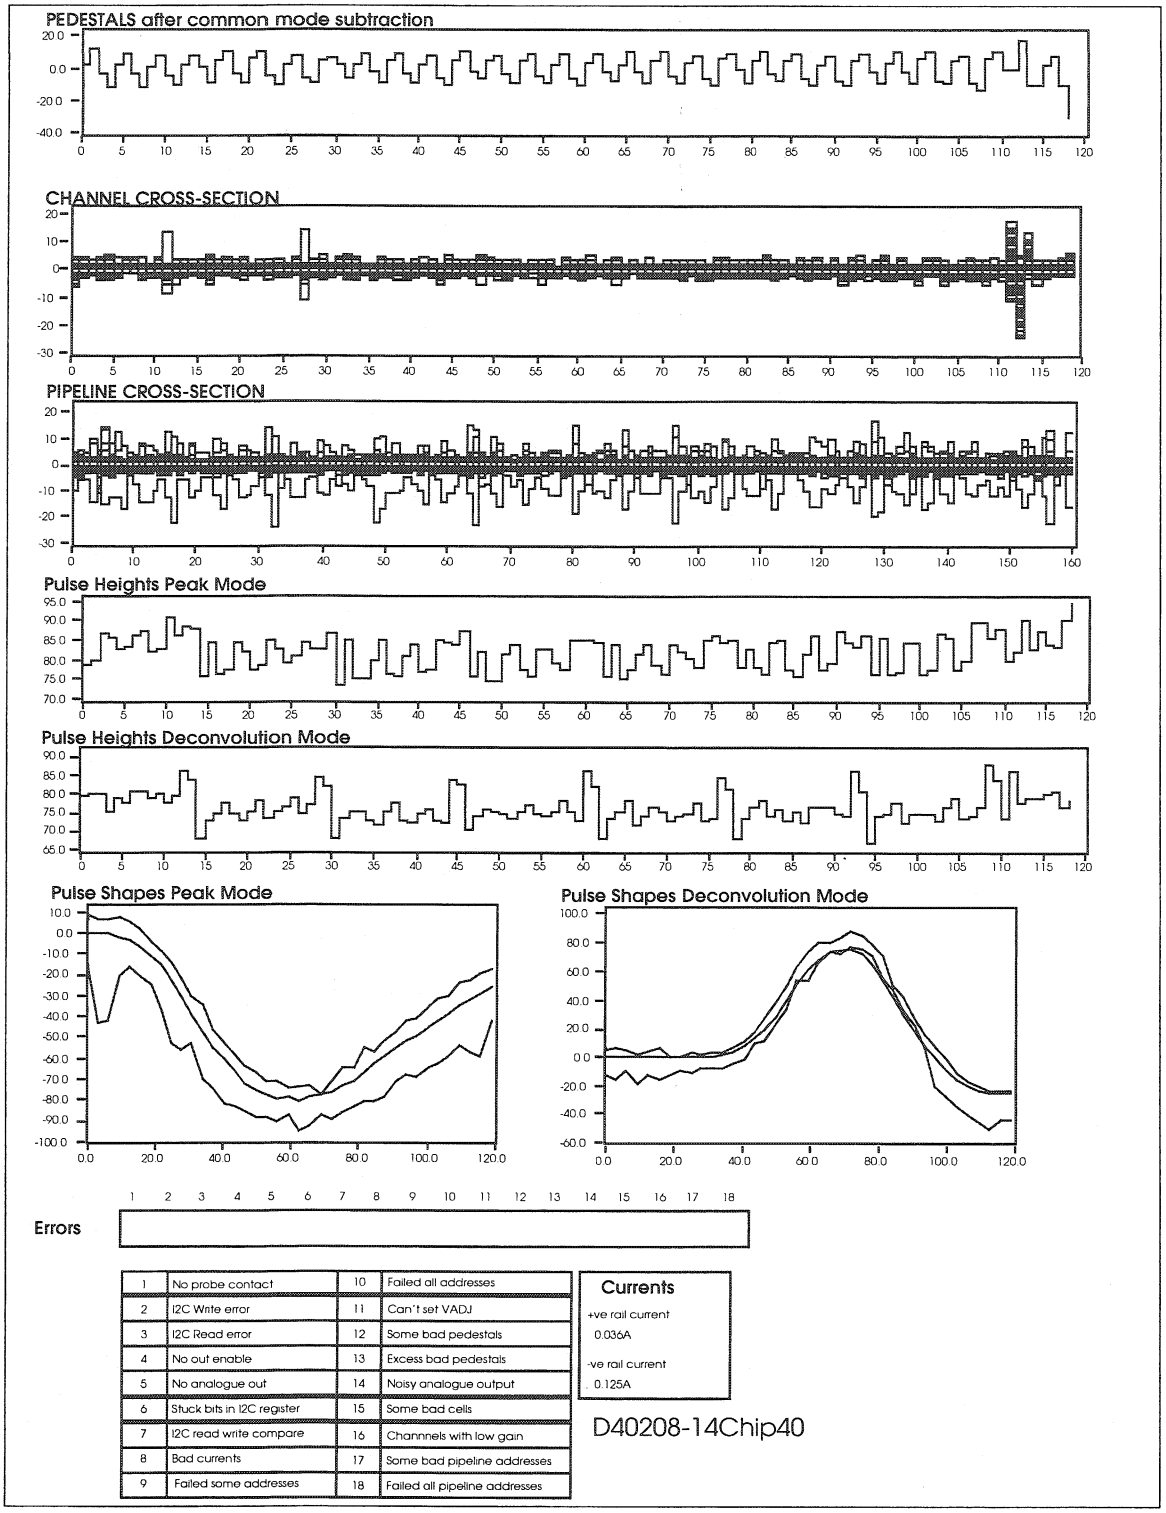

Figure 2.4 Data sheet from chip 40, wafer 14 from the first APV6 run.

2.4.2 Example APV6 Wafer Map

For every wafer tested the software produces a map displaying the recommended chip failures and passes. Figure 2.5 shows the result of a full wafer scan.

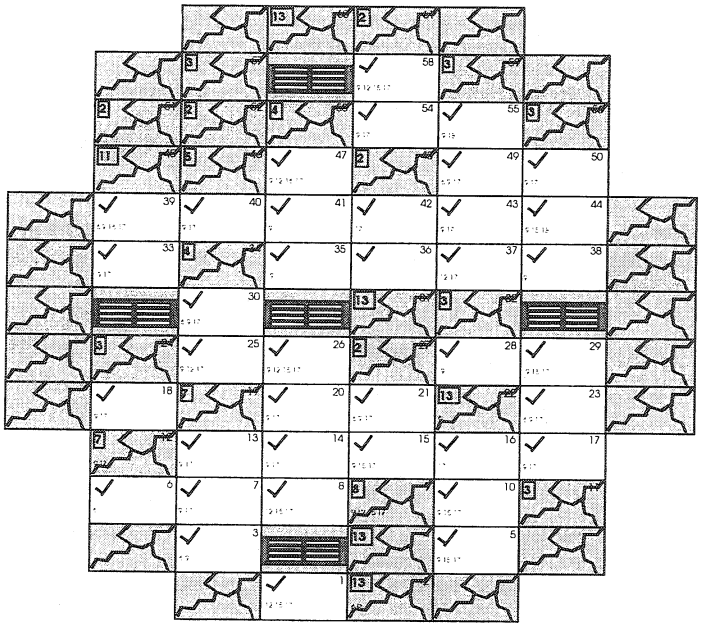

Figure 2.5 Map of wafer 14 from APV6 run D40208A. Note incomplete sites at edge of wafer are never probed, also the 5 sites which contain transistor test structures.

Each of the sites contains a summary of the information gleaned during testing. The failed sites contain the number of the fatal tests that they failed. The good sites contain a list of the non-fatal test that they failed. The availability of this information is helpful in speeding up the process of selecting chips for distribution and also for passing on the information to the recipient of the chip.

A similar map for testing of the APV25 is under development and will be included in the testing software in the next few months.

2.4.3 Availability of data on the web

Presently all the data is available for access through a web site at the url: http://albert.hep.ph.ic.ac.uk/silicon/projects/apv/apvdat.html. The data for each chip are stored as a single spreadsheet file, which is available for download. The web site contains descriptions of the tests and special information relating to each type of chip tested. All APV6 and APVM data is now available and data from the APV25 testing will be posted soon.

The plan is to develop the site into an on-line database of all the test data from all the chips. Work is in progress at CERN to develop the database software. Eventually this database will contain all the information that exists about every chip in the CMS tracker.

2.5 Yields

The yield of a wafer or run of chips is the percentage of working chips. Ideally chips would be tested and then either marked as working or non-working. However, the yield of the APV6 and APVM was very bad on some wafers. A lot of chips were not perfect, but still partially functional and useful. Since there was a shortage of working chips it was decided that rather than simply labelling chips as good or bad, they would be divided into four categories. The categories are shown in Table 2.3

Category	Description	Tests Passed
A	Perfect	All
B	Minor Faults	All Except 9 12 15 16
C	More Significant Faults	All Except 6 9 10 12 14 15 16 17 18
D	Only useful for testing	None

Table 2.3 Categorisation of tested APV6 and APVM chips.

The yield for the APV25 was considerably better and therefore categorisation was unnecessary.

2.5.1 APV6 Yields

Table 2.4 shows a breakdown of the yields for the APV6 wafers. From the D40208A run 16% of the chips are class A and from the DALND run only 6%. Chips in categories A, B and C are considered usable in some way and so contribute to the overall yield. However, only class A chips would be used in the final system and so these numbers are more interesting from the point of view of CMS. The uniformity across wafers is not very good. Some wafers have a class A yield as high as 50% and overall yield as high as 67% where the most common Class A yield is ~ 0% and overall as low as 16%. In two wafers the overall yield is virtually 0%. The average class A yield of 10% is extremely low and would increase the predicted manufacturing cost by a factor of 5. This would be of great concern if the final version of the APV were to be fabricated by the same manufacturer. However, one should note that these wafers came from a series of fabrication runs during which the process was still being adjusted to a new foundry.

Chip ID	Batch	Date	Scribe	Cut?	Yield	A	B	C	D
APV6	D40208A	Nov-98	2	Yes	40	33	7	0	21
			3	Yes	32	2	0	30	29
			8	Yes	46	0	0	46	15
			9	Yes	28	1	27	0	33
			12	Yes	10	0	0	10	51
			14	Yes	17	3	14	0	44
			15	No	0	0	0	0	61
			18	Yes	37	32	5	0	24
			20	Yes	28	19	9	0	33
Sub-total out of 549					**238**	**90**	**62**	**86**	**311**
Percentage of total					**43%**	**16%**	**11%**	**16%**	**57%**
APV6	DALND	Dec-98	2	No	12	0	6	6	49
			3	Yes	37	1	22	14	24
			6	Yes	36	0	21	15	25
			7	Yes	41	2	22	17	20
			8	Yes	12	12	0	0	49
			9	Yes	40	0	24	16	21
			10	Yes	36	0	23	13	25
			12	Yes	42	0	20	22	19
			13	No	3	0	0	3	58
			14	Yes	16	2	12	2	45
			15	No	41	0	5	36	20
			16	Yes	26	17	8	1	35
			17	Yes	30	0	14	16	31
			18	No	31	8	9	14	30
			19	Yes	12	11	1	0	49
			20	Yes	36	7	4	25	25
Sub-total out of 976					**451**	**60**	**191**	**200**	**525**
Percentage of total					**46%**	**6%**	**20%**	**20%**	**54%**
Total for APV6 out of 1525					**689**	**150**	**253**	**286**	**836**
Percentage					**45%**	**10%**	**17%**	**19%**	**54%**

Table 2.4 Breakdown of APV6 wafers.

A large number of the chips had problems with the calibrate pulse and some had no pulse at all. The lack of a calibrate pulse does not mean that these chips are completely useless, since they may still have working amplifiers and be perfectly normal in all other respects. Because of the nature of the test setup it was not possible to inject a real input pulse, therefore the only way of testing the amplifier and shaper was with the calibrate. Also a large number of the chips required very high settings for VADJ in order to produce a reasonable output. Many chips were failed because a reasonable level could not be obtained. Under these circumstances it is clear that selection of the

chips becomes far more difficult and not as clear-cut as would be required in order to be confident in an automatic selection procedure. The cuts that would be made automatically would have to rely upon some consistency between chips and wafers.

2.5.2 APVM Yields

The same problem occurred with the APVM wafers. Again the yield of class A chips is very low at an average of 15%. The same problem with inconsistency across wafers is seen, some wafers have a class A yield of ~40% and a most common class A yield of ~0%.

Chip ID	Batch	Date	Scribe	Cut?	Yield	A	B	C	D
APVM	DAP4P	Feb-99	3	Yes	43	15	10	18	18
			4	Yes	38	0	0	38	23
			7	Yes	46	24	10	12	15
			8	Yes	49	9	24	16	12
			10	Yes	38	4	20	14	23
			13	Yes	50	1	36	13	11
			16	Yes	38	0	0	38	23
			17	Yes	44	0	0	44	17
			18	Yes	27	0	0	27	34
Sub-total out of 549					**373**	**53**	**100**	**220**	**176**
Percentage of DAP4P total					**68%**	**10%**	**18%**	**40%**	**32%**
APVM	DAP9D	Jul-99	5	Yes	29	9	6	14	32
			6	Yes	39	18	8	13	22
			11	Yes	43	15	10	18	18
			12	Yes	37	9	14	14	24
			14	No	40	0	3	37	21
			16	Yes	53	23	13	17	8
			17	Yes	48	22	13	13	13
Sub-total out of 427					**289**	**96**	**67**	**126**	**138**
Percentage of DAP9D total					**68%**	**22%**	**16%**	**30%**	**32%**
Total for APVM out of 976					**662**	**149**	**167**	**346**	**314**
Percentage of APVM total					**68%**	**15%**	**17%**	**35%**	**32%**

Table 2.5 Breakdown of APVM wafers.

2.5.3 APV25 Yields

Since all the tested APV25 chips were already sawn from the (multi project/shared) wafer and the uniformity and yield are much better, there has been no separation

between wafers. The improvements on the APV6 and APVM yield and uniformity are staggering. The yield is so high that it is unnecessary to categorize the chips further than just as 'passed' or 'failed'. 419 out of 501 chips passed all tests, giving a yield of 84%, with no significant difference between wafers. Table 2.6 gives a breakdown of the APV25 failures. Since the yield was so high the cuts were very tight. No chips would be passed with any minor errors. All of the passed chips match the criteria for class A chips from the APV6 and APVM testing. Therefore only those chips with no bad channels are passed.

Test Description	**# of fails**
Digital functionality	33 (1)
Power supply rails out of range	6
Pipeline defects (>0 bad cells)	12
Channel defects (pedestals or pulse heights)	26 (5)
Physically damaged during probing	2
Rejected due to visible cutting damage	3
Total number of failures	82
Yield	**84%**

Table 2.6 APV25s0 failures summary (numbers in brackets indicate no. of chips in that category which showed cutting damage).

It is clear from these results that the APV25 has been fabricated in a more reliable and consistent technology. With such a high yield large savings would be made on the manufacture from CMS, and even tighter cuts could be considered to ensure confidence in the chips going into the final system. Obviously these results are, at this stage, indicative only of the production yield. Nevertheless they are very encouraging. First measurements on full wafers are presently underway.

2.6 Analysis of Test Data

2.6.1 Currents

Positive and negative rail current values from every tested chip have been histogrammed in Figure 2.6 to Figure 2.9 to show distributions. Note the high frequency of chips with zero currents, which is due to chips failing before the procedure reached the current test. Chips in this category are assigned a value of zero for both current rails.

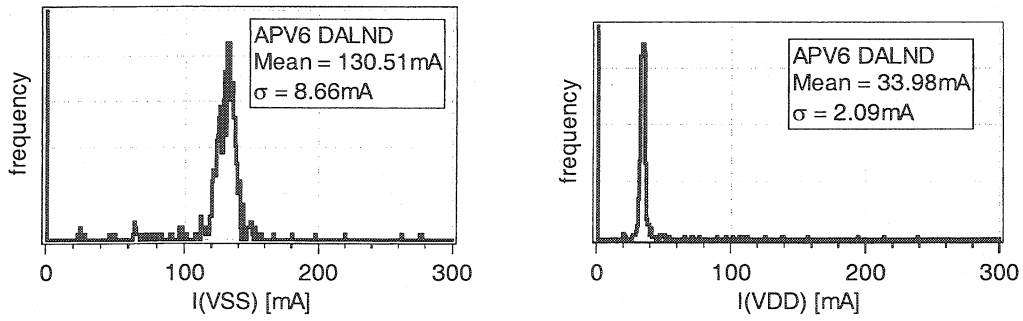

Figure 2.6 Current distributions for APV6.

The distribution is wider for the negative rail than the positive rail. By inspecting the distributions one can set sensible limits for the current test. However, it is prudent to consider these distributions with one eye on the overall functionality of the chip. Later the question of any correlation between particular test failures and currents is discussed.

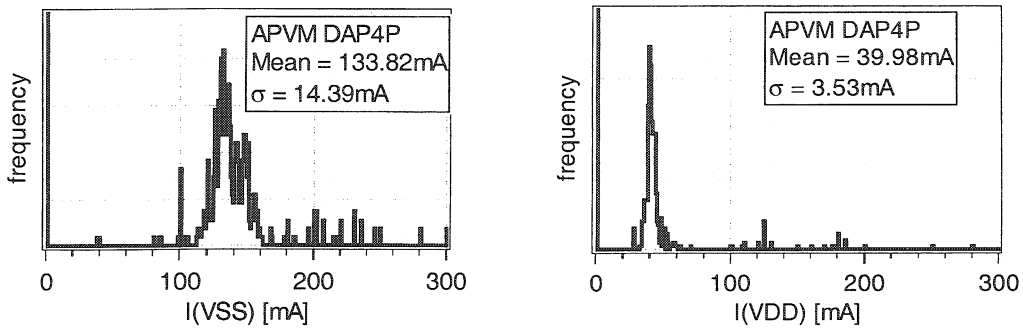

Figure 2.7 Current distributions for APVM DAP4P.

The APVM DAP4P distributions are much the same as the APV6. The average currents are slightly larger, which is due to the increase in the power supply voltages from ± 2V to ± 2.25.

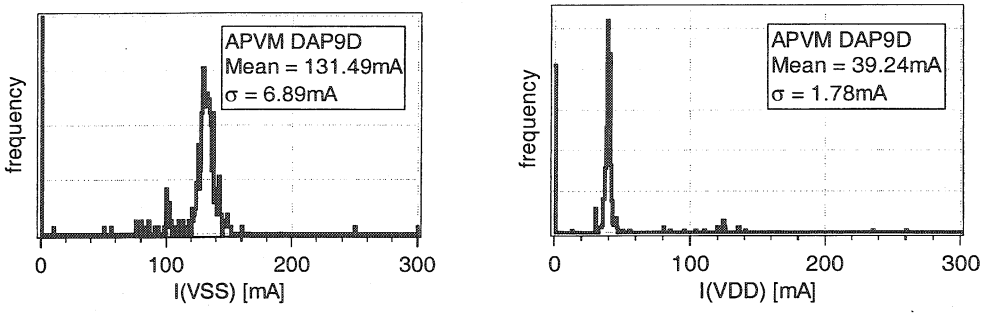

Figure 2.8 Current distributions for APVM DAP9D.

Some slight differences can be seen between the two APVM runs, DAP4P has slightly wider distribution with a fair number of chips showing higher than normal

currents, whereas DAP9D has a number of chips with lower than normal currents. This again emphasizes the inconsistency between runs.

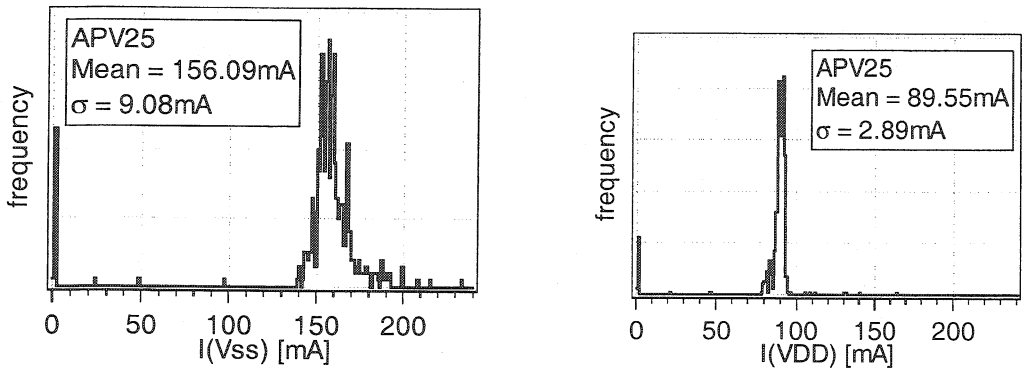

Figure 2.9 Current distributions for APV25.

The width of the APV25 distributions is comparable to those from the previous distributions. However, the frequency of chips with abnormal currents is lower, noticeably the reduction in the number of chips showing zero currents, which indicates the higher yield.

Chip	Positive rail		Negative rail	
	Mean [mA]	σ [mA]	Mean [mA]	σ [mA]
APV6	34.0	2.1	130.5	8.7
APVM Dap4p	40.0	3.5	133.8	14.4
APVM Dap9d	39.2	1.8	131.5	6.9
APV25	89.6	2.9	156.1	8.9

Table 2.7 Summary of current mean and standard deviation.

2.6.2 Scatter Plots in Current Space

By plotting I(VDD) against I(VSS) one has a useful view of the full current distribution. The rectangles drawn onto the plots represent the cuts made by the current test.

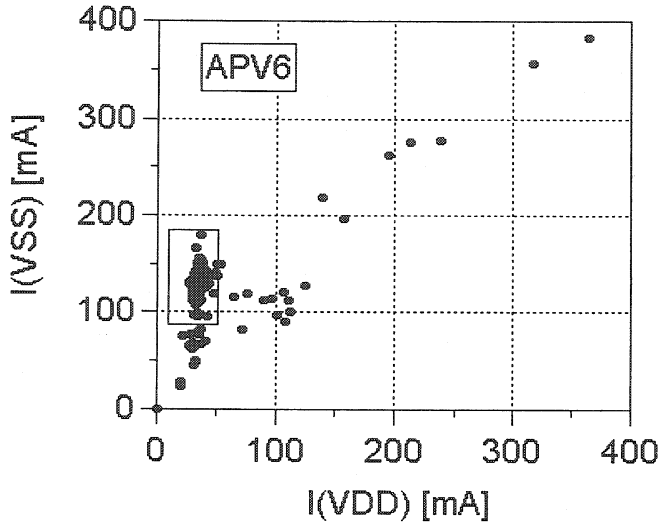

Figure 2.10 I(VDD) vs I(VSS) for APV6.

The lower limit on I(VDD) in Figure 2.10 could be increased with no reduction in the yield, while the remaining cuts appear to be well chosen.

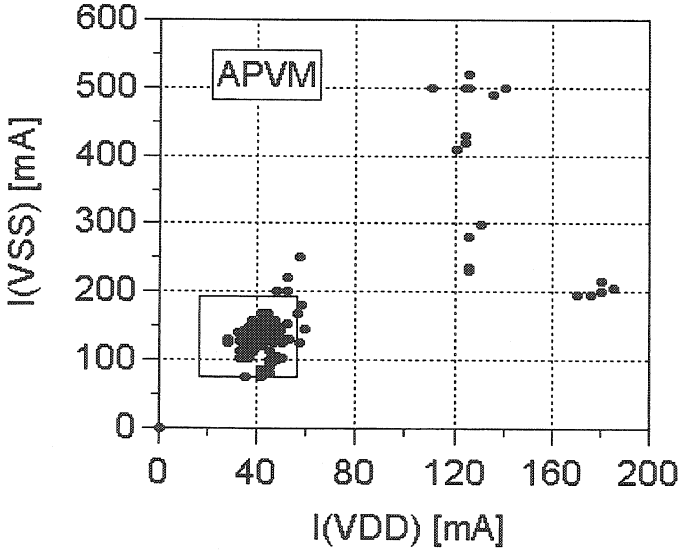

Figure 2.11 I(VDD) vs I(VSS) for APVM.

Both APVM and APV25 limits show room for contraction, with the distributions looking tight in comparison to the APV6.

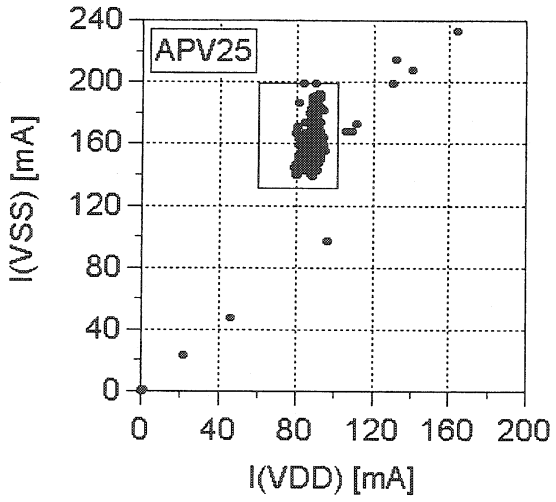

Figure 2.12 I(VDD) vs I(VSS) for APV25.

2.6.3 APV6 Pipeline Noise and its Correlation with Currents

During the testing the noise in each channel from throughout the pipeline is measured and the average over all 128x160 is calculated. Figure 2.13 shows the average channel noise for the APV6. Test 13 was originally set with an upper limit on this noise of 100, which is clearly very generous. This cut can now be reduced to about 7 with very little affect on the yield. The standard deviation is small, which demonstrates good uniformity between wafers and runs. However, this does not necessarily mean that the functionality of the chip will also be uniform, to investigate this one must look at the gain distributions shown in section 2.6.4.

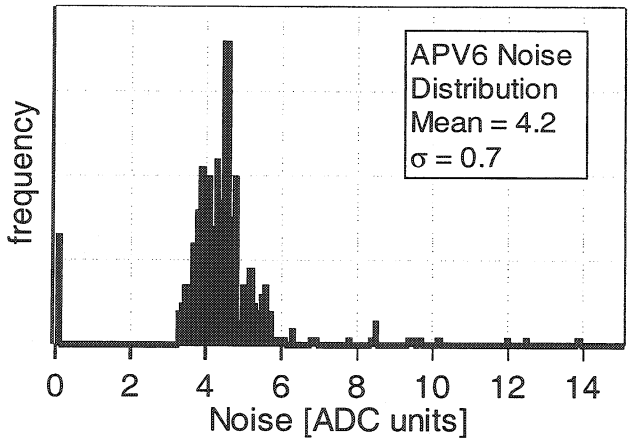

Figure 2.13 Distribution of average channel noise from all APV6 chips.

It is interesting to see if there is any correlation between different features of the chips. By plotting the noise distribution against the current distribution, as in Figure 2.14, we find no correlation between the noise and the current, which indicates that the

variation in the current is mainly due to circuits in the chip that are unrelated to the performance of the amplifiers.

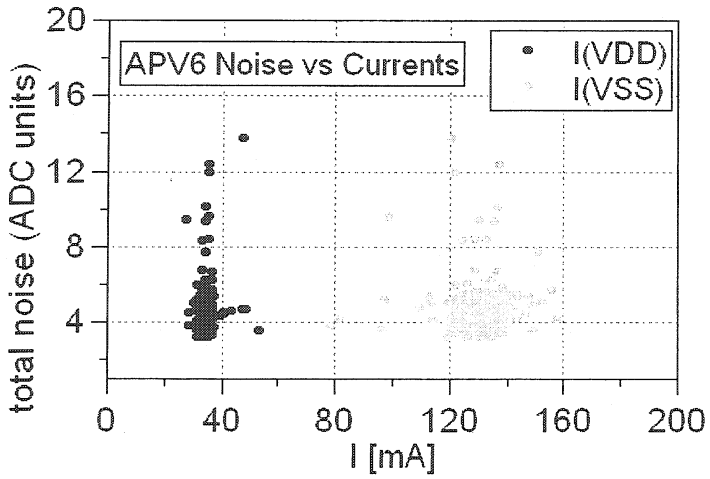

Figure 2.14 Noise vs Current for APV6.

2.6.4 Gain and its Correlation with Currents

One thing to look for is a consistency in the gain across wafers and runs. Collated data from all tested chips have been histogrammed. The high frequency of low gain values in both the APV6 and APVM (Figure 2.15 and Figure 2.16) are indicative of the large number of chips that failed test 16. Failure of this test could be cause by either true low channel gain, or faulty calibrate circuits. Therefore this picture could be misleading. Of the chips that did display working calibrate pulses, the gain distribution is fairly wide, also clear evidence for large differences between runs is manifested as multiple peaks in the distribution, especially for the APVM. The lower second peak for the APV6 is due to the large number of failures in the second run (DALND), which exhibited serious problems with the calibrate circuitry.

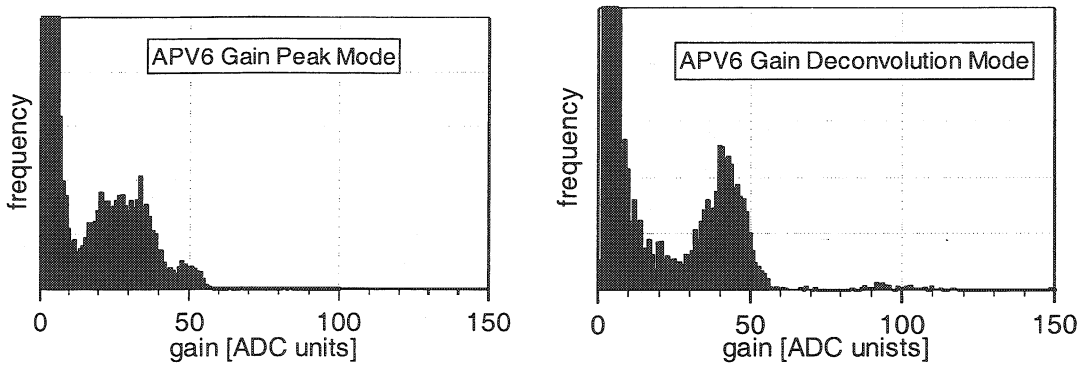

Figure 2.15 APV6 gain distribution in peak and deconvolution mode.

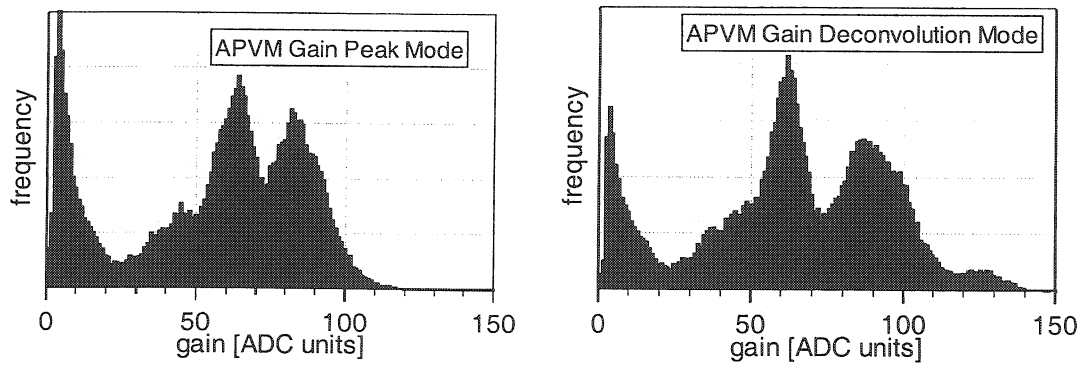

Figure 2.16 APVM gain distribution in peak and deconvolution mode.

In contrast the APV25 distribution is more encouraging with a tight distribution indicative of very good uniformity across wafers.

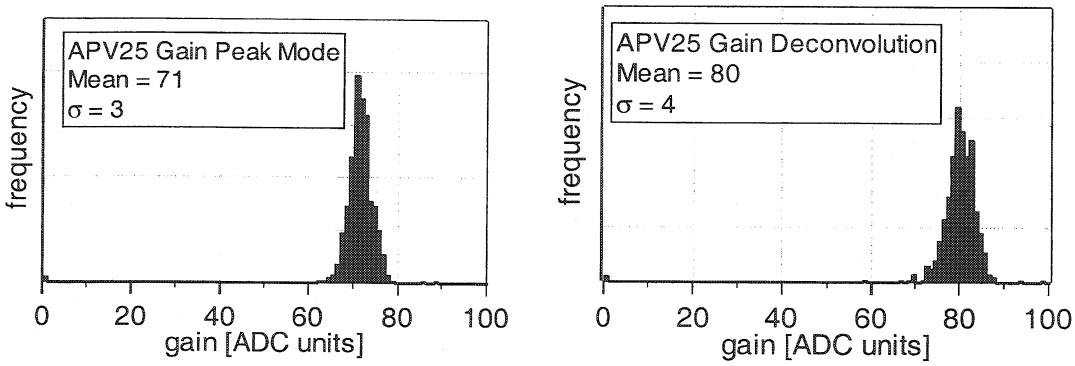

Figure 2.17 APV25 gain distribution in peak and deconvolution mode.

Again one can look for correlation, this time between the gain and the current distributions. Figure 2.18 shows gain against currents for the APVM and APV25, and no correlation is observed.

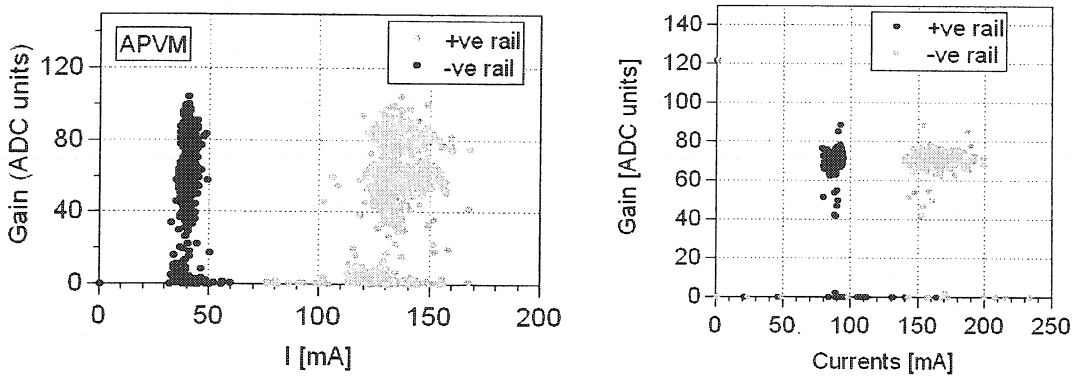

Figure 2.18 Gain vs Current for APVM and APV25.

2.6.5 Correlation Between Current and Test Failures in APV6/M

One purpose of this analysis is to look for ways to reduce the testing time. If correlation between specific test failures and chip currents was present it may be possible to reduce the number of necessary tests. For example, if chips that failed a test also drew abnormal currents then perhaps tighter cuts on the currents could eliminate these chips and also remove the necessity for that specific test. To investigate whether such correlations exist one can plot the current distributions of chips that passed a specific test against those which failed. Since the yield for the APV25 was so high very few failures are available to compile such plots, therefore only the data from the APV6/M testing has been used.

A thorough investigation of all the errors and correlation between them, the currents, and inter error correlation has been performed. However it is only necessary to include some of these results, since all show similar behavior. Figure 2.19 shows the total error distribution; since the best statistics have been gathered for the errors with higher frequencies we will investigate errors 9, 10, 12 and 15, all of which show high frequencies for either the APV6 or APVM.

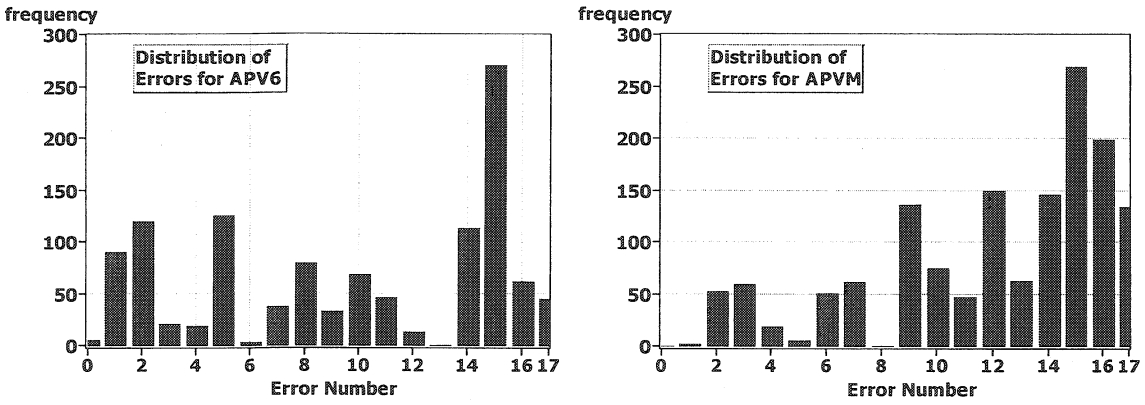

Figure 2.19 Distribution of Errors for APV6 and APVM.

2.6.5.1 I^2C Address Test (Errors 9 and 10)

In Figure 2.20 and Figure 2.21 the negative rail current distributions for chips that passed and those that failed are plotted separately. At first glance the current distributions for both passed and failed chips does not appear to be affected by the separation between passes and failures indicating a lack of correlation. A reasonable number of chips that failed the test display normal currents and a number of chips that have current values around the edge of the distribution passed the test, which again

shows a lack of correlation. However, if one plots the percentage of chips that fail out of all the tested chips, then a clear dip is present in the center of the current distribution. Therefore it is less likely that a chip will fail the test if it is drawing a normal current. However, since the percentage of failures does not drop significantly until near the center of the distribution, where it is ~ 10% for both APV6 and APVM, tighter cuts on the current would not be beneficial and the address test would still be necessary to remove the 10%. The same correlation is present in the positive rail distributions and the same conclusions apply.

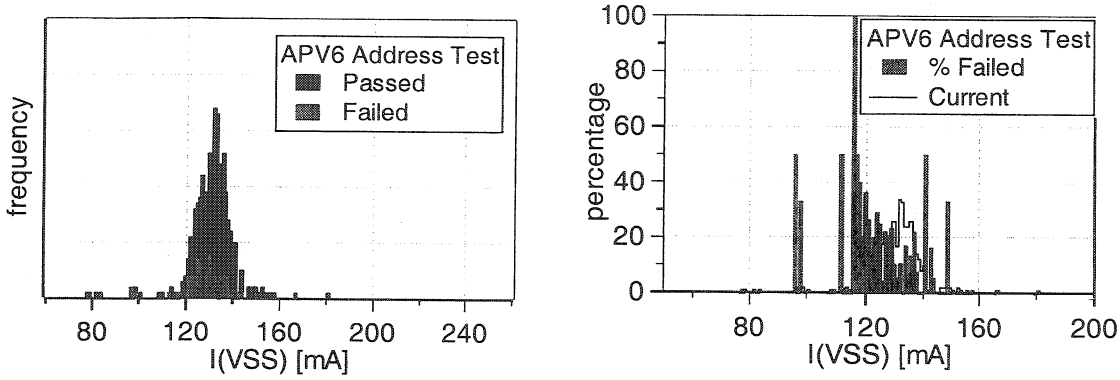

Figure 2.20 Current and failure percentage distributions of APV6 chips that passed and failed I^2C address test.

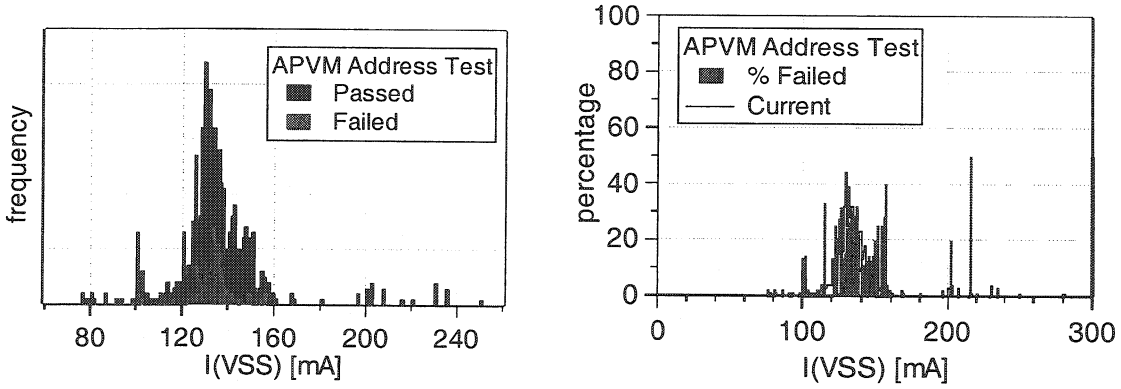

Figure 2.21 Current and failure percentage distributions of APVM chips that passed and failed I^2C address test.

2.6.5.2 Pipeline Cell Test

Figure 2.22 examines the current distributions of the chips that passed and failed the pipeline cell test. The distributions are similar to those from the address. One thing to note is the high percentage of failures outside the central peak. This again displays the expected correlation, but again the failure percentage does not drop far enough in the central region to warrant tightening of the current test.

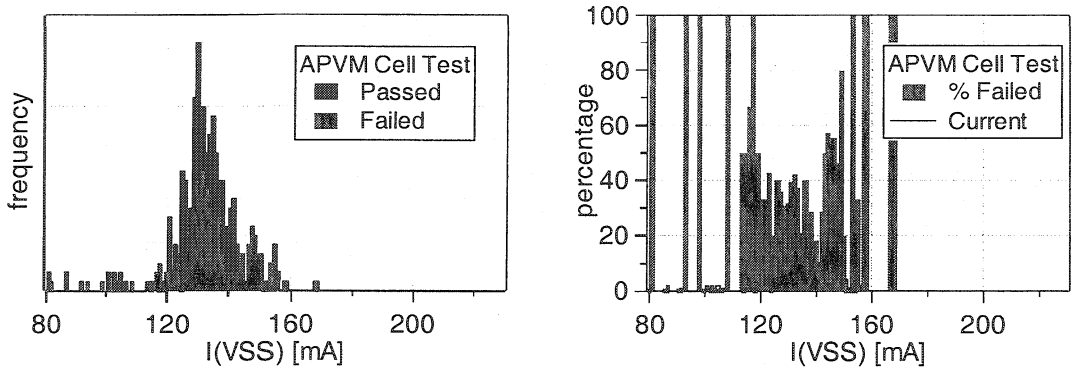

Figure 2.22 Current and failure distributions of APVM chips that passed and failed pipeline cell test.

2.6.5.3 Pedestal test

Figure 2.23 shows the same distributions for the pedestal test, which displays similar results yet again. Nearly all chips that had abnormally high currents failed this test, but also some chips with normal currents.

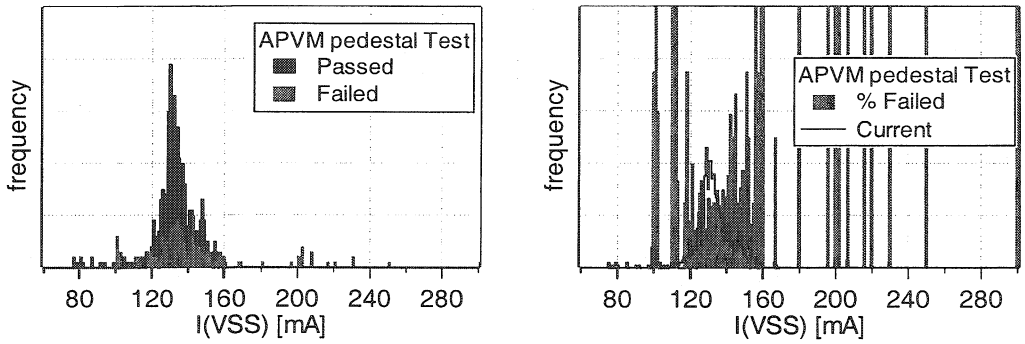

Figure 2.23 Current and failure distributions of APV6 chips that passed and failed pedestal test.

2.6.6 Uniqueness of APV25 Pedestal Characteristics

In this last section, we will investigate the degree of uniqueness possessed by the channel pedestals of each chip. It has been proposed that the pedestal values of a chip could be used as a fingerprint of the chip. By comparing the pedestals of a random sample chip with those of all other chips and calculating a measure of the similarity, it is possible to establish the degree of uniqueness. We will only look at the APV25 since it is now clear that this is the chip that will be used in the CMS tracker.

There are various ways of comparing chip pedestals, one of which is to simply compare the individual channels directly, setting some threshold for which the difference between the two channels must be less. A chip is considered to be a match if

the difference in all the channels is less than the threshold. Initially this method was tested giving rise to the conclusion that no other chip resembled the pedestals of the sample chip if the threshold was set at 4 ADC units; this was true for both peak and deconvolution modes. However, this justification holds only if we look at all 128 channels; comparing 95% of the channels, with the same threshold, began to yield matching chips. By reducing the threshold to 2 ADC units, the uniqueness of the original chip was regained. However, a threshold of 2 ADC units is very small and the chance of not being able to recognise the original chip is increased, since the error between measurements of the pedestals can be larger than this.

A more efficient way of looking at the problem is to construct a measure of the overall deviation across all channels. A simple way to do this is to sum the square of the differences between each channel, which gives an abstract measure of the difference δ.

$$\delta = \sqrt{\sum_{i=1}^{128} (E_i - x_i)^2} \qquad [2.1]$$

Where i is the channel number; E_i is the sample chip pedestals, x_i is the compared chip pedestals. Four samples were taken from four different wafers, δ was calculated for each comparison between the sample chip and the remaining chips. The resulting distributions are shown in Figure 2.24 and Figure 2.25.

The single entry at 0 in each distribution is the comparison of the sample chip itself. The distributions show a clear gap between 0 and about 130, leaving plenty of room to set a threshold.

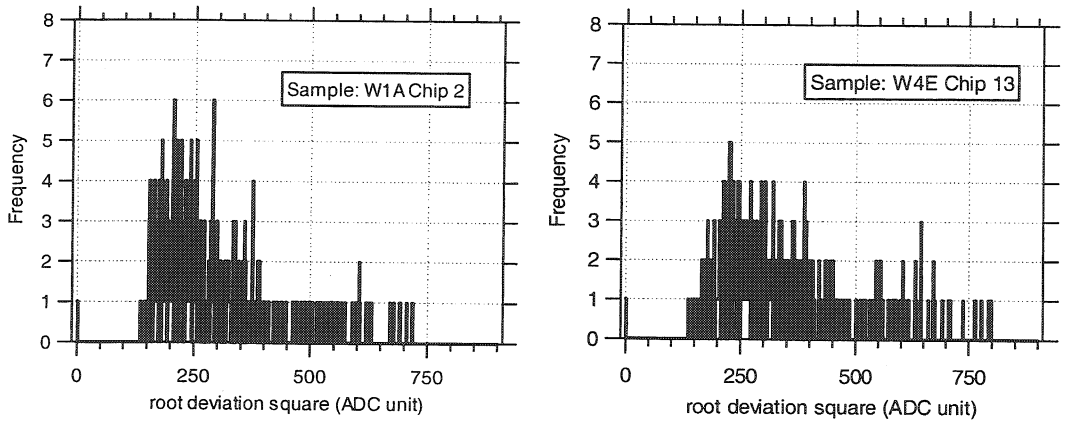

Figure 2.24 Distribution of δ for two sample chips.

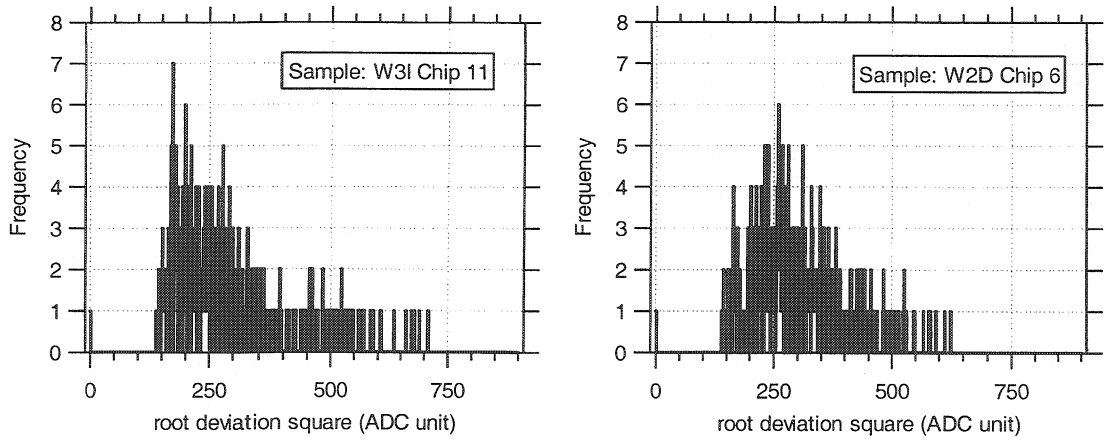

Figure 2.25 Distribution of δ for two sample chips.

From this study it appears that it should be possible to use the pedestals of a chip as a good identifier.

2.7 Conclusions

To date the wafer screening setup is capable of automatically testing all chips on an entire wafer, making an exhaustive study of the digital functionality and comprehensive characterisation of the analogue components. All data retrieved during the test procedures is then stored and a map of the wafer, indicating the good and bad sites and the degree of their functionality, is produced.

The comparison of the yields from the APV6 and the APV25 shows a considerable difference in the uniformity and reliability between the two processes; the APV6 showing large variations between runs and on the wafer, but the APV25 showing very good uniformity and tight distributions of currents and other statistics. The pipeline uniformity in the APV25 is greatly improved and displays average variations throughout of no more than ±1 ADC count.

It is important that careful consideration of the test threshold settings is made in order to ensure the optimisation of chip performance and minimisation of production costs. From the results shown in this chapter there is now solid information on which to base the failure criteria and acceptable limits in which the chips should fall; some correlation between currents and failure modes has also been discovered, but it would appear that this correlation is not strong enough to warrant tightening of the current limits and removal of certain tests, in order to bring down the testing time.

Work has been done on developing a comprehensive data storage and recovery system, which will be required later. All data are presently saved in binary format and later converted into a text version for distribution on the Web. The development of a database for the test data is an important task and will commence relatively soon. The exact format of this database and the media, which it will employ, is under discussion. It has also been shown that the pedestal uniqueness of individual chips could be used as a useful chip identifier, which could also be included in the final database.

Chapter 3

Radiation Effects on Microelectronics and the APV

Semiconductor devices have been consistently decreasing in size since the invention of the first transistor. The special properties of semiconductors have enabled production of extremely sophisticated electronic systems and to design devices capable of capturing the ionized charge from interactions between particles and the semiconductor. Unfortunately, transistors are also susceptible to unwanted effects produced by ionizing radiation. Moreover, the regular crystal lattice and specific properties, which make semiconductors unique, also make them vulnerable to the physical effects of radiation damage. Interactions between radiation and the transistor can introduce defects and trapped charge, giving rise to subtle changes in device characteristics. However, these "total dose" radiation damage mechanisms only occur over long periods of exposure. An understanding of the mechanisms behind the effects enables design of transistors in a way that can reduce the impact of the damage on circuit operation. Nearly all commercial transistors produced today are fabricated on silicon substrates and so a large proportion of the interactions occur within the silicon lattice and surface oxide layer. Understanding how particles interact with silicon is essential if we are to predict how they will affect transistor operation.

In the mid to late eighties the feature sizes of CMOS processes had shrunk to sizes of order 10 μm, and the charge involved in the operation of certain digital devices was becoming so small that it was possible for individual charged ions to deposit comparable amounts of charge. Under such conditions individual interactions can produce single event effects; these come in different forms, some transient and non-fatal

and others devastating and permanent. This chapter will investigate the different forms of radiation effects on silicon electronic devices. Today feature sizes of the most advanced processes are as small as a fraction of a micron, with surface oxide thickness even lower than 5.5 nm. Such small sizes actually reduce the impact of some radiation effects, but on the other hand they may give rise to new and, as yet, poorly understood effects. Some of these considerations are discussed in this chapter, but the main focus of attention is Single Event Upset and its effect on the operation of the APV chip.

3.1 Ionizing Radiation and Energy Deposition in Silicon

When a charged particle passes through matter it loses energy by means of three principal processes, ionization (excitation), radiative effects (Bremsstrahlung), and Rutherford scattering[3.1]. High energy hadrons can also lose energy by means of inelastic nuclear interactions.

3.1.1 Ionization

Mediated by the Coulomb interaction, ionization is the stripping of orbital electrons from the atoms of the material. Excitation is when an electron is given enough energy to ascend to a higher energy level after which it, or another electron, will emit a photon and return to the lower state. This is the principal means by which a fast moving charged particle loses kinetic energy when moving through a material.

3.1.2 Bremsstrahlung

When a charged particle undergoes acceleration, it transfers energy to the surrounding medium by emitting a photon. For electrons in matter, this process is important since the relatively low electron mass makes it easy for them to be accelerated. Within silicon, this process is the dominating mechanism by which high-energy electrons (K.E. > 40 MeV) lose energy to the surroundings. For low energy electrons, Coulomb interactions dominate.

3.1.3 Rutherford Scattering

Rutherford scattering is the Coulomb interaction of a charged particle with the electric field of the nucleus. For high energy incident particles, sometimes enough

energy will be transferred to displace the atom from its location within the material, which can then go on to heavily ionise within the local vicinity. For particles with kinetic energy greater than 1MeV, the energy loss is small compared to the ionization energy loss.

3.1.4 Inelastic nuclear interactions

Processes in which hadrons, such as protons, neutrons and pions, interact within the nuclei of atoms in the material are known as inelastic nuclear interactions, or simply hadronic interactions. Such interactions will typically produce a shower of secondary particles and one nuclear recoil – in some cases there can be several fragments. The secondary particles and fragments will interact futher within the material. In the context of this thesis the heavier fragments and recoil particles are important since they can go on to produce localized ionization, which can cause unwanted effects.

3.1.5 Stopping powers

Later in this chapter we will begin considering specific energy depositions within sections of a device. In order to calculate the expected energy deposition by incident ions we can simply use the Bethe Bloch equation[3.2] as shown in equation [3.1].

$$-\frac{dE}{dx} = K\rho z^2 \frac{Z}{A}\frac{1}{\beta^2}\left(\frac{1}{2}\ln\left(\frac{2m_e c^2\gamma^2\beta^2 T_{\max}}{I^2}\right) - \beta^2 - \frac{\delta}{2}\right) \qquad [3.1]$$

E = projectile energy, β = projectile velocity(in units of c), $\gamma=(1-\beta^2)^{1/2}$, T_{max} = maximum kinetic energy which can be imparted to a free electron in a single collision, z = projectile charge, x = path length, $K = 4\pi r_e^2 m_e c^2 N_A$, r_e = classical electron radius, m_e=electron rest mass, N_A = Avogadro's number, Z = atomic number of medium, A = atomic weight of medium, ρ = mass density of medium, I = average ionization potential, δ= density correction.

$-\frac{dE}{dx}$ is the energy loss with distance and is known as the stopping power in units of eV/cm; to a good approximation it is proportional to the electron density in the medium (given by $\rho Z N_A/A$) and to the square of the projectile charge, and otherwise depends mainly on the projectile velocity. It decreases with $1/\beta^2$ for increasing velocity until reaching a minimum around $\beta\gamma$ = 3 to 4 (*minimum ionization*), then starts to rise

logarithmically (*relativistic rise*) levelling off finally at a constant value (the *Fermi plateau*). The numerical value of the minimum ionization (more precisely: of minimum energy loss) is $dE/\rho dx \sim 2$ MeVcm^2/g. In this form, the Bethe-Bloch equation describes the energy loss of pions in a material such as copper to about 1% accuracy for energies between about 6 MeV and 6 GeV. At lower energies the term –C/Z must be added for tightly bound atomic electrons and other effects, and at higher energies radiative effects begin to be important. These limits of validity depend on both the effective atomic number of the absorber and the mass of the slowing particle.

An example of some measurements can be seen in Figure 3.1, obtained in a time projection chamber, taken from reference[3.3]. The measured particles have a distribution of stopping powers with a mean represented by the Bethe Bloch equation.

Figure 3.1 Example measurements of the stopping powers of the proton, K-meson, pion electron and muon.

At lower projectile energies, the stopping power rapidly increases until it reaches a peak, known as the Bragg peak, after which it decreases to zero. Figure 3.2 shows the Bragg peak for a silicon ion calculated using the SRIM code[3.4].

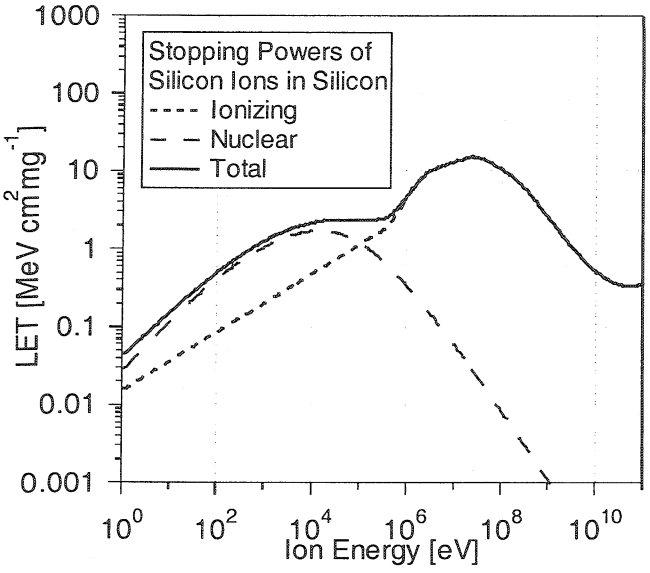

Figure 3.2 Energy loss of silicon in silicon as a function of the kinetic energy.

It is also common to measure stopping powers in a slightly different form under the name Linear Energy Transfer (LET), is a measure of a particle's energy transfer per unit of path length in a particular medium and is given by, $LET = \frac{dE}{dx}.\frac{1}{\rho}$. All reference to particle stopping powers in the remainder of this thesis is done so in terms of LET.

3.2 Total dose radiation damage

Silicon can be affected by radiation in two ways: surface effects and damage to the bulk. CMOS electronics operate purely through effects at the surface and junctions and operation of these devices can be affected by build up of trapped charge, created mainly by ionizing radiation. On the other hand, detector devices such as silicon microstrips, operate by collecting charge created right through the bulk. In this case, changes in the bulk, such as variations in doping, will lead to changes in charge collection.

3.2.1 Bulk Damage

Bulk damage has a negligible effect on CMOS devices, so I discuss it very briefly. Physical damage to the bulk occurs when incident particles transfer enough energy to a lattice site to displace it from its original site. Such atoms can then go on to cause secondary displacements and produce a cluster of irregularity in the lattice. These

clusters are stable complexes of silicon[3.5] which possess slightly different characteristics than the bulk lattice.

The properties of a semiconductor depend crucially on the symmetry of the lattice. Asymmetries can lead to new energy levels within the band gap.

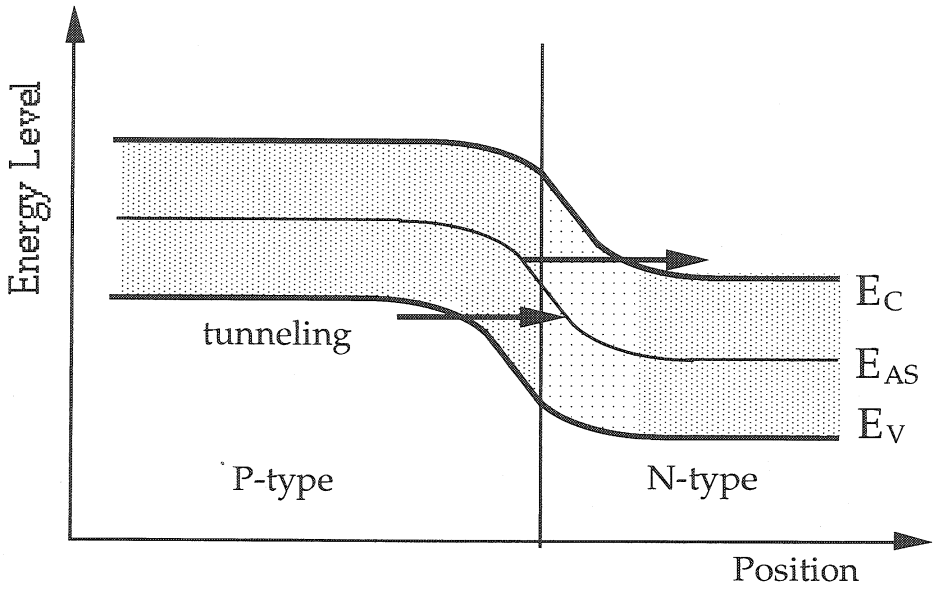

Figure 3.3 An effect of energy levels in the band gap of a p-n junction.

Figure 3.3 gives an example of how a new energy level, E_{AS}, in the band gap makes it easier for an electron to tunnel through to the conduction band. In some cases the required energy for such an event can be reduced by as much as half depending on the position of the energy level within the band gap. Other effects caused by intermediate energy levels are increased leakage currents due to generation and recombination effects[3.6], and an increase in noise[3.7].

3.2.2 Ionizing radiation effects

Surface changes are primarily caused by ionizing radiation and internal photoemission creating electron hole pairs in the surface oxide, as shown in Figure 3.4[3.6]. The three main effects are: a build up of charge in the oxide caused by trapped holes, which contributes to the field in the oxide and causes a threshold voltage shift; an increase in the number of interface traps, which can capture carriers from the channel and thus alter the threshold voltage and increase noise; an increase in the number of bulk oxide traps, which can drift with the electric field and, in the case of NMOS transistors, can form new interface traps.

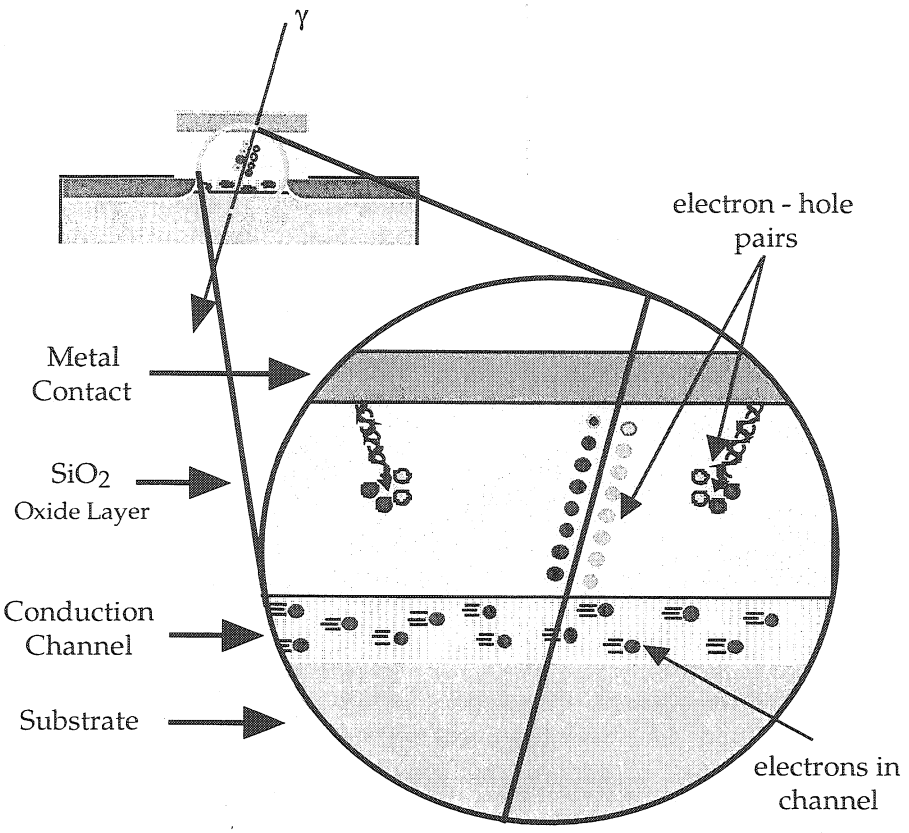

Figure 3.4 Electron hole pair production.

MOS transistors have significantly different electron and hole mobilities in the oxide;

$$\mu_e \sim 20 cm^2/V.s \quad [3.2]$$

$$\mu_h \sim 10^{-6} - 10^{-4}\ cm^2/V.s \quad [3.3]$$

The holes are created because after ionization the high mobility electrons are swept away by the electric field (typically 10^7-10^8 V/m), more rapidly than the low mobility holes. In general after irradiation the net charge of the SiO_2 is typically positive. Due to the opposite polarities of PMOS and NMOS devices the charge collects at the bulk interface in NMOS and at the gate electrode in PMOS. Any extra charge within the gate oxide will contribute to the electric field across the channel causing a change in threshold voltage.

As the remaining holes drift under bias they may be trapped in the oxide resulting in a net positive charge in the oxide, causing a negative shift in the transistor threshold

voltage. In NMOS the holes are trapped predominantly near the Si/SiO_2 interface. However, holes trapped within 2-5 nm of the interface will be removed by a tunnelling process, which is also thought to be responsible for long term recovery of the device properties, which is also known as annealing. The electron trapping cross section is 3 to 6 orders of magnitude lower than that of holes, which combined with their higher mobility makes the probability of trapping an electron very much lower.

Holes created at the interface, or those that reach the interface in NMOS, may be collected, or form an interface trap, which may capture an electron from the channel. This process lasts for between hours and years as long as a potential is maintained. These traps give rise to a positive shift in the threshold voltage for NMOS, which is independent of the silicon surface potential. The charge traps in PMOS that are capable of trapping holes from the channel are responsible for a negative threshold shift. The interface states are more stable than trapped holes.

An interface trap is an electronic level situated at the Si/SiO_2 interface which may capture or emit electrons or holes (see Figure 3.5[3.6]). They occur normally because of lattice discontinuities at the interface, loose chemical bonds and impurities.

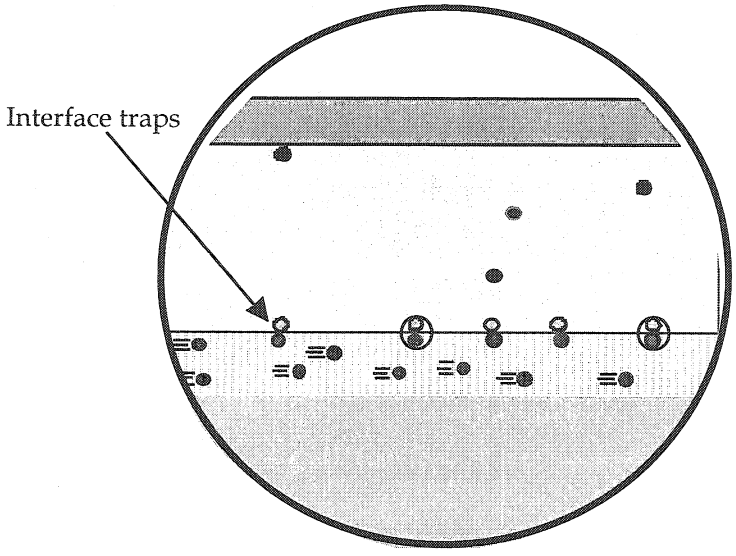

Figure 3.5 Trapped charge in the oxide and charge traps at the interface.

Only those traps with energy levels within the bandgap are of interest. The others are regarded as fixed charge within the oxide. There are two types of traps: donor-like traps which are neutral when filled with an electron or hole, and acceptor-like traps which are

negative when filled with an electron and positive when filled with a hole. These traps act to remove the carriers in the channel, for NMOS electrons are trapped and for PMOS holes.

The crystal structure of SiO_2 is a face centred cubic, with each silicon atom making 4 covalent bonds to oxygen atoms[3.7]. Chemical bonds in the SiO_2 may be broken by ionizing radiation passing through this crystal. Electron hole pairs can then be formed from the broken bonds, which can then recombine in the vicinity of the defect. The remaining broken bonds give rise to electrically active traps; either trap sites for carriers within the oxide or interface traps after migration. Some impurities in the lattice, which contain hydrogen or hydroxyl groups, will release these when their bonds are broken, these impurities become mobile defects within the SiO_2. In migrating towards the Si/SiO_2 interface they can be captured as interface traps[3.8].

The polarity of the interface trapped charge and the oxide trapped charge are the same for PMOS devices and opposite for NMOS devices. The oxide trapped charge always decreases with time, and the interface trapped charge usually increases with time. It is therefore possible in NMOS devices for annealing of the oxide charge to be so large as to cause a *rebound*, when the reduction in the oxide trapped charge is such that the trapped oxide charge is small compared to the trapped interface charge and the interface charge dominates. Figure 3.6 and Figure 3.7 show the changes in threshold voltage in NMOS and PMOS transistors respectively, due to oxide trapped charge and interface trapped charge during irradiation and after irradiation has ceased.

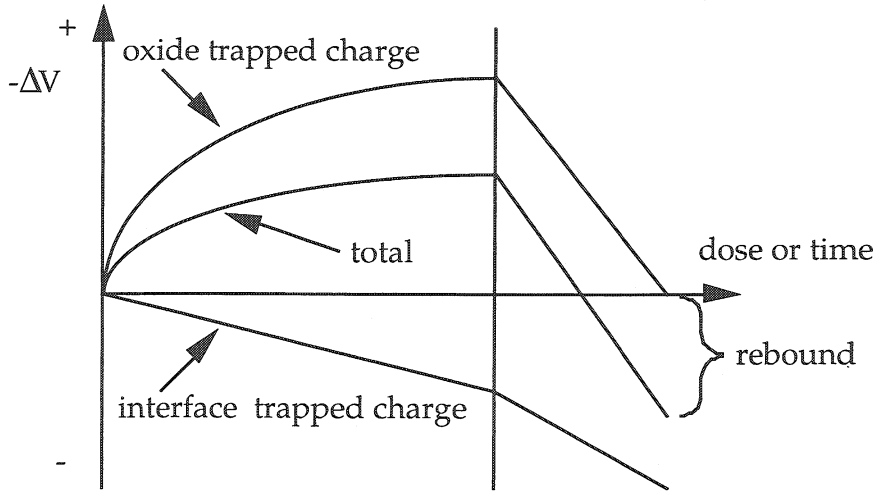

Figure 3.6. The theoretical threshold shifts in an NMOS transistor before and after irradiation.

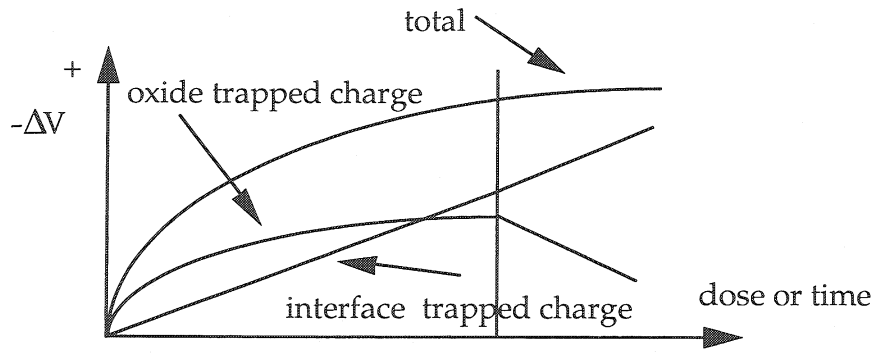

Figure 3.7 The theoretical threshold shifts in a PMOS transistor before and after irradiation.

The boundary between the two periods is indicated by the vertical line. In both PMOS and NMOS transistors, the interface trapped charge also reduces the charge carrier mobility by increasing coulomb scattering. The carrier density is reduced through surface recombination. This reduces the transconductance and the saturated drain current. Changes are observed in the sub-threshold transfer characteristics due to charging and discharging of the interface states.

3.2.3 Effects on noise

Ionizing radiation will result in changes in the noise characteristics of MOS transistors. Thermal noise increases because of the increase in interface traps which increases the width of the random thermal motion distribution, by trapping and releasing charge carriers in the channel; this effect is greater in NMOS since a larger number of interface states are created by irradiation. Experiment shows that 1/f noise is also expected to increase. The overall increase in noise spectral density for PMOS transistors is known to be less than in NMOS, for this reason the input transistor to the amplifier stage in the APV was chosen to be PMOS.

3.3 Single Event Effects

Radiation can also cause transient effects which can result in soft digital errors or even circuit damage[3.9]. There are many device conditions and failure modes due to single event effects (SEE), depending on the incident particle and the specific device. It may be convenient to think of two types of SEE: soft errors and hard errors. Soft errors are non-destructive and may appear as a bit flip in a memory cell or latch, or as transients occurring on the output of a device. Also included are conditions that cause a

device to interrupt normal operations and either perform incorrectly or halt. Hard errors may be (but are not necessarily) physically destructive to the device, but are permanent functional effects. Different device effects, hard or soft, may or may not be acceptable for a given design application. The following list describes the effects that will be discussed later.

- Single Event Latchup (SEL) - a condition which causes loss of device functionality due to a single event induced high current state. An SEL may or may not cause permanent device damage, but requires the power to be cycled to resume normal device operations.
- Single Event Gate Rupture (SEGR) - a single ion induced condition in MOSFETs which may result in the formation of a conducting path in the gate oxide because of the high electric field. Such an event will rupture the gate and permanently damage the device
- Single Event Upset (SEU) - a change of state or transient induced by an energetic particle such as a cosmic ray or proton in a device. This may occur in digital, analogue, and optical components or may have effects in surrounding interface circuitry (a subset known as Single Event Transients (SETs)). SEUs are "soft" errors in that a reset or rewriting of the device causes normal device behavior thereafter.
- Single Hard Error (SHE) - an SEU which causes a permanent change to the operation of a device. An example is a stuck bit in a memory device.

3.3.1 Single Event Latchup

Circuits are made in silicon by combining adjacent p-type and n-type regions. Paths other than those chosen to form the desired transistor can sometimes result in so-called parasitic transistors, which, under normal conditions, cannot be activated (see Figure 3.8). Latchup occurs when a spurious current spike, such as that produced by a heavy ion, activates one of a pair of these parasitic transistors, which combine into a circuit with large positive feedback. The result is that the circuit turns fully on and causes a short across the device until it burns up or the power is cycled.

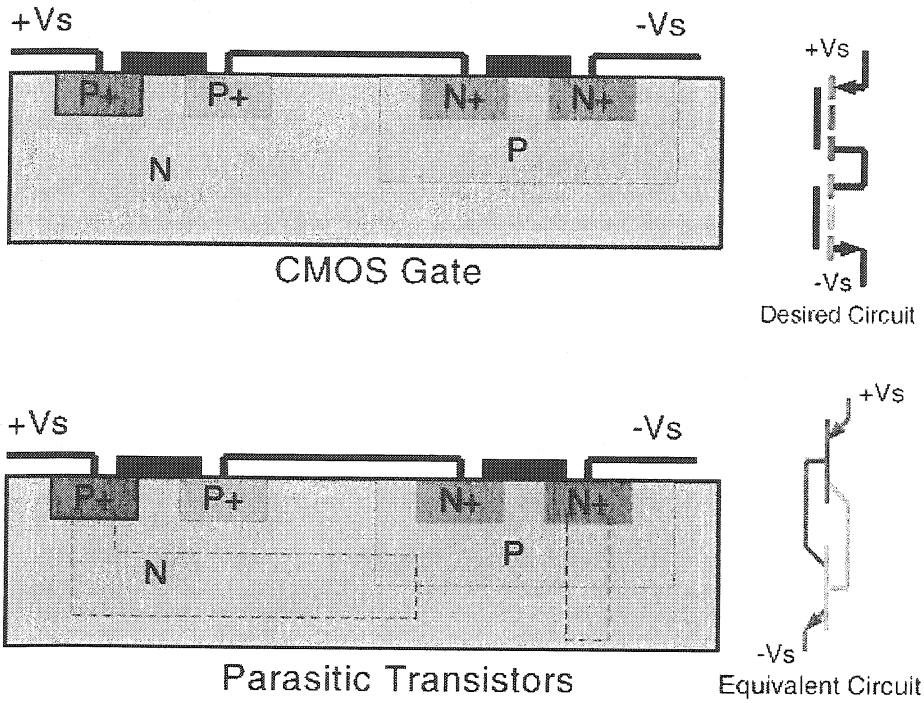

Figure 3.8 Two parasitic transistors in positive feedback loop that can give rise to SEL.

Much research has been done into this subject and it has been well understood for some time. The APV has been carefully designed to avoid SEL with the use of guard rings and design layouts[3.10][3.11][3.12][3.13].

3.3.2 Single Event Gate Rupture

SEGR has also been studied[3.14][3.15]. However, its effect in the APV has not, and work is under way at Imperial College to understand the mechanism more thoroughly. SEGRs are normally associated with power MOSFETs operating under high electrical fields. However, their importance has to be assessed for the CMS environment, therefore some preliminary measurements have been made in an attempt to evaluate the susceptibility of the APV to SEGR, but more detailed experiments on other subsystems will also be necessary. The results are presented in Chapter 4.

SEGR is caused when a highly energetic ionizing particle traverses the gate oxide and induces ionised charge along its track, which can lead to an electrical short through the oxide, causing permanent damage. Various models have been developed in an attempt to explain the SEGR mechanism. One of the models which fits experimental data well is based on the concept of a conducting cylinder surrounding the track of the ionizing particle.

The conducting pipe model is a simple semi-empirical model which points to a linear dependence of inverse electric field on LET:

$$E_{CR} = \frac{E_0}{1 + L/B} \qquad [3.4]$$

Where E_{CR} is the critical field to rupture, E_o is the breakdown field, L is the particle LET and B is a fitting parameter, which is technology dependent.

The increase in oxide electric field as the thickness decreases raises the concern that single event gate rupture (SEGR) could become a severe problem for advanced integrated circuits, such as those proposed for CMS. The dependence of SEGR on oxide thickness has been explored by Sexton et al.[3.14] for oxide thicknesses of 6 - 18 nm. They report that the critical field, defined as the field at which gate rupture is observed at a given LET, increases with decreasing oxide thickness at any given LET. At an LET of 80 MeVcm^{2}mg^{-1}, the thinnest oxides broke down under fields above 7 MV/cm, leaving room for improvement on SEGR tolerance in advanced technologies.

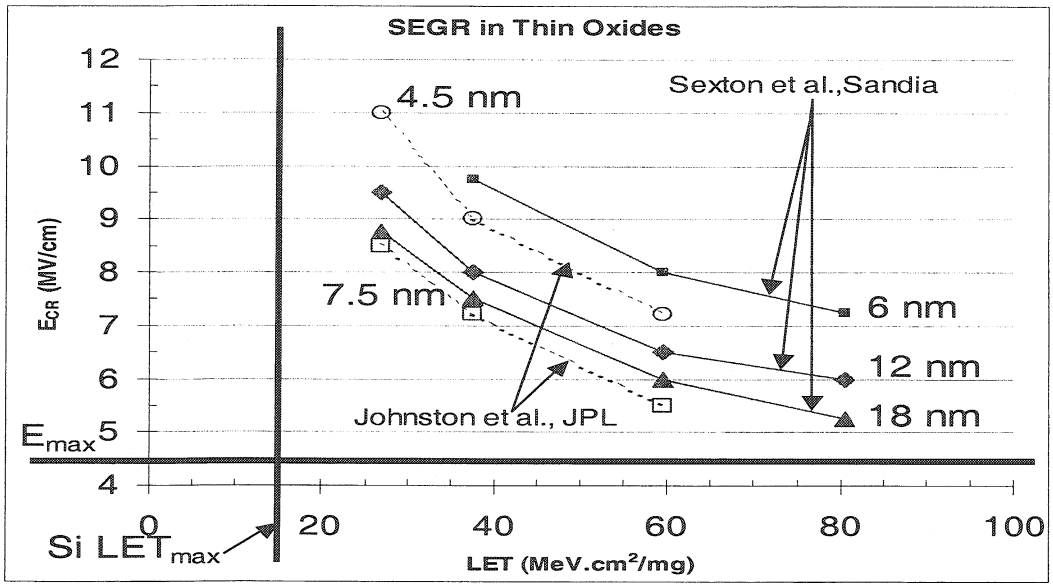

Figure 3.9 Critical electrical field as a function of particle LET. The horizontal bold solid line indicates the electric field across the oxides in an APV25 operated under normal conditions. The vertical bold solid line indicates the maximum LET expected in the CMS tracker.

Figure 3.9[3.16] gives some examples of measurements by other people illustrating the fact that the APV25 should not be susceptible to SEGR.

3.3.3 Single Event Upset

SEU is a non-destructive phenomenon which affects both dynamic and static memory registers that temporarily store logic states. It manifests itself as a soft error appearing in a device and is caused by the deposition of charge by an ionizing particle.

Figure 3.10 shows a simple 1-bit storage device illustrating the effect of an SEU, or "bit-flip".

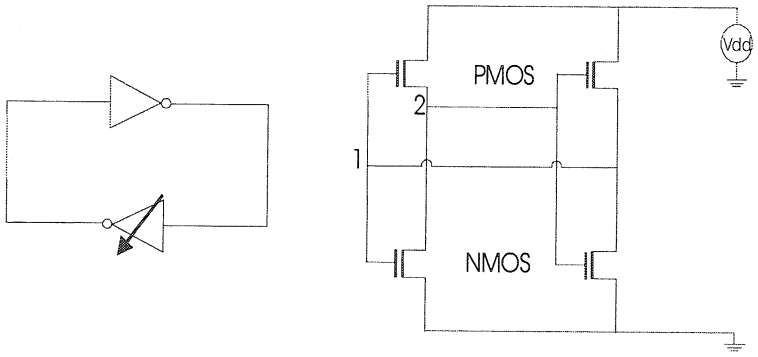

Figure 3.10 Schematic representation of a memory cell composed of two cross coupled inverters, and its circuit description.

The circuit is designed so that it has two stable states, one representing a stored '0' and one a stored '1.' In each state, two transistors are turned on and two are turned off. A bit-flip occurs when an energetic particle causes the state of the transistors in the circuit to reverse. Depending on the state of the cell, a sufficient injection of charge at points 1 or 2 will cause the state of the cell to invert. For an upset to occur two basic criteria must be met. The incident particle must provide a high enough LET in order to deposit a charge larger than the critical charge Q_{crit}, which is the minimum required to cause the circuit to invert logic state. The particle must also strike close enough to the sensitive part of the circuit, so that charge can be collected fast enough and an upset inducing path is present for the charge to flow. The sensitive volume (SV) of one memory element is defined as being that volume in which the incident particle must strike to cause an upset. A closer investigation of these sensitive volumes is discussed in section 3.5.1.

Figure 3.11 illustrates how an energetic particle produces a spurious electrical signal by inducing charge along its path, in the form of electrons and holes. These are collected by the electric field in the depletion region and charge funnel[3.17][3.18],

marked in Figure 3.11, and a current pulse appears, which can be large enough to produce an effect like that of a normal signal applied to the transistor.

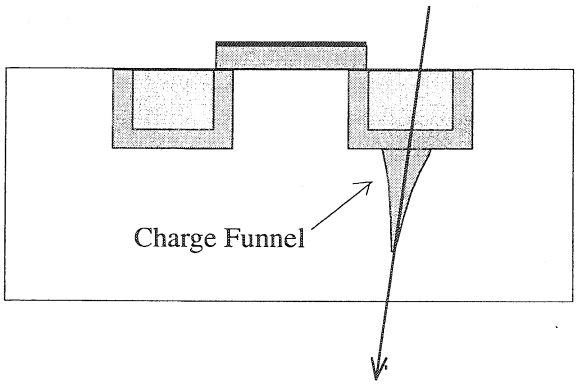

Figure 3.11 Interaction of an ion in an implant depletion region.

Some of the deposited charge will recombine, and some will be collected at the junction contacts. The net effect is a very short duration pulse of current that induces transient charges at internal circuit nodes. The magnitude of the charge, which is generally much larger for ions with high atomic numbers, depends on the energy and ion type, as well as the path length over which the charge is collected. Figure 3.12 shows a rough representative time response of charge collected from a single ion strike. The prompt charge is collected in much less than 1 ns, which is shorter than the response time of most MOS transistors.

Time [ns]

Figure 3.12 Illustration of collected current pulse shape.

This was predicted for a general very simple transistor simulated in EVEREST[3.19].

3.4 Making Experimental Measurements

In order to measure the upset cross-section of a specific memory device one must place the device in a beam of known particle species and energy. In general there are two experimental measurements that can be made to test susceptibility to SEU. One is to measure a single upset rate by placing the device in a radiation environment equivalent to the expected working environment. The other is to place the device in a beam of heavy ions and evaluate the cross-section of the device for many different incident LETs.

3.4.1 Individual upset rate measurements

It is often useful to measure a single upset rate which will give a good prediction of that expected in the final system. This can be done by putting the device into an environment equivalent to the expected radiation environment, such as a beam of light particles which are themselves incapable of causing upsets directly since their ionizing LET is too small. Commonly used particles are protons and pions. In these cases, upsets only occur by virtue of nuclear interactions between the beam and the silicon lattice. Knock-on silicon ions within the device may have enough LET to cause upsets. However, sometimes it is difficult to reproduce the radiation environment in a controlled experiment. One example is the space environment, where there are many particle species, including heavy ions and light particles such as protons and neutrons. In these cases it becomes more useful to use the second method. In other cases the low SEU rate means poor statistics in reasonable exposure times.

3.4.2 Use of heavy ions to scan LET

In real applications, the radiation environment that gives rise to SEUs usually produces a spectrum of energy deposition that covers a wide range of LET. Therefore upsets could be caused by small charge deposits just large enough to cause an upset or by extremely heavily ionizing particles, which deposit charge well in excess of the upset threshold. By measuring the response to specific LET values, one can predict the behaviour in any environment by simply calculating the probability of a specific LET being produced either directly or by a silicon lattice interaction with one of the particles in the environment. If the cross-section for this LET is already known, one can predict the upset cross-section in that environment.

Therefore, in order to examine behaviour with regard to SEU, it is necessary to know the upset probability for precise values of deposited charge. This is achieved by using mono-energetic heavy ion radiation containing a carefully selected ion species. These ions have a well-known surface LET, which allows an accurate calculation of the deposited charge. By selecting different charge states of an ion and adjusting the beam energy, it is possible to irradiate a device with a range of LETs, each giving rise to a corresponding upset cross-section. The cross-section, σ, is defined at normal incidence as:

$$\sigma = \frac{N_{events}}{\Phi}\left[cm^2\right] \qquad [3.5]$$

Where Φ is the total incident particle fluence, and N_{events} is the number of events (SEU) counted during the test.

3.4.3 The shape of the cross-section curve

Figure 3.13 shows a typical heavy ion cross-section curve where the upset cross-section σ of the device is plotted as a function of the normally incident ion LET. The curve has a distinct threshold, after which the energy is high enough to cause upsets, and then a relatively slow rise to saturation. It is not immediately obvious what causes the curve to have such a slow rise to saturation and we investigate some possible explanations.

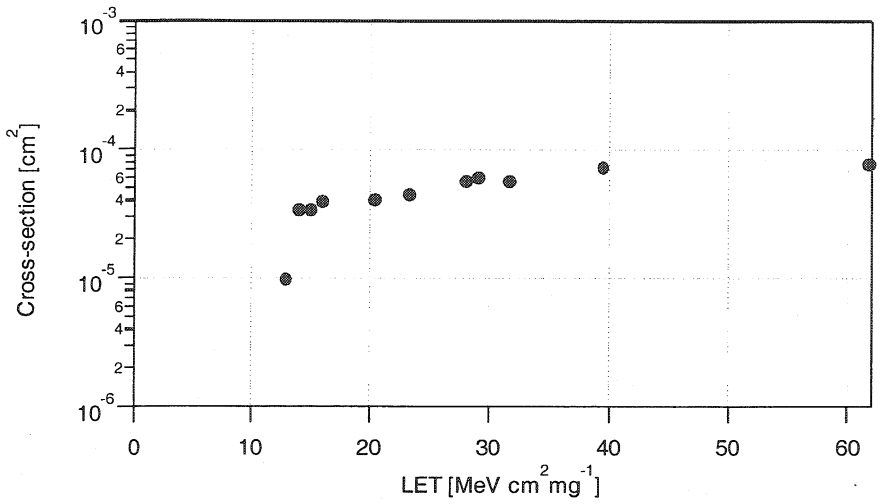

Figure 3.13 A typical set of experimental measurements of SEU cross-section vs LET.

It has become widely accepted to fit heavy ion cross-section curves with the Weibull probability distribution.

$$\sigma = \sigma_{sat}\left\{1-\exp\left[-\left(\frac{LET - LET_{th}}{W}\right)^{S}\right]\right\} \qquad [3.6]$$

Figure 3.14 shows the shape of this distribution. There is no physical significance, as yet, of this fit in terms of actual device physics. However, it has provided a reasonable method for extracting the threshold LET and saturating cross-section.

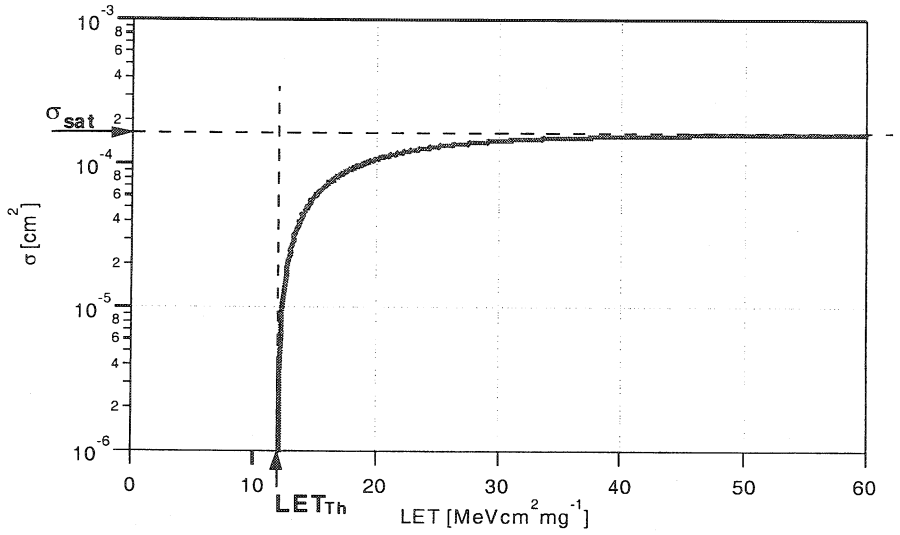

Figure 3.14 The Weibull fit.

LET_{th} is the minimum LET for which upsets can occur and σ_{sat} is the saturating cross-section for high values of LET. These two defining features of the device behaviour can then be used to reconstruct the curve analytically, which enables an interpolated prediction of the upset rate for other forms of radiation.

3.4.4 Possible Causes of the Cross-section Curve Shape

Consider first the case where all cells within a memory are identical and where the charge deposited by normally incident ions lies on an infinitesimally thin line. The sensitive volume is taken to be a cuboid so one would expect the cross-section curve to be a step function, since any ion that traverses the sensitive volume will deposit a set charge. Therefore if this charge is greater than the critical charge, an upset will always be caused. This is not the case in reality as can be seen from Figure 3.13; some explanations have been offered, but none have yet been accepted as the true cause.

Peterson et al. initially attributed the cross section curve shape to a distribution of memory cell sensitivities and cell-to-cell variations in critical charge resulting from processing[3.20]. In a subsequent article, some of the same authors recognize that other

physical effects can also be significant[3.21]. Massengill et al. have modelled cross section curves of SOI memories based on statistical effects of parasitic bipolar gain and critical charge distributions[3.24]. On the other hand, Langworthy has been able to model the upset cross section curves of several technologies by ignoring critical charge variations altogether[3.25]. This model is based on assuming the cell has a range of collection depths, which leads to the concept of an LET-dependent sensitive volume size. However, for modern technologies, such as the one used to make the APV25, variations in processing between cells in one chip should be very small, since the technology is extremely precise. Only a small proportion of the slowly rising cross-section could be attributable to such variations, since the continuing increase in saturation cross-section can be as much as one order of magnitude. Large enough variations in the collection depth of the sensitive volume are also unlikely because of the uniform nature of the implantation process.

I will show that in the case of normally incident ions, the slow rise in the cross-section curve can be attributed to switching on of different modes of upset within one cell. Each of these modes possesses a distinct threshold LET and cross-section, which combine to give the observed curve shape. Furthermore, different types of cells possess different numbers of modes and an individual cross-section curve shape.

3.5 Single Event Upset in the APV

3.5.1 Sensitive circuits in the APV

In the APV the sources of SEUs are digital pointers of the pipeline, the FIFO address memory, the I^2C control logic and data registers and other main control logic. Three types of digital memory element comprise all of these, they are the D Flip-Flop (DFF) (see Figure 3.15), D Flip-Flop with set (DFF-set) and the D Flip-Flop with reset (DFF-reset) (see Figure 3.16) and have been designed, using closed geometry transistors, at CERN[3.12]. Each of these circuits responds differently to deposited charge with characteristic upset thresholds and cross-sections, but the mechanisms by which SEU is manifested are identical in each case. SEUs can be induced in both the master and slave sections of each type of DFF, depending on whether the clock is high or low.

Both the master and the slave of the DFF are two cross-coupled inverters with two transmission switches. The DFF-reset is the same, except there are more transistors in

order to couple the reset input, rather than two cross-coupled inverters; there is one inverter and one NAND gate. The DFF-set is very similar to the DFF-reset with the same number of transistors, but a slightly different layout.

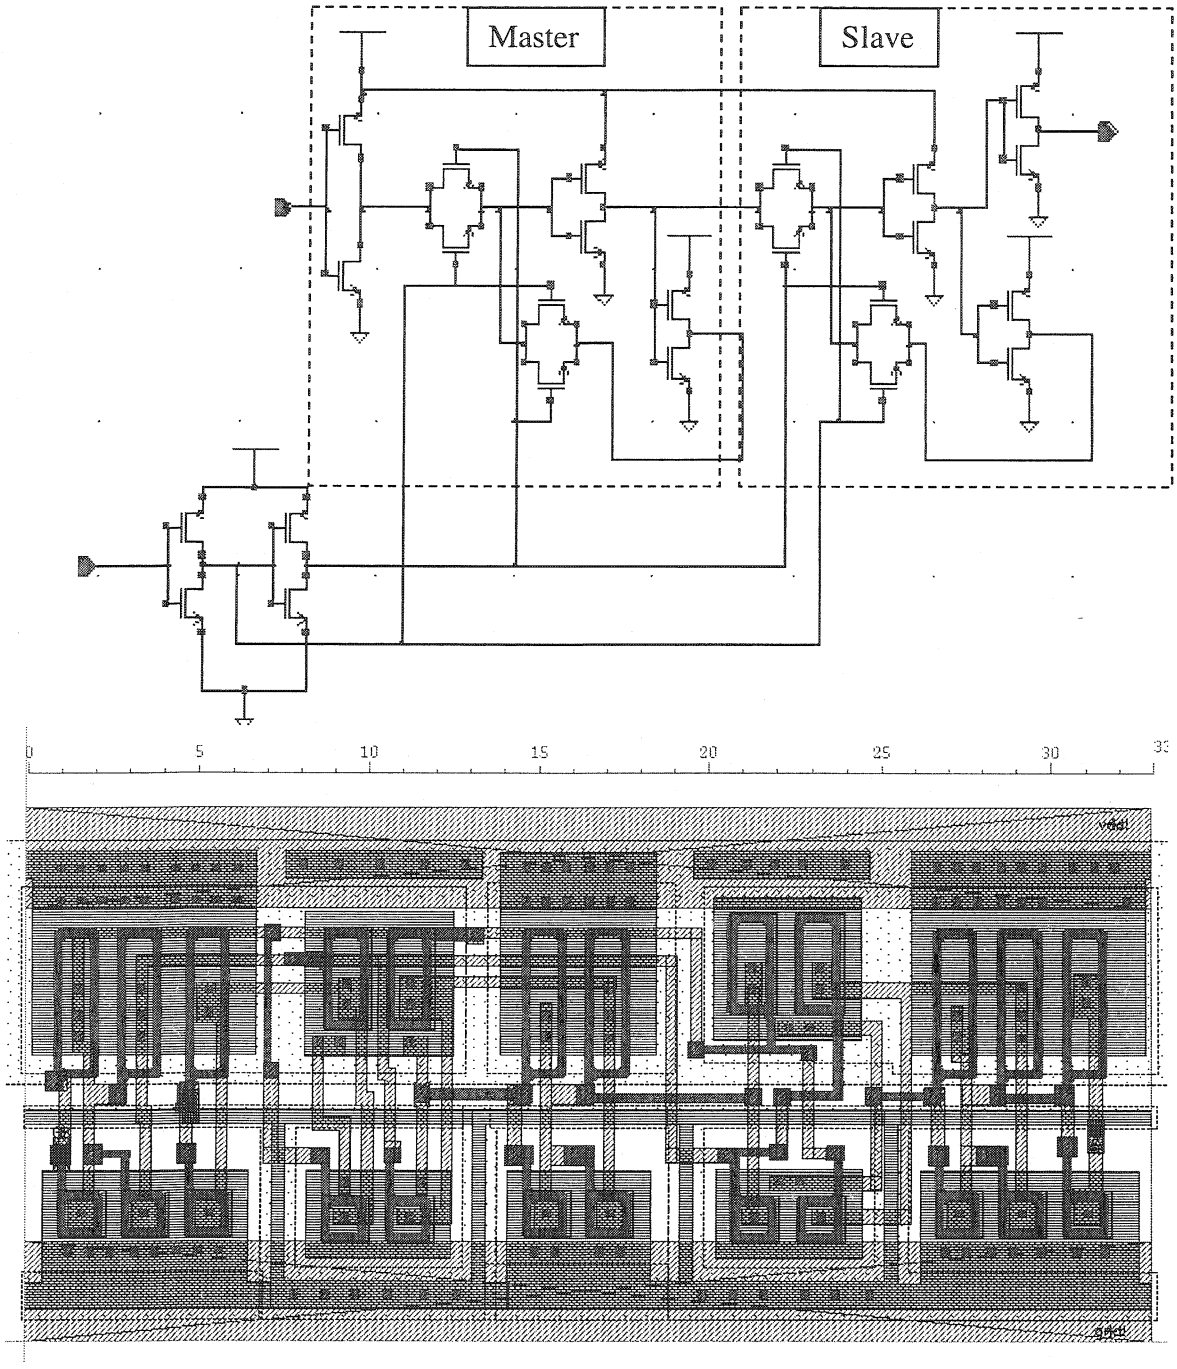

Figure 3.15 DFF Schematic and layout diagram.

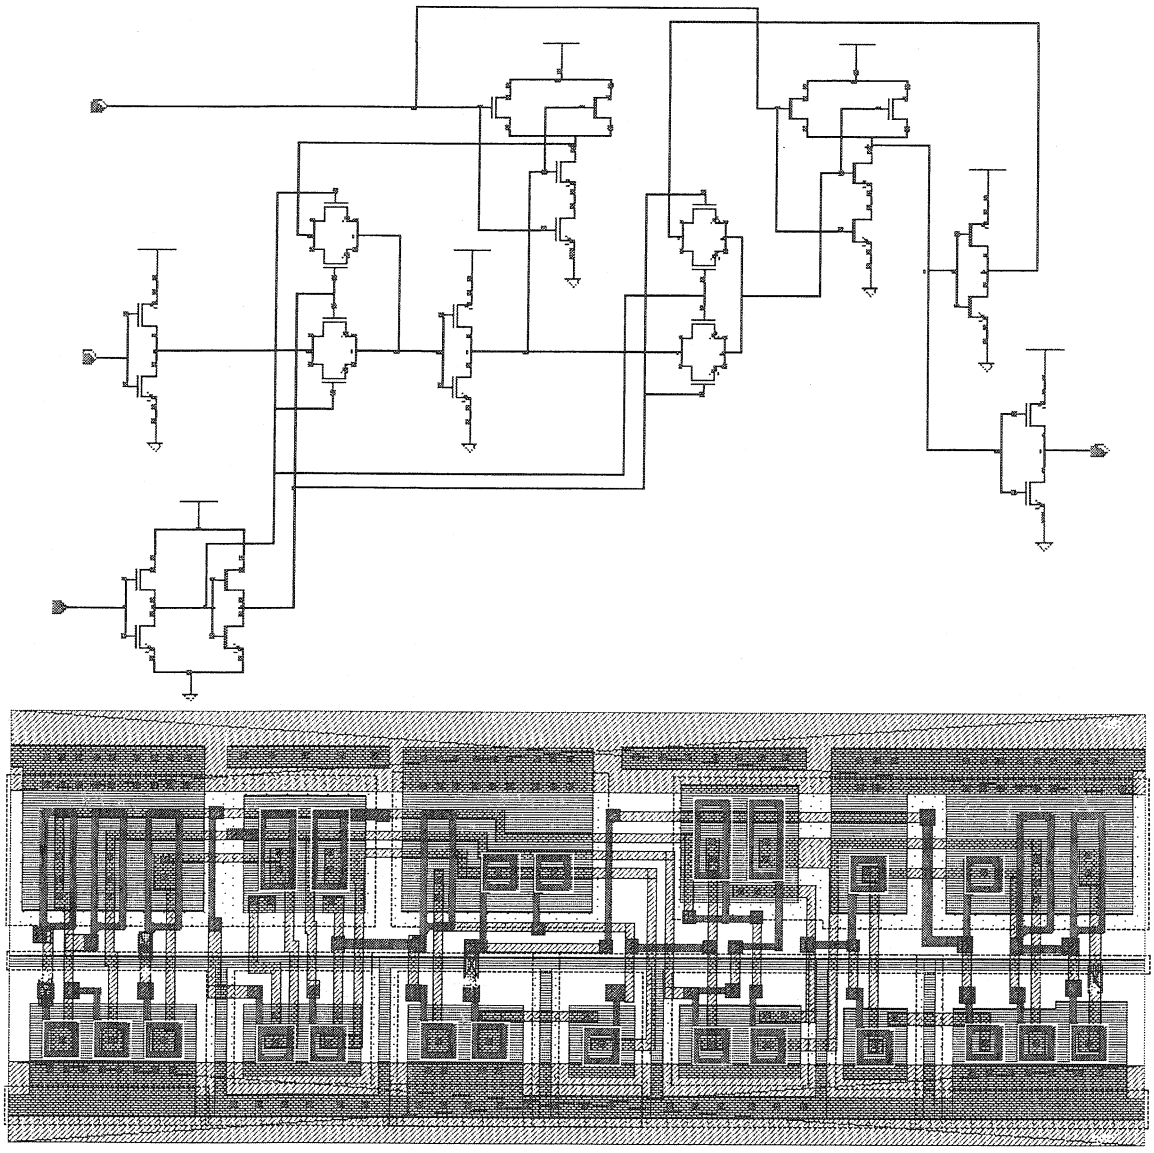

Figure 3.16 DFF-Reset schematic and layout diagram.

3.5.2 Simulations of SEU in APV25 circuits

To get a better understanding we will investigate the behaviour of the DFF. Since the bias registers in the APV25, which are made up of simple DFFs, are normally un-clocked, we only need to consider the slave of the DFF in order to predict their behaviour. Similar investigations have been performed for master and slave in each of the three types of DFF, but we need only consider one circuit, since the principles are identical for all.

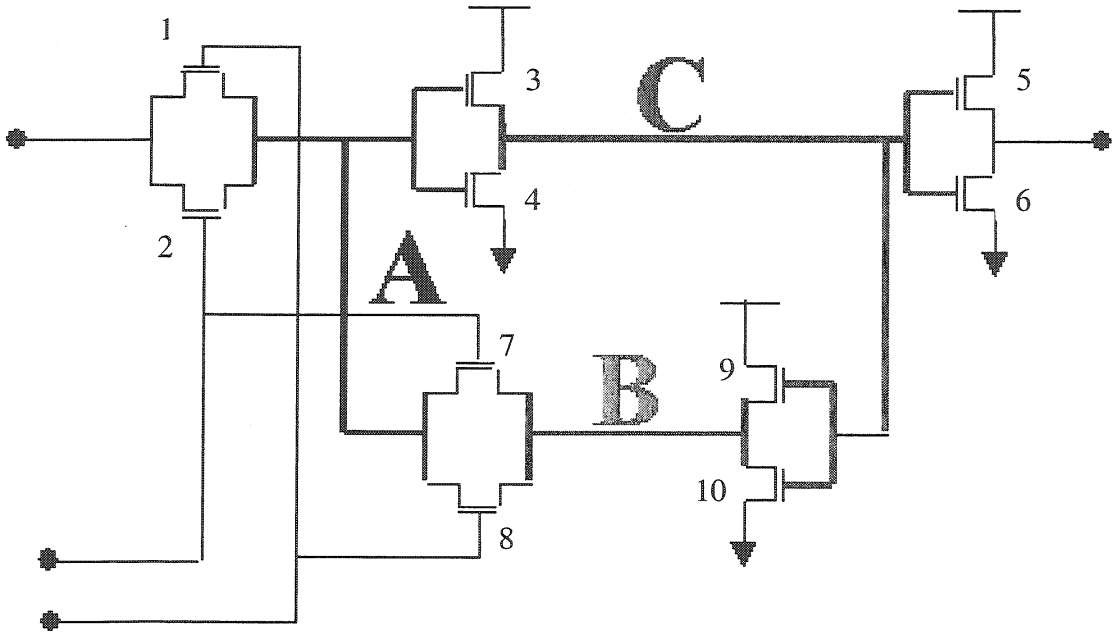

Figure 3.17 schematic of DFF slave.

There are three nodes in which charge collection can give rise to an upset, which are labelled A, B and C in Figure 3.17. Each node has two critical charges, one for each state transition, 1 to 0 and 0 to 1. The sensitive parts of these nodes are the depletion regions surrounding the highly doped n^+ and p^+ implants, which form the drains and sources of the FETs. The n^+ implants are capable of collecting electrons, therefore charge collected here can only cause an upset if the state of the node, to which the implant belongs, is high. The opposite is true for the p^+ implants. In total there are 5 sensitive n and 5 sensitive p implants, which are shared by the three nodes in the following way: nodes A and B contain 2 of each, and node C contains one of each.

Sensitive Implant	Implant Type	Associated Transistor	Associated Node	Surface Area [μm^2]
1	n-source	2	A	3.85
2	n-drain	4	C	0.65
3	n-drain	8	A	0.65
4	n-source	8	B	3.85
5	n-drain	10	B	0.65
6	p-source	1	A	4.97
7	p-drain	3	C	2.77
8	p-drain	7	A	1.96
9	p-source	7	B	4.97
10	p-drain	9	B	2.77

Table 3.1 Sensitive implants of DFF.

Table 3.1 gives a breakdown of all the sensitive implants, the transistor and node to which they belong, and their surface area. By summing these areas in the correct way, an estimate of the normal incidence upset cross-section can be made

3.5.3 Simulating SEU in the APV25

In order to establish the critical charge for each mode the circuit was constructed in SPICE and the injection of a simple piecewise linear current pulse (50ps rise and 100ps fall) in the three nodes was used to simulate the charge deposition of an ion strike. EVEREST simulations have been performed in order to investigate the charge collection efficiency of both the n and p implants, and have shown that the p implants have 50% the efficiency of the n implants, which collect more than 95% of the charge deposited in the sensitive volume. This work has also provided a realistic charge pulse shape, which was comparable to the simple current pulse used in the SPICE simulations.

To find the critical charge one must inject pulses with an increasing amount of charge, checking the response of the circuit after each. The current pulse in the 'on' NMOS transistor looks like that shown in Figure 3.18.

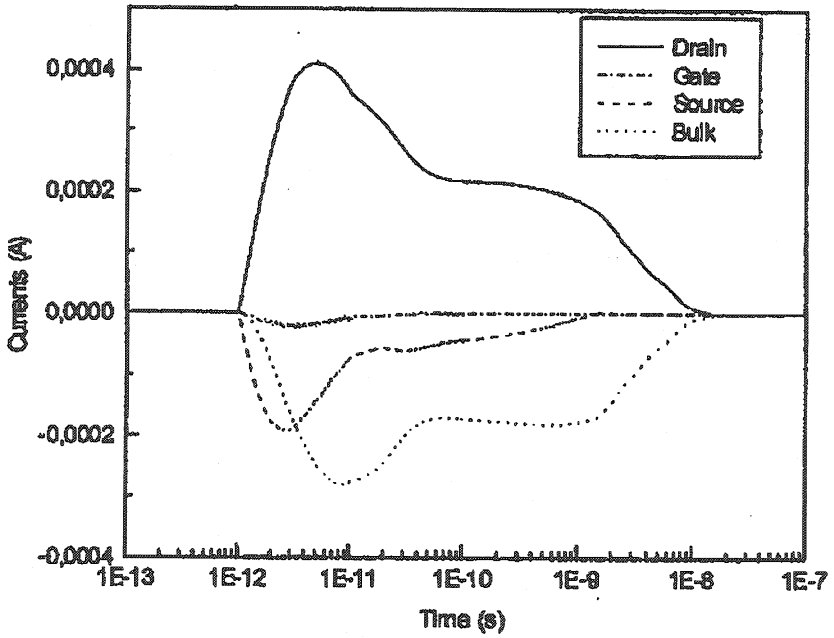

Figure 3.18 Currents of the on NMOS transistor.

To see the effect on the state of the DFF one can look at the voltages, this time for the 'off' NMOS transistor. These figures were taken from similar simulations performed in [3.23].

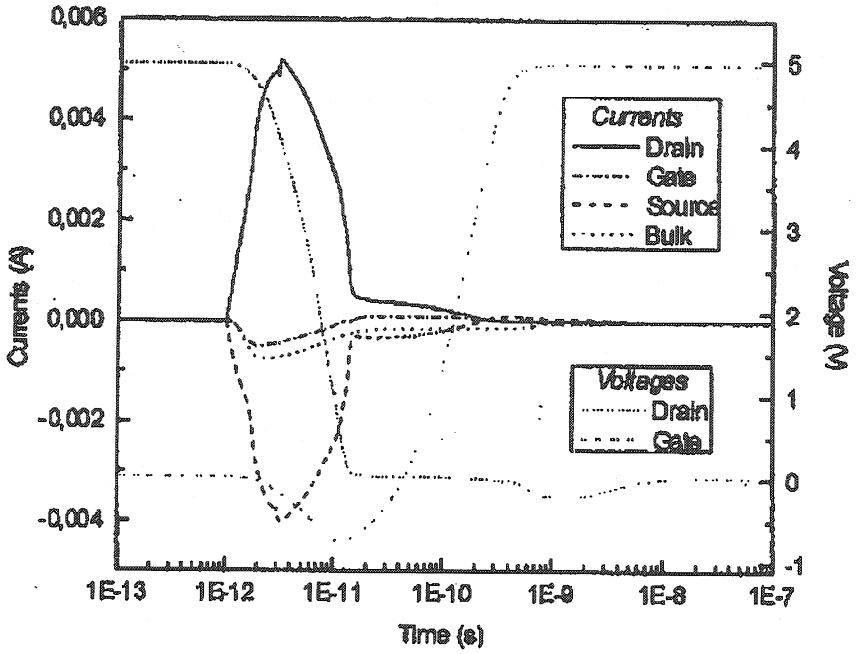

Figure 3.19 Currents and voltages during a charge injection greater than critical.

This work has been performed for all upset modes in the three types of DFF, establishing the critical charges. Ongoing work will include more Spice simulations with a more realistic pulse shape, but this should not affect the predicted critical charges to a large degree.

Table 3.2 shows the simulated critical charge along with the sensitive areas, and the implant type in which the charge must be collected, for each mode of upset.

	Upsets from 1-0			Upsets from 0-1		
	Q_{crit} [fC]	n or p	σ [μm^2]	Q_{crit} [fC]	n or p	σ [μm^2]
Hit on A	206	n	7.7	442	p	9.94
Hit on B	319	n	1.3	638	p	4.73
Hit on C	592	p	2.77	274	n	0.65
Total	-		11.77	-		15.32

Table 3.2 Sensitive area and critical charge for six modes in a DFF slave.

The lower charge collection efficiency increases the critical charge for the p-modes. If one converts the critical charges into equivalent LET values for ions in silicon, it is possible to predict the shape of the heavy ion upset cross-section curves.

The simple DFF is only used in the un-clocked registers of the APV, so no upsets can occur in the master except when the memory is being loaded. The time it takes to load is only one clock cycle (25 ns), so it is not worth considering the DFF master in the

analysis of the APV. However, the pipeline logic and control logic are continually clocked and therefore both the master and the slave are sensitive to upsets. Simulations for master and slave in both the DFF-set and DFF-reset have been performed, the results of which are summarised in Table 3.3 and Table 3.4. The number of modes is slightly larger for these circuits, since they are more complicated and have more sensitive nodes and implants. The modes have been arbitrarily labelled.

		Upsets from 1-0			Upsets from 0-1		
		Qcrit [fC]	n or p	σ [μm^2]	Qcrit [fC]	n or p	σ [μm^2]
Slave	A	638	P	2.77	293	N	0.65
	B	304	N	2.61	532	P	3.25
	C	161	N	7.40	406	P	10.20
	D	428	N	0.65	-	-	-
	Total			13.43			14.10
Master	E	308	N	1.30	586	P	3.50
	F	371	N	7.40	690	P	10.20
	G	720	P	4.75	356	N	1.30
	H	540	P	10.20	248	N	7.40
	Total			23.65			22.40

Table 3.3 Sensitive area and critical charge for all modes in a DFF-set.

The total sensitive areas for both 1-0 and 0-1 transitions in the master of both circuits are similar, with an average of ~23μm^2. However, there is some difference between the 1-0 and 0-1 transitions in the slave. To make an accurate prediction of the cross-section of a particular circuit made up of a number of either of these DFFs it would be necessary to know the number of bits normally at 1 and the number of those at 0.

		Upsets from 1-0			Upsets from 0-1		
		Qcrit [fC]	n or p	σ [μm^2]	Qcrit [fC]	n or p	σ [μm^2]
Slave	A	480	P	1.30	285	N	0.65
	B	326	N	1.30	690	P	4.70
	C	206	N	7.70	496	P	10.00
	Total			10.30			15.35
Master	D	345	N	1.30	802	P	4.75
	E	394	N	7.70	1350	P	10.00
	F	562	P	3.30	353	N	1.30
	G	436	P	10.50	237	N	7.50
	Total			22.80			23.55

Table 3.4 Sensitive area and critical charge for all modes in a DFF-reset.

3.6 Modelling the Heavy Ion Upset Cross-Section in the APV25

3.6.1 Total Sensitive Area of APV25

The circuits in the APV25 are made up of a certain number of the different types of DFF. By using the sensitive areas of 1-0 and 0-1 transitions and accounting for the normal state of each of the bits, one can make a rough estimate of the saturated upset cross-section for each circuit.

The bias registers do not have a definite state for each bit, since this depends upon what they are programmed with, but we know what they are set to in any experiment, therefore we can predict the cross-section for that experiment. The experiments performed on the APV25 presented in the next chapter were done with a total of 29 bits set to 1, so it follows that:

$$\sigma_{registers} = 29 \times 11.77 + 80 \times 15.32 = 1567 \mu m^2 \quad [3.7]$$

The pipeline logic is also straightforward, since it is made entirely from the DFF-reset and has all but a few bits set to 0, therefore it is a close enough approximation to assume that all upsets are 0-1 transitions. We must also account for the fact that the pipeline logic is being clocked and so the master and slave are sensitive for half of the time, which is equivalent to them having half the sensitive area for the same amount of time:

$$\sigma_{pipeline} = \frac{768 \times 10.3 + 768 \times 23.8}{2} = 13094 \mu m^2 \quad [3.8]$$

The FIFO is more difficult, since it is made up of both the DFF and DFF-reset. Also the number of bits which are sensitive is not necessarily the same as the number of bits in which upsets can be detected. Since we are attempting to predict the behaviour of the APV25 in a heavy ion beam under controlled conditions, we will again consider the precise case relating to this test, during which the APV is triggered once and then reset, thus leaving only the output data frame vulnerable to upsets occurring in the FIFO read pointer and the first 8-bit address memory in the buffer. This results in a total number of 8 DFFs and 32 DFF-resets. In this case none of the cells are being clocked except for loading of the buffer, so only the slaves are sensitive. The pointers are all, except one,

set to 0 as before and the address buffer contains the address of the pipeline location that was read out, which on average contains four 1s and four 0s. It follows that:

$$\sigma_{FIFO} = 4\times15.32 + 4\times11.77 + 31\times10.3 + 15.3 = 443\mu m^2 \quad [3.9]$$

The control logic is the most difficult circuit to predict, because it is a mixture of all three DFFs and it is constantly being clocked and changing the states of the bits. The best estimate one can make is to assume that on average half of the bits are 1 and half 0. The exact distribution of the different DFFs was also not established. In this case the best estimate is made by using the average of all the possible sensitive areas, which is ~17 μm^2. It follows that:

$$\sigma_{control} = 167\times17 = 2839\mu m^2 \quad [3.10]$$

Table 3.5 summarises these results giving the total sensitive area of each circuit in cm^2, which is a more common unit for measurements of cross-sections. The main contributor to the sensitive area is the pipeline logic.

	Pipeline logic	FIFO	I^2C Registers	Control logic
No. of DFFs	768	40	109	167
Sensitive Area (cm^2)	13.1×10^{-5}	0.43×10^{-5}	1.57×10^{-5}	2.8×10^{-5}

Table 3.5 Number of sensitive bits in APV25 sub-circuits and the total cross-sections.

3.6.2 Converting toThreshold LET

Using 3.6 eV per electron-hole pair (energy losses to phonons leads to a mean required energy larger than the silicon band-gap), the critical energy is given by:

$$E_{crit} = \frac{Q_{crit}}{e}\times3.6 \quad [3.11]$$

We can convert this value into a measure of LET using equation[3.12]. The depth z is a variable and represents the sensitive depth of the implants; this is the sum of the depletion region depth and the ion impact resultant funnel length. This value can be determined by fitting a theoretical curve to the data and extracting the threshold LET. LET_{th} is the minimum LET for which upsets of a particular mode can occur.

$$LET_{th} = \frac{E_{crit}}{z\rho} \quad [3.12]$$

3.6.3 Modifying the Weibull fit to include different modes

In previous research[3.24][3.25] the existence of modes has been ignored; it was assumed that there was just one characteristic critical charge. Previous measurements of σ have been fitted with the Weibull function, which, as we have seen, gives the impression that σ rises slowly as the incident LET increases. From section 3.5.3 it is clear that there should be a number of distinct modes of upset each with a different threshold LET.

If we now go back to the initial assumption of a more abrupt step-like function for each mode, we can sum the cross-sections for each mode to give the full curve:

$$\sigma = \sum_{i=0}^{Nn} \sigma_n^i(LET) + \sum_{i=0}^{Np} \sigma_p^i(LET) \qquad [3.13]$$

Where N_n is the number of n-modes and N_p is the number of p-modes. Here we are not distinguishing between 1-0 and 0-1 transitions, but there is a difference and therefore there should be a separate equation for each. The components of σ_n^i and σ_p^i can be represented by a single Weibull curve, as in equations[3.14] and [3.15]. Each mode has its own values of LET_{th} and σ_{sat} , which can be taken directly from the simulations.

$$\sigma_n^i = \sigma_{nsat}^i \left(1 - \exp\left\{ -\left[\frac{LET - LET_{nth}^i}{W_n} \right]^{s_n} \right\} \right) \qquad [3.14]$$

$$\sigma_p^i = \sigma_{psat}^i \left(1 - \exp\left\{ -\left[\frac{LET - LET_{pth}^i}{W_p} \right]^{s_p} \right\} \right) \qquad [3.15]$$

The standard shaping parameters W and s are left in for generality. However, further work should establish the exact shape of the individual mode curves, by careful consideration of the implant profile and more simulations of charge collection.

Using these equations, taking the threshold values directly from Table 3.2 and using equation[3.12] with a sensitive depth of 1μm, a sensible estimate[3.26], one can create the theoretical cross-section curves for a simple DFF for both 0-1 and 1-0 transitions(see Figure 3.21. and Figure 3.20). Variations in the true sensitive depth will

have the effect of contracting or expanding the curve in the LET domain. At this stage the important aspect is the general shape of the curve, in the next chapter is it shown how the curve can be fitted to the data and thus the sensitive depth can be extracted.

It is easy to see from Figure 3.20 that any steps in the experimental curves, could quite easily be lost due to statistical errors, and reproducing a curve like this, in practice, would require a very thorough scan of incident LET, which is not a trivial thing to achieve in a small amount of time. However, Figure 3.21 shows a more distinctive step, which may provide us with a better chance to see the structure experimentally.

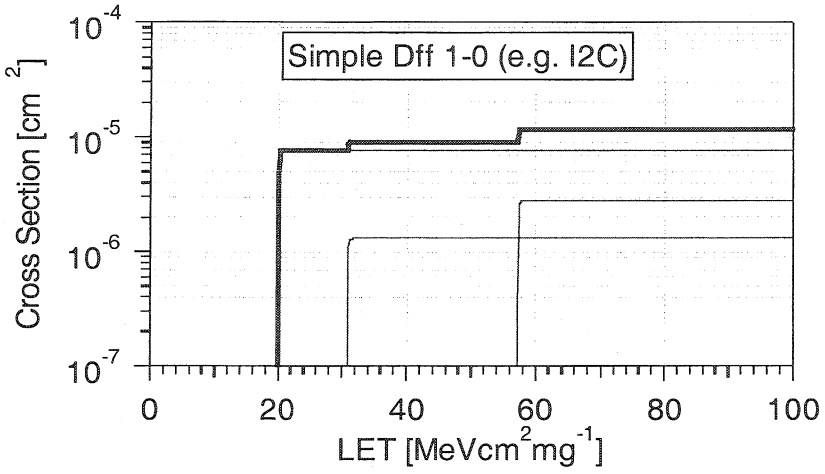

Figure 3.20 Theoretical cross-sections for 1-0 transitions in a simple DFF.

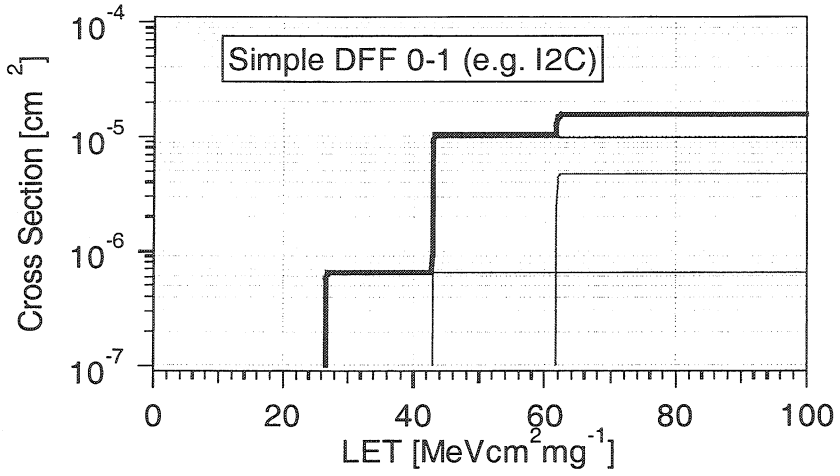

Figure 3.21 Theoretical cross-sections for 0-1 transitions in a simple DFF.

Using the same method one can predict σ of the DFF-reset, which is used in the pipeline. In this case we only need look at 0-1 transitions, since these will dominate the experimental curve. The first three steps are all n-modes with the initial mode at

~$3x10^{-5}$ cm^2. The fourth step is the first p-mode, which increases σ to ~ $7x10^{-5}$ cm^2, a step of $4x10^{-5}$ cm^2.

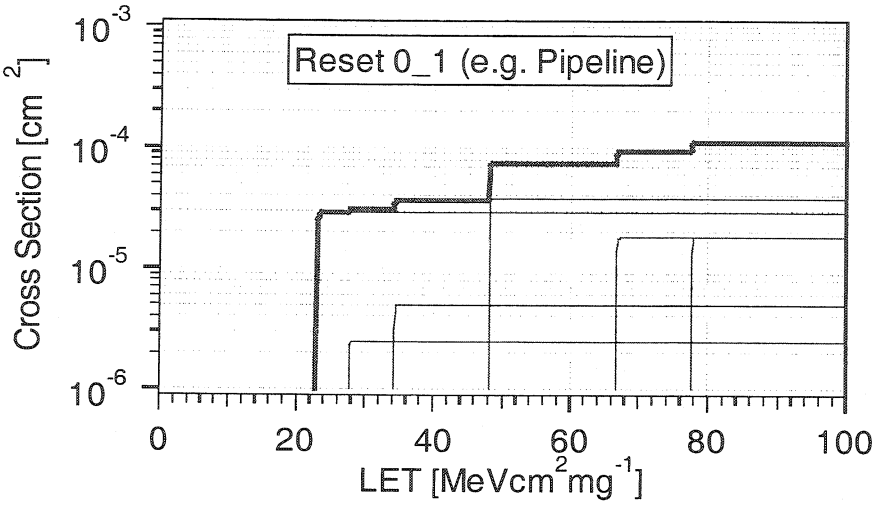

Figure 3.22 Theoretical cross-section for pipeline logic.

Circuits such as the control logic would exhibit a smoother rise to saturation because of the large number of modes present due to contributions from master and slave, and both 1-0 and 0-1 transitions.

3.7 Modelling the CMS Environment

By fitting the experimental data with a curve, it is possible to interpolate predictions of the upset rate for other forms of radiation. In the case of the CMS tracker, the required calculations are complicated since the incident particles are typically of single charge and therefore only cause large enough ionization by virtue of interactions with silicon lattice sites. Monte-Carlo simulations are required and have been developed at CERN[3.22] The following is a brief description of the method used.

At LHC the heavy ions that will be produced by hadronic interactions mainly include protons, neutrons and pions. Such interactions will generally produce a shower of secondary particles and either one nuclear recoil or sometimes several fragments. The typical energies of the recoils rarely exceed 10MeV, therefore their range is generally below about 10μm. Thus for such recoils to cause SEU they must be created locally in the vicinity of the sensitive volume. Figure 3.2 shows that a 10MeV silicon ion deposits about 30 GeV/cm, which is 4 orders of magnitude above minimum ionizing.

The basic method employed in the simulations is to consider all of the most important reactions which can give rise to recoils and fragments with enough energy to

cause an SEU. Once the energy spectrum for each recoil type has been established it remains to investigate the probability distributions of their energy deposition, which is ultimately responsible for SEU. Various different simulation codes have been employed and adapted for simulation. To predict the spectrum of the particle emissions from inelastic nuclear reactions, the FLUKA simulation package[3.22] has been used. For low-energy neutrons (E<20MeV), the information contained in the ENDFB-VI cross-section tabulations is used, which also contains values for low-energy neutron elastic scattering. At energies above 20MeV, a relatively simple optical model based on Glauber theory[3.27] is used to calculate the elastic differential cross-section. For protons below 95MeV, fits to experimental data are used instead of the Glauber model. However, the contribution of elastic scattering, at energies above 20MeV, to the SEU rate is very small.

A large part of the work concerns the stopping of the produced nuclear fragments. The theory describing the energy loss of slow ions in solids is provided in section 3.1.5. Figure 3.23 shows the ranges of some ion species as a function of the energy. For the energies and ion types of interest, the ranges are typically a few microns. The short ranges, and much higher dE/dx, of heavy fragments suggest that the main emphasis must be put on events in the close vicinity of the SV.

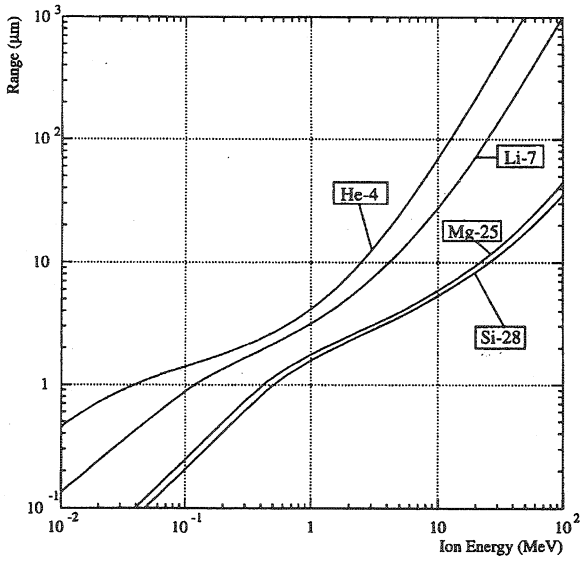

Figure 3.23 Ranges of a few ion types in silicon.

In order to calculate the charge deposited, it is necessary to choose a geometry for the SV. Since the exact shape of the SV in the APV25 is not easy to define, a simple cubic model has been used, with some different SV sizes for comparison, namely:

1x1x0.5 μm^3, 1x1x1 μm^3, 1x1x2 μm^3 and 2x2x2 μm^3, where the last value always gives the depth of the sensitive layer.

Figure 3.24 shows the simulated energy spectra of recoils for various incident neutron energies above 20 MeV. As expected the maximum energy of fragments increases with the energy of the incident neutrons.

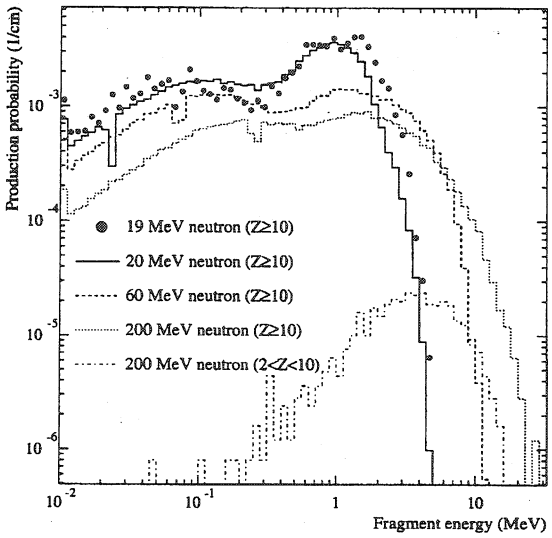

Figure 3.24 Fragment energy spectra from n-Si scattering at energies above 20MeV.

In order to predict the total probability distribution of energy depositions within the SV for a particular primary radiation spectrum, one must run a Monte-Carlo simulation including the full energy spectrum of incident particles, and calculate the probability of each fragment passing through the SV and depositing energy greater than, or equal to an energy (E_{DEP}). However, these simulations are extremely complicated and take many days to run, and no plots of the total probability distribution in CMS have yet been produced. Nevertheless, we can gain insight into these curves by considering a few different incident energies. Figure 3.25 shows the energy spectra of neutrons and charged hadrons in the tracker. The APV25 will be located in the outer tracker where the neutron flux peaks at about 1MeV and charged hadrons at about 100MeV.

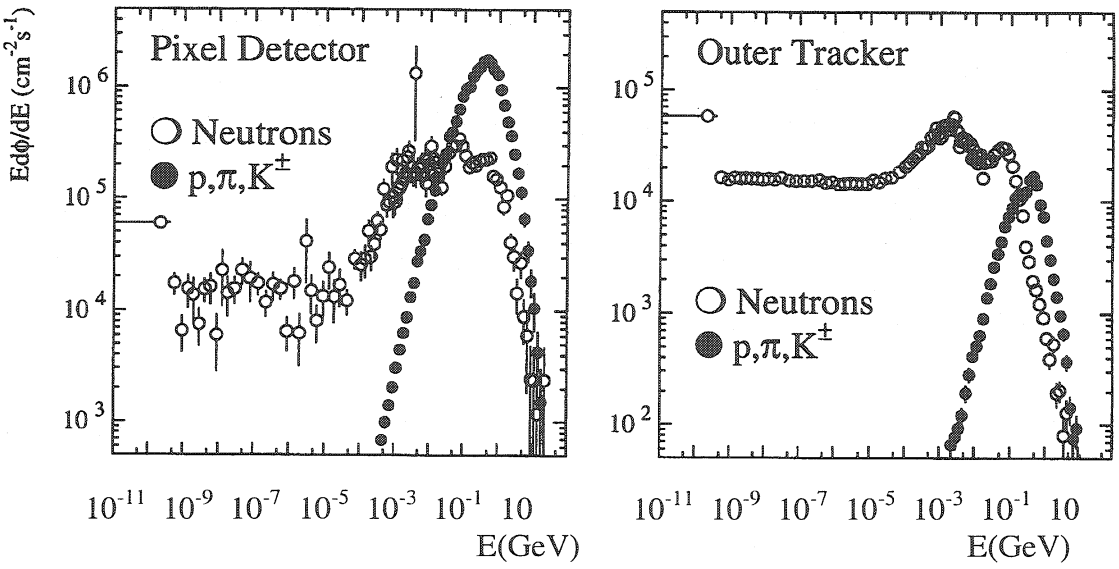

Figure 3.25 Energy spectra of neutrons and charged hadrons in the inner and outer tracker.

Figure 3.26 shows the probability distribution of deposited energies within a 1x1x1 μm^3 SV for four incident proton energies, which are representative of the energies shown above. There is a strong dependence on E_{DEP} with a maximum tending towards ~ 5 MeV for 200MeV protons.

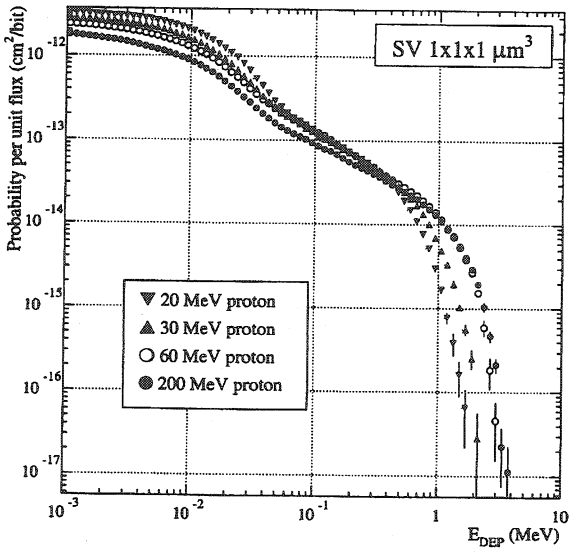

Figure 3.26 Energy deposition probabilities for protons of different energies. The curve shows the probability to have, within the SV an ionizing deposition greater or equal to the indicated E_{DEP} value.

The typical critical energy for SEU in the APV25 is predicted to be ~ 4 MeV, in this region there is a very strong dependence on E_{DEP} , therefore small variations in the true critical energy could have a large effect on the upset cross-section under such radiation.

A similar plot for neutrons is shown in Figure 3.27 with similar strong dependence around the same value of E_{DEP}.

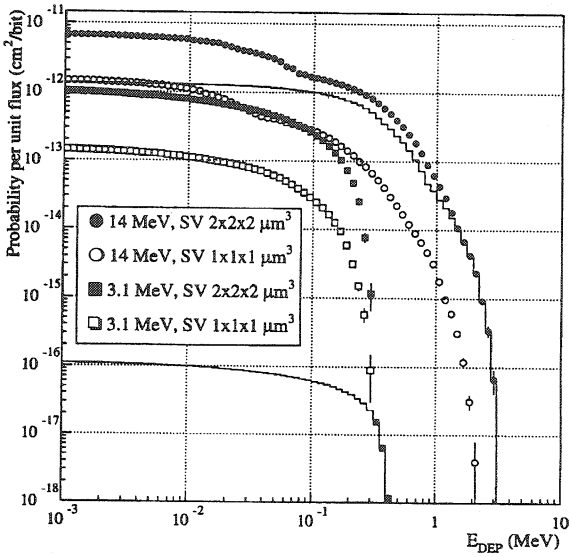

Figure 3.27 Energy deposition probabilities for 3.1 MeV and 14 MeV neutrons with two different SV sizes. The solid lines show the contribution of (n,α) reactions for the 2x2x2μm³ cases.

These simulations are encouraging since the expected critical energy is so close to the edge of the spectrum for both neutrons are protons. There is a chance that the actual sensitive volume in the APV25 could be larger than the simulated sizes shown above. In this case the upset cross-sections for the APV in the CMS tracker will strongly depend on exactly where the true critical energy lies in relation to the E_{DEP} probability spectra.

3.8 Conclusions

The variety of radiation effects has only been touched upon in this chapter with the main focus of attention on total dose effects in MOS transistors and SEU in DFFs. It is clear that with such a variety of effects it is very important to thoroughly investigate electronic systems intended for operation in radiation environments such as that of CMS. With a better understanding of the mechanisms behind the effects the chances of being able to design radiation hard systems is much improved.

Chapter 4

Radiation Testing Results

Both individual test transistors and the APV chip, from different fabrication processes and foundries, have now been tested to establish their degree of tolerance to the radiation effects discussed in Chapter 3. The APV6, fabricated in the Harris bulk AVLSIRA process, and some of the accompanying test transistors were irradiated to 100 Mrads(SiO_2). The APV25 and test structures, fabricated in two different foundries, both employing the same 0.25μm process, have been irradiated to 50 Mrads(SiO_2). Heavy ion irradiations of the APV6 and APV25 chips have been performed in order to establish their susceptibility to single event effects. Results from all of these studies are presented in this chapter.

4.1 Total Dose Studies of APV Test Transistors

Irradiations of Harris APV6 devices were performed in 1998, and the more recent studies on the APV25 devices in 1999 and 2000. The results from the two irradiation studies are presented and comparisons between the two technologies are made in a separate section following the results. The method for both sets of irradiations, however, was broadly the same.

4.1.1 Measurement of Transfer Characteristics

The characteristics of transistors can be found by measuring variation of the source-drain current (I_{DS}) with gate voltage (V_{GS}). The I_{DS} v. V_{GS} transfer characteristics were measured using the circuit in Figure 4.1.

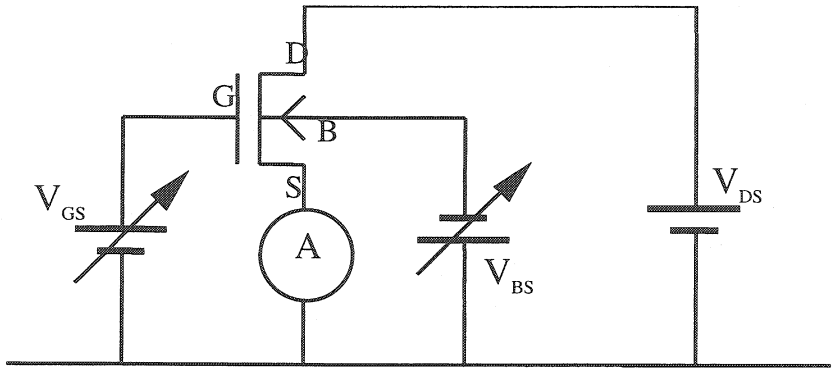

Figure 4.1 Circuit diagram for transistor test setup.

The bias conditions for these measurements are shown in Table 4.1. The chosen voltages were the normal operating voltages of the devices within the APV chip.

	Harris	0.25μm
$\|V_{DS}\|$	1 Volt	1 Volt
$\|V_{GS}\|$	Swept from sub-threshold region to strong inversion	Swept from sub-threshold region to strong inversion
$\|V_{BS}\|$	Fixed values of 0 and 2.0 Volts	Fixed values of 0 and 1.25 Volts

Table 4.1 Bias conditions for irradiated test transistors.

From these measurements it is possible to calculate the threshold voltage and transconductance of the transistor. To define these terms it is necessary to consider a plot of drain current vs gate voltage for a typical transistor. Figure 4.2 shows square root of the current vs gate voltage taken from a NMOS transistor.

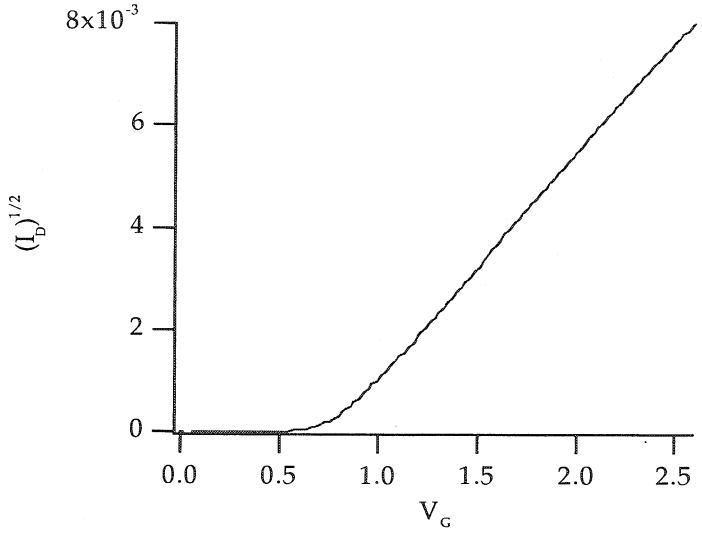

Figure 4.2 $\sqrt{I_{DS}}$ v. V_{GS} for an NMOS transistor.

The threshold voltage is defined as the point at which an extrapolated straight-line fit crosses the voltage axis. The transconductance is defined as dI_{DS}/dV_{GS} and is denoted g_m, which varies with the current.

4.1.2 Noise Measurements

The voltage noise spectra of the transistors as a function of frequency (f) were obtained using a network/spectrum analyser (Hewlett Packard model HP4195A). The system is described in more detail in [4.1]. The drain voltage (V_{DS}) on the device under test (DUT) was set to 500 mV and a self biasing circuit regulated the gate voltage (V_{GS}) to a value allowing for a drain current (I_{DS}) of 500 μA.

4.1.3 Harris Devices

Studies up to 10 Mrads plus an annealing step were completed for four different processing runs. Measurements were made on the following device geometries:

Type	Label	Width [μm]	Length [μm]
PMOS	P1	2000	1.4
PMOS	P5	10	10
NMOS	N1	2000	1.2
NMOS	N5	10	10

Table 4.2 Irradiated Harris test structures.

These geometries are representative of devices used in the APV6. The gate referred noise spectral density for the widest/shortest devices have also been measured under the conditions shown in Table 4.3.

Transistor	I_{DS} [μA]	V_{DS} [Volts]	V_{BS} [Volts]
P1	500	-1	0 and 2
N1	500	1	0

Table 4.3 Operating conditions of transistors with noise measurements.

Irradiation and annealing was performed under bias and devices were kept under bias (as much as possible) between irradiations and afterwards. Irradiations were carried out using the Brunel University ^{60}Co gamma-ray source and subsequent measurements made independently at Brunel and Imperial College for cross-checking. Four die from each of three processing runs were split into two batches as shown in Table 4.4.

Batch	Run	Die number	Irradiated in three stages to
1	APV5	5 and 7	300 krads, 1Mrads, 3 Mrads
	APV6	2 and 5	
	MX	1 and 2	
2	APV5	6 and 8	500 krads, 3 Mrads, 10 Mrads
	APV6	6 and 6a	
	MX	4 and 5	

Table 4.4 Categorisation of Harris devices.

4.1.4 0.25μm Devices

The tested transistors are listed in Table 4.5 As with the APV6 test structures, these have varying conducting channel dimensions. These dimensions are representative of the transistors that comprise the APV25.

	Type	Label	Width [μm]	Length [μm]
Foundry A	PMOS	P1	2000	0.36
	PMOS	P2	2000	0.50
	PMOS	P3	2000	0.64
	NMOS	N1	2000	0.36
Foundry B	PMOS	P4	25	2.5
	PMOS	P5	2000	0.24
	PMOS	P6	2000	0.36
	PMOS	P7	2000	0.48
	PMOS	P8	2000	0.60

Table 4.5 Type and dimensions of test transistors.

Two different foundries owned by the same company and employing the same 0.25μm process provided the transistors. The gate oxide of this process is approximately 5.5 nm, which provides a degree of intrinsic radiation hardness (see section 3.2.2).

The irradiation was performed by means of a 10 keV X-ray machine (manufactured by SEIFERT) at Imperial College, instead of the more usual ^{60}Co gamma-ray source. Research has shown that these two sources are equivalent0 The total dose received by the devices was 50 Mrad(SiO_2) at a rate of approximately 500 krad/h.

Typical bias operating conditions were applied to the devices during all stages of irradiation and annealing; approximately 500 mV across the gate. Following this the devices were annealed by placing them in an oven at a temperature of 100° C for a

period of one week. Static parameters and transistor noise were measured after each irradiation step and after annealing. The gate referred noise spectral density was measured for devices P1 and P6, under the conditions shown in Table 4.6.

Transistor	I_{DS} [μA]	V_{DS} [Volts]	V_{BS} [Volts]
P1 and P6	500	-0.5	0

Table 4.6 Bias condidions of transistors with noise measurements.

4.1.5 Harris Irradiation Results

A substantial volume of data was obtained. Since the main focus of attention in this thesis is the APV chip, results are presented for the transistors that most closely represent the input transistor of the APV pre-amplifier, demonstrating differences between processing runs. It was found that the differing channel dimensions did not affect the behaviour of the devices qualitatively. There are some small quantitative differences between the narrow and the wide gate devices. However, these differences are insignificant and would not affect conclusions drawn from one transistor size. Therefore only results for the 2000 μm wide devices are presented.

4.1.5.1 Transfer characteristics (I_{DS} vs V_{GS})

Figure 4.3 and Figure 4.4 show the dependence of $|I_{DS}|$ on V_{GS} for the widest PMOS and NMOS transistors respectively. The apparent dip in the current just before threshold, is due to the current magnitude being plotted on a logarithmic scale, requiring a negative current to appear positive. This is a pre-threshold leakage current, due to leakage through to the source from the bulk. The bulk-source junction is reverse biased and so its leakage current is opposite in polarity to the drain-source current.

One can see clear differences between the APV5 and subsequent APV6, MX and APV6c runs. A more significant negative shift along the V_{GS} axis is visible in the PMOS devices. There is also a significant change of sub-threshold current slope in the NMOS devices after irradiation, which is indicative of an increase in charge traps at the gate oxide - silicon interface.

The characteristics after annealing show a significant recovery in the APV6 and MX runs. The lack of original degradation in the APV5 devices leaves little room for improvement after annealing, with no observable adverse effects.

The effect of interface charge traps can be distinguished from that of traps in the oxide by analysing the sub-threshold slope in the I_{DS}-V_{GS} curve. The sub-threshold slope is measured in the linear region of operation of the transistor, where I_{DS} increases linearly with $|V_{GS}|$. Charge traps at the interface can cause a reduction in mobility (reducing $\partial I_{DS}/\partial V_{GS}$). When the reduction in mobility is high, the slope is visibly less steep. This effect is visible in the NMOS devices, but less so for the PMOS.

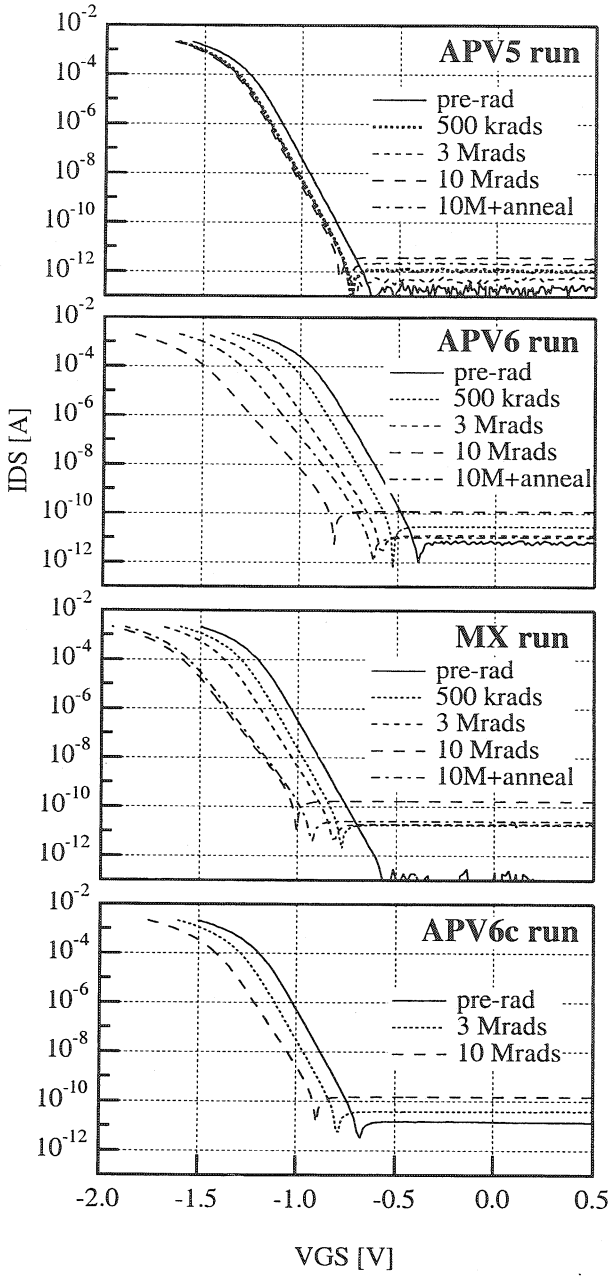

Figure 4.3 PMOS 2000/1.4 I_{DS} vs. V_{GS} characteristics up to 10 Mrads and after annealing.

An important point to note here is the inconsistency between runs from the same manufacturer. For a system such as the CMS tracker, which will contain ~ 100,000 APV chips, one must be confident that the radiation tolerance is uniform across all processing runs.

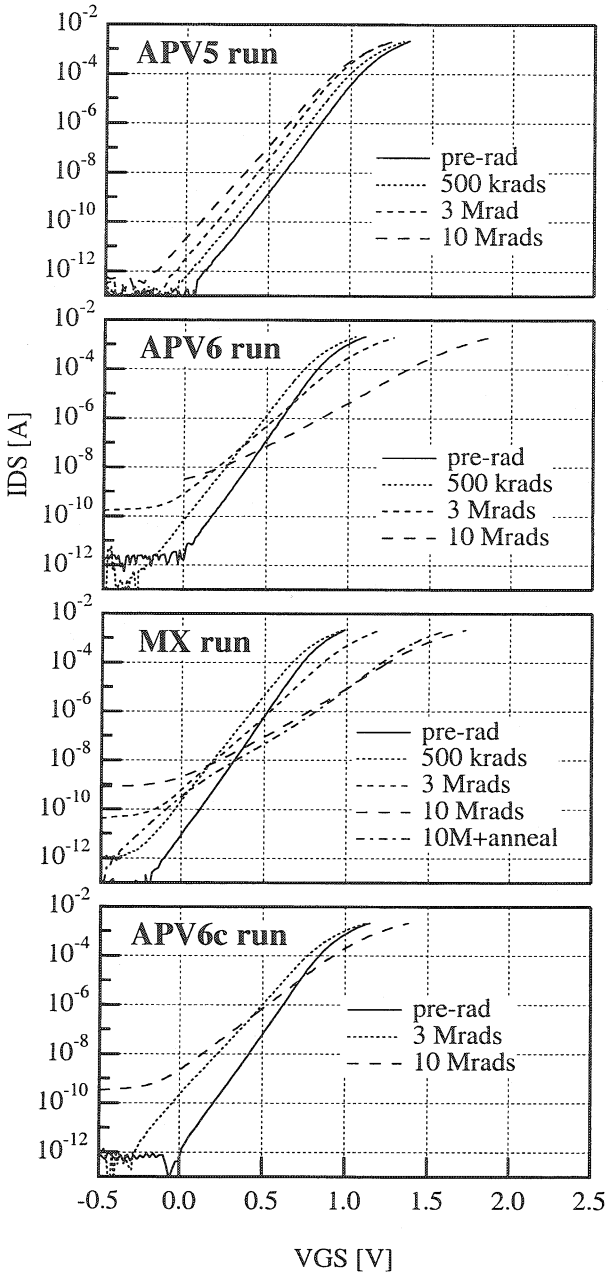

Figure 4.4 NMOS 2000/1.2 I_{DS} vs. V_{GS} characteristics up to 10 Mrads and after annealing.

4.1.5.2 Transconductance

Table 4.7 and Table 4.8 show the transconductance g_m (determined at $|I_{DS}|=500\mu A$) dependence on dose for the PMOS 2000/1.4 and NMOS 2000/1.2 devices respectively.

In order to indicate relative changes after irradiation, data are compared to the pre-irradiation g_m value. For the PMOS devices in Table 4.7 the APV6 and MX runs show a reduction in g_m of approximately 20% after 10 Mrads, compared with less than 5% for the APV5 run device. The APV6c run shows better performance than the APV6 and MX runs with a reduction of 10%. Annealing after 10 Mrads had no significant effect on the APV5 run device but a greater beneficial effect on the APV6 and MX run devices, where about half the loss of g_m prior to anneal is recovered.

run	3 Mrads	3 Mrads + anneal	10 Mrads	10 Mrads + anneal
APV5	-5	-5	-5	-5
APV6	-10	-5	-20	-10
MX	-10	-5	-17	-10
APV6c	-5	-	-10	-

Table 4.7 PMOS 2000/1.4 g_m reductions after irradiation, expressed as a percentage change from pre-irradiation values.

run	3 Mrads	3 Mrads + anneal	10 Mrads	10 Mrads + anneal
APV5	-16	-15	-19	-
APV6	-30	-24	-47	-
MX	-30	-22	-48	-36
APV6c	-14	-	-34	-

Table 4.8 NMOS 2000/1.2 g_m reductions after irradiation, expressed as a percentage change from pre-irradiation values.

The NMOS device g_m dependence on dose in Table 4.8 shows a substantial reduction in g_m of almost 50% after 10 Mrads for the APV6 and MX run devices compared with 20% for the APV5 run device and about 35% for the APV6c run device. The APV5 and MX run devices were mistakenly biased incorrectly during annealing and so these results are omitted.

4.1.5.3 Threshold voltages

Figure 4.5 and Figure 4.6 show plots of threshold voltage dependence on dose for the same PMOS and NMOS devices. The results after annealing have, for clarity, been indicated by offsetting the point to the right, which does not mean that the device has been irradiated further.

Both interface traps and oxide trapped charge cause a negative threshold voltage shift in PMOS transistors. However, in NMOS oxide traps cause a negative threshold voltage shift, whereas interface traps cause a positive threshold voltage shift. The two types of traps compensate and the dominant type will cause a shift in the relevant direction (Cross-reference to theory Section).

For the APV5 run the threshold voltage shifted by about 100mV for the PMOS and 200mV for the NMOS after 10 Mrads, most of the shift occurring before the first 1 Mrad.

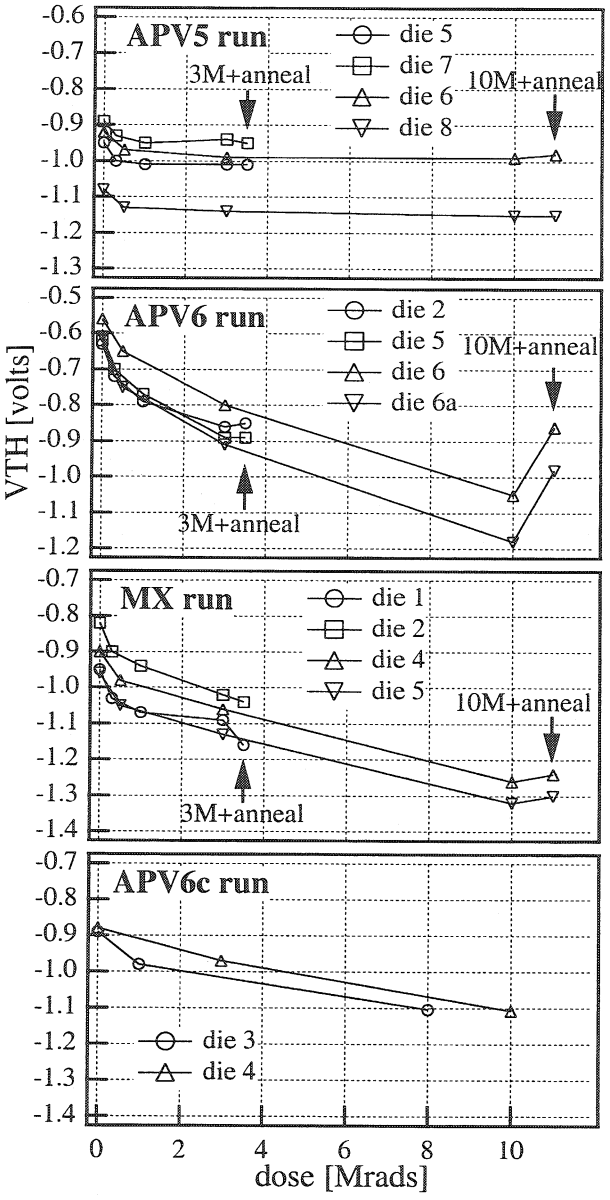

Figure 4.5 PMOS 2000/1.4 threshold voltages (V_{BS}=0).

The APV6 and MX run devices show much greater shifts of up to 600 mV, for both PMOS and NMOS, after 10 Mrads, which build up gradually with dose. A shift of this amount is dangerously high. Although it is hard to quantify the exact amount of acceptable shift, anything greater than 600mV will start to cause operational errors in most chips. The APV6c run devices show a negative shift of 200 mV after 10 Mrads for PMOS and a positive shift of 50 mV for NMOS. The initial shifts with dose, for the NMOS devices, are negative, but subsequently rebound (see section 3.2.2) occurs for the APV6, MX and APV6c runs, leading to an overall positive shift after 10 Mrads.

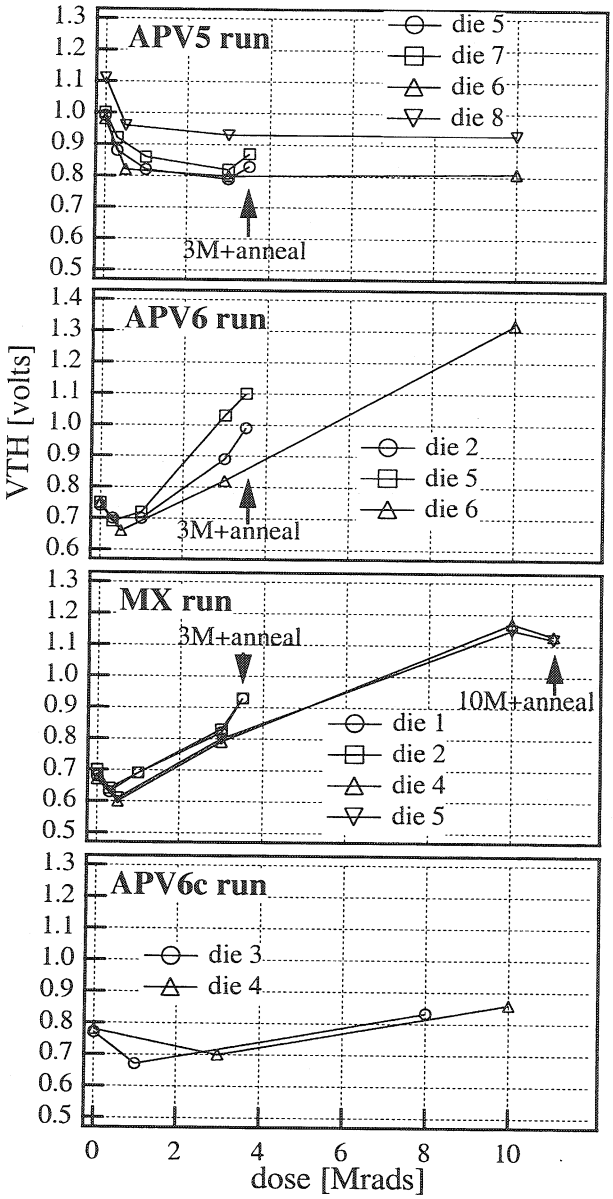

Figure 4.6 NMOS 2000/1.2 threshold voltages (V_{BS}=0).

The APV5 run exhibits a lack of rebound. This threshold voltage behaviour, like the subthreshold current slope behaviour in section 4.1.5.1, indicates a build up of interface charge traps for the APV6, MX and APV6c runs, and conversely a build up of oxide trapped charge, up to 10Mrads, for the APV5.

The PMOS devices in Figure 4.7 show some recovery, after annealing, for the devices irradiated to 10 Mrads, particularly for the APV6 run devices.

4.1.5.4 Noise Spectra Measurements

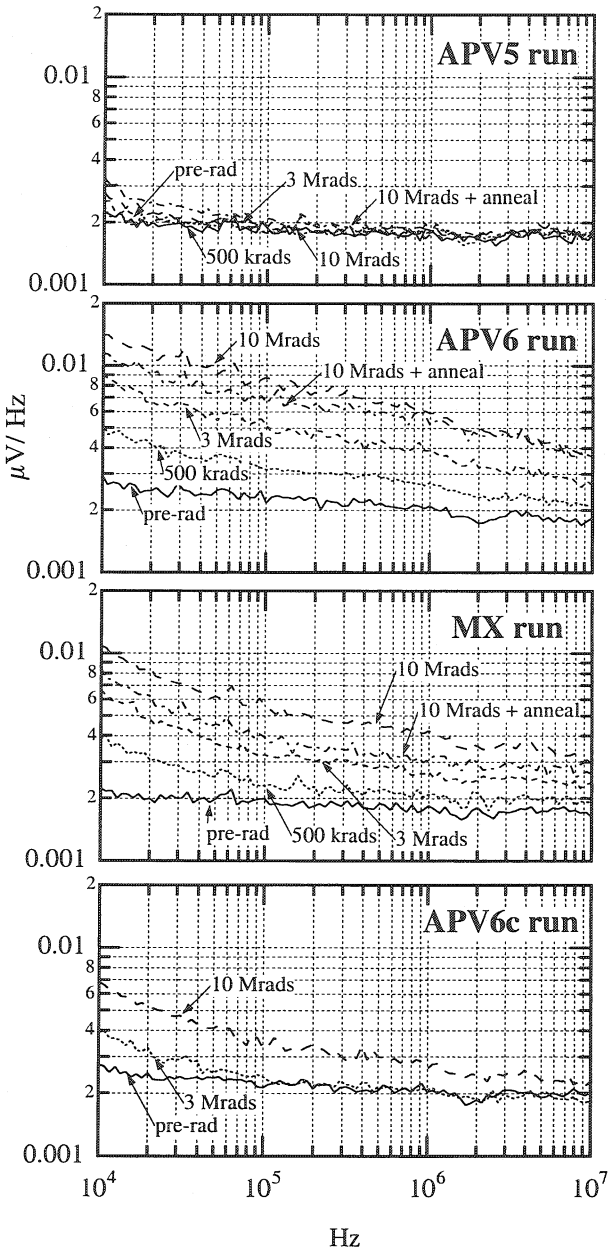

Figure 4.7. PMOS noise spectral density measurements (I_{DS}=-500μA, V_{BS}=2V).

The annealing results after 10 Mrads are missing for the NMOS devices in Figure 4.6, but examination of the results for the 10/10 devices shows no significant recovery for the APV5 run devices and results very similar to the MX run for the APV6 run devices[4.1], where a small recovery is apparent. Some improvement after annealing is observed but not at a level which would allow the APV to remain fully functional after 10 Mrads.

Figure 4.7 shows typical noise spectral density measurements for the same PMOS transistor. PMOS device noise performance is more critical for the APV chip, since the pre-amplifier input transistors are PMOS. One can again see the radiation hardness of the APV5 run. The APV6 and MX run devices both show a significant increase in the noise after irradiation, but less so for the MX run. At 3 MHz (the central frequency of the APV amplifier filter) an increase from around 2 nV/√Hz to about 4.5 for the APV6, and to 3.5 nV/√Hz for the MX run.

The APV6c run device shows a less significant increase from 2 to 2.5 nV/√Hz at this frequency. For the APV6 and MX run devices some reduction of the noise can be seen after annealing. Again the overall consideration is the lack of consistency between different runs, but also the very bad impact of the noise increase on the signal to noise ratio of the final system. It is clear that there is a serious problem with the radiation hardness of these transistors, therefore other technologies had to be investigated. Work was underway on test structures and an APV chip designed in DMILL. However, it was still uncertain whether it would be a suitable process. Soon after the Harris results were published the first data from the 0.25 μm process had emerged and it was known that there was a good chance it would exhibit much better tolerance to radiation because of its thin oxide. Work commenced straight away on test structures from the new process in the search for a more suitable technology.

4.1.6 Results of 0.25μm Irradiations

As with the Harris test structures, study of the 0.25μm test structures yielded a large set of results from all the irradiated devices. The following set of results is representative of the most relevant devices in the context of the APV25 chip.

4.1.6.1 Transfer characteristics (I_{DS} vs V_{GS})

Figure 4.8 and Figure 4.9 show the dependence of $|I_{DS}|$ on V_{GS} for two representative PMOS transistors, one from each foundry for comparison. One can see that there is much less degradation in the behavior of both when compared to the Harris devices. A fairly modest shift along the V_{GS} axis is observed after 50 Mrads, with a slight improvement after annealing.

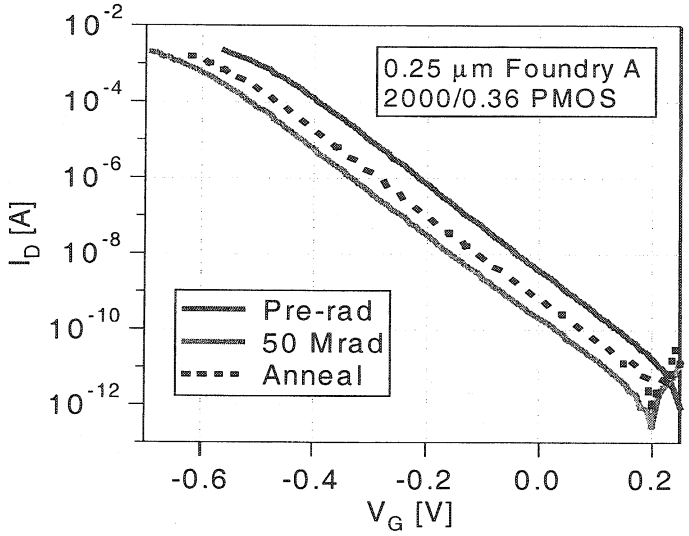

Figure 4.8 Foundry A PMOS (2000/0.36) I_{DS} vs. V_{GS} characteristics up to 50 Mrads and after annealing.

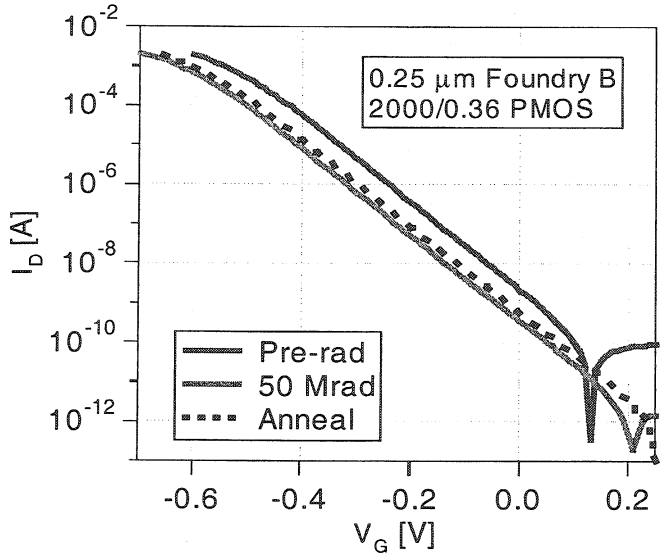

Figure 4.9 Foundry B PMOS (2000/0.36) I_{DS} vs. V_{GS} characteristics up to 50 Mrads and after annealing.

The shift is very slightly less for the device from foundry B, but the difference is relatively small, and on the whole a good agreement between the two devices can be

seen. Both devices exhibit no discernible change in sub-threshold slope. Figure 4.10 shows the same transfer characteristics for the NMOS device from Foundry A. Here a slight decrease is visible in the gradient of the sub-threshold current slope due to creation of interface charge traps, but it is clearly less significant than for the Harris devices. Very little shift along the V_G axis is evident from these results.

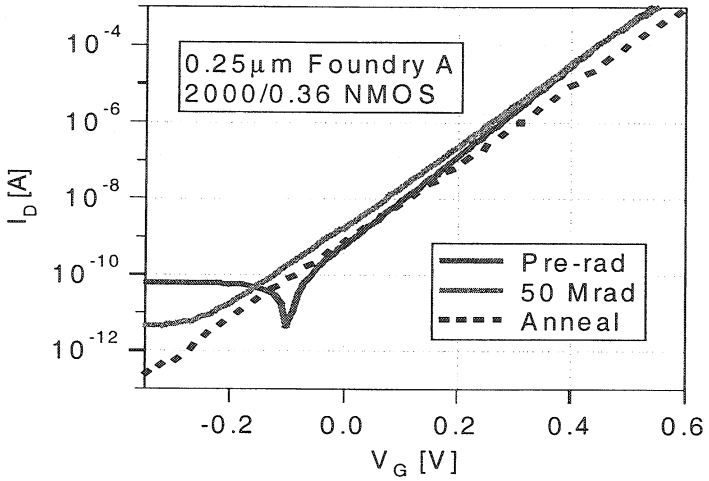

Figure 4.10 Foundry A NMOS (2000/0.36) I_{DS} vs. V_{GS} characteristics up to 10 Mrads and after annealing.

4.1.6.2 Transconductance

The lack of a significant change in the current slope, for PMOS transistors (Figure 4.8 and Figure 4.9), would suggest that very few interface states are generated, and that there is no significant change in transconductance.

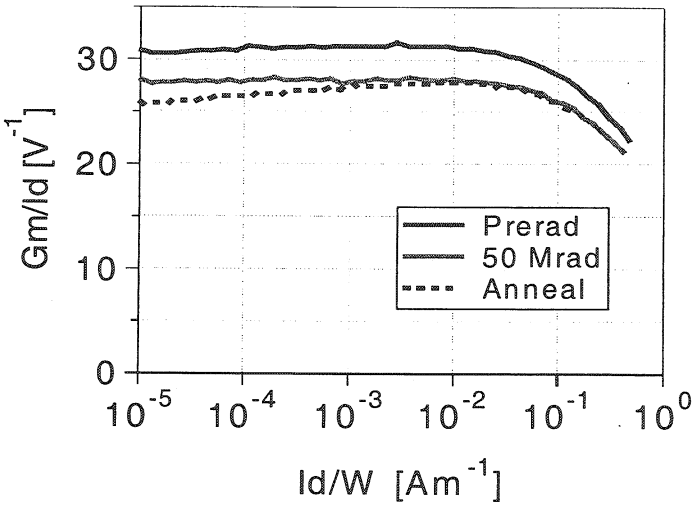

Figure 4.11 Transconductance of NMOS after irradiation and annealing.

However, a reduction in the current slope of the NMOS device is visible after irradiation and after annealing (Figure 4.10). This is a clear indication that the irradiation process has generated interface states, which cause a reduction in the mobility of charge carriers within the thin conducting channel beneath the oxide. The reduction in mobility can also be seen in a plot of the normalised transconductance (g_m/I_{DS}), (Figure 4.11). g_m/I_{DS} decreases by 9% after 50 Mrad and by 15 % after annealing in the weak inversion region.

4.1.6.3 Threshold Voltage Shift

Figure 4.12 and Figure 4.13 show threshold voltage dependence on dose for irradiated PMOS devices from both foundries. As before, the results after annealing have been indicated by offsetting the point just to the right.

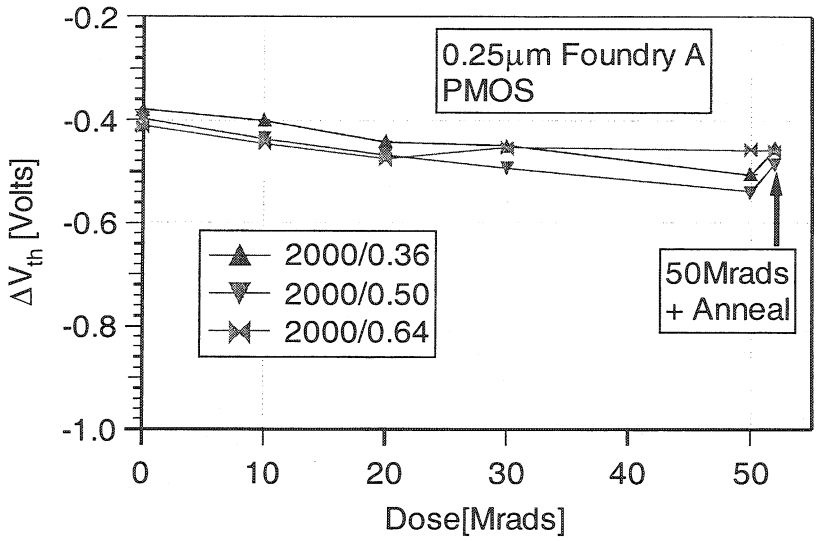

Figure 4.12 Threshold voltage shifts for Foundry A PMOS.

The maximum shift after a dose of 50 Mrad(SiO_2) is around 120 mV for foundry A and 80 mV for foundry B for transistors of comparable dimensions (W/L=2000/0.36). All shifts are again negative as expected. After annealing, most transistors show some partial recovery. Shifts of this magnitude are well within acceptable limits; although it is again hard to predict the exact limits for shifts in the APV25, one would expect the chip to remain functional up to shifts of 500 mV.

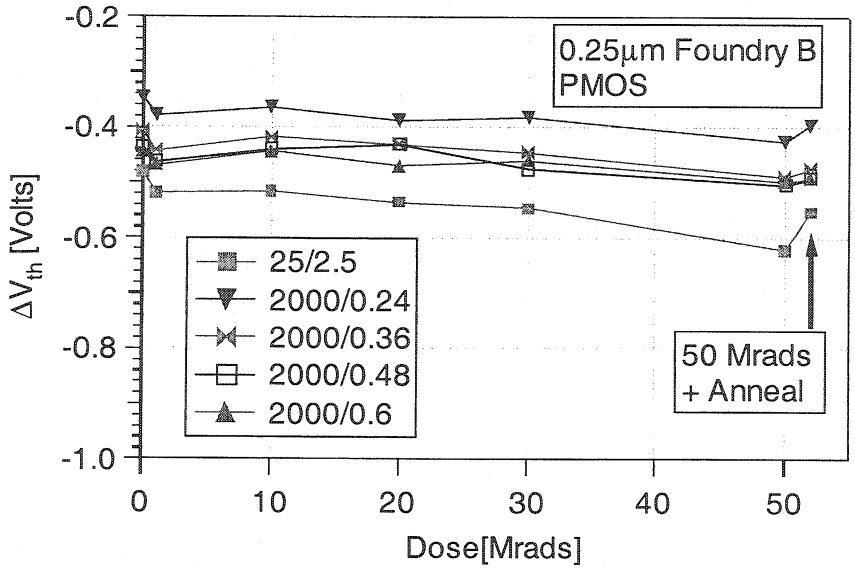

Figure 4.13 Threshold voltage shifts for Foundry B PMOS.

Figure 4.14 shows the threshold voltage shift for the only irradiated NMOS device. Very little shift is evident up to 50 Mrads, with a small but noticeable shift of about 30mV after annealing. This positive shift indicates a build up of interface charge traps after annealing, which could be due to the improved transport of oxide trapped charge towards the interface under high temperature conditions.

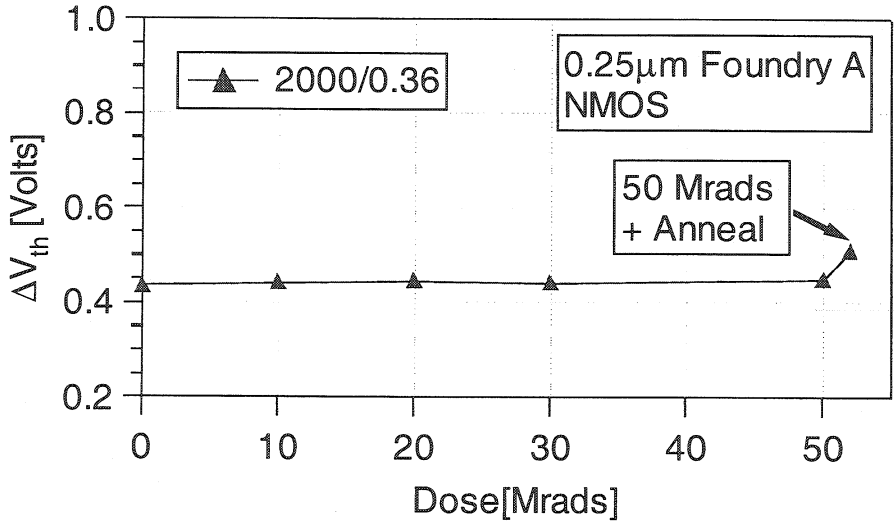

Figure 4.14 Threshold voltage shifts for Foundry A NMOS.

4.1.6.4 Noise in PMOS transistors

Figure 4.15 and Figure 4.16 show the noise spectral density in the PMOS devices from both foundries. Transistors from foundry A show a lower noise level at low frequency. However, there is no significant difference between the two foundries at the

central operating frequency of the APV amplifier filter (3 MHz). An increase in the noise of no more than 0.2 nV/√Hz is evident after 50 Mrads in both foundries.

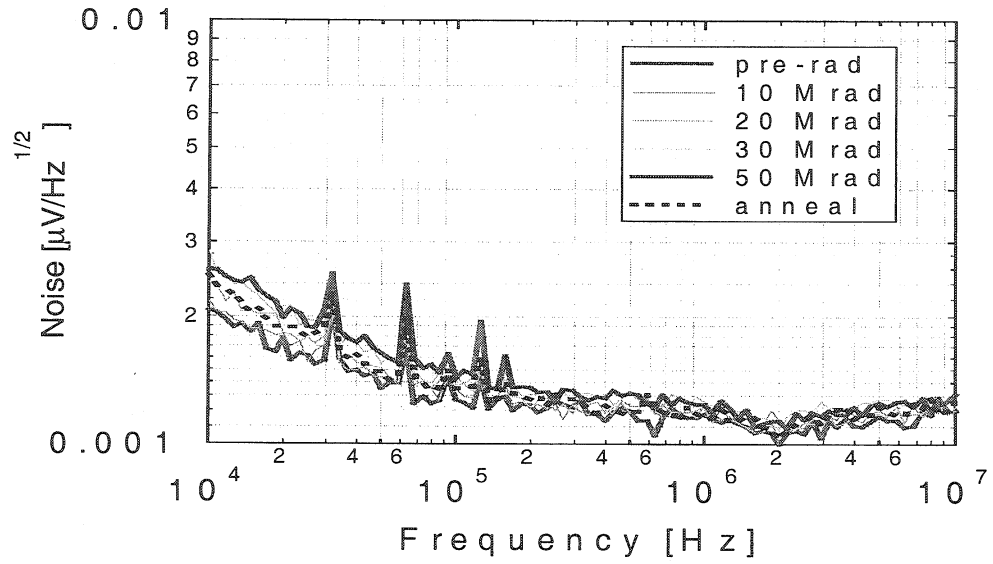

Figure 4.15 Noise in Foundry A PMOS.

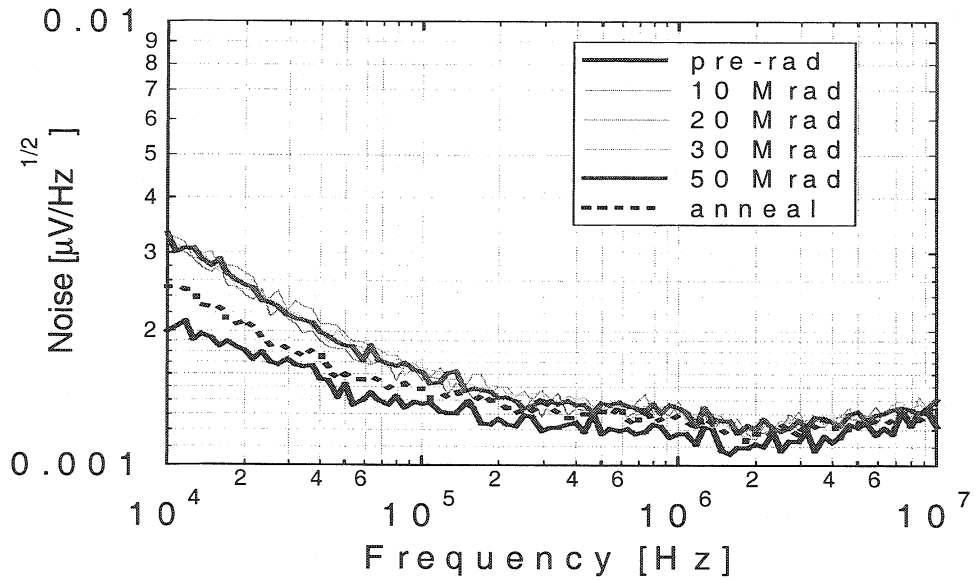

Figure 4.16 Noise in Foundry B PMOS.

4.1.6.5 Noise in the NMOS transistor

Although NMOS transistors are not used as input transistors in the amplifying stage, it is interesting to include some NMOS transistor noise measurements. A comparison between PMOS and NMOS transistors confirms that the NMOS transistor has a much higher spectral noise density. In the case of the NMOS transistor (Figure 4.17), the white noise is higher after irradiation than it is for PMOS. This confirms the fact that a

significant number of interface states are being generated, since white noise is related to the random fluctuations of carriers in the channel which is adjacent to the interface.

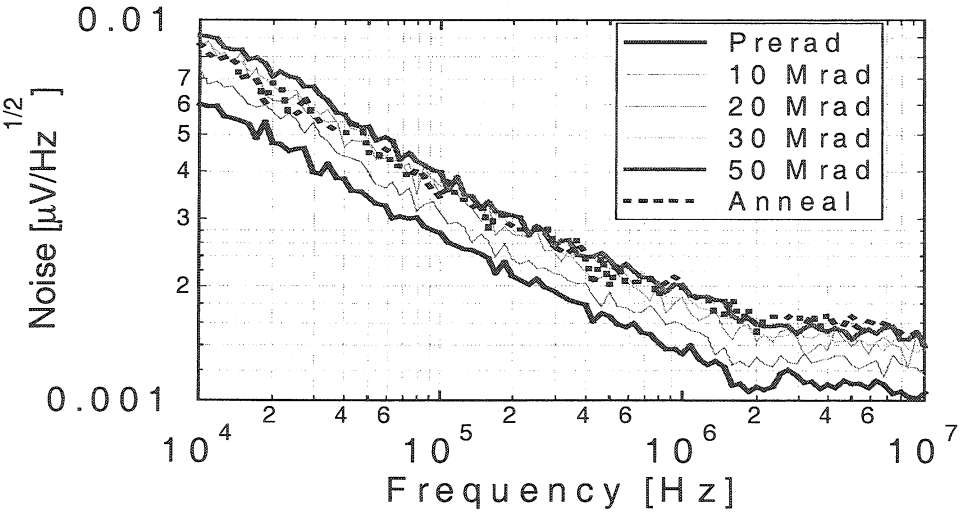

Figure 4.17 Noise in Foundry A NMOS.

4.1.7 Comparison between Harris and 0.25μm Technologies

Clear differences have been observed between the behaviors of the two technologies under the effects of irradiation.

Technology	Dose	Transistor type	Max g_m decrease	Threshold voltage shift	Noise increase
Harris	10Mrads	PMOS	20%	-600mV	2.5 nV/√Hz
		NMOS	48%	600mV	-
0.25 mm	50Mrads	PMOS	0%	-120mV	0.2 nV/√Hz
		NMOS	9%	~0mV	-

Table 4.9 Summary of results of irradiation effects on both Harris and 0.25μm technologies.

Table 4.9 Shows that the 0.25m technology displays a clear improvement in the achievable radiation tolerance. Most of its success is attributable to the remarkably thin gate oxide, which makes it a more intrinsically radiation hard technology.

4.1.8 APV25 X-ray Irradiation Results

Four APV25 chips were irradiated with the X-ray source up to 20 Mrad (SiO_2). All four show very good radiation tolerance. The results obtained on one of the chips are shown below.

The behaviour of the APV25 up to 20 Mrad(SiO_2) is illustrated in Figure 4.18 to Figure 4.20. The pulse shapes do not change significantly after irradiation for a wide range of input signals. The reduction in gain is well below 10 % for the highest input signal and the similarity of the pulse shapes between pre- and post-irradiation measurements is convincing demonstration of the tolerance of the process to doses comparable with those of the LHC.

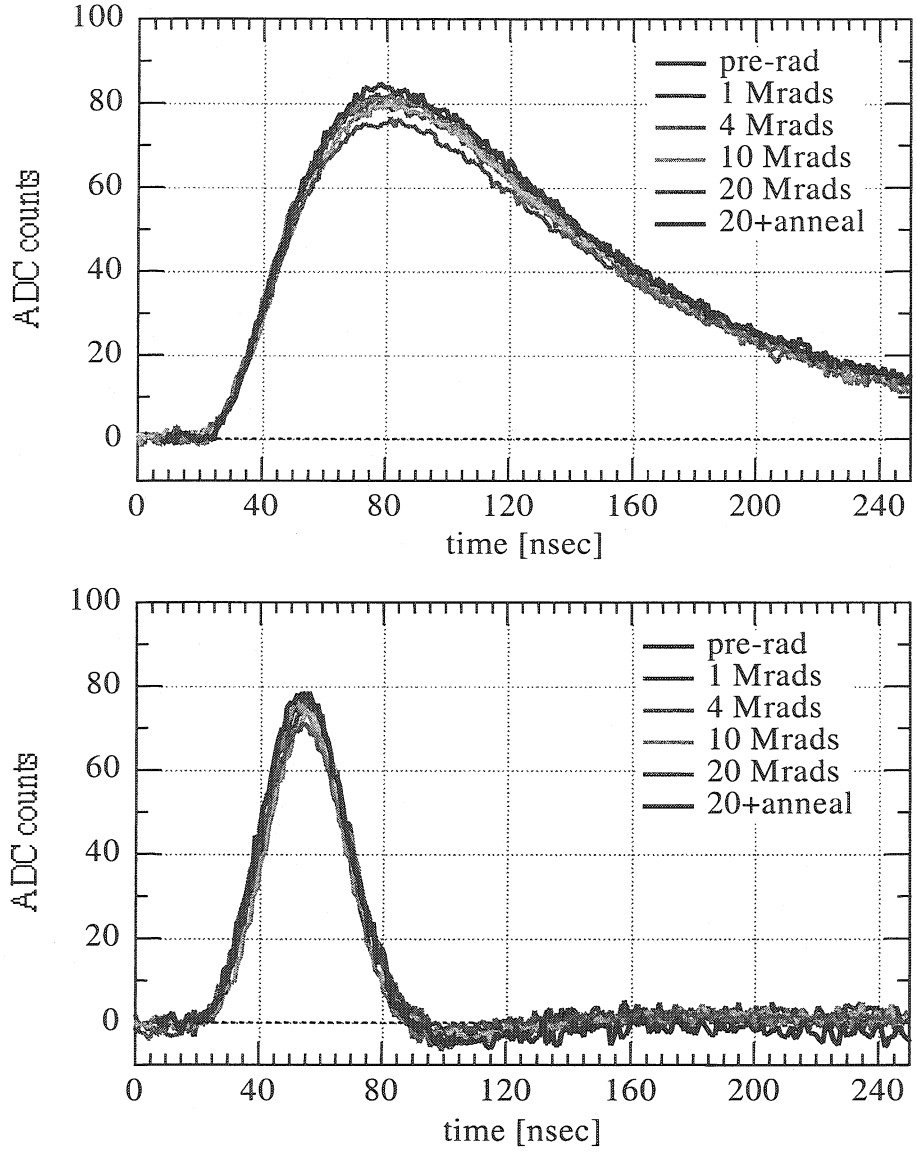

Figure 4.18 APV25 pulse shapes before and after irradiation.

There is no significant difference in the noise behaviour of the chip after irradiation. The three peaks in Figure 4.20 correspond to the three channels which were bonded out for testing and which therefore have extra capacitance associated with them.

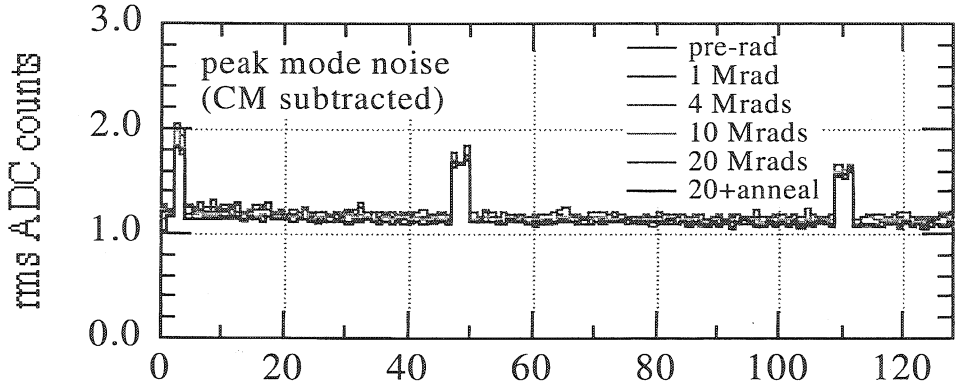

Figure 4.19 Noise in peak mode before and after irradiation, CM subtracted.

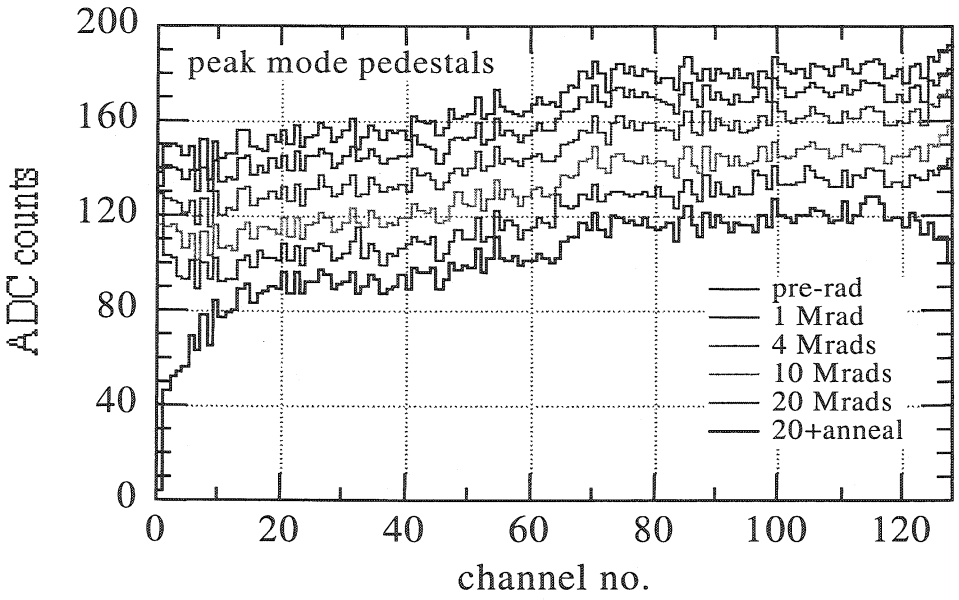

Figure 4.20 Peak mode pedestals before and after irradiation. The offset between irradiations has been artificially introduced to separate the different pedestal sets for clarity.

The uniformity of the pedestals over the different steps of irradiation is extremely good. The offset between the irradiations has been artificially included to separate the pedestal sets for clarity. There were no adverse affects after irradiation apart from a

slight droop in the pedestals at either end after annealing. This effect is under investigation, but is not thought to be of any concern for CMS.

The overall performance of the APV25 after irradiation is very impressive and well within the requirements of CMS.

4.1.9 Gate Rupture in the APV6 and APV25

During tests designed to assess SEU susceptibility, 4 APV25 chips were subjected to high fluences from particles with varying LETs. This is the first evaluation of the susceptibility of the APV25 to SEGR.

4.1.9.1 APV25 Results

Table 4.10 summarises the fluences to which the different chips were exposed. The total fluence is given as well as the fluence for the highest LET irradiation (61.8 $MeVcm^2mg^{-1}$).

	High LET Fluence [cm^{-2}]	Total Fluence [cm^{-2}]
Pipeline Logic	6.02E+07	7.65E+09
FIFO and I2C registers	9.42E+07	8.53E+09
Control Logic	2.17E+08	3.33E+09
Whole Chip	9.03E+07	2.82E+09

Table 4.10 Fluences received by the four APV25 chips. The high LET fluence column indicates values reached for the iodine irradiation (LET = 61.8 $MeVcm^2mg^{-1}$).

The probability for an SEGR to occur is highest for the highest LET irradiations. Taking the pipeline logic as an example, the approximate gate area is 5×10^{-3} cm^2 and the iodine fluence was 6.02×10^7 cm^{-2}. Therefore the total number of ions incident on the gate area was 3×10^5. During this test no device failures were observed and hence it can be concluded that no detectable SEGRs occurred. This first test is by no means exhaustive and simply sets a lower limit to the mean time to failure. Further investigation into the susceptibility of the APV25 to SEGR will continue.

The highest LETs expected in the tracker should not exceed 15 $MeVcm^2mg^{-1}$. This, combined with the low electric field across APV25 oxides (4.5 MV/cm), makes it very unlikely that SEGR will pose a problem for CMS electronics.

4.2 Single Event Upset Testing

SEU testing has been performed on both the APV6 and the APV25; the experimental method for both tests was virtually identical. The experimental apparatus and method are described in this section along with the changes and additions that were made between the first and second tests.

4.2.1 Hardware

Before setting up the system in the beam area a preparatory system was set up in the lab at IC. The system reflected all the requirements of the final system, including long cables between the beam area and the barracks, and vacuum prepared interface cards and cables. The system was tested fully, making data taking runs as if in the beam. Control was performed by a PC running LabVIEW, which communicated with the VME crate via a PCI VME interface. The trigger sequence for the APV was provided by a SEQSI sequencer, and control of the APV performed by a VI^2C slow control interface. The output from the APV was digitized by a flash ADC and subsequently read out by the PC software.

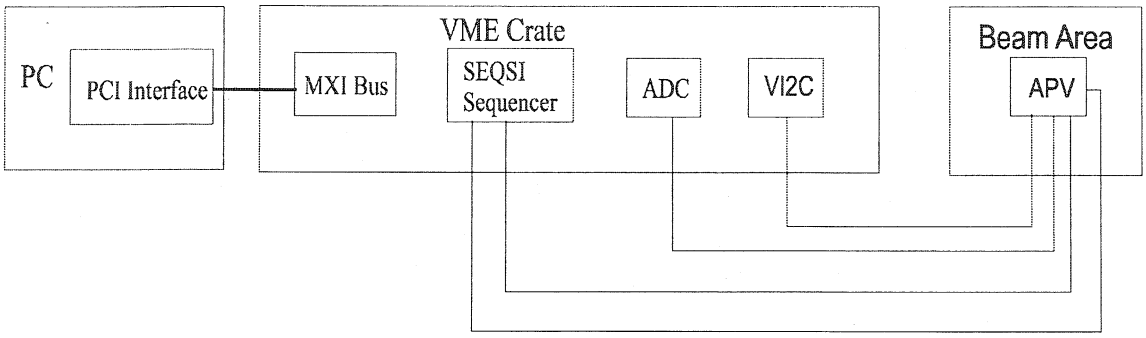

Figure 4.21 Schematic representation of the hardware.

Figure 4.21 is a schematic representation of the experimental setup. In the APV6 test there was one APV on the test board; this number was increased to four in the APV25 test.

Below is a photograph of the APV6 test board. The copper masks were aligned, under a microscope, to expose all the digital parts of the chip except for the control logic. Section 3.5.1 gives a description of the parts of the chip under test. The test board for the APV25 is shown in Figure 4.23. Four chips are mounted in a line, each with a carefully engineered mask. Both of the test boards were prepared for operation

in a high vacuum. In order to achieve this the materials that were used had to be capable of withstanding vacuums in excess of 10^{-7} mbar, with little or no out-gassing.

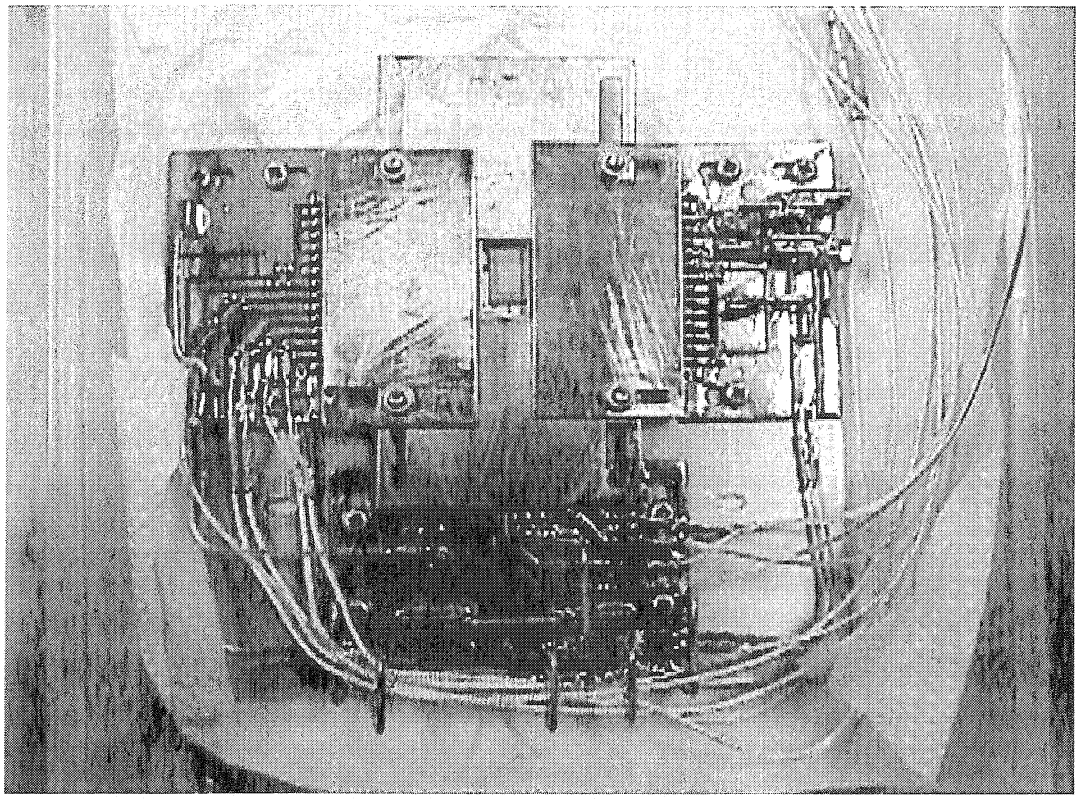

Figure 4.22 APV6 SEU test board containing interface electronics, one APV6 and a set of simple copper masks.

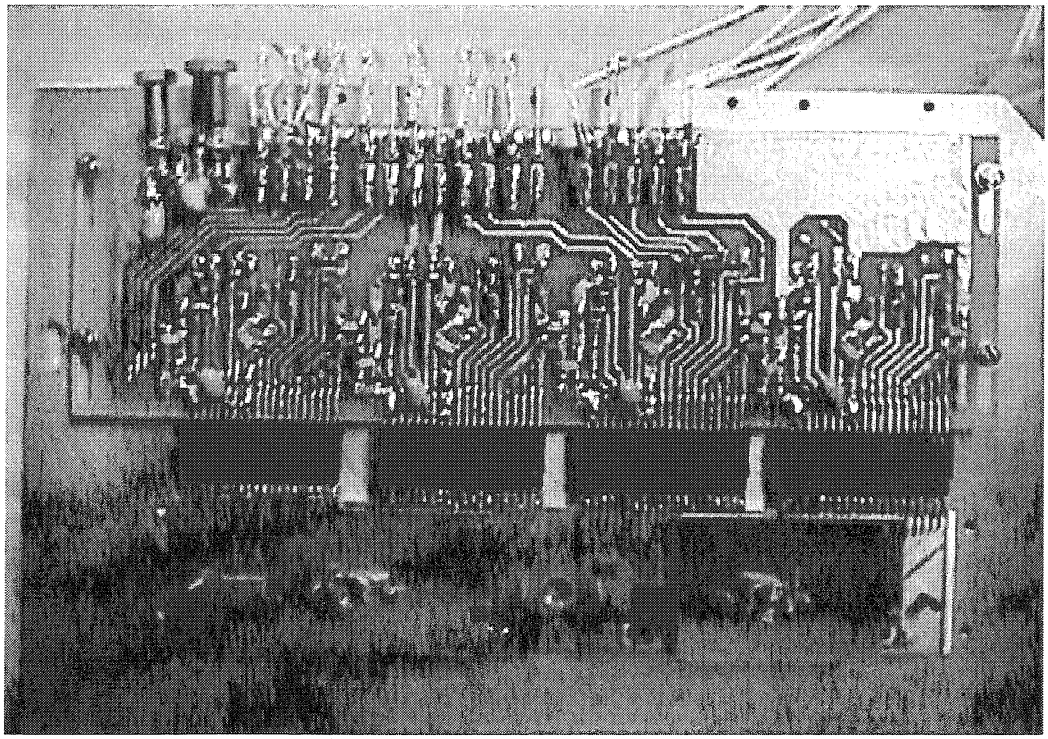

Figure 4.23 APV25 SEU test board containing interface electronics, four APV25s and precision machined masks.

4.2.2 Masking APV sections

One of the requirements was the ability to mask off sections of the APV. The masks for the APV6 test were made from copper plates approximately 1mm thick to provide an adjustable aperture, while the masks for the later APV25 test were precision engineered to expose specific predetermined areas of the chip. With four chips in the beam at the same time, it was possible to have four different masks, which enabled selection of the parts of the chip to test without the necessity to break the vacuum.

The heavy ions used in the tests are easily stopped by a thin sheet of metal. Table 4.11 and Table 4.12 give the range of these ions in copper and aluminium respectively. The range depends on the ion species and kinetic energy, with a maximum of 127 μm for lithium. The copper masks were ~500 μm thick and the aluminium masks were ~1 mm.

4.2.3 Software

The control and data acquisition was carried out by custom designed software developed in LabVIEW. The main tasks were to provide resets and triggers, via the SEQSI, to capture the digitised APV output data frame, via the ADC, to perform rudimentary on-line analysis and to save data to disk. The on-line counting of events was designed to provide visual real-time proof that SEUs were occurring and that the data were worth taking. It also enabled fine-tuning of the sensitive time in order to ensure that event counting was non-saturated (see section 4.2.5.1). Other on-line information included an overall count of upsets and a total sensitive elapsed time counter. The software also included I^2C control for testing the APV static registers.

The second edition of the software for the APV25 test was virtually identical with the addition of a hard reset for testing the control logic.

4.2.4 The TANDEM Van der Graaf heavy ion accelerator

The TANDEM is located at the INFN laboratory in Legnaro. Table 4.11 and Table 4.12 show the full range of surface LETs that were chosen for these tests along with the ion species, corresponding beam energy and range in the mask material.

Ion	A	LET [MeVcm^{2}mg^{-1}]	Energy [MeV]	Range in Copper [μm]
Li	3	0.38	55	127
C	6	1.5	92	57
O	8	2.9	106	38
F	9	3.8	116	33
Si	14	8.7	153	22
Ni	28	28.7	214	13
Ag	47	54.9	259	11

Table 4.11 Ions used in APV6 test.

Ion	A	LET [MeVcm^{2}mg^{-1}]	Energy [MeV]	Range in Aluminium [μm]
Si	14	9	145	50
		10.4	100	32
Cl	17	12.9	160	49
		14	130	32
		15	107	26
		16	87	21
Ti	22	20.4	178	32
		23.2	115	21
Ni	28	28	237	33
		29	200	28
		31.6	138	20
Br	35	39.4	100	15
I	53	61.8	250	25

Table 4.12 Ions used in APV25 test.

The typical beam flux was in the range 10^4–10^7 cm^{-2} s^{-1} and maintained at a stable value as far as possible. Due to accelerator problems during the experiment it was only possible to measure the APV6 SEU cross-section for two values of LET: oxygen, at a typical flux of $2\text{x}10^6$ cm^{-2} s^{-1} and silicon, at a flux of $2\text{x}10^5$ cm^{-2} s^{-1}. However, APV25 SEU cross-sections were measured for all the values of LET shown above.

4.2.5 Measuring errors in the APV

In the event of an upset in the pipeline logic or FIFO, there are two possible outcomes: either the error bit in the output data frame is set, or the pipeline address in the output data frame is incorrect. The error bit is set if an upset in the pointer logic causes the latency of the trigger pointer to change. Only a tiny proportion of upsets produce both outcomes simultaneously.

One can also test for events in the I^2C registers by writing defined values, reading out the values after a set period of time, and comparing them with initial values. In this case it is possible to detect individual cells, which have been upset. Upsets in the control logic block can have a more drastic effect, such as incorrect pipeline addresses, loss of digital header, or loss of entire data frame. When measuring upsets in the pipeline and FIFO, one only needs to reset the chip using a soft 101 reset. However, when measuring upsets in the control logic one requires a hard power-on reset, since upsets in the control logic can cause the chip to lock up and subsequent data frames all exhibit errors, caused by the original upset, thus making it impossible to distinguish further upsets. A power-on reset recovers the operation of the control logic. Following this one must apply a soft reset and then readout.

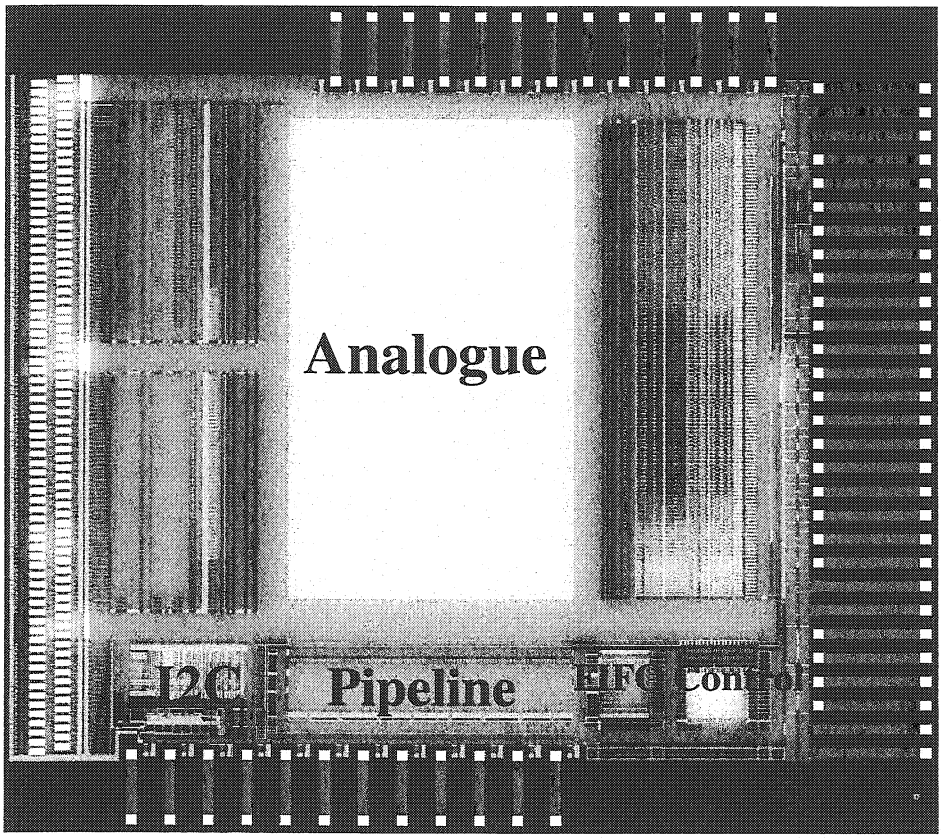

Figure 4.24 Location of tested circuits.

It is important that the average number of upsets per time interval is less than one so one can be confident that when an error is detected within one sensitive time interval it was in fact only one error and not more; this prevents undercounting. The running time per measurement, when the chip is sensitive to upsets, is defined as the sensitive time (ST). When measuring the upset rate there is a readout period following each ST, which is repeated many times to measure a reasonable number of upsets. For the

pipeline, FIFO and control logic a typical run consists of 100,000 STs, and for the I^2C, 100. In the case of the I^2C test, instead of a reset and trigger, one writes a simple 8-bit pattern followed by the ST and then a read. The pattern can be varied to establish the cross-section for both 1 to 0 and 0 to 1 transitions. Figure 4.25 shows the definition of ST, in the case of the control logic sequence, the hard reset occurs just prior to the soft reset and ST is measured from the hard reset. Equation[4.1] gives the definition of $T_{sensitive}$ the total sensitive time for one run.

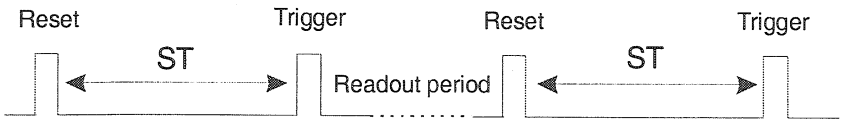

Figure 4.25 Definition of ST for the pipeline logic.

$$T_{sensitive} = N_{triggers} . ST \qquad [4.1]$$

4.2.5.1 Non-saturated measurements of events in the pipeline logic

One important issue is undercounting of events in the pipeline logic. Only one event can be detected within each sensitive time interval, as more than one error still only produces symptoms consistent with a single error. If we lose events by undercounting then the measured upset rates appear to saturate, hence one must ensure that the number of events is small in comparison with the number of sensitive time intervals. In order to achieve this condition one can adjust the length of the sensitive time interval (the achievable range for ST was from 10 to 1500 μs), or the flux of the ion beam, but this is not as trivial.

4.3 APV6 SEU Results

Measurements were made of the SEU cross-section of both the pipeline and I^2C logic at two values of LET using oxygen and silicon ions. For oxygen a large number of individual data taking runs with high fluence were made and no upsets were observed. Thus it can be deduced that the threshold for both pipeline and I^2C logic is above 2.9 $MeV.cm^2.mg^{-1}$. Silicon, on the other hand, produced a significant upset rate, giving the following results.

4.3.1 Error distribution

In order to be confident that events were genuine SEUs, their time distribution was measured; for random processes the distribution should be Poisson.

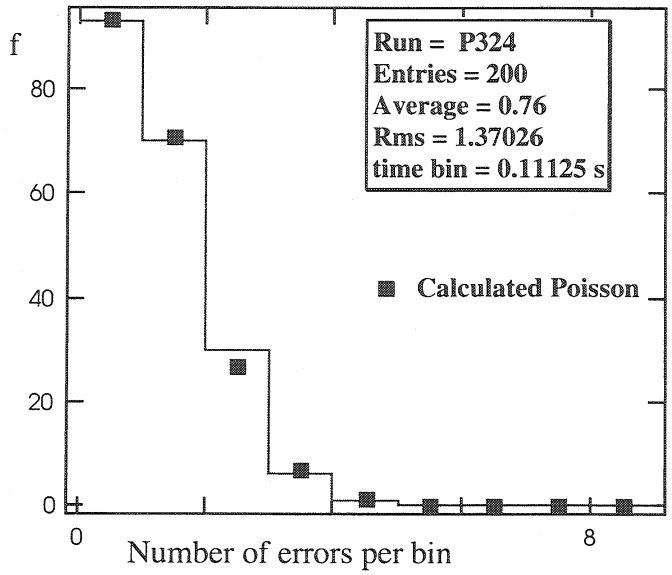

Figure 4.26 SEU distribution for pipeline logic.

Figure 4.26 shows a typical distribution of 200 SEUs observed in the pipeline logic. All data were checked in this way and showed good agreement with Poisson statistics.

4.3.2 Pipeline cross-section

These data represent errors in the pipeline logic only, which are identified as having caused the error bit to be set in the APV6 data header. All results that follow were obtained with a ratio of SEUs to number of sensitive time intervals ~ 0.001. Under these conditions one would expect the upset cross-section to show no dependence on the sensitive time.

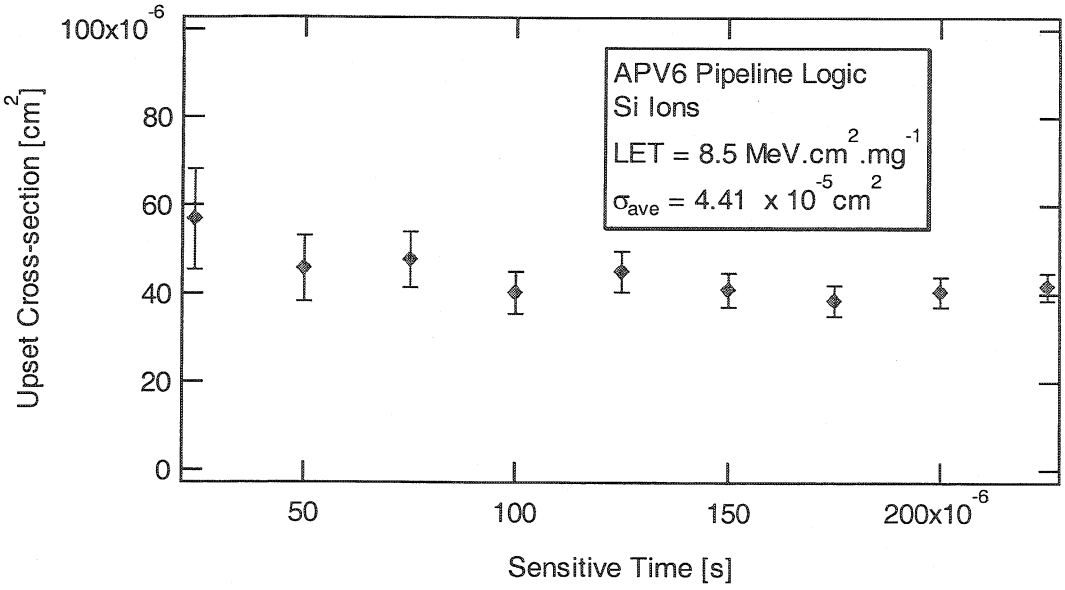

Figure 4.27 Variation of cross-section with sensitive time.

Figure 4.27 confirms this expectation, exhibiting no significant variation, with an average cross-section of $(4.4 \pm 0.2) \times 10^{-5}$ cm^2. All errors plotted are statistical errors taken as ± one standard deviation, which is calculated from Poisson statistics as $N^{1/2}$, where N is the number of upsets.

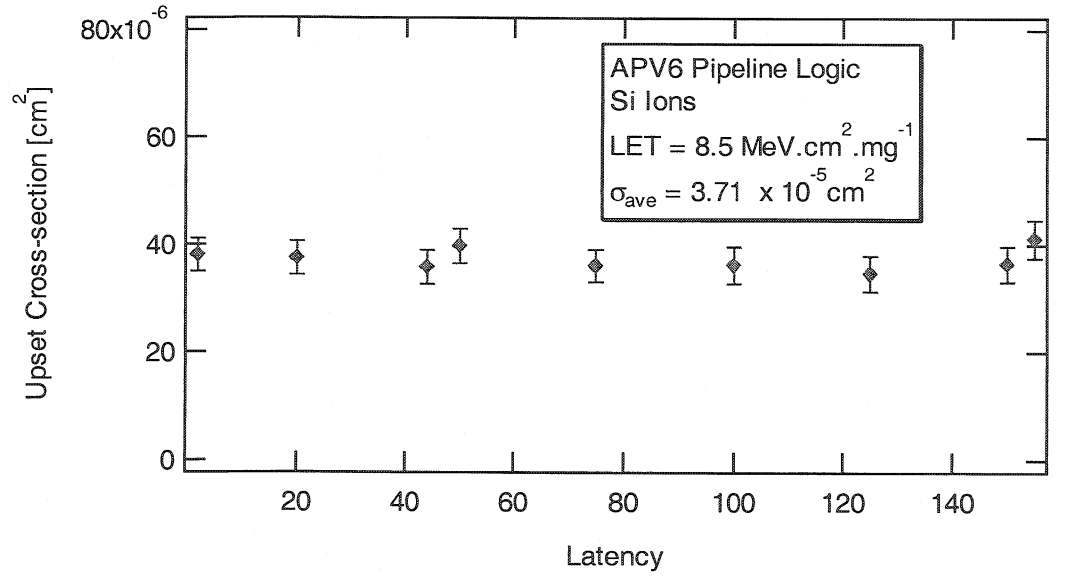

Figure 4.28 Variation of cross-section with latency.

A similar test was made by varying the latency of the trigger pointer. Again one would expect to see no major variation in the cross-section, and Figure 4.28 shows that the measured value varies very little, with an average of $(3.7 \pm 0.1) \times 10^{-5}$ cm^2.

If one assumes the cross-section to be constant at one value of LET, the average of all measurements for the cross-section of the pipeline logic can be taken as a first estimate, which gives $(4.1 \pm 0.1) \times 10^{-5}$ cm^2.

4.3.3 I^{2}C registers cross-section

Similar measurements made of the I^{2}C register cross-section show consistent results with variation in the sensitive time. Figure 4.29 shows the average cross-section to be about $(1.5 \pm 0.9) \times 10^{-5}$ cm^2.

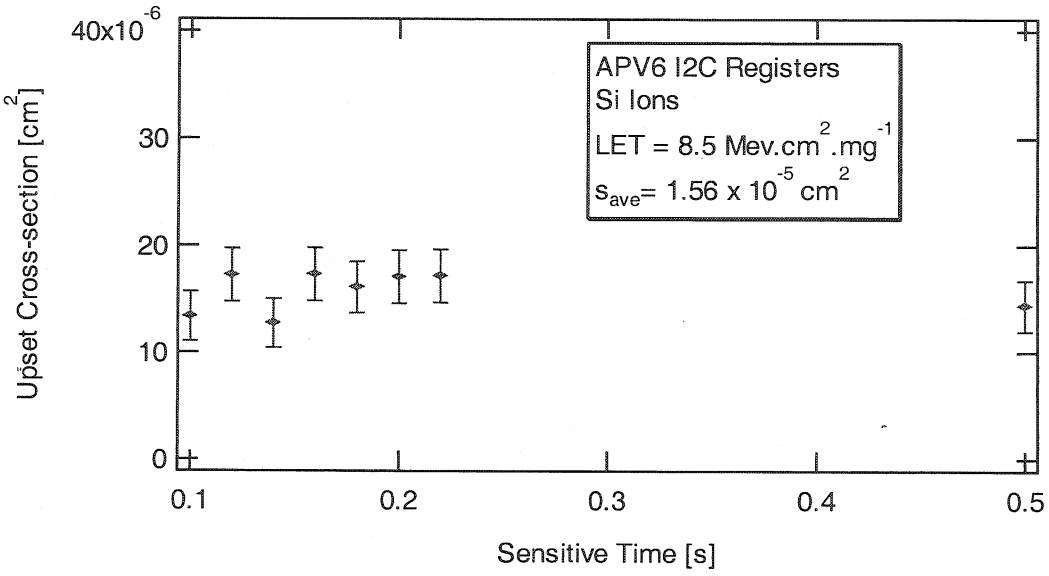

Figure 4.29 Variation of cross-section with sensitive time.

4.3.4 Discussion and future measurements

With only two measurements of cross-section it is impossible to draw a curve of cross-section against LET. Figure 4.30 shows a possible range in which the curve could lie, illustrating the uncertainty that still remains.

It is clear that more points are needed in order to make an accurate measurement of the saturated cross-section, and the threshold LET. The purpose of these measurements is to establish the sensitivity of the final system to SEUs,. However, since the decision has already been made to adopt the APV25 in the final system, the measurement of the APV6 SEU cross-section has become less important.

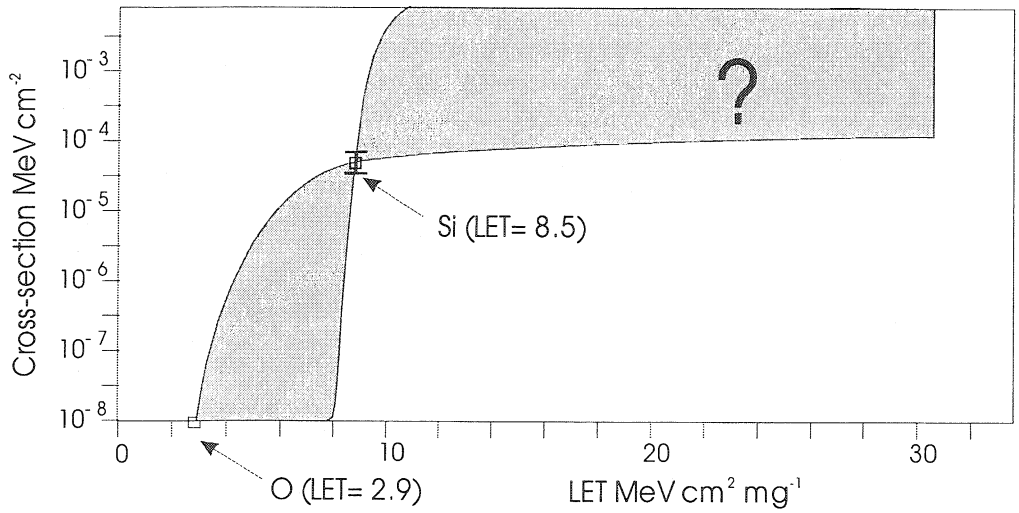

Figure 4.30 Possible range for the APV6 cross-section curve.

Table 4.13 shows a comparison between the predictions in Chapter 3 and the measurements above. Since it is extremely difficult to quantify the systematic errors in the simulations such as variations in device characteristics and implant dimensions, the agreement is very good. In most cases, if the predictions are within an order of magnitude, agreement is considered to be good. If the prediction for the threshold LET is accurate then the APV6 would be sensitive to upsets caused by knock-on silicon ions. The maximum LET for these ions in Silicon is ~ 15 MeV.cm^2.mg^{-1}, which is well above the threshold.

	Predicted	Measured
Threshold LET [MeV.cm^2.mg^{-1}]	2.9	2.9 – 8.5
I^2C Cross-section [cm^2]	1.5 x 10^{-5}	1.5 ± 0.9x 10^{-5}
Pipeline Cross-section [cm^2]	1.2 x 10^{-5}	4.1 ± 0.1 x 10^{-5}

Table 4.13 Predicted and measured upset cross-sections for APV6.

This is the current state of measurements on the APV6. From these measurements it is impossible to make a prediction of the expected upset rate of the APV6, if it were to have been used in the CMS tracker, since a full cross-section curve is required. However, we can note that the threshold lies between 2.9 and 8.5 MeV cm^2 mg^{-1}, and saturating cross-section is >10^{-4} cm^2. In order to make any useful prediction, one must have an upper limit on the cross-section, but the lower limit that we do have will enable us to make a comparison with the APV25, providing the cross-section of the APV25 is lower than that of the APV6.

4.4 APV25 SEU Cross-sections

The measurement of the SEU effect in the APV25 and its subcircuits is of more interest in this thesis, since the final CMS tracker system will contain ~ 100,000. The main emphasis of these measurements is on the shape of the cross-section curves and prediction of the upset rate in the final system. Note the theory and simulations described in Chapter 4 predicted steps in the SEU cross-sections. Measurements were made of the SEU cross-section of the pipeline, FIFO, control and bias registers.

4.4.1 Effective LET

For the APV6 test the incident ion LET values were taken to represent the LET of the ions in the SV. However, the effective LET in the SV is adjusted by the slight loss of energy in the material above the SV. In the APV25 results this loss of energy has been accounted for and an estimate of 6μm has been used for the SV depth, which is taken to be the thickness of the passivation and field oxide layer above the implant. Table 4.14 shows the incident ion LET and the corresponding adjusted LET.

Ion	Si	Cl				Ti		N			B	I
Surface LET [MeVcm^{2}mg^{-1}]	10.4	12.9	14.0	15.0	16.0	20.4	23.2	28.0	29.0	31.6	39.4	61.8
Effective LET	11.2	13.5	14.8	16.1	17.4	22.0	25.2	29.3	30.7	33.3	36.0	58.8

Table 4.14 Incident Ion LET and effective LET after 6μm of surface material.

Note the lower value LETs are increased by a reduction in the energy, whereas the two highest values are reduced. The error in the LET is taken directly from the error in the ion energy at the accelerator, which is 0.01%. The error in the Effective LET is the same plus a systematic error associated with the accuracy of the estimated SV depth. It is not easy to quantify this error, since the true SV depth is unknown, but a conservative estimate would be ~ ±1μm, which gives a 1% error in the effective LET.

4.4.2 Bias Registers

The results in Figure 4.31 and Figure 4.32 show cross-sections for the bias registers for both 1-0 and 0-1 transitions. As expected, it is difficult to see any structure in the curve for upsets from 1 to 0. However the overall shape is similar to expectations. The data for 0-1 transitions displays clear steps very close in likeness to the predicted curve.

Note that the saturated cross-section of 0-1 transitions is almost an order of magnitude lower than that of 1-0 transitions; this will be discussed again in a moment. The saturating cross-section for the 1-0 transitions is very close to the predicted value, but this may be misleading, since the total cross-section is a sum of the modes. The proportion by which the cross-section is divided between the n and p modes is important if we are to understand the true shape of the curve.

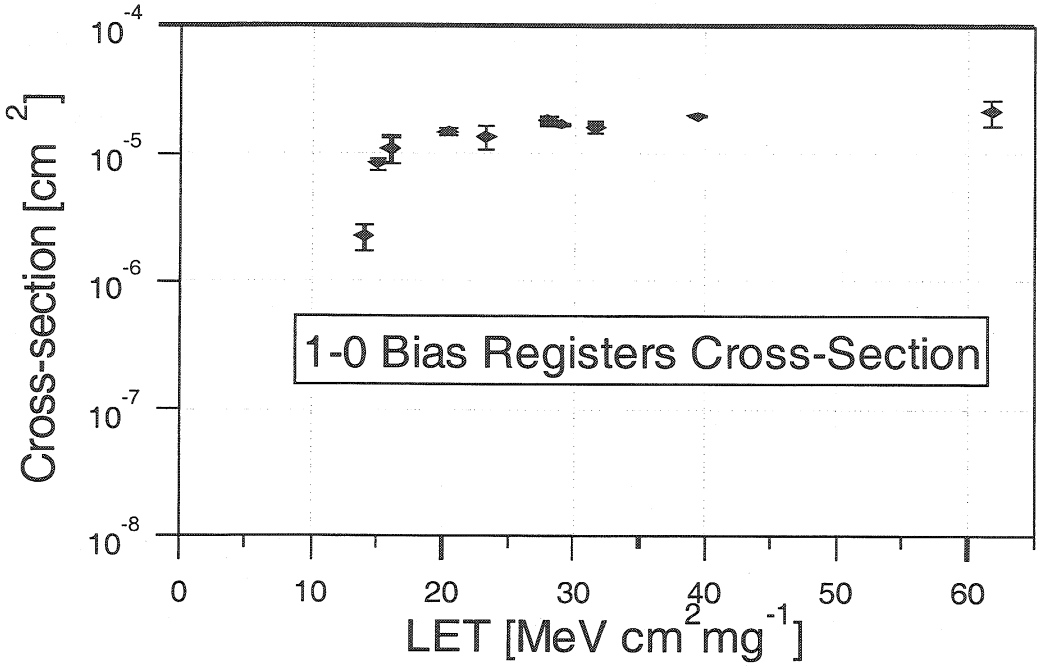

Figure 4.31 1-0 upset cross-section for APV25 bias registers.

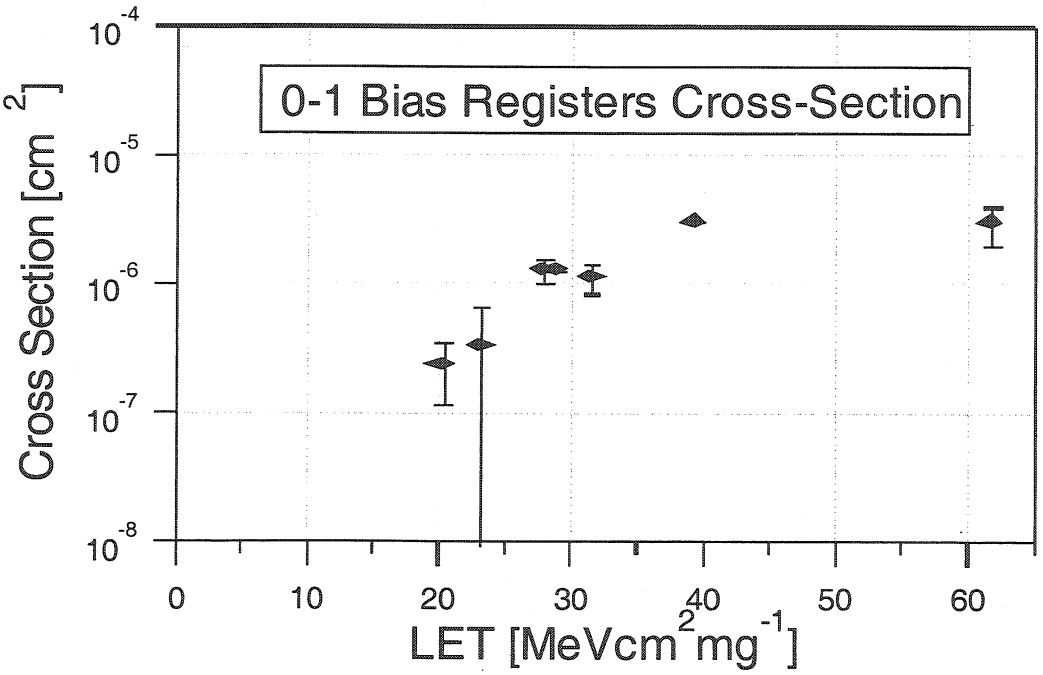

Figure 4.32 0-1 upset cross-section for APV25 bias registers.

4.4.3 Fitting Theoretical Curves to Data

By fitting both a single Weibull curve and the full mode curve defined in chapter 4, we will be able to make a comparison of the two fits. Recapping from chapter 3, the cross-section represented by the Weibull function is described by equation[4.2].

$$\sigma = \sigma_{sat}\left(1-\exp\left\{-\left[\frac{LET-LET_{th}}{W}\right]^{s}\right\}\right) \qquad [4.2]$$

W and s are shaping parameters used in the fit, which is performed by varying these parameters along with the values of LET_{th} and σ_{sat}, so as to minimise the χ^2 value.

The full mode description of the curve is more complicated and can be represented by equation[4.3] (see section 3.6.3).

$$\sigma = w_2\sum_{i=0}^{Nn}\sigma_n^i\left(LET, w_0\right) + w_3\sum_{i=0}^{Np}\sigma_p^i\left(LET, w_1\right) \qquad [4.3]$$

Where N_n is the number of n-modes and N_p is the number of p-modes. Each of the $N_{n/p}$ components of σ_n^i and σ_p^i can be represented by a single Weibull curve, as in equations[4.4] and [4.5]. Each mode has its own values of LET_{th} and σ_{sat} , taken directly from the simulations.

$$\sigma_n^i = \sigma_{nsat}^i\left(1-\exp\left\{-\left[\frac{LET-\left(w_0 LET_{nth}^i\right)}{W_n}\right]^{s_n}\right\}\right) \qquad [4.4]$$

$$\sigma_p^i = \sigma_{psat}^i\left(1-\exp\left\{-\left[\frac{LET-\left(w_1 LET_{pth}^i\right)}{W_p}\right]^{s_p}\right\}\right) \qquad [4.5]$$

The fitting parameters w_0 and w_1 act to vary the LET thresholds for the n and p modes individually, while w_2 and w_3 do the same for the saturating cross-section. Any variation in the LET threshold can be attributed to a difference between the true sensitive depth and the estimation of 1μm. Therefore the true sensitive depth, for both n and p modes, can be extracted after fitting and is given by the inverse of the weighting parameters w_0 and w_1. W_n , W_p , s_n and s_p are fixed parameters from the original

Weibull curve, whose values should be fixed to give the individual Weibull curves the closest approximation to the shape of an individual mode curve. In this instance these parameters are fixed to make the mode curves almost step functions.

It makes sense to vary all modes of the same type by the same amount, but separately from the opposite type of modes for the following reasons:

- The only thing that could vary the relative LET thresholds between modes would be variations in the depth of, or amount of material above, the sensitive volume, between the different sensitive implants. There may well be a difference between the sensitive depth of an n and a p implant, but differences between implants of the same type should be small, since a high degree of uniformity in the processing must be expected.
- Variation in the expected cross-section must be implant type dependent, since the charge collection region and efficiency depend heavily upon the charge carrier type being collected. Also, since there is no major difference in the sensitive areas of different implants of the same type, one would expect the modification to the physical surface area, due to charge collection effects, to be the same for each implant.

In fitting the full mode curve, one varies the four parameters: w_0, w_1, w_2 and w_3, again minimising the χ^2 value. Using a standard curve-fitting algorithm to vary w_2 and w_3 produces weight values and an associated error. However, it is not possible to use a standard algorithm to vary w_0 and w_1 because the curve is composed of step functions, which prevent the fitting algorithm from converging. To circumvent this problem the fit was performed by setting w_0 and w_1 to position the steps in the expected places; for instance, if the cross-section was measured to be zero at 15 $MeVcm^2mg^{-1}$, and 10^{-6} cm^2 at 21 $MeVcm^2mg^{-1}$, then it is clear that the first step must lie between these two points. Certain subtle variations in the data can also force the step to lie exactly upon a point, as in the data shown in Figure 4.31. Therefore it becomes very difficult to assign an error to these weights; in the instance where the step passes through a point the error would be extremely small, but in the other case the step could lie anywhere between the two points, so a range of values for the weight could be assigned. During the calculation of χ^2, w_0 and w_1 were set to a sensible position and the curve fitting algorithm was

performed by varying only w_2 and w_3. In the following sets of data, only three curves are fitted with the full mode fit and in each case the error is considered separately.

The definition in the data of the bias registers 1-0 cross-section in Figure 4.33 is not very high, due to the low number of measured upsets and hence large errors. Even with better statistics it would still be hard to distinguish the individual mode plateaus since the steps in the curve are expected to be small (see section 3.6.3). In this instance a visually good fit can be made by both the simple Weibull, and the full mode curves, Figure 4.34. This may explain why many previous investigations into SEU have not yielded clear steps in the data set. Note also that χ^2 for the mode fit is very high in this instance; because of the problems described with the fitting procedure, the first point can be seen to lie to the left of the first step, which greatly increases the value of χ^2.

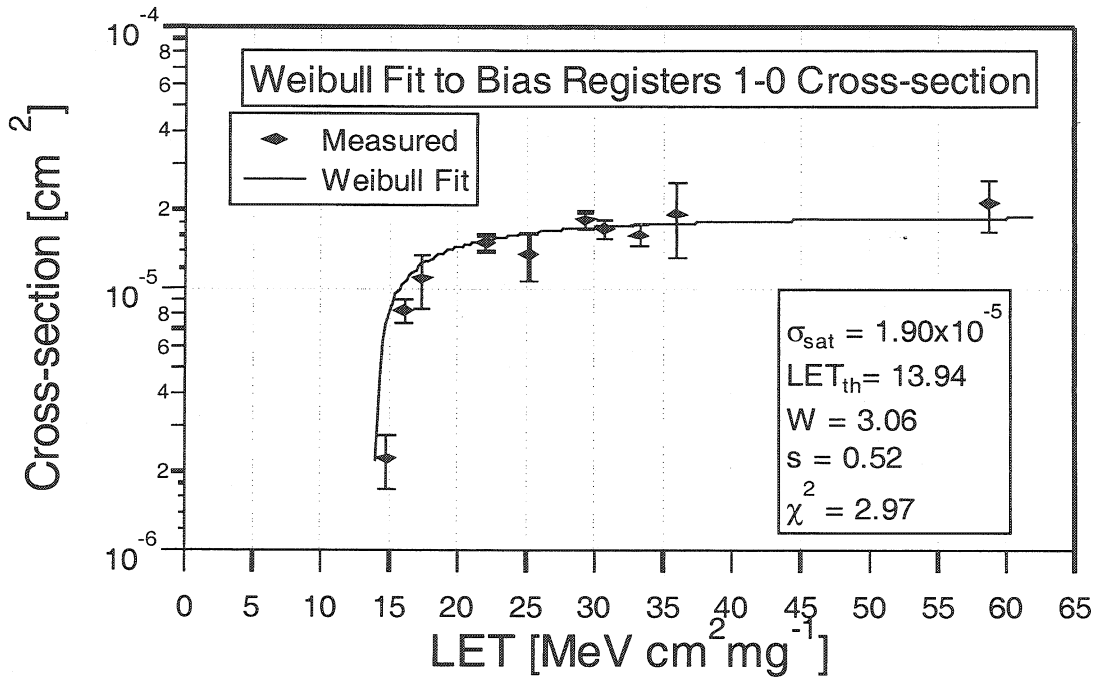

Figure 4.33 Fitting simple Weibull curve to bias registers 1-0 cross-section.

With 6 degrees of freedom, χ^2 per degree of freedom is 0.50 for the Weibull fit, and 7.06 for the full mode fit, which gives a probability that the fit is representative of the data of 0.85 and < 0.01 respectively, it is hard to draw any statistical conclusions about which fit is more likely to be a true representation.

Taking the inverse of the weights w_0 and w_1 in the modes fit, the extracted sensitive depths are 1.33μm for the n modes and 1.69μm for the p modes.

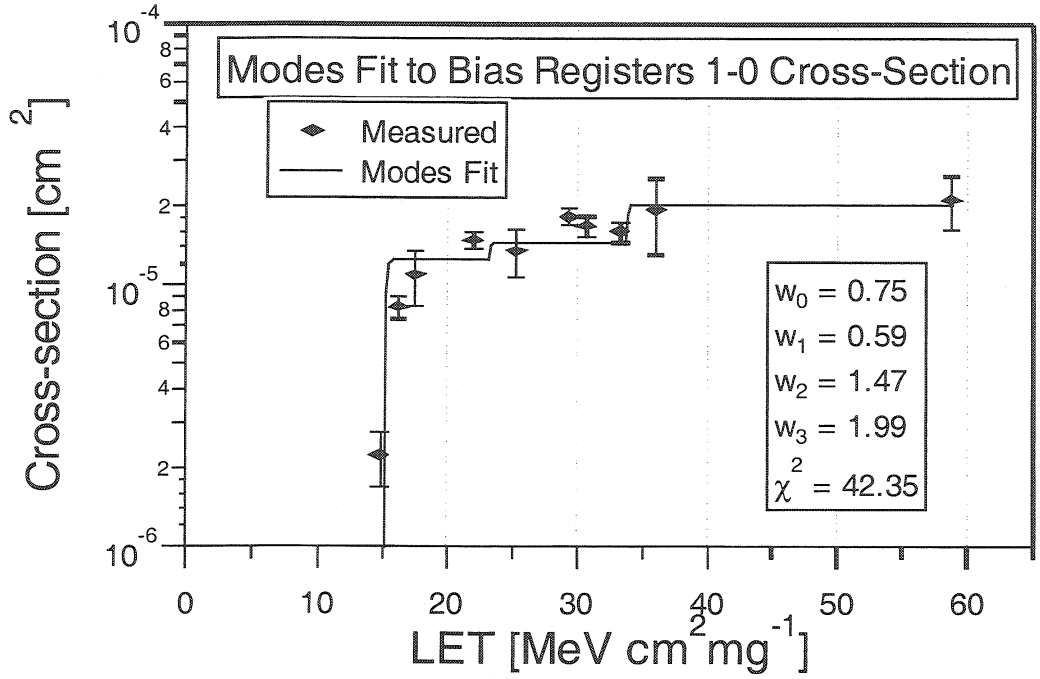

Figure 4.34 Fitting full mode curve to bias registers 1-0 cross-section.

The error in w_0 from the fit is quite small, since the first step has to be close to the first point. The error in the LET of this point is 1%, so this should be a good estimate for the error in w_0, giving an error in the sensitive depth of ± 0.02μm. The first p mode step occurs between 33.3 and 36 MeVcm^{2}mg^{-1}, hence the step has a window of 2.7 MeVcm^{2}mg^{-1}. Taking the error in the position of the step to be equal to about half of the window gives an error of ~ 5%, and hence and error in the sensitive depth of ±0.09μm.

Figure 4.35 shows the Weibull fit to the bias register 0-1 cross-section. With 3 degrees of freedom, χ^2 per degree of freedom is 2.53, yielding a probability that the fit is representative of the data of 0.07. Comparing this value with 0.94, obtained from the full mode fit shown in Figure 4.36, one appears to see an improvement with the new fit. However, a probability of 0.07 is not low enough to statistically exclude the Weibull fit, but it is hard to ignore that visually one can see that the measured points display clear plateaus which match the theory.

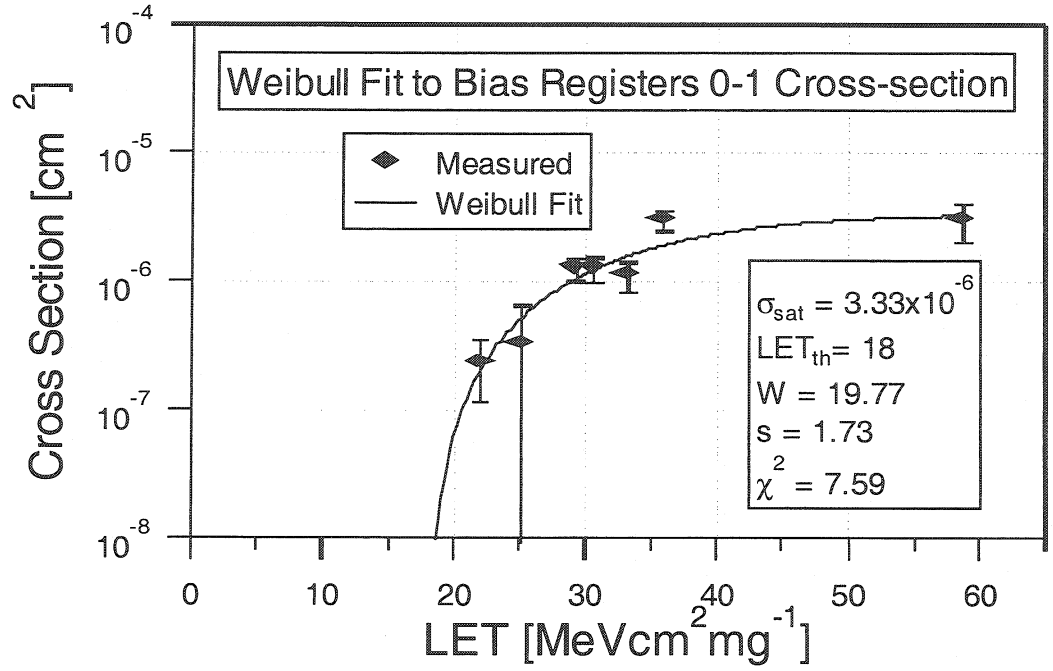

Figure 4.35 Fitting simple Weibull curve to bias registers 0-1 cross-section.

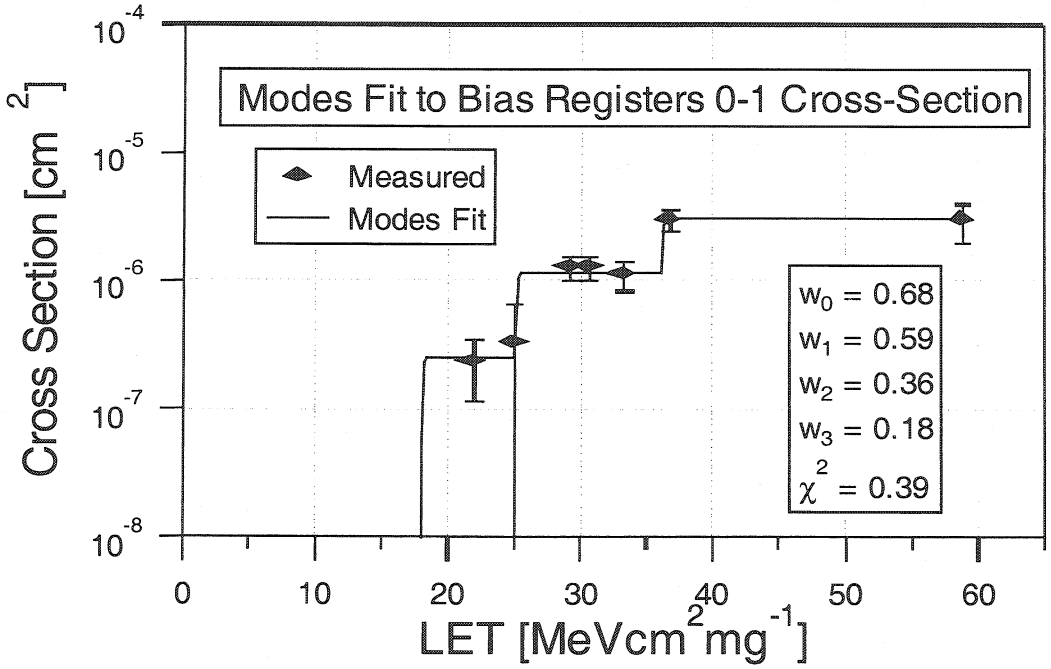

Figure 4.36 Fitting full mode curve to bias registers 0-1 cross-section.

Again extracting the sensitive depths and using the same considerations for the errors gives 1.47 ± 0.15μm for the n modes and 1.69 ± 0.09μm for the p modes, comparing well with those extracted from the curves for upsets from 1-0, as would be expected.

4.4.4 Pipeline

Figure 4.37 shows the results for the pipeline logic along with a fitted Weibull curve. The statistics gathered for this circuit were much better, because of the large cross-section, and even though the steps in the curve are quite small, they are visible. Both the Weibull and full mode fits have been performed on these data and again the new fit in Figure 4.38 has a lower χ^2 value and the plateaus in the data visibly match the predicted steps. With 7 degrees of freedom, χ^2 per degree of freedom is 51.14 for the Weibull fit and 7.56 for the mode fit, which yields a low probability, <0.01, indicating that more thought may need to go into the theory and fitting procedure in order to produce a more consistent description of the upset mechanism.

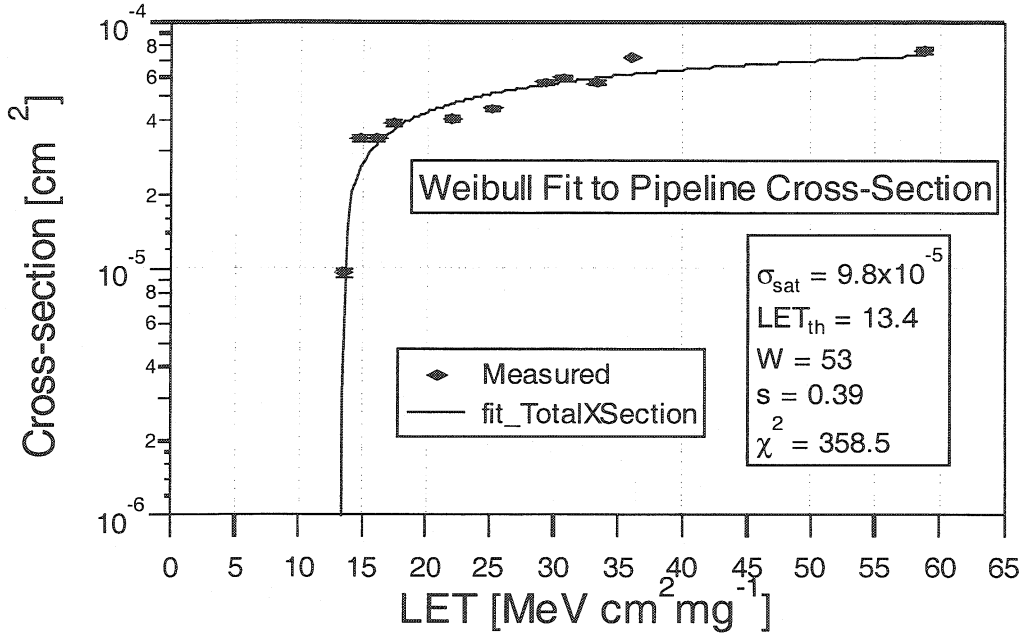

Figure 4.37 Fitting simple Weibull curve to pipeline cross-section.

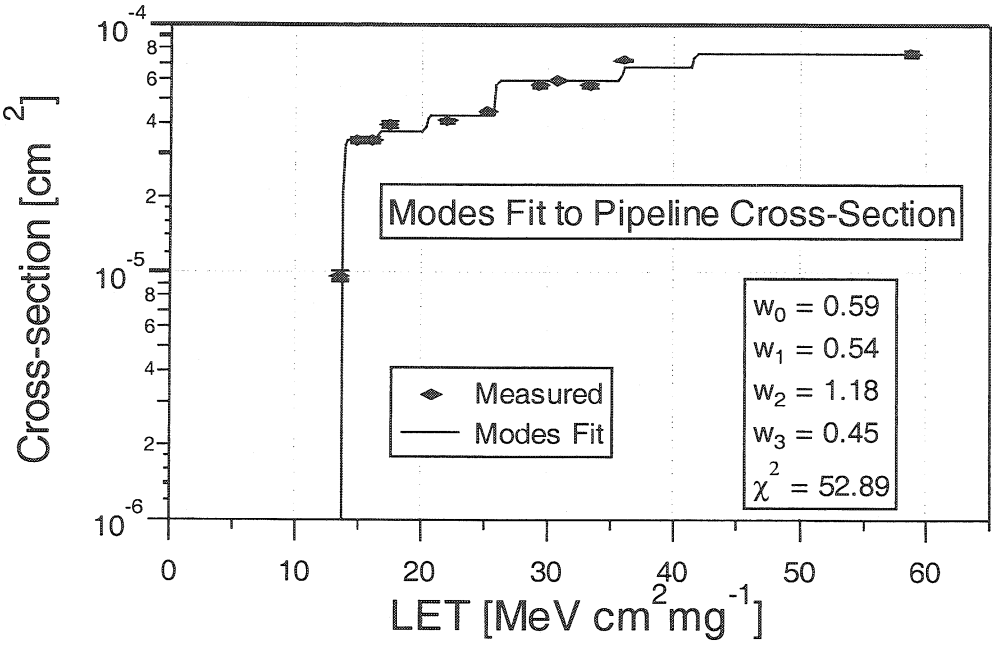

Figure 4.38 Fitting full mode curve to pipeline cross-section.

Again extracting the sensitive depths gives 1.69±0.02μm for the n modes and 1.85±0.09μm for the p modes, slightly larger than the values extracted for the simple DFF, but still in good agreement.

4.4.5 A Closer Look at the Modes

To get a better impression of how good the predictions of the mode cross-sections and LET thresholds were, Table 4.15 and Table 4.16 summarise the original predictions of the thresholds and cross-sections for both pipeline and bias register 0-1 transitions. In both cases w_0 and w_1 introduce offsets for both n and p modes, which account for the error in the predicted sensitive depth.

Mode Type	LET^i_{th} [MeVcm^{2}mg^{-1}]	σ^i [μm^2]	W_0	W_1	W_2	W_3
n	22.9	3.75	0.59 ±0.01	-	1.18 ±0.01	-
	27.5	0.325				
	34.1	0.65				
p	47.9	5	-	0.54 ±0.03	-	0.45 ±0.02
	66.6	2.35				
	77.5	2.375				
Full curve	22.9	14.45	-	-	-	-

Table 4.15 Pipeline predicted mode thresholds and cross-sections with theoretical fitting parameters.

The cross-sections for the pipeline p modes are a factor of 0.45 lower than expected, while the n modes were a factor of 1.18 higher. Errors could be extracted for the cross-section weights from the curve fitting algorithm and those for the threshold weights come from the arguments in the previous section.

Mode Type	LET^{i}_{th} [MeVcm^{2}mg^{-1}]	σ^i [μm^2]	W_0	W_1	W_2	W_3
n	26.47	0.65	0.68 0.01±	-	0.36 ±0.16	-
p	42.10	9.94	-	0.59 ±0.03	-	0.18 ±0.03
	61.63	4.73				
Full curve	26.47	15.35	-	-	-	-

Table 4.16 Bias registers 0-1 predicted mode thresholds and cross-sections with theoretical fitting parameters.

The cross-sections for the bias registers 0-1 transition p modes are a factor of 0.18 lower, which is in agreement with those of the pipeline. However, the cross-section of the first n mode is a factor of 0.36 lower than predicted, unlike the n modes of the pipeline. The origin of this discrepancy is unknown. It is worth noting that this first n mode is due to a very small implant with an expected surface area of 0.65 μm^2, whereas the first n mode for the pipeline is due to an implant with an expected surface area of 3.75 μm^2. There is a chance that the actual implants do not have the expected surface areas, or perhaps some subtlety in the doping profile has an effect on the effective charge collection surface area. This will be investigated in future studies.

4.4.6 FIFO and Control Logic

Predictions of the shape of the FIFO and control logic curves can be made in the same way as those of the bias registers and pipeline. However, these circuits are more complicated in that they are comprised of combinations of all different types of DFFs. In order to predict the shape of these curves one must know the exact number of each type of DFF and then sum all the modes from each with weightings to account for their distribution. Including these curves would add nothing to this thesis, so for brevity they have been omitted.

The more circuits that become involved in the upset mechanism, the smoother the cross-section curve becomes. In these cases it is easier to fit the data with a simple

Weibull curve in order to extract the values of LET_{th} and σ_{sat}. From these and the previous Weibull fits it is then possible to predict the upset rates in CMS.

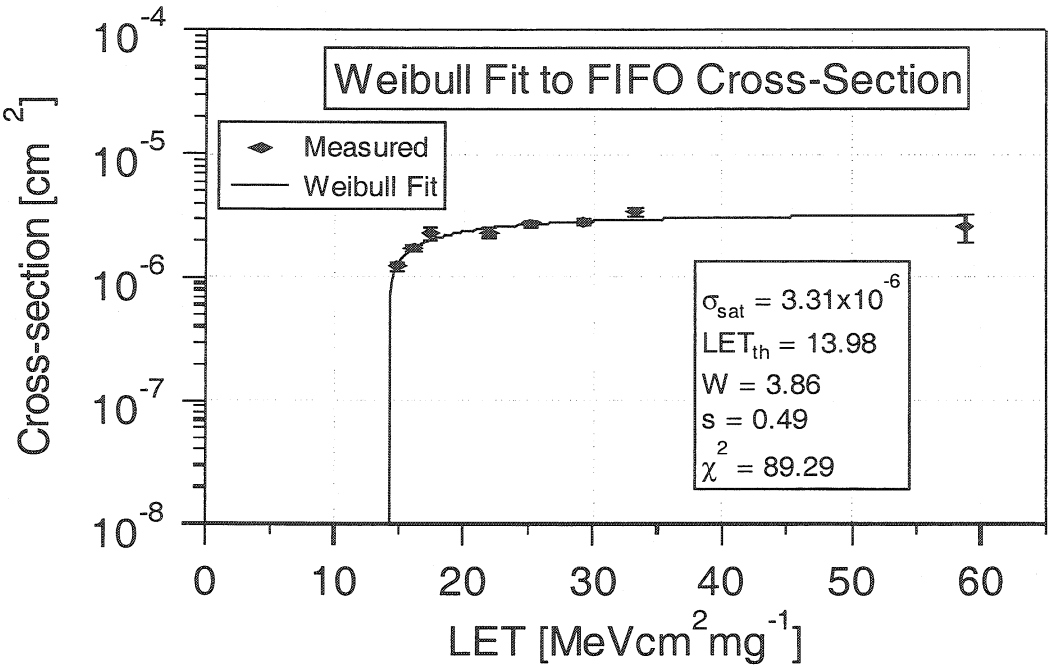

Figure 4.39 Upset cross-section for the APV25 FIFO.

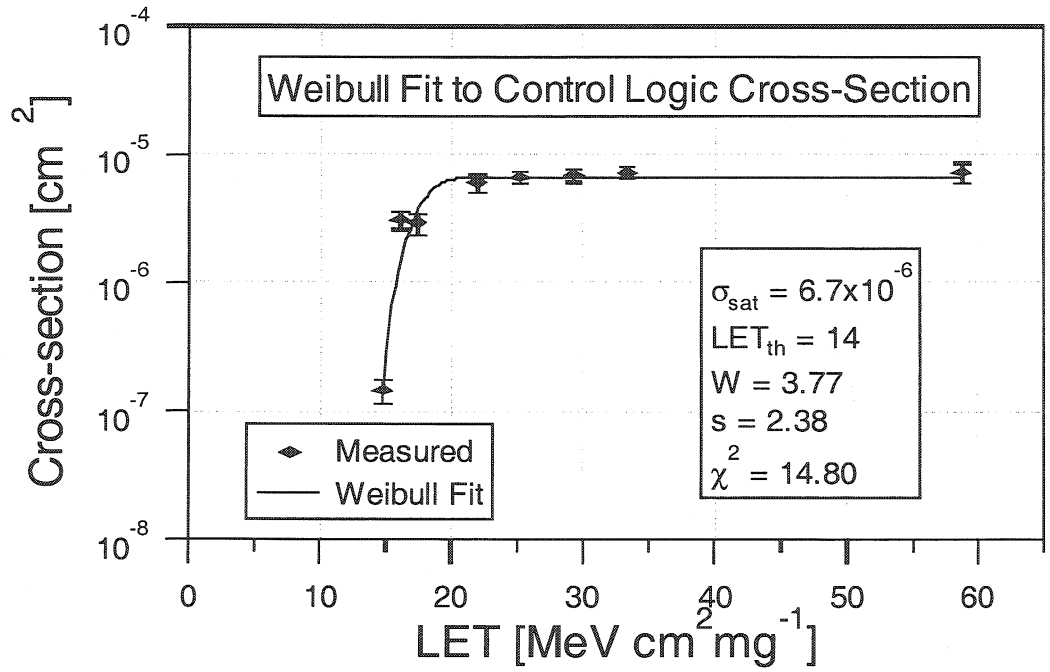

Figure 4.40 Upset cross-section for APV25 control logic.

All the extracted values of σ_{sat} are shown in Table 4.17 along with the geometrical predictions, which are taken as the sum of the cross-sections from all modes.

σ_{sat} [cm^2]	Pipeline	FIFO	I^{2}C (1-0)	I^{2}C (0-1)	Control
Predicted	1.1×10^{-4}	8×10^{-6}	1.3×10^{-5}	1.7×10^{-5}	2.3×10^{-5}
Measured	9.8×10^{-5}	3.3×10^{-6}	2.1×10^{-5}	3.3×10^{-6}	6.7×10^{-6}

Table 4.17: Comparison of predicted and measured cross-sections.

In general the predictions are close to the measured values with the exception of the I^{2}C 0-1 transitions. However, this could be misleading because in the case of the pipeline cross-section the weights act to reduce the p modes but increase the n-modes. It is clear that ongoing work to understand the origin of the difference between n and p modes is necessary. Table 4.18 shows the sensitive depths extracted from three fitted curves, which are in good agreement with each other. The two I^{2}C measurements are comparable, as would be expected since they come from the same circuit.

Sensitive Depth [μm]	Pipeline	I^{2}C (1-0)	I^{2}C (0-1)
N modes	1.69 ± 0.02	1.33 ± 0.02	1.47 ± 0.15
P modes	1.85 ± 0.09	1.69 ± 0.09	1.69 ± 0.09

Table 4.18 Extracted sensitive depths for n and p modes from three fitted curves.

The differences between them and those of the pipeline are not very large and not unexpected since they come from a completely different circuit.

4.5 Upset Rate Predictions for the CMS Tracker

Table 4.19 gives a breakdown of the predicted upset rates in the CMS tracker at full LHC luminosity. These figures have been calculated using the method described in Section 3.7, and summing the effect of the four circuits to produce a prediction of the behaviour of the whole chip. The average predicted SEU cross-section of the APV25 in the tracker is ~ 10^{-12}cm^2. It is extremely difficult to give a quantitative estimate of the error in this prediction; at this stage the simulation code has not been adapted to include a treatment of the errors. It is important that some error should be assigned to these figures since it is likely that it may be large due to the strong dependence of the probability on E_{DEP}, which can be seen in Figure 3.26 and Figure 3.27.

Tacker region	Radius cm	No. APVs	No. SEU/Layer/ s	Seconds / SEU	No. SEU/hour	Fraction chips/hour
Inner Barrel	24-52	14400	1.46×10^{-2}	68.6	52	0.36%
Outer Barrel	60-115	29232	4.1×10^{-3}	243.7	15	0.05%
Inner Endcap	22-41	4416	5.15×10^{-3}	194.2	19	0.42%
Forward Endcap	33-106	30208	8.58×10^{-3}	116.5	31	0.10%
Total	-	78256	3.24×10^{-2}	30.9	116	0.15%

Table 4.19 Upset rates in the CMS tracker.

With no estimate of the error it was necessary to make a measurement of the cross-section in a representative radiation environment. To this end the APV25 was very recently exposed to a 200MeV pion beam at the PSI in Switzerland. Preliminary results indicate that the SEU cross-section was measured as $\sim 2 \times 10^{-12}$ cm^2, which would suggest that even though there remains an uncertainty in the error, the prediction is still very accurate.

4.5.1 Resetting Errors in the Final System

As discussed in section 4.2.5, all upsets in the APV can be recovered by one or more of three methods depending on the location of the upset: an I^2C slow control register re-write is required for upsets in the bias registers, a soft reset is required for upsets in the pipeline logic or FIFO and a hard (power-on) reset followed by I^2C register re-write is required for upsets in the control logic. Each time an error is detected one of the ~100,000 APVs will be flagged, after which it will resume normal function following a reset. The dead time incurred whilst resetting the tracker is shorter for soft resets, since there is not also the need to reset the bias registers via the slow control.

The dominant upset cross-section is that of the pipeline logic, by nearly an order of magnitude. The dead time incurred during a soft reset is approximately the latency + 10 clock cycles, roughly 130 × 25 ns 3.25μs, during this time no data can be collected and data which was stored in the pipeline for readout is lost (on average this is 1 to 2 events). If the reset is applied towards the end of the regular 10μs gap in the proton bunches of the LHC, occurring once every 89μs, the average data loss would be close to one event per reset, since each event takes 7μs to read out. To keep within the CMS

requirement of less than 1% data loss, resets would have to be limited to a maximum rate of once every 700μs. Even if a reset were only applied once every 31s, which is the same as the average time between upsets, only ~ 0.0005 % of data would be lost due to SEUs. In reality the reset rate of the tracker will be less than once every 31s - the exact frequency is under discussion.

Errors in the pipeline and control logic are detectable and hence the final system will be capable of flagging the data from erroneous chips. Errors in the bias registers will go unnoticed until an I^2C read operation is performed. However, since the I^2C upset cross-section is an order of magnitude less than the pipeline cross-section, a regular reset will limit the number of unknown errors to within acceptable limits. Resetting once every 300s would yield $t_d/3$ % loss of data due to the dead time (t_d), and 0.00005% loss due to SEUs. In this instance it is clear that keeping the dead time to a minimum is important. At the moment the exact time it will take to reset all the bias registers in the tracker is not fixed, but from the above evidence it would be sensible to ensure a dead time of less than a few hundred milliseconds.

4.6 Conclusions

Transistor test structures have been irradiated and measured from both Harris and deep sub-micron processes. The overwhelming evidence shows that Harris devices exhibit some degradation and inconsistency between different runs under irradiation totalling 100Mrads. Whereas the deep sub-micron devices show encouraging tolerance up to 50 Mrads for both test structures and the whole APV chip.

The APV25 was irradiated by a heavy ion beam and no evidence of SEGR was observed for fluences up to 2.82×10^9.

Both the APV6 and APV25 were irradiated by heavy ions to evaluate their susceptibility to SEU. The results from the APV25 measurements have been used to predict the upset rates in the CMS tracker giving a value of 0.15% of chips upset per hour, which is well within acceptable limits.

Values for the sensitive depth of both n and p modes have been extracted from the curve fitting and have shown good consistency between measurements, with that of the n modes being slightly smaller than the p modes.

Future work will include developing methods to use the full mode fits to predict upset rates in different radiation environments. However, the only available method to date, which was described in section 4.10, uses the Weibull fit. Also work will continue on calculating the errors in the predicted upset rates, which is important if we are to have more confidence in this method.

Chapter 5

Summary and Conclusions

Throughout this thesis comparisons have been made between two technologies and their suitability for use in the fabrication of the APV front-end readout chip evaluated. Both general manufacture quality and uniformity and radiation tolerance have been thoroughly investigated. In particular the research into the mechanism behind Single Event Upset yielded the explanation for the shape of the normally incident heavy ion cross-section curves. Experiments have also yielded the evaluation of the sensitivity of both technologies to total dose and transient radiation effects.

5.1 Wafer Testing

The comparison of the yields from the APV6 and the APV25 shows a considerable difference in the uniformity and reliability between the two processes; the APV6 showing large variations between runs and on the wafer, but the APV25 showing very good uniformity and tight distributions of currents and other statistics.

It is important that careful consideration of the test threshold settings is made in order to ensure the optimisation of chip performance and minimisation of production costs. Solid information on which to base the failure criteria and the acceptable limits in which the chips should fall have been established; some correlation between currents and failure modes has also been discovered, but it would appear that this correlation is not strong enough to warrant tightening of the current limits and removal of certain tests, in order to bring down the testing time.

Work has been done on developing a comprehensive data storage and recovery system, which will be required later. All data is presently saved in binary format and later converted into a text version for distribution on the web. The development of a database for the test data is an important task and will commence relatively soon.

The exact format of this database and the media, which it will employ, is under discussion. It has also been shown that the pedestal uniqueness of individual chips could be used as a useful chip identifier, which could also be included in the final database.

The main conclusion to draw from these studies is the clear superiority of the deep sub-micron process both in uniformity on the wafer and across different wafers. With an overall yield of 84% compared to an average of ~13% for chips of a similar quality from the Harris process.

5.2 Radiation Effects

A variety of radiation effects have been examined carefully with the main focus of attention on total dose effects in MOS transistors and SEU in DFFs. It is clear that with such a variety of effects it is very important to thoroughly investigate electronic systems intended for operation in radiation environments such as that of CMS. With a better understanding of the mechanisms behind the effects the chances of being able to design radiation hard systems is much improved. To this end a thorough investigation of the SEU mechanism has yielded a novel method for predicting the behaviour of devices under heavy ion irradiation.

Transistor test structures were irradiated and measured from both Harris and deep sub-micron processes. The overwhelming evidence shows that Harris devices exhibit some degradation and inconsistency between different runs under irradiation totalling 100Mrads. Whereas the deep sub-micron devices show considerable tolerance up to 50 Mrads for both test structures and the whole APV chip.

The APV25 was irradiated by a heavy ion beam and no evidence of SEGR was observed for fluences up to 2.82×10^9.

Both the APV6 and APV25 were irradiated by heavy ions to evaluate their susceptibility to SEU. The results from the APV25 measurements have been used to predict the upset rates in the CMS tracker giving a value of 0.15% upset chips per hour, which is well within acceptable limits. This result is supported by recent experimental measurement of the APV25 cross-section in a 200MeV pion beam. However, work to establish the error in the predictive method is necessary.

Values for the sensitive depth of both n and p modes have been extracted from the curve fitting and have shown good consistency between measurements, with that of the n modes being slightly smaller than the p modes.

Future work could include closer examination of the charge collection mechanism in order to try to understand the difference between the SEU cross-section of n and p modes, also developing methods to use the full mode fits to predict upset rates in different radiation environments could improve the accuracy of such calculations.

References

Chapter 1 and Introduction

[1.1]F. Halzen, A. D. Martin, "Quarks and Leptons, An Introductory Course in Modern Particle Physics", Published by John Wiley and Sons, 1984

[1.2]CERN, The Compact Muon solenoid Technical Proposal. CERN/LHCC 94-38

[1.3]CMS Letter of Intent. CERN/LHCC 92-3 (October 1992)

[1.4]CMS Technical Proposal. CERN/LHCC 94-38 (December 1994)

[1.5]Atlas Technical Proposal. CERN/LHCC 94-43 (December 1994)

[1.6]M. Huhtinen, P.A. Aarino, "Neutron and Photon Fluxes and Shielding Alternatives for the CMS Detector at LHC", CERN CMS TN/94-241 (1994)

[1.7]Addendum to CMS Tracker TDR. CERN (2000). CERN/LHCC 2000-016

[1.8]CMS Tracker Technical Design Report. CERN/LHCC 98-6 (1998)

[1.9]RD20 Status Report 1994. CERN/DRDC 94-39 (1994)

[1.10]S. Gadomski, G. Hall, T. Hogh, P. Jalocha, E. Nygard and P. Weilhammer, "The Deconvolution Method of Fast pulse Shaping at Hadron Colliders", Nuclear Instruments and Methods in Physics Research A320 (1992) 217-227

[1.11]M.Raymond et al, "The APV6 readout chip for CMS microstrip detectors", Third workshop on electronics for LHC experiments, London, September 1997. CERN/LHCC/97-60, p.158

[1.12]M.Raymond et al, "The CMS tracker APV25 0.25 Micron CMOS Readout Chip", Sixth workshop on electronics for LHC experiments, Krakow, Poland, October 2000. CERN/LHCC/2000-041, p 130

[1.13]Harris Corporation, Semiconductors section, P.O. Box 833, Melbourne, Florida, U.S.A. http://www.harris.com/

[1.14]Philips data handbook, "I^2C peripherals for microcontrollers". http://www.semiconductors.philips.com/handbook/handbook_38.html

[1.15]J. Mavor, M. A. Jack and P.B. Denyer, "Introduction to MOS LSI Design", Published by Addison-Wesley, 1983.

[1.16]Phillip E. Allen and Douglas R. Holberg. Holt, "CMOS Analogue Circuit Design", Reihart and Winston 1987.

[1.17]Martin Millmore, "Studies of Radiation Hardened Electronics for use in Inner Tracking Systems at the Large Hadron Collider" Thesis, RAL-TH-96-009.

[1.18]Zhong Yuan Chang. Willy M.C. Sansen, "Low-Noise Wide-Band Amplifiers in Bipolar and CMOS Technologies", Published by Kluwer Academic Publishers, 1991

Chapter 2

[2.1]LabVIEW Instrumentation development software, National Instruments, http://www.ni.com/

[2.2]Lawrence Jones, "APV6 User Guide version 2", http://www.te.rl.ac.uk/med/

[2.3]Lawrence Jones, "APV25-S1 User Guide version 1.1", http://www.te.rl.ac.uk/med/

Chapter 3

[3.1]Glenn F. Knoll, "Radiation Detection and Measurement" (Second Edition), John Wiley and Sons, 1979

[3.2]D.E.Groom et al. "Particle Physics Booklet", Extracted from the review of Particle Physics, the European Physical Journal C3,1 (1998)

[3.3]P. Abreu et al., "DELPHI Collaboration, Performance of the DELPHI detector", Nucl. Instrum. Meth. Phys. Res. A378 (1996) 57.

[3.4]J.F.Ziegler. "Stopping and Range of Ions in Solids", vol.1 Publisher : Pergamon Press, New York 1985

[3.5]V.A.J.Van Lint et al. IEEE Transactions on Nuclear Science. NS-19(6) (1972) 181.

[3.6]Martin Millmore, "Studies of Radiation Hardened Electronics for use in Inner Tracking Systems at the Large Hadron Collider", Thesis. RAL-TH-96-009.

[3.7]Andrew S. Grove, "Physics and technology of Semiconductor devices", John Wiley & sons, Inc 1967

[3.8]T.P.Ma & P.V.Dressendorfer, "Ionizing radiation effects in MOS Devices and circuits", Wiley-Interscience 1989

[3.9]A.H. Johnston, "Radiation Effects in Advanced Microelectronics Technologies", IEEE Trans. Nucl. Sci., Vol. 45, No. 3, 1998.

[3.10]G. Anelli, F. Anghinolfi, M. Campbell, J. Christiansen, W. Dabrowski, F. Faccio, E.H.M. Heijne, P. Jarron, K. Kloukinas, M. Letheren, A. Marchioro, P.M. Signe, W. Snoeys, "Rad-Tolerant Design using Commercial VLSI Technologies", Addendum to P63 proposal for studying radiation tolerant ICs for LHC, CERN/LHCC 97-24, LEB P63 add. 1, 27 February 1997.

[3.11]W. Snoeys, F. Faccio, M. Burns, M. Campbell, E. Cantatore, N. Carrer, L. Casagrande, A. Cavagnoli, C. Dachs, S. Di Liberto, F. Formenti, A. Giraldo, E. Heijne, P. Jarron, M. Letheren, A. Marchioro, P. Martinengo, F. Meddi, B. Mikulec, M. Morando, M. Morel, E. Noah, A. Paccagnela, I. Ropotar, S. Saladino, W. Sansen, F. Santopietro, F. Scarlassara, G.F. Segato, P.M. Signe, F. Soramel, L. Vannuci, K. Vleugels, "Layout Techniques to Enhance the Radiation Tolerance of Standard CMOS Technologies Demonstrated on a Pixel Detector Readout Chip", Nucl. Instr. and Meth.A 439 (2000) 349-360.

[3.12]G. Anelli, M. Campbell, M. Delmastro, F. Faccio, S. Florian, A. Giraldo, E. Heijne, P. Jarron, K. Kloukinas, A. Marchioro, P. Moreira, W. Snoeys, "Radiation Tolerant VLSI Circuits in Standard Deep Submicron CMOS Technologies for the LHC Experiments: Practical Design Aspects", IEEE Trans. Nucl. Sci. Vol. 46, No. 6, December 1999.

[3.13]C. Dachs, F. Faccio, A. Giraldo, E. Heijne, P. Jarron, K. Kloukinas, A. Marchioro, A. Paccagnella, "Total Dose Behaviour of Commercial Submicron

VLSI Technologies at Low Dose Rate", Proceedings of the Third Workshop on Electronics for LHC Experiments, London, September 1997. CERN/LHCC/97-60. p 144

[3.14]F.W. Sexton, D.M. Fleetwood, M.R. Shaneyfelt, P.E. Dodd, and G.L. Hash, "Precursor Ion Damage and Angular Dependence of Single Event Gate Rupture in Thin Oxides", IEEE Trans. Nucl. Sci., Vol. 45, No. 6, 1998.

[3.15]A.H. Johnston, G.M. Swift, T. Miyahira and L.D. Edmonds, "Breakdown of Gate Oxides During Irradiation with Heavy Ions", IEEE Trans. Nucl. Sci., Vol. 45, No. 6, 1998.

[3.16]E.Noah, "Transfer Report, for transfer from Mphil to PhD at Imperial College", 2000.

[3.17]P.E.Dodd, M. R. Shaneyfelt, and F. W. Sexton, "Charge Collection and SEU from Angled Ion Strikes", IEEE transactions on Nuclear Science, 1997 Vol. 44, No.6 Pt1, pp.2256-2265

[3.18]L. W. Connell, P. J. McDanial, A. K. Prinja, F. W. Sexton, "Modeling the Heavy Ion Upset Cross-section", IEEE transactions on Nuclear Science, Vol 42, no 2, April 1995.

[3.19]J. V. Ashby et al, "Everest. Version 4.0", December 1998, Mathematical Software Group, Department for Computation and Information, Rutherford Appleton Laboratory, Chilton, Didcot.

[3.20]E. L. Peterson, J.C. Pickel, J. H. Adam, Jr. and E.C. Smith, "Rate prediction for single event effect - a critique", IEEE trans. Nucl. Sci. 39, pp. 1577-1599, December 1992

[3.21]E.L. Peterson, P.Rudeck, J.C.Pickel and J.Letaw, "Geometrical Factors in SEE Rate Calculations", IEEE Transactions on Nuclear Science, 40, Dec 1993

[3.22]M.Huhtinen, F.Faccio, "Computational method to estimate Single Event Upset rates in an accelerator environment", NIM A 450 (2000) 155-172.

[3.23]V.Pouget et al. "Spice modeling of the transient response of irradiated MOSFETs", IEEE Transactions on Nuclear Science, Vol. 47, No. 3. June 2000.

[3.24]L.W.Massengill, M.L.Alles,S.E.Kerns and K.Jones, "Analytical Correlation between observed non-ideal SEU cross section data and process parameter distributios." IEEE Trans. Nucl. Sci. Dec 1993

[3.25]J.B. Langworthy, "Depletion region geometry analysis applied to single event sensitivity," IEEE Trans. Nucl. Sci 36, pp. 2427-2434, Dec 1989.

[3.26]F. Faccio, M. Huhtinen, "First evaluation of the single event upset (SEU) risk for electronics in the CMS experiment", CMS Note1998/054

[3.27]R.J.Glauber and G.Mattiae, Nucl. Phys. B21 (1970) 135.

Chapter 4

[4.1]M. Raymond, G. Hall, M. Millmore, R. Sachdeva, M. French, E. Nygard, K. Yoshioka, "Measurements of transistors and silicon microstrip detector readout circuits in the Harris AVLSIRA rad-hard CMOS process", Nucl. Instr. and Meth. A 351 (1994) 449.

[4.2]C.M. Dozier et al., "An Evaluation of Low Energy X-ray and Cobalt-60 Irradiation on MOS Transistors", IEEE Trans. Nucl. Sci, Vol. 34, 1987.