



Pennicard, David (2009) *3D detectors for synchrotron applications*. PhD thesis.

<http://theses.gla.ac.uk/694/>

Copyright and moral rights for this thesis are retained by the author

A copy can be downloaded for personal non-commercial research or study, without prior permission or charge

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

3D Detectors for Synchrotron Applications

David Pennicard



Department of Physics & Astronomy
Faculty of Physical Sciences

*Submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy*

April 24 2009

© D. Pennicard April 24 2009

Abstract

3D detectors are a novel variety of photodiode radiation detector, invented by Parker, Kenney and Segal (1997). Instead of having n- and p-type contacts on the front and back surfaces of a silicon substrate, like a standard photodiode, they have columns of doped material passing through the thickness of the silicon. This structure means that the detector can combine a reasonable substrate thickness with a very small electrode spacing, resulting in a low depletion voltage, fast charge collection and low charge sharing.

These detectors have a couple of promising applications. Their fast charge collection and low depletion voltage should make them very radiation-tolerant. So, they could be used for future particle physics experiments at the Super Large Hadron Collider (SLHC), where high levels of radiation damage are expected. Also, their low charge sharing means they could potentially improve X-ray diffraction measurements at synchrotrons such as Diamond Light Source. This would allow these experiments, for example, to determine the structures of biological molecules more accurately.

However, before 3D devices can be used in practical experiments, their design and fabrication must be optimised to ensure that reliable, high-performance detectors can be produced on a reasonably large scale. The aim of this thesis is to evaluate and understand the behaviour of a variety of 3D detectors using a combination of lab tests and computer simulations. Using these results, future fabrication runs can then be re-designed to improve their performance.

Firstly, the “Synopsys TCAD” simulation package was used to determine the optimum design for 3D detectors at the SLHC. It was found that the device behaviour depends strongly on the electrode spacing, and the choice of spacing requires a trade-off between different effects. Using a smaller spacing reduces the detector’s operating voltage, and improves the charge collection efficiency by reducing carrier trapping. However, reducing the spacing also increases the capacitance, resulting in greater noise, and also increases the insensitive volume occupied by the columns. At SLHC radiation damage levels, the optimal electrode spacing was found to be 40–55 μm .

CNM (Centro Nacional de Microelectronica) in Barcelona have produced a set of “double sided” 3D detectors. The n- and p-type columns in these devices are etched

from opposite sides of the substrate and do not pass through the full substrate thickness. Computer simulations show that these detectors should give similar performance to full-3D detectors. The main difference is that these devices have slower charge collection around their front and back surfaces. Basic electrical characterisation of the detectors showed that they have low depletion voltages. However, the guard ring current varied a great deal between detectors, though this was fixed by using better guard structures. Charge collection tests on these detectors using beta particles gave mixed results. A heavily-irradiated detector gave a relatively high collection signal, similar to the simulated value, which demonstrated the structure's radiation hardness. However, an unirradiated detector gave an unexpectedly low collection signal. This was perhaps due to poor coupling between this detector and the readout chip.

Three of these “double-sided” 3D detectors were bonded to Medipix2 pixel readout chips. These chips are specifically designed for X-ray detection, and can count individual photon hits. The detectors worked successfully, and initial lab tests demonstrated that they depleted extremely rapidly. The detectors were then tested in an X-ray beam at Diamond Light Source. These tests showed that the detectors have lower charge sharing than a standard planar photodiode. For example, 24% of the hits on a double-sided 3D detector at 22V were shared, compared to 40% on a planar detector at 100V.

A set of devices with a simplified “single-type-column” structure, fabricated by FBK-IRST in Trento, were also tested. Simulations showed that although this structure will have a low depletion voltage and fast electron collection, the hole collection will be slow. This will result in poorer behaviour than full- and double-sided 3D detectors. This was confirmed by lab tests, which showed that when the detector was coupled to fast readout electronics, the charge collection efficiency was reduced due to ballistic deficit.

Contents

1	3D silicon detectors and synchrotrons	1
1.1	Operating principles of silicon photodiode detectors	1
1.1.1	Basic semiconductor device physics	2
1.1.2	The p-n junction	6
1.1.3	P-n junction photodiodes	8
1.1.4	Readout electronics	13
1.1.5	Further design considerations	17
1.1.6	CCDs and CMOS detectors	19
1.2	3D detectors	21
1.2.1	The 3D detector structure	21
1.2.2	Fabrication of 3D detectors	24
1.2.3	Active edges	26
1.3	Introduction to synchrotrons	27
1.4	Radiation-hard detectors for high-luminosity colliders	28
1.4.1	Particle physics at the Large Hadron Collider	28
1.4.2	Detectors at the LHC	30
1.4.3	The Super-LHC and radiation hardness	32
1.5	Detectors for X-ray synchrotrons	33
1.5.1	Synchrotron light sources	33
1.5.2	Diamond: a third-generation light source	34
1.5.3	Applications of synchrotron X-rays	36
2	Simulation of 3D detectors for high-luminosity colliders	38
2.1	Simulation with Synopsys TCAD	39
2.1.1	Basic finite-element analysis in semiconductors	39
2.1.2	Mesh building with MESH	41
2.1.3	Simulation with Sentaurus Device	43
2.2	Basic simulation of full-3D detectors	45
2.2.1	Mesh structure	45

2.2.2	Depletion and electric field behaviour	46
2.2.3	Typical charge collection speed	48
2.2.4	Weighting field	49
2.3	Radiation damage	53
2.3.1	Mechanisms of radiation damage	53
2.3.2	Microscopic defects and macroscopic damage	54
2.3.3	Modelling radiation damage in Synopsys TCAD	59
2.3.4	Comparing the damage model to experiment	62
2.4	ATLAS 3D detectors	64
2.4.1	ATLAS 3D pixel structures	64
2.4.2	Introduction to simulating radiation-damaged ATLAS 3D	65
2.5	Comprehensive simulation of radiation-damaged ATLAS 3D detectors	67
2.5.1	Depletion voltage and high-field behaviour	68
2.5.2	Average charge collection efficiency at $10^{16}n_{eq}/\text{cm}^2$	70
2.5.3	Uniformity of the charge collection across the pixel	71
2.5.4	Capacitance and noise	73
2.5.5	The new ATLAS pixel chip	79
2.6	Conclusions	80
3	Experimental techniques for silicon detectors	82
3.1	Current-voltage (I-V) testing	83
3.2	Capacitance-voltage (C-V) testing	83
3.3	Charge collection—pad detector	84
3.3.1	Basic principles	84
3.3.2	Test setup	85
3.3.3	Data analysis	88
3.4	Charge collection—LHC-speed strip detector	89
3.4.1	The Beetle readout chip	90
3.4.2	The detector module	93
3.4.3	The TELL1 readout board	95
3.4.4	Data analysis with Vetra and ROOT	96
4	Double sided 3D detectors from CNM	101
4.1	Double-sided 3D detectors	101
4.1.1	Structure and fabrication	101
4.1.2	Devices produced by CNM	103
4.1.3	New fabrication run	105
4.2	Simulation of double-sided 3D detectors	106

4.2.1	Simulation methods and device structure	106
4.2.2	Depletion, electric field and weighting field	107
4.2.3	Charge collection	111
4.2.4	C-V characteristics	114
4.2.5	Breakdown behaviour	115
4.2.6	Effects of altering the substrate thickness	118
4.2.7	Behaviour after radiation damage	119
4.3	Experimental results	120
4.3.1	I-V tests	120
4.3.2	C-V tests	123
4.3.3	Unirradiated strip detector charge collection efficiency	127
4.3.4	Strip detector irradiated to 5×10^{15} 1MeV- n_{eq}/cm^2	133
4.4	Conclusions	138
5	3D detectors for synchrotron X-ray experiments	141
5.1	X-ray crystallography	141
5.1.1	Theory, experiment and applications	141
5.2	Medipix2 single-photon-counting detectors	144
5.2.1	The Medipix2 readout chip	144
5.2.2	Advantages and disadvantages of single-photon-counting	146
5.2.3	Charge sharing, X-rays and 3D detectors	147
5.3	3D Medipix detectors	148
5.3.1	Detector production	148
5.3.2	Data acquisition	149
5.4	Initial 3D Medipix tests	150
5.4.1	Digital write / read test	150
5.4.2	Detector biasing and leakage current	150
5.4.3	Threshold equalisation	151
5.4.4	Basic image acquisition	153
5.4.5	Detection efficiency versus voltage	154
5.4.6	Electronic noise, pixel masking, and radiation damage	154
5.5	Tests on Diamond beamline B16	157
5.5.1	Preparation for tests	158
5.5.2	Beam spectra measurements and calibration	159
5.5.3	Measuring Line Spread Function using an edge	163
5.5.4	Silicon powder diffraction experiment	166
5.6	Alpha particle tests	171
5.7	Conclusions	173

6	Single-type-column 3D detectors from FBK-IRST	175
6.1	Single-type column 3D (3D-STC) detectors	176
6.1.1	Structure and fabrication	176
6.1.2	Devices produced by FBK-IRST	177
6.2	Simulation	179
6.2.1	Simulation methods and device structure	179
6.2.2	Depletion and electric field behaviour	180
6.2.3	Weighting field in pad and strip detectors	183
6.2.4	Charge collection behaviour	186
6.3	Experimental results	189
6.3.1	Pad detector I-V tests	189
6.3.2	Pad charge collection efficiency	192
6.3.3	Pion beam test of a single-type-column 3D strip detector . . .	196
6.4	Conclusions	202
7	Conclusions	205

List of Figures

1.1	Tetrahedral covalent bonding structure in silicon.	2
1.2	Band structures of a semiconductor, an insulator and a metal.	3
1.3	Diagram of a p-n junction in equilibrium.	7
1.4	Structure of a typical p-i-n photodiode	9
1.5	Diagram of a bump-bonded hybrid pixel detector.	15
1.6	The charge-sensitive preamplifier circuit.	16
1.7	A typical guard ring structure.	19
1.8	Structure of a full-3D detector.	22
1.9	Typical electrode layout of a 3D detector.	22
1.10	SEM images showing electrode column fabrication.	25
1.11	Accelerators and experiments at CERN.	29
1.12	Computer image of the ATLAS detector.	32
1.13	Aerial photograph of Diamond Light Source.	35
1.14	Schematic of the Diamond synchrotron.	36
2.1	Region of a 2D mesh, with a network of nodes which form triangular elements.	39
2.2	Structure of a full-3D simulation mesh.	43
2.3	Electric field in a horizontal cross-section of a full-3D detector.	47
2.4	Electrostatic potential in a vertical cross-section of a full-3D detector.	48
2.5	Simulated collection current from a 3D detector.	49
2.6	Weighting potential in a planar strip detector.	50
2.7	Weighting potential in a 3D pixel detector.	51
2.8	Weighting potential in a 3D pad detector.	52
2.9	Effects of radiation-induced traps.	55
2.10	Change in effective doping with radiation fluence in an n-type detector.	57
2.11	Illustration of the 3-trap model from Ref. [67].	61
2.12	Simulated and experimental depletion voltages in n-in-p planar pad detectors.	62

2.13	Simulated and experimental charge collection in n-in-p planar strip detectors.	64
2.14	Possible ATLAS 3D pixel layouts.	65
2.15	Structure used in the “3-column” ATLAS pixel detector simulations.	66
2.16	Simulated and experimental charge collection in a “3-column” ATLAS pixel detector.	67
2.17	Simulated depletion voltages and high-field behaviour in ATLAS pixel devices at $10^{16}n_{eq}/\text{cm}^2$	68
2.18	Surface behaviour in a 4-column ATLAS pixel device at $10^{16}n_{eq}/\text{cm}^2$	69
2.19	Simulated charge collection in ATLAS pixel devices at $10^{16}n_{eq}/\text{cm}^2$	71
2.20	Simulated charge collection versus bias in the 5-column ATLAS 3D detector.	72
2.21	Charge collection versus position in different ATLAS 3D detectors at $10^{16}n_{eq}/\text{cm}^2$	74
2.22	Electric field distributions in different ATLAS 3D devices.	75
2.23	Simulated capacitances of ATLAS 3D devices.	76
2.24	Estimated signal-to-noise ratios in ATLAS 3D after $10^{16}n_{eq}/\text{cm}^2$	78
2.25	Signal-to-noise comparison between ATLAS 3D and planar n-on-p detectors.	79
3.1	A single-type-column 3D pad detector mounted on a chip carrier and PCB.	85
3.2	Block diagram of the pad detector CCE setup.	86
3.3	Oscilloscope traces of signals produced by the pad detector CCE setup.	87
3.4	The Landau distribution.	89
3.5	Calibration graph for the pad detector CCE setup	90
3.6	Diagram of the strip detector CCE setup.	91
3.7	Schematic of the Beetle detector.	92
3.8	The pulse produced by the Beetle’s shaping amplifier.	92
3.9	A strip detector module.	94
3.10	Pulse shape from a planar strip detector.	99
3.11	Signal spectrum from a planar strip detector.	100
4.1	Structure of a double-sided 3D detector.	102
4.2	A 4” n-type double-sided 3D detector wafer produced by CNM.	104
4.3	Detail of the short strip detector structure produced by CNM.	105
4.4	Structure of a simulated double-sided 3D detector.	106
4.5	Depletion behaviour of a double-sided 3D detector.	108

4.6	Electric field in a p-type substrate, n-type readout double-sided detector	109
4.7	Electric field in an n-type substrate, p-type readout double-sided detector	110
4.8	Weighting potential in a double-sided 3D detector	111
4.9	Simulated current signals from different 3D detectors.	112
4.10	Variation in charge collection time with depth in a double-sided 3D detector.	113
4.11	Simulated C-V characteristics of a double-sided 3D detector.	114
4.12	Electric field around the column tip in a double-sided 3D detector.	116
4.13	Electric field distributions in double-sided 3D detectors with different substrate thicknesses.	118
4.14	Electric field in a double-sided 3D detector following $10^{16}n_{eq}/\text{cm}^2$ radiation damage.	120
4.15	Simulated charge collection signals in double-sided and full-3D detectors with equal column lengths following radiation damage.	121
4.16	Current-voltage characteristics of a 3D strip detector.	121
4.17	Inverse capacitance characteristics of a 3D pad detector.	124
4.18	C-V characteristics of a 3D strip detector.	126
4.19	Signal spectrum from a double-sided 3D strip detector at 20V.	128
4.20	Signal spectrum from a double-sided 3D strip detector at 6V.	128
4.21	Experimental and simulated CCE of the double-sided 3D strip detector.	129
4.22	Mean pulse shapes measured on the double-sided 3D strip detector.	130
4.23	Noise on the double-sided 3D strip detector.	131
4.24	Charge sharing in double-sided 3D and planar detectors.	133
4.25	Noise on the irradiated double-sided 3D strip detector at room temp.	134
4.26	Noise on the irradiated double-sided 3D strip detector after cooling.	135
4.27	Signal spectrum from a double-sided 3D strip detector irradiated to $5 \times 10^{15} \text{ 1MeV-}n_{eq}/\text{cm}^2$	135
4.28	Charge collection comparison between double-sided 3D, full-3D and planar detectors.	136
4.29	Signal-to-noise comparison between double-sided, full-3D and planar detectors.	137
5.1	Diffraction image taken with a PILATUS 6M detector.	143
5.2	Circuitry within each Medipix2 pixel.	145
5.3	Front surface of a Medipix2 3D sensor.	149
5.4	Medipix2 chipboard connected to IEAP's USB interface.	150
5.5	Total detector current vs voltage for 3D and planar Medipix2 detectors.	151

5.6	Distributions of noise centroid positions in each pixel of the 3D1 detector.	152
5.7	Image of a PCB, taken with detector 3D2 and an X-ray tube.	153
5.8	X-ray count rate on 3D1 against detector bias.	155
5.9	Noise results from 3D1.	156
5.10	Noise results from a planar detector.	156
5.11	Change in the equalised THL adjust values for detector 3D1 due to beam damage.	158
5.12	Detectors mounted in the B16 beamline at Diamond Light Source.	159
5.13	Calibration plot for the 3D1 sensor.	160
5.14	Spectrum from the 15keV beam measured by the 3D detector.	161
5.15	Comparison of 15keV beam spectra measured by 300 μ m 3D and planar detectors.	162
5.16	Edge Spread Function measured by the 3D1 detector.	164
5.17	Line Spread Function measured by the 3D1 detector.	165
5.18	Line Spread Functions seen on the planar detector with different threshold settings	166
5.19	Diagram showing how charge sharing and varying thresholds affect the LSF from a single-photon-counting chip.	167
5.20	Variation in the FWHM of LSFs measured on 3D and planar detectors at different threshold settings.	167
5.21	A silicon powder diffraction ring imaged with 3D1.	169
5.22	Seven diffraction rings imaged with 3D1.	169
5.23	Silicon powder diffraction intensity versus angle, measured with 3D1.	170
5.24	Profile of the first diffraction ring obtained from images taken by 3D sensor.	171
5.25	Alpha clusters seen on detector 3D1.	172
5.26	Histogram of alpha cluster sizes on 3D1.	172
5.27	Variation in alpha cluster size with threshold setting for the 3D and planar detectors.	173
6.1	Structure of a single-type-column 3D detector.	176
6.2	A single-type-column 3D pad detector.	178
6.3	Depletion behaviour of a single-type-column 3D pad detector.	181
6.4	Structure of the simulated single-type-column 3D strip detector.	182
6.5	Electric field strength in a 3D-STC strip detector at 100V.	183
6.6	Electric field strength in a 3D-STC strip detector at 200V.	184
6.7	Weighting potential of a 3D-STC strip detector.	185

6.8	Weighting potential of a 3D-STC pad detector.	186
6.9	Simulated current signals in different 3D-STC devices.	188
6.10	Simulated collected charge in different 3D-STC devices.	189
6.11	I-V characteristics of 300 μ m-thick MCz 3D-STC pad detectors. . . .	190
6.12	I-V characteristics of 500 μ m-thick FZ 3D-STC pad detectors.	191
6.13	I-V behaviour of the guard ring of each 3D-STC pad detector.	192
6.14	Signal spectrum obtained from a 3D-STC pad detector at 30V.	193
6.15	Charge collection results from the 3D-STC beta particle test.	194
6.16	Signal spectrum obtained from a 3D-STC pad detector at 0V.	195
6.17	Photo of the 3D-STC test beam setup.	198
6.18	Pulse shapes from 3D-STC and planar strip detectors.	199
6.19	Signal spectrum obtained from the 3D-STC strip detector.	200
6.20	Average pulse shapes on the two neighbouring strips to the hit strip. .	201
6.21	Hit efficiency map for a unit cell of the 3D-STC strip detector.	203

List of Tables

2.1	Original and modified radiation-induced trap models.	60
4.1	Simulated capacitances of double-sided 3D and full-3D detectors.	115
4.2	Measured capacitances of double-sided 3D pad detectors.	123
4.3	Capacitances of new double-sided 3D pad detectors.	124
5.1	Results from spectra measurements on planar and 3D detectors.	163
5.2	Calculated and measured silicon powder diffraction angles.	170

Acknowledgements

Firstly, I would like to thank my supervisor, Chris Parkes, for all his help and guidance over the last few years, not to mention the opportunities that his project has given me. I knew that I'd be going to Oxfordshire during the PhD: California and Siberia were a bit more of a surprise! Thanks also to my secondary supervisor, Val O'Shea.

Throughout most of my PhD I've worked with Celeste Fleta, who deserves thanks for all her work on the 3D project, her help and advice with lab tests, simulation and paper-writing, and also many useful discussions. Also, thanks for being great company, particularly during a couple of sleep-deprived testbeams at CERN and Diamond.

I would like to thank Diamond Light Source for sponsoring my EPSRC CASE studentship. I'm grateful to Victoria Wright, Nicola Tartoni and Julien Marchal for supervising me over the course of my PhD, and for making me feel very welcome whenever I've visited Diamond.

Thanks to Giulo Pellegrini, Manolo Lozano and other staff at CNM, not only for providing me with detectors to test but for their advice on the work. Thanks also for inviting me to Barcelona! This work was supported by the Spanish Ministry of Education and Science through the GICSERV program "Access to ICTS integrated nano- and microelectronics cleanroom".

Thanks to Maurizio Boscardin, Gian-Franco Dalla Betta, Claudio Piemonte, Alberto Pozza, Sabina Ronchin and Nicola Zorzi at FBK-IRST, for proving me with their 3D-STC detectors.

Many thanks to all the staff and students at Glasgow for their help with the PhD work and for making the whole experience so enjoyable. In particular, Richard Bates organised a lot of the 3D project work, and helped ensure that I actually had detectors to test! Thanks to Lars Eklund for many hours of help with readout electronics, and for remaining cool under pressure. Thanks also to Fred, Fiona and John for their help with the lab equipment and assembling the detectors.

It's hard to list all the people in the group who have helped me, and who I've shared good times with, but I'll give it a go. Laurence, Dima, Franciole, Kenny and Derek have all shared the stresses and frustrations of the PhD and thesis writing. Over the years I've had a series of great office mates—Andy, James, Isaac and

Liam. Thanks also to Keith, Aaron, Andrew, Marco, Barinjaka, Debbie, Tomasz and everyone else.

I'd also like to show my appreciation to the other friends who've helped remind me that there is a life outside of physics. Thanks to my flatmates China, Liz and Eleanor, who've been such good company over the years; Colin, Chris and Stuart from undergrad, for cinema trips and all-you-can-eat buffets; Paul and Eamonn for some great Friday nights during their Masters'; and Gary, Sarah, Lynn, Dave, Niall and others from Lesmahagow for staying in touch.

Finally, many thanks to my family for their love and support—Ian, Trixie, Lisa, Bappa and Maya.

Declaration

I declare that except where explicit reference is made to the work of others, this dissertation is the result of my own work. This work has not been submitted for any other degree at the University of Glasgow or any other institution.

David Pennicard

Chapter 1

3D silicon detectors and synchrotrons

As well as being used for applications like computation and signal processing, semiconductor devices can be used to detect ionising radiation. Photodiodes are a variety of semiconductor detector which are commonly used in particle physics experiments and related areas of research. These detectors are compact, have high position resolution, are relatively fast, and can provide good energy sensitivity.

Section 1.1 describes the basics of semiconductor device physics, and explains how photodiodes function as radiation detectors. The 3D detector structure, which is a specialised variety of photodiode with electrode columns passing through the thickness of a silicon substrate, is discussed in section 1.2. This includes a description of the detector's fabrication, which requires specialised micromachining tools.

A synchrotron is a variety of particle accelerator, which can be used both for high-energy-physics experiments and for photon science. The basic operating principles of a synchrotron are described in section 1.3. 3D detectors could potentially be used in two main synchrotron applications. Firstly, they could be used as radiation-hard pixel detectors in future high-energy-physics experiments at the Super-LHC (section 1.4). Secondly, they could improve X-ray diffraction experiments at synchrotron light sources such as Diamond (section 1.5).

1.1 Operating principles of silicon photodiode detectors

This section discusses the aspects of semiconductor physics that make semiconductors useful as radiation detectors. In particular, it focuses on the p-n junction

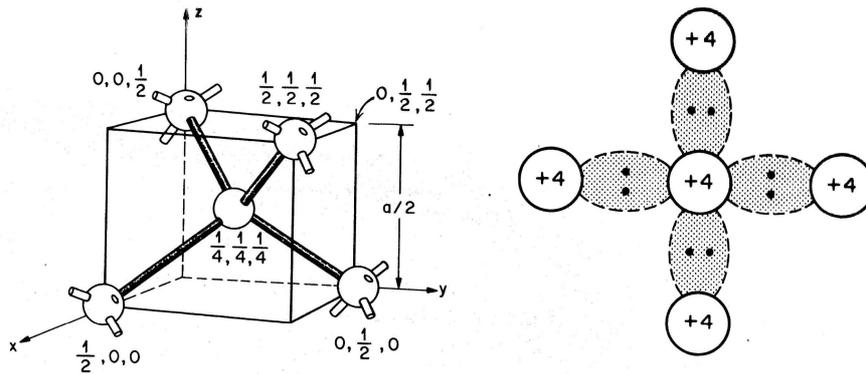


Figure 1.1: (a) The tetrahedral bonding structure in silicon. (b) A representation of the covalent bonding in silicon. Taken from Sze 1985, page 8 [2].

structure, and how this forms the basis of a photodiode detector. This section also describes the readout circuitry used to operate these photodiodes. Additionally, two alternative types of semiconductor detector, CCDs and CMOS, are briefly discussed for comparison. The material in this section is taken from textbooks by G. Lutz, 1999 [1], S. Sze, 1985 [2] and E. Yang, 1988 [3].

1.1.1 Basic semiconductor device physics

Semiconductor crystals and their band structure

Most semiconductors are single crystals, typically with either a diamond structure, like Si and Ge, or zinc blende structure, like GaAs and other compound semiconductors. Both of these structures have a tetrahedral bonding structure, with each atom sharing its outermost electrons (valence electrons) with its four neighbours, forming covalent bonds. The bonding structure of silicon, which has four valence electrons, is illustrated in Fig. 1.1.

In a gas, each atom has its own set of discrete energy states which can be occupied by electrons. However, when atoms are brought close together to form crystalline solids, the wavefunctions of the valence electrons begin to overlap. This interaction means that the energy levels split into a large number of energy states, each of which corresponds to an electron wavefunction which is spread out over many atoms. These states are very closely spaced in energy and momentum, and form virtually continuous energy bands, which may be separated by energy gaps.

The solid's electrical behaviour will depend on its band structure. The structure of a semiconductor, insulator and conductor are shown in Fig. 1.2. In a semicon-

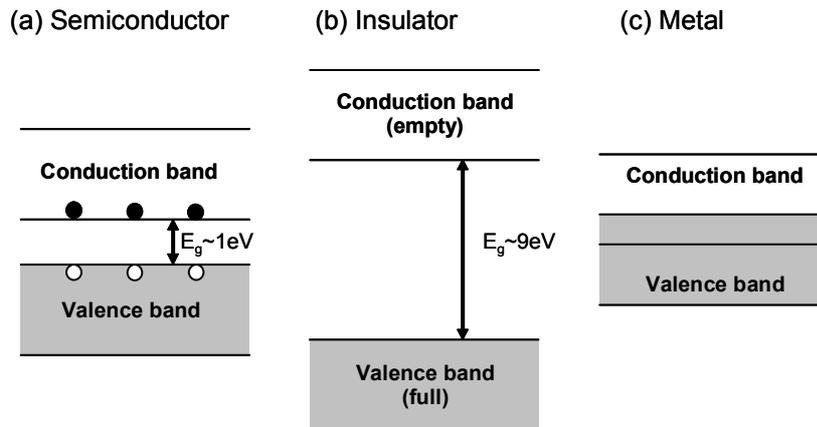


Figure 1.2: Band structures of (a) a semiconductor, (b) an insulator and (c) a metal.

ductor, at extremely low temperatures all the valence electrons will occupy the lower-energy valence band. There is an energy gap between this and the conduction band; in silicon, the width of the bandgap is 1.12 electron volts. The fully-occupied valence band has zero total momentum, so there can be no net current flow. However, with increasing temperature some of the electrons become thermally excited to the conduction band. Both the free electrons in the conduction band and the holes in the valence band are able to gain kinetic energy and momentum, and there can be current flow. Although the band structure arises from quantum mechanics, both the electrons and the holes behave quite similarly to classical particles.

An insulator has a similar band structure to a semiconductor, except that the wider bandgap means that there are virtually no free electrons and holes under normal conditions. In a conductor, there is no band gap between the occupied and unoccupied states, which means that the electrons are able to move freely even at low temperatures.

Intrinsic and doped semiconductors

The electrical properties of a semiconductor are strongly affected by the presence of impurities in the crystal.

An intrinsic semiconductor contains no (or, in practice, very few) impurity atoms. Since electrons and holes are produced in pairs by thermal excitation, their concentrations will be equal. So, where n is the electron concentration and p the hole concentration, we can write

$$n = p = n_i \quad (1.1)$$

where n_i is the intrinsic carrier concentration. The value of n_i at a particular temperature can be found by considering the density of states available in the conduction and valence bands, and the Fermi-Dirac distribution $f(E)$ which gives the probability of a level at energy E being occupied by an electron. In intrinsic silicon, the Fermi level E_F , which is defined as the energy at which the probability of occupation is 0.5, is close to the centre of the bandgap.

However, the presence of impurities in a semiconductor will create additional energy states, altering its behaviour. Group V elements such as phosphorus have 5 valence electrons rather than 4, and are referred to as donors. The extra electron is not involved in bonding, and occupies an energy state just below the edge of the conduction band. Due to the small energy difference, virtually all of these electrons become ionised due to thermal excitation. If we have a concentration of N_d donors, then provided that $N_d \gg n_i$ the free electron concentration will become $n = N_d$. The ionised donors also act as fixed positive charges in the silicon. This material is referred to as n-type silicon. Similarly, group III materials such as boron are “acceptors” which create an additional empty state just above the valence band. These become occupied by electrons from the valence band, so the hole concentration increases and the donors become negatively charged. If the acceptor concentration is N_a , then the hole concentration will be $p = N_a$. This material is referred to as p-type.

The carrier concentrations in silicon are affected both by the thermal generation of electron-hole pairs and their recombination. If, for example, n-type doping is introduced into the silicon, some of these extra electrons will recombine with holes, reducing the hole concentration. In thermal equilibrium, the rates of generation and recombination will be equal, and the carrier concentrations will obey the mass-action law:

$$np = n_i^2 \quad (1.2)$$

Also, the presence of dopants will alter the Fermi level. In n-type silicon, the Fermi level is close to the conduction band, since states in the conduction band have a higher probability of being occupied by electrons than in intrinsic material. Conversely, in p-type material the Fermi level is closer to the valence band edge.

Carrier transport

Free electrons and holes in a semiconductor are constantly undergoing random thermal motion, but this alone does not result in net current flow. Current flow can be

caused by two effects; drift in an electric field, and diffusion of carriers from regions of high to low carrier concentration.

Firstly, there is carrier drift. Naturally, if there is an electric field \mathbf{E} within the semiconductor, electrons and holes will be accelerated by it. The carriers scatter frequently, losing momentum in each collision, so overall the carriers will travel at an average drift velocity given by:

$$\mathbf{v}_n = -\mu_n \mathbf{E} \quad \mathbf{v}_p = \mu_p \mathbf{E} \quad (1.3)$$

where μ_n and μ_p are the electron and hole mobilities. At lower electric fields, the drift velocity is small compared to the thermal velocity, and the scattering rate is independent of the field strength. So, the drift velocity increases linearly with the field strength, and the mobilities will be roughly constant. However, as the field strength and drift velocity get high, scattering occurs more frequently, and eventually the velocity saturates. (This occurs around 10^7 cm/s in silicon for both electrons and holes.) Note also that the electron and hole mobilities may be different; for example, the electron mobility in silicon is about 3 times the hole mobility.

The second source of net carrier motion is diffusion. The thermal motion of a individual carrier is random. However, if neighbouring regions of the device have different carrier concentrations, then the probability of carriers moving from the high-concentration region into the lower-concentration region is higher than the probability of the reverse process. The fluxes of charge carriers due to the carrier concentration gradients are given by:

$$\mathbf{F}_n = -D_n \nabla n \quad \mathbf{F}_p = -D_p \nabla p \quad (1.4)$$

where D_n and D_p are the diffusion constants for electrons and holes.

Overall, the electron and hole currents will be given by:

$$\mathbf{J}_n = q\mu_n n \mathbf{E} + qD_n \nabla n \quad \mathbf{J}_p = q\mu_p p \mathbf{E} - qD_p \nabla p \quad (1.5)$$

Carrier continuity and Gauss' law

The relationship between current flow and the generation and recombination of carriers is described by the carrier continuity equations. The rate of change in electron or hole concentration in an infinitesimal volume is given by the net rate of

generation, plus the net rate at which carriers enter from outside the volume. So:

$$\frac{\partial n}{\partial t} = G - R + \frac{1}{q} \nabla \cdot \mathbf{J}_n \quad \frac{\partial p}{\partial t} = G - R - \frac{1}{q} \nabla \cdot \mathbf{J}_p \quad (1.6)$$

where G is the carrier generation rate and R is the recombination rate.

The behaviour of the electromagnetic field is governed by Maxwell's equations. Typically, carrier currents are not large enough to generate significant magnetic fields. Also, since semiconductor devices are small, and carrier drift velocities are low compared to the speed of light, the time for changes in the electric field to propagate can be ignored. So, we only need to consider Gauss' law:

$$\epsilon_s \nabla \cdot \mathbf{E} = \rho \quad (1.7)$$

where ρ is the charge density and ϵ_s is the permittivity of silicon. The permittivity is given by the electric constant ϵ_0 multiplied by the relative permittivity of silicon, which is 11.8.

The left-hand side can be re-written in terms of the electrostatic potential ψ , where $\mathbf{E} = -\nabla\psi$, and the charge density will be determined by the concentrations of electrons, holes and ionised dopants. This leads to Poisson's equation:

$$\epsilon_s \nabla^2 \psi = q(n - p + N_a - N_d) \quad (1.8)$$

1.1.2 The p-n junction

In the previous section, the concept of n- and p-type doping was introduced. The most basic semiconductor device structure that can be produced is a junction between p-type and n-type material. This is used in many applications, including radiation detection.

P-n junction electrostatics

If we have separate pieces of p and n-type material, the p-type will contain a high concentration of holes, and equal numbers of ionised acceptors, giving overall charge neutrality. Likewise, n-type material has a high electron concentration and positively charged donor atoms.

At a junction between the two types of material, the differences in carrier concentration will lead to diffusion of holes from the p-type material into n-type, and electrons from the n-type to the p-type. These electrons and holes will tend to recombine with each other, creating a "depletion region" with much lower carrier

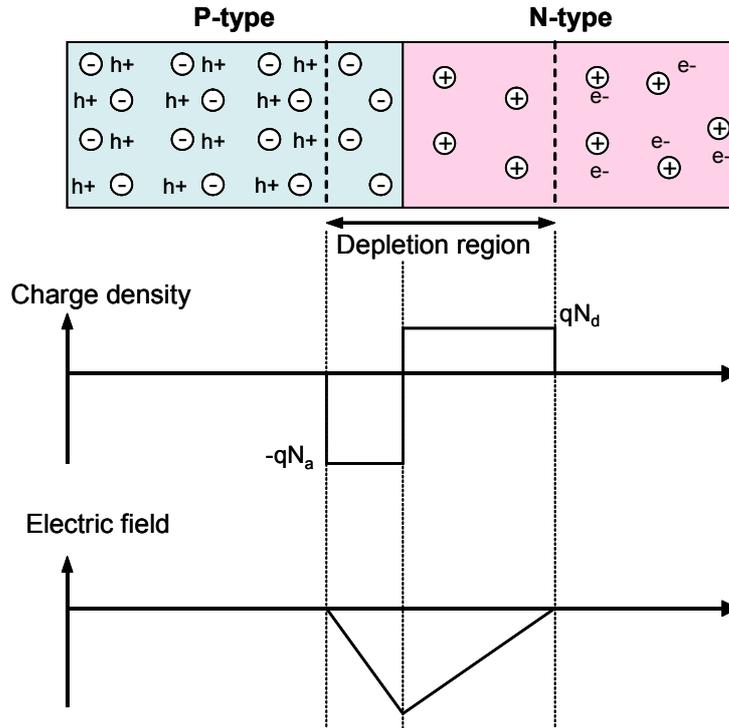


Figure 1.3: Diagram of a p-n junction in equilibrium, showing the charge distribution and electric field in the depletion region.

concentrations than in the bulk material. Within the depletion region, there will be uncompensated negatively-charged acceptors in the depleted p-type material, and positively-charged donors in the n-type—see Fig. 1.3. This space charge will produce an electric field in the depletion region, which will cause carrier drift in the opposite direction to the diffusion. If no external voltage is applied, then the device will reach a state of equilibrium where the net current flow is zero. The fixed charge in the depletion region will mean that the device has a built-in potential difference.

Note that when the device is in equilibrium, with no external bias, the Fermi level E_F will be equal in both p- and n-type regions. (If the Fermi levels were unequal, then net current flow would occur.) Using this, it is possible to calculate the built-in voltage V_{bi} :

$$V_{bi} = \frac{kT}{q} \ln \frac{N_a N_d}{n_i^2} \quad (1.9)$$

where T is the temperature and k is Boltzmann's constant. This also implies that the built-in voltage will be smaller than the bandgap, which is 1.12V in silicon.

Reverse biasing the p-n junction

Generally, semiconductor radiation detectors work using reverse-biased p-n junctions, where a positive bias is applied to the n-type region and a negative bias to the p-type region. Under these conditions, the external bias has the same polarity as the built-in potential of the depletion region. Since the depletion region contains few charge carriers, the external bias produces very little current flow. Instead, electrons in the n-type and holes in the p-type will be attracted away from the edges of the depletion region, causing it to widen. Consider a simple, one-dimensional p-n junction, with a reverse bias V_R applied to it. The total positive charge from donors in the depleted n-type region will match the negative charge in the depleted p-type, and the total potential difference across the junction will be $(V_{bi} + V_R)$. Using the Poisson equation, the depleted width x_d is found to be:

$$x_d = \sqrt{\frac{2\epsilon_s(N_a + N_d)}{qN_aN_d}} (V_{bi} + V_R) \quad (1.10)$$

Since the total charge in the depletion region is zero, if the doping concentrations are unequal the depletion region will extend further into the lightly-doped region, as indicated in Fig. 1.3. In particular, if (say) $N_a \gg N_d$, then the depletion region will appear almost entirely in the lightly-doped material, and the width of the depletion region will be proportional to $\sqrt{V_R/N_d}$ (ignoring the relatively small built-in voltage).

The capacitance of the depleted p-n junction, which is a function of voltage, will be given by $C = dQ/dV$, where the charges stored in the two sides of the depletion region are Q and $-Q$. It can be shown that the capacitance per unit area depends quite simply on the width of the depletion region:

$$C = A \frac{\epsilon_s}{x_d} \quad (1.11)$$

1.1.3 P-n junction photodiodes

When photons and high-energy particles deposit energy in a semiconductor, this energy can excite electrons from the valence band to the conduction band, creating free electron-hole pairs. The specific behaviour of different forms of ionising radiation is discussed later. If we have a reverse-biased p-n junction, then in a steady state the current flow in the junction is low. However, when excess electron-hole pairs have been generated, these will be swept out of the depletion region, producing a current. This allows the structure to be used as a radiation detector.

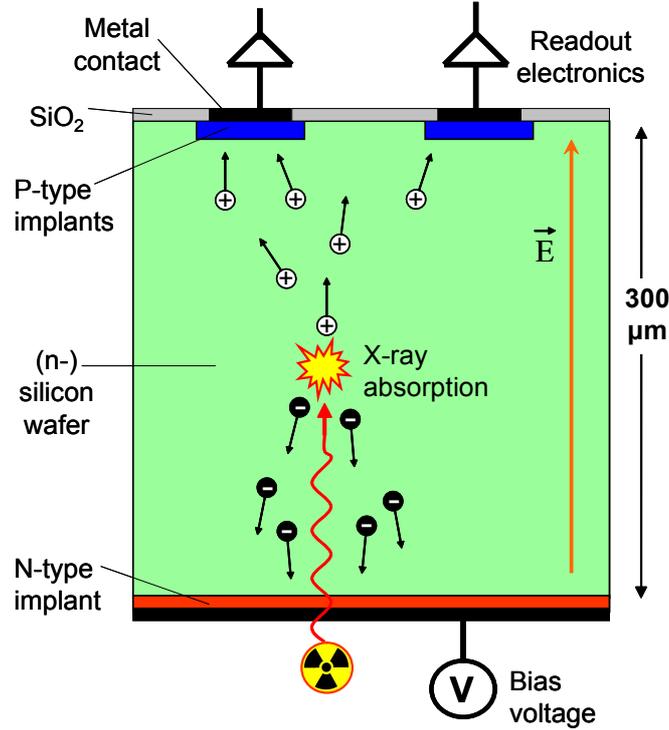


Figure 1.4: Structure of a typical p-i-n photodiode, showing charge collection from an X-ray hit.

A basic photodiode structure is shown in Fig. 1.4. The photodiode is fabricated on a silicon wafer with very low doping—a high-quality detector wafer will have a concentration of order 10^{12}cm^{-3} . In this particular example, an n-type wafer is used, but devices can also be fabricated on p-type. On the front and back surfaces, there are narrow, highly-doped p- and n-type regions, with thicknesses of a few microns and peak doping of over 10^{18}cm^{-3} . This is referred to as a p-i-n structure (p-type, lightly doped, n-type). Metal electrodes, typically aluminium, make contact with these heavily doped regions. To achieve position sensitivity the doping and electrodes on the front surface will be segmented, forming an array of p-n junctions, and each electrode will be connected to readout electronics. The back surface contact is used to apply a reverse bias, so that the full volume of the lightly-doped substrate is depleted. Using Equation 1.10, the full depletion voltage of the detector will be:

$$V_{dep} = \frac{qN_d x_d^2}{2\epsilon_s} \quad (1.12)$$

For example, a typical detector with a substrate thickness of $300\mu\text{m}$ and 10^{12}cm^{-3} doping will become fully depleted at 70V.

When electron-hole pairs are generated in the depletion region by ionising radiation such as X-rays, the electric field will cause the electrons to drift to the n-type electrode and the holes to drift to the p-type electrodes. Since the carriers are charged, as they move through the device their field will induce current flow in the electrodes, producing a measurable signal. For example, as a hole approaches the p-type electrode, it will repel the other holes in this region, resulting in current flow into the readout electronics connected to the electrode. These induced currents can be calculated quite efficiently using Ramo's theorem [4]. Suppose we have a single electron-hole pair, generated in a device with a single n-type back contact and segmented p-type readout electrodes. Ramo's theorem shows that, although transient currents will be generated in the different p-type electrodes as the carriers move, the total charge signal (i.e. the integral of the current) will be q on whichever p-type electrode the hole arrives at, and zero on the others.

When a particle hits a detector, the carriers it generates will often be shared between multiple electrodes. Obviously, this will occur if the initial distribution of charge carriers generated by the particle covers an area containing more than one electrode. However, after the charge carriers are generated they will also spread outwards by diffusion, increasing the amount of charge sharing. Naturally, charge sharing is more likely to occur when the spacing between the segmented readout electrodes is small. Charge sharing will also increase when the time to collect the carriers is long, since more diffusion can take place. In particular this means that thicker detectors, where the charge has to travel further to reach the electrodes, will have more charge sharing.

Interactions between ionising radiation and semiconductors

The quantity and distribution of the electron-hole pairs generated by ionising radiation will depend on the type of radiation. In general, semiconductor detectors require a relatively low input of energy to create an electron-hole pair—for example, the mean energy in silicon is 3.6eV, compared to 30eV or so for gas detectors. This, combined with their reasonably high density, means that these detectors can be very thin, compact and highly segmented, while still achieving good sensitivity.

Photons will primarily interact with semiconductors by the photoelectric effect, where the photon is absorbed in a single interaction and an electron gains its energy. To generate an electron-hole pair, the photon energy must exceed the width of the bandgap. So, for example, since silicon has 1.12eV bandgap it is transparent to infrared light at wavelengths of beyond 1100nm. Photons of visible light will produce a single electron-hole pair. X-ray photons, however, have energies of thousands

of electron volts, and will produce large numbers of electron-hole pairs in a small spatial region where the photon was absorbed. The absorption of these photons is probabilistic, so if we have a beam of X-rays the flux reaching a depth d will be described by a decaying exponential, $F = F_0 \exp(-d/l)$. In the keV energy range, the absorption length l generally increases with energy. For example, in silicon 12keV photons have $l=116\mu\text{m}$, and a $300\mu\text{m}$ substrate will absorb 92% of the photons, whereas at 20keV only 46% will be absorbed. So, this limits the energy range these silicon detectors can be used for. Additionally, at higher energies other interactions such as Compton scattering will occur. Other semiconductors with higher atomic numbers can provide better absorption at higher photon energies.

Unlike a photon, a charged particle will undergo a series of Coulomb interactions as it passes through silicon. These can be regarded as a series of collisions between the particle and the electrons in the silicon. Additionally, the charged particles can transfer energy to the crystal lattice itself, potentially displacing atoms from their lattice sites—the resulting radiation damage is considered in section 2.3.

When a charged particle passes through matter, its rate of energy loss with distance is given by the Bethe-Bloch formula [5], assuming that the particle is not deflected by the collisions:

$$\frac{dE}{dx} = 2\pi N_0 r_e^2 m_e c^2 \rho \frac{Z}{A} \frac{z^2}{\beta^2} \left[\ln \left(\frac{2m_e \gamma^2 v^2 W_{max}}{I^2} \right) - 2\beta^2 - \delta - 2\frac{C}{Z} \right] \quad (1.13)$$

The terms in this equation are:

N_0 : Avogadro's number	r_e : classical electron radius
m_e : electron mass	I : mean excitation potential
Z : atomic number of absorber	A : atomic weight of absorber
ρ : density of absorber	
$\beta = v/c$ of incident particle	$\gamma = 1/(\sqrt{1 - v^2/c^2})$
δ : density correction	C : shell correction
z : charge of incident particle in units of e	
W_{max} : maximum energy transfer in a single collision	

At non-relativistic particle energies, the energy loss rate is inversely proportional to the energy. This is because as the particle's velocity (βc) increases, the time over which the particle's Coulomb force acts on the electrons in the material is reduced, resulting in a smaller energy transfer. In turn, this will mean that as the particle loses energy to the medium, its loss rate will get higher and higher.

However, at relativistic energies, where the particle's speed approaches c , the rate of energy loss reaches a minimum. For electrons, for example, this point is reached at energies of about 1MeV. Beyond this point, the rate of loss increases slowly with energy. In a thin material, the total energy loss will be small compared to the total particle energy. In practical terms, this means that the high-energy particles produced in colliders like the LHC (see section 1.4) will only lose a small fraction of their energy as they pass through a silicon detector, and will produce nearly uniform carrier generation along their path. The density of generated charge will be nearly independent of the particle energy, at about 80 electron-hole pairs per micron for a singly-charged particle [6].

Since these interactions between charged particles and the semiconductor are statistical in nature, the total energy deposited by each particle will vary. However, the energy distribution produced over a large number of events is predictable, and follows a Landau distribution. This is discussed in section 3.3.

Aside from these direct detection processes, it is possible to cover a silicon detector with converter material to allow it to detect other particles. For example, a scintillator will absorb high-energy X-rays and gamma rays and emit visible photons, which are then measured by the silicon. Similarly, a suitable converter material such as ${}^6\text{Li}$ will absorb neutrons and decay to high-energy ${}^3\text{H}$ and ${}^4\text{He}$, which can be detected by the ionisation they produce.

Photodiode fabrication

The processes used to fabricate these detectors are broadly similar to those used to produce integrated circuits and other silicon devices. As a result, it is possible to produce large arrays of extremely small photodiodes. In practice, the photodiode size will be set by design considerations (such as the required resolution) and the limitations of the readout electronics, as discussed later.

Fabricating a very simple photodiode array only really requires three processes: doping the entire back surface of the detector; producing segmented doped regions on the front side of the detector to form an array of p-n junctions; and adding segmented metal contacts so that the elements can be connected to readout electronics. In practice, the fabrication tends to be more complicated—for example, extra doped regions are required to carry out extra functions such as electrode isolation, and multiple metal layers may be needed for the readout connections. However, this can all be achieved by repeating the same basic fabrication processes.

If we want to dope specific regions of the surface of a silicon wafer, we first need to use a masking process. First, the silicon is covered in a layer of silicon dioxide,

either by oxidising the surface or depositing the silicon dioxide using chemical vapour deposition. Then, a layer of light-sensitive polymer called photoresist is added to the oxide surface. The photoresist is exposed to light through a patterned mask, and then developed to remove the exposed regions (or unexposed regions, depending on the resist type). This leaves a patterned layer of resist. Then, the wafer is placed in hydrofluoric acid. This etches away the silicon dioxide in the unmasked regions, exposing the silicon beneath.

To create the p-n junctions, we need to add dopants to the surface of the silicon. This is normally done by ion implantation or diffusion. In ion implantation, ions of the appropriate element (e.g. boron for p-type) are generated, then accelerated onto the surface of the wafer by an electric field. The ions scatter from the silicon atoms, and come to rest within the crystal. After the implantation the silicon is heated to anneal out the damage caused to the crystal lattice. The alternative method, diffusion, involves heating the silicon in the presence of the dopant, which allows it to diffuse into the crystal. This process is simpler, but gives less control over the doping profile. During both of these processes, the oxide layer will prevent the dopant from reaching the masked regions.

Typically, aluminium is used to form the contacts. A slightly different masking process is used here. After the photoresist is deposited and developed as normal, the entire surface is covered in aluminium using aluminium vapour or chemical vapour deposition. Then, the photoresist is dissolved with a suitable solvent, removing any aluminium overlying the resist in the process.

1.1.4 Readout electronics

During operation, each electrode of a photodiode detector will be connected to an input channel on a readout chip. The readout chip carries out a few major functions. Firstly, it amplifies the signal from each element of the detector, to allow it to be measured accurately. Secondly, it can perform basic signal processing on the data. Frequently, the processing will reduce the quantity of data, for example by applying a threshold to each signal to give a binary value. This allows the data to be transferred from the readout chip to the external electronics at a reasonable rate. The readout chip will also temporarily store the data, so that it can be passed on to the external electronics when it is required. As an example, section 5.2 describes the Medipix2 readout chip, which is designed for X-ray detection.

Strip, pixel and pad detectors

Physically, there are three main configurations that can be used. In a strip detector, the p-n junctions on the readout chip consist of long, narrow strips of dopant covered with a metal layer. These strips can have a spacing of tens of microns—for example, the 3D strip detectors tested in section 4.3.3 have a spacing of $80\mu\text{m}$ —and they can be many centimetres long. The readout chip is placed at the edge of the sensor, and contacts at the end of each strip are connected to channels on the readout chip using wire bonds. The detector provides one-dimensional position resolution over a reasonably large area, and if a particle passes through multiple strip detector layers with different orientations the data can be combined to calculate its 3D path. However, if multiple particles pass through the strip detectors simultaneously, the results may be ambiguous.

In pixel detectors, the p-n junctions form a 2-dimensional array. For example, the Medipix2 pixel sensor discussed in section 5.2 has a 256×256 array of $55\mu\text{m}$ by $55\mu\text{m}$ pixels [7]. This 2-dimensional position resolution is vital for many applications. However, a very large number of connections are required to read out all the pixels. This is achieved by giving the readout chip the same pixel structure as the sensor, with each pixel containing the electronics required to process one readout channel. Each pixel on the sensor and the readout chip has a metal pad, and the two chips can be connected together face-to-face with metal bumps as shown in Fig. 1.5. This process is called bump bonding, and the finished assembly is referred to as a hybrid pixel detector. The size of the bump bonds and factors such as alignment accuracy will determine the minimum possible pixel size.

Pad detectors use only a single photodiode. As a result, no specialised readout chip is necessary—a readout chain can be built using individual components such as amplifiers, analogue-to-digital converters etc. These detectors can be used for applications such as spectroscopy. Additionally, position resolution can be achieved by moving the photodiode, though scanning a large area like this will be slow. Due to their simplicity, pad detectors are also useful as test devices.

Signal processing

The first function of the readout chip is to amplify the signal from each strip or pixel.

As shown by Ramo's theorem, the current signal produced at each electrode when an electron-hole pair is collected will vary over time, depending on the hit location and device structure, but the total charge received will always be one electron. So,

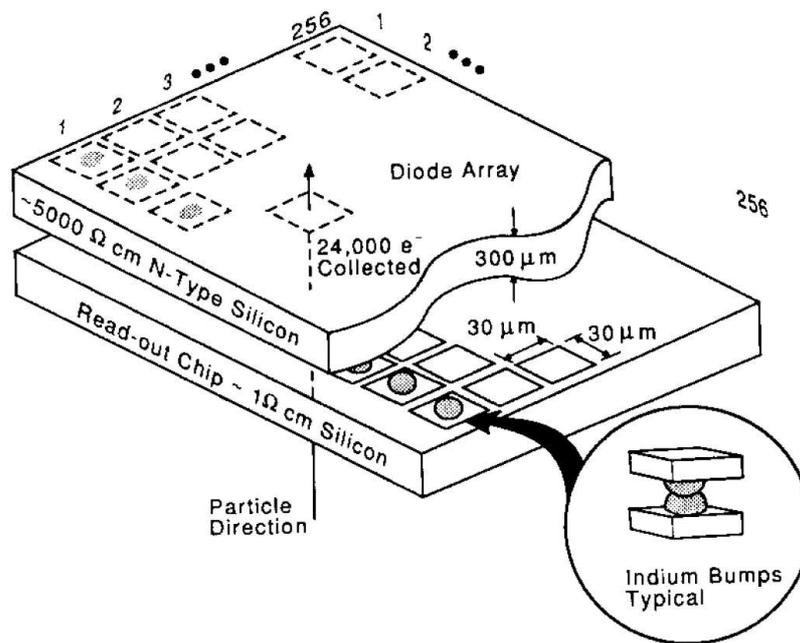


Figure 1.5: Diagram of a bump-bonded hybrid pixel detector. Reproduced from Shapiro 1989 [8] with permission from Elsevier.

a charge-sensitive preamplifier is normally used, as shown in Fig. 1.6. In this circuit almost all the current from the detector, i_{in} , will flow into the feedback capacitor C_f and charge it up. The output voltage will be given by $V_{out} = -Q_i/C_f$ where Q_i is the total charge. After the current pulse, the capacitor will discharge through appropriate circuitry—typically a transistor biased to behave as a high-resistance channel, represented by R_f .

If the detector is connected directly to the preamplifier (which is referred to as a “DC coupled” configuration) any leakage current will flow into the capacitor. If the capacitor cannot discharge quickly enough, then it may charge up until the output of the preamplifier saturates, preventing it from working. This is a problem for radiation-damaged detectors, which have higher leakage currents as discussed in section 2.3.2. To deal with this, some detectors use an “AC coupled” configuration. In this configuration, the detector channel is connected to the readout chip via a capacitor, and connected to ground via a resistor. This means that the leakage current from the reverse-biased detector will pass through the resistor, but fast signals due to hits on the detector will pass through the capacitor to the readout electronics, due to the capacitor’s low impedance at high frequency. Commonly, the capacitor is formed by an oxide layer sandwiched between two metal layers on the detector, and the resistors are also fabricated on the detector using polycrystalline

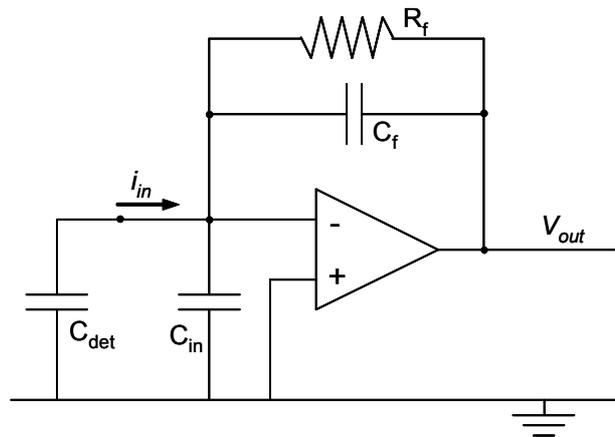


Figure 1.6: The charge-sensitive preamplifier circuit, commonly used in readout chips.

silicon (polysilicon).

After the preamplifier, additional shaping amplifiers can also be used to increase the amplification and to filter the signal. Filtering very high-frequency components of the signal will reduce the noise. Filtering low-frequency components will make the signal pulse return to zero more quickly. If two hits occur on a detector in quick succession, a pile-up effect can occur where the second hit is added to the tail of the first hit signal, distorting its amplitude. So, giving the detector a faster return to zero will increase the maximum hit rate. However, if the charge collection in the device is slow, the low-frequency filtering can reduce the readout signal—this is referred to as ballistic deficit.

A major aim of the amplifier design is to maximise the signal-to-noise ratio. There are two separate problems to be considered; interference, where the equipment picks up unwanted signals generated by nearby apparatus (e.g. crosstalk between two current-carrying wires) and noise sources which are intrinsic to the detector and readout electronics. After amplification, any additional interference or noise added by the electronics should be small compared to the amplified signal, so effects occurring before the input of the amplifying stage are most important. To minimise the interference, the preamplifier should be placed as close to the detector as possible, and where possible shielding should be used.

The system's intrinsic noise comes from three major sources: thermal voltage noise, due to random thermal motion of electrons in the circuitry; low-frequency noise, due to trapping and detrapping of charge carriers by defects; and current noise,

due to statistical variations in the current flow into the charge-sensitive amplifier.

Both the thermal and low-frequency noise effectively act as voltage sources at the input of the amplifier. However, since the preamplifier is sensitive to charge, the noise at the output of the amplifier will be affected by the impedance Z seen at the input, since the resulting noise current will be $I=V/Z$. This means that the noise is proportional to the total capacitance at the input of the amplifier, i.e. the combined capacitance of the amplifier (C_{in}) and the detector (C_{det}). So, limiting the capacitance of the photodiode detector is an important design consideration.

The statistical variation in the current flow is proportional to \sqrt{i} . Even if a detector is AC-coupled, high-frequency fluctuations in the leakage current will pass through the capacitor to reach the preamplifier, so high leakage currents will still lead to higher noise.

The later stages of signal processing vary depending on the application. The signal processing electronics in a strip detector designed for particle tracking in high-energy physics experiments, and a pixel detector designed for X-ray detection, are described in sections 3.4 and 5.2 respectively.

1.1.5 Further design considerations

So far, this section has described the basic working principles of photodiode detectors. However, a photodiode detector also needs to be designed to minimise certain problems.

Avalanche breakdown

When electrons and holes drift in an electric field, they continually gain energy from the field, then lose it in collisions. If the electric field is high enough, some of the collisions become energetic enough to excite new electron-hole pairs. These new carriers can in turn excite yet more carriers, resulting in a multiplication of the current. Some devices, such as silicon photomultipliers, actually exploit this effect in order to amplify weak current signals. However, in a photodiode this effect means that when the bias applied to the detector is too large, avalanche breakdown will occur and large currents will flow through the detector. This will prevent any signal from being measured, and can also damage the readout electronics or the detector itself.

In silicon, breakdown effects tend to occur when the maximum electric field in a device reaches around 3×10^5 V/cm. In a detector, there tend to be regions where the field is particularly high, for example at the edges of the doped implants on the

front surface, and in particular around the edges of the detector chip. So, the design should be optimised to reduce the field in these regions.

Oxide charge, and electron accumulation at the silicon-oxide boundary

The silicon dioxide layers on a detector contain a certain amount of fixed positive charge. This is because when electron-hole pairs are generated in silicon dioxide by ionising radiation, the electrons are fairly mobile and have a high chance of diffusing out of the oxide, whereas holes have low mobility and can be trapped by defects in the oxide. The oxide charge density increases with the dose of ionising radiation, and saturates when all the defects are occupied by holes. This effect is produced by all forms of ionising radiation, and at the fluences considered here saturation is reached very quickly [9]. Typical oxide charge densities are $1\text{--}4 \times 10^{11} \text{cm}^{-2}$ before irradiation, and $1 \times 10^{12} \text{cm}^{-2}$ after saturation [10].

The positive charge in the oxide layer will attract a layer of electrons to the silicon-oxide interface, which will affect the electrical behaviour in this region. In a detector with p-type readout electrodes, a higher-field region will be created where each p-type implant meets the electron layer. In a detector with n-type readout electrodes, the electron layer will short the n-type implants together, leading to unwanted signal sharing. To counteract this, these devices use additional p-type implants to overcompensate the electron layer. The two common approaches are to use a p-spray, which is a uniform p+ layer across the entire surface (with much lower concentration than the n+ electrodes, of course), or a p-stop, which consists of strips or rings of p+ implant between the n+ electrodes [11]. Once again, having these p-type implants close to the n-type electrodes can lead to high-field regions.

The effects of using these forms of electrode isolation with different 3D detector structures are simulated in later chapters.

Edge effects and guard structures

A silicon wafer may contain many detector devices, each of which will consist of an array of photodiodes. After fabrication, the wafer must be diced to separate the chips, and this is typically done using a diamond saw. The saw-cut edges will contain many defects and dangling bonds. These defects make the edge more conductive than the bulk material, and will also allow increased electron-hole pair generation. (The effects of defects in silicon are discussed in section 2.3.) If the depletion region from a photodiode reaches the edge, then there can be a large current flow from the edge to the junction, preventing it from working properly. Additionally, since the

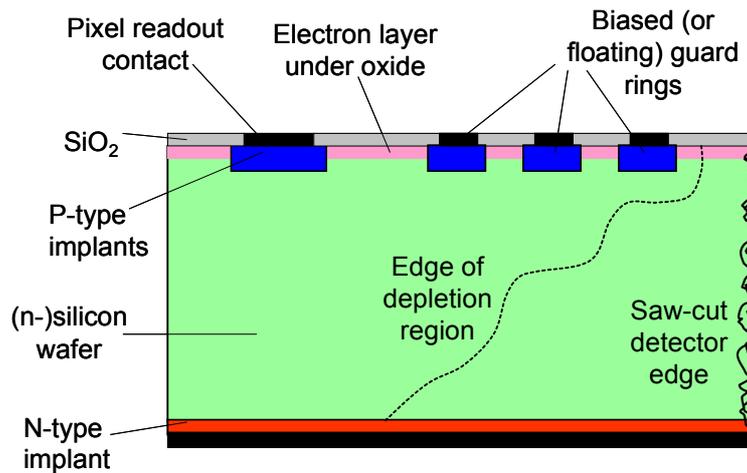


Figure 1.7: A typical guard ring structure at the edge of a detector with p-type readout electrodes. Generally at least one guard ring is biased, so that it will collect any surface leakage current.

edge provides a path from the biased back surface to the electron layer (or p-spray) on the front, this can lead to high-field regions around the edge of the photodiode array. To prevent this, there is typically an insensitive area at the edge of the chip, generally at least $100\mu\text{m}$ wide, containing guard ring structures [12]. These are rings of doped implant of the same type as the readout electrodes, as shown in Fig 1.7. Generally, at least one of the guard rings will be biased, so that any surface current will flow through the guard ring without reaching the detector array. The guard rings also prevent the depletion region from reaching the edge, and reduce the maximum field at the surface by allowing the potential to drop gradually towards the edge.

1.1.6 CCDs and CMOS detectors

Although this thesis deals with photodiode detectors, two other common silicon detector technologies are described here for comparison.

Charge coupled devices (CCDs)

In a highly segmented pixelated photodiode detector, each pixel has its own readout connection on the front surface, and carriers generated within each pixel are immediately swept to the electrodes and produce a signal in the readout electronics. The front surface of a CCD, on the other hand, is almost entirely covered in a oxide layer, over which there is an array of metal electrodes, forming a series of metal-oxide-semiconductor structures. Every third electrode is kept at the same

potential. The biased electrodes create a series of potential minima near the surface of the silicon, which will store up any charge carriers generated by radiation. At the edge of the pixel array, there are one or more readout electrodes. By cycling the potentials on the three sets of electrodes, it is possible to shift the collected carriers from pixel to pixel, so that the charge collected in each pixel can be sequentially read out at the readout electrode [13].

The big advantage of this structure is that no separate readout chip is required, making the device much simpler and cheaper. Since the metal-oxide-semiconductor structures in each pixel are fairly simple, a small pixel size can be achieved. Additionally, the readout electrode(s) can be made small, to reduce the capacitance, and large numbers of pixels will be read out by the same amplifier. So, both the noise and the pixel-to-pixel variation will be small. As a result, CCDs are very commonly used in consumer digital cameras, and in science and industry.

The main disadvantages of a CCD are that the readout process is slow, each pixel can only store a limited quantity of charge, and the depleted thickness is generally smaller than in a photodiode. The slow collection also means that the device has poor radiation hardness. Additionally, by integrating the generated charge we lose the opportunity to do more sophisticated signal processing. CCDs cannot be used for many particle physics applications, where a high readout speed is needed to obtain event-by-event information, and high radiation tolerance is required. Likewise, they are unable to exploit the very high brightnesses of modern synchrotron X-ray sources.

CMOS pixel sensors

In a photodiode hybrid pixel detector, the readout electronics for each pixel are contained on a separate readout chip. In a CMOS pixel sensor, the charge is still collected by a photodiode, but the readout electronics are fabricated within each pixel on the sensor wafer. This is done using the CMOS (Complementary Metal Oxide Semiconductor) technology used in most integrated circuits [14]. At the edge of the chip there are a series of readout channels, and the pixels can be selectively read out through these—typically, an entire pixel array will be read out one row at a time. Minimally, each pixel in a CMOS detector needs to have a photodiode, and a transistor to select the pixel for readout. In an active pixel sensor, the performance is improved by including a buffer in each pixel, which stores the collected charge and, when the pixel is selected, produces an amplified readout current proportional to this charge.

CMOS pixel sensors can be cheaper and more compact than CCDs, and can have lower power consumption. As a result, they have been widely used in camera phones.

Also, it is possible to include more sophisticated functionality within the pixels. For example, the readout rate could be greatly improved by selectively reading out the regions of the detector which contain a useful signal. However, these sensors tend to suffer from a great deal of pixel-to-pixel variation in their response. Including more electronics in a pixel can reduce its sensitive volume. Also, while they are more radiation-tolerant than CCDs due to their faster collection, they are still less radiation-tolerant than standard photodiodes. So, although CMOS detectors are promising, most experiments still use CCD or photodiode detectors.

1.2 3D detectors

1.2.1 The 3D detector structure

A 3D detector is a variety of photodiode detector where the p-i-n structure is formed by columns of p- and n-type material passing through the thickness of the substrate. The structure was proposed by S. Parker et al. in 1997 [15], and is illustrated in Fig. 1.8. One set of electrodes are connected separately to readout electronics, like in a standard planar photodiode, and the other set are connected together and used to bias the device. Typically, the electrodes form a diamond pattern, producing a series of square pixels with a readout column in the centre and four bias columns at the corners, as shown in Fig. 1.9. When the device is depleted, electron-hole pairs generated by ionising radiation will be swept horizontally to the neighbouring electrodes.

In a “full 3D” detector, shown here, the electrode columns pass through the full thickness of the substrate, and are connected out on the front surface. However, the detectors tested in this thesis use two alternative configurations. In the double-sided 3D detectors produced by CNM (Chapter 4) the two sets of columns are fabricated from opposite sides of the silicon, and do not pass through the full substrate thickness. In the single-type-column 3D detectors from FBK-IRST (Chapter 6) there is only one set of columns, which extend from the front surface and do not pass through the full substrate thickness, and the bias contact is simply a doped layer on the back surface. The fabrication process is similar in each case. Each set of doped columns is produced by etching deep holes in the silicon, then filling these with doped polysilicon, as described later.

The main advantage of the 3D structure is that, unlike in a planar detector, the spacing between p- and n-type electrodes is not determined by the substrate thickness. So, it becomes possible to achieve an electrode spacing of as little as

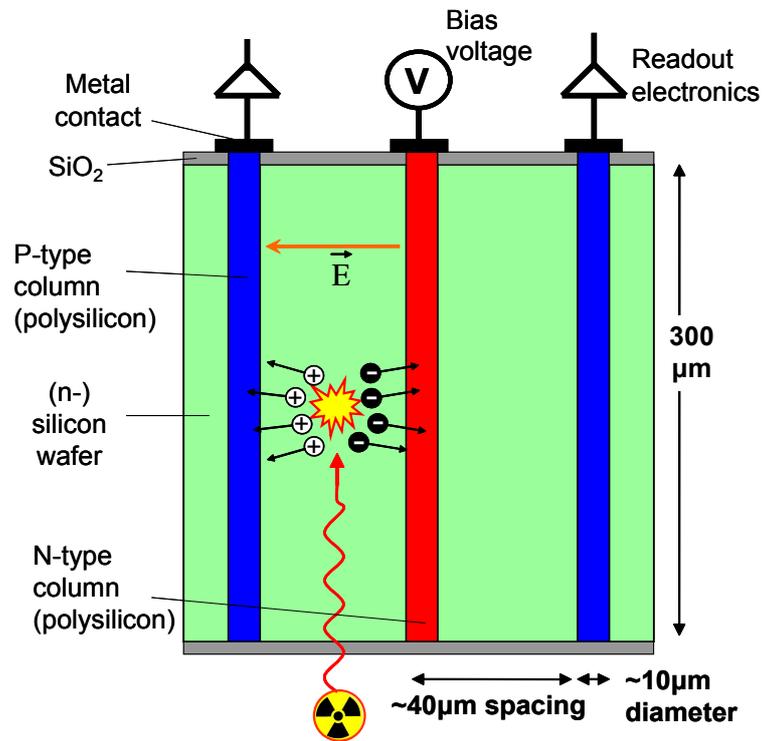


Figure 1.8: Structure of a full-3D detector, showing charge collection from an X-ray hit.

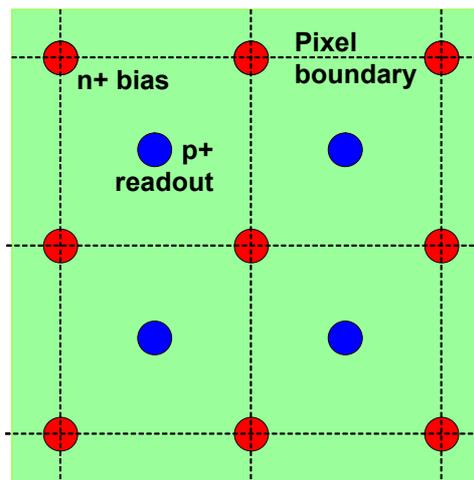


Figure 1.9: Typical electrode layout of a 3D detector. Dotted lines indicate the edges of the pixels.

$30\mu\text{m}$, while still maintaining a sensitive thickness of a few hundred microns. Firstly, the reduced electrode spacing dramatically reduces the depletion voltage. As shown by Equation 1.12, the depletion voltage in a planar detector is proportional to the square of the electrode spacing. In a 3D detector, the situation is more complicated. The depletion region around each column will initially grow as a cylinder, but the electrodes are arranged in a square grid. If we simply consider a cylindrical depletion region growing around a single column we get the following expression:

$$V_{dep} = \frac{qN_d}{2\epsilon_s} \left(r_d^2 \left[\ln \left(\frac{r_d}{r_c} \right) - 0.5 \right] + 0.5r_c^2 \right) \quad (1.14)$$

where r_d is the radius of the depleted region, and r_c is the radius of the doped column itself. So, the 3D detector's full depletion voltage is roughly proportional to the square of the electrode spacing. Hence, by reducing the electrode spacing by up to a factor of 10, the 3D detector structure can reduce the depletion voltage by up to a factor of 100!

Similarly, the time to collect charge carriers is determined by the collection distance, which is dramatically reduced by the 3D structure, and the field strength, which will be high in a 3D detector (at a given bias). So, the 3D detector achieves fast charge collection, which is particularly important for radiation hardness as discussed in section 2.3.

The 3D structure also reduces charge sharing between adjacent pixels [16]. Firstly, the fast collection time means that the carriers have less opportunity to diffuse outwards before being collected. Secondly, the 3D detector structure makes the carriers drift horizontally to the columns, keeping them away from the pixel boundaries. In contrast, in a planar detector the carriers drift vertically through the detector towards the electrodes, and are free to diffuse horizontally across the pixel boundaries. The 3D detector's reduced charge sharing isn't an advantage in every application. When charge is shared between pixels (or strips), it is possible to determine the hit position with sub-pixel resolution by comparing the relative signal sizes on the pixels, whereas with zero charge sharing it is only possible to determine which pixel was hit. However, reduced charge sharing does improve the chances of getting an unambiguous hit on the detector above the noise level, and improves energy measurements.

It is also possible to add an "active edge" electrode to the edge of a 3D detector chip. This reduces the insensitive area at the edge. This is discussed more fully in section 1.2.3.

Aside from the complexity of their fabrication process—see the following section—

3D detectors have a couple of disadvantages. Firstly, any electron-hole pairs produced inside the electrode columns themselves will produce low or zero signal [17]. This occurs because the electric field will be negligible within the doped polysilicon (due to the high carrier concentration and high conductivity), and a signal will only be produced by the small number of carriers which diffuse out of the polysilicon into the depletion region. Furthermore, some 3D detector designs partially fill the column with silicon dioxide, which is not sensitive. Another disadvantage is that the small electrode spacing can lead to a high detector capacitance, which will increase the noise in the readout signal (see section 1.1.4).

1.2.2 Fabrication of 3D detectors

The basic fabrication methods that are used to produce planar photodiodes can only alter the regions around the surfaces of a silicon wafer. To etch and re-fill the electrode columns in a 3D detector, specialised micromachining techniques are required. These fabrication tools are becoming more widely used in a variety of applications. For example, they can be used to make microscopic inertial, pressure and chemical sensors on a silicon chip, along with appropriate readout electronics [18].

The first 3D detectors were produced at the Stanford Nanofabrication Centre, as described in Ref. [19], and used the “full-3D” structure described above. Although the detectors tested in this thesis had different structures, their columns were still fabricated in much the same way [20]. Any differences in the fabrication methods used to produce the double-sided and single-type-column 3D structures are discussed in sections 4.1.1 and 6.1.1 respectively.

The first step of the column fabrication is to etch deep holes in the silicon, using Inductively Coupled Plasma (ICP) etching. An ICP machine uses a set of induction coils to generate an intense plasma, which can then be accelerated onto a silicon wafer by an electric field [21]. Before the etching is carried out, a metal mask is deposited on the silicon, so that only the exposed areas will be etched. The etching process itself is a two-stage cycle. Firstly, the ICP machine uses SF_6 to generate a plasma containing fluorine ions. These are driven down into the wafer and react with the silicon, etching away the base of the hole. However, if this process is continued for too long, unreacted fluorine will accumulate and etch away the sides of the hole. To avoid this, after several seconds of etching the SF_6 is pumped out of the chamber and replaced with C_4F_8 . This produces CF_2 , which forms a Teflon-like coating on the silicon, protecting the sides of the holes during the next phase of etching. This

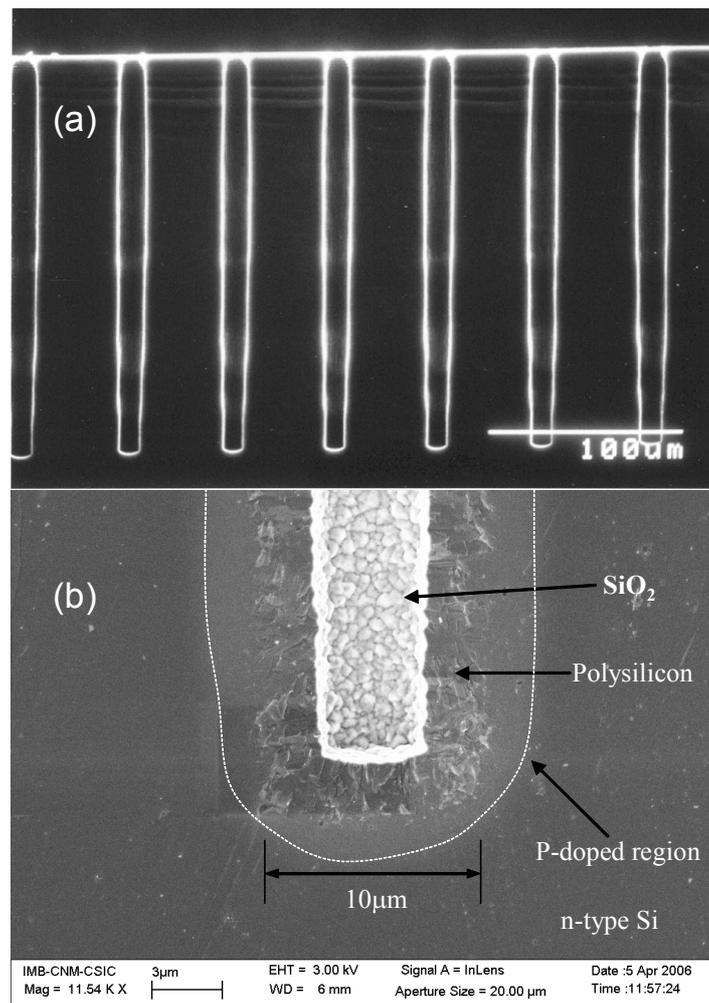


Figure 1.10: (a) SEM image of etched holes in a silicon substrate. (b) SEM image showing detail of a doped column, containing a layer of polysilicon and oxide passivation. A lighter region can be seen around the column, where dopants have diffused through the poly into the silicon; this region has been marked by a dotted line. Images provided by G. Pellegrini, CNM-IMB.

is repeated to produce a set of deep holes, as shown in Fig. 1.10. Ideally, the holes should be very narrow, so that the total volume occupied by the columns is small. In practice, there are limits on how deep a hole of a given diameter can be etched. For example, with current etching machines $10\mu\text{m}$ -diameter holes cannot be etched much deeper than $250\mu\text{m}$. This is acceptable for sensing high-energy particles, but for X-ray detection we would ideally want to use as thick a substrate as possible.

The next fabrication phase is to dope and fill the holes using doped polysilicon. The polysilicon is deposited using SiH_4 gas at low pressure and a high temperature (about 620°C), which produces a very uniform layer of polysilicon within the holes

and across the surface of the wafer. Since the polysilicon is formed under different conditions to the original wafer, the polysilicon layer can place the wafer under stress. This, combined with the fact that the wafer has holes in it, means that there is a risk of bending or even breaking the wafer. So, the conditions need to be optimised to reduce this stress. After depositing 2–3 μm of polysilicon, the columns can then be doped through the polysilicon by diffusion from a solid source. After this, the columns can be completed in different ways. In the method used at Stanford [19], the columns are completely filled with polysilicon. However, at CNM-IMB [20] a layer of silicon dioxide is deposited inside the holes as passivation, but the column is not completely filled—see section 4.1.1. Figure 1.10 shows an SEM image of a column produced at CNM-IMB. After the columns are produced, excess polysilicon can be removed from the surface of the wafer by etching.

These processes need to be repeated to produce the two sets of columns. Then, metal layers can be added to provide readout and bias connections, much like in a standard photodiode detector.

1.2.3 Active edges

As discussed in section 1.1.5, the edges of a saw-cut silicon chip will have large numbers of defects, and possibly cracks, which can cause various problems such as high surface currents. In particular, it is important to prevent the depletion region from reaching the edges. So, a standard photodiode detector chip needs to have an insensitive area at its edge, typically at least 100 μm wide, containing various structures like guard rings.

In the active-edge process [22], instead of saw-cutting the wafer, a wide trench is etched around the edge of each detector. During this process, the sensor wafer is mounted on a carrier wafer, to prevent it from falling apart. By depositing a layer of polysilicon within the trench and doping it, an electrode is fabricated around the edge of each detector array. This means that the detector has virtually no dead area, except the thickness of the electrode itself, which is about 5 μm . The sensors can then be separated by removing the support wafer.

This active edge electrode can be added to both 3D detectors and planar detectors [23]. However, since the same fabrication tools are used to fabricate 3D electrode columns and active edges, they are a convenient combination. Also, in a 3D active edge detector, the electrode spacing and electric field pattern around the edge will be much the same as in the rest of the device. In contrast, a planar ac-

tive edge detector will have an unwanted high-field region around the front surface, where the active edge is close to the doped implants.

The obvious appeal of the active edge is that it can allow a large number of small detectors to be tiled together without any dead areas in between. For example, this would be useful in X-ray diffraction experiments, where it is important not to miss any spots in the diffraction pattern [24]. This would also improve detector yields, since it is a lot easier to produce a large number of small detectors, and reject those with flaws, than it is to directly produce a single, large, flawless device. Active edges could also be useful for applications where we need to place a detector extremely close to a beam, for example in X-ray beam monitoring [25] or to measure proton diffraction in the TOTEM experiment at the LHC [26].

Additionally, since a 3D detector can be operated at a relatively low voltage, and its readout columns are enclosed by the array of bias columns, its depletion region won't extend very far outwith the detector array. So, even if conventional guard structures are used with a 3D detector, it may be possible to achieve a smaller dead area than in a typical planar detector.

1.3 Introduction to synchrotrons

A synchrotron is a variety of particle accelerator, invented in 1943 [27]. A synchrotron can be used to accelerate beams of charged particles to very high energies, for use in particle physics experiments, or can be used to generate intense beams of light for use in other experiments. These two distinct applications are discussed in more detail in the following sections.

Particle accelerators use electric fields to increase the particles' energy. This can be done by using a series of hollow radio-frequency (RF) cavities, each of which has an alternating voltage applied to it. The alternating voltage is timed so that when the bunches of particles pass through the cavity, the electric field will accelerate them and increase their energy. In a linear accelerator, the RF cavities simply form a line. In a synchrotron, the RF cavities form a circle, and a series of bending magnets are used to steer the particle beam. So, the particles repeatedly travel round the ring, gaining energy with each revolution. A full particle accelerator system will typically consist of multiple accelerators with different functions. Figures 1.11 and 1.14 show the layouts of two synchrotrons; the Large Hadron Collider and Diamond Light Source respectively.

When a charged particle travels with velocity \mathbf{v} through a magnetic field \mathbf{B} , it experiences a force which alters its momentum \mathbf{p} :

$$q\mathbf{v} \times \mathbf{B} = \frac{d\mathbf{p}}{dt} \quad (1.15)$$

In a synchrotron, the particles travel perpendicular to the magnetic field, and undergo circular motion. By considering the rate of change in momentum as a particle travels a circular path of radius ρ , we get:

$$\frac{d|\mathbf{p}|}{dt} = |\mathbf{p}|\frac{d\theta}{dt} = \frac{|\mathbf{p}|}{\rho}|\mathbf{v}| \quad (1.16)$$

Hence

$$B\rho = \frac{p}{q} \quad (1.17)$$

This has two implications. Firstly, since the synchrotron's radius is fixed, when particles are being accelerated the magnetic field must be carefully ramped up to ensure the above condition is satisfied. Secondly, the maximum possible momentum that can be achieved is limited by the radius of the ring and the strength of the magnets. Standard electromagnets have a maximum strength of about 2T (limited by heating in the coils), and superconducting magnets are limited to about 10T, so larger rings are needed to achieve increasingly high energies.

To prevent the bunches of particles in the synchrotron from dispersing, it is necessary to apply focusing, both in terms of spatial position and momentum. Along the direction of motion, the RF cavities have a focusing effect. The phase of the sinusoidal voltage applied to the RF cavity is chosen so that particles lagging behind the bunch pass through the cavity when the electric field is stronger, boosting their energy and causing them to catch up with the main bunch. Conversely, particles leading the bunch are accelerated less. To provide focusing perpendicular to the particles' motion, additional magnets are used, as described in Ref. [28].

1.4 Radiation-hard detectors for high-luminosity colliders

1.4.1 Particle physics at the Large Hadron Collider

The Large Hadron Collider (LHC) at CERN in Geneva is a high-energy hadron accelerator and collider, which will be used to investigate physics beyond the Stan-

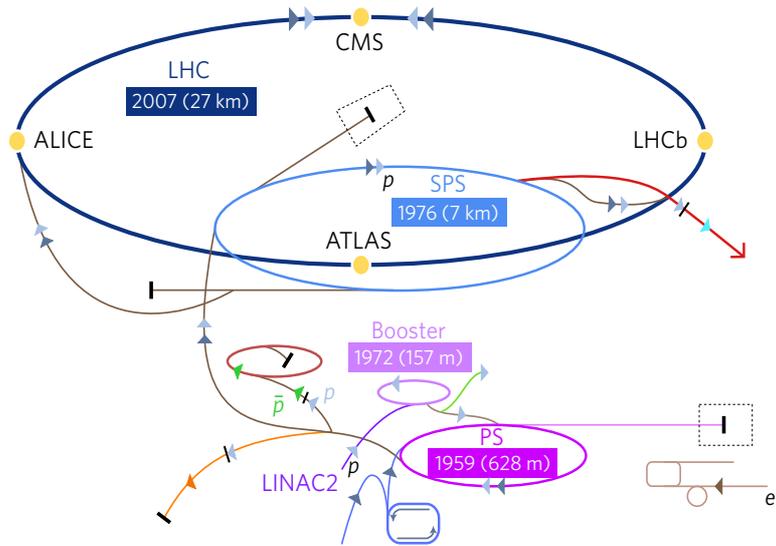


Figure 1.11: A schematic of the accelerators and experiments currently running at CERN. Reproduced from Ref. [30] with permission from Nature Publishing Group.

standard Model. The LHC has two rings inside a 26.7km underground tunnel, which can produce two beams of 7 TeV protons travelling in opposite directions. At four interaction points, the beams cross and produce collisions with a centre-of-mass energy of 14 TeV, making it the highest-energy collider in the world. The LHC can also be used to collide lead ions.

As discussed in section 1.3, the maximum particle energy and momentum that a synchrotron can achieve is limited by its radius of curvature and the field strength of its bending magnets. As well as having an extremely large circumference, the LHC uses superconducting bending magnets to provide a high field of around 8T, allowing it to reach such a high energy.

Figure 1.11 shows a schematic of the accelerators at CERN. The protons are accelerated to their final energy in a series of steps. First, a linear accelerator LINAC2 produces protons at 50 MeV, which then have their energy raised to 1.4 GeV by a multi-ring booster synchrotron, then to 26 GeV by the Proton Synchrotron (PS). The Super Proton Synchrotron (SPS) accelerates the beam further to 450 GeV, and these protons are then injected into the two counter-rotating rings of the LHC itself [29]. The injection process is repeated multiple times to fill the LHC rings.

During operation, the proton beams within the LHC consist of a series of 2808 bunches, spaced by intervals of 25ns. (For lead ions, the spacing is 125ns [31].) By using two colliding beams rather than firing a single beam at a target, the LHC achieves much higher-energy collisions. In fact, when a particle beam hits a

stationary target, the energy in the centre-of-mass frame is only proportional to \sqrt{E} ; due to conservation of momentum, the total kinetic energy after the collision will still be large, and so the energy available for creating new particles is reduced. The main downside of using colliding beams is that the collision rate will be reduced. The collision rate will be proportional to the collider's luminosity. In the simple case of two beams with uniform cross-sectional density meeting head-on, the luminosity is:

$$L = \frac{N^2 f_b}{A} \quad (1.18)$$

where N is the number of particles per bunch, f_b the bunch frequency, and A the cross-sectional area of the beams at the crossing [27]. The LHC is designed to have a luminosity of $10^{34} \text{cm}^{-2} \text{s}^{-1}$ [29].

When particles with high kinetic energies collide, enough energy is available to produce new, unstable particles with a higher rest mass, which cannot normally be observed in nature. Studying these particles gives us more information about the basic laws that govern our universe, and in particular could help us understand what happened in the early Universe, when it was in an extremely hot, dense state.

The LHC will be the first accelerator to investigate physics at TeV energies, and it is expected that various forms of new physics could be discovered at this scale. Firstly, the Higgs particle [32], which is the only undiscovered particle in the Standard Model, is expected to be found at this scale. In the Standard Model, the Higgs explains the origin of particle masses, particularly the W and Z particles associated with the weak force. Other possibilities could be the discovery of supersymmetric partners of existing particles, or the discovery of new weakly-interacting massive particles, which could explain dark matter [33]. These particles will only be produced very rarely during collisions, which is why the LHC's high luminosity is necessary for their discovery. Aside from this, the high collision rates at the LHC will produce previously-discovered particles in large quantities, for example making it possible to study asymmetries between matter and anti-matter by looking at the behaviour of hadrons containing B-quarks.

1.4.2 Detectors at the LHC

At each of the four interaction points where collisions occur, a series of detectors are used to observe the collisions. Since the new particles decay very rapidly into stable, lower-mass particles, their properties have to be determined indirectly from measurements of their decay products. For example, by measuring the paths, momenta and energies of the particles produced in a collision, it is possible to identify

if some of these particles were produced by the decay of a single particle, and then determine the rest mass and other properties of the particle.

The four main experiments at the LHC are ATLAS, CMS, LHCb and ALICE—see Fig. 1.11. ATLAS [34] (“A Toriodal LHC ApparatuS”) is a general-purpose experiment, illustrated in Fig. 1.12. The detector covers the full solid angle surrounding the interaction point, except along the beampipe where the beams enter and leave. The detector can be split into two main sections. The inner detector consists of a silicon pixel detector close to the interaction point, strip detectors, and then the TRT (transition radiation tracker), each of which consists of multiple layers. The particles produced in the collision will pass straight through these detectors, losing relatively little of their energy, allowing measurements of their paths. This makes it possible to reconstruct the positions of the collisions and decays taking place around the interaction point. Since the inner detector is surrounded by magnetic coils, the momentum and charge of the particles can be found by measuring the curvature of their paths as they move through the field. Then, a series of calorimeters will absorb the electrons and hadrons and determine their energy, before the muon chambers finally detect any muons produced in the collision.

As previously mentioned, proton-proton bunch crossings occur with a period of 25ns. Each full detector system must be fast enough to distinguish between the particles produced in each separate bunch crossing, which places major demands on the design of the sub-detectors. For the pixel detector and semiconductor tracker, silicon photodiodes are essential, due to their high speed, high position resolution, and low thickness and density (which minimises scattering of the particles as they pass through the inner detector).

Given the short period between bunch crossings and the huge number of readout channels, it isn’t possible to read out the data from every bunch crossing in a detector like ATLAS. Instead, each detector system will temporarily store all the data, and a limited subset of the data will be used to identify the bunch crossings that contain interesting events. Specifically, data from the calorimeters (read out with reduced position resolution) will be used to identify particles with a large momentum transverse to the beam, which will indicate a “head-on” collision between protons with large amounts of energy available for creating new particles. A trigger signal will then be sent, and the corresponding event data will be read out from the full detector. In ATLAS, the full detector will only be read out at a rate of about 75kHz, and after further trigger logic is applied events will be written to disk at a rate of below 3.5kHz.

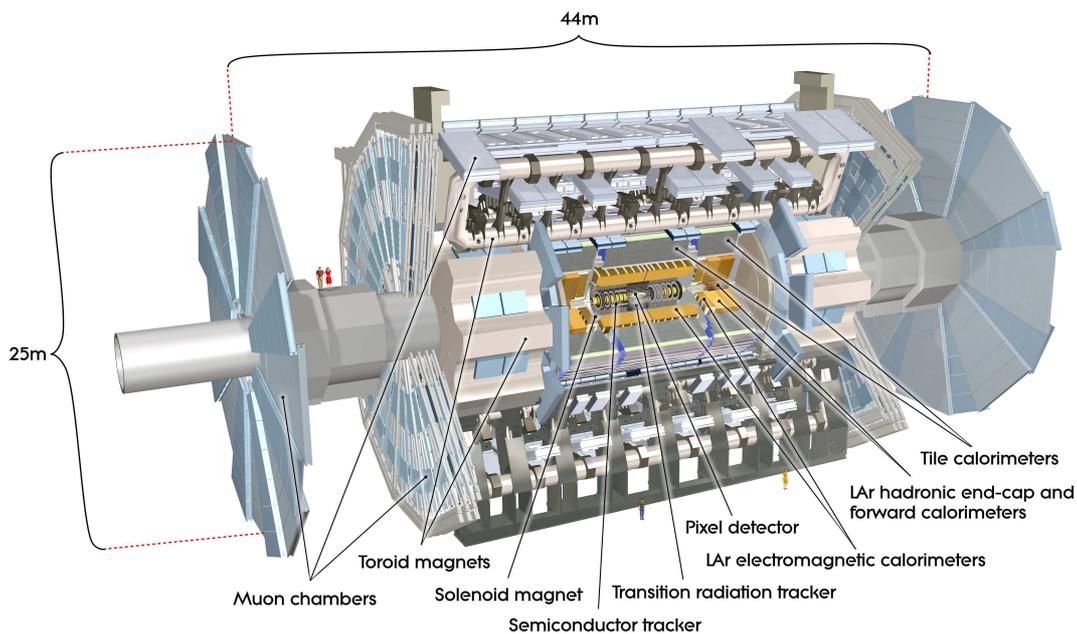


Figure 1.12: Computer image of the ATLAS detector, reproduced from Ref. [34] with permission from IoP Publishing Ltd.

1.4.3 The Super-LHC and radiation hardness

Even before the LHC's switch-on in 2008, plans have been made for an upgrade to the collider around 2017, referred to as "Super LHC" (SLHC). This would involve altering various aspects of the machine to increase its luminosity by a factor of ten [35]. This tenfold increase in the collision rate could allow more precise measurements of the properties of any newly-discovered particles, like the Higgs. Additionally, the increase in the number of rare events would effectively increase the LHC's mass reach by 20-30% [36].

However, the increase in the number of collisions would also put greater demands on the detectors used in the experiments. Firstly, the number of particles produced in each bunch crossing would increase by a factor of 10 (or more, since the bunch crossing rate might be reduced in the upgrade). In order to reliably distinguish between the tracks produced by these particles, the granularity of many of the detectors would need to be increased. This would also place greater demands on the pattern-recognition software used to assemble the hits on the detectors into full particle tracks.

The increase in luminosity would also lead to a tenfold increase in the radiation damage received by the detectors. As shown in Fig. 1.12, the innermost section of

ATLAS is a silicon pixel detector, which is important for accurately reconstructing the paths of particles produced by the collisions. Due to its closeness to the collision point, the pixel detector receives the highest level of radiation damage; its innermost layer, the b-layer, is just 5cm from the beamline and will need to be replaced after about 3 years of operation at the LHC's design luminosity. Following the SLHC upgrade, the b-layer would receive an equivalent fluence of about 1×10^{16} 1MeV- n_{eq}/cm^2 over the SLHC's running time, mostly in the form of charged hadrons [37]. Since the current pixel detector is only designed to withstand 1×10^{15} 1MeV- n_{eq}/cm^2 , a much more radiation tolerant design is needed. So, the CERN-RD50 collaboration was started, to develop more radiation-hard detectors [38]. Currently, 3D detectors are a promising technology for the b-layer at the SLHC, and are also being considered for the earlier b-layer replacement [39].

There are three main effects of radiation damage. Firstly, the leakage current increases, leading to higher power dissipation, which in turn places greater demands on the cooling system. Secondly, the effective doping concentration of the substrate increases, which makes it more difficult to deplete the detector. Thirdly, many of the charge carriers generated by a particle will be trapped by radiation-induced defects, reducing the collected signal. These radiation effects are discussed in detail in Chapter 2.

Because 3D detectors have a very short distance between electrodes, their depletion voltage is substantially reduced, and it should be possible to fully deplete them after high radiation doses. Since charge trapping causes the number of free carriers to fall exponentially with time, the 3D detector's fast collection speed should reduce the number of trapped carriers. So, the 3D structure can counteract two of the main problems of radiation damage. The increasing leakage current must be dealt with by cooling the detector, as with standard photodiodes.

In Chapter 2, different 3D detectors are simulated, in order to find the optimum design for a detector at the SLHC. In Chapter 4, a set of double-sided 3D detectors are simulated and tested, including tests of collection efficiency after radiation damage.

1.5 Detectors for X-ray synchrotrons

1.5.1 Synchrotron light sources

When the path of the particle beam in a synchrotron is bent by each set of dipole magnets, this acceleration causes the particles to emit electromagnetic radiation,

which is referred to as synchrotron radiation. If we consider a reference frame with the same instantaneous velocity as the particle, then the radiation is uniformly distributed. However, since the particle is travelling at relativistic speeds in the laboratory frame, the radiation is mostly emitted in a narrow cone in the particle's direction of travel, and it is also shifted to shorter wavelengths [40].

The power radiated by the particle travelling at relativistic speed is given by:

$$P = \frac{1}{6\pi\epsilon_0} \frac{e^2 c}{\rho^2} \gamma^4 \quad (1.19)$$

Since $E = \gamma m_0 c^2$, this means that the power radiated by the particle is proportional to E^4/m_0^4 . This means firstly that the power radiated has a very strong energy dependence, and secondly that electrons radiate much more than protons, by a factor of $(m_p/m_e)^4 = 1.13 \times 10^{13}$!

Above energies of about 1 GeV, electron synchrotrons require a rapidly increasing input of RF power simply to compensate for the loss due to synchrotron radiation, limiting the maximum energy that can be achieved. However, the synchrotron radiation itself can be used as a powerful research tool [41]. It has a spectrum that extends from infra-red up to X-ray energies, which can be used in a variety of experiments. The spectrum peaks in the X-ray region, providing a vertically collimated X-ray source orders of magnitude more intense than conventional sources.

1.5.2 Diamond: a third-generation light source

Diamond Light Source is a new synchrotron light source, located at the Rutherford Appleton Laboratory in Oxfordshire. It opened to its first users in February 2007. It has a 3 GeV beam with a circumference of 158m, and is optimised to produce X-rays with energies between 100eV and 20keV [42]. Figure 1.13 shows a picture of Diamond.

The earliest experiments with synchrotron light in the 1960s ran parasitically on synchrotrons used for high-energy physics experiments. However, in the late 1970s and onwards, second- and third-generation synchrotrons have been purpose-built to produce synchrotron light, allowing more optimisation of their design [44].

In high-energy physics experiments, the aim is to produce a particle beam at as high an energy as possible, often using multiple synchrotrons to successively increase the energy. In light sources, however, the aim is to maintain the electron beam at a specific energy, using the RF cavities to replenish the energy emitted as synchrotron radiation. So, modern light sources generate their photon beams using a large



Figure 1.13: Aerial photograph of Diamond Light Source, provided by Diamond Light Source Ltd. [43].

“storage ring” operating at a fixed energy [28]. At Diamond, a linear accelerator accelerates electrons to 100 MeV, a smaller booster synchrotron is used to increase the energy to 3 GeV, and then the electrons are injected into the storage ring. The ring accumulates pulses of electrons from the booster, until the required operating current is reached, at which point the injection system can be switched off. The electrons in the ring continue to circulate at an energy of 3 GeV. The synchrotron light emitted by the storage ring is transmitted via beamlines to various experimental stations. See Fig. 1.14.

Since the storage ring operates at a single energy, the magnetic fields in the ring can be kept constant, which greatly improves the stability of the beam. This makes it possible to inject a reasonably large electron current into the storage ring, then keep it circulating for many hours. Diamond, for example, can have a beam current of around 300mA circulating for 10–20 hours. This ensures that the output spectrum and intensity are kept fairly constant and the useful operating time between beam refills is maximised. This stability also makes it easier to focus the beam. This focusing reduces the spot size and divergence of the X-ray beams produced by the synchrotron, increasing their brilliance.

A major advance in third-generation sources is the use of insertion devices to produce light, rather than bending magnets alone. An insertion device is a straight section of the ring containing sets of dipole magnets with alternating polarity that produce a periodic field. As the electrons pass through the alternating fields they

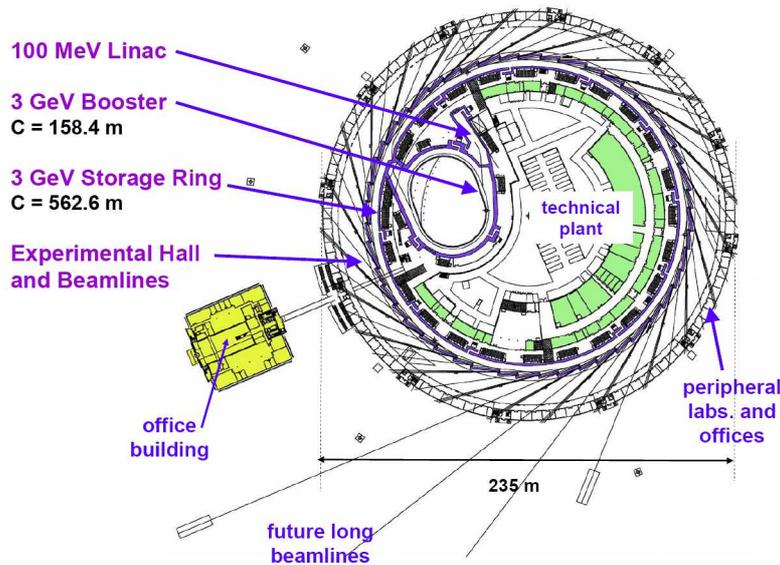


Figure 1.14: Schematic of the Diamond synchrotron [43], provided by Diamond Light Source Ltd.

follow an oscillating path and emit electromagnetic radiation, like when they pass through a bending magnet. However, since the electrons are following an approximately straight path rather than a curve, the insertion device creates a single intense beam of photons rather than an arc, greatly increasing the brightness. Additionally, different designs of insertion device can produce beams with different properties [41]. “Wigglers” use high-field magnets to generate X-rays at higher energies than the bending magnets can produce. “Undulators” have a large number of alternating magnets with moderate field. The radiation produced by each successive oscillation will produce constructive interference at certain photon energies (determined by the period of the undulator), so, in addition to normal synchrotron light, the undulator produces intense, coherent radiation at a specific wavelength [28].

1.5.3 Applications of synchrotron X-rays

The X-ray beams produced by synchrotrons can be used for three general applications [43].

- *Diffraction and scattering:* X-rays with energies of order 10keV have wavelengths around 1\AA , which is similar to the spacing between the atoms in a crystal. When these X-rays pass through a crystal, a diffraction pattern is produced, containing information about the crystal’s structure. This technique

can determine the structure of many materials including biological samples, minerals, and electronic and magnetic materials. A particularly important variant is macromolecular crystallography, which is the main technique used to obtain structural information about proteins and other biological molecules. This information can be used to understand how the body functions and how diseases develop, and to design new drugs. The very intense X-ray beams produced by new light sources like Diamond make it possible to measure weak diffraction patterns more accurately, gaining more information from samples that are small, weakly diffracting, or have complicated structures. These experiments need fast, accurate position-sensitive detectors to record the diffraction patterns.

- *Spectroscopy*: Synchrotrons produce a wide spectrum of X-ray energies. The absorption, reflectivity or fluorescence of a sample at different X-ray energies provides information about its elemental composition, chemical state and physical properties. This can be used to determine the composition of unknown samples, or to investigate the properties of new materials. The high intensity and tunability of synchrotron X-rays mean they can detect a wide range of elements in extremely low concentrations. Also, by scanning an X-ray microbeam across a sample, this information on chemical composition can be obtained with fine position resolution, too.
- *Microscopy and imaging*: Samples can be imaged in an X-ray beam either by illuminating the whole sample, or by scanning a small beam spot across a sample. Imaging an entire sample by absorption can provide additional information about samples tested using other techniques, and the wide spectrum from a synchrotron allows contrast measurements at a range of energies. With spot scanning it becomes possible to apply diffraction or spectroscopic techniques with position resolution.

To match this wide range of experimental techniques, a wide variety of X-ray detectors are used in synchrotrons. With appropriate readout chips, photodiode detectors can be used to count individual photons hits at high speed and with good spatial resolution, making them useful for diffraction experiments. 3D detectors could potentially improve these experiments further, due to their superior performance—particularly their low charge sharing, which can improve the image quality. Chapter 5 discusses these issues more fully, and presents results from X-ray tests on a set of double-sided 3D detectors connected to single-photon-counting readout chips.

Chapter 2

Simulation of 3D detectors for high-luminosity colliders

It is possible to model the behaviour of a semiconductor detector using computer simulation. This can be a useful way of evaluating and optimising the design of a detector before fabrication, or testing a range of different detector structures to see which is the most promising. Simulations can also be helpful for understanding the behaviour of a device.

In this chapter, the simulation package “Synopsys TCAD” [45] is used to model the behaviour of 3D detectors. Firstly, the principles of semiconductor simulation are described in section 2.1, along with details of the Synopsys software. This is followed in section 2.2 by a few illustrative simulations of 3D detectors, investigating their basic electric field pattern, collection speed and weighting field.

Section 2.3 then discusses the physical effects of radiation damage on semiconductor detectors, and how these effects can be modelled by Synopsys TCAD. This damage model is then used in sections 2.4 and 2.5 to simulate the behaviour of a range of possible 3D pixel structures after high levels of radiation damage. These simulations consider depletion and breakdown behaviour, charge collection efficiency and noise, in order to find the optimal detector design for use in future high-luminosity colliders such as the Super-LHC. This accelerator and its experiments were discussed in section 1.4.3.

In later chapters, the methods introduced here are also used to simulate alternative 3D structures, to compare their behaviour to full-3D. The alternative structures are double-sided 3D detectors, fabricated by CNM-IMB (Chapter 4) and single-type-column 3D detectors, fabricated by FBK-IRST (Chapter 6).

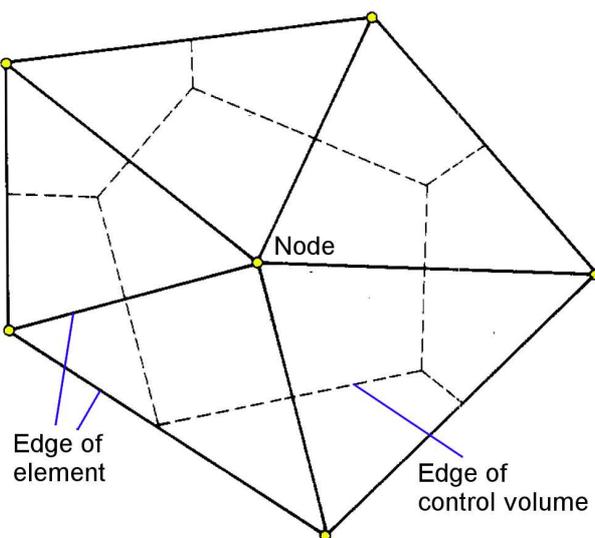


Figure 2.1: Region of a 2D mesh, with a network of nodes which form triangular elements.

2.1 Simulation with Synopsys TCAD

All of the following simulations were done with “Synopsys TCAD” version Z-2007.03 [45], a finite-element semiconductor simulation package.

Back in section 1.1, the basic differential equations of semiconductor physics were introduced. By solving these equations with appropriate boundary conditions, the behaviour of a semiconductor device can, in principle, be found analytically. However, in practice this can only be done for relatively simple devices and conditions. The alternative approach involves representing the device structure by a mesh of discrete nodes, and applying the semiconductor equations to each point in an approximate form. Instead of partial differential equations, we now have a large system of equations, written in terms of the electrostatic potential and carrier concentrations at each node. These can then be solved to an acceptable level of accuracy by iteration [46].

2.1.1 Basic finite-element analysis in semiconductors

In finite-element analysis, the device is approximated by a 2D or 3D mesh of connected nodes which may have an irregular arrangement—unlike, say, finite difference methods which rely on a regular grid [47]. A simple example is given in Fig. 2.1.

Firstly, we need to be able to represent the state of the semiconductor using the mesh. The fixed properties of the device itself are its basic geometry, for example the substrate thickness and positions of contacts, and the various doping profiles

within the silicon. Looking at the semiconductor equations in section 1.1, there are three important variables that describe the state of the device at any moment—the electrostatic potential, and the electron and hole concentrations. Other quantities, such as the carrier currents and electric field, are simply functions of these three variables. During the simulation, each node will have its own electrostatic potential, hole concentration, and carrier concentration. The volume between the nodes is split up into a series of elements. The values of the three variables are defined throughout each element by taking the values at the surrounding nodes and applying a linear interpolation process [48]. As an aside, the electron and hole concentrations are usually represented by the electron and hole “quasi-Fermi potentials”, which are proportional to the natural logs of the electron and hole concentrations [1]. Since the carrier concentrations in a device can vary by many orders of magnitude, this approach makes the interpolation more realistic and reduces rounding errors.

Next, we need to apply the semiconductor equations in a discrete, approximate form to each node in the mesh. Consider applying the Poisson equation—Eq. 1.8 in section 1.1.1—to a particular node. The carrier concentration terms n and p are simply given by their values at the node. However, we also have the term $\nabla^2\psi = \nabla \cdot (\nabla\psi)$, which involves spatial derivatives. To deal with this, we consider a “control volume” around the node, constructed by taking perpendicular bisectors from each line leading to the adjacent mesh points [47]. See Fig. 2.1. Then, we integrate this $\nabla \cdot (\nabla\psi)$ term over the control volume, V , and normalise the result by dividing by the volume. By applying the divergence theorem, we get:

$$\frac{1}{V} \iiint_V \nabla \cdot (\nabla\psi) dV = \frac{1}{V} \iint_S (\nabla\psi) \cdot \mathbf{n} dS \quad (2.1)$$

where S is the surface of the volume, and \mathbf{n} is the normal to the surface. Then, the linear interpolation referred to above can be used to find ψ and hence $\nabla\psi$ over the surface, reducing this term to a fairly simple summation. The discrete form of the Poisson equation is then:

$$\epsilon_s \frac{1}{V} \sum_j \frac{A_{ij}}{d_{ij}} (\psi_i - \psi_j) = q \left(n_i - p_i + N_{a(i)}^- - N_{d(i)}^+ \right) \quad (2.2)$$

where subscript i refers to the node we’re dealing with, j refers the adjacent nodes, d_{ij} is the distance between the nodes i and j , and A_{ij} is the area of the face of the control volume between the elements. The carrier continuity equations can be handled in the same way.

The resulting system of equations (plus boundary conditions) can be solved by a variety of iterative methods. In Synopsys TCAD, a global approximate Newton method is used [45]. The problem can be expressed in the form $g(\mathbf{z}) = 0$, where \mathbf{z} is a vector representing a possible solution and g the system of equations. Suppose we start with a guess solution \mathbf{z}_n , giving us a system of equations $g(\mathbf{z}_n)$ whose value is nonzero. If we know the gradient $(\nabla g)_{\mathbf{z}=\mathbf{z}_n}$, then this allows us to generate a new solution that should be closer to the correct one:

$$\mathbf{z}_{n+1} = \mathbf{z}_n - \lambda \frac{g(\mathbf{z}_n)}{|(\nabla g)_{\mathbf{z}=\mathbf{z}_n}|} \hat{\mathbf{k}} \quad (2.3)$$

where λ is a constant of less than 1 (to give more reliable convergence) and $\hat{\mathbf{k}}$ is a unit vector in the direction of $(\nabla g)_{\mathbf{z}=\mathbf{z}_n}$.

This iterative process can be repeated until the both the error $|g|$ and the correction term are sufficiently small. The main challenge of using this method is that since $(\nabla g)_{\mathbf{z}=\mathbf{z}_n}$ is not known precisely, a suitable approximation needs to be found. This is discussed in Ref. [49].

2.1.2 Mesh building with MESH

The first step in the simulation process is to build a network of nodes that approximate the structure of the device. In Synopsys, this is done using the program MESH. The input to MESH is a set of text files which specify the following:

- The basic device structure—the dimensions of the different materials in the device, such as the silicon substrate and dielectric layers, and also the contacts.
- Doping profiles within the device. For example, a doped strip will be defined by the surface region where the doping takes place, and a Gaussian profile describing how the doping concentration varies with depth in the silicon.
- The rules that MESH should follow when choosing the mesh spacing in different regions of the device. These can include maximum and minimum allowed spacings in different regions, and special rules that adapt the mesh spacing to match doping profiles accurately.

The simulation process assumes that, between nodes, the electrostatic potential and quasi-Fermi potentials vary linearly. So, if the value of some variable changes rapidly across a region containing few nodes, the accuracy of the simulation will be reduced. However, as the number of nodes increases, the resulting system of

equations gets larger and the solving process becomes slower and more difficult. The solution is to use a high mesh density only in the regions where the doping concentration changes rapidly, the electric field is high, or high levels of carrier generation will occur. In practice, it is often useful to run a quick simulation using a widely-spaced mesh, and then use the crude results from this to decide how the mesh should be improved.

Fig. 2.2 shows an example of a mesh used to simulate a full-3D detector. The basic structure is simple, consisting of a p-type silicon substrate with two contacts and oxide layers on the surfaces. The doping profiles consist of the two cylindrical columns, and also a p-spray, as discussed in section 1.1.5. Within the columns, the doping concentration is high and takes a fixed value, and at the edges of the columns the doping concentration falls off with distance with a error function. The mesh spacing is highest around the columns, where the doping concentration varies rapidly with position, and the front and back surfaces, where the oxide/silicon interface affects the device behaviour. In the central region of the device, the electric field varies rapidly with the horizontal position, but not along the z-direction. So, the vertical mesh spacing can be much wider.

Another important aspect of mesh design is choosing how large a region of the device to simulate, since simulating a larger volume will of course increase the complexity of the mesh. Generally, a detector will have a repetitive structure, with many pixels or strips, and in a steady state it will have a repetitive field pattern, carrier densities etc.

At the contacts, the boundary conditions generally consist of a fixed electrostatic potential, and charge neutrality. At the boundaries of the device mesh, the default conditions are that the electric field and current density normal to the boundary are zero. These are referred to as reflecting or Neumann boundary conditions. In a real detector in a steady-state, these conditions will naturally occur along planes of symmetry in the device. This means that basic steady-state simulations can be done using just the simplest repeating unit of the detector. In a 3D detector, this will be a quarter of a pixel, as shown in Fig. 2.2. However, in order to simulate capacitive coupling or charge sharing between adjacent pixels, a larger region must be simulated. This will be discussed in section 2.2.

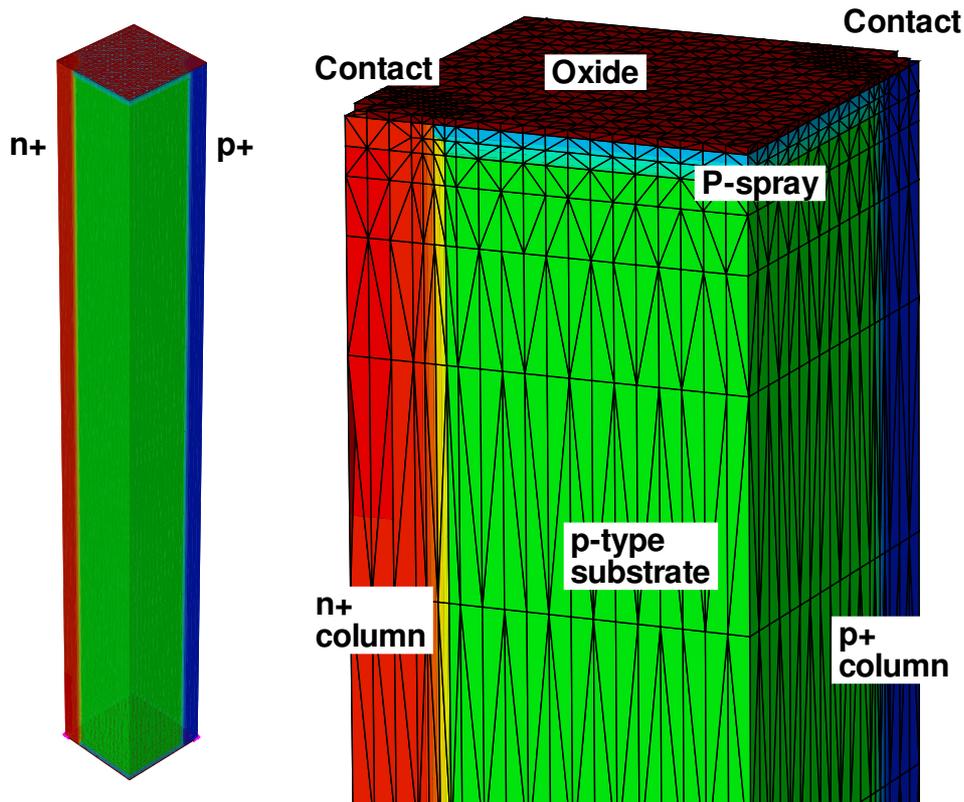


Figure 2.2: Structure of a full-3D simulation mesh, showing the grid of connected nodes.

2.1.3 Simulation with Sentaurus Device

Physics models

After creating the mesh, the Sentaurus Device program can be used to run the simulation. Once again, this is controlled by a text file. The file contains the details of the physics models used in the simulation, and a series of commands fully specifying the conditions that need to be simulated.

Section 1.1 gives the basic semiconductor physics equations that Synopsys uses. However, many of the terms in these equations are dependent on more complicated physics models, which can be altered by the user. For example, in simple simulations (without radiation damage) the carrier generation and recombination rates are based on a Shockley-Read-Hall model [50] and depend on the electron and hole concentrations and the doping density. However, it is also possible to choose an avalanche multiplication model, where the generation rate increases in high-field regions. Similarly, a “heavy ion” model will generate extra electron-hole pairs in some region of the device, to simulate the effects of ionising radiation. In section 2.3.3

onwards, the effects of radiation-induced defects are simulated by including carrier motion between the conduction and valence bands and a series of extra trap levels within the bandgap.

In all the following simulations, the mobility of the carriers was taken into account using doping-concentration-dependent and high-field-saturation-dependent physical models. This means that the electron and hole velocities cannot increase without limit as the electric field gets stronger, but will instead reach a saturation velocity. This is particularly important when simulating photodiode detectors, where the fields tend to be high. Unless otherwise stated, the simulations were done at 300K temperature, and used the Shockley-Read-Hall generation model above.

Specific physics models can be applied in different regions of the device. The main additional model used here was the introduction of a fixed positive charge of $4 \times 10^{11} \text{ e.cm}^{-2}$ at the interface between the silicon substrate and the oxide layer. This is due to the presence of trapped holes within the oxide layer, as discussed in section 1.1.5. The outer surfaces of the oxide layers used Neumann boundary conditions, which is a good approximation to the oxide behaviour in a “clean” wafer [51].

Simulation conditions

The most basic form of simulation simply applies a set of boundary conditions, typically a set of electrode voltages, and finds the solution for a device in a steady state. Under these conditions, the time-dependent terms in the semiconductor equations are zero. These simulations can, for example, be used to find the electric field pattern in the device.

The next variety is “quasi-stationary”. The device is first solved in a stationary state, as described above. Then, some of the boundary conditions such as the electrode voltages are changed by a small amount, and the device is re-solved in a steady state. This is repeated over a series of steps, in order to find how the device behaviour varies with some parameter. This can be used to find the current-voltage or capacitance-voltage characteristics of the device. At each step, an initial guess at the correct solution is found by extrapolation from the previous solutions, speeding up the process.

Lastly, there are “transient” simulations, where the device is simulated over time. The initial state of the detector is found in a steady-state, as above, and then the simulation proceeds in a series of small time steps. At each step, the rates of change in carrier concentrations and potential are found for each node, and these are then used to find the state of the device at the next step [52]. This can be used to find

the current signal produced by a particle interacting with a detector. To ensure that the process is accurate, the step sizes must be small compared to the time scales of the processes occurring during the simulation (e.g. the collection time).

2.2 Basic simulation of full-3D detectors

Here, a few basic simulations with a standard full-3D detector structure are presented, to establish general behaviour such as the depletion process and electric field pattern. Most of the devices tested during this thesis actually used the alternative “double-sided” 3D structure. In Chapter 4, a more thorough series of simulations on standard and double-sided 3D detectors are presented side-by-side for comparison, along with experimental results. Likewise, in Chapter 6 the third structure, single-type-column 3D, is simulated and tested.

2.2.1 Mesh structure

Firstly, the 3D detector mesh was designed—see Fig. 2.2. The mesh had a p-type substrate and used n-type readout. The specific details of the simulated structure were as follows:

- The p-type substrate was $230\mu\text{m}$ thick, with $7 \times 10^{11} \text{ cm}^{-3}$ boron doping.
- The columns were cylindrical and $5\mu\text{m}$ in radius. To simplify the structure, the columns consisted of doped silicon, whereas polysilicon is used in real devices. The n- and p-type columns had $6 \times 10^{18} \text{ cm}^{-3}$ phosphorus and boron doping respectively. The doping profile at the edge of each columns was an error function, which increased the radius of the junction around the n-type columns to about $6.5\mu\text{m}$.
- The n+ columns were isolated using a p-spray [53] covering the front and back surfaces. This used a total boron dose of $1.2 \times 10^{12} \text{ cm}^{-2}$, and had a peak concentration of $2.5 \times 10^{16} \text{ cm}^{-3}$. The p-spray was chosen to have a high enough dose to compensate for the electron layer (preventing it from linking the n+ columns together) and to have a fairly realistic profile. The need for electrode isolation is discussed in section 1.1.5.

The substrate thickness was chosen to match the thickness of a set of ATLAS 3D sensors tested by another institution, as discussed later. For this particular simulation, a pixel size of $55\mu\text{m}$ was chosen, so the simulation mesh had a volume

of $27.5\mu\text{m}$ by $27.5\mu\text{m}$ by $230\mu\text{m}$. This is a relatively small pixel size, corresponding to the Medipix2 readout chip tested in Chapter 5.

2.2.2 Depletion and electric field behaviour

Firstly, a quasi-stationary simulation was run, where the bias applied to the detector was steadily increased from 0V to 100V, to find how the depletion region and electric field pattern develop. Initially, the depletion region appears where the n+ column meets the p-type substrate, then grows cylindrically outwards. Due to the small electrode spacing, the full volume of the detector is depleted at just 2V, and so the detector can be operated at an extreme overbias.

Figure 2.3 shows the electric field strength in a horizontal cross-section of the device, taken at a depth of $115\mu\text{m}$ (i.e. halfway through the substrate) and with a bias of 100V. The field direction is indicated by the black lines. In the immediate vicinity of each of the columns, the field is at its strongest, and points radially outwards from the column. Although at low biases the field will be stronger around the n-type column, where the n-p junction is formed, when the device is overbiased like this both columns have similar field strengths. In the region between the columns, the small electrode spacing means that the field is strong; over $25000\text{V}/\text{cm}$. However, at the corners of the simulated region, which will correspond to the midpoints between the readout columns in two adjacent pixels, the field becomes weaker, falling to zero at the very corner due to the device symmetry. So, charge deposited in this region is likely to be collected more slowly. The electric field pattern also means that when electrons drift to the readout electrode, they will tend to drift away from the pixel boundaries, which will reduce the effects of charge sharing.

Throughout most of the device volume, the electric field has no vertical component, so carriers will just drift horizontally to the readout electrodes. This can be seen in Fig. 2.4, which shows the electrostatic potential in a vertical cross-section through the detector, passing through both the n+ and p+ columns (i.e. diagonally across the pixel). The electrostatic field only changes around the front and back surfaces of the detector. The p-spray at the surfaces makes contact with the p-type columns, so its potential matches the applied bias. This adds a vertical component to the field around the surface, and also creates a high-field region where the p-stop meets the n+ column. In detectors with p-type readout and no p-spray, the electron layer attracted by the oxide charge is at the same potential as the n-type bias columns, so a similar effect occurs.

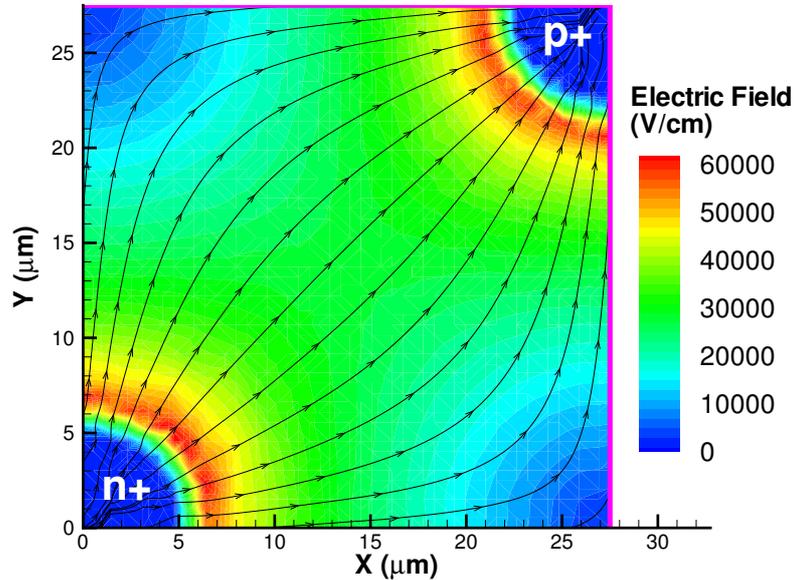


Figure 2.3: Electric field in a horizontal cross-section of a full-3D detector at 100V bias. The black “streamtrace” lines show the direction of the electric field, but their spacing is not proportional to the field strength.

When the electric field in a detector becomes too high, avalanche breakdown can occur—see section 1.1.5. Although avalanche breakdown models can be used in Sentaurus Device, problems can occur when these models are combined with the radiation damage models used later in this chapter. Instead, a simulation was run which recorded how the maximum electric field within each device varied with the applied bias. Since the breakdown field in silicon is $3 \times 10^5 \text{V/cm}$, a maximum field of $2.5 \times 10^5 \text{V/cm}$ can be regarded as a reasonably safe operating condition. Because the high-field region occurs where the p-stop meets the n-type column, the point where this field strength is reached depends on the oxide charge. With the oxide charge at its saturated value of $1 \times 10^{12} \text{e.cm}^{-2}$, this field is reached at 170V. However, with $4 \times 10^{11} \text{e.cm}^{-2}$ oxide charge, this “high-field” bias falls to 55V. This effect is typical for p-spray isolation. With higher oxide charges, the electrons attracted by the oxide charge tend to compensate for the p-spray doping, reducing the maximum field [11]. Similarly, when the p-spray dose was doubled from $1.2 \times 10^{12} \text{cm}^{-2}$ to $2.4 \times 10^{12} \text{cm}^{-2}$, the electric field between the n+ column and the p-spray increased, lowering the safe operating voltage.

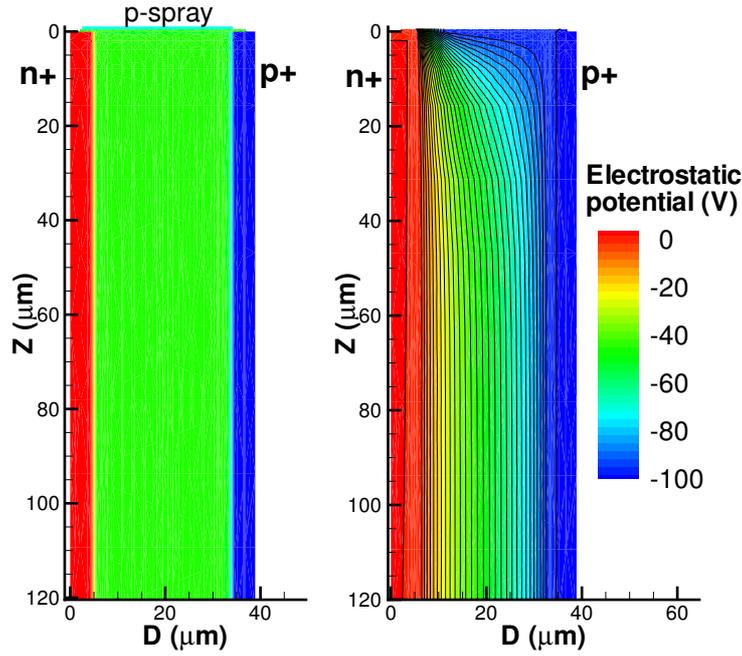


Figure 2.4: Electrostatic potential in a vertical cross-section of a full-3D detector at 100V bias. The cross-section passes through neighbouring n- and p-type columns, as shown by the diagram on the left. Only the top $120\mu\text{m}$ of the full $230\mu\text{m}$ thickness is shown; since the columns pass through the full device thickness, the behaviour is the same around the front and back surfaces.

2.2.3 Typical charge collection speed

Next, a basic charge collection simulation was done with this structure. At the start of the simulation, free electron-hole pairs were created uniformly throughout the device volume, and then the currents at the electrodes were simulated over time. The total number of electron-hole pairs was chosen to match the charge generated by a minimum ionizing particle passing through the detector, which is about 18500 (i.e. 80 pairs/ μm) [54]. Once again, the detector was biased to 100V.

This charge sharing simulation, and others that follow, only use a quarter of a device pixel, in order to improve the simulation speed. This introduces a couple of limitations. Firstly, charge sharing between adjacent pixels will not be taken into account. However, charge sharing has been previously simulated in 3D detectors [55], and should be low. Secondly, the signals on the electrodes will be slightly altered. Moving charge carriers induce signals on the electrodes due to their electric fields, and these fields will be affected by the presence of neighbouring readout electrodes. However, the “Weighting field” section below demonstrates that this effect will be small in 3D detectors.

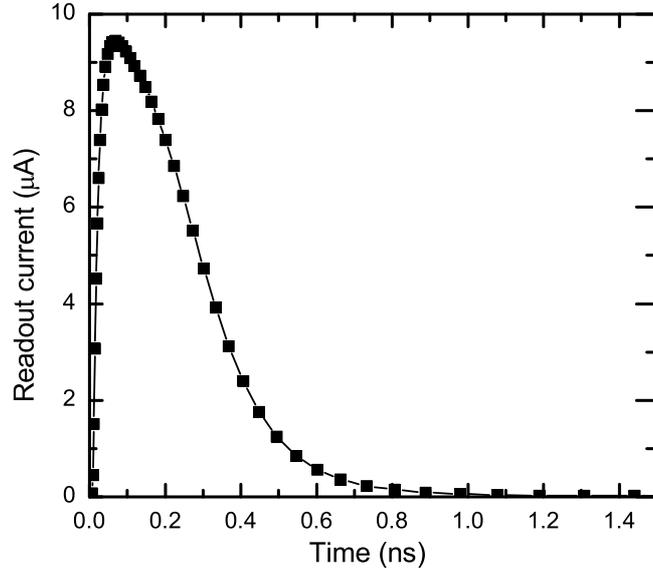


Figure 2.5: Simulated current signal produced when the 3D detector at 100V is uniformly flooded with electron-hole pairs.

The resulting current signal at the n-type readout electrode is shown in Fig. 2.5. Virtually all the signal is collected within 1ns, and the total charge collected matches the charge deposited. Given that LHC-speed electronics are read out at 25ns intervals, this means that the 3D detector’s collection time is much smaller than the shaping time of the readout electronics. As discussed in the following sections, a fast collection time is beneficial for reducing charge trapping effects following radiation damage.

2.2.4 Weighting field

The signals generated on the readout electrodes by moving charge carriers can be calculated using Ramo’s theorem [4]. A “weighting field” \mathbf{E}_w can be calculated for each electrode in the device, such that the current signal induced on that electrode by a moving charge carrier will be:

$$I_i = q\mathbf{v} \cdot \mathbf{E}_w \quad (2.4)$$

where \mathbf{v} is the velocity of the carrier.

The weighting field is given by $-\nabla\Psi_w$, where $\Psi_w(x)$ is the weighting potential. The weighting potential for a given electrode can be calculated by finding the

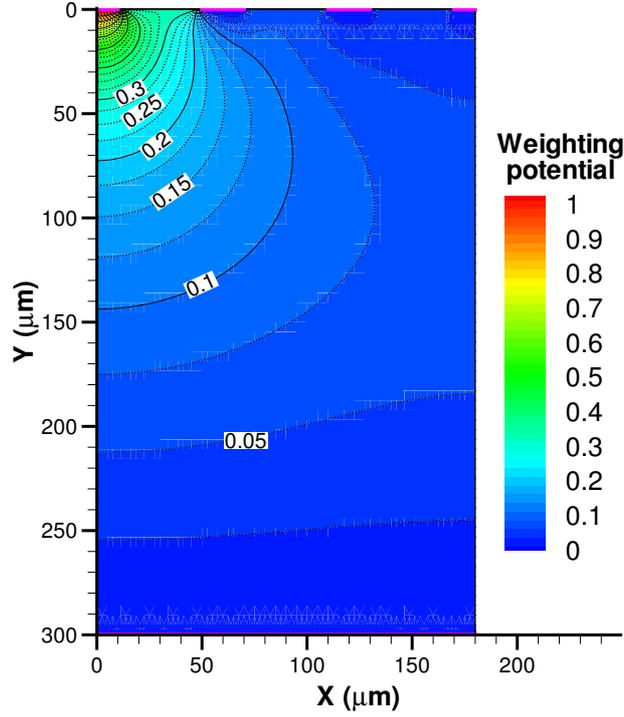


Figure 2.6: Weighting potential in an n-on-p planar strip detector. The potential is calculated for the leftmost strip. The potential gradient is highest near to the strip, which means that carriers moving close to the strip will generate large current signals.

incremental change in the real potential throughout the device in response to an incremental change in the bias applied to that particular electrode [56]:

$$\Psi_w(\mathbf{x}) = \frac{d\Psi(\mathbf{x})}{dV_i} \quad (2.5)$$

It also follows that the total charge signal induced on the electrode when the charge carrier moves from position x_A to position x_B will be $q(\Psi_w(x_B) - \Psi_w(x_A))$.

The weighting potential can be found quite easily in Synopsys by biasing the device to the required potential, making a small change to the voltage on one electrode, and calculating the change in the potential distribution. Figure 2.6 shows the calculated weighting potential for an electrode in a fully depleted planar strip detector. The strip pitch is $60\mu\text{m}$ and the substrate thickness is $300\mu\text{m}$. In this case the readout electrodes are n-type and the substrate is p-type, though the weighting field pattern would be the same for a p-on-n device. Inevitably, the weighting potential is 1 on the appropriate electrode, and 0 on the others.

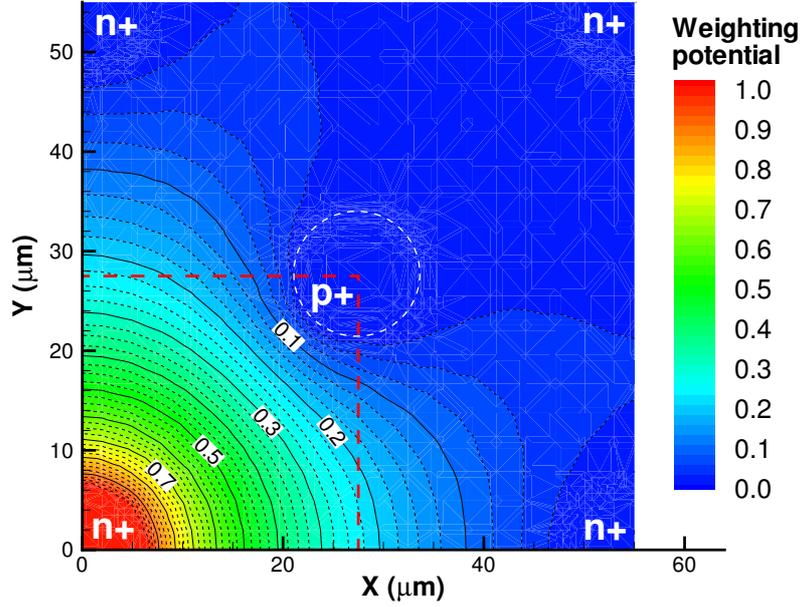


Figure 2.7: Weighting potential in a cross-section of a 3D pixel detector. The potential is calculated for the n-type electrode in the bottom-left corner, and the area of the corresponding pixel is indicated by the dotted red line.

In the case of the strip detector, the weighting potential is high in the vicinity of the readout strip, but drops rapidly with distance from it. So, whichever carrier is collected at the segmented electrode will generate a high proportion of the total signal. For example, if an electron-hole pair is generated at $150\mu\text{m}$ depth beneath the strip, where the weighting potential is 0.1, the electron will induce a signal of $-0.9e$ at the electrode when it is collected, and the hole will induce just $-0.1e$ as it drifts to the back contact. Since electrons are more mobile than holes, this will have an important effect on the collection speed.

The weighting field was then calculated for a 3D detector. Here, the device structure included 4 adjacent n-type readout electrodes. By symmetry, this is enough to include the effects of the nearest-neighbouring pixels to the readout electrode. As would be expected, the weighting field varies with horizontal position, but not with depth. Figure 2.7 is a horizontal cross-section through the detector showing the weighting potential of one electrode. The dotted red line indicates the pixel surrounding the readout electrode. The average weighting field in the pixel is less than 0.5, so on average electrons drifting to the n-type readout electrode will induce a somewhat larger proportion of the readout signal than holes drifting to the p-type bias electrodes.

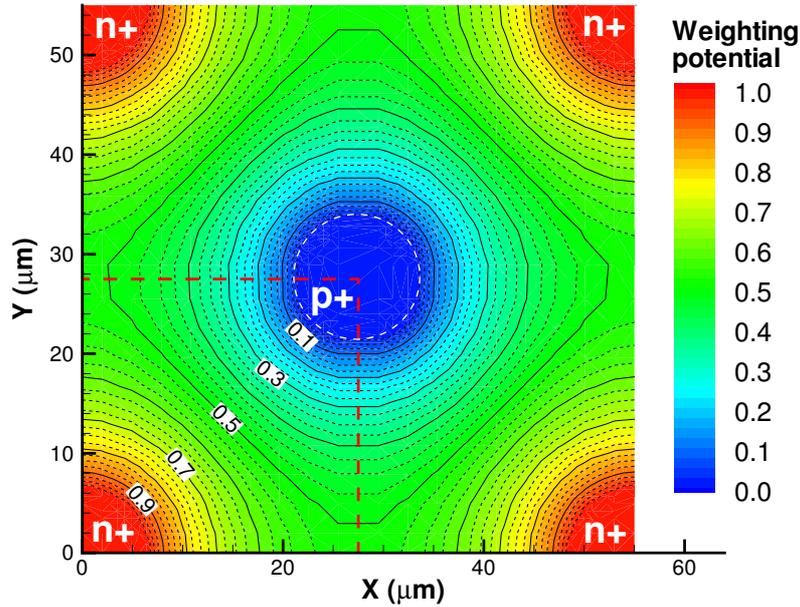


Figure 2.8: Weighting potential in a cross-section of a 3D pad detector. In this configuration all the readout electrodes are connected together.

Figure 2.8 then shows the weighting potential obtained when all of the readout electrodes are connected together to form a single pad detector. Naturally, in this case we will get a readout signal from our single pad regardless of which specific “pixel” is hit, so the weighting field extends throughout the simulated structure. Also, the average weighting potential is 0.5, and the weighting field is the same around the readout and bias columns, meaning that electrons and holes will contribute equally to the readout signal. This is due to the symmetry between the readout and bias columns.

If we look specifically at the region in Fig 2.7 and Fig 2.8 corresponding to a single pixel, as indicated by the red dotted line, there isn’t too much difference in the weighting field. This is because the p-type bias columns surround the readout electrode, meaning that the other n-type readout electrodes have relatively little effect on the weighting field. This in turn implies that the “quarter-pixel” simulation structure used previously should give an approximately accurate weighting field.

2.3 Radiation damage

2.3.1 Mechanisms of radiation damage

When high-energy particles pass through a semiconductor, they can transfer energy to the material in two ways [57]. Firstly, energy can be transferred to electrons in the material, causing ionisation and creating free pairs of electrons and holes, as described in section 1.1.3.

Secondly, high-energy particles can transfer energy to atoms in the semiconductor crystal, by Coulomb interactions or by scattering directly from the nucleus. This is referred to as Non-Ionizing Energy Loss (NIEL). If enough energy is transferred to an atom, it can be ejected from its site in the semiconductor lattice, leaving behind a vacancy. The ejected atom will then lose any excess kinetic energy by colliding with other atoms, which can potentially cause further displacements and so on, creating a cluster of defects. After this damage has taken place, the thermal energy of the lattice will allow these defects to move randomly through the crystal, and many pairs of vacancies and interstitials (atoms that are not located at a lattice site) will recombine. However, a proportion of them will form stable complexes, either with other defects or with other impurities in the crystal [58].

The total number of vacancy-interstitial pairs initially produced is proportional to the non-ionising energy loss. This depends on both the particle type and energy. For example, at energies of below 1MeV, protons have much higher NIEL than neutrons, because the proton can transfer energy to the lattice by Coulomb effects. To compare the damage caused by different particle types and energies, it is generally assumed that the effects of radiation damage are proportional to the NIEL. So, given a radiation fluence in particles per cm^2 passing through a material, we can convert this to the equivalent fluence of 1MeV neutrons per cm^2 (n_{eq}/cm^2).

However, as well as affecting the NIEL, the particle type and energy also affect the distribution of defects caused by radiation damage. The subsequent processes of carrier migration and the formation of complexes will be affected by this. For example, high-energy electrons generally produce single point defects rather than large, dense clusters. For a given NIEL, the macroscopic effects of electron irradiation are much less severe than that for protons or neutrons, because these single defects are less likely to combine to produce stable defect complexes [57].

Since the number of stable traps increases linearly with the radiation fluence, we can express the trap concentration after irradiation as follows, where Φ is the

radiation fluence in $1\text{MeV-n}_{eq}/\text{cm}^2$ and η (cm^{-1}) is the trap introduction rate.

$$N_t = \eta\Phi \quad (2.6)$$

2.3.2 Microscopic defects and macroscopic damage

In a perfect silicon crystal, there is an energy gap between the conduction and valence bands, as described in section 1.1.1. Crystal defects will introduce additional energy states between these two bands, which can either be occupied by electrons or left empty (“occupied by a hole”). Acceptor defects are negatively charged if occupied by an electron, and neutral if empty, and donors are neutral if occupied by an electron, and positive if empty, much like acceptor and donor dopants [59].

There are four processes that can change the state of a trap. An electron can fall from the conduction band to the trap level or from the trap level to an empty state in the valence band, removing a hole. The rates of these processes largely depend on the relative occupation of the conduction band, trap level and valence band. Also, electrons can be thermally excited from the valence band to the trap level, producing a hole, or from the trap level to the conduction band. The rates of these processes are strongly dependent on the energy differences between the levels and the temperature.

The net rate of electron flow from the conduction band to a particular trap level is:

$$R_{c \rightarrow t} = v_{th}^e \sigma_e N_t (n(1 - f_t) - n_i \exp(E_t/kT) f_t) \quad (2.7)$$

The net rate of electron flow from the valence band to the trap level is:

$$R_{v \rightarrow t} = v_{th}^h \sigma_h N_t (n_i \exp(-E_t/kT) (1 - f_t) - p f_t) \quad (2.8)$$

This second expression can also be seen as the hole flow from the trap to the valence band. In the above expressions, v_{th} is the thermal velocity of electrons or holes as indicated by the superscript, $\sigma_{e/h}$ is the trap’s electron/hole capture cross section, N_t is the number of traps, n_i the intrinsic carrier concentration, and E_t the energy of the trap relative to the midgap (i.e. positive if the trap is above the middle of the bandgap, negative if it’s below). f_t is the value of the Fermi-Dirac distribution at the trap level, as discussed in section 1.1.1.

In terms of the macroscopic behaviour of the silicon detector, these traps have three main effects, which are illustrated in Fig. 2.9.

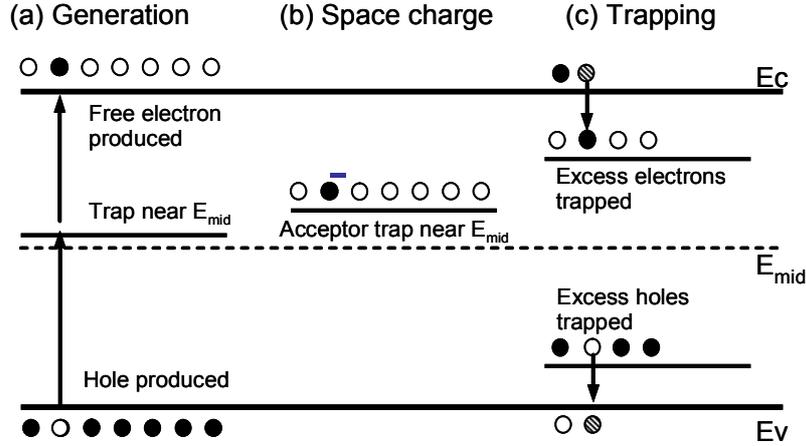


Figure 2.9: Energy diagram illustrating how traps produce leakage current, space charge and trapping of excess carriers.

Increase in leakage current under bias.

In a two-step process, random thermal excitation can cause an electron to jump from the valence band to an unoccupied defect level, and then from this level to the conduction band, producing an electron-hole pair as shown in Fig. 2.9. Under depletion, these carriers will be swept away by the electric field before they can recombine, producing a leakage current.

Using the equations above, the generation rate produced by a trap level in the depletion region will be:

$$G = N_t n_i \frac{v_{th}^e \sigma_e v_{th}^h \sigma_h}{v_{th}^e \sigma_e \exp\left(\frac{E_t}{kT}\right) + v_{th}^h \sigma_h \exp\left(\frac{-E_t}{kT}\right)} \quad (2.9)$$

If the trap level is far from the midgap, then one of the steps in the generation process will involve a large energy jump, and will occur only rarely. As a result, the generation rate falls exponentially as the energy difference between the trap's energy and the midgap increases. Additionally, the intrinsic carrier concentration n_i itself has a strong temperature dependence, so the leakage current increases rapidly and predictably with temperature. Generally, measurements of leakage current are quoted at a temperature of 20°C, using the following scaling factor [2]:

$$I(T_R) = I(T) \left(\frac{T_R}{T}\right)^2 \exp\left(-\frac{E_g}{2k} \left[\frac{1}{T_R} - \frac{1}{T}\right]\right) \quad (2.10)$$

where T is the measurement temperature and T_R is the required temperature.

Even in an unirradiated detector with few defects, this effect is the main source of bulk leakage current [50]. Following damage, the leakage current increases linearly with the radiation fluence, and can be parametrised by α , where:

$$\frac{I}{\text{Vol}} = \alpha\Phi \quad (2.11)$$

If the radiation fluence is scaled by the NIEL, then the increase in leakage current is identical for both neutrons and protons, and independent of the substrate type [60]. Following the damage, the leakage current falls over time due to annealing, so the particular value of α depends on the thermal history of the silicon.

From a practical point of view, the increasing leakage current will increase the power dissipation and current noise in the detector. Potentially, high power dissipation can cause a thermal runaway effect, where both the temperature and current flow increase until the device is destroyed. So, radiation-damaged detectors are generally cooled to reduce their current flow and to remove the excess heat they generate. However, in high-energy physics experiments the material in the cooling system can scatter the particles produced in the experiment and degrade the results. So, the system must be designed to provide adequate cooling while using as little material as possible.

Increase in effective doping concentration.

When a detector is unbiased, it will have zero net space charge. When reverse-biased, the substrate will deplete (see section 1.1.2), forming a space-charge region. In undamaged semiconductors, the charge density, and hence the depletion voltage, will be proportional to the substrate doping. Some radiation-induced traps can also act as donors or acceptors and contribute to space charge in the depletion region, altering the effective doping concentration.

Experimentally, standard p-type silicon wafers show a linear increase in depletion voltage with the radiation fluence, because the radiation-induced defects act as acceptors. In n-type substrates, the depletion voltage initially drops, then the substrate effectively becomes p-type, due to the creation of acceptor defects and the removal of the donor levels which were originally present. After this type inversion occurs, the depletion voltage then grows linearly with fluence as the substrate becomes increasingly p-type. This is shown in Fig. 2.10. At fluences of 1×10^{15} $1\text{MeV-n}_{eq}/\text{cm}^2$, the effective doping concentration is roughly 100 times its original value!

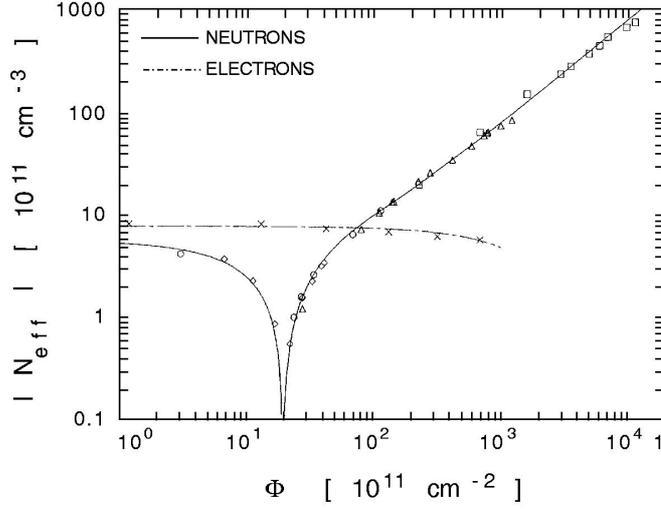


Figure 2.10: Change in effective doping concentration with radiation fluence in an n-type detector, showing type inversion. Reproduced from Ref. [61] with permission.

The effective doping concentration is affected by various factors. Following the damage, annealing causes short-term improvement in the effective doping (beneficial annealing), but in the long-term the effective doping becomes higher (reverse annealing). The increase in effective doping concentration can be lessened by using oxygen-rich substrates [62, 63]. For some substrates, the change in effective doping (for a given 1MeV neutron equivalent fluence) differs between proton and neutron irradiation, too.

Now, consider the behaviour of the defects. If there are acceptor defects above the midgap, most of them will be unoccupied by electrons; the small proportion that are occupied will have negative charge. In the depletion region, the density of negatively-charged traps will be:

$$N_d^- \approx N_t \exp\left(-\frac{E_t}{kT}\right) \left(\frac{n}{n_i} + \frac{v_{th}^h \sigma_h}{v_{th}^e \sigma_e} \exp\left(-\frac{E_t}{kT}\right)\right) \quad (2.12)$$

As with the leakage current, the linear increase in effective doping concentration with fluence can be explained in terms of linearly increasing trap concentration. However, the more complex behaviour of the effective doping concentration implies that more complicated defect chemistry is involved.

When operating a detector, the increasing full depletion voltage with damage leads to higher power dissipation. Generally, it is impractical to bias planar detectors beyond about 1000V, largely due to breakdown effects. Additionally, the

services for the ATLAS detector are only designed to carry 500V, and these services will be re-used for Super-ATLAS if possible [64]. So, at high fluences (typically above 1×10^{15} 1MeV- n_{eq}/cm^2) it can become impossible to fully deplete a $300\mu\text{m}$ substrate, reducing the detector's collection efficiency. Also, if we have a detector with segmented p-type readout electrodes and an n-type substrate, after the substrate type-inverts the p-n junction will appear at the back surface rather than at the readout contacts. To avoid this problem, radiation-hard planar detectors generally use n-type readout.

Trapping of free charge carriers.

In a state of equilibrium, there will be no net change in the occupation of the defect states. However, when extra electron-hole pairs are generated by ionising radiation, they can become trapped by the defects. For example, any defect levels above the midgap will be virtually unoccupied in equilibrium. When extra free electrons are generated, they can lose energy and fall into the unoccupied traps, as illustrated in Fig. 2.9. This will happen at a rate given by:

$$\frac{\partial n}{\partial t} = -nv_{th}^e \sigma_e N_t \quad (2.13)$$

Similarly, free holes will be trapped by energy states below the midgap—to be specific, electrons in these defect states will drop to the valence band and fill these holes.

This will mean that the number of free carriers will decay exponentially over time. Most of the trapped carriers will not de-trap within the time limit set by LHC-speed electronics, which means that the carriers will not be fully collected and the detector's CCE will be reduced. At fluences approaching 1×10^{16} 1MeV- n_{eq}/cm^2 , the mean drift distance will be much smaller than the substrate thickness, and trapping will dominate the CCE performance.

Experimentally, this exponential decay in the number of free electrons can be parameterised using the effective electron lifetime, τ_{eff_e} , as follows:

$$\frac{\partial n}{\partial t} = -\frac{n}{\tau_{\text{eff}_e}} \quad (2.14)$$

A similar result is seen for holes. Furthermore, up to fluences of 1×10^{15} 1MeV- n_{eq}/cm^2 , these effective lifetimes have been shown to vary inversely with the radia-

tion fluence Φ_{eq} [65]. This can be parametrised by:

$$\frac{1}{\tau_{\text{eff}_e}} = \beta_e \Phi_{eq} \quad (2.15)$$

Then, since the defect concentrations should increase linearly with the radiation fluence, we can relate the parameter β_e to the trap parameters by:

$$\beta_e = \sum v_{th}^e \sigma_e \eta \quad (2.16)$$

The summation is done over all the traps above the midgap. Similar equations apply to hole trapping by states below the midgap [66].

Experimentally, in silicon the trapping parameters β_e and β_h are fairly similar [65]. However, since electrons have about three times the mobility of holes, they can travel much further in a given time, and so are less susceptible to trapping effects. As shown by the weighting field calculations in section 2.2.4, the carriers drifting to the segmented electrodes on the front surface of a planar detector make a much larger contribution to the total signal than the carriers drifting to the back surface. So, by using n-type readout, the electrons will contribute more to the total signal, the effects of trapping will be reduced, and the signal will be improved. In 3D detectors, there is much less difference in the weighting field at the readout and bias electrodes, so the readout type will not have such a strong effect on the signal.

2.3.3 Modelling radiation damage in Synopsys TCAD

Synopsys simulates bulk radiation damage by directly modelling the dynamics of the radiation-induced traps [45]. So, it is necessary to select a set of traps which produce the correct macroscopic behaviour and are reasonably consistent with direct experimental measurements of trap types and concentrations.

The radiation damage models used here are based on work done at the University of Perugia [67, 68, 69]. As noted above, the change in depletion voltage with damage is dependent on the source of the radiation damage, the substrate doping type and the substrate production process. Reference [67] contains a model for proton-irradiated, p-type float zone silicon. Close to the interaction point in a collider like the LHC, the fluence will consist mostly of charged hadrons, so this is appropriate [37]. At the very high fluences expected at the SLHC, n-type silicon will type-invert rapidly [61], so the results from simulating p-type at high fluences should be also representative of n-type behaviour, too. The choice of a float-zone substrate is somewhat more arbitrary. Single-crystal silicon wafers are produced by melting

Table 2.1: Original P-type float zone silicon trap model from Ref. [67], and the modified model used in the following simulations.

Type	Energy (eV)	Defect	$\sigma_e(cm^2)$		$\sigma_h(cm^2)$		η (cm ⁻¹)
			Original	Modified	Original	Modified	
Acc.	$E_C - 0.42$	VV	2×10^{-15}	9.5×10^{-15}	2×10^{-14}	9.5×10^{-14}	1.613
Acc.	$E_C - 0.46$	VVV	5.0×10^{-15}	5.0×10^{-15}	5.0×10^{-14}	5.0×10^{-14}	0.9
Don.	$E_V + 0.36$	C_iO_i	2.5×10^{-14}	3.23×10^{-13}	2.5×10^{-15}	3.23×10^{-14}	0.9

and recrystallising high-purity polysilicon. In the Czochralski process, this takes place in a silica crucible, whereas in the float-zone process a rod of polysilicon is recrystallised in an inert atmosphere by slowly passing RF heating coils along its length. Float-zone wafers are commonly used when producing detectors, due to their low impurity concentration [2]. However, standard float zone does have a low oxygen content, whereas substrates with a higher oxygen concentration show a smaller increase in bias voltage with damage. So, basing the model on standard float zone will give depletion voltages at the higher end of the possible range.

The original Perugia damage model uses 3 trap levels as described in Ref. [67]. The parameters of these traps are given in Table 2.1. The trap parameters have already been discussed in the previous section, though in the table the trap energy levels are specified relative to the conduction and valence bands rather than the midgap. Note that the width of the bandgap is 1.12eV. The trap concentrations increase linearly with fluence, as specified by η —see Equation 2.6. The two acceptor levels are slightly above the midgap, and so will generate leakage current, increase the effective p-type doping, and trap excess electrons from the conduction band. The one donor level is significantly below the midgap, and will make little contribution to leakage current or effective doping concentration—its main effect is trapping excess holes. This is illustrated in Fig. 2.11.

The damage model was developed by selecting traps based on direct measurements of trap properties from techniques such as Deep Level Transient Spectroscopy [70], then modifying the levels to give a better match to the macroscopic damage effects seen in detectors. The model gives values of leakage current and effective doping concentration that are a good match to experimental results [67].

However, the free carrier trapping behaviour of this model isn't well-matched to experimental results. Experimentally, the trapping rates for electrons and holes have been shown to increase linearly up to 1×10^{15} 1MeV-n_{eq}/cm² [65], as parametrised by $\beta_{e/h}$. Experimentally, these have been measured as $\beta_e = 4.0 \times 10^{-7}$ cm²s⁻¹ and $\beta_h = 4.4 \times 10^{-7}$ cm²s⁻¹. However, using Equation 2.16, the parameters used

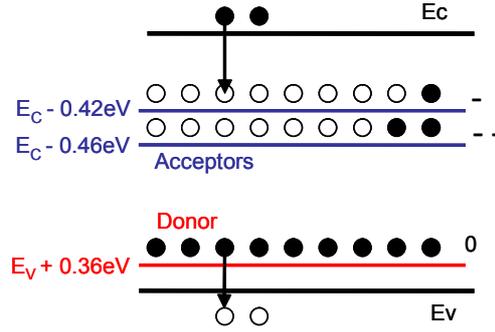


Figure 2.11: Illustration of the 3-trap model from Ref. [67].

in the standard Perugia model give values of $\beta_e = 1.56 \times 10^{-7} \text{cm}^2 \text{s}^{-1}$ and $\beta_h = 3.49 \times 10^{-8} \text{cm}^2 \text{s}^{-1}$. Aside from these values being too low, they will lead to a big discrepancy between the electron and hole trapping rates, invalidating comparisons between different device structures.

So, the model needs to be modified to produce the correct trapping rates, without disrupting the leakage current or effective doping concentration. As shown in Equation 2.12, the charge state of a trap above the midgap is dependent on σ_h/σ_e , whereas the rate of electron trapping will be proportional to σ_e . So, provided that σ_h/σ_e remains constant, the cross-sections can be modified to set the trapping rate without altering the effective doping concentration. Although this will affect the leakage current somewhat, since the leakage current is primarily generated by traps near the midgap, adjusting the cross-sections of the two traps further from the bandgap won't have too big an effect. For the purposes of this simulation, the leakage current only needs to be similar to the correct value, rather than finely-tuned, so small changes are acceptable.

So, the cross sections of the acceptor level at $E_C - 0.42\text{eV}$ and the donor at $E_V - 0.36\text{eV}$ were altered as shown in Table 2.1, to reproduce the trapping rates quoted above. An important point here is that the data on trapping times is only available for fluences below $1 \times 10^{15} \text{ n}_{eq}/\text{cm}^2$. This is because the ‘‘charge correction’’ method used to measure the lifetimes will only work with fully depleted detectors, and at high fluences a detector’s depletion voltage becomes very high, preventing full depletion. So, this model works on the assumption that the linear relationship between fluence and trapping in Equation 2.15 can be extrapolated all the way up to $1 \times 10^{16} \text{ n}_{eq}/\text{cm}^2$.

Since the behaviour of the traps is affected by the width of the bandgap, all the simulations were done using the same bandgap model—Slotboom [71]. These trap

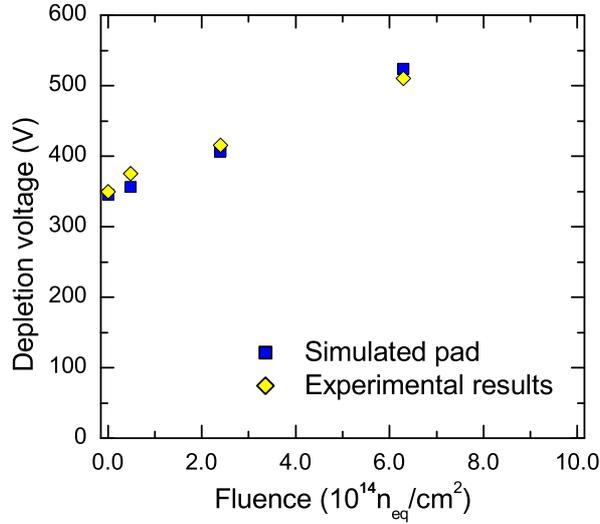


Figure 2.12: Comparison between simulated and experimental depletion voltages in n-in-p planar pad detectors. Experimental results are taken from Ref. [73].

models were designed to work using Synopsys TCAD’s default temperature setting of 300K, whereas during lab tests a detector may operate at room temperature or be cooled to 263K or so depending on the particular test setup. (When they are not being tested, detectors are stored at low temperatures to prevent unwanted annealing.) Experimentally, the leakage current generated by the traps increases predictably with temperature, and leakage current values are normally scaled to 293K to allow a fair comparison [72]. Finally, to model the effects of surface damage, a layer of uniform positive charge was introduced at the interface between the silicon substrate and the oxide layer. This had an area density of $4 \times 10^{11} \text{cm}^{-2}$ before irradiation, and $1 \times 10^{12} \text{cm}^{-2}$ after irradiation [67].

2.3.4 Comparing the damage model to experiment

First, this radiation damage model was applied to planar detectors, to ensure that it gives accurate results. A $280\mu\text{m}$ -thick n-in-p pad detector was simulated with different damage fluences, to determine the variation in depletion voltage and leakage current. The structure and substrate doping of these devices matched those tested in Ref. [73]. The resulting depletion voltages in Fig. 2.12 show a good match between the simulation and experiment.

The leakage current after irradiation is parametrised by α , as shown in Equation 2.11. The simulation gives $\alpha = 5.13 \times 10^{-17} \text{A}\cdot\text{cm}^{-1}$, whereas the experimental value

is $\alpha = (3.99 \pm 0.03) \times 10^{-17} \text{A}\cdot\text{cm}^{-1}$, measured at 20°C following an 80 minute anneal at 60°C [60]. So, the simulated value is about 30% higher than experiment. Given that these simulations are primarily intended to model effective doping concentration and trapping rather than leakage current, and the experimental value of the leakage current can change by more than 30% under different annealing conditions, this result is acceptable.

Next, the charge collection behaviour of a $280\mu\text{m}$ -thick n-in-p strip detector was simulated at different levels of damage. The charge collection was found by starting each simulation with charge deposited along a track passing through the full thickness of the detector, in order to simulate a minimum ionising particle. The resulting readout current was integrated over 10ns, after subtracting the leakage current, to find the total charge collected. The structure and simulation conditions were chosen to match the tests done in Ref. [74]—in particular, all the simulations were done at a high bias of 900V.

The charge collection results, in thousands of electrons, are shown in Fig. 2.13. The simulated CCE values follow the same trend as the experimental values. However, at higher fluences the simulations give substantially lower charge collection. This effect is also seen in other simulation work [75, 76]. Since experimental trapping rates are unavailable above $1 \times 10^{15} \text{ n}_{eq}/\text{cm}^2$, the rates used in this simulation were extrapolated linearly from results at lower fluences. However, experimental charge collection results at very high fluences provide evidence that the trend in the trapping rate is less than linear at high fluences.

In principle, it would be possible to alter the trap parameters to fit these high-fluence experimental results more closely. However, since the charge collection is dependent on different factors—electron and hole trapping rates, and changes in the effective doping—there is the risk of “overfitting” and producing a model that matches a single detector well but generalises poorly to different structures. While it would be possible to take experimental results from a wide variety of device structures, simulate them all, and tune the model to match all the results as well as possible, this would be a time-consuming task. So, the model was left unchanged, and the simulated CCE values should be regarded as a pessimistic estimate.

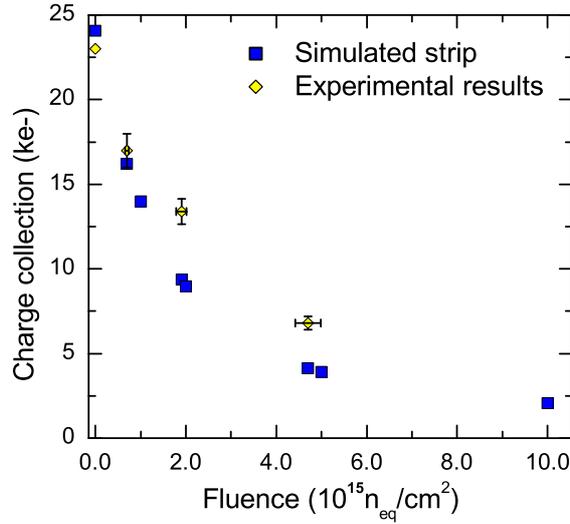


Figure 2.13: Comparison between simulated and experimental charge collection in n-in-p strip detectors. Experimental results are taken from Ref. [74]. Both the simulation and the experimental results used 900V bias.

2.4 ATLAS 3D detectors

2.4.1 ATLAS 3D pixel structures

The ATLAS pixel detector readout chip [77] has $400\mu\text{m}$ by $50\mu\text{m}$ pixels, arranged in a 18×160 matrix. The full ATLAS detector [78] is built using over 23000 of these chips. Each pixel contains an integrating amplifier to measure the total charge produced by a hit on the pixel. This signal is then compared to a programmable threshold, and if a hit has occurred then information on the hit is stored in buffers at the edge of the chip. This information consists of the hit pixel address, a hit time stamp, and also the time for which the signal exceeded the threshold, which gives an approximate measurement of the hit amplitude. If an external trigger is received, all the hits from the corresponding event can then be read out.

In a simple 3D device with small square pixels, each pixel would have a single readout electrode at its centre, and the bias electrodes would be placed at the corners of the pixels. However, the elongated ATLAS pixel size makes it possible to use a variety of different layouts, as illustrated in Fig. 2.14. Sets of n+ readout electrodes are joined together, effectively grouping several smaller sub-pixels to form each ATLAS pixel. The number of n+ columns per pixel can be varied from 2–8 by altering the spacing between the electrodes. Note that the ATLAS chip was only

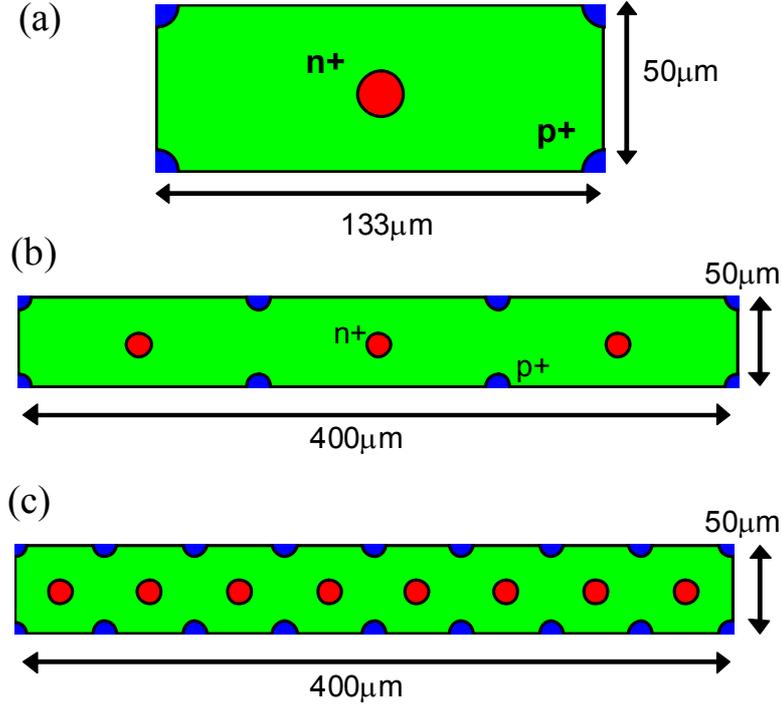


Figure 2.14: 3D pixel layouts showing: (a) a simple sub-pixel with one n+ readout column; (b) an ATLAS 3D pixel consisting of 3 of these simple sub-pixels, as discussed in section 2.4.2; (c) an example of another ATLAS pixel layout using 8 n+ readout columns.

designed to work with n-type readout detectors, so structures with p-type readout cannot be used.

In the next section, results from the 3-column structure (Fig. 2.14b) are compared to experimental results. Then, in section 2.5, all these possible structures are simulated with $10^{16}n_{eq}/\text{cm}^2$ damage, to compare their performance.

The meshes used in the simulations have the same basic structure as the mesh in Section 2.2, except that the electrode spacing was changed as appropriate. As discussed previously, the simulated region consists of the simplest repeating unit of the device, containing 1/4 of an n+ column and 1/4 of a p+ column. Figure 2.15 shows the simulated 3-column structure as an example.

2.4.2 Introduction to simulating radiation-damaged ATLAS 3D detectors

Reference [39] presents recent experimental charge collection results from a 3D ATLAS pixel detector with 3 n+ columns per pixel, following radiation damage.

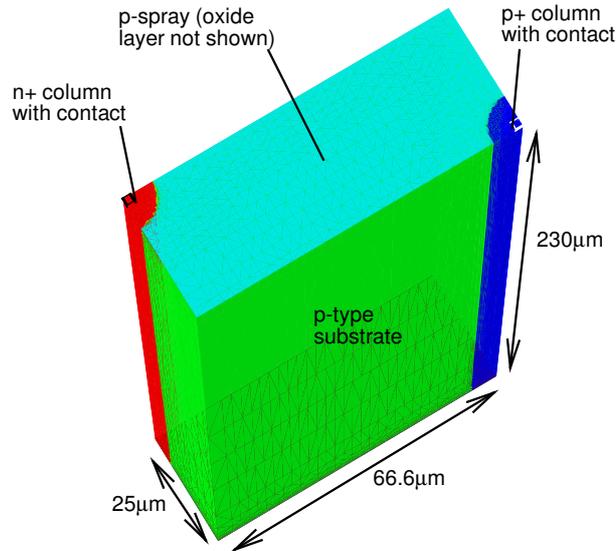


Figure 2.15: Simulated structure used in the “3-column” ATLAS pixel detector simulations. The oxide layer is not shown.

The experimental tests were done using an n-type float-zone substrate device with n+ readout electrodes. After irradiation with neutrons, the substrate type-inverted, becoming effectively p-type. The following simulations, however, have been done using a p-type substrate and n+ readout, as shown in Fig. 2.15. Since the radiation fluences here are very high, the difference in substrate type shouldn't have too much effect, since these n-type substrates will have type-inverted. Aside from this difference, the simulation mesh closely matched the real device.

This mesh was used to do a series of charge collection simulations. The experimental tests had been done using a de-focused, pulsed 1060nm IR laser, which should provide fairly uniform charge deposition throughout the pixel volume. The resulting charge collection was then scaled to find the corresponding signal that would be produced by a minimum ionizing particle. In the test, a large number of pixels were connected together, which means charge sharing would have had no effect, and the current signals measured were averaged over 1000 pulses. So, the simulation used uniform charge generation throughout the device volume, in order to match these conditions. The total number of electron-hole pairs generated was 18500, to match the total charge that would be generated by a MIP passing through the detector. As in the experiment, the current signal was integrated over 10ns, after correcting for the leakage current. Lastly, since the experimental tests were done using different biases at different fluences, the simulations also used varying biases. The charge collection results are shown in Fig. 2.16, and each data point is

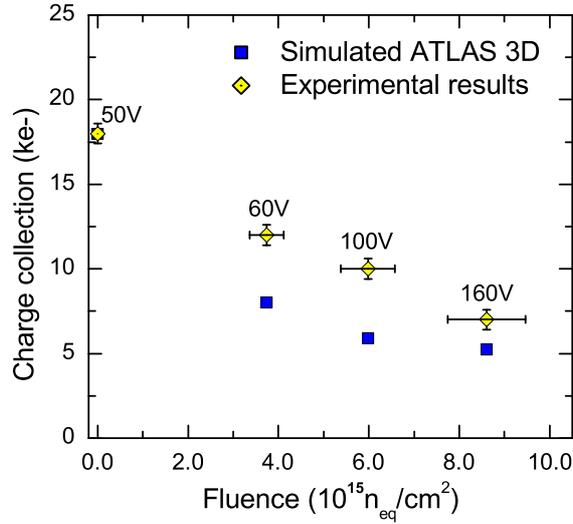


Figure 2.16: Comparison between simulated and experimental charge collection in a “3-column” ATLAS pixel detector. Experimental results are taken from [39]. The labels indicate the bias used in both the experiments and the simulations.

labelled with the corresponding bias voltage.

Once again, the simulation results show a similar trend to the experimental results, but give lower charge collection. The likely reasons for this have been discussed in section 2.3.4. In the planar detector simulation (Fig. 2.13) the simulated collection efficiencies at high fluence are about 60%–70% of the experimental values, and the same is true for the 3D detector, even though the two devices have very different structures. This indicates that these simulations are useful for comparing and understanding the behaviour of different device structures, even if the absolute value of the CCE is lower than it should be at high fluences.

2.5 Comprehensive simulation of radiation-damaged ATLAS 3D detectors

In this section, the different possible ATLAS 3D structures (with 2 to 8 columns per pixel) are simulated at radiation damage levels of $10^{16} n_{eq}/cm^2$, to compare various aspects of their performance such as depletion voltage, charge collection behaviour, and capacitance.

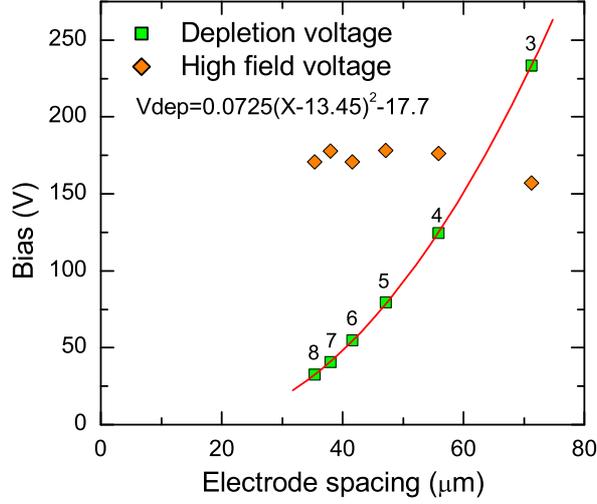


Figure 2.17: Simulated depletion voltages and high-field behaviour in ATLAS pixel devices at $10^{16}n_{eq}/\text{cm}^2$. The “high field voltage” refers to the voltage at which the maximum field in the device reaches $2.5 \times 10^5 \text{V}/\text{cm}$. The data points are labelled to show the number of n+ columns per pixel in each device.

2.5.1 Depletion voltage and high-field behaviour

Using a fluence of $10^{16}n_{eq}/\text{cm}^2$, a steadily increasing bias was applied to each of the 3D structures. The full depletion point of each detector was found using the point where the simulated leakage current levelled out. (The generation current in the simulations is dependent only on the depleted volume, and not the field strength.) The resulting depletion voltages are shown in Fig. 2.17. In this figure, the x-axis gives the distance between the centres of neighbouring n+ and p+ electrodes, as shown in Fig. 2.14. Each data point is also labelled with the corresponding number of n+ columns per full ATLAS pixel. The result for the 2-column structure is omitted, because it is excessively high. The devices with spacing of less than $40\mu\text{m}$ are fully depleted at less than 50V. A quadratic fit has been made to the data, and the depletion voltage increases with $(\text{electrode spacing} - 13.5\mu\text{m})^2$. Given that each column is $5\mu\text{m}$ in radius, and has a doping profile extending a further $1.5\mu\text{m}$ or so, this quadratic factor is roughly equal to the distance from the edge of the n-type doped column to the p-type.

Of course, although the devices with more widely-spaced columns have higher depletion voltages, it might also be possible to apply a higher bias to them without breakdown occurring. As an estimate of this, Fig. 2.17 shows the point where the maximum field in each device reached $2.5 \times 10^5 \text{V}/\text{cm}$ (the breakdown field in silicon is

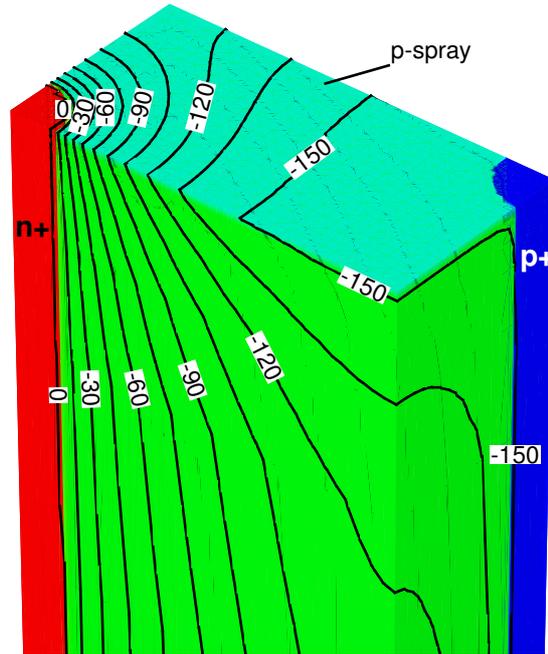


Figure 2.18: Surface behaviour in a 4-column ATLAS pixel device at $10^{16}n_{eq}/\text{cm}^2$. The contour lines show the electrostatic potential, and shading indicates the doping. The n+ readout electrode is held at ground, and a bias of -150V is applied to the p+ electrode. Notice how the bias applied to the p+ column also falls across the p-spray.

$3 \times 10^5 \text{V}/\text{cm}$). Somewhat surprisingly, this point is reached around 170V in all of the devices, regardless of electrode spacing. This is also the same as the value obtained for an unirradiated 3D detector with saturated oxide charge back in section 2.2.2.

As can be seen in Fig. 2.15, the p-spray in this device makes contact with both the n+ and p+ electrodes. As a result, the bias applied to the p+ column is also applied to the p-spray, resulting in a high potential gradient where the p-spray meets the n+ column. This can be seen in more detail in Fig. 2.18, which shows the electrostatic potential in the 4-column ATLAS 3D device with a fluence of $10^{16}n_{eq}/\text{cm}^2$ and a bias of 150V. Ultimately, this means that the behaviour of the high-field region around the n+ column is much the same regardless of how far away the p+ columns are. In 3D detectors with p+ column readout and no isolation, a comparable effect is seen; the layer of electrons at the oxide interface makes contact with both sets of columns, the bias on the n+ columns is also applied to the electron layer, and so a high-field region develops where the electron layer meets the p+ column—see section 4.2.5.

As discussed in section 2.2.2, when the oxide charge is reduced to a typical pre-irradiation value, the “high-field voltage” is reduced dramatically, to just over 50V.

Reference [39] also reports that an unirradiated ATLAS 3D sensor could not be biased far beyond 50V.

Overall, these results show that using a larger number of columns—at least four per pixel—will allow the 3D devices to be fully depleted at very high fluences without the electric field becoming excessively large. More broadly, since the operating voltage is limited by the high-field region at the edge of the n+ column, it may be possible to improve the breakdown behaviour by altering the isolation. A field plate around the edge of the n+ column might reduce the maximum field at this point. Moderated p-spray, which acts like a combination of low-dose p-spray plus a p-stop, gives improved breakdown voltages in planar sensors [11]. When moderated p-spray is used, after high irradiation the low-dose p-spray does not fully compensate the electron layer, and the high-field region shifts to the edge of the higher-dose p region. In a 3D detector, this effect could be particularly useful; there will no longer be an uninterrupted p-type layer linking the n+ and p+ electrodes, and the breakdown region will appear at a greater radius from the column, where the field will be lower due to the cylindrical geometry. More ambitiously, guard-ring-like structures might give a more gradual drop in potential across the surface between the n+ and p+ columns, particularly in devices with a larger column spacing.

Lastly, it should be noted that the breakdown behaviour of the double-sided 3D detectors in Chapter 4 is quite different. This is because one set of columns makes contact with only the front surface, and the other only the back surface.

2.5.2 Average charge collection efficiency at $10^{16}n_{eq}/\text{cm}^2$

The previous section established that 150V is a reasonable choice of bias for all the different 3D ATLAS devices. So, a series of charge collection simulations were done at this voltage with $10^{16}n_{eq}/\text{cm}^2$ radiation damage, using the same methods as in section 2.4.2. Fig. 2.19 shows that the average collection efficiency increases as the electrode spacing is reduced. This is due to the decreasing collection distance, and the increasing electric field strength. The improvement is substantial, with the charge collected almost doubling as the number of columns per pixel is increased from 3 to 5. The “error bars” in this figure give an estimate of the variation in the collection efficiency with position for the devices with 8, 6, 4 and 3 n+ columns per pixel, as described in section 2.5.3. Devices with a small number of n+ columns show poorer uniformity, relative to the average collection.

Additionally, Fig. 2.20 gives an example of how the charge collection in the 5-column ATLAS 3D device varies with the applied bias, with the results being shown

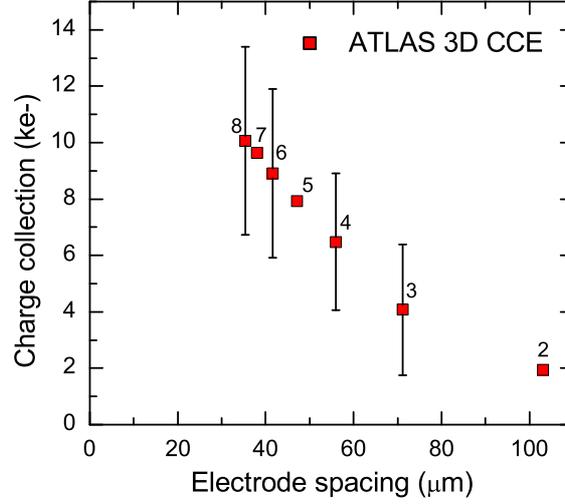


Figure 2.19: Simulated charge collection in different ATLAS pixel layouts at $10^{16}n_{eq}/\text{cm}^2$ fluence. The data points give the average collection efficiency, and the “error bars” give an estimate of the variation in the charge collection with lateral position—see section 2.5.3. This variation was only calculated for the 8-, 6-, 4- and 3-column devices. The applied bias was 150V in all cases, and the charge deposited was 80 electron-hole pairs per micron.

at both $10^{16}n_{eq}/\text{cm}^2$ and $5 \times 10^{15}n_{eq}/\text{cm}^2$. There is a clear change in the gradient of the curves when the depletion voltage is reached, although the charge collection continues to rise as the device is overbiased, due to the increasing electric field. This once again shows the benefits of choosing a device with more columns, in order to reduce the depletion voltage.

2.5.3 Uniformity of the charge collection across the pixel

Using a 3D detector structure means that the electric field and the carrier drift distances vary with the horizontal position across each pixel, rather than with depth. As a result, the charge collection efficiency may vary with horizontal position, too. The significance of this effect will vary depending on how the detector is used; if all the tracks are travelling parallel to the electrode columns, the effect will be particularly important.

Additionally, there is the effect of the electrode columns themselves. As shown in Ref. [26], the collection efficiency is substantially reduced within the columns. Within a column the lack of an electric field means that no signal will be produced unless some of the carriers escape from the electrode by diffusion. Furthermore,

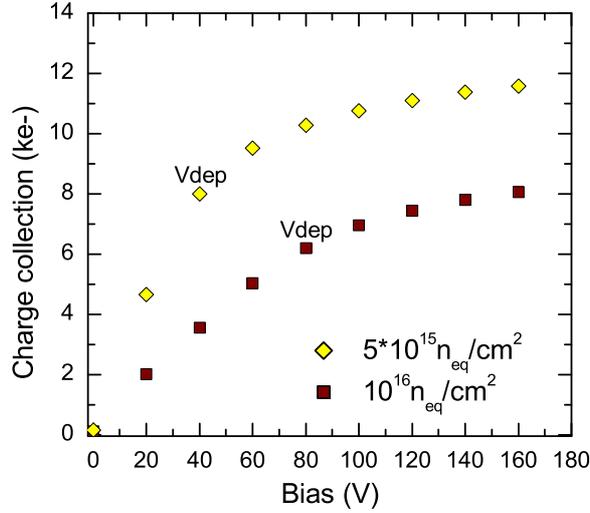


Figure 2.20: Simulated charge collection in the 5-column ATLAS 3D detector against bias at two different fluences. The depletion voltages obtained from earlier simulations are indicated.

some 3D detectors use only partially-filled columns. As the number of columns per cell increases, the total area they occupy will obviously increase. The columns in these simulations are $5\mu\text{m}$ in radius, which is the smallest radius that can currently be achieved with $230\mu\text{m}$ columns, and their doping profile extends a further $1.5\mu\text{m}$. With these conditions, the columns will occupy 5% of the device volume in the 4-electrode device, and 10% if 8 n+ electrodes per pixel are used. If the radius of the columns is increased by just $2\mu\text{m}$, their total area will be doubled. Furthermore, unlike some of the other effects (e.g. trapping, increasing N_{eff}), the column area will reduce the device's effectiveness even before irradiation. So, although the average collected signal at $10^{16}n_{eq}/\text{cm}^2$ steadily increases as more columns are used, it may be better to use a device with a lower average CCE and a smaller column area.

The variation in the collection efficiency with hit position was studied by simulating the effects of minimum ionising particles passing through the detectors. Because this was time-consuming to simulate, only the 8, 6, 4 and 3-column devices were considered. Each minimum ionising particle passed vertically through the device, generating $80\text{ e-h pairs}/\mu\text{m}$. The tracks had a Gaussian lateral distribution, with a standard deviation of $1\mu\text{m}$. These MIPs were simulated at 25 different positions in each detector structure, forming a regularly-spaced 5×5 rectangular grid. The first MIP was deposited $3.5\mu\text{m}$ from the centre of the n+ column, and hence fell within the column, and the final MIP was deposited at the equivalent position within the

p+ column. As before, these simulations were done at $10^{16}n_{eq}/\text{cm}^2$, with 150V bias. Also, this simulation didn't consider the effects of charge sharing at the very edges of each pixel.

For each structure, the charge collection values from the 25 MIPs were averaged. These average values were close to the results obtained earlier in section 2.5.2, with the biggest difference being 6%. So, this shows that while the 5×5 pattern doesn't map the collection efficiency with great resolution, the results are still reasonably representative of each device's behaviour. Next, the standard deviation of each set of MIP results was calculated, as a measure of the variation in collection efficiency with position. The standard deviation is shown in Fig. 2.19 by the "error bars". Relative to the average CCE, the standard deviation is larger in the devices with fewer n+ columns.

Figure 2.21 shows the charge collection with position for different ATLAS 3D detectors. In all of these ATLAS devices, negligible charge is collected when the MIP falls within the n+ or p+ column, as expected. (This confirms that the average collection simulated earlier includes the effects of the columns.) Aside from the columns, the lowest collection efficiencies are seen around the null points, and to some extent around the edges of the cell—particularly the short edges in the devices with fewer columns. Generally, the highest collection signals are seen somewhere midway between the n+ and p+ columns.

Figure. 2.22 shows the electric field distributions in the 6-, 4- and 3-column ATLAS 3D devices for comparison. On the whole, the field is less uniform than the charge collection efficiency. As the electrode spacing increases, the field becomes less uniform along the length of the pixel, with the field being weaker around the p-type column. In particular, the 3-column device is not fully depleted around the p-column, which explains the lower charge collection in this region.

2.5.4 Capacitance and noise

Capacitance simulation

Next, the capacitances of the ATLAS 3D structures were simulated. While the previous simulations used a relatively small region of the device, so that the charge collection simulations would be faster, the capacitance simulations were done using larger meshes. This meant that the C-V simulations were able to find the capacitance between the n+ readout columns and the p+ bias columns, and also the inter-pixel capacitance between one set of n+ columns and those in the two neighbouring pixels. Because ATLAS pixels are so elongated, the interpixel capacitance is much greater

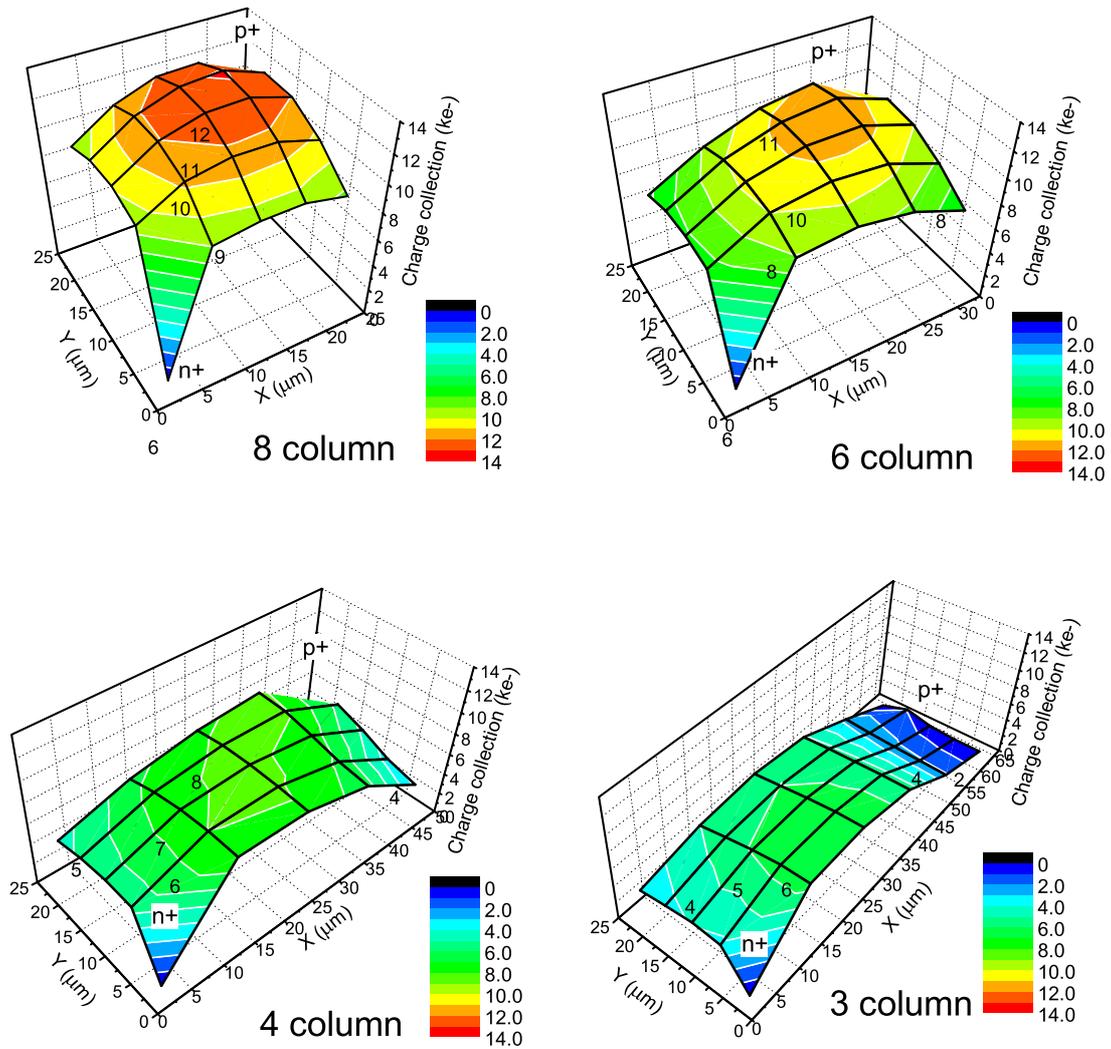


Figure 2.21: Charge collection with position in 8-, 6-, 4- and 3-column ATLAS 3D detectors at $10^{16} n_{eq}/\text{cm}^2$ and 150V bias. The collection is indicated by both colour and height, and the black grid shows the positions of the 25 MIP simulations. Both the n+ and p+ columns show low sensitivity in each device.

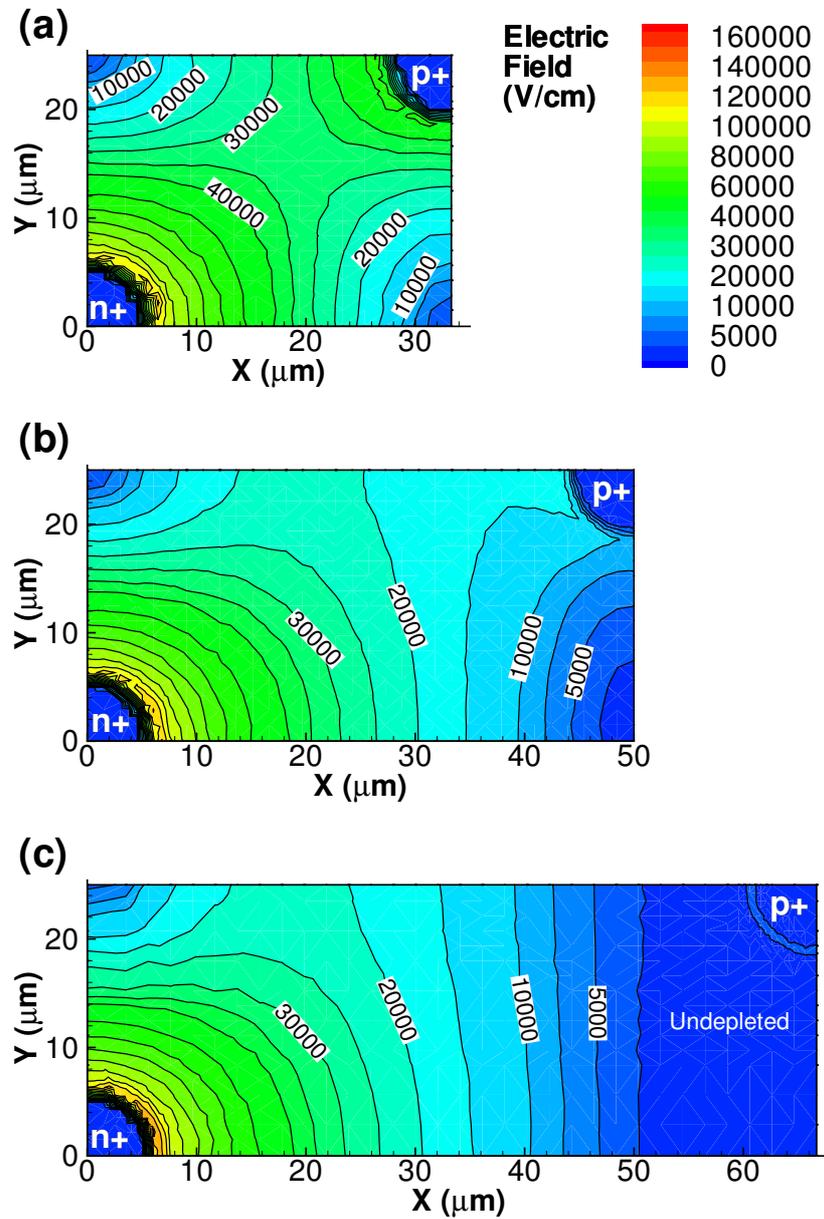


Figure 2.22: Electric field distributions in ATLAS 3D devices with: (a) 6 columns per pixel; (b) 4 columns; (c) 3 columns. In each case the radiation fluence is $10^{16} \text{ n}_{eq}/\text{cm}^2$ and the bias is 150V. The cross-sections are taken from the $z=115 \mu\text{m}$ plane, midway through the thickness of the substrate.

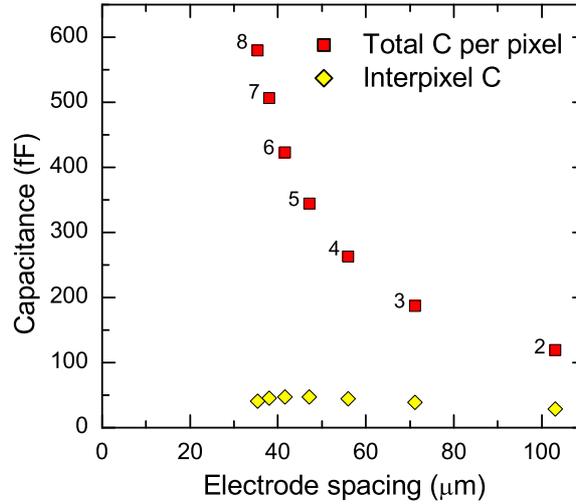


Figure 2.23: Simulated capacitance per pixel and interpixel capacitance in ATLAS 3D devices. No bulk damage was included in these simulations.

between pixels sharing their $400\mu\text{m}$ -long sides than those sharing their $50\mu\text{m}$ -long sides. The simulation does not account for the effects of the n+ and p+ columns in further-away pixels, but these effects should be fairly small. These capacitance simulations were done using saturated surface charge, but no bulk damage, because the capacitance simulations become unreliable when combined with trap modelling. A frequency of 1MHz and a bias of 150V were used.

Fig. 2.23 shows the total readout capacitance seen at each pixel, and also the capacitance between two neighbouring pixels. The total capacitance increases dramatically as the number of columns is increased, e.g. from 265fF/pixel in the 4-column device to 580fF/pixel with 8 columns. This increase occurs not only because there are more n+ columns per pixel to contribute to the capacitance, but also because the capacitance seen at each individual column rises as the n+ and p+ columns become closer together.

The interpixel capacitance in Fig. 2.23 is an order of magnitude smaller than the total capacitance, and doesn't increase very much as the number of columns is increased. Unsurprisingly, the capacitance between adjacent n+ and p+ columns dominates the total capacitance. In contrast, the n+ and p+ contacts in a planar detector are separated by $300\mu\text{m}$ or so of silicon, and the interpixel capacitance tends to dominate.

Effects of capacitance on signal-to-noise ratio

In a pixel detector with binary readout, like the ATLAS chip, a charge signal must exceed the readout chip's threshold value to be registered. The threshold setting itself must be set high enough to ensure that the rate of fake hits due to noise is low. The readout noise of a detector increases with its capacitance, as discussed in section 1.1.4. So, although using a greater number of columns will improve a detector's average collection efficiency, the increasing noise may counteract this.

Aside from this, there are two additional effects that need to be considered. Firstly, the transistors in each channel on the readout chip will not be perfectly identical, meaning that the effective threshold level will vary from pixel to pixel. In addition to having a global threshold setting, the ATLAS pixel chip also allows fine-tuning of the threshold in each individual pixel in order to reduce this effect. However, some variation—referred to as threshold dispersion—will remain [77]. This can be treated like an additional noise source.

Secondly, when a hit occurs on a channel, the channel's amplifier will produce a pulse with a fixed rise time. This means that the slope of the rising edge will vary depending on the pulse height. So, the time at which the amplifier signal exceeds the threshold will depend on the amount of charge collected. In particular, hits which are only slightly above the threshold will be registered later. In the ATLAS experiment, the signal must exceed the threshold within about 20ns to ensure the hit is assigned to the correct bunch crossing. So, for any given threshold setting, the hit must actually exceed a higher “in-time threshold” to be registered correctly [79]. The in-time threshold will equal the threshold level plus an additional “overdrive” signal. Experimentally, the overdrive signal increases with the detector's capacitance [80]. Detailed figures for the relationship between the overdrive signal and the capacitance are not available. To some extent, this is because the overdrive is also strongly dependent on the chip's settings. However, if we assume the overdrive increases with capacitance in a similar way to the noise level [6], then the signal-to-noise ratio will still give a reasonable comparison between detectors.

Tests of unirradiated ATLAS detectors show a linear increase in noise with capacitance [81], with $\text{noise}(\text{electrons}) \simeq 60e^- + 39e^- \cdot (C/100fF)$. Reference [82] reports that the noise in a typical ATLAS module increases only by 15% after irradiation to $10^{15}n_{eq}/\text{cm}^2$, provided that it is cooled to reduce leakage current, so the relationship above was taken as an acceptable estimate of the noise in an irradiated ATLAS 3D sensor. The noise was calculated for each ATLAS layout, assuming that the preamplifier and stray capacitances contribute an extra 100fF, then added in

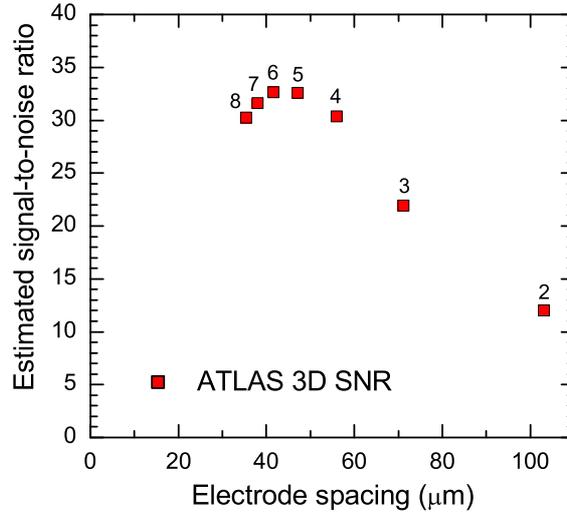


Figure 2.24: Estimated signal-to-noise ratios in ATLAS 3D detectors after $10^{16}n_{eq}/\text{cm}^2$ irradiation, based on previous CCE and capacitance simulations.

quadrature with a threshold dispersion of $70e^-$. Then, the signal-to-noise ratios in these detectors were estimated by combining these values with the average collection efficiencies at $10^{16}n_{eq}/\text{cm}^2$. The results are shown in Fig. 2.24. Although the signal-to-noise ratio increases substantially going from the 3-column to the 5-column device, increasing the number of columns further does not improve the SNR, which actually falls slightly when using 8 columns.

Comparison to simulated planar devices

The collection signal for the 5-column ATLAS 3D device was simulated at varying levels of radiation damage, using 150V bias in each case. The signal-to-noise ratio at each fluence was calculated as described above. Then, the signal-to-noise ratio was found for the n-on-p planar detector simulated in section 2.3.4. This was calculated using a typical planar ATLAS pixel detector noise of 185 electrons [82]. Figure 2.25 shows the signal-to-noise of each simulated device against the fluence. At high fluences, this optimised 3D structure has a substantially higher simulated signal-to-noise ratio.

It is important to note that at low fluences the planar detector has a higher signal-to-noise ratio, due to the greater substrate thickness and the lower capacitive noise. Since it is more difficult and expensive to produce 3D detectors, it will only make sense to use them in environments where their extra radiation hardness

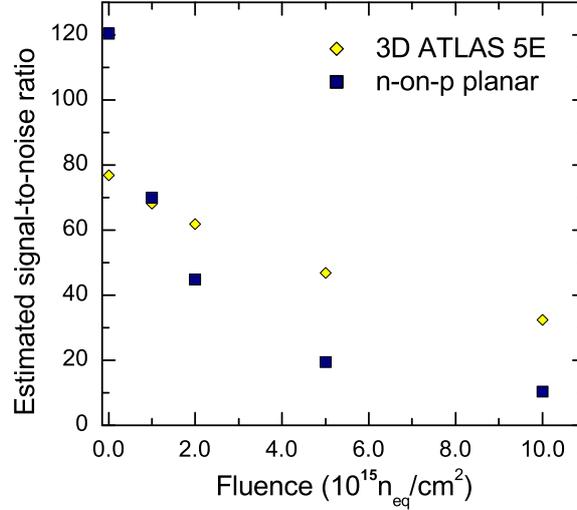


Figure 2.25: Simulated signal-to-noise ratios in an ATLAS 3D 5E detector at 150V and a planar n-on-p detector at 900V at different fluences.

is required—for example, the inner layer of the ATLAS detector at Super-LHC. Additionally, few planar detectors have actually been CCE tested at extremely high fluences of $10^{16} n_{eq}/cm^2$, and some recent tests on n-on-p planar detectors have shown surprisingly high collection efficiencies at this fluence. This is discussed later, in section 4.3.4.

2.5.5 The new ATLAS pixel chip

Recently, design has started on a new ATLAS pixel readout chip for use at the Super-LHC [83]. One change in the chip design is that it will have a pixel size of $250\mu m \times 50\mu m$, rather than $400\mu m \times 50\mu m$. Using smaller pixels will allow the detector to cope better with the increased number of particles produced in each SLHC bunch crossing.

Most of the specific electrode spacings used in the above simulations can no longer be used with this new pixel layout. However, it will be still be possible to vary the electrode spacing by using different numbers of readout electrodes per pixel. So, the trends in the device behaviour with varying electrode spacing seen here will still be valid. By applying interpolation to these simulations it would be possible to estimate properties like the depletion voltage for new electrode spacings.

The reduced pixel size of the new chip will mean that the pixel capacitance will be reduced proportionately, giving a lower noise level for a given electrode

spacing. The improvement in the signal-to-noise ratio will be greater for the high-capacitance devices with a small electrode spacing. So, the optimal electrode spacing will become smaller. This, however, relies on the assumption that the variation in noise with capacitance will remain the same. If the noise behaviour of the new chip is different, then this will also affect the choice of 3D layout. In any case, the capacitance simulation results shown here can still be interpolated and scaled to find the capacitance of a pixel with arbitrary length and electrode spacing, and this in turn can be used to find the noise level.

2.6 Conclusions

The “Synopsys TCAD” software package is able to model the behaviour of semiconductor devices through finite-element simulation. Using this package, the basic behaviour of a 3D detector was simulated. The device showed an extremely low depletion voltage. The electric field is strong at a moderate applied voltage of 100V, due to the small electrode spacing. The electric field itself lies entirely in the horizontal plane between the electrodes, except around the very front and back surfaces. When free carriers are generated in the device, they are rapidly swept to the electrode columns, producing a short signal pulse. For example, in a device with $55\mu\text{m}$ pixels biased to 100V, the pulse lasts for less than 1ns. The weighting field of a particular readout electrode is largely confined to the surrounding pixel, due to the presence of the bias electrodes at each corner of the pixel. Also, the weighting field throughout the pixel is much more uniform than in a planar strip detector. The relative contribution to the total signal from electron and hole drift will depend on the horizontal position of the hit.

3D detectors could potentially be used in high-radiation environments, such as the inner layer of a pixel detector at the Super-LHC, due to their expected radiation tolerance. When silicon is irradiated, defect complexes are created in the crystal, which introduce extra energy states into the bandgap. These defects have three main effects; an increase in effective p-type doping concentration, trapping of free electrons and holes, and an increase in the reverse leakage current. To evaluate the 3D structure’s radiation hardness a damage model based on work done at the University of Perugia was introduced into the simulations. This modelled all three of these effects by directly simulating the behaviour of the radiation-induced defect states. The Perugia model was modified to make the trapping rates of electrons and holes match the measured values more accurately.

This radiation damage model was tested against experimental results from n-on-p planar and full-3D detectors. The model worked reasonably well, but at very high fluences it underestimated the charge collection from both types of detector by 30–40%. This implies that at very high fluences, the rates of carrier trapping no longer increase linearly with fluence, as they do below $10^{15}n_{eq}/\text{cm}^2$. If trapping rates up to $10^{16}n_{eq}/\text{cm}^2$ were measured experimentally, then this model could be improved.

This model has been used to simulate the behaviour of a variety of possible 3D ATLAS pixel architectures at damage levels of $10^{16}n_{eq}/\text{cm}^2$. Overall, the simulations show improved depletion and charge collection behaviour compared to planar detectors. Comparisons between the different 3D pixel layouts have shown that if the electrode spacing is large, then the depletion voltage, the average collection efficiency and the uniformity of the collection across the pixel are poor. The surface effects which can cause breakdown are not strongly affected by the electrode spacing, which means that using more widely-spaced columns doesn't noticeably improve the breakdown voltage. So, getting an acceptable depletion voltage and uniformity at high fluence requires an electrode spacing of about $55\mu\text{m}$ or less. In the $400\mu\text{m}$ by $50\mu\text{m}$ ATLAS pixel, this corresponds to having at least 4 n+ readout columns per pixel. However, using a large number of columns per pixel will mean that the columns themselves will take up a significant proportion of the device volume (particularly if the columns have a greater radius than the $5\mu\text{m}$ used here), and the rapidly increasing capacitive noise limits the signal-to-noise ratio. The improvements in the absolute charge collection and depletion voltage obtained by using an electrode spacing of less than $40\mu\text{m}$ are relatively small. So, for detectors operating at $10^{16}n_{eq}/\text{cm}^2$, the best trade-off is likely to be achieved by having an electrode spacing of $40\text{--}55\mu\text{m}$. This corresponds to 4–6 columns per $400\mu\text{m}$ by $50\mu\text{m}$ pixel.

Chapter 3

Experimental techniques for silicon detectors

This chapter describes a variety of experimental techniques that can be used to evaluate the performance of photodiode detectors, and to investigate their internal behaviour. In Chapters 4 and 6, these techniques are then applied to different 3D detectors.

Section 3.1 describes current-voltage (I-V) testing. Since large leakage currents can degrade a detector's performance and make it more difficult to operate, the magnitude of the leakage current is an important test of device performance. In particular, if avalanche breakdown occurs at a very low voltage, the device may be inoperable.

Capacitance-voltage (C-V) testing is then described in section 3.2. A detector's capacitance can affect some aspects of its behaviour, such as noise. Also, since the capacitance is created by the device's depletion region, the C-V curve provides useful information about the depletion behaviour, for example the full depletion voltage.

Ultimately, of course, these devices are designed to detect particles. To measure the performance of silicon detectors for high-energy-physics, a standard approach is to measure the signal charge collected by a detector when it is hit by a high-energy charged particle. In this chapter, two different charge collection test setups are discussed. Section 3.3 describes a setup which can test simple pad detectors, using a single readout chain built from Nuclear Instrumentation Module standard elements. Section 3.4 describes a second setup for testing strip detectors using electronics from the LHCb experiment.

Different techniques are used to test the imaging quality of X-ray detectors, since X-rays and high-energy charged particles interact differently with silicon. These techniques are described and used in Chapter 5.

3.1 Current-voltage (I-V) testing

At Glasgow, a Cascade Microtech probe station is used to make I-V tests. The probe station uses finely-tipped probe needles to make contact with the metal pads on the detector's front surface. The detector sits on a metal chuck during the test, which provides a contact to the back surface. A Keithley 4200 measurement system acts as a combined voltage supply and ammeter to measure the variation in the current through each needle with the applied voltage. An example of an I-V test on a double-sided 3D strip detector can be seen in Fig. 4.16 in the next chapter.

The most basic function of an I-V test is to confirm that a device actually works as a photodiode, by checking that the reverse-bias current is small compared to the forward-bias current. Assuming that the device can be successfully reverse-biased, there will be two main sources of current flow—bulk current generated within the depleted substrate, and current flow at the surfaces of the detector. As described in section 1.1.5, a detector will usually have a guard ring structure at its edge to collect this surface current. So, the bulk and surface currents can be measured independently, assuming that the guard ring performs adequately.

From a practical point of view, the current flow in a detector should ideally be as low as possible. This is discussed in section 1.1.4. If a detector is DC-coupled, then the bulk current flow in each readout electrode will pass into the readout electronics, and will be integrated by the charge-sensitive preamplifier. If the current flow is too high, the preamplifier will saturate, making it inoperable. In both DC-coupled and AC-coupled detectors, statistical fluctuations in the current will increase the noise level. Also, leakage current flow will lead to power dissipation within the detector, which places greater demands on the power supply and the cooling system. As discussed in section 2.3.2, the leakage current will increase with radiation damage.

When the electric field gets excessively high, avalanche breakdown will occur. So, the I-V tests will also determine the detector's maximum operating voltage.

Aside from these practical concerns, the I-V test does give some information about the device's internal behaviour. Since bulk leakage current is generated in the depletion region, the shape of the I-V curve can give some information about how the depletion region develops.

3.2 Capacitance-voltage (C-V) testing

Capacitance-voltage testing can also be done using a probe station, in a similar way to I-V testing. At Glasgow, the Keithley 4200 system described above controls

an Agilent 4284A LCR meter which performs the C-V measurement. The probe needles simultaneously bias the device and apply a small-amplitude AC voltage to two of the contacts. The amplitude of the resulting small-signal AC current flow can be used to calculate the capacitance between the two contacts, given that $I_c = V_c/Z_c = j\omega CV_c$. A frequency of 10kHz was used, to match RD50 recommendations [84]. An example of a C-V curve from a double-sided 3D pad detector can be seen in Fig 4.17 in the next chapter. Note that many real devices have multiple contacts, and a capacitance will be present between every pair of contacts. For example, if we consider a particular strip on a strip detector, then there will be a capacitance to the back surface, inter-strip capacitances to each neighbouring strip, and also small capacitances to further-away strips.

When a detector is read out with a charge-sensitive preamplifier, the noise level will be proportional to the sum of the detector capacitance and the preamplifier's input capacitance. Additionally, capacitive coupling between neighbouring strips can lead to crosstalk. So, high capacitances will degrade the device performance.

Since the “capacitor” in a photodiode is the depletion region, the variation in capacitance with voltage gives information about the detector's depletion behaviour. In the case of a 3D detector, the depletion behaviour will be more complicated than that of the simple planar diode behaviour described in section 1.1.2. Nevertheless, the point where the C-V curve reaches its minimum indicates the full depletion voltage, and through simulation it is possible to relate the specific shape of the curve to the device's internal behaviour. See, for example, the tests on double-sided 3D detectors in section 4.3.2.

3.3 Charge collection—pad detector

3.3.1 Basic principles

As discussed in section 1.1.3, when a singly-charged particle travelling close to the speed of light passes through a silicon detector, it will deposit a uniform amount of energy per unit distance along its path, and the energy deposited will not vary strongly with the particle's rest mass or energy. So, it's possible to get a good idea of a detector's performance simply by testing it with one type of high-energy particle. By measuring a detector's charge collection efficiency (CCE) versus the bias voltage, it is also possible to find out more about the detector's depletion behaviour, since only the carriers generated in the depletion region will be collected.

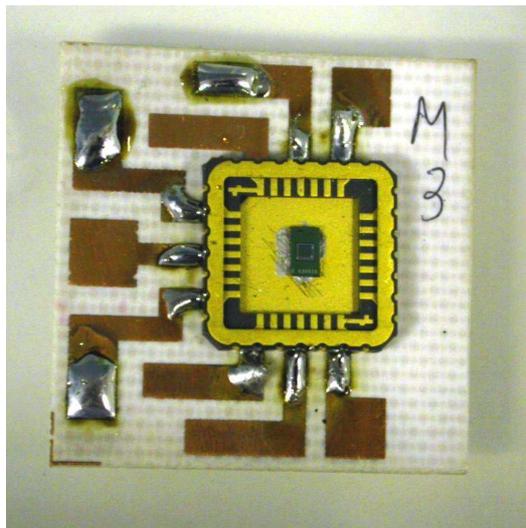


Figure 3.1: A single-type-column 3D pad detector mounted on a chip carrier and PCB.

Although most silicon detectors used in experiments have a strip or pixel structure, for the purposes of basic lab testing it can be convenient to use a pad detector. Since this only has a single, relatively large detection element plus a guard ring, it is comparatively simple to test the detector. Figure 3.1 shows the pad detector tested in section 6.3.2, mounted on a chip carrier. The contacts on the pad detector in the centre are wire-bonded to pads on the chip carrier, which then are connected to the printed circuit board (PCB). The readout chain for testing the pad detector can be built from individual modules such as amplifiers, ADCs and so on, which means that the setup can be used with a variety of different detectors. However, the pad will have relatively high capacitance, and the connections between the different components in the readout chain will be relatively long, so the pad test setup may have greater noise and interference.

3.3.2 Test setup

Figure 3.2 shows a diagram of the pad detector test setup used at Glasgow.

The setup uses a ^{90}Sr source that emits betas with a spectrum of energies up to 2.83MeV. Betas with an energy above about 1MeV will pass straight through the silicon detector and generate 80 electron-hole pairs per micron along their path, much like any other singly-charged particle with high enough energy. These betas can be referred to as minimum ionising particles (MIPs) because their energy loss in the detector is minimised at this energy. The source will also emit lower-energy

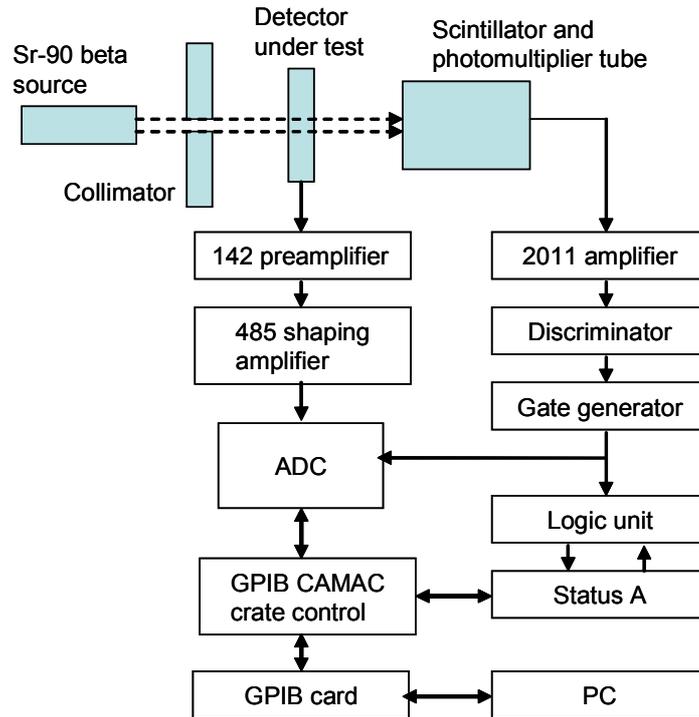


Figure 3.2: Block diagram of the pad detector CCE setup.

betas which will be stopped by the detector and generate a larger, variable quantity of charge carriers.

In the test setup, a collimator is placed between the source and the detector under test, to give a narrower beam of betas. Behind the detector, there is a scintillator, connected to a photomultiplier tube. When the source emits a high-energy beta, it passes through the detector, resulting in a current signal, and is then absorbed by the scintillator. The light the scintillator produces is then detected by the photomultiplier tube, producing a fast signal pulse. This provides the system with a trigger signal, making it possible to measure the signal from the detector under test each time a hit occurs. This trigger signal is also used to exclude the lower-energy beta hits. Firstly, most of the low-energy betas will not pass through the detector to reach the scintillator. Secondly, the scintillator and photomultiplier tube provide a rough measurement of the particle's energy, so betas below about 1MeV can be ignored. These parts of the test setup are housed inside a lead-lined box, to absorb the betas.

The pad detector being tested has its back side connected to a bias supply, and its guard ring connected to ground. The pad itself is connected to a 142 preamplifier,

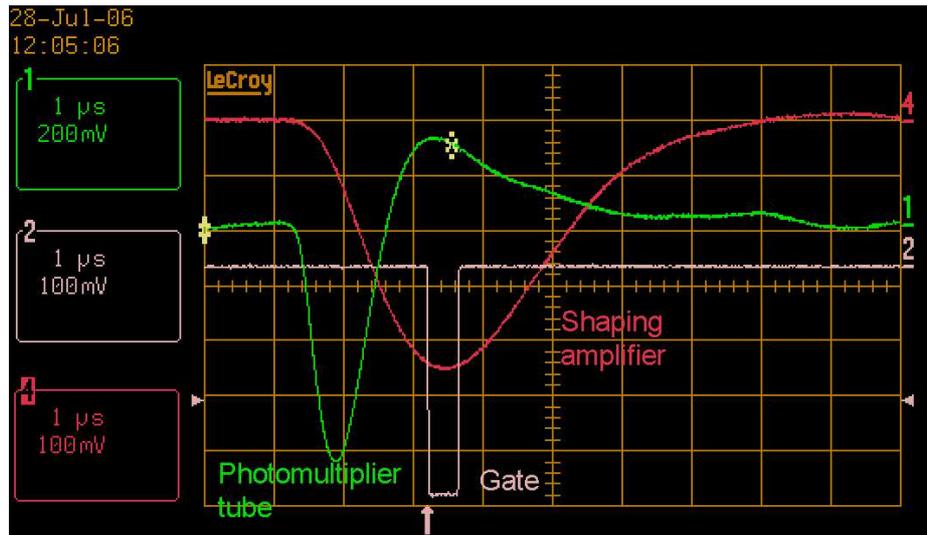


Figure 3.3: Oscilloscope traces of signals produced by the pad detector CCE setup. The figure shows the fast photomultiplier tube signal, the signal produced by the shaping amplifier (with amplitude proportional to the charge collected by the detector) and the gate used to control the sampling of the shaping amplifier’s signal. The time scale is $1\mu\text{s}$ per division.

just outside the test box. This is a charge-sensitive preamplifier which integrates the current it receives from the detector onto a capacitor, as described in section 1.1.4. After this, the capacitor discharges through a resistor. The signal produced by the preamplifier will be a sudden step as the carriers are collected, followed by a slow exponential decay with a $100\mu\text{s}$ time constant.

The following stages of signal processing are carried out by a series of modules. All of these modules are designed to match the Nuclear Instrumentation Module (NIM) standard, which means that they can all be connected together safely and conveniently, they all use the same signal levels, and they can be supplied with power using a standard NIM crate [5].

The signal from the preamplifier passes to a shaping amplifier. Aside from amplifying the signal further, it also applies lowpass and highpass filtering, to reduce the noise and to shorten the long tail-off produced by the preamplifier. The output of this amplifier is semi-Gaussian, and peaks approximately $2\mu\text{s}$ after the initial hit, as can be seen from the oscilloscope trace in Fig. 3.3. (This is still relatively slow, compared to the readout electronics used at the LHC.) This signal then passes to an analogue-to-digital converter.

At the same time, the signal from the photomultiplier is processed. First it is amplified, increasing its amplitude while retaining a reasonably fast rise time. It

is then passed to a discriminator, which will produce a digital signal if the signal amplitude exceeds a threshold. The threshold is set to ensure that only betas with energy of approximately 1MeV or greater are accepted. This digital signal then activates a gate generator. This sends a gate signal to the analogue-to-digital converter, making it sample the amplified signal from the detector. The gate is $0.5\mu\text{s}$ wide, and is delayed so that it coincides with the peak of the pulse. The ADC measures the maximum signal from the shaping amplifier within the gate period, and this digital value is sent to the control PC via a GPIB interface. The interface is controlled by Labview software, which writes the ADC values to a file. The gate generator also produces a “veto” signal, which lasts for about $10\mu\text{s}$. While this veto is active, any further trigger signals that occur will be ignored, in order to avoid pile-up effects.

3.3.3 Data analysis

When a high-energy charged particle passes through a detector, it deposits energy through a series of collisions with electrons in the material. This is a statistical process, which means that the amount of energy deposited in the detector, and hence the signal generated, will vary from hit to hit. However, over a large number of hits, the quantity of energy deposited will follow a predictable distribution. In a thick detector, the number of collisions will be large, and so this distribution will be a Gaussian about the mean value. For a typical silicon detector, which is relatively thin, the number of collisions will be smaller. The energy loss will follow a Landau distribution [85] as shown in Fig. 3.4. This has a long high-energy tail, since collisions can occasionally result in a very large energy transfer to the detector. As a result, the most probable energy loss is different from the mean energy loss.

The standard approach when measuring CCE is to plot a histogram of the measured signal amplitude, and then to fit the Landau distribution to it to find the most probable energy loss. In silicon, the most probable charge signal will be 80 electron-hole pairs per micron if the collection efficiency is 100% [54]. In these pad detector tests, the graphing and analysis program Origin was used to produce the histogram and apply the Landau fit. The Landau distribution itself was generated using code available from CERN [86]. Figure 6.14 in section 6.3.2 shows a typical distribution produced by a single-type-column 3D detector.

The system was calibrated by using a precision pulse generator to inject known charge pulses into the preamplifier, and measuring the signal arriving at the ADC. Due to noise, the recorded ADC values followed a Gaussian distribution, so the peak of the distribution was taken. After repeating this using charge pulses with different

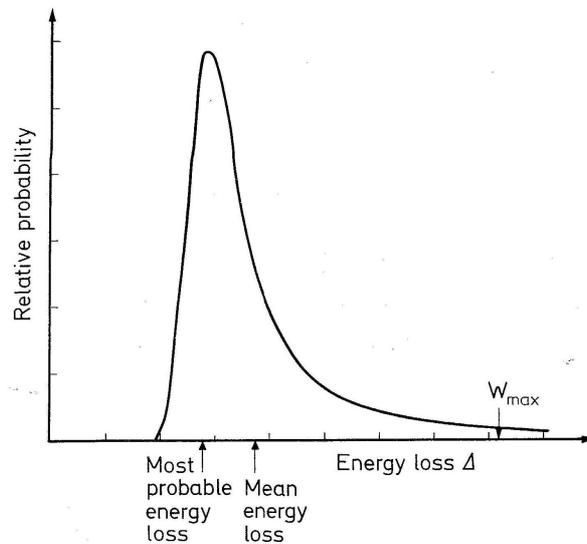


Figure 3.4: The Landau distribution, which describes the energy deposition by a high-energy charged particle passing through a thin detector. W_{max} is the maximum energy that can be deposited in a single collision. Reproduced from Ref. [5] with permission from Springer.

magnitudes, a linear fit was made to this data to find the conversion factor from ADC values to charge collected. The results are shown in Fig. 3.5. The linear fit gives $\text{ADC value} = 108.7 \times \text{Charge}(\text{fC}) + 76.0$. An input charge of zero gives a positive ADC value, which means that the ADC can measure small negative signals too.

3.4 Charge collection—LHC-speed strip detector

The charge collection test system described in the previous section is suitable for testing simple pad devices. However, it is preferable to test full strip or pixel detectors, using readout electronics for a real experiment, in order to evaluate the signal size, noise behaviour and reliability of these detectors. In particular, the readout electronics for LHC experiments need to be fast, to deal with the 40MHz bunch crossing rate. This can lead to effects such as ballistic deficit (loss of signal due to slow collection) which would not appear in the relatively low-readout-speed pad detector CCE setup.

So, a strip detector test setup has been built using electronics from the LHCb experiment [87], which is one of the smaller experiments at the LHC. (Rather than directly searching for new particles, LHCb will measure the production and decay

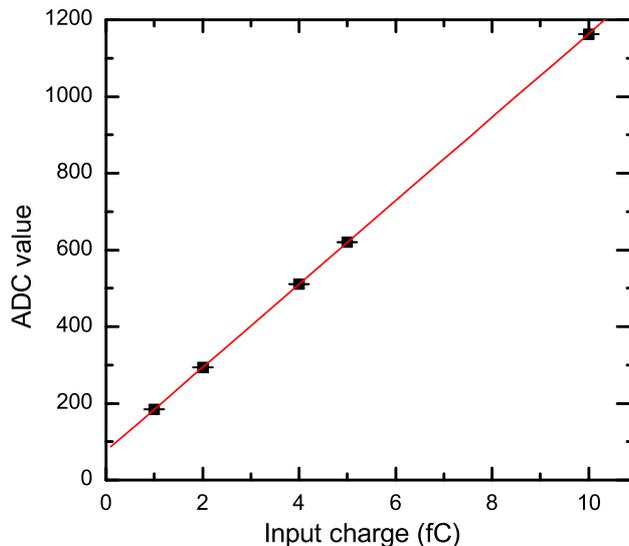


Figure 3.5: Calibration graph for the pad detector CCE setup, obtained by using a precision pulse generator to inject known amounts of charge into the preamplifier and reading out the resulting ADC signal.

rates of hadrons containing B-quarks in order to find indirect evidence of physics beyond the standard model. In particular, it will look for sources of CP violation, which could explain why the universe consists of matter, rather than equal quantities of matter and antimatter.) The setup uses the Beetle readout chip [88], which is used in the silicon vertex locator and silicon tracker, and the TELL1 readout board [89], which is used in most of the LHCb sub-detectors.

A diagram of the strip detector test system is shown in Fig. 3.6. It works by the same general principles as the previous MIP test setup. The beta source is placed in front of the detector, with a scintillator coupled to a photomultiplier behind it. This equipment sits inside an environmental chamber, so that the detector can be cooled. When a MIP passes through the detector and scintillator, a trigger signal is generated, and the TELL1 board reads out the detector and sends the information to the control PC. However, since we have a full detector system rather than a single readout channel, various aspects of the setup are more complicated.

3.4.1 The Beetle readout chip

The Beetle readout chip is described in detail in its reference manual [90]. A schematic is shown in Fig. 3.7. The main section of the figure shows the circuitry within each of its 128 readout channels. At each channel's input, there is a pream-

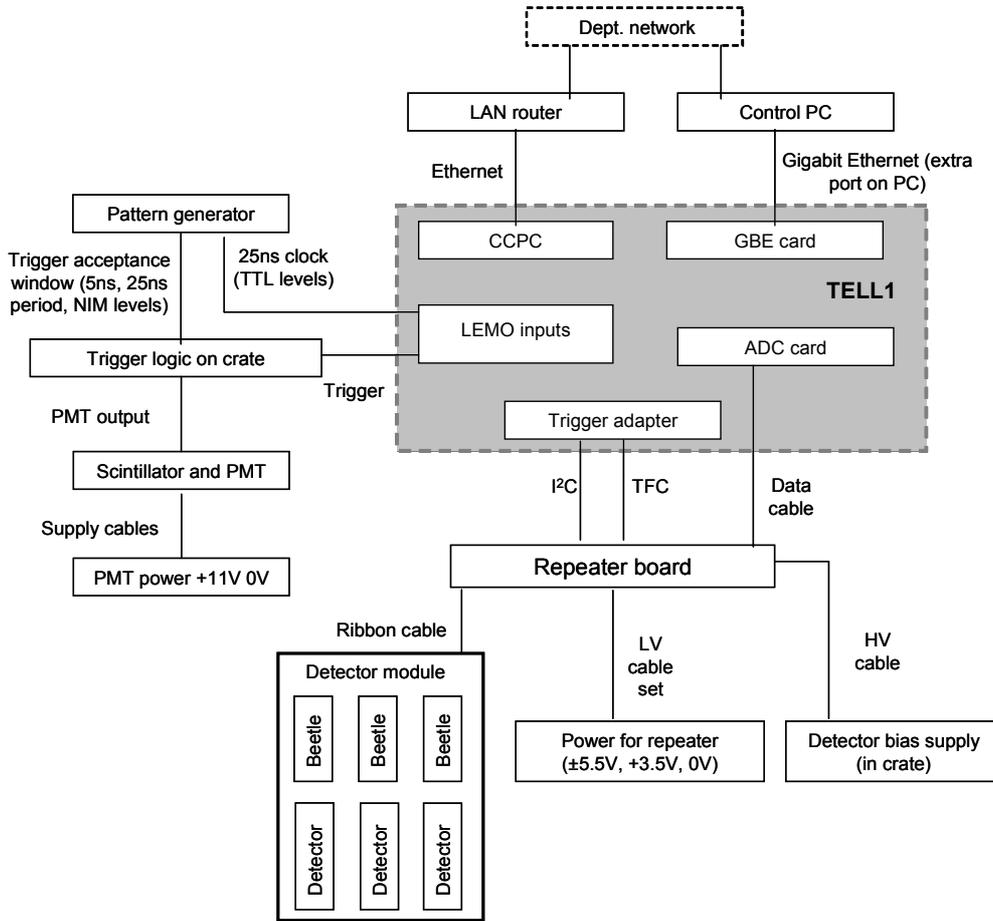


Figure 3.6: Diagram of the strip detector CCE setup, built using LHCb electronics.

plifier followed by a CR-RC pulse shaper. Fig. 3.8 shows the typical pulse shape (taken from the Beetle manual) which has a peaking time of order 25ns, depending on the chip's settings. So, the detector's charge collection time must be shorter than this to avoid ballistic deficit. The amplifiers are designed to cope with input signals of positive or negative polarity.

The chip samples the analogue value of the shaper's output on the clock edge of the Beetle, every 25ns. Each of these samples is then passed into an analogue pipeline, which can store up to 160 consecutive samples. Like the ATLAS experiment (see section 1.4.2), the LHCb experiment does not read out all the detector data from every bunch crossing. Instead, it uses a limited subset of the data from the calorimeters, muon detectors and pile-up system to identify interesting events. Then, it sends a trigger signal to all the detectors, making them read out the stored data from the appropriate bunch crossing. The Beetle's readout time is 900ns. Data

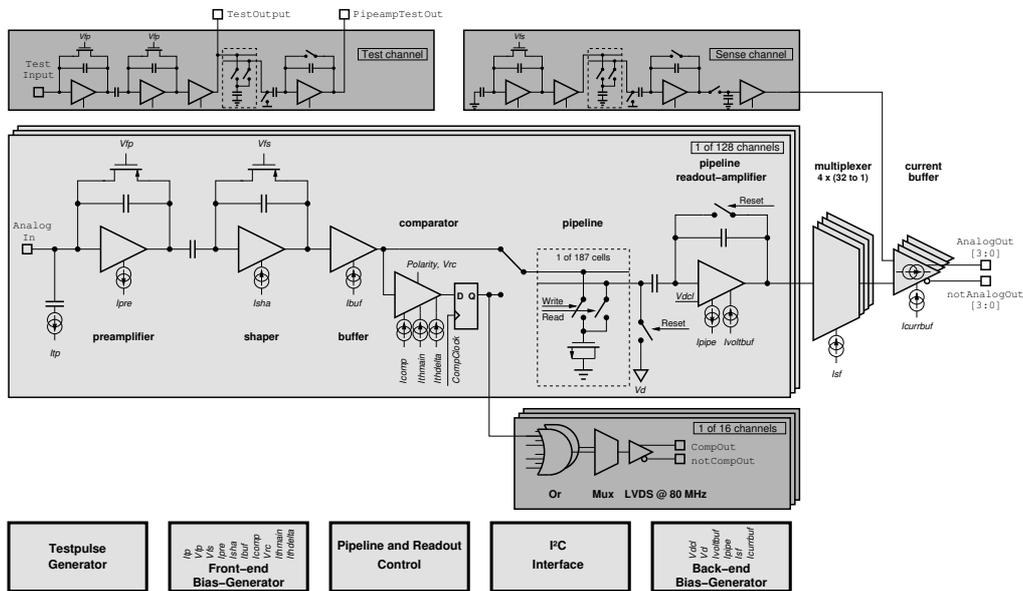


Figure 3.7: Schematic of the Beetle detector, showing the circuitry in each readout channel and the overall architecture. Reproduced from the Beetle reference manual [90].

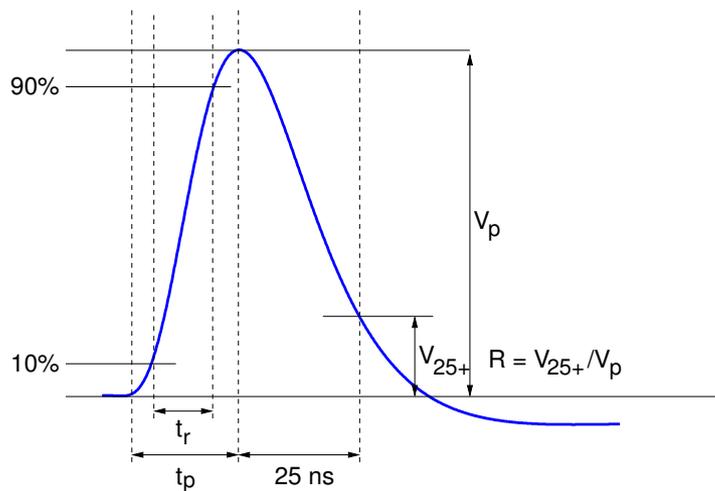


Figure 3.8: The semi-Gaussian pulse produced by the Beetle’s shaping amplifier. Reproduced from the Beetle reference manual [90].

processing can then be done by the TELL1 readout board or in software, as described later.

The chip's use of analogue signal sampling, rather than applying a threshold to create a digital signal, is convenient for measuring charge collection efficiency. However, the Beetle does not have any circuitry to detect the peak of the signal pulse—it simply samples the shaper's output every 25ns. In the LHCb experiment this is not a problem, because the bunch crossings always occur at 25ns intervals, and the Beetle's clock can be synchronised with the collisions. However, in a lab setup, the beta particle hits on the detector will be random with respect to the clock. So, the sample may be taken at any time within a 25ns range around the peak. The solution to this problem is to measure the time of the hit signal on the photomultiplier, relative to the clock, and to only accept the hits which occur with the required timing. In most cases the timing will be chosen to ensure that the peak of the pulse is sampled, but it's also possible to vary the timing in order to study the shape of the pulse. This time discrimination is achieved by using the same pattern generator to produce both the system's clock signal and a "trigger acceptance window", which is a 5ns-wide pulse, repeated with every 25ns, with a fixed phase relative to the clock. The trigger logic on the NIM crate will only generate the trigger signal if the pulse from the photomultiplier coincides with this acceptance window.

Various aspects of the chip's behaviour are adjustable—for example, currents and feedback voltages in the preamplifier and shaping amplifier can be adjusted to alter the pulse shape. These are set by digital values stored in a series of registers on the chip. During these tests, the settings matched those reported in Ref. [91]. The only exception was the latency setting. This must be chosen to match the time delay (in clock cycles) between a hit occurring on the detector and the corresponding trigger signal reaching the Beetle, so that the correct data can be retrieved from the memory pipeline.

3.4.2 The detector module

The strip detectors, Beetle readout chips and various essential electronics are built into a detector module, as shown in Fig. 3.9. This particular module contains an irradiated double-sided 3D strip detector, tested in section 4.3.3. The module is designed to hold three detectors, along with their readout chips and other essential electronics. The module has a thermal baseboard made from carbon fibre, which

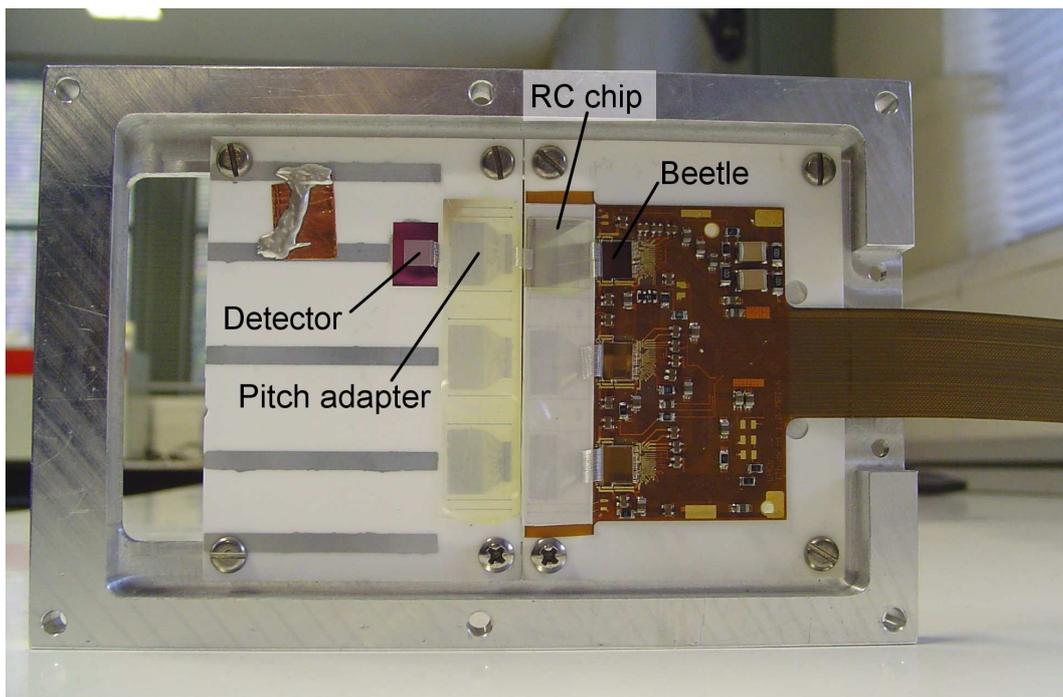


Figure 3.9: Photograph of a strip detector module with an irradiated double-sided-3D detector.

conducts heat away from the components. Each detector is attached to a metal strip by electrically conductive glue, so that the back side of the sensor can be biased.

Each strip on a detector has a contact pad, which needs to be connected to a channel on the Beetle chip. All the strip detectors tested had a strip pitch of $80\mu\text{m}$, whereas the analogue input pads on the Beetle chip are spaced by $40\mu\text{m}$. So, at least one pitch adaptor was required to alter the pitch. The Beetle chip is designed to work with AC coupled sensors, where the connections from the strips to the readout chip are made via capacitors, and connections from the strips to ground are made through resistors. This prevents the detector's leakage current from flowing into the readout chip, as described in section 1.1.4. Although these components are often built into the sensor chip itself, the CNM double-sided 3D devices were not AC coupled. In the irradiated module shown above, an extra chip was added between the sensor and Beetle to provide AC coupling. This chip was provided by J. Härkönen at the Helsinki Institute of Physics, and has a $1\text{M}\Omega$ resistor and 67pF capacitor for each channel. The other modules tested did not use an RC chip. One module used planar strip detectors, which already had built-in AC coupling. Likewise, the set of single-type-column strip detectors tested in section 6.3.3 also had built-in AC coupling. Finally, an unirradiated double-sided 3D detector was tested

without AC coupling. In each case, all the connections between the detectors, the Beetle chips and the other circuitry in-between were made by wire bonds. This wire bonding was done by technicians at Glasgow.

When a module is being tested, it is connected to a “repeater board” by a ribbon cable. This repeater board contains the circuitry required to amplify the control signals being sent from the TELL1 to the module, and the data being sent back from the module to the TELL1. The power supplies required to run the Beetles and to bias the detectors are also sent via the repeater board.

3.4.3 The TELL1 readout board

The TELL1 readout board [89, 92] is designed to control and read out virtually all the sub-detectors at the LHCb experiment. In the LHCb experiment, the boards are located about 60m from the interaction point, behind a shielding wall, and act as the second stage of the data acquisition system, after the readout chips.

In the lab test setup, the TELL1 board can read out up to four detector modules. The board’s main functions are:

- *ADC conversion:* When the TELL1 reads out an event, it needs to read 128 analogue values from each Beetle chip, corresponding to the signal amplitudes on each strip. Each Beetle chip has 4 coaxial connections to the TELL1. So, the strip signals are divided into 4 blocks of 32, and the signals from each block are then sent serially to the TELL1. ADC cards then convert these signals to 10-bit digital values.
- *Data transfer to the control PC:* The digital data values are sent back to the control PC via a Gigabit Ethernet connection, which can give a much higher data rate than standard Ethernet. This is a one-way connection, with the PC passively receiving the data. A piece of code called the “event builder” writes this data to a file.
- *Credit-Card sized PC (CCPC):* The TELL1 is controlled by the CCPC, which is a miniature PC mounted on the board. The CCPC runs Linux, and can be controlled remotely by a normal PC via an Ethernet connection. This allows the user to set the TELL1’s configuration, send signals such as triggers or test pulses, and so on.
- *Programming the Beetles via I²C:* The various registers on each Beetle that control the chip’s performance are programmed by the TELL1 using an I²C connection [93], which passes through the repeater board.

- *Trigger generation:* In the test setup, NIM logic is used to produce a trigger signal every time a hit occurs on the photomultiplier at the correct time with respect to the clock. This trigger signal, and the external clock signal, are passed to the TELL1 via LEMO ports. The TELL1 uses these to generate its own, well-synchronised clock and trigger signals, which are sent to the detector module via the repeater board using the TTC (Trigger, Timing and Control) connection.

In the LHCb experiment itself, the TELL1s will use FPGAs to pre-process the data before sending it to CPU farms via Gigabit Ethernet. (Different algorithms can be used in the TELL1s controlling different subsystems.) This will substantially reduce the quantity of data needing to be transferred and stored, and speed up later event reconstruction. However, in the lab test setup the TELL1 simply sends out the raw data, and the processing is done by the control PC.

3.4.4 Data analysis with Vetra and ROOT

In the lab test setup, data processing is mostly carried out using the program Vetra [94]. This software was originally developed to emulate the data processing algorithms that would be used on the TELL1 board, in order to improve their design. However, it has also become widely-used in lab tests, beam tests and detector commissioning for the LHCb vertex locator (VELO).

The phases of data processing in Vetra are as follows. Each stage is described in more detail in Ref. [92].

- *Finite Impulse Response (FIR) filtering:* The analogue strip signals are transferred serially from the detector module to the TELL1, and consecutive samples can interfere with each other. For example, when a large amplitude signal returns to zero it tends to undershoot, which makes it appear as if there has been a small negative signal on the next strip being read out. This stage applies a correction to each value, calculated from a weighted sum of the current value and the two previous values.
- *Pedestal subtraction:* Since the electronics in each Beetle chip channel are not perfectly identical, the voltage level corresponding to zero signal will vary between them. This level can also vary over time, for example due to temperature changes. The pedestal subtractor calculates the zero level for each strip by taking the average of a large number of previous samples (1100 in this setup). This is then subtracted from the signal.

- *Reordering*: In the LHCb vertex locator, the order in which the strips are read out doesn't match the physical positioning of the strips. This stage is designed to correct for this. For the lab test setup, this section was altered so that the default is to have no reordering.
- *Linear Common-Mode Suppression (LCMS)*: The Beetle strip signals can be affected by interference. For example, the sensor itself may pick up radio frequency interference, or fake signals can be induced by variations in the power supplies. Typically, an interference signal will appear on many strips simultaneously and be dependent on strip position, making it possible to recognise and remove it. For each event, the LCMS stage takes the data from each block of 32 strips, and applies a linear fit to a plot of strip signal versus strip channel. Any outliers from this fit are assumed to be genuine hits, and the linear fit is repeated with these strips excluded. Any difference between this linear fit and the zero-signal level is assumed to be due to common-mode noise, and is subtracted from the strip values.
- *Clustering*: The signal from a particle hit may be shared across multiple strips due to charge sharing, or due to the particle actually passing through more than one strip. The clustering algorithm identifies these clusters, and combines the strip signals in each cluster into a single hit. In lab tests, this is important for measuring the detector's collection efficiency accurately. Initially, all the strips which exceed a certain threshold are found (the default is 3 times the noise level). Out of these, the strip with the highest signal is taken as the starting point of the first cluster. Neighbouring strips are added to the cluster if their signal exceeds 10% of the signal on the central strip. Once the cluster is formed, if its total signal fails to reach a second, higher threshold (4.5 times the noise level) it is discarded. This process is repeated until all the initial seeding strips have been used.

After each stage of the processing, Vetra can write output files containing the signal values, either in the form of histograms or in a more detailed format where each hit is recorded separately. These output files are compatible with ROOT, a C++ based data analysis framework for high-energy physics developed by CERN [86]. These files can be used to find the signal spectrum, and by applying a Landau fit the most probable charge collection can be found. The noise level on each strip can also be determined by applying a Gaussian fit to the data from each strip (ignoring hits) and finding the sigma of the distribution. As mentioned above, in this test

setup it is possible to control the time at which the signal pulses are sampled. So, by doing multiple tests with different timing, it is also possible to find the shape of the signal pulse from the Beetle.

Changes to Vetra

Although Vetra was written by Tomasz Szumlak, some modifications needed to be made to use it with the lab setup. Primarily, this involved altering the output data formats. Firstly, the software is designed to work with the 2048-strip LHCb Vertex Locator, which has a complex strip ordering, so changes had to be made to get detector-by-detector information without any re-ordering. In the LHCb experiment, processes like clustering are used to minimise the amount of data needing to be transferred and stored. So, the output stages were modified to increase the amount of data available. Also, the version of Vetra used in this tests was designed to process the data from each bunch crossing independently, making it difficult, for example, to look at the full pulse shape on a particular strip over 5 consecutive samples. So, Vetra was modified to add additional data to each entry of the output file, which would make it possible to carry out this kind of processing using ROOT scripts afterwards.

When the strips are being read out to the TELL1 in blocks of 32, the data transfer begins with some “header” information. The header signals have a large amplitude compared to the analogue strip values. As a result the first two samples of the block, which are sent to the TELL1 immediately after the header, suffer from a great deal of interference. So, it was necessary to modify Vetra to exclude these strips from the later stages of analysis such as LCMS correction and clustering.

Calibration tests with a planar detector

To calibrate the system’s response, it was used to test a planar detector. The planar detector had a $300\mu\text{m}$ -thick p-type float zone substrate, with n-type readout. There were 128 strips, each 1cm long, and the inter-strip spacing was $80\mu\text{m}$. During the tests, the detector was biased to 120V, to ensure that it was fully depleted. The detector was mounted on a module as described above. Due to a problem with the pitch adaptor, 1 out of every 4 strips had to be left unbonded. During analysis, only hits occurring on the central strip out of each set of 3 bonded strips were used.

Initially, the beta source tests were used to find the pulse shape of the detector. Five test runs were taken, varying the phase of the sampling clock relative to the hit time in 5ns increments. Whenever a hit occurred, 5 samples were taken at 25ns

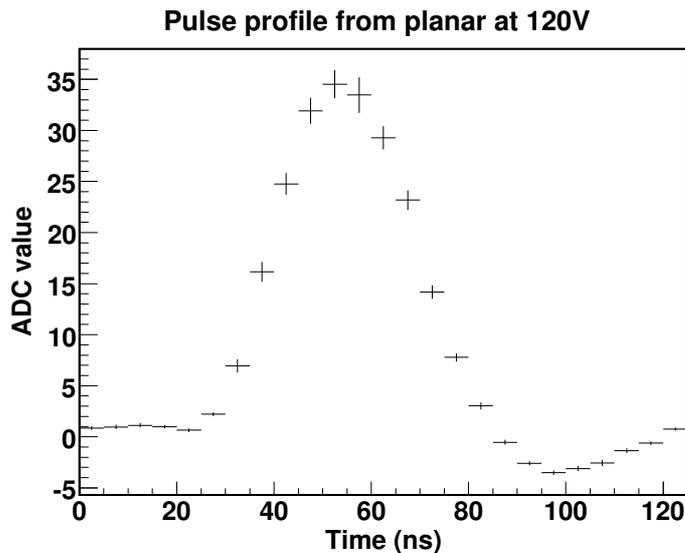


Figure 3.10: Pulse shape from beta tests on a $300\mu\text{m}$ -thick n-on-p planar strip detector, biased to 120V. The offset of the time axis is arbitrary.

intervals. So, these five runs provided measurements of a large number of pulse shapes over a period of 125ns, with 5ns resolution. In each time bin, the average signal amplitude was calculated. The resulting pulse shape is shown in Fig. 3.10. The shape of the signal pulse is very close to the expected shape shown in Fig. 3.8, with a peaking time of 30ns and a clear undershoot 65ns after the hit. (Note that these tests used a slightly lower level of amplification than the following signal spectrum tests on the detector. In all the 3D detector tests done later, the Beetle settings were kept constant to ensure accurate calibration.)

After establishing the pulse shape, the setup's timing was chosen so that the pulse would only be sampled at its peak. After measuring a large number of hits, a histogram of signal sizes was plotted, as shown in Fig. 3.11. Note that this histogram shows the signal sizes obtained after the clustering algorithm in Vetra was used to combine the strip signals from charge-shared hits. The spectrum was then fitted with a Landau distribution, which was convolved with a Gaussian to compensate for any broadening of the spectrum due to noise. The fit matches well with the data. By assuming that the most probable value of the ADC signal obtained here corresponds to 100% charge collection (24000 electrons in a $300\mu\text{m}$ detector), these results were used to calibrate the 3D detector tests in the next chapter.

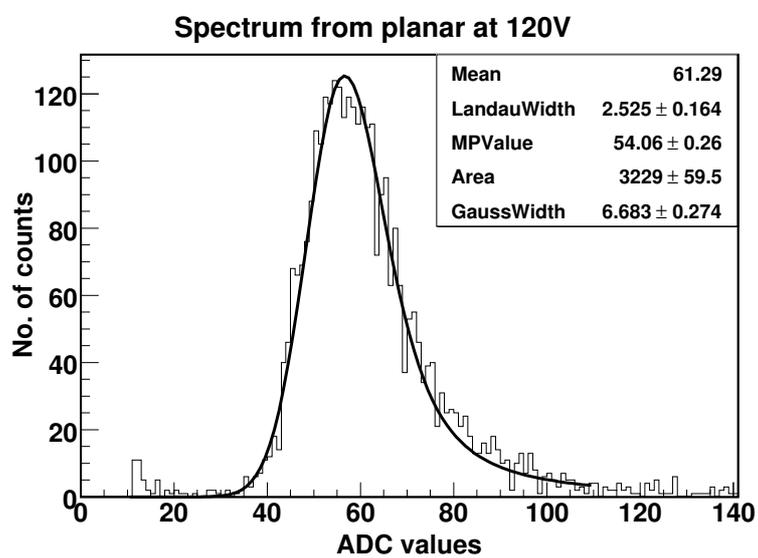


Figure 3.11: Signal spectrum from beta tests of a $300\mu\text{m}$ -thick n-on-p planar strip detector, biased to 120V. The spectrum is fitted with a Landau distribution convolved with a Gaussian.

Chapter 4

Double sided 3D detectors from CNM

CNM-IMB (Centro Nacional de Microelectronica, Instituto de Microelectronica de Barcelona) have proposed and fabricated an alternative “double-sided” 3D detector structure, where the two sets of electrode columns are etched from opposite sides of the substrate [20]. This structure is intended to make certain aspects of the fabrication process easier. CNM have worked with Glasgow to design and test these detectors.

Section 4.1 describes this structure and how it is fabricated, and gives details about the devices that CNM have produced. Section 4.2 then presents a series of simulations focusing on the differences between this structure and full-3D. This includes its charge collection behaviour before and after radiation damage, and its breakdown behaviour. The results of experimental tests performed on these detectors at Glasgow are given in section 4.3, including basic electrical characterisation and charge collection tests on strip detectors connected to LHC-speed readout electronics.

The X-ray detectors tested in Chapter 5 were also produced using this structure.

4.1 Double-sided 3D detectors

4.1.1 Structure and fabrication

The structure of a double-sided 3D detector is shown in Fig. 4.1. The two sets of columns are etched from opposite sides of the substrate, and neither set of columns passes through the full substrate thickness. The columns from the front side can

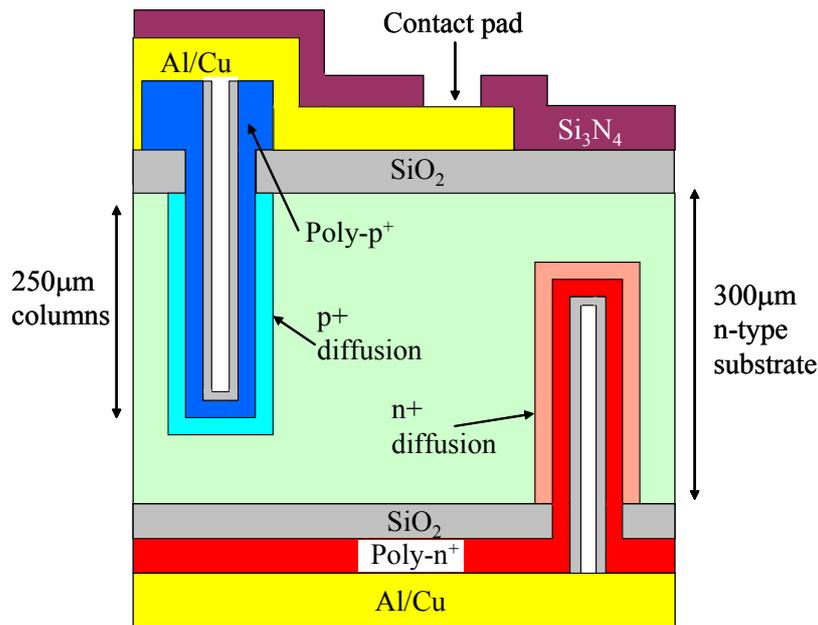


Figure 4.1: Structure of a double-sided 3D detector with p-type readout and an n-type substrate. The columns have been etched and partially filled with polysilicon. The interiors of the columns have been doped by diffusion through the polysilicon, so the highly-doped regions extend into the bulk silicon. The interiors of the columns have then been passivated with silicon dioxide.

be individually coupled to readout electronics. The columns from the back side are shorted together, and are used to bias the detector.

Generally, the fabrication process is very similar to the full-3D process described in section 1.2.2. Alignment marks must be added to the front and back surfaces of the wafer to ensure correct alignment of the two sets of columns, which introduces extra steps. However, since the double-sided approach avoids producing both sets of columns from the same side of the substrate, the column fabrication process can be made simpler, as described in more detail below. Additionally, depositing polysilicon onto a wafer after etching holes through it is a risky process; the polysilicon will place the wafer under stress, which can lead to it bending or breaking. The double-sided 3D process reduces this risk by only etching partway through the substrate, using thinner polysilicon layers (see below), and depositing the polysilicon once on each side of the substrate rather than twice on the same side. In turn, this makes it possible to do the fabrication without using a support wafer.

The fabrication process is described fully in Ref. [20]. After the alignment marks have been added, the columns on the back side of the device are produced. Like with a full-3D detector, deep holes are etched into the silicon with an inductively-

coupled plasma etcher. In these first devices produced by CNM, they are $250\mu\text{m}$ long, compared to the $300\mu\text{m}$ substrate thickness, and have a diameter of $10\mu\text{m}$. This 25:1 ratio of hole depth to diameter is about the maximum that can currently be achieved by ICP etching. Then, a $3\mu\text{m}$ -thick layer of polysilicon is deposited onto the wafer, which coats the inside of the columns and covers the oxide layer on the back surface. Next, the columns are doped by diffusion through the polysilicon layer. However, the columns are not completely filled. Instead, a layer of silicon dioxide is deposited inside them for protection, as shown in Fig. 4.1. In contrast, in a full-3D detector the columns need to be filled to leave a flat surface; otherwise, problems could occur when etching and filling the second set of columns [19]. Since the polysilicon layer covers the back surface of the substrate, this means that the bias columns are already connected together as required.

After this, the columns on the front side are etched, partially filled, doped, and passivated with silicon dioxide. The layer of doped polysilicon is selectively etched away, so that the columns are separated. A metal readout pad can then be added to each strip or pixel, before passivating the rest of the surface with silicon nitride. Since only one polysilicon layer is used on the front surface, it is easier to ensure that the front surface remains reasonably flat for bump-bonding. The back surface is also coated with aluminium, to allow the device to be biased. In a full-3D detector, the connections to both sets of columns need to be made on the front surface, so the metal layers and readout connections will be more complicated.

A downside of this double-sided process is that it is less compatible with active edges (which are discussed in section 1.2.3). The active-edge fabrication process inevitably requires a support wafer, and involves etching trenches through the full substrate thickness, so adding an active edge will remove many of the advantages of the double-sided approach. So far, the devices produced by CNM have not used active edges. However, a 3D detector's depletion region will tend to extend less far from the detector array than that of a planar detector. So, even without active edges, it may be possible to obtain a smaller dead area with a double-sided 3D detector than with a planar detector.

4.1.2 Devices produced by CNM

In December 2007, CNM finished their first fabrication run of these detectors. This consisted of two 4" n-type wafers with p-type readout, matching the structure shown earlier in Fig. 4.1. (Two more wafers broke during fabrication.) The wafers included pixel detectors that matched the Medipix2, Pilatus and ATLAS readout chips, along

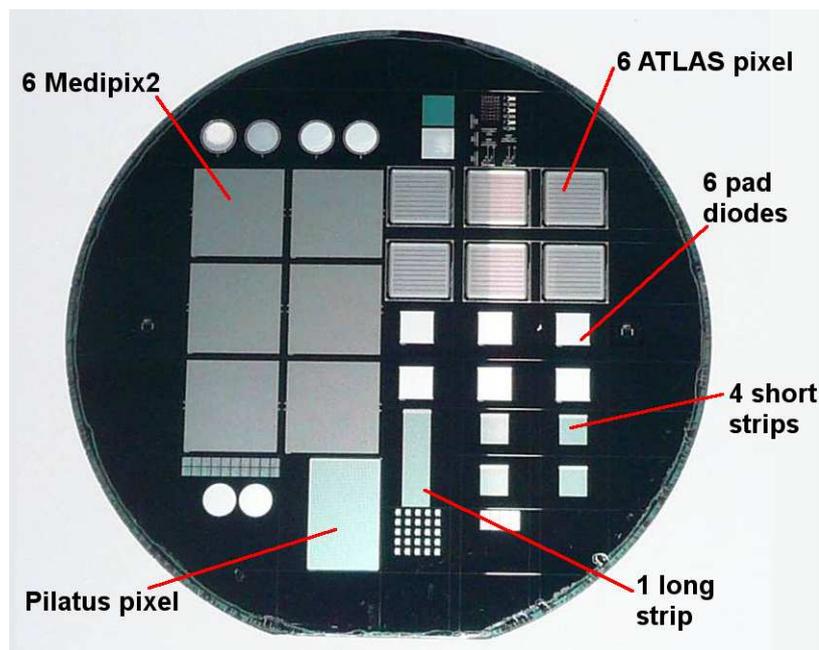


Figure 4.2: A 4" n-type double-sided 3D detector wafer produced by CNM.

with strip detectors, pad diodes and various test structures. Figure 4.2 shows an annotated photo of one of the wafers.

One of the wafers was diced without undergoing any special treatment. The pad detectors from this wafer were used for basic electrical characterisation, and the short strips were used to make charge collection tests with beta particles, as described later in this chapter. The other wafer had its pads coated with under-bump metallisation, so that the devices could be bump-bonded. The Medipix2 devices, which are designed for X-ray detection, were bonded and tested as described in the next chapter. Because this first run used p-type column readout, the ATLAS devices could not be used, since the ATLAS readout chip is only designed to use n-type readout [77]. The Pilatus detector, which is designed for X-ray crystallography [95], was not tested because no suitable test setup was available.

Pad detector structure

The pad detector structure has a 92 by 92 array of p-type readout columns with $55\mu\text{m}$ spacing, to match the pixel size of Medipix2. The 90 by 90 columns in the centre are connected by a metal layer to form a pad, and the surrounding ring of readout columns are connected together to form a guard ring.

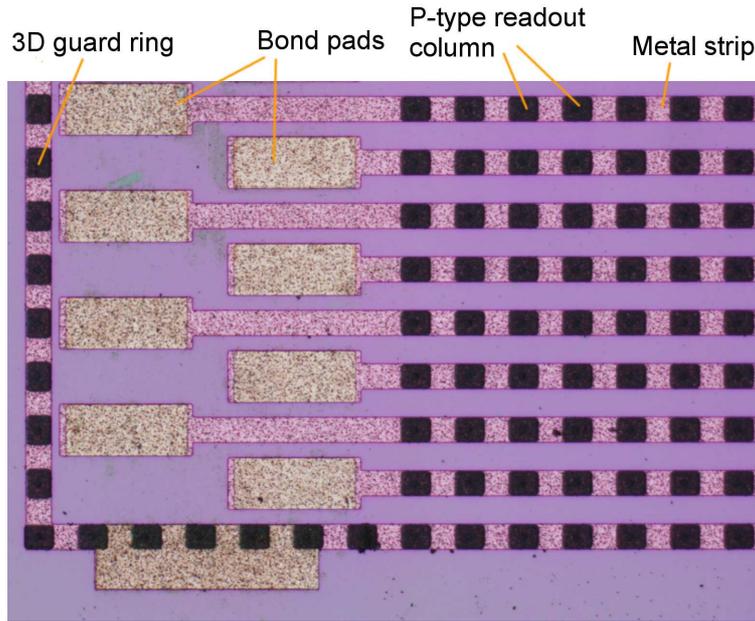


Figure 4.3: Detail of the short strip detector structure produced by CNM.

Strip detector structure

The short strip detectors have 50 strips, with $80\mu\text{m}$ spacing between strips. Each strip consists of a row of 50 p-type readout columns with $80\mu\text{m}$ spacing, connected together by a metal track. So, the detector covers a $4\text{mm} \times 4\text{mm}$ area. At the end of each strip is a larger metal contact pad. The entire structure is surrounded by a guard ring of readout electrodes, as shown in Fig. 4.3.

4.1.3 New fabrication run

A second run of detectors was finished in November 2008. This consisted of 6 n-type wafers with p-type column readout (identical to the previous run) and 8 p-type wafers with n-type column readout. Two further n-type wafers broke during fabrication. On the p-type wafers, the n-type readout columns are isolated by rings of p-stop around each column. In addition to pad, strip and Medipix2 detectors, these p-type wafers will provide ATLAS pixel devices with the required readout polarity.

The completed devices have been I-V and C-V tested at CNM, and in section 4.3 a few results have been quoted. Since large numbers of both types of strip detector are available, Glasgow, CNM and Freiburg are planning to test their charge collection

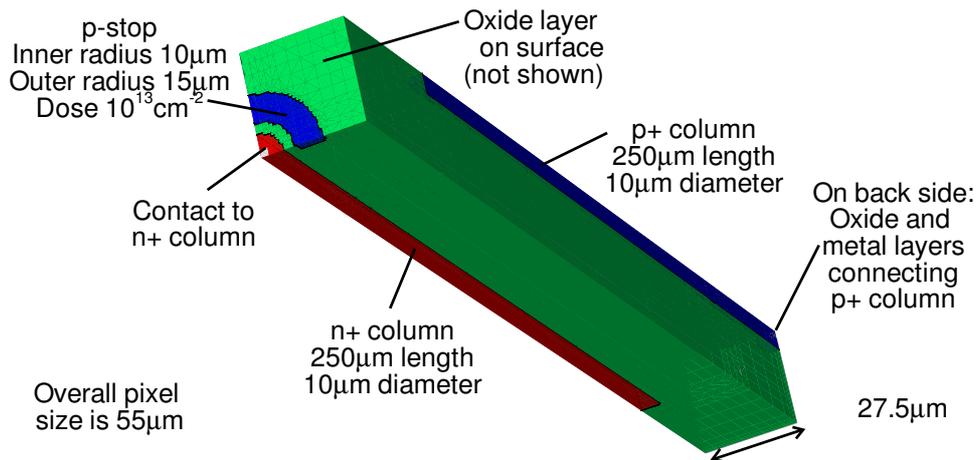


Figure 4.4: Structure of the simulated double-sided 3D detector with a p-type substrate and n-type readout. The dimensions and doping concentrations used in the simulation match the devices fabricated by CNM. Other simulations were also done using a device with an n-type substrate, p-type readout, and no p-stop isolation.

efficiencies at a range of radiation fluences. Also, Medipix2 3D detectors of both types will be bump-bonded to readout chips, and tested at Diamond Light Source.

4.2 Simulation of double-sided 3D detectors

4.2.1 Simulation methods and device structure

Double-sided 3D detectors are simpler to fabricate than full-3D detectors. However, changing the device structure may also degrade the performance. To investigate this, a series of simulations were done using Synopsys TCAD. The simulation process is described in Chapter 2.

Two different types of double-sided 3D structure were simulated. One had an n-type substrate, and used the p-type columns for readout, like the first set of devices produced by CNM. The other type used a p-type substrate and n-type readout columns, like the devices in the second run. This simulated p-type substrate device is shown in Fig. 4.4. The simulations used 10 μm -diameter, 250 μm -long columns in a 300 μm -thick substrate, like the real devices. The substrate doping concentration was $N = 7 \times 10^{11} \text{cm}^{-3}$. The pixel size was 55 μm by 55 μm , to match the Medipix2 readout chip. Note that in most of the simulations, only a quarter-pixel was simulated, much like in the previous 3D simulations. Also, filled polysilicon electrodes were used, rather than the partially-filled columns actually produced by CNM, in order to simplify the simulation.

In the p-substrate, n-readout device, there is a ring of p-stop around each n-type column for isolation. The rings cover a radius of $10\mu\text{m}$ to $15\mu\text{m}$, and have a boron dose of 10^{13}cm^{-2} . Previously, CNM have shown that this boron dose provides reliable electrode isolation in strip detectors, and that higher doses result in a reduced breakdown voltage [73]. The majority of the simulations shown below use this structure, in order to determine the effects of using the p-stop. However, any notable differences between the results from the two structures are described.

In particle physics experiments it is important to limit the material budget (i.e. the quantity of material the particles have to pass through) in order to reduce particle scattering. So, a double-sided device with $250\mu\text{m}$ columns and a $300\mu\text{m}$ substrate should ideally be compared to a $300\mu\text{m}$ -thick full-3D detector. However, currently there are practical limits on how deep a set of columns with a given diameter can be etched, so in this respect a $250\mu\text{m}$ full-3D device would be a better comparison. In the simulations below, both of these comparisons are made.

4.2.2 Depletion, electric field and weighting field

A simulation was carried out where the bias applied to the device was increased from zero in a series of small steps, and the resulting electric field, current density etc. were plotted at regular intervals.

Figure 4.5 shows the hole concentration in the n-readout, p-substrate device at 0V, 1V and 10V bias. When the substrate is depleted, its hole concentration drops dramatically.

At 0V, the depletion region is present at the junction between the n-type column and the substrate. As the bias increases, the depletion region grows cylindrically outwards from the n+ column, much like in a standard 3D detector. Due to the small electrode spacing, most of the device volume depletes very quickly. At 1V, the depletion region has almost reached the p-type bias column, and by 2V most of the device volume is depleted. However, because the n-type columns only extend to a depth of $250\mu\text{m}$, the region at the base of the device depletes more slowly. As can be seen in the “1V” image, the depletion region grows downwards from the tips of the n-type columns to the back surface. The full device volume becomes depleted at 8V. So, although most of the device volume depletes as quickly as a standard 3D detector, full depletion requires a higher bias. Of course, the full depletion voltage is still much lower than that of a planar detector fabricated on the same substrate, which is around 50V. Note that the very front and back surfaces of the detector

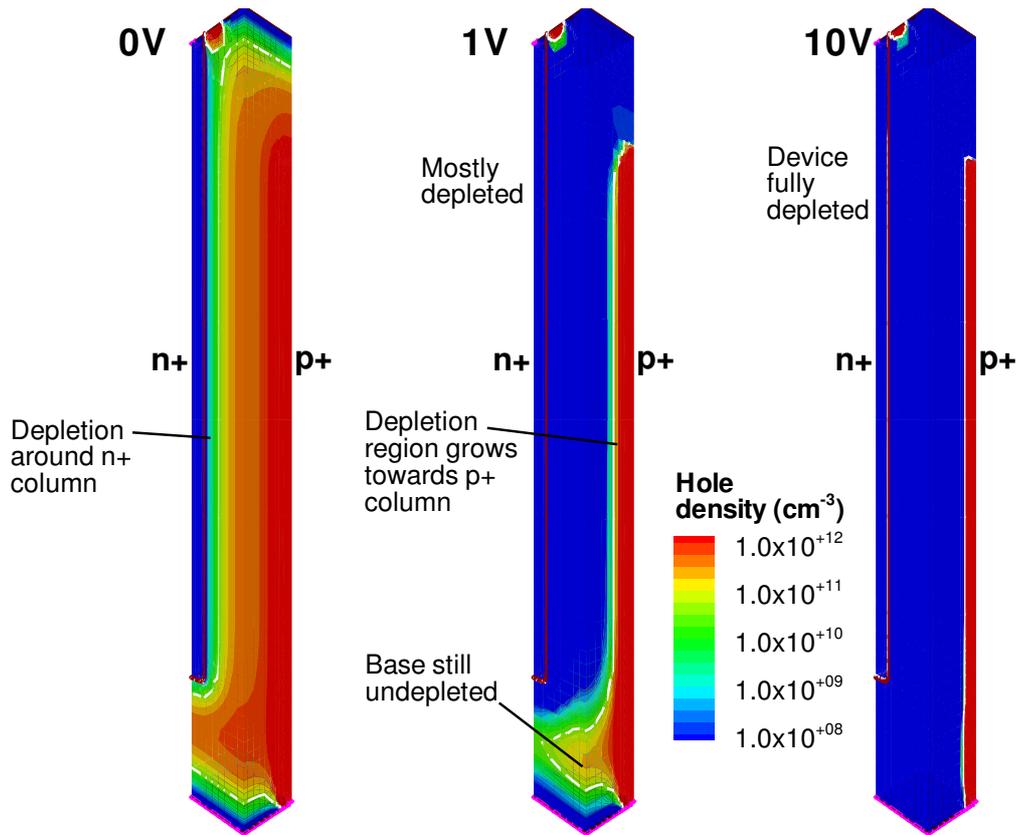


Figure 4.5: Depletion behaviour of a double-sided 3D detector. The hole concentration in the device is shown at 0V, 1V and 10V bias. The substrate has 7×10^{11} cm⁻³ p-type doping, and as the substrate depletes the hole concentration drops dramatically.

appear to be depleted even at zero bias. This is because the oxide layer attracts a layer of electrons, which reduces the hole concentration due to recombination.

The depletion behaviour is basically the same in the n-substrate, p-column-readout device, except of course that the depletion region grows from the p-type columns towards the n-type columns and the back surface.

The electric field behaviour of the double-sided device was simulated over a range of applied voltages. Figure 4.6 shows the electric field strength in a vertical cross-section through the detector, passing through the adjacent n-type and p-type columns. The field around the front surface is shown in detail. In this simulation, the device is biased to 100V.

In the region where the columns overlap, extending from a depth of $50\mu\text{m}$ to $250\mu\text{m}$, the electric field is very similar to that in a full-3D detector. In fact, in a horizontal cross-section of the device at a depth of $150\mu\text{m}$, the electric field is identical to the full-3D detector results shown earlier in Fig. 2.3. However, near the

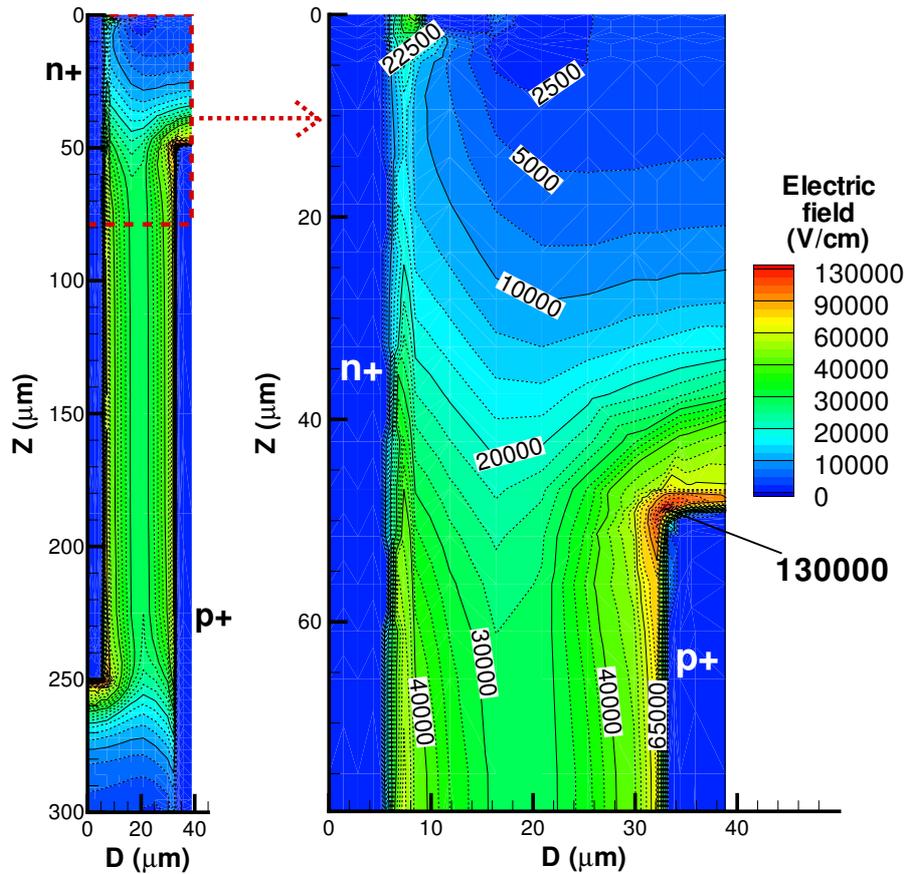


Figure 4.6: Detail of the electric field (V/cm) in a vertical cross-section of the p-type substrate, n-type readout double-sided detector, passing through adjacent n+ and p+ columns. The detector is at a bias of 100V.

front and back surfaces of the detector, the behaviour changes. Moving towards the surfaces, the field gets progressively weaker, with the field strength dropping from 30000 V/cm and higher at $60\mu\text{m}$ depth to less than 5000 V/cm at $10\mu\text{m}$ depth. A similar pattern is seen at the back surface of the detector. Carriers generated in these regions will also have a greater drift distance to the electrodes. So, we can expect slower charge collection from these regions.

The pattern of high-field regions also differs from a full 3D detector. In a full-3D detector both sets of columns make contact with the front and back surfaces of the substrate. As described in section 2.5.1, this creates high-field regions around the surfaces, since the electron or p-spray layers will connect the columns together. In contrast, the surface fields in the double-sided 3D device are relatively low, even around the p-stop. However, high-field regions appear around the tips of the columns, which can lead to breakdown as discussed later.

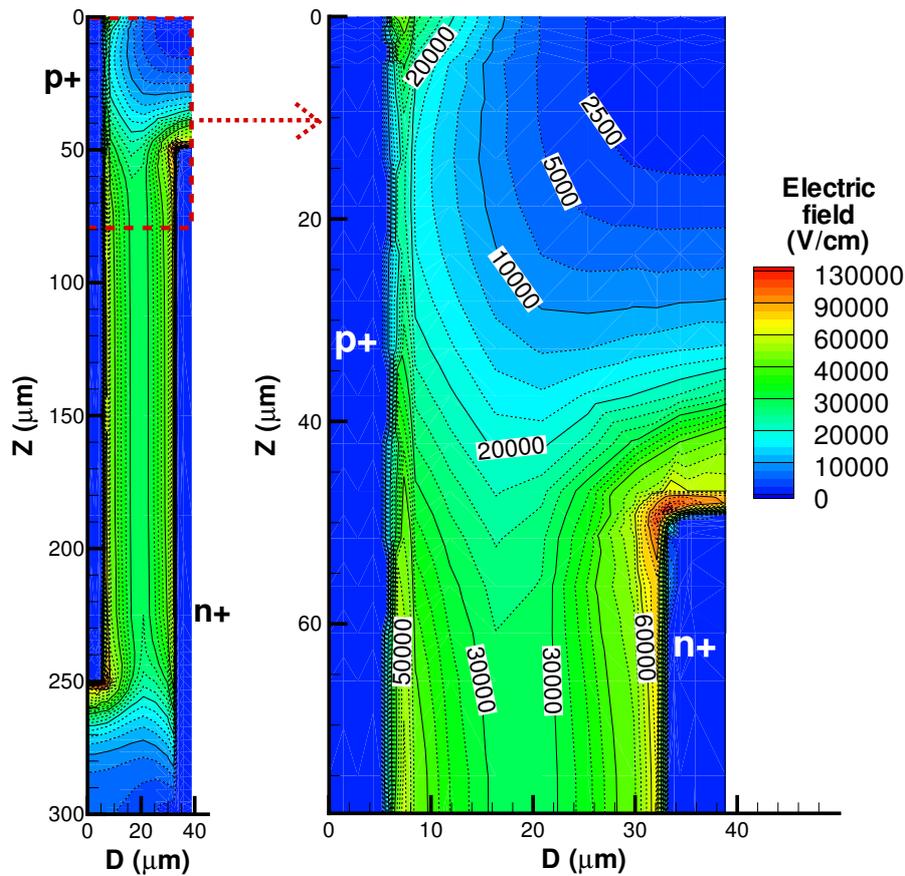


Figure 4.7: Detail of the electric field (V/cm) in a vertical cross-section of the n-type substrate, p-type readout double-sided detector, passing through adjacent n+ and p+ columns. The detector is at a bias of 100V.

The device with an n-type substrate and p-type readout gives a very similar field pattern, as can be seen in Fig. 4.7. The only noticeable difference is that there is no distortion of the field around the front surface due to the p-stop, but there is a slightly higher-field region where the p-type column meets the electron layer at the front surface.

Using the method described in section 2.2, the double-sided 3D detector's weighting field was calculated. Since the weighting field of one readout electrode is affected by the presence of the other readout electrodes, the simulated device structure used four readout electrodes plus the one bias electrode between them. Once again, in the region where the columns overlap, the detector's weighting field matches that of a full-3D detector, which was shown in Fig. 2.7. However, the weighting field does vary around the front and back surfaces. Figure 4.8 shows the weighting potential of an n-type readout electrode, in a cross-section passing diagonally through the readout

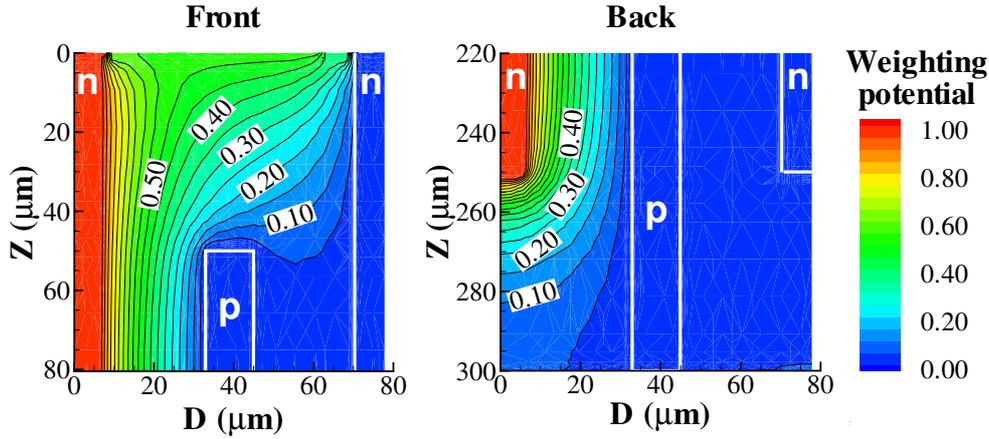


Figure 4.8: Weighting potential around the front and back surfaces of the n-type substrate, p-type readout double-sided detector. The plots show a vertical cross-section passing through the readout column under consideration, the neighbouring p-type bias column, and the next readout column.

electrode, the adjacent bias column, and the next n-type electrode. The presence of the p-type electrode means that in the region where the columns overlap, the readout electrode's weighting field is largely confined to the surrounding pixel, and is relatively uniform. In the region around the back surface, the weighting potential is lower, meaning that electrons drifting to the n-type column will make a greater contribution to the total charge collection than holes drifting to the p-type. At the front surface, the weighting potential is higher, so the holes will make a greater contribution. Overall, the differences in the weighting field are minor and should have little effect on device performance.

4.2.3 Charge collection

A series of transient simulations were used to investigate the charge collection behaviour in these devices. Like in section 2.2, the simulations were done by flooding each device with a uniform concentration of electron-hole pairs, chosen to give the same total number of carriers as a minimum ionising particle, then simulating the current flow in the device over time. The simulations were done for both double-sided 3D structures. For comparison, $250\mu\text{m}$ and $300\mu\text{m}$ -thick full-3D detectors with n-type readout and p-type substrates were also simulated. All of these simulations were done at 100V bias. The resulting signal currents are shown in Fig. 4.9.

It can be seen that all of the devices produce a fast current signal, lasting for around 0.5ns. However, the double-sided 3D detectors also show a longer tail-off.

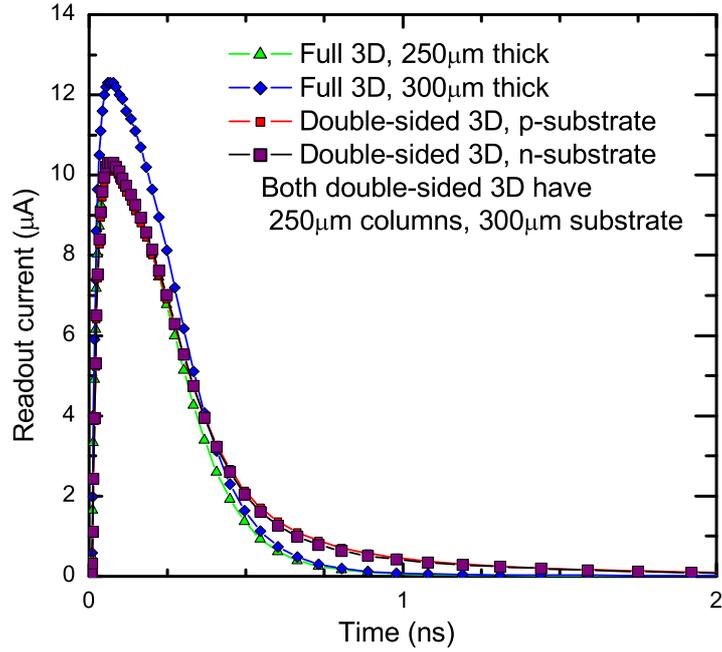


Figure 4.9: Current signals produced by different detector structures. In the simulations, each device is uniformly flooded with excess electron-hole pairs. The simulated structures are the two double-sided 3D structures described in the text (p-type substrate with n-type readout and n-type substrate with p-type readout) and 250 μm and 300 μm -thick full-3D detectors.

This tail-off is still relatively quick compared to the collection time of 5ns or longer for a planar detector. Interestingly, the 250 μm -thick full-3D detector shows an extremely similar pulse shape to the double-sided detectors for the first 0.5ns. In effect, this suggests that the central region of the double-sided 3D detector produces the same collection signal as the full-3D, and the extra 50 μm substrate thickness gives additional signal current with a slower collection time. However, the 300 μm full-3D detector gives superior performance, since it gives the same total signal in a shorter time than the double-sided 3D.

To investigate the effects of the low-field regions around the front and back surfaces more fully, simulations were done to map the variation in collection speed with depth. This was done by depositing short tracks of charge at different depths in the device. Each of these short tracks was midway between the n- and p-type columns, ran parallel to these columns, and had a Gaussian lateral distribution with a sigma of 1 μm . The device with n-type substrate and p-type readout was used, and the bias was 100V as before. A single long track was used from 100 μm to 200 μm ,

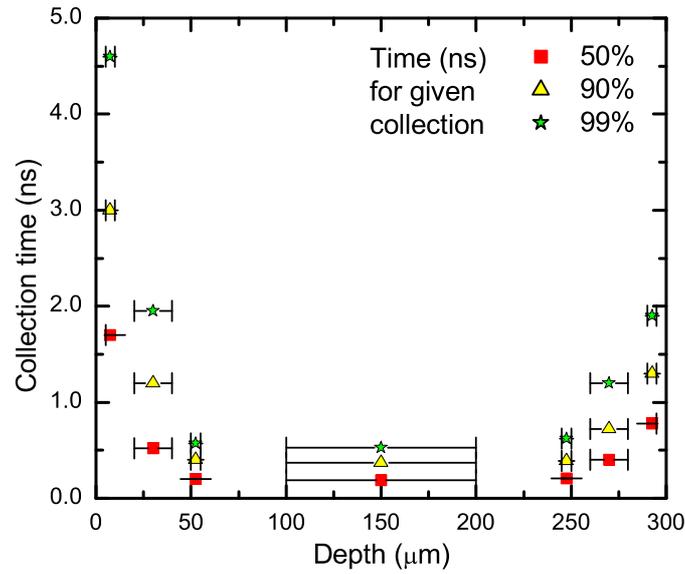


Figure 4.10: Simulated charge collection times for charge deposited at different depths in the double-sided device (100V bias). In each case, the charge was deposited along a short path running parallel to the p+ and n+ columns, lying midway between them. The length of each path is indicated by the “error bars”.

because the electric field did not vary significantly with depth in this region. Also, the lengths of the short charge tracks were adjusted to ensure that they passed through several mesh elements, to improve the accuracy of the charge integration.

The resulting variation in the collection time with depth is shown in Fig. 4.10. The charge deposited in the region where the columns overlap ($50\mu\text{m}$ to $250\mu\text{m}$ depth) is collected in about 0.5ns, matching the main signal pulse seen in Fig. 4.9. The charge collection time increases towards the front and back surfaces, as would be expected due to the weakening field and the increasing collection distance. The collection time is longer at the very front surface than at the very back. The weighting field simulation shows that collection signal from the back surface of the n-type readout device is dominated by the electron drift, whereas around the front surface hole drift will make a bigger contribution. Since electrons have higher mobilities, this could explain the difference. The distortion of the field around the front surface due to the presence of the p-stop may also make the collection slower.

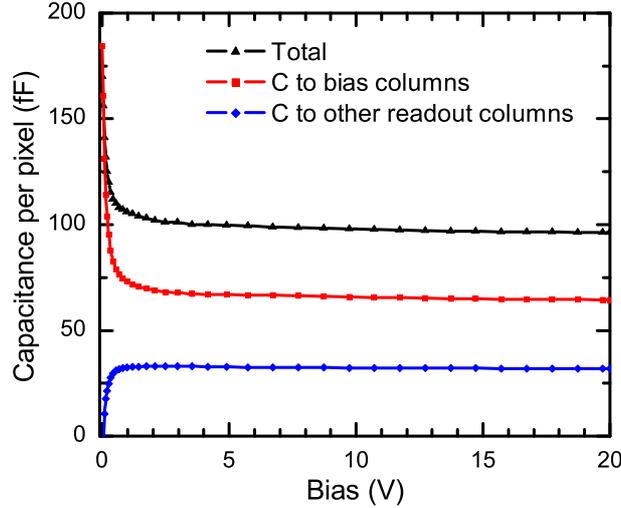


Figure 4.11: Simulated C-V characteristics of a double-sided 3D detector with n-type readout and a p-type substrate. The curves show the total capacitance seen at a readout column, the capacitance to the bias columns, and the capacitance to the neighbouring readout columns.

4.2.4 C-V characteristics

To simulate the C-V characteristics of the device, a mesh was created with four n+ quarter-electrodes, one full p+ electrode between them, and a p-type substrate. This larger mesh was used so that the capacitance between the readout electrode and the neighbouring pixels would be taken into account. This simulation was repeated with the p-readout, n-substrate double-sided 3D detector, and with two equivalent 250 μ m-thick full-3D detectors.

The C-V curves in Fig. 4.11 show the results from the p-type substrate, n-type readout double-sided detector. The capacitance between the readout column and the bias columns, which dominates the total capacitance, drops rapidly as the depletion region grows. Most of this decrease happens over the first 2V, as the depletion region around the readout column grows to meet the bias columns. There is a further small drop in the capacitance as the region at the back surface depletes. The capacitance to the neighbouring readout columns starts at zero, due to the undepleted substrate between them, and rises as the device becomes depleted.

The saturation capacitances of the four detector structures are shown in Table 4.1. In each case, the total capacitances are extremely similar, at around 95fF per column. The main difference is that in the double-sided 3D devices the capacitance to the bias columns is smaller, and the capacitance to the neighbouring

Table 4.1: Simulated capacitances of double-sided 3D and full-3D detectors. The total capacitance per column is given, along with its components, i.e. the capacitance to the bias columns and the capacitance to the neighbouring readout columns.

Structure	Substrate	Total C (fF/column)	Bias columns (fF/column)	Neighbouring readout (fF/column)
Double	p-type	96.3	64.5	31.8
Double	n-type	94.8	63.6	31.2
Full	p-type	95.5	70.0	25.5
Full	n-type	94.5	76.6	17.9

readout columns is larger. This will be because the readout and bias columns are offset from each other in the double-sided structure. While the results from the two different double-sided detectors are virtually the same, the difference between the two full-3D detectors is larger, perhaps because the differences in their surface isolation have a larger effect.

Additionally, a capacitance simulation was done for a double-sided 3D strip detector with an n-type substrate, p-type readout and $80\mu\text{m}$ column spacing, to match the strip detectors tested later. This had a lower total capacitance of 68fF per column, mainly due to its wider pitch. This corresponds to a strip capacitance of 8.5pF/cm, or 3.4pF for the 4mm strips tested later. In contrast, the capacitance of an equivalent strip detector would typically be 1.0–1.5pf/cm, depending on the width of the strip implant [96]. So, the double-sided 3D detector’s large capacitance is a disadvantage.

4.2.5 Breakdown behaviour

Breakdown at the column tips

As shown by the electric field simulations, the highest-field regions in a double-sided 3D detector appear at the tips of each column. To investigate the device’s avalanche breakdown voltage, an impact ionisation model was added to the simulations. For a device with $250\mu\text{m}$ columns, avalanche breakdown occurs at both sets of column tips at 230V. Fig. 4.12 shows the electric field around the p+ tip at 215V, just before breakdown. The high-field region appears around the edge of the cylindrical column, where the curvature of the column causes a large electric flux to pass through a small region.

The overall “shape” of the doped column, and hence the field distribution, is affected by two factors: the geometry of the etched hole and the diffusion profile of the dopant added to it. This simulation used a realistic diffusion profile, which tends

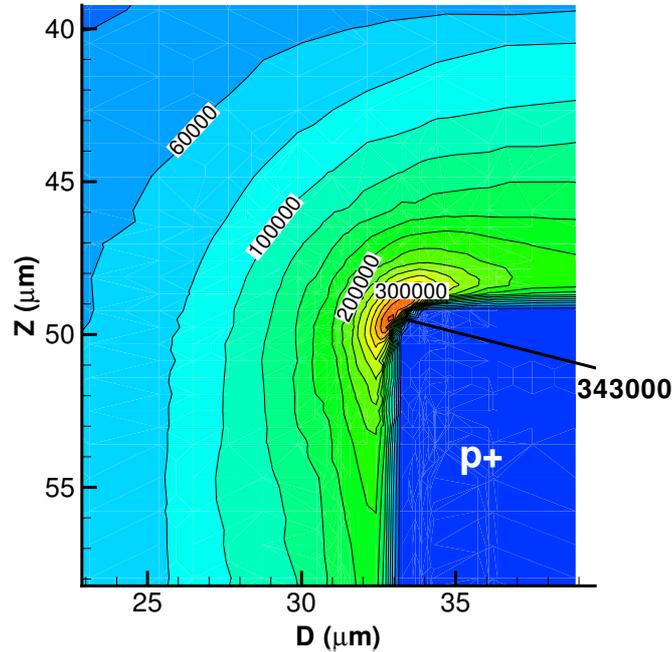


Figure 4.12: Electric field around the tip of the p+ column in a double-sided 3D detector at 215V, just before breakdown.

to smooth out the curved edge of the column, and a cylindrical hole. In practice, the base of each etched hole is slightly rounded (see Fig. 1.10), which may reduce the field slightly. Unfortunately, MESH is not flexible enough to produce a suitable simulation structure. As an alternative test of the effect of the column geometry on the breakdown, the simulation was repeated using columns with square cross-sections. The square columns break down at 180V, with the greatest field appearing at the corners. So, even with a distinctly suboptimal design, the breakdown voltage is of order 20 times greater than the unirradiated depletion voltage.

Effects of increasing oxide charge

So far, the breakdown simulations have been performed with lower levels of oxide charge, corresponding to an unirradiated detector. The simulations were repeated with a typical saturated oxide charge of 10^{12}cm^{-2} [10].

In the p-substrate, n-readout device with p-stops, high-field regions develop at the front surface between the p-stop ring and the n+ column, and at the back surface where the accumulated electron layer meets the base of the p+ column. However, the

electric field is greatest around the tips of the electrode columns, and the eventual breakdown at 210V is primarily due to avalanche breakdown at the column tips.

Similarly, in the n-substrate, p-type readout devices, there is a high-field region where the electron layer at the front surface meets the p-type column. Once again, the field is greatest around the column tips. The breakdown voltage is 210V, which is only a slight decrease compared to the case with low oxide charge.

Comparison with full-3D detectors

It was found that the impact ionization models used here couldn't be reliably combined with the radiation-damage model used in Chapter 2. So, in section 2.5.1, an estimate of a safe operating voltage was determined by finding the bias at which the maximum field in a full-3D detector reached $2.5 \times 10^5 \text{V/cm}$. When this test is applied to this double-sided 3D detector, the corresponding voltage is 165V. This is close to the value obtained for most of the full-3D detectors, which used n-type readout and p-spray isolation. However, in the full-3D detectors the high-field region occurs at the front surface, where the p-spray meets the n-type column. Also, this "high-field" voltage is substantially lowered if the p-spray dose is not well matched to the charge concentration in the oxide, for example if the p-spray dose is increased or the oxide charge is below its saturation value.

Another breakdown simulation was run for a full-3D detector with p-stop isolation, using rings of p-stop matching those in the double-sided 3D detector. With 10^{12}cm^{-2} oxide charge, the device breaks down at 85V, with the high-field regions appearing both at the edge of the n-type column and the inner edge of the p-stop. This is unsurprising, since a p-stop generally gives lower breakdown voltages than a well-chosen p-spray when the oxide charge is saturated [11].

When the breakdown behaviour of a p-readout, n-substrate full-3D device was simulated, without electrode isolation, the breakdown voltage decreased as the oxide charge was increased, falling to just 55V with 10^{12}cm^{-2} . While this may be sufficient for applications such as X-ray detection, operating the detector following radiation damage could present problems.

Overall, these results demonstrate that the double-sided 3D structure's breakdown point is largely determined by the basic device geometry, whereas the full-3D detector is sensitive to the specific design of the surface isolation.

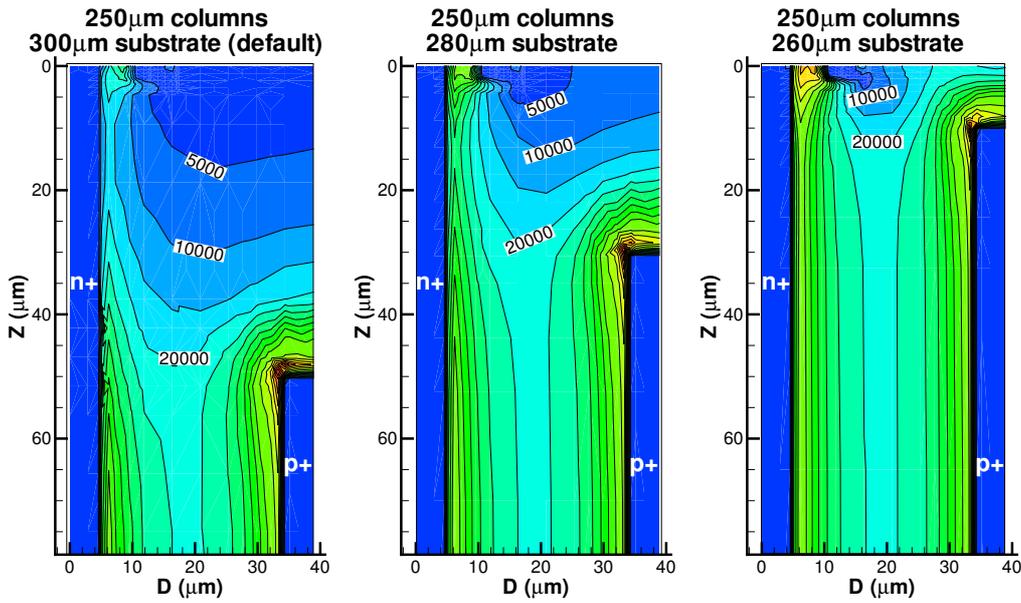


Figure 4.13: Detail of the electric field distribution in double-sided 3D detectors with $250\mu\text{m}$, $270\mu\text{m}$ and $290\mu\text{m}$ columns at 100V . Labelled contours indicate the field strength in V/cm , showing the extent of the low-field region.

4.2.6 Effects of altering the substrate thickness

By reducing its substrate thickness for a given column length, a double-sided 3D detector can be made more similar to a full-3D device. To investigate this, further simulations were done where the columns remained $250\mu\text{m}$ long but the p-type substrate's thickness was reduced to $280\mu\text{m}$ or $260\mu\text{m}$. Fig. 4.13 shows the electric field distribution near the front surfaces of these structures, alongside the results with the $300\mu\text{m}$ substrate. (Fewer contour levels are shown, for the sake of clarity.) It can be seen that reducing the substrate thickness reduces the size of the low-field regions. The charge collection behaviour was also simulated. The tail-off signal becomes smaller as the thickness is reduced, though this does reduce the total charge collected. In particular, the pulse shape from the $260\mu\text{m}$ substrate is virtually identical to that from the full-3D detector, apart from the marginally higher total charge.

However, the substrate thickness also affects the breakdown behaviour at the column tips. As the substrate thickness is reduced, and the tip of n-type column approaches the back surface, the field gets stronger. So, the breakdown voltage falls from 230V with a $300\mu\text{m}$ substrate, to 175V with a $260\mu\text{m}$ substrate. However, the field doesn't particularly increase at the p-type column tip, which is near the front

surface. This difference is seen because the voltage is applied to the bias columns via layers of polysilicon and metal on the back surface, which overlie the oxide layer. So, when the n-type column approaches the back surface, the potential difference applied to the detector falls across the relatively small thickness of silicon and oxide between the column and the metal layer, increasing the field. When these simulations were repeated with this conductive material removed (apart from the contacts to the p-type columns) this problem did not occur. So, in order to fabricate a double-sided detector with a column length similar to the substrate thickness, it would be worthwhile to remove most of the polysilicon layer covering the back surface (see Fig. 4.1), and connect the bias columns together with metal tracks.

4.2.7 Behaviour after radiation damage

Following radiation damage, the double-sided detector's effective doping concentration will increase, altering its electrostatic behaviour, and carrier trapping will also affect the charge collection.

The electric field in the p-type substrate, n-type readout device was simulated following $10^{16}n_{eq}/\text{cm}^2$ radiation damage, using the model from section 2.3.3. Figure 4.14 shows the electric field around the front and back surfaces of the device at 100V bias. As before, in the region from $50\mu\text{m}$ to $250\mu\text{m}$ depth where the columns overlap, the electric field behaviour is like that in a full-3D detector. In the case of the radiation-damaged device, this means that the lateral depletion voltage increases with fluence, and the electric field is higher around the n-type column. The region around the front of the device is depleted, but the electric field decreases more rapidly with distance from the n-type column, as well as decreasing toward the front surface. The biggest change in the behaviour is seen around the back surface, which is undepleted beyond a depth of about $275\mu\text{m}$. This is because a higher voltage is required to grow the depletion region from the tip of the n-type column to the back surface. So, any carriers generated near the back surface will be lost.

The charge collection performance was then simulated for this device, and for the equivalent $250\mu\text{m}$ -thick full 3D detector. As before, the entire device volume was uniformly flooded with electron-hole pairs, and the resulting current was integrated to find the total signal. This was done at a range of fluences up to $10^{16}n_{eq}/\text{cm}^2$, with 100V bias. The results for the two devices are shown in Fig. 4.15. Since the double-sided detector has a greater substrate thickness than a full-3D detector with the same column length, its collection signal is higher before radiation damage. As the damage level increases, the collection efficiency falls more quickly in the double-

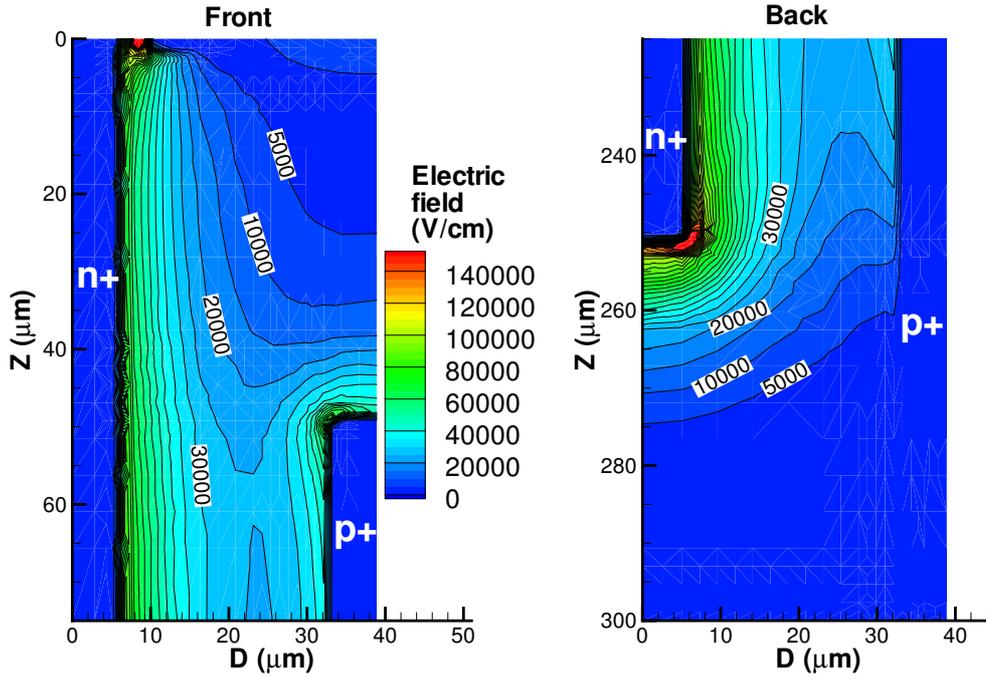


Figure 4.14: Simulated electric field in a double-sided 3D detector following $10^{16}n_{eq}/\text{cm}^2$ radiation damage. A vertical cross-section was taken, passing through adjacent n-type and p-type columns; here, the field strengths around the front and back surfaces are shown. The device is biased to 100V.

sided detector than in full-3D. This will be due to incomplete depletion, and also greater trapping of carriers from the front and back surfaces due to the lower field. However, even at $10^{16}n_{eq}/\text{cm}^2$, the double-sided detector still achieves the same collection signal as the full-3D detector.

4.3 Experimental results

4.3.1 I-V tests

During the I-V tests on each strip detector, probe needles were placed on three adjacent strips and the guard ring, and held at ground, while the back contact was biased. So, the central strip out of the three should experience relatively normal bias conditions. Figure. 4.16 shows the I-V results from one of the strip detectors, measured at 21°C.

The four strip detectors showed large variations in the guard ring current. Two of the detectors had acceptable guard ring currents of 30nA and $1\mu\text{A}$ at 50V. (The figure shows the I-V characteristics of the latter.) In these detectors, each strip that

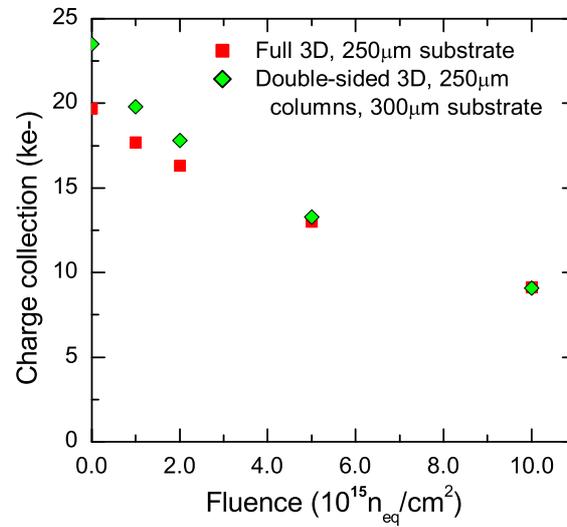


Figure 4.15: Simulated charge collection signals in double-sided and full-3D detectors with equal column lengths following radiation damage.

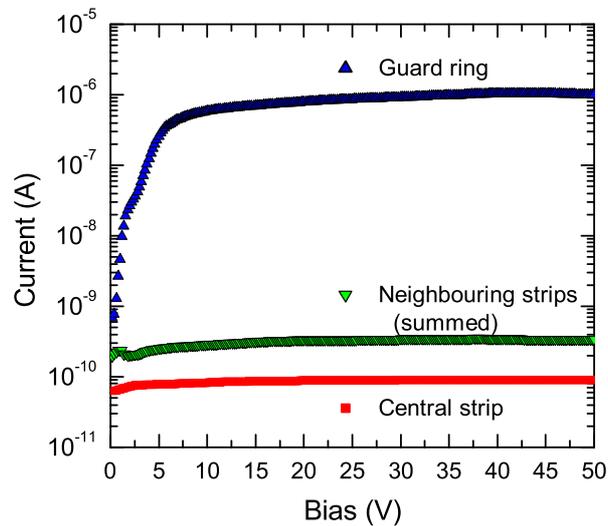


Figure 4.16: Current-voltage characteristics of a 3D strip detector with $80\mu\text{m}$ pitch and 4mm strip length, tested at 21°C . Three strips and the guard ring were biased during the test.

was tested had a current of around 100pA at 21°C, and could be reliably biased to 50V without breakdown. The other two strip detectors had poor behaviour. In one detector, breakdown occurred at 10V, with the strip current rapidly rising from less than 1nA to over 1 μ A. This detector also had a high guard ring current of 10 μ A at 10V. The other detector had an extremely high guard ring current, which reached the probe station's 20 μ A current limit at just 6V. The strip itself only showed about 100pA current during the test, though given the extremely low bias this is unremarkable.

All four of these detectors were fabricated on the same wafer, so their differences must be due to some inherent variability in the fabrication process. On the whole, it appears that the surface leakage current measured in the guard rings is much more variable than the bulk current in the strips. This could possibly be due to variations in the saw-cut edges of the detector; for example the detectors with poor performance may have more cracks.

I-V tests following irradiation

Following irradiation to 5×10^{15} 1MeV-n_{eq}/cm² with reactor neutrons, one of the strip detectors was retested. The I-V tests were done at 21°C as before, due to the lack of cooling in the probe station, and the high leakage current prevented the detector from being biased to full depletion. At the highest bias of 18V, the leakage current per strip was found to be 8.4 μ A. From simulations, lateral depletion in this device is expected to occur at around 50V, and if we linearly extrapolate the current flow to this value we can estimate the current flow at 20°C at full depletion to be $\sim 25\mu$ A per strip.

As discussed in section 2.3.2, the leakage current per unit volume in an irradiated detector increases linearly with fluence, and this can be parametrised by α . Using this estimate of the current flow at full depletion, $\alpha \approx 5 \times 10^{-17}$ A/cm, which falls within the range of values reported in Ref. [60].

New detectors with p-type substrates and n-type readout

The first set of 3D strip detectors used n-type substrates and p-type readout columns. A newer batch of strip detectors has recently been completed by CNM, with p-type substrates, n-type readout columns and p-stop electrode isolation. CNM staff have performed I-V tests on these detectors. Out of 16 new detectors tested, 14 could be biased to 200V without breakdown. The remaining two broke down very early, below 10V. Nearly all of the 14 detectors had guard currents at 50V ranging from

Table 4.2: Measured capacitances of double-sided 3D pad detectors.

Pad	Total C (nF)	C per unit area (nF/cm ²)	C per readout column (fF)
1	0.294	1.20	36.3
2	0.278	1.13	34.3
3	0.231	0.94	28.5
Mean	0.268	1.09	33.0

45–135nA, which is a big improvement over the first set of devices. (There was one new device with a somewhat higher guard ring current of $4\mu\text{A}$ at 50V, which is still acceptable.) As well as using rings of p-stop to isolate the readout columns, the new p-substrate detectors also use p-stops to reduce the surface current. In particular, there are p-stop implants along the dicing lanes on the wafer, where the individual detectors are diced apart. So, this explains these devices' superior behaviour.

4.3.2 C-V tests

Three pad detectors were C-V tested in the probe station. During each test, the pad was held at ground, the back contact was biased, and a 10kHz AC signal was applied between the two. The guard ring was also grounded, but the AC signal was not applied to it.

The capacitance decreases as the device depletes. Fig. 4.17 shows the inverse of the capacitance measured on one of the pad detectors versus bias; this was chosen to make small changes in capacitance more visible. The curve falls into three distinct phases. Up to 2.4V, the curve rises steeply, as the capacitance falls rapidly. This will correspond to the depletion regions around each column growing laterally to reach the bias columns. After 2.4V, there is a kink in the curve, and the capacitance changes at a slower rate. This will correspond to the slower growth of the depletion region to the back surface of the substrate. Then, at about 9V, the curve flattens out, corresponding to the device becoming fully depleted and the capacitance reaching a minimum. So, the general shape of the curve corresponds to the expected depletion behaviour, and the specific voltages at which lateral and full depletion occur are close to the simulated values of 2V and 8V.

The final saturation capacitances were different in the 3 detectors. The saturation capacitances are shown in Table 4.2, along with the capacitance per square centimetre and per column (given that we have a 90 by 90 array with $55\mu\text{m}$ pitch).

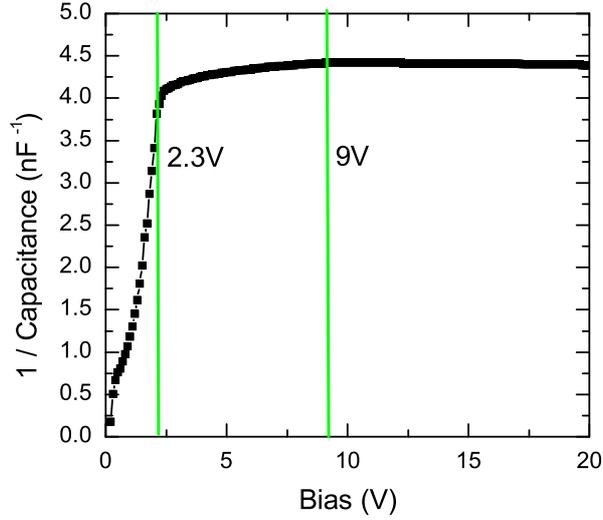


Figure 4.17: Inverse capacitance characteristics of a 3D pad detector with 90 by 90 columns and $55\mu\text{m}$ pitch.

Table 4.3: Capacitances of new double-sided 3D pad detectors, tested at CNM.

Substrate	Column diameter (μm)	Total C (nF)	C per unit area (nF/cm^2)	C per readout column (fF)
n-type	10	0.7	2.9	85
n-type	13	1.1	4.5	135
p-type	10	0.6	2.4	75
p-type	13	0.8	3.3	100

C-V tests on a pad detector will measure the capacitance between the bias columns and the readout columns, but not the capacitance between different readout columns which would be seen in a strip or pixel device. From the simulations, the predicted pad capacitance is 64fF per column, which is roughly double the mean experimental value of 33fF . The variation in capacitance from device to device is also quite large. However, when these detectors were retested at a later date the same results were obtained, so experimental errors are unlikely.

These pad detectors have n-type substrates and p-type readout columns. In the new batch of pad detectors produced by CNM, there are more of these devices, and also devices with p-type substrates and n-type readout. Additionally, both varieties of device were produced with $10\mu\text{m}$ and $13\mu\text{m}$ diameter columns. These detectors were tested by CNM staff, and their results are reported in Table 4.3. The $13\mu\text{m}$ -diameter devices were also retested at Glasgow and the same results were obtained, demonstrating that the test setups behave consistently.

The new $10\mu\text{m}$ -diameter n-type substrate pad devices have a capacitance per column of 85fF , and the p-type substrate have 75fF per column. These values agree better with the simulated value of 64fF , and are more than double the mean value measured on the old n-type detectors. Although these new results are more reasonable, the difference between the old and new n-type devices is surprising, given that they should have the same structure. In particular, it is difficult to explain how the old devices could have substantially lower capacitance than simulated, since most unwanted effects (such as stray capacitances) tend to increase the overall capacitance rather than reduce it. In principle, if the columns were much shorter than expected, this could lead to lower capacitance. However, this would have major effects on other aspects of the device performance. For example, depleting to the backplane of the detector would require a higher voltage than is seen in the C-V curve here.

Looking at Table 4.3, the p-type substrate devices have lower capacitance than the n-type. The p-substrate devices have p-stops around each n-type readout column, and around the detector array, whereas the n-type have no isolation. So, there may be less capacitive coupling across the surfaces in the p-type devices. The devices with $13\mu\text{m}$ column radius have greater capacitance, which is unsurprising, given that the wider columns have a larger surface area and smaller spacing between them.

Strip detectors

A strip detector was also C-V tested. However, it was only possible to bias one strip plus the backplane while making the C-V measurements. Ideally, multiple strips would have been biased, and the AC signal could then have been applied between a strip and the bias columns, or between adjacent strips to find the interstrip capacitance. (If the detector had been AC coupled, all the strips could have been biased simultaneously.)

The resulting C-V curve is rather strange, as can be seen in Fig. 4.18. Initially, the capacitance drops rapidly as expected, reaching a minimum of 3.1pF (about 62pF per readout column) at 2.8V . However, beyond this the measured capacitance increases, briefly passes through a plateau from $7\text{-}9\text{V}$, then rises still further before levelling out at about 10.5pF , i.e. 210fF per readout column.

Since this device has an $80\mu\text{m}$ pitch, one would expect it to have a lower capacitance per column than a $55\mu\text{m}$ pitch device. From the simulation of the strip detector structure, the capacitance between a strip and the backplane is expected to be 2.6pF per strip, corresponding to 52fF per readout column. (The capacitance to the neighbouring strips is 0.8pF .) This is close to the minimum value of the C-V

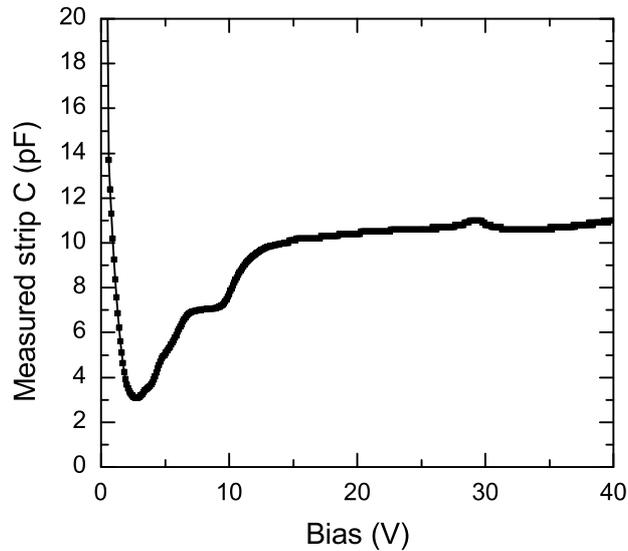


Figure 4.18: Capacitance-voltage characteristics of a single biased strip on the 3D strip detector. Note that several adjacent strips should have ideally been biased simultaneously; the poor bias conditions may explain the odd results.

curve at 2.8V. However, the 210fF capacitance per column after the curve levels out is much higher than this value. This is also much higher than the capacitances per column obtained from the pad detectors.

Normally, when a device is biased, the depletion region will grow outwards from all the readout columns simultaneously. In this test setup only one strip is biased, and the depletion region produced by the columns in this strip will grow outwards without limit, eventually meeting the n-type columns in the neighbouring strips. These neighbouring strips will then begin to deplete, too. It's perhaps possible that these neighbouring strips contribute to the measured capacitance, which would explain why the capacitance rises as more of the device becomes depleted. As well as measuring the capacitance, the test measured the parallel conductance between the readout strip and the backplane. It was found that between the minimum at 2.8V and the first plateau at 9V, the conductance increased by almost an order of magnitude from 1.2×10^{-8} S to 1.0×10^{-7} S. So, this also suggests that the growth of the depletion region led to increased coupling between the strip and the bias electrodes.

4.3.3 Unirradiated strip detector charge collection efficiency

The double-sided 3D strip detectors were tested using the beta test setup built from LHCb electronics described in section 3.4. For these tests, two detectors were left unirradiated, and two were irradiated to 5×10^{15} 1MeV- n_{eq}/cm^2 .

One of the unirradiated detectors was built into a module. No AC coupling was used; the DC-coupled strips were connected to the Beetle chip via a pitch adaptor. Although the Beetle is designed to use AC coupling, the DC-coupled configuration was expected to work, due to the relatively low strip current. The main problem with operating a detector in DC-coupled mode is that the leakage current will tend to charge up the integrating capacitor in the preamplifier. If the capacitor doesn't discharge quickly enough, this accumulation of charge can alter the amplifier's operating conditions, or saturate it completely. According to a calculation made by the Beetle's designer, a strip with about 600nA leakage current would completely saturate its preamplifier. Since the strips in the detector have a current of order 0.1nA, in principle this configuration should work OK.

The charge collection behaviour was tested with betas at a variety of biases. Figures 4.19 and 4.20 show the spectra measured at 20V and 6V respectively. Each spectrum is fitted with a Landau convolved with a Gaussian, to find the most probable signal value. Compared to the planar strip detector results in section 3.4.4, the signal sizes here are generally lower, and there appear to be excess counts at low energies. The lower-energy hits could be due to betas which partially pass through the electrode columns, since charge loss is expected from these regions.

Figure 4.21 then shows the charge collection versus bias. This was found by assuming that the planar detector results from section 3.4.4 provide 100% CCE, and then using this to calibrate the setup. It can be seen that the charge collection is relatively low compared to the expected 24000 electrons in $300\mu\text{m}$ of silicon, even well above the depletion voltage.

The low CCE is unusual, given that this problem is not seen in the 3D Medipix2 devices tested in Chapter 5. As mentioned previously, charge loss in the columns can potentially explain why a small proportion of hits give low CCE. However, the charge collection is determined by the most probable value of the Landau fit. Since we get a reasonably well-shaped Landau with a low most probable value, this implies that every hit on the detector has low collection efficiency.

One possible explanation could be ballistic deficit—the Beetle chip has a relatively short peaking time of around 30ns, so if the charge collection is not fast enough, charge will be lost. However, according to simulations this should be un-

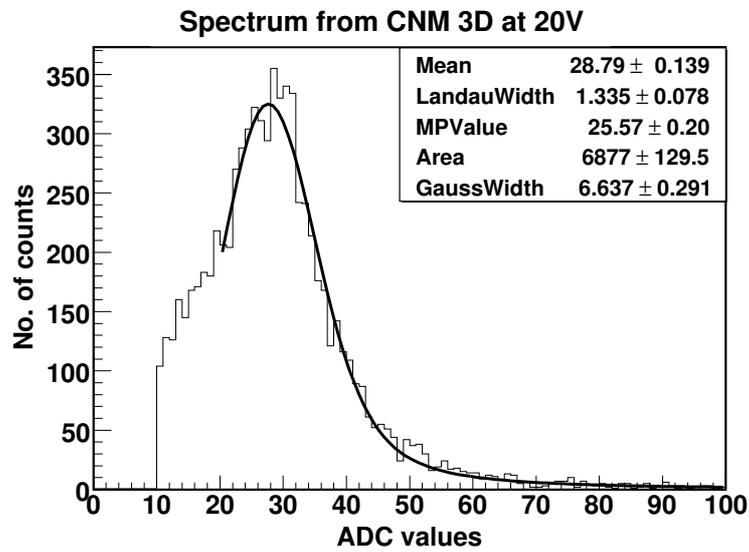


Figure 4.19: Signal spectrum from beta tests of the unirradiated double-sided 3D strip detector at 20V, fitted with a Landau convolved with a Gaussian.

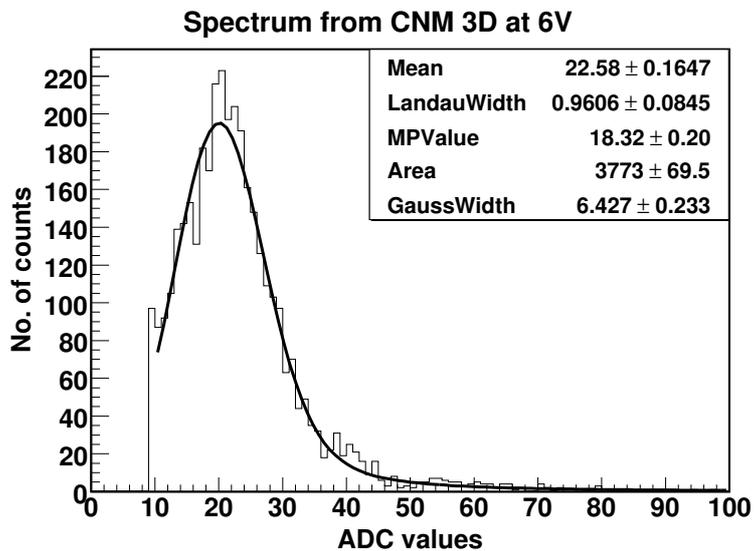


Figure 4.20: Signal spectrum from beta tests of the unirradiated double-sided 3D strip detector at 6V, fitted with a Landau convolved with a Gaussian.

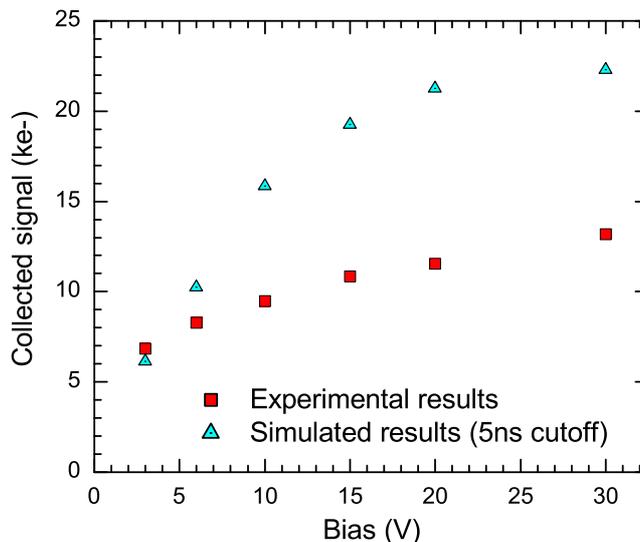


Figure 4.21: Experimental and simulated charge collection efficiency of the double-sided 3D strip detector. The simulation uses a short integration time of 5ns, to exaggerate the effects of ballistic deficit. The experimental CCE is much lower than expected.

likely. The collection behaviour of the detector was simulated at different biases, and the readout current was integrated over just 5ns following the hit, to emulate a much shorter peaking time than the Beetle chip. The predicted charge collection at 20V is still much higher than the experimental value, as shown in Fig. 4.21. To test this further, the average Beetle front-end pulse shapes were measured at 20V and 3V, as shown in Fig. 4.22. The pulse is not only smaller at 3V, but it peaks later and shows less of an undershoot, indicating slow charge collection at this voltage. However, the pulse shape at 20V is basically the same as the planar detector's pulse shape in Fig. 3.10, apart from its lower amplitude. In particular, the peak occurs in the same time bin for both detectors. So this suggests that the 20V results are not significantly affected by ballistic deficit.

Another possibility is that the DC coupling is causing unexpected problems with the readout chip. Each strip on the Beetle chip can generate test pulses and send them to the preamplifier, so this was used to test the front-end. As noted previously, the Beetle chip has 128 readout channels, but the strip detector only has 50, so many channels on the Beetle were not bonded to the detector. When test pulses were applied to each strip, the readout channels bonded to the detector gave the same response as the unbonded channels. So, the test pulses don't demonstrate

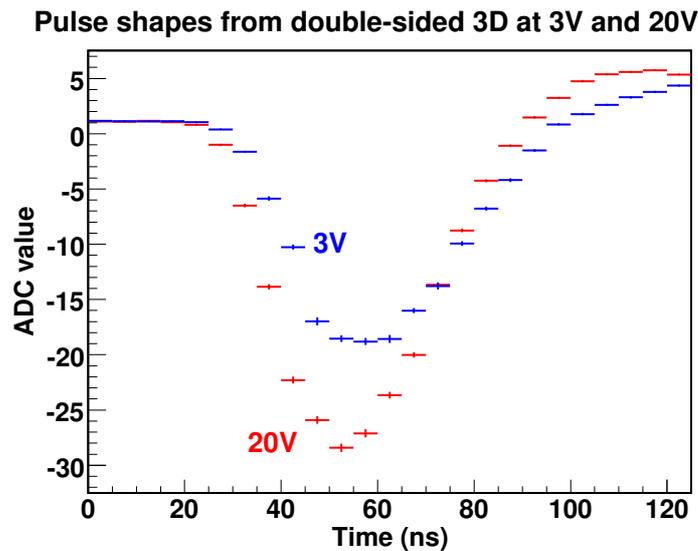


Figure 4.22: Mean pulse shapes measured on the double-sided 3D strip detector at 20V and 3V. As well as having a smaller magnitude, the pulse at 3V is slower to peak, and shows less undershoot. The offset of the time axis from zero is arbitrary, but is consistent between the two sets of measurements.

any problem with the preamplifiers.

Given that the Medipix2 devices show much better charge collection, the low CCE appears to be specific to the strip detectors. Furthermore, the heavily irradiated strip detectors, which use proper AC coupling, show relatively high charge collection—see below. So, it still appears likely that the DC coupling is causing the problem.

Unirradiated detector - Noise

The noise level on each strip was also measured with the detector biased to 20V. A large number of samples were taken without the source in place, using the TELL1 to generate the triggers internally. A Gaussian fit was applied to the results from each strip, and the sigma was taken as the noise level. The results are shown in Fig. 4.23. The position of the Beetle chip on the module means that the 128 readout channels correspond to channels 384–511 on the TELL1. Also, the strip detector itself is only 50 strips wide, extending from channels 422–471. The figure shows the noise level after the pedestal subtraction, and after both the pedestal subtraction and the linear common-mode suppression (LCMS). These processing steps are described in section 3.4.4.

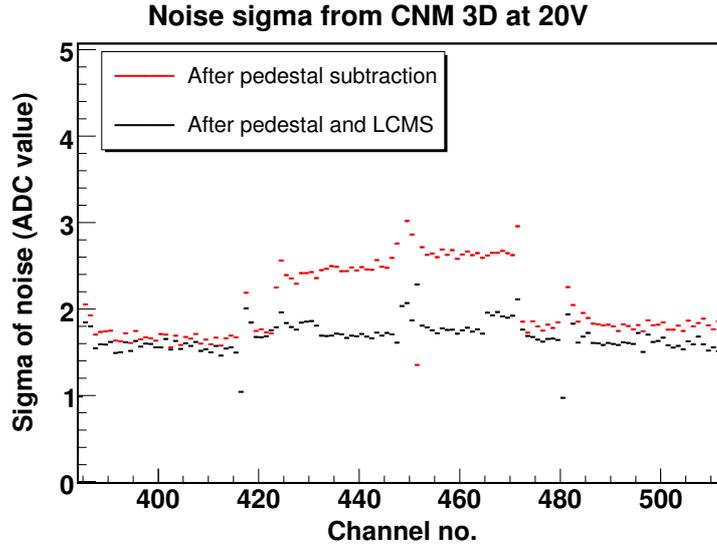


Figure 4.23: Noise sigma on each strip of the double-sided 3D strip detector at 20V bias. The results are shown after the pedestal subtraction phase, and after the subsequent Linear Common-Mode Suppression phase.

After the pedestal subtraction, but before the LCMS, the channels attached to the detector clearly show higher noise than the unbonded channels. At the start of each block of 32 strips, the noise is also higher due to header interference. Following the LCMS subtraction, the noise level on the strips is reduced, but the pattern of the noise levels becomes unusual. The LCMS subtraction is applied to the Beetle channels in blocks of 32, and because the detector has 50 strips we end up with bonded and unbonded channels within the same block. So, if a common-mode signal appears only on the bonded channels, the LCMS correction will be wrong. In particular, this means that the LCMS subtraction will be less accurate for the channels at the edge of the detector. (However, any common-mode signals appearing directly on the readout chip will be compensated correctly.)

The readout noise of a detector increases with its capacitance. However, the Beetle chip is designed to work with strip detectors that are much longer than the 4mm strips here, so after common-mode subtraction there isn't much difference in the noise level between the bonded and unbonded strips. The unbonded strips have a noise level of 1.57 ADC counts, corresponding to 710 electrons. The settings of the Beetle chip used during these tests were taken from Ref. [91], which reported that the unbonded channel noise was 730 ± 50 electrons. So, the results here are reasonable. Looking at the bonded strips near the centre of the detector (to reduce the LCMS discrepancy mentioned above) the mean noise level on the 3D detector after the

LCMS stage is 1.69 ADC counts, corresponding to about 760 electrons signal. The noise level seen on the 1cm-long planar detector is 1.73 ADC counts, corresponding to 780 electrons. So, there is negligible difference in the detectors' noise levels. According to the Beetle Reference Manual [90], the Beetle's noise increases by 50 electrons per picofarad of capacitance. Given the small difference between the noise on bonded and unbonded strips, it isn't possible to use these results as an accurate measurement of capacitance, but they imply that the detectors' capacitances aren't dramatically higher than 1pF per strip.

At 30V bias, the collection signal on the double-sided 3D detector was 11600 electrons, giving a signal-to-noise ratio of 15.

Unirradiated detector - Charge sharing

By looking at the cluster sizes seen on the double-sided 3D and planar strip detectors, it is possible to compare their charge sharing. Figure 4.24 shows histograms of the cluster sizes seen on the double-sided 3D detector at 20V and the planar detector at 120V. The 3D detector had fewer strips than the planar detector, so only the hits occurring on the central region of the planar detector were used in the analysis. This was to ensure that any angular spread in the beta particles from the source would affect both detectors equally. The figure shows that a large majority of the hits on the 3D detector are seen on single strips, whereas the majority of the hits on the planar detector are multi-strip clusters. This demonstrates the 3D detector's lower charge sharing.

The double-sided 3D detector gave substantially lower charge collection than the planar detector, which creates problems when comparing cluster sizes. The clustering algorithm used here searched for initial seeding strips with a signal-to-noise of 3 or greater, before adding any neighbouring strips whose signal was at least 10% of the signal on the seeding strip. This approach will tend to compensate for the 3D detector's low signal size, whereas using a fixed inclusion threshold for neighbouring strips would make it harder to form clusters on this detector. An additional complication arises because one in every four strips was unbonded on the planar detector. Only clusters with their peak signal on the central bonded strip in each group of three were analysed. As a result of these effects, these results cannot be considered to provide a quantitative charge sharing analysis. However, various charge-sharing tests with X-rays are reported in Chapter 5, which do provide a more reliable comparison between planar and double-sided 3D detectors.

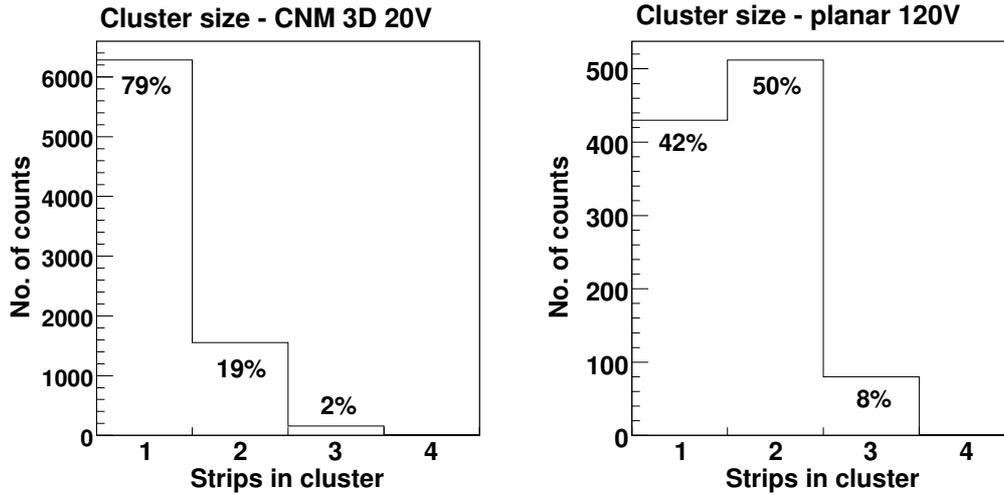


Figure 4.24: Histograms of cluster sizes obtained during beta tests of a double-sided CNM 3D detector at 20V, and a planar detector at 120V. The smaller cluster sizes in the double-sided 3D detector demonstrate its lower charge sharing.

4.3.4 Strip detector irradiated to 5×10^{15} $1\text{MeV-n}_{eq}/\text{cm}^2$

Another strip detector was irradiated to 5×10^{15} $1\text{MeV-n}_{eq}/\text{cm}^2$ using neutrons from the TRIGA Mark II reactor at the Jozef Stefan Institute in Ljubljana. The irradiations were performed by Vladimir Cindro. When testing this detector, an RC decoupling chip was connected between the detector and the Beetle chip, since the leakage current following irradiation was far too high for the Beetle front-end to cope with. Additionally, the detector had to be cooled during tests, to ensure that the leakage current was small enough for the bias supply to handle. The source, detector module, scintillator and photomultiplier tube were held inside a cooling chamber during the tests.

The detector showed strange noise behaviour. When the device was at room temperature and unbiased, the noise behaviour was fairly normal. This can be seen in Fig. 4.25. The channels bonded to the detector via the RC chip run from 270 to 319, and the noise level on these channels is only slightly higher than on the unbonded channels. However, when the device was cooled in the environment chamber, the noise on many of the strips increased dramatically. Figure 4.26 shows the noise when the chamber was cooled to -30°C and the detector was biased to 200V. (Note that the temperature of the detector itself is not measured directly, and will be higher than the chamber temperature during operation.) It can be seen that the strip noise is extremely large and variable, with a large common-mode component. The cooled detector's noise was high even when the detector was unbiased, which

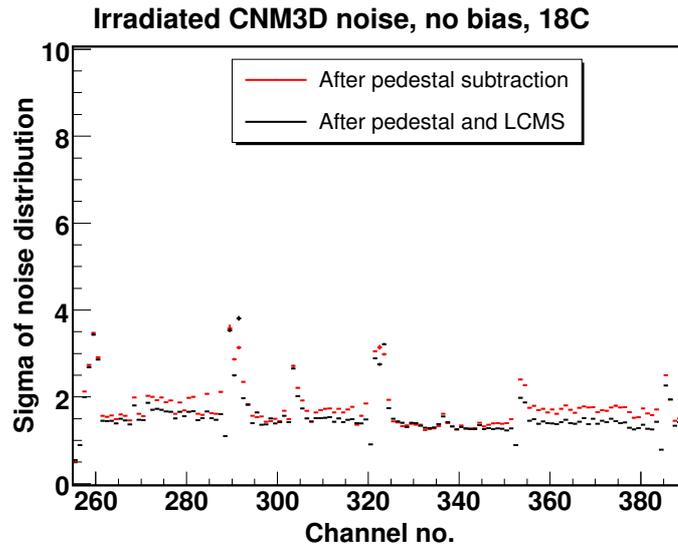


Figure 4.25: Noise level seen on the irradiated 3D strip detector, measured at room temperature with no applied bias. The noise on each strip is shown after pedestal subtraction and after both pedestal subtraction and LCMS. The detector is bonded to strips 270–319. The noise spikes every 32 channels are caused by the Beetle header crosstalk effect described in section 3.4.4.

shows that the effect is specifically due to the cooling rather than the biasing.

This noise was problematic during the MIP tests. To find the signal spectrum, only the signals on a set of strips with lower noise running from 308 to 316 were used. The resulting spectrum is shown in Fig. 4.27. The most probable signal measured on the detector at 200V corresponds to 12800 electrons. The signal spectrum seen on the detector is broader than in the tests on the unirradiated device. This will in part be due to the higher strip noise. However, it could also result from variations in the collection efficiency with position.

When the charge collection behaviour of this strip detector structure was simulated in Synopsys TCAD at 200V, using the methods described previously, the collection signal was 13300 electrons. In section 2.4.2, it was shown that at high radiation fluences the simulation model normally underestimates the charge collection. So, the experimental collection signal of 12800 electrons is a bit lower than would be expected. Nevertheless, the agreement between simulation and experiment is far better than for the unirradiated detector, which suggests that the use of AC coupling has improved the results.

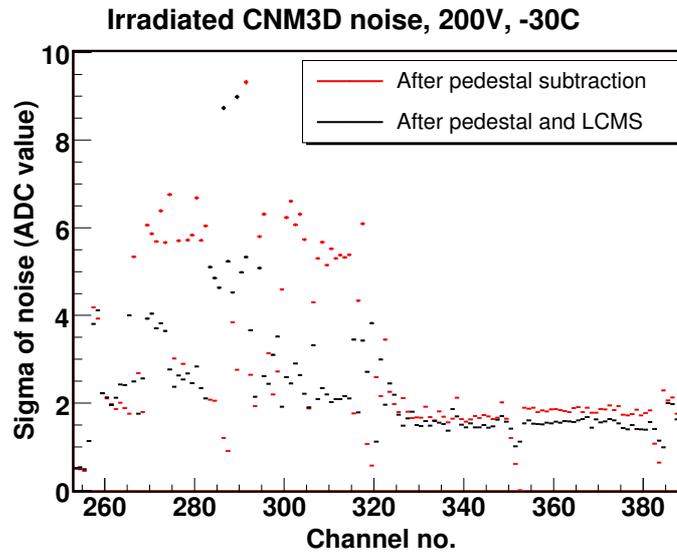


Figure 4.26: Noise level seen on the irradiated 3D strip detector, measured with the environment chamber at -30°C and with 200V bias. The noise on each strip is shown after pedestal subtraction and after both pedestal subtraction and LCMS. The detector is bonded to strips 270–319.

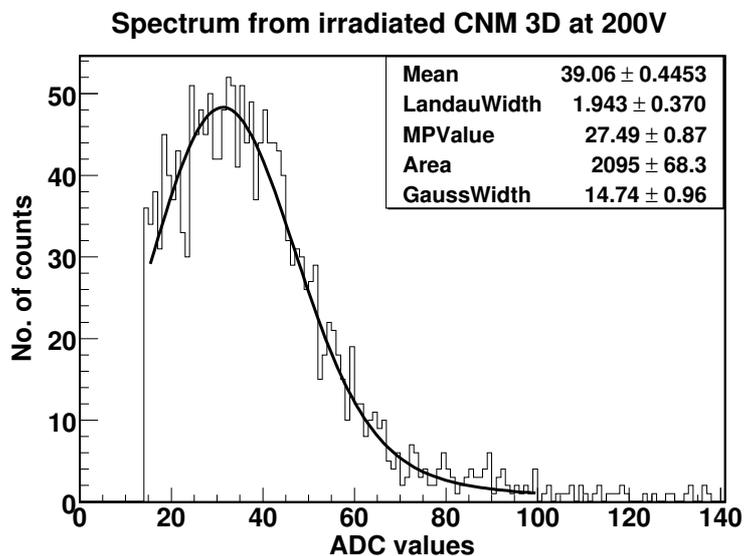


Figure 4.27: Signal spectrum from beta tests of a double-sided 3D strip detector irradiated to 5×10^{15} $1\text{MeV-n}_{eq}/\text{cm}^2$ and biased to 200V. The spectrum is fitted with a Landau convolved with a Gaussian. The events are taken from strips 308–316, where the noise is lower.

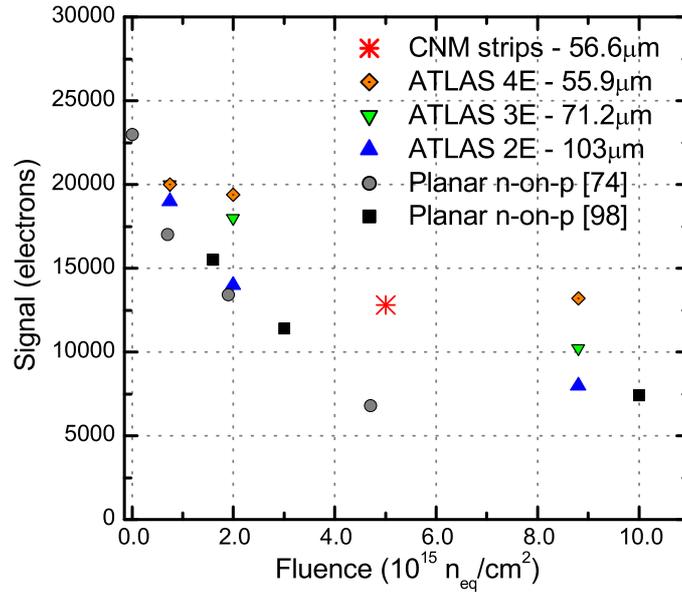


Figure 4.28: Charge collection for the double-sided 3D strips from CNM-IMB, full-3D detectors [97] and n-on-p planar sensors [74, 98]. The signal collected from MIPs is plotted against the irradiation fluence for each device.

Comparison to other detectors

Figure 4.28 compares the collected signal from this double-sided 3D detector with results from different full-3D detectors and two sets of results from planar n-on-p detectors.

The full-3D detectors were fabricated at the Stanford Nanofabrication Centre, and tested at Manchester [97]. Like the double-sided detectors, they have $250\mu\text{m}$ long columns. However, the full-3D structure means that the substrate is $250\mu\text{m}$ thick. Also, these detectors used n-type readout and had p-type substrates, whereas the detectors from CNM have p-type readout and n-type substrates. The full-3D devices used three different ATLAS-compatible electrode layouts, with 2, 3 and 4 n-type readout columns per ATLAS pixel—see section 2.4. These are referred to as 2E, 3E and 4E in the figure. The spacings between adjacent n- and p-type electrodes are also shown. The first set of n-on-p detector results are from Ref. [74], and were used back in section 2.3.4 to test the simulation of radiation damage. These were $280\mu\text{m}$ -thick strip detectors, with a pitch of $80\mu\text{m}$. These CCE tests were performed at a high bias of 900V. The second set of n-on-p results come from a recent paper which demonstrated unexpectedly high charge collection at high fluences [98]. These results were also obtained at 900V.

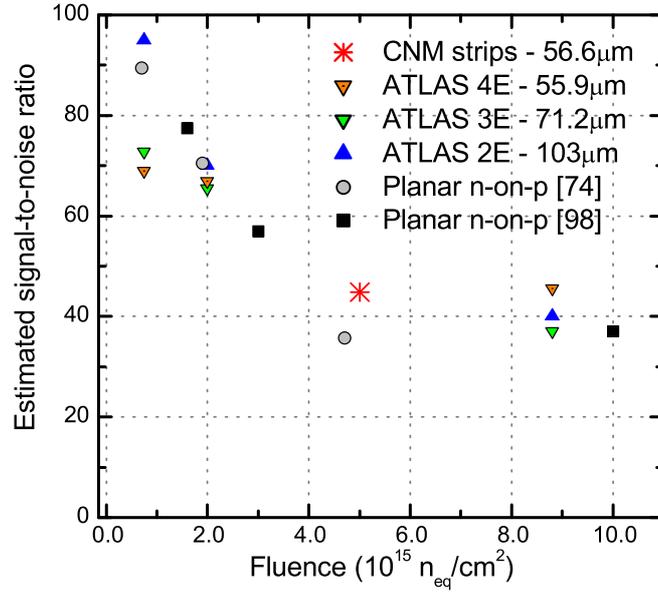


Figure 4.29: Signal-to-noise for the double-sided 3D strips from CNM-IMB, full-3D detectors [97] and n-on-p planar sensors [74, 98]. The noise values used here were not obtained directly from the irradiated devices; see the text for details.

Overall, the charge collection from the double-sided 3D detector is comparable to the results from the full-3D detectors. The double-sided 3D detector has a relatively small electrode spacing, matching most closely to the “ATLAS 4E” configuration, but its collection signal is more like that of a device with a wider electrode spacing. So although the charge collection is substantially higher than the n-on-p planar detector results from Ref. [74], and is achieved at a relatively low voltage, it is a bit lower than expected.

In Fig. 4.29, signal-to-noise ratios are plotted for each of these detectors. Note that none of these detectors had their signal-to-noise measured directly; indeed, the current ATLAS pixel readout chip isn’t radiation-hard to $1 \times 10^{16} n_{eq}/cm^2$, so such a measurement would be impossible. The noise levels for the three full-3D structures are from measurements on three unirradiated ATLAS-3D detectors [97]. These noise levels were higher than the values predicted from the simulations in section 2.5.4, possibly because the measurements came from prototypes devices. For the double-sided 3D detector, the noise level was assumed to match that of the “4E” full-3D detector, which has a very similar electrode spacing. For the n-on-p detectors, the typical noise level of an ATLAS module (185 electrons) was used [82].

Once the signal-to-noise ratio is taken into account, the differences between the

3D structures are reduced, since structures with smaller electrode spacing have both higher collection and higher noise. Also, the performance gap between the planar and 3D detectors is reduced, with the results from Ref. [98] roughly matching the 3D. However, as noted above the prototype 3D detectors had higher noise than expected, and future runs may give better noise performance.

In Ref. [99], the charge collection performances of a variety of detectors are compared, including the 3D detectors from Ref. [97], thin planar detectors, and n-on-n and n-on-p planar detectors. Using these, the authors obtain an empirical model relating electrode spacing to collection efficiency, which fits well to most of the available data. However, the new n-on-p sensors tested in Ref. [98] give double the collection efficiency predicted by this model at $1 \times 10^{16} \text{ n}_{eq}/\text{cm}^2$. There doesn't appear to be any dramatic difference between these new n-on-p sensors and those in Ref. [74], which agree with the model. Although it's possible that the new devices might have somewhat different substrate chemistry, they were intended to provide a comprehensive test of radiation damaged n-on-p (e.g. the effects of different annealing conditions) rather than testing new structures or materials. So, these results need further investigation.

4.4 Conclusions

Centro Nacional de Microelectronica (CNM-IMB) have produced a set of detectors with an alternative “double-sided” 3D structure. In this structure, the two sets of doped electrode columns are fabricated from opposite sides of the substrate, and neither set of columns passes through the full substrate thickness. Using this approach simplifies the fabrication in various ways: there is less risk of wafer breaking, no support wafer is required, there is less need to fill the columns with polysilicon, and the contact to the bias columns can be made more easily. However, this process does require double-sided alignment, and it is also less compatible with the active edge fabrication process.

This first fabrication run included pad, strip and pixel devices. (The test results from the Medipix2 pixel devices are presented in the next chapter.) All of these devices have $300\mu\text{m}$ -thick n-type substrates and $250\mu\text{m}$ -long columns, and use p-type readout.

Simulations of this device structure show that throughout most of the device volume, where the columns overlap, the electric field and charge collection behaviour match that of a full-3D detector. However, the regions around the very front and back surfaces have weaker fields and slower collection. Additionally, while the region

around the readout column depletes quickly, a somewhat higher bias is required to grow the depletion region to the back surface. These disadvantages are offset by the fact that the double-sided 3D detector can have a greater substrate thickness for a given column length. In simulations at very high radiation damage fluences, double-sided and full-3D detectors with the same column length give similar charge collection efficiencies.

In double-sided 3D detectors, avalanche breakdown is caused by high fields at the tips of the columns. In contrast, full-3D detectors break down due to high surface fields. Although these effects occur at comparable voltages (e.g. the field reaches $2.5 \times 10^5 \text{V/cm}$ at about 170V in these devices and in the 3D detectors in Chapter 2), the full-3D simulations assume that the p-spray isolation is well-matched to the oxide charge. With sub-optimal isolation, the full-3D detectors break down much earlier, whereas double-sided detectors are relatively insensitive to the design of the isolation. In principle, the full-3D detectors' breakdown behaviour might be improved further by using more sophisticated isolation (e.g. moderated p-spray), whereas the double-sided 3D detectors have less room for improvement. However, since both the front and back surfaces of a full-3D detector require isolation, anything more than a simple p-spray would make the fabrication substantially more complicated.

I-V tests on four strip detectors showed quite a lot of variation in their performance. One broke down below 10V, and the guard ring currents varied dramatically from device to device, with one detector reaching a guard current of $20 \mu\text{A}$ at just 6V. However, the strips within each detector gave reasonably low currents of around 100pA. A newer set of detectors fabricated by CNM, with p-type substrates, n-type readout and p-stop isolation, gave lower and more consistent guard currents. This is probably because these devices used p-stops to reduce the surface current flow.

The C-V curves obtained from three pad detectors indicated a two-stage depletion process, with the capacitance dropping rapidly up to 2.3V as the region between the columns depleted, then decreasing more slowly until around 9V as the depletion grew to the back surface. The saturation capacitances of the detectors corresponded to 36.3fF, 34.3fF and 28.5fF per column. Not only are these results quite variable, they are also substantially lower than the simulated value of 64fF. However, the newer set of pad detectors produced by CNM gave saturation capacitances which were much closer to the simulated value.

Charge collection efficiency tests with betas on the double-sided 3D strip detectors produced mixed results. An unirradiated strip detector produced a signal spectrum with a clear Landau peak, even at low biases, demonstrating that the detector works. However, the charge collection efficiency after depletion was unex-

pectedly low, at around 50%. In contrast, the Medipix2 pixel detectors tested in section 5.5.2 didn't show this kind of signal loss. The signal pulse measured from the detector at 20V gave no sign of slow charge collection, and the simulations also showed that the collection should be fast, so the signal loss probably wasn't due to ballistic deficit. The strip detector was DC-coupled to the readout chip, but the chip was designed for AC-coupling, so this may have altered its behaviour. However, in principle the chip should be able to operate in DC-coupled mode, and there was no clear evidence of a problem with the chip.

The CCE tests were also performed with an AC-coupled strip detector irradiated to 5×10^{15} 1MeV- n_{eq}/cm^2 . After it was cooled, the irradiated detector had some problems with noisy strips. However, once the noisy strips were masked, a high signal of 12800 electrons was achieved at 200V bias. Although this charge collection is slightly lower than the simulated value, the values are in far better agreement than they were for the unirradiated detector. So, this also supports the idea that the low signal on the unirradiated detector was due to problems with the DC coupling. The irradiated double-sided detector's charge collection is comparable to experimental results from full-3D detectors, and superior to typical planar detectors. So, this confirms that double-sided 3D detectors are a promising option for high-radiation environments.

A new, larger run of double-sided 3D detectors has been fabricated at CNM, and will be tested at Glasgow, CNM and Freiburg. An important first step will be to test an unirradiated strip detector with AC coupling, to see if this solves the low-signal problem seen here. Since a large number of strip detectors are available, it will be useful to test their collection efficiency at a range of damage fluences, up to an extremely high fluence of 2×10^{16} 1MeV- n_{eq}/cm^2 . Since some of the detector wafers have p-substrates and n-type readout, this new run will also provide double-sided 3D ATLAS pixel detectors. Signal-to-noise measurements from these devices would provide a more direct comparison between double-sided 3D and other detector technologies for the Super-LHC.

Chapter 5

3D detectors for synchrotron X-ray experiments

Synchrotron light sources such as Diamond generate extremely bright X-ray beams which can be used in a variety of experiments, as discussed in section 1.5.

In this chapter, section 5.1 describes how position-sensitive X-ray detectors are used in crystallography experiments to determine the structures of biological molecules and other materials. Section 5.2 then describes the Medipix2 readout chip. When coupled to this, or a similar chip, photodiode detectors can operate in a single-photon-counting mode, which brings various advantages compared to standard integrating detectors. 3D detectors offer further advantages over planar photodiodes, due to their low charge sharing and the possibility of using active edges. Section 5.3 describes a set of three Medipix2 3D detectors, which were produced at CNM-IMB and use a double-sided structure. These detectors were initially tested with an X-ray tube, as described in section 5.4. Following this, the detectors were tested with monochromatic X-rays on a beamline at Diamond Light Source. The spectral response and line spread function of a 3D and a planar detector were measured, to compare their charge sharing behaviour—see section 5.5. Also, a standard powder diffraction experiment was done with the detectors. Finally, section 5.6 describes further charge-sharing tests carried out with an alpha particle source.

5.1 X-ray crystallography

5.1.1 Theory, experiment and applications

A crystal has a lattice structure, where the most basic unit of the crystal, called the basis, is repeated over and over to form a regular array. The basis may range from

a single atom (for example in sodium) to a large organic molecule such as a protein.

In an X-ray crystallography experiment, a narrow beam of X-rays is fired at a crystal sample. Interference effects will occur when the X-rays scatter from the atomic electrons in the crystal, producing a diffraction pattern. Due to the crystal's lattice structure, the pattern will consist of a series of distinct diffraction spots. Each spot is produced when the X-rays reflect from a specific set of crystal planes and give constructive interference. This will occur when Bragg's Law is satisfied, i.e.:

$$n\lambda = 2d_{hkl} \sin \Theta \quad (5.1)$$

where n is an integer, λ is the X-ray wavelength, d_{hkl} is the spacing between the family of planes causing the reflection (which can be described by indices h k and l) and Θ is the angle of both the incoming and reflected X-rays with respect to the planes (hkl). Note that whether or not this condition will be met for each set of planes will depend on the orientation of the crystal relative to the X-ray beam [100].

While the positions of the diffraction spots are determined by the spacing and angles of the crystal planes, their relative intensities are determined on the electron density distribution of the basis. This electron density depends on the position and element of each atom, so by measuring the relative spot intensities accurately the structure of the basis can be found. A 12keV photon will have a wavelength of around 1\AA , so X-ray diffraction can determine the structure at an atomic scale.

Both the spot positions and intensities will be affected by the orientation of the crystal with respect to the beam. So, to get the full 3D structure of a crystal, a series of 2D diffraction patterns must be taken with the sample at different angles. An example of a diffraction pattern from a thaumatin crystal is shown in Fig. 5.1. The image was recorded with a PILATUS hybrid pixel detector [101].

One variant of this technique, macromolecular crystallography, is the most commonly used method for determining the structure of large biological molecules such as proteins and nucleic acids. In fact, out of the approximately 50000 structures reported in the Protein Data Bank, 85% were found using X-ray crystallography [102]. Proteins consist of long chains of amino acids, and a protein's function is heavily dependent on how this chain folds up into a three-dimensional structure—for example, the structure of an enzyme allows it to catalyse a specific chemical reaction. So, macromolecular crystallography is important for understanding how the body functions at a cellular level, and how different drugs and diseases affect the body.

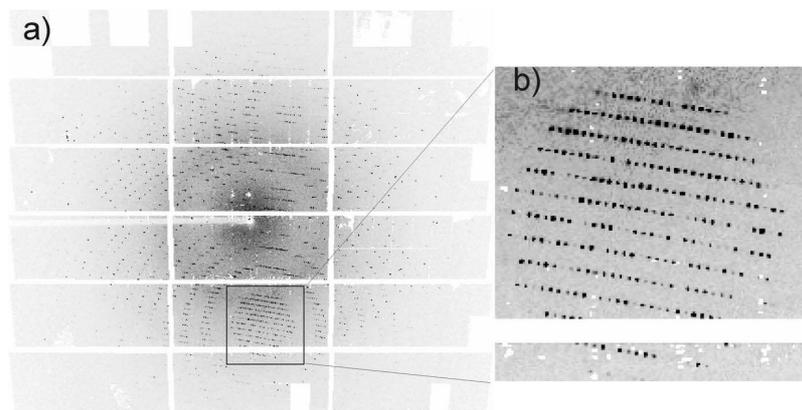


Figure 5.1: (a) Diffraction image of a thaumatin crystal taken with a PILATUS 6M detector. (b) Zoom of a small region. Reproduced from Ref. [101] with permission from IUCr (<http://journals.iucr.org>).

Detector requirements

The quality of the structural information obtained during a diffraction experiment will depend on how accurately the detector can measure the positions and intensities of the spots in the diffraction pattern. This is particularly important when applying X-ray crystallography to large, complicated structures such as viruses [103].

The signal-to-noise ratio in any X-ray pattern is fundamentally limited by statistical variations in the number of photons arriving at each point in the image. Suppose that we acquire a series of X-ray images under identical conditions with a perfect detector. If a particular region contains an average of N photons, then the standard deviation of the number of photons will be \sqrt{N} , giving a signal-to-noise ratio of \sqrt{N} [104].

New synchrotrons such as Diamond Light Source provide extremely bright X-ray beams. This increases the intensity of the diffraction spots, improving the intrinsic signal-to-noise ratio. The high brightness also makes shorter exposure times possible, allowing faster experiments, more acquisitions from each sample (improving the data quality [105]) and offering the possibility of time-resolved measurements of changes in structure. However, this in turn places greater demands on the detector, which must meet the following requirements [106]:

- When measuring weak diffraction spots, any electronic noise introduced by the detector will degrade the signal to noise ratio. So, low-noise detectors are needed.

- The brightest diffraction spots will be hit by a large number of photons during an acquisition, so the detector must have a reasonably high count rate and a large, linear dynamic range to measure the intensity accurately.
- To take images at a high rate without an unacceptably long dead time between acquisitions, the detector should have a readout rate of 10ms or less.
- When the X-ray beam interacts with the crystal, Compton scattering or fluorescence can occur. The lower-energy photons produced by these processes contain no diffraction information, and worsen the signal-to-noise ratio. Ideally, energy discrimination should be used to reject these photons [107].
- The detector should have good spatial resolution, in order to distinguish the diffraction spots and measure their positions.

5.2 Medipix2 single-photon-counting detectors

5.2.1 The Medipix2 readout chip

Medipix2 [7] is a pixel readout chip designed specifically for X-ray detection, particularly in medical applications such as digital radiography [108]. It has an array of 256×256 pixels, each of which is $55\mu\text{m} \times 55\mu\text{m}$, giving an active area of 14.08 by 14.08 mm. Like other hybrid pixel readout chips, it is bump-bonded to a sensor chip. It operates in a “single photon counting” mode, where it counts all the individual X-rays hitting the sensor which fall within a certain energy range.

Figure 5.2 shows a schematic of the circuitry within each Medipix2 pixel, taken from the Medipix2 manual [109]. The first stage of the processing is signal amplification. Medipix2 uses a single-stage charge preamplifier, which integrates the current signal produced when a particle hits the detector. The preamplifier can accept signals of positive or negative polarity, allowing it to be used with a wider variety of sensors. It is designed to be able to compensate for positive or negative leakage currents. Following a hit, the preamplifier signal returns to the baseline level in under $1\mu\text{s}$, giving a maximum possible count rate of 1MHz. In most Medipix2 applications, hits on the sensor will occur at random times, so the actual hit rate needs to be substantially lower than this to ensure that the proportion of events experiencing pile-up is low.

The output of the preamplifier feeds into a pair of discriminators, each of which has an adjustable threshold, $V_{\text{th Low}}$ or $V_{\text{th High}}$. If the amplitude of the pulse falls between these two thresholds, then the Basic Window Discriminator (BWD)

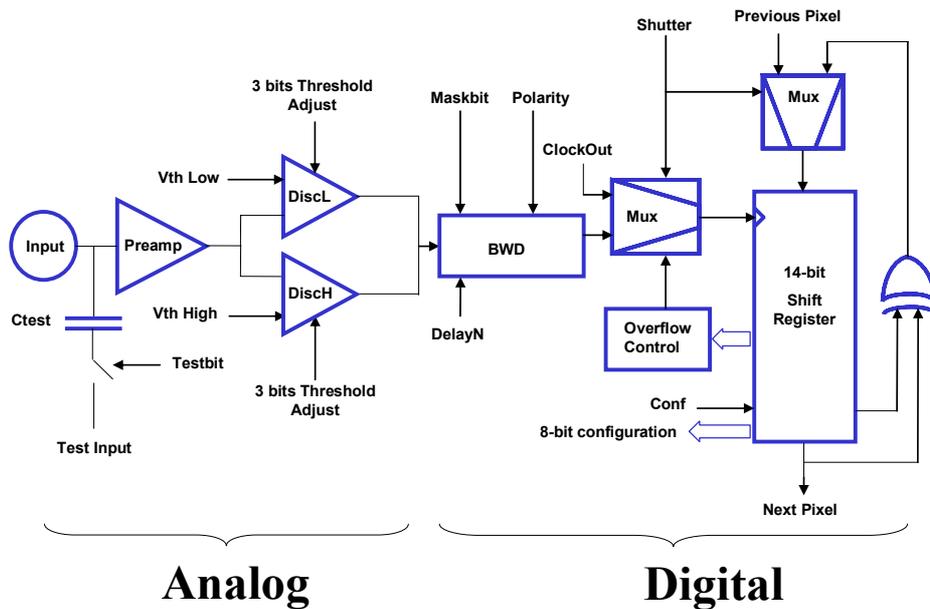


Figure 5.2: Circuitry within each Medipix2 pixel. A description of the component blocks is provided in the text. Reproduced from [109] with author's permission.

produces a digital output pulse. So, the circuit is able to accept hits on the detector falling within a specific energy window, and reject those that don't. The threshold values are adjusted by programming the chip with digital low threshold (THL) and high threshold (THH) values, which are then converted to actual voltage levels by digital-to-analogue converters. Each threshold has a 14-bit range.

The device operates in two modes; acquisition and readout. During acquisition mode, whenever the discriminator accepts a hit, the 14-bit shift register is incremented, acting as a counter. (Each register has a maximum range of 11810 counts, after which any additional hits will be ignored.) Then, during readout mode, the total count from each pixel is read out by using the shift registers to pass the data along each column of 256 pixels to a fast shift register at the edge of the chip.

During operation, the Vth Low and Vth High thresholds are set globally across the detector. However, there will inevitably be some variation between the preamplifiers and discriminators in different pixels, meaning that the effective threshold energies will vary from pixel to pixel. To compensate for this, each pixel contains an 8-bit register which allows adjustment of the thresholds. Each of the two discriminators has three adjustment bits, which alter the threshold level by connecting three independent current sources to the discriminator. The procedure for performing this adjustment is described section 5.4.3. The remaining two bits are a mask bit, which can switch off noisy pixels entirely, and a test pulse select bit. The chip can apply

test pulses to each pixel using a capacitor (see Fig. 5.2)

5.2.2 Advantages and disadvantages of single-photon-counting detectors for X-ray detection

The single-photon-counting approach relies on the fact that each photodiode pixel has its own readout channel, amplifier and other circuitry, which will immediately process the signals from the sensor. Most other silicon detectors such as CCDs [13] and CMOS active pixel sensors [110] will integrate all the charge collected in each pixel during the acquisition time. These detectors generally have a thinner active region, and so instead of absorbing the X-rays directly they are normally coupled to a segmented scintillator [111], which produces visible light when hit by X-rays. These detectors are discussed more fully in section 1.1.6.

While single-photon-counting photodiodes are comparatively expensive and complex, they offer superior performance for X-ray crystallography. Their advantages and disadvantages are discussed below.

- *Electronic noise*

In an integrating detector, any electronic noise is integrated over the acquisition time, along with the signal. In contrast, a single-photon-counting detector will reject this noise, since its output is simply the total number of hits detected. However, the single-photon-counting detector's noise will affect its energy resolution.

- *Energy resolution.*

In an integrating detector, energy resolution of single photons is not possible. In X-ray crystallography, a single-photon-counting detector can distinguish lower energy photons from Compton scattering or fluorescence from the original monochromatic beam. In applications which use a source with a wide spectrum, such as medical imaging, each photon energy will carry different information. This will be exploited by newer chips such as Medipix3 [112], which will obtain basic spectral information by using multiple thresholds to define a series of energy ranges.

- *Readout time*

Reading out from Medipix2 can take as little as $266\mu\text{s}$ with a parallel readout connection or 8ms with serial [7], so these detectors can meet the 10ms readout time requirement for X-ray crystallography at new light sources. In contrast,

the process of reading out a CCD will take at least a few tenths of a second, as discussed in section 1.1.6. However, CMOS detectors can potentially achieve high readout speeds.

- *Count rate and dynamic range*

A single-photon-counting detector has the advantage of a large, linear dynamic range, limited only by the size of the counter in each pixel. However, if the hit rate on a single-photon-counting detector is too high, then the preamplifier pulses produced by the photons will pile up, meaning that individual hits will not be correctly distinguished. Conversely, a pixel in an integrating detector will have a limited dynamic range but there is no limit on the actual hit rate. On the whole, a single-photon-counting detector is better for sensing intense but approximately continuous X-ray beams in standard crystallography experiments. However, in new pulsed X-ray sources such as XFEL [113], integrating detectors will be essential.

- *Spatial resolution.*

Due to the simpler pixel structure of a CCD, it can achieve a smaller pixel size of down to $5\mu\text{m}$ [111], whereas single-photon-counting photodiodes are limited to around $50\mu\text{m}$. Since the incoming image is effectively being sampled at a spacing equal to the pixel size, this places a fundamental limit on the spatial information a detector can record. According to Fourier analysis, any image can be built up from a series of sine waves of varying spatial frequency, amplitude and phase. A detector with a pixel size of d can only accurately reconstruct an image containing spatial frequencies below the Nyquist limit of $1/2d$ [114]. So, for example, a detector with $55\mu\text{m}$ pixels can only accurately reconstruct features with a spatial frequency of less than 9.1 line pairs per mm.

However, the resolution of a detector is dependent on the entire imaging system. In a system based on a scintillator coupled to a CCD, the additional stages will introduce additional blurring (such as scattering of light inside the scintillator crystal before it reaches the CCD), reducing the image quality compared to a direct-conversion detector like a photodiode [115].

5.2.3 Charge sharing, X-rays and 3D detectors

In a single-photon-counting detector, charge-sharing between pixels can reduce the image quality in various ways. If the threshold is close to the photon energy, then

the signal might not be large enough to produce a “hit” in any of the pixels, reducing the detection efficiency and worsening the signal to noise ratio. Alternatively, if the threshold is low, then the hit may be registered in more than one pixel, blurring the image and reducing the detector’s ability to resolve fine features. If the incident beam contains a wide spectrum, both these effects may occur. This will also mean that higher-energy photons will effectively have a higher “weighting” than low-energy photons, which can reduce the image quality. Additionally, charge-sharing reduces the detector’s spectral resolution, for example making it harder to exclude Compton-scattered photons.

As described in section 1.2, the 3D detector structure offers reduced charge sharing compared to a standard photodiode, making it potentially useful for single-photon-counting applications.

Higher-energy X-rays have a greater absorption lengths. So, photodiodes with greater substrate thicknesses would be required to detect these X-rays with a high efficiency. While planar detectors suffer from increased charge sharing as the substrate is made thicker, 3D detectors will not, so in principle thick 3D X-ray detectors would be particularly useful. However, as discussed in section 1.2.2, there are practical limits to the column length that can be achieved with a given column radius. So, producing 3D detectors with thick substrates and small pixels would require improved fabrication technologies.

Although none of the detectors tested in this thesis have active edges, this is another potential advantage of 3D detectors. Reducing the dead area at the edge of each 3D detector would make it possible produce large-area X-ray detectors with very little dead area, as discussed in section 1.2.3. In X-ray crystallography, diffraction spots falling across dead areas of the detector will be lost, so minimising the dead area is an important consideration.

5.3 3D Medipix detectors

5.3.1 Detector production

A set of 3D Medipix2 detectors were produced by CNM as part of the double-sided 3D fabrication run described in Chapter 4. The devices were fabricated on a $300\mu\text{m}$ -thick n-type substrate. Figure 5.3 shows the front surface of a detector after fabrication. The detector has a $55\mu\text{m}$ spacing between adjacent columns of the same type, to match the pixels on the Medipix2 chip. Like in the 3D strip detectors, the p+ columns are etched from the front surface and used for readout. At the outer

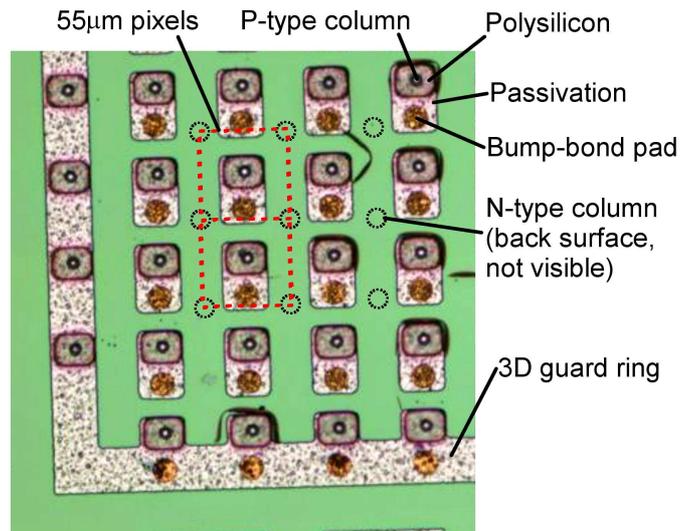


Figure 5.3: Part of the front surface of the 3D detector before bump bonding to Medipix2. The n-type columns etched from the back are not visible, but the positions of a few have been indicated by circles.

edge of the chip, a ring of p+ columns are connected by a metal track to form a 3D guard ring, which protects the pixels from any surface leakage current. Within each pixel there is a bump-bond pad, which is used to connect the detector to the readout chip. After fabrication, the sensor was Indium bump-bonded to a Medipix2 chip by VTT. VTT have previous experience of producing simpler “semi-3D” sensors and bonding them to Medipix2 [116]. After bump-bonding, the Medipix2 chip was attached to a standard Medipix2 chipboard (with the back surface of the 3D detector facing upwards), and the readout connections on the chip were wirebonded to the chipboard.

In this first production run, three Medipix2 3D sensors were bump-bonded. In the following sections, they are referred to as 3D1, 3D2 and 3D3. A planar p-on-n Medipix2 sensor, which was also $300\mu\text{m}$ thick, was tested for comparison.

5.3.2 Data acquisition

The Medipix2 chipboards were programmed and read out using a Medipix2 interface produced by IEAP, Czech Technical University, Prague [117]. Figure 5.4 shows one of the detectors connected to the interface. The interface communicates with a control PC via USB. It incorporates a power supply for the Medipix2 chip, and a variable high-voltage power supply which can be used to bias the sensor. Along



Figure 5.4: Medipix2 chipboard connected to IEAP's USB interface.

with the interface, IEAP have developed “Pixelman” software to control the system [118].

5.4 Initial 3D Medipix tests

5.4.1 Digital write / read test

Firstly, a basic test of the Medipix2 chips' digital circuitry was performed, where values were written to the shift registers in each pixel then read back. This test had already been performed on the bare Medipix2 chips after fabrication, and the chips were known to each have 2 columns of 256 pixels which had failed. (The highest-quality readout chips are generally reserved for more developed detector systems, rather than prototypes like these.) When this test was done with the bump-bonded 3D sensors, 2 columns of dead pixels were seen on each chip as expected. These dead columns can be seen in Fig. 5.7 later.

5.4.2 Detector biasing and leakage current

The 3D detectors were biased using the built-in high voltage source on the USB interface, which was connected to the Medipix2 chipboard via a LEMO cable.

On the Medipix2 chipboard, the high voltage passes through an R-C network before it is applied to the sensor, to limit the current that can flow through the sensor and to smooth out sudden changes in voltage. The resistor is $100\text{k}\Omega$, and it was found that when biasing the 3D detectors the current was large enough to produce a substantial voltage drop across this resistor. By measuring this voltage drop while varying the supply voltage, it was possible to determine the voltage applied to the sensor itself, and the corresponding current flow through the detector.

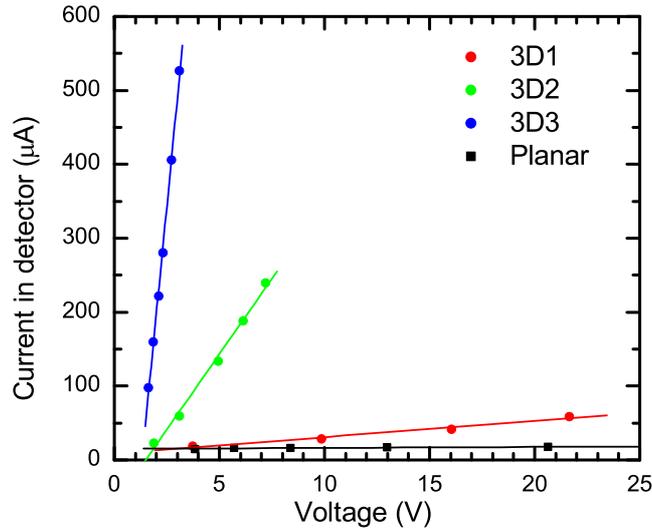


Figure 5.5: Total detector current vs voltage for 3D and planar Medipix2 detectors.

The resulting I-V curves for the three 3D detectors are shown in Fig. 5.5. All three curves are linear, but with a wide range of gradients. In section 4.3.1 it was found that the double-sided 3D strip detectors had large, extremely variable guard ring currents. It is likely that the current draw in these Medipix2 devices is also dominated by surface leakage effects.

5.4.3 Threshold equalisation

Next, threshold equalisation was applied to each detector, using the Pixelman software. As described previously, each pixel of the readout chip has 3 low threshold (THL) and 3 high threshold (THH) adjustment bits, which can correct for variations in the effective threshold. The appropriate adjustment is found by using a reference such as a known X-ray source, the intrinsic noise in each channel or test pulses. In the tests performed with the detector, only the low threshold THL was used, so only the THL equalisation needed to be performed.

The equalisation was carried out using the centroid of the noise on each channel. Since the Medipix2 chip allows bipolar readout, the “zero” signal level lies around the middle of the preamplifier’s dynamic range. Firstly, the THL adjust bits were set to 0 in every pixel. The global THL threshold was scanned from a level corresponding to a positive-polarity signal to a level corresponding to negative polarity, taking a 0.1s acquisition at each step. As the threshold passed through the noise level, the

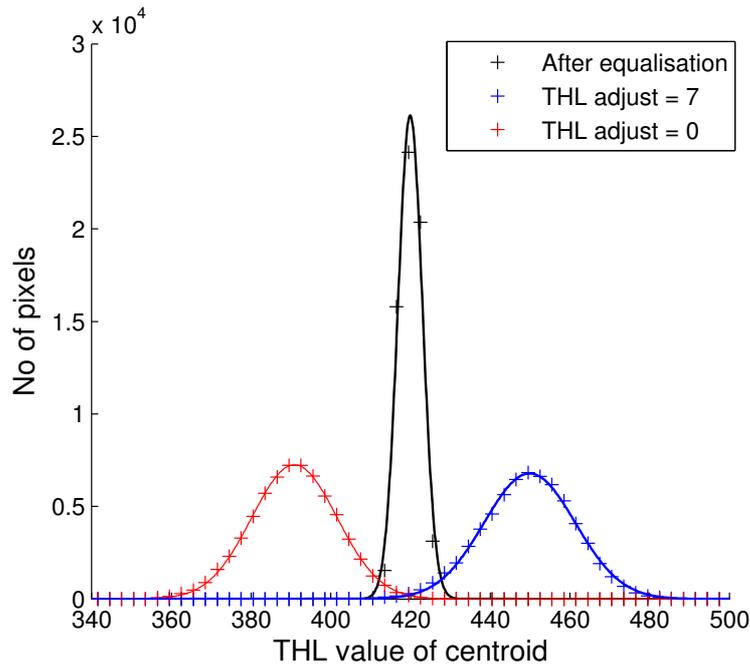


Figure 5.6: Distributions of noise centroid positions in each pixel of the 3D1 detector. Results are shown with all THL bits set to 0 or 7, and then after equalisation.

number of counts seen on each pixel rose then fell. Using this data, the Pixelman software calculated the centroid of the noise in each pixel, for this THL adjust value.

This process was repeated to find the noise centroid in each pixel with THL adjust set to 7, the maximum value. Using these two sets of data, the software calculated the THL adjust setting for each pixel that would make the noise centroid level as uniform as possible across the detector. The THL scan was repeated for a final time, to find the centroid positions after equalisation. Figure 5.6 shows a histogram of the noise centroid positions found during the THL scans on detector 3D1. During the equalisation, the 3D detector was biased to 20V.

A Gaussian fit was made to each distribution. With all the THL values set to 0 or 7, the threshold distribution has a sigma of 11.1 THL, whereas after equalisation it has narrowed substantially to 3.1 THL. Using the energy calibration found in section 5.5.2, this corresponds to a reduction from 1.9keV to 0.55keV.

The equalisation process was applied to each of the 3D detectors, and also a 300 μ m-thick planar p-on-n Medipix2 sensor. (The planar sensor was biased to 100V during equalisation, to ensure full depletion.) The widths of the pixel threshold distributions were approximately the same in each case. However, there were chip-to-chip variations in the THL value corresponding to the noise centre after equalisation. For example, while 3D1's noise centroid was at THL 419.9, the planar detector tested

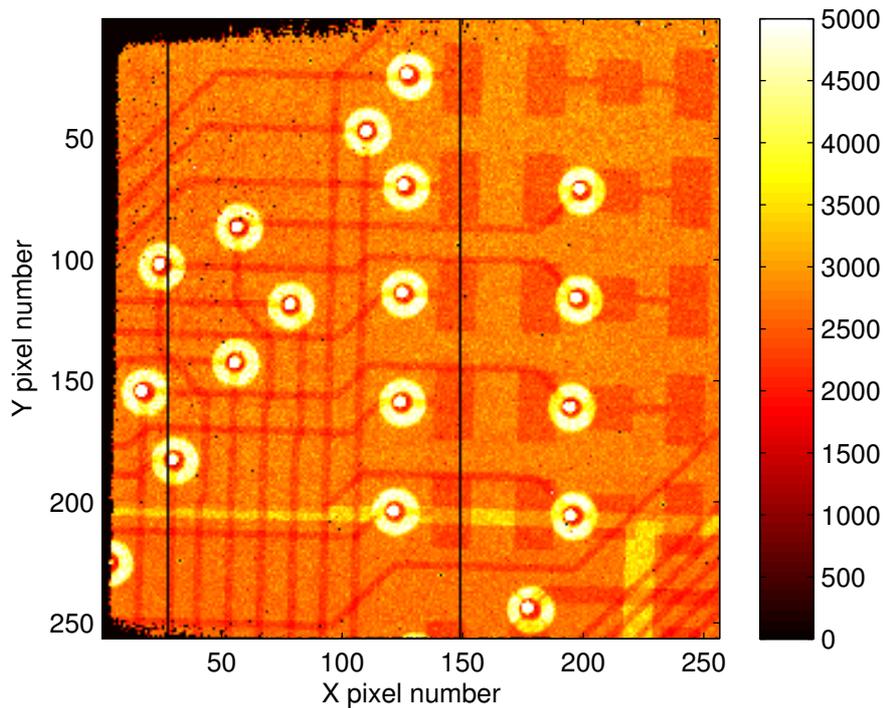


Figure 5.7: Image of a PCB, taken with detector 3D2 and an X-ray tube. The scale shows the number of photon hits per pixel.

for comparison gave THL 447.8. This kind of chip-to-chip variation has also been seen between planar sensors [116], so it is not a specific feature of these 3D devices.

5.4.4 Basic image acquisition

After equalisation, the detectors were used to take basic test images. A PCB was held over the detector, and was imaged using an X-ray tube. The tube used a Tungsten target and a tube voltage of 60kV to produce a wide Bremsstrahlung spectrum, with peak intensity at around 30keV. During the acquisitions, only the low threshold (THL) was used. Figure 5.7 shows an image of the PCB taken using detector 3D2. The X-ray illumination lasted for 0.84s, and the threshold setting was 300 for this particular acquisition, which is approximately 20keV (using the calibration from section 5.5.2).

It was found that 3D2 and 3D3 had a narrow dead region running along the entirety of one edge, and extending into the adjacent edges. The dead region on 3D2 is about 8 pixels wide. These dead regions will be caused by detached bump bonds. For successful bump bonding, both the detector and readout chips need

to be extremely flat, and VTT reported that the 3D chips were slightly bent. It's possible that this bending occurred during fabrication when the polysilicon was being deposited, since this stage places the wafer under stress. However, damage during the saw-cutting could also have this effect. The 3D1 sensor did not suffer from these dead pixels. This sensor also had the lowest current draw, so it was used for most of the tests that follow.

5.4.5 Detection efficiency versus voltage

Next, using the X-ray tube the count rate on the detector 3D1 was tested at different voltages. At each voltage setting the detector was illuminated with the tube for 0.84s. Initially, the detector's threshold was set to 300, approximately 20keV.

As described above, the threshold equalisation on 3D1 had been performed at 20V, and the position of the noise centroid was found to be 419.9 after equalisation. However, repeating the equalisation at 40V caused a small shift in the noise centroid position, by a THL value of about 5. This was presumably due to the change in leakage current. So, when operating the detector at higher voltages the THL value had to be adjusted slightly, to ensure the effective energy threshold was kept constant.

As seen in Fig. 5.8, the count rate increases very rapidly as the voltage is increased from zero to 2V. It then increases more slowly up to about 9V, and then levels out. The preamplifiers in Medipix2 have a peaking time of about 150ns [109], which is slow compared to the 3D detector's collection time, so the count rate will basically be determined by the depleted volume. The simulations in section 4.2 showed that the region in the double-sided detector where the columns overlap should deplete at just a couple of volts, but a higher voltage is required to deplete the region at the back surface of the detector. So, there is good agreement between the simulation and experiment. The C-V curves from the CNM pad detectors in section 4.3.2 also followed this pattern.

5.4.6 Electronic noise, pixel masking, and radiation damage

Since it counts individual photons, rather than integrating the detector current over the acquisition time, the Medipix2 chip is not affected by electronic noise in a conventional way. However, if the noise fluctuations are large enough, a pixel can register false hits. This will of course depend on the threshold setting.

As a test of detector noise, the low threshold was scanned through a range of values without any beam present, with 0.1s acquisitions at each setting. Figures 5.9

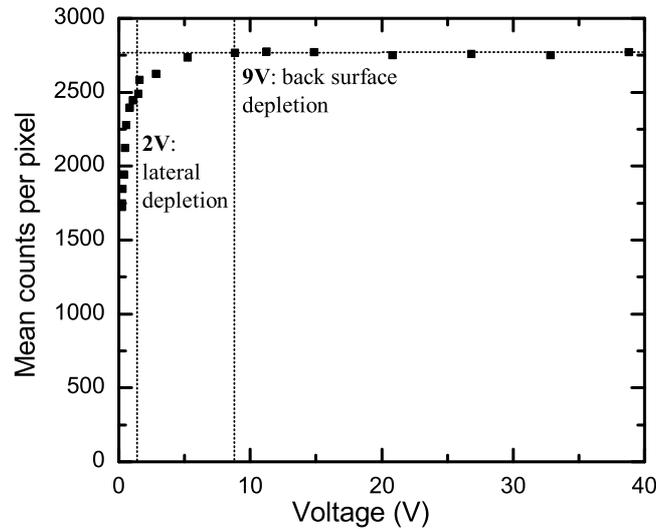


Figure 5.8: Count rate per pixel detected on 3D1 at different bias voltages, using the X-ray tube as illumination. The voltages at which lateral depletion between the columns and depletion to the back surface occur are indicated.

and 5.10 show the results for the 3D1 and planar sensor respectively. The number of pixels with greater than or equal to 1, 10, 100 and 1000 counts in 0.1s are shown versus the threshold energy. (Section 5.5.2 describes how the calibration of THL values to energy was found.) At each threshold setting, the 3D detector has more noisy pixels. To take an example, at an energy of 6keV the 3D detector has 66 pixels with over 100 noise counts, whereas the planar detector has only one.

However, the spectral measurements in section 5.5.2 show that, once this relatively small number of noisy pixels are masked, both detectors give a similar peak width when tested with monochromatic X-rays. Since the peak width will be proportional to the mean electronic noise (plus the effects of threshold dispersion), this suggests that most of the pixels on the 3D detector are not particularly noisier than those on the planar. Instead, it seems that the 3D detector has a higher proportion of pixels with poor noise performance. This may be related to some non-uniformity in the device or the bump bonding. However, the noise can also depend on the individual Medipix2 readout chip. During the data analysis, noisy pixels were appropriately masked.

During tests of the 3D1 and planar detectors at Diamond—see section 5.5—radiation damage occurred to both Medipix2 chips. The spectra measurements in section 5.5.2 and the Line Spread Function (LSF) measurements in section 5.5.3

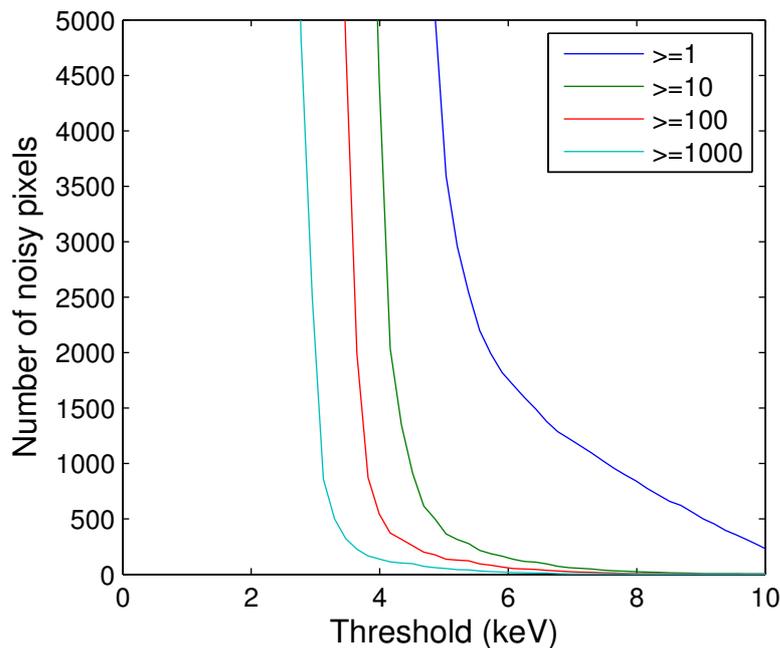


Figure 5.9: Noise results from 3D1. The lines show the number of pixels with greater than or equal to 1, 10 etc. noise counts in 0.1s, versus the threshold energy setting.

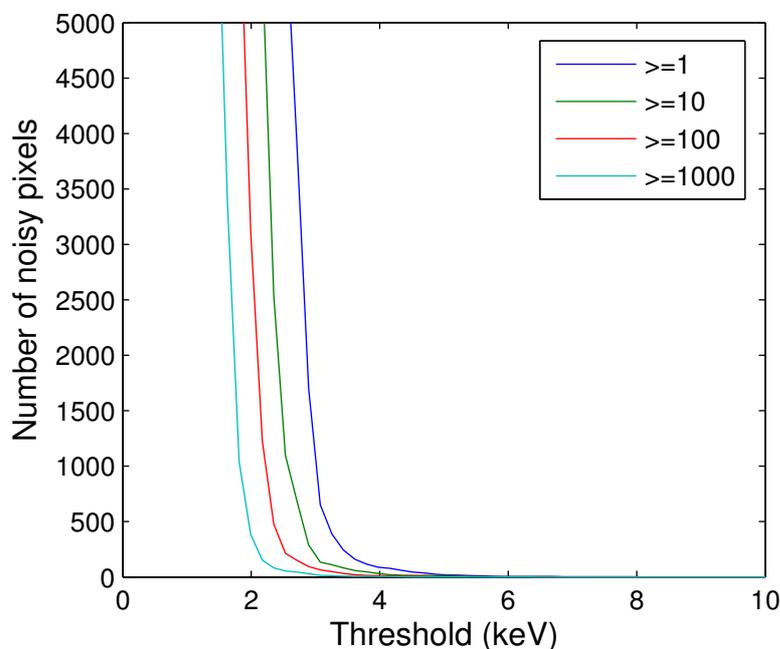


Figure 5.10: Noise results from the planar detector. The lines show the number of pixels with greater than or equal to 1, 10 etc. noise counts in 0.1s, versus the threshold energy setting.

were done during the course of a single day, using relatively high X-ray fluxes, and the damage was observed when testing the detectors at the start of the following day.

Firstly, the number of noisy pixels had increased on each chip. Secondly, when the threshold equalisations were repeated, the THL adjust values had changed for some of the pixels, and the global “zero signal” THL level had also shifted, indicating that the effective threshold levels had changed. These damage effects will be due to the production of fixed charges in gate oxides on the readout chip, which will alter the behaviour of some of the transistors.

Figure 5.11 shows the change in the THL adjust values in each pixel on 3D1 caused by the damage. A broad band of pixels across the centre of the detector, where the beam was most intense, show large and highly variable changes in THL adjust, whereas outwith this region the changes are small. The global threshold level had changed by +3.77 on the 3D detector, and +7.77 on the planar detector. This change in threshold was taken into account when selecting the threshold levels in experiments following the damage.

The spectral tests in section 5.5.2 were done at 15keV, 12keV and then 20keV. Although the number of noisy pixels increased slightly during the lower-energy tests, the biggest increase was seen at 20keV. Additionally, the spectral results indicate that the shift in the pixel thresholds occurred during the 20keV test, implying that most of the damage occurred at this energy. This would make sense, given that the flux was higher during this test (more photons were able to pass through the attenuator) and more photons would have passed through the silicon substrate to reach the Medipix2 chip underneath.

Nevertheless, Medipix2 is designed to be used with medical X-rays, which typically operate at higher energies, so the level of damage seen here was unexpected. It’s perhaps possible that the detectors might have been accidentally exposed to the direct X-ray beam when changing attenuators.

5.5 Tests on Diamond beamline B16

A series of tests on the 3D detectors were performed at Diamond Light Source. The tests were done on beamline B16, which is designed to test new detectors, experimental techniques and optics [43]. The beamline can provide white or monochromatic beam; during the tests, only the monochromatic beam was used, in the range 12-20keV. Since the beam was intense, highly monochromatic, and extremely direc-

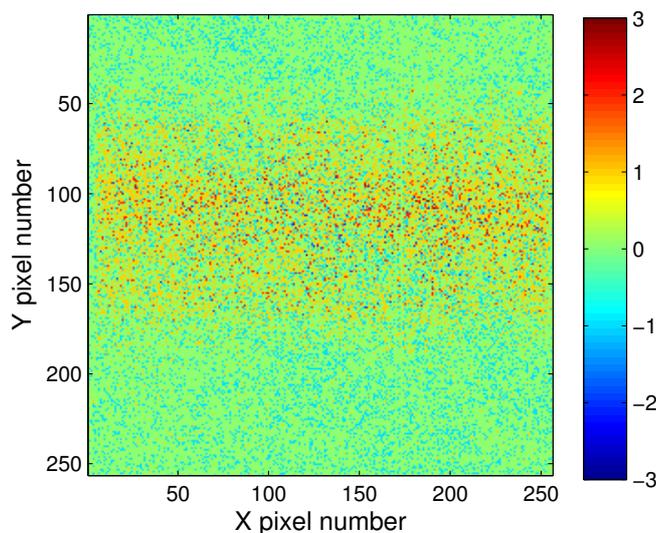


Figure 5.11: Change in the equalised THL adjust values for detector 3D1 due to beam damage.

tional, it could be used for tests that weren't possible with the X-ray tube or lab sources.

5.5.1 Preparation for tests

The detector 3D1 was tested on the beamline, along with the $300\mu\text{m}$ -thick p-on-n planar detector for comparison. Throughout the tests the 3D detector was biased to 21.5V and the planar detector was biased to 100V, to ensure that both were fully depleted.

The Medipix2 detectors were each mounted in an aluminium test box with a $50\mu\text{m}$ -thick aluminized Mylar window in front of the detector active area. During the tests, they were mounted on a remotely-controlled XY stage as shown in Fig. 5.12, allowing them to be moved in and out of the beam without having to re-open the beam area. The two USB interfaces controlling the detectors were connected to a control PC outside the beam hutch, using Cat 5 cables and sets of USB to Ethernet converters.

As described previously, threshold equalisation was performed on each detector before testing. Only the low threshold THL was used during the tests.

When the beam was switched on, the beam profile was found to be non-uniform. So, motorised tungsten slits were used to give a square beam spot, slightly smaller than the area of a detector. This ensured that the same area of the beam was viewed when switching between the two detectors.

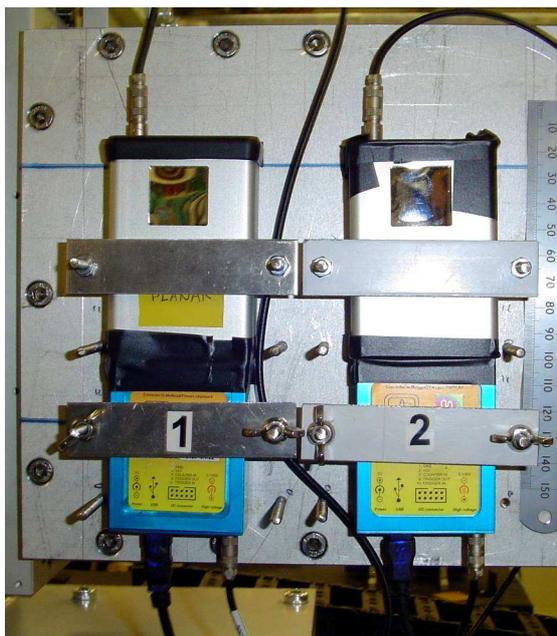


Figure 5.12: Planar (left) and 3D1 (right) detectors mounted in test boxes and equipped with USB interfaces, on an XY stage in the B16 beamline at Diamond Light Source.

5.5.2 Beam spectra measurements and calibration

After setting the X-ray beam to 15keV, the response spectrum of each detector was tested. This was done by taking images at a series of consecutive THL values ranging from above the beam energy down to the noise level. An acquisition time of 0.5s was used at each step—this was chosen to ensure a high count in each pixel, without reaching the Medipix2 chip’s maximum count.

Then, the total number of counts was calculated for each image (excluding noisy pixels), producing an integral spectrum of counts vs threshold setting. This was then differentiated to find the differential spectrum measured by each detector. This spectral measurement was repeated a total of 10 times, and the mean spectrum found. The standard error on each point in the spectrum was also calculated from the spread of the 10 results.

At this point, the spectrum simply gave the signal at each THL setting. In order to calibrate the detector, the tests were repeated at 12keV and 20keV. The peak in each spectrum was found by using a Gaussian fit. The THL value corresponding to the centre of the channel noise, obtained from the threshold equalisations, was also taken as the “zero signal” level.

A linear least-squares fit was then made to these data points, in order to find

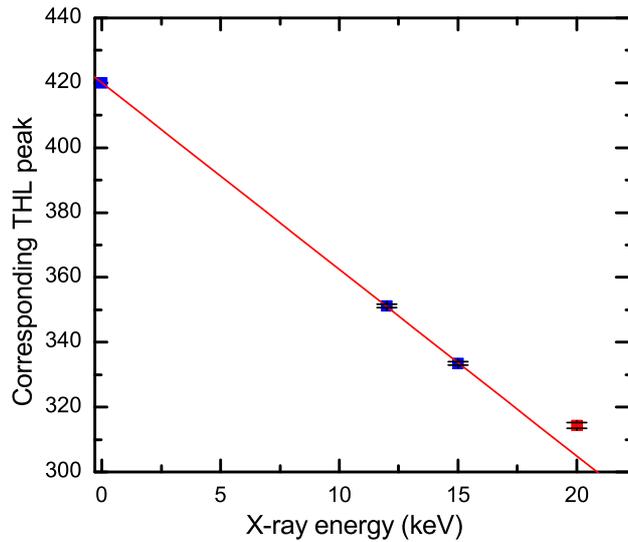


Figure 5.13: Calibration plot for the 3D1 sensor, using the spectral peak results. The linear fit ignores the data point at 20keV, as discussed in the text.

the calibration for each detector. However, it was found that for both detectors, the fit was greatly improved by excluding the measurements taken at 20keV. In Fig. 5.13, the calibration graph for the 3D1 sensor is shown, along with the linear fit excluding the 20keV results. The errors in the THL value were found by applying the peak fitting algorithm to each separate test run at a given energy, then finding the standard deviation of the results. It can be seen that the fit matches the 0, 12 and 15keV results to within their error values, which are small, whereas the 20keV results are well above the fit. In contrast, if the 20keV results are included, none of the points match the fit to within the calculated error. As discussed in section 5.4.6, it was found that radiation damage occurred during the tests. The anomalous results for 20keV can be explained by the shift in the pixel thresholds caused by this radiation damage. So, the calibrations for the two detectors used the fit without the 20keV results.

In section 5.4.3, it was shown that the zero signal THL values on the two detectors differed by about 30. (This difference in zero level is often seen on Medipix2 [116].) The calibration results showed that a step in threshold of 1 THL corresponded to a change in threshold energy of 0.174keV on the 3D1 sensor, and 0.180keV on the planar sensor. So, the energy response of the two Medipix2 hybrids is very similar, as would be expected. During the analysis of later tests—done after the radiation damage—it was assumed that this energy scaling from keV to THL remained the

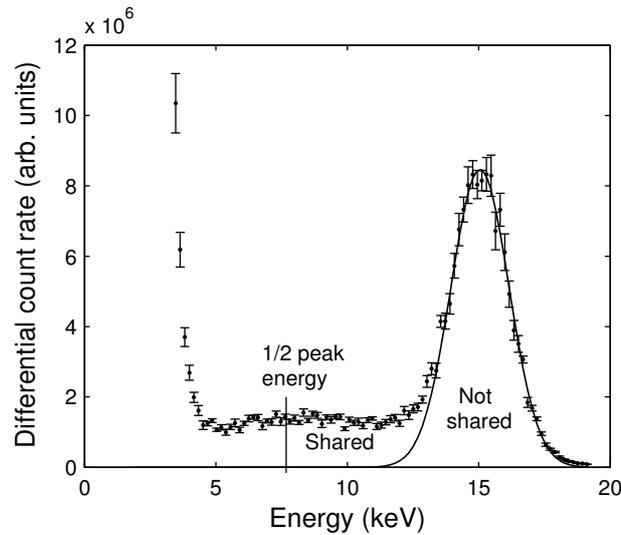


Figure 5.14: Spectrum from the 15keV beam measured by the 3D detector. The counts corresponding to the peak energy deposition have been fitted with a Gaussian. The point at half the peak energy has been indicated.

same, and the “zero signal” level was taken from threshold equalisation tests.

The detector calibrations were applied to the spectra results. Figure 5.14 shows the spectrum obtained from the 3D detector with the 15keV beam, including the Gaussian peak fit and the standard error on each data point. In addition to the peak, lower energy “hits” can be seen, primarily due to some X-rays being charge-shared between pixels. (In the 3D detector, some lower-energy hits may also occur if some of the charge generated by an X-ray falls within an electrode column.) Before fitting the Gaussian to the peak, the steepest point on the low-energy edge was found, and the fit was only applied to the data at higher energies. This prevented the charge-shared signal from affecting the fit.

Figure 5.15 shows the 15keV results from the 3D and planar detectors superimposed, with a Gaussian fit to each peak. It can clearly be seen that the planar detector has a smaller peak and a greater charge shared signal, giving a poorer spectrum as a result.

The onset of false counts due to noise occurs sooner on the 3D detector than on the planar, as would be expected from the results in section 5.4.6. However, in all of these tests the 3D and planar detectors give very similar peak widths. Since the beam is very monochromatic, the peak width reflects aspects of detector response such as electronic noise and threshold dispersion. So, the earlier onset of noise on the 3D detector might only be due to poor noise performance in a small fraction of the pixels.

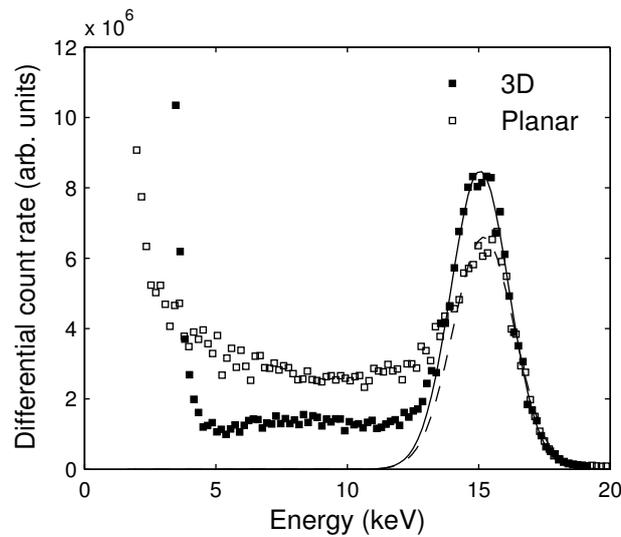


Figure 5.15: Comparison of 15keV beam spectra measured by 300 μ m 3D and planar detectors, showing greater charge sharing on the planar sensor.

To determine the total number of photons detected by each sensor, the count rate was found with the low threshold set to half the beam energy. This value will include as many of the charge-shared hits as possible, without any hits being counted twice in adjacent pixels. Only the small number of hits shared between 3 or 4 pixels will be missed. By integrating the total number of hits in the Gaussian fitted to the peak, the number of non-charge-shared hits was also found. So, the proportion of charge-shared hits on the detector could be found for each data set. Fig. 5.14 illustrates this.

For each beam energy, Table 5.1 shows the proportion of charge-shared events on each detector, and the relative number of hits seen on the two detectors at half the peak energy. To make the latter comparison, the total counts were adjusted to take into account small changes in the beam intensity as the electron current in the synchrotron decayed. Errors were estimated by looking at the variation in the number of hits in the peak when the Gaussian fit range was shifted by 0.4keV in either direction, and the variation in the number of hits at half the beam energy when this energy value was shifted by 0.2keV. The errors are larger on the planar detector, primarily because the higher charge sharing makes the peak-fit more variable. The mean results for the three tests are also shown.

The 3D detector has a substantially lower proportion of charge-shared hits; 24%, compared to 40% for the planar sensor. The results show no clear pattern with energy. There is greater variation in the results from the planar sensor, probably because the larger number of charge-shared events makes it harder to get an accurate

Table 5.1: Results from spectra measurements on planar and 3D detectors.

Energy (keV)	3D - % charge shared	Planar - % charge shared	(Counts 3D) / (Counts planar) at $E_{\text{beam}}/2$
12	26.1 ± 1.5	37.1 ± 3.0	0.856 ± 0.02
15	24.0 ± 1.5	45.0 ± 3.0	0.892 ± 0.02
20	22.0 ± 1.5	37.4 ± 3.0	0.832 ± 0.02
Mean	24.0	40.0	0.86

Gaussian fit to the peak.

Additionally, when the threshold is set to half the peak energy, the count rate on the 3D detector is only 86% compared to the planar detector. This suggests that the 3D detector has a smaller sensitive volume. This may at least partly be due to loss of signal from hits occurring in the electrode columns, which occupy about 5% of the device volume, or 10%, if we include the heavily-doped region around the columns. Variations in substrate thickness could also affect the results. The nominal substrate thickness is $300 \pm 15 \mu\text{m}$, which would correspond to a variation of $\pm 5\%$ in the active volume of each detector.

5.5.3 Measuring Line Spread Function using an edge

Next, the Line Spread Function (LSF) was measured for the two detectors, using a 12keV beam. The LSF is the response of the detector to an input signal consisting of an extremely narrow line. It is equivalent to the response of a single pixel when a narrow line is scanned across it. So, it quantifies how the detector “blurs” the pattern in the X-rays falling on the detector. The LSF was found by measuring the Edge Spread Function (ESF)—the response to a sharp edge—then differentiating it, using a similar method to Ref. [119].

Firstly, a thin lead edge on a glass slide was placed over each detector, at a shallow angle to the horizontal (about 2 degrees). Across any individual pixel, the edge is virtually horizontal, but the tilt is large enough to ensure that the distance from the edge to the centre of each pixel varies across the detector. This makes it possible to obtain the “oversampled” ESF and LSF with a step size much smaller than the pixel spacing itself.

Ideally, this experiment should be carried out using a perfectly uniform X-ray source. To compensate for the non-uniformity of the X-ray beam used here, images were taken both with the lead edge mounted on the detector and with it removed. During analysis, the count rate seen in each pixel in the edge image was divided by

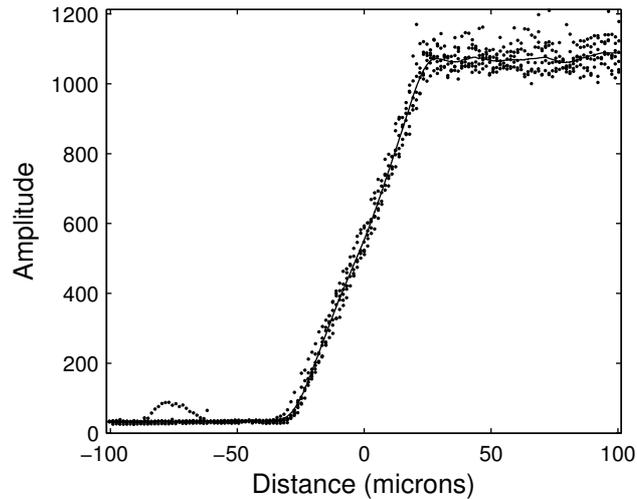


Figure 5.16: Edge Spread Function measured by the 3D1 detector at threshold setting THL385 (6.7keV) with 12keV beam, showing individual pixel values (dots) and the smoothed fit (solid line) against the distance from the edge.

the count rate without the edge in place. As well as helping to compensate for the beam's non-uniformity, this will also compensate for variations in the response of individual pixels (flat-field correction). These images were taken for both the 3D and planar detectors at a range of different threshold settings. Each measurement consisted of 100 images taken with a 2.5s acquisition time, which were then averaged.

Firstly, the position of the edge in the image was found. This was done by looking at each column of pixels in the image and finding the Y-position where the count rate was halfway between the maximum and minimum signal level, using interpolation to get sub-pixel precision. Then, a linear fit was made to the resulting data points.

Next, the perpendicular distance from the line to the centre of each pixel was calculated. From this, a scatter plot of the count rate against the distance from the edge was constructed—an example from the 3D detector is shown in Fig. 5.16. Pixels at the same distance from the line can have different values, due to any variations in pixel response or beam intensity still present after applying the correction. So, to obtain the Edge Spread Function from this data, smoothing was applied using a locally weighted linear regression (LOESS) fit. For each point in the original data, this applies a linear fit to the surrounding data points, with the weighting of each point falling rapidly with distance. The smoothing also assigns zero weight to points more than 6 standard deviations from the line, to reject outliers. As can be seen in Fig. 5.16, this gives a smoothly varying ESF.

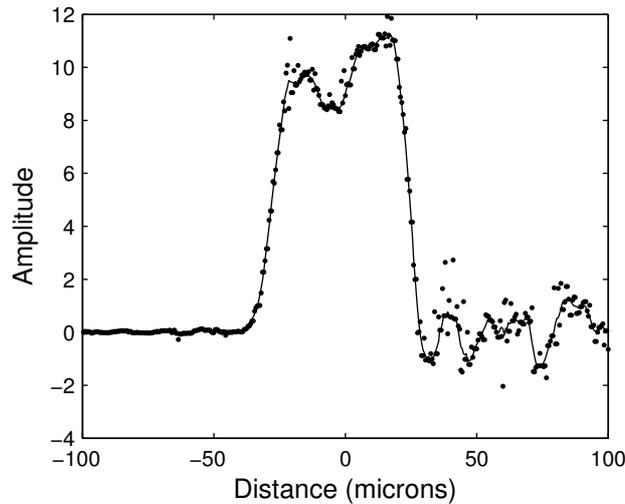


Figure 5.17: Line Spread Function measured by the 3D1 detector at threshold setting THL385 (6.7keV), before and after smoothing.

The ESF was interpolated, to give equally-spaced data points, then differentiated, to obtain the Line Spread Function. Once again, some LOESS smoothing had to be applied, because the differentiation enhanced small variations in the ESF data. An example of the LSF before and after smoothing is shown in Fig. 5.17. This shows the response of a pixel, with steep rising and falling edges and a plateau in the centre. Note that both the plateau and the region to the right show some undulation. This is most likely due to variations in the beam profile which could not be fully corrected. Unfortunately, this prevents us from seeing if there is any fine structure in the “plateau” region due to the readout column. However, the rising and falling edges of the LSF are a good source of information.

The LSF was found for the two detectors at a range of threshold settings. Some results from the planar sensor are shown in Fig. 5.18. Each data set from the detectors gave a LSF with steep rising and falling edges, and a plateau, but the width of the LSF tended to increase as the threshold setting was lowered. Consider the response of the detector system to monochromatic X-rays. The signal amplitude produced by a pixel of a photodiode detector will vary smoothly with the X-ray hit position as shown in Fig. 5.19. The steepness of the rising and falling edges will depend on the amount of charge sharing. When the signal is passed to the Medipix2 chip, it will either register a hit or not, depending on the signal amplitude and the threshold setting. So, the LSF obtained from the Medipix2 with a monochromatic beam will tend to have sharp edges, regardless of charge sharing. However, the width of the LSF will vary with the threshold setting, and this variation will be

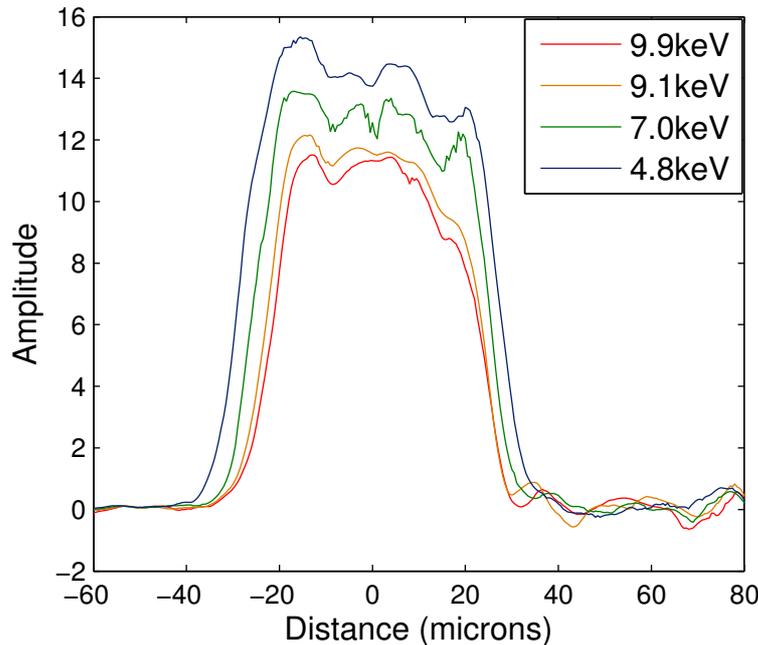


Figure 5.18: Line Spread Functions seen on the planar detector with different threshold settings, using a 12keV beam.

greater if there is high charge sharing.

So, for each data set, the full-width-half-maximum (FWHM) of the LSF was found. As an estimate of the error in this measurement, the widths of the LSF at 40% and 60% of the plateau level were also calculated. The resulting full width half maxima, with these errors, are plotted against the threshold setting in Fig. 5.20.

The FWHM of the 3D detector shows very little variation with the threshold setting—in the range of thresholds tested here, it only varies by $1.5\mu\text{m}$, with a value of $52.5\mu\text{m}$ when the threshold is set to half the 12keV beam energy. In contrast, the planar detector’s response varies by $15\mu\text{m}$ across the threshold range, with $53.5\mu\text{m}$ width at half the beam energy. So, this shows that the 3D detector has much lower charge sharing at the edges of the pixel. On both detectors, the width at half the beam energy is a little narrower than the $55\mu\text{m}$ pixel size, which is reasonable. Since charge sharing will be greater at the corners of the pixel, this will make the 1-dimensional LSF profile narrower than the pixel width.

5.5.4 Silicon powder diffraction experiment

After making these direct measurements of the two detectors’ performance, the detectors were then used in an experiment to measure powder diffraction patterns.

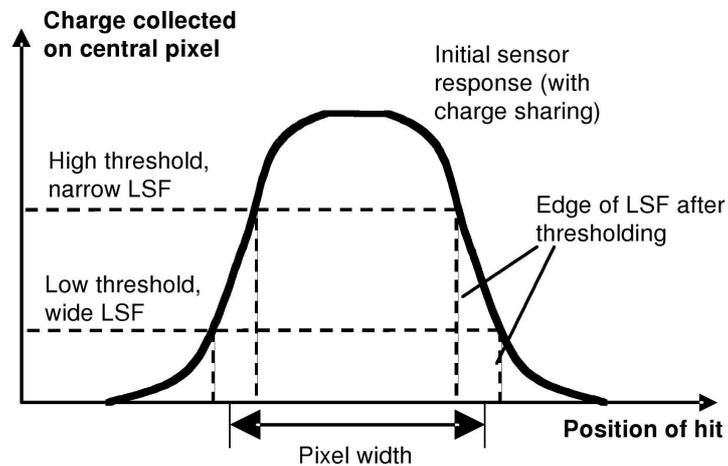


Figure 5.19: Diagram showing how charge sharing and varying thresholds affect the LSF from a single-photon-counting chip.

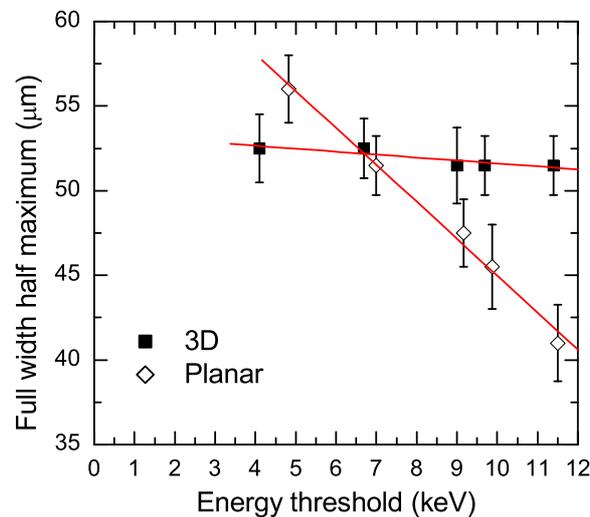


Figure 5.20: Variation in the full width half maxima of LSFs measured at different threshold settings using a 12keV beam, for both 3D and planar detectors. The planar detector has greater variation, indicating higher charge sharing.

As described in section 5.1, when a narrow, monochromatic X-ray beam is fired at a single crystal, the X-ray photons will scatter from the crystal planes, producing a series of diffraction spots at angles predicted by Bragg's Law (equation 5.1).

Now, suppose we have a fine powder sample, consisting of a large number of small, randomly oriented crystals. For each set of planes, there will be some crystals with the correct orientation with respect to the beam to satisfy Bragg's Law. In fact, because of the rotational symmetry of the setup, each set of planes will produce a cone of X-rays, composed of a large number of individual spots. If an X-ray detector is placed behind the sample, we will see a series of rings, and by measuring their angle and intensity we can gain information about the crystal structure [100].

On the beamline, a capillary tube containing a reference sample of powdered silicon (SRM 640c) was placed in a focused 15keV beam, producing a series of diffraction rings. The sample was 24cm from the detector stage. Since the diffraction rings covered a large angular range, and each Medipix2 image covered only 14mm by 14mm, a series of images were taken with the detectors while moving the stage in 10mm increments along the X-direction. If a ring was seen on a detector, 10 images were taken (each with a 10s shutter time) and averaged; if not, only one image was taken. During these measurements, the detection threshold on each detector was set to 7.5keV, half the beam energy.

An example image, showing the innermost diffraction ring as seen by the 3D detector, is shown in Fig. 5.21. Noisy pixels, identified by taking images without the beam present (see section 5.4.6), were replaced by the median count rate from the neighbouring pixels. The detector is only large enough to image a segment of the ring, but the curvature is still visible. Also, it can clearly be seen that the ring is composed of a series of diffraction spots, produced by individual crystals in the powdered silicon sample. To get the response of the detector across the full angular range, the images were cropped and stitched together appropriately. The result for the 3D detector is shown in Fig. 5.22, with the left-hand side of the image having the smallest angle relative to the incoming beam. Seven diffraction rings can be seen, along with a background which decreases as the diffraction angle increases.

Next, a projection was made of the mean number of counts per pixel versus the distance from the centre of the diffraction pattern. This was found by making a circular fit to the innermost diffraction ring, calculating the distance of each pixel from the centre of this ring, and binning the results in steps of half a pixel width. This "distance" variable was then used to find the corresponding angle of the diffracted beam, relative to the incoming beam (i.e. 2θ).

The diffraction results for the 3D detector are shown in Fig. 5.23. Using NIST

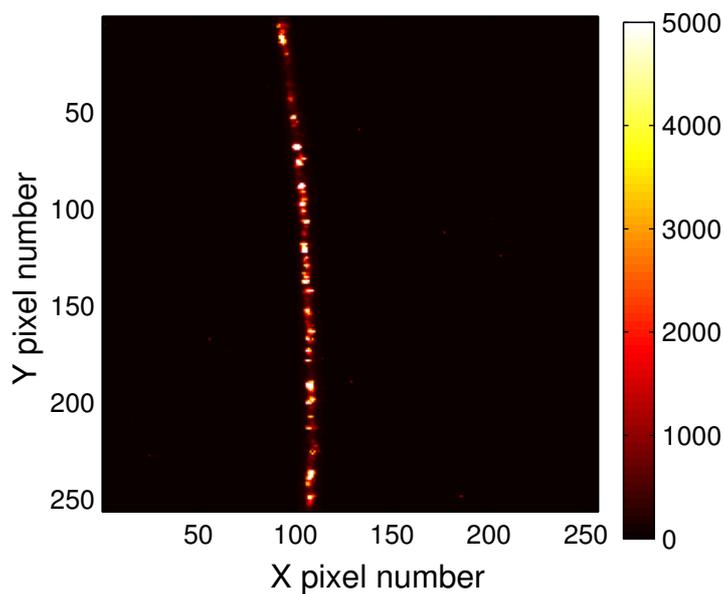


Figure 5.21: Image of the innermost diffraction ring from a powdered silicon sample, taken with the 3D detector. The colour scale shows the counts per pixel.

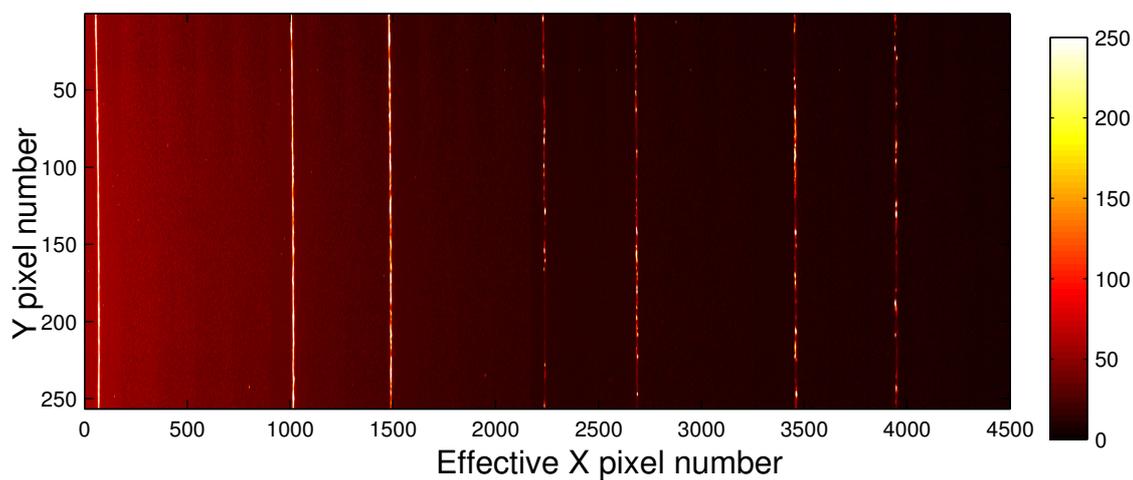


Figure 5.22: Image showing 7 diffraction rings from a powdered silicon sample, produced by combining multiple images taken by the 3D sensor. The colour scale uses a limited range, to make the weaker rings more visible. Note that the X and Y axes are scaled differently.

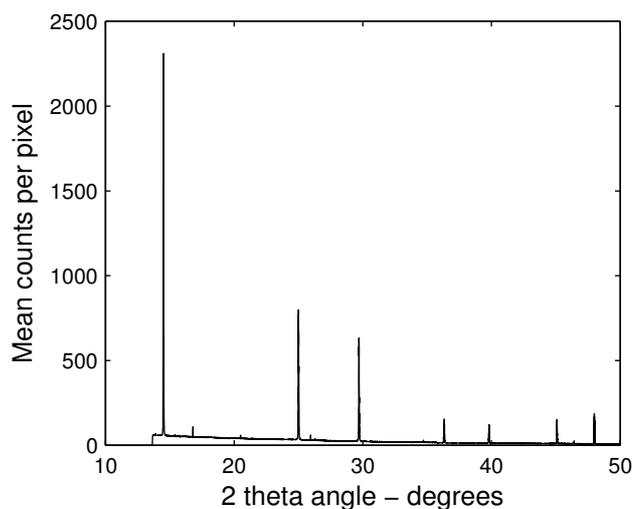


Figure 5.23: Powder diffraction ring intensity from a silicon sample vs 2θ angle, obtained from images taken by 3D sensor.

Table 5.2: Calculated and measured silicon powder diffraction angles.

h k l	d-spacing (Å)	Predicted 2θ (degrees)	Experimental 2θ (degrees)
1 1 1	3.1355	15.16	14.51
2 2 0	1.9201	24.87	25.00
3 1 1	1.6374	29.26	29.71
4 0 0	1.3577	35.46	36.32
3 3 1	1.2459	38.77	39.81
4 2 2	1.1085	43.80	45.08
3 3 3	1.0451	46.61	47.99

data [120], the positions of the first 7 expected diffraction peaks from silicon were calculated. The results, and the measured diffraction ring angles, are shown in Table 5.2. In the table, the planes causing the reflections are referred to by their (h k l) indices. There is fairly good agreement between the angles, showing that all of the reflections have been found successfully. The small differences between the angles, which increase steadily as the diffraction angle increases, will be due to inaccuracies in the measured position of the detector on the XY stage relative to the sample.

Fig. 5.24 shows the first peak in the projection, expressed in terms of the distance measured on the detector rather than the angle. The full-width half maximum of this peak is 3.5 pixels. Given the relatively large width of the peaks, when the experiment was repeated using the planar sensor there was no visible difference between the peak

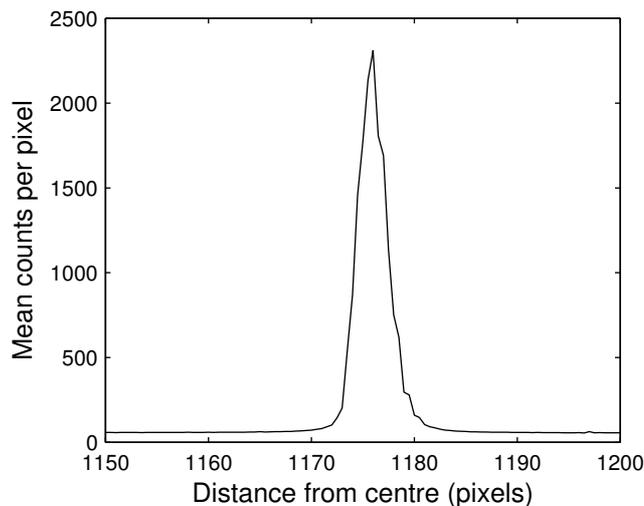


Figure 5.24: Profile of the first diffraction ring obtained from images taken by 3D sensor.

shapes seen on the two sensors. The peak width is affected by both spot size and sample size, which were comparatively large in this test experiment, so if a more focused beam and a smaller capillary of powdered silicon were used then a better comparison between the detectors may be possible.

5.6 Alpha particle tests

When an alpha particle hits a silicon detector, it will generate high concentrations of electrons and holes in a small volume near the surface [121]. Because electrons have higher mobility than holes, they will diffuse outward more quickly, and the charge cloud will become surrounded by a layer of electrons. Due to their high concentration, these electrons will screen out the external electric field, giving carriers within the cloud more time to diffuse. This “plasma effect”, and the high total number of carriers, will mean that the alpha produces a large, circular cluster on the detector.

Both the 3D1 and planar detectors were tested using an ^{241}Am source, which emits 5.637 MeV alphas. The source was placed 5mm from the back side of the sensor, to minimise the energy loss of the alphas in air. (Since the front surface of the sensor is bump-bonded to the Medipix2 chip, only the response from the back side could be tested.) The 3D detector was biased to 20V, and the planar to 80V. Two hundred images were taken with each of the two sensors, using a short acquisition time to minimise the number of overlapping alpha clusters. Figure 5.25 shows part of one of these images, containing clusters of different sizes. The number

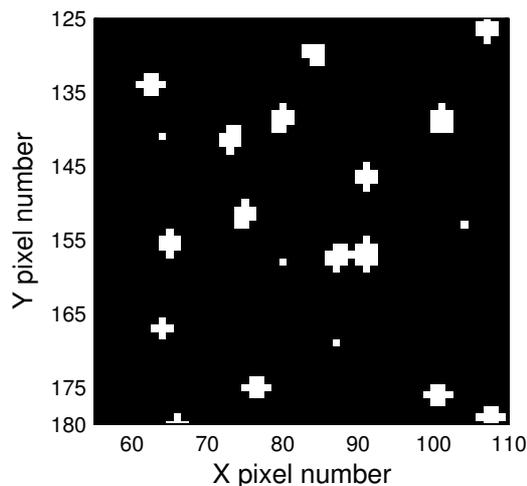


Figure 5.25: Part of an image of alpha clusters, taken on the 3D detector with the threshold at 10keV. At this low threshold, single-pixel hits due to gammas can also be seen.

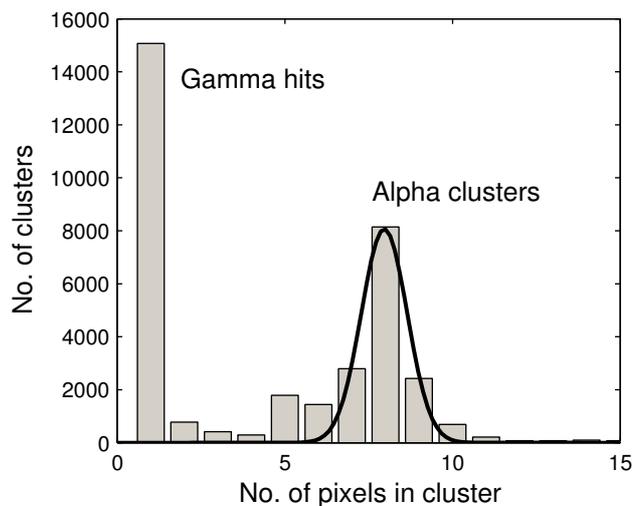


Figure 5.26: Histogram of alpha cluster sizes seen on the 3D detector, with the threshold set to 10keV. A Gaussian fit has been applied to the alpha cluster distribution, to find the typical cluster size.

of pixels in each cluster was found, and a histogram of cluster sizes was produced as shown in Fig. 5.26. Then, a Gaussian fit was applied, and the peak value was taken as the typical cluster size. This process was repeated using a range of threshold settings. Since ^{241}Am also emits gammas (at 13.9, 26.3 and 59.5keV), one- and two-pixel gamma hits were seen at lower threshold settings. These were excluded from the Gaussian fit.

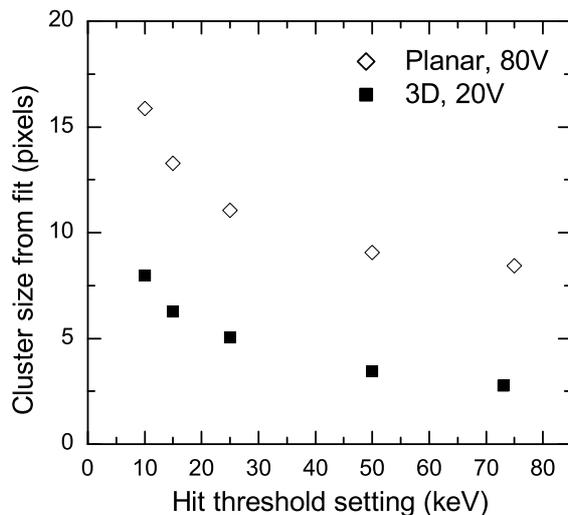


Figure 5.27: Variation in alpha cluster size with threshold setting for the 3D and planar detectors. The smaller clusters on the 3D detector show that it has lower charge sharing.

Figure 5.27 shows the resulting cluster sizes on the two detectors, at a range of threshold settings. At each threshold setting, the 3D detector shows much smaller cluster sizes, demonstrating its reduced charge sharing. Because the alpha particle is deposited at the back surface of the detector, and the planar detector collects the charge at the front surface, the difference in collection distance between the 3D and planar detectors will be particularly dramatic in this case. Also, in addition to the 3D detector's fast collection time and the self-shielding effect of its electric field pattern, the 3D electrodes could perhaps disrupt the plasma effects occurring in the charge cloud. For example, once the charge cloud reaches an electrode, the carriers at the outer edge of the cloud will be quickly removed.

Both detectors show a fairly similar drop in cluster size as the threshold is increased, because using a high threshold will exclude the pixels at the outer edges of the charge cloud.

5.7 Conclusions

Pixel detectors for X-ray diffraction experiments have various requirements such as a large, linear dynamic range, fast readout, low noise and high resolution. Hybrid pixel detectors, where a silicon sensor is bump-bonded to a readout chip, can offer various advantages over other technologies such as CCDs. In particular, a single-

photon-counting chip such as Medipix2 can give a very low noise, a good dynamic range (provided the X-ray beam is relatively continuous) and the possibility of energy discrimination. Using a 3D detector instead of a planar silicon sensor could provide the additional advantages of reduced charge sharing, which improves image quality and energy resolution, and potentially active edges.

Three 3D detectors have been successfully bump-bonded to Medipix2 readout chips at VTT. All three work successfully, though two of the sensors have some dead pixels along the edge. This is possibly a side-effect of wafer bending during fabrication, though damage during saw-cutting could have the same effect. If the saw-cutting is to blame, then CNM could correct this fairly easily. If not, it will be important to experiment with fabrication conditions to reduce this problem.

Initial tests of collection efficiency against bias show that these detectors reach lateral depletion around 2V, at which point most of the detector is depleted, and full depletion around 9V. This is in agreement with simulations and previous C-V tests.

The 3D detectors were tested with monochromatic X-rays at Diamond Light Source. Measurements of their spectral response show that they give substantially less charge sharing than an equivalent planar detector—24%, compared to 40%. Likewise, the width of the 3D detector's Line Spread Function of the 3D detector varies by only $1.5\mu\text{m}$ when the threshold setting is adjusted from just below the beam energy to half the beam energy. In contrast, the variation on the planar detector is $12.5\mu\text{m}$. This also indicates low charge sharing on the 3D sensor. Similarly, the 3D sensor also produces much smaller cluster sizes when hit by alpha particles.

Currently, more Medipix2 detectors are available from CNM's new fabrication run. Glasgow and Diamond plan to test these using a microfocused X-ray beam with a spot size of a few microns. This will be scanned across a pixel to directly measure the charge sharing at its edge and the charge loss within the readout column. If possible, it would also be useful to use a uniform X-ray beam and an edge or slit to measure the detective quantum efficiency [114], which quantifies how the signal-to-noise ratio of the detector varies with the spatial frequency of the input signal. This would make it possible to directly compare the image quality from 3D detectors to that obtained from other technologies.

Chapter 6

Single-type-column 3D detectors from FBK-IRST

Another variety of 3D detector is the “single-type column” 3D structure [122] developed by FBK-IRST (Fondazione Bruno Kessler, Il centro per la Ricerca Scientifica e Tecnologica) in Trento, Italy. This has a simpler structure with only one set of electrode columns, which are etched from the front surface and do not pass through the full substrate thickness. The bias contact is provided by a planar implant on the back surface. This structure was partly intended as a test of the 3D fabrication process. However, it still offers some of the benefits of the standard 3D structure.

The simulations and tests of this detector structure are broadly similar to those done with the double-sided 3D structure back in Chapter 4. These devices were actually fabricated and tested before the double-sided detectors. However, the single-type-column structure differs more radically from the original full-3D design, and is also less useful, so the chapter ordering was chosen to reflect this.

Section 6.1 describes the single-type-column structure and its fabrication process, followed with details of the devices that FBK-IRST provided to Glasgow. Then, section 6.2 contains a series of simulations of this structure, with a particular focus on how its carrier collection behaviour differs from full-3D and double-sided-3D detectors. The experimental results are shown in section 6.3. As well as basic electrical characterisation, and charge collection tests on pad detectors with betas, some results from a CERN beam test on a strip detector are also presented.

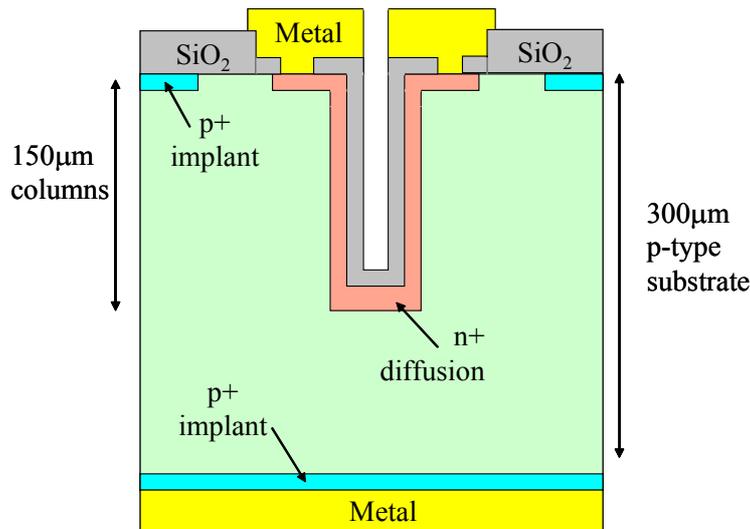


Figure 6.1: Structure of a single-type-column 3D detector with n-type readout and a p-type substrate. After etching, the columns have been doped by diffusion, and passivated with silicon dioxide. The surface around the column has also been doped, to provide a better electrical contact to the metal readout pad.

6.1 Single-type column 3D (3D-STC) detectors

6.1.1 Structure and fabrication

The structure of a single-type-column 3D detector is shown in Fig. 6.1. A set of readout columns, with opposite doping type to the substrate, are etched from the front surface of the detector partway into the substrate. On the back surface of the detector a simple planar implant, with the same doping type as the substrate, forms the bias contact.

The fabrication process is described fully in Ref. [123]. Once again, the process is broadly similar to that used to produce full-3D and double-sided 3D devices. The main advantages of the process are that only one set of columns needs to be produced, and that no double-sided alignment is required since we only use a simple planar contact on the back surface. Most of the benefits of the double-sided 3D structure described in section 4.1.1 also apply to single-type-column 3D; for example, no support wafer is needed.

Once again, the first step in the column fabrication is to etch holes into the silicon using ICP etching. In FBK-IRST's first set of devices, the column etching was done at CNM-IMB. Although it would be possible to completely or partially fill the etched columns with polysilicon, as with the other 3D devices discussed in this

thesis, FBK-IRST's fabrication process does not use any polysilicon filling. Instead, the columns are doped by diffusion straight after the etching step, and then a layer of silicon dioxide is deposited inside the columns to passivate them. This simplifies the fabrication process. However, this approach may increase the effective column radius, since enough dopant must be diffused into the silicon to produce a layer with high conductivity (which would normally be provided by the layer of doped polysilicon deposited inside the column). During the doping step, a region at the front surface around each column is also doped, to provide a good electrical contact to the metal readout pad.

When these devices were produced, CNM-IMB hadn't yet optimised their etching process. So, these devices have relatively short columns, with a depth of $150\mu\text{m}$ and $10\mu\text{m}$ diameter. All the devices have n-type readout columns, a p-type substrate, and a p-type back contact.

6.1.2 Devices produced by FBK-IRST

FBK-IRST have provided Glasgow, and other institutions such as the University of Freiburg, with a variety of 3D single-type column (3D-STC) pad and strip detectors. These detectors have been fabricated on two different types of p-type wafer; $300\mu\text{m}$ -thick Magnetic Czochralski (MCz), and $500\mu\text{m}$ -thick Float Zone (FZ). Section 2.3.3 explains the difference between these two substrate types. Most of the tests and simulations focus on the $300\mu\text{m}$ MCz devices, since in the $500\mu\text{m}$ FZ devices the column length is small compared to the device thickness and the behaviour is expected to resemble a planar detector. The resistivity of the FZ wafer is quoted as being $5\text{k}\Omega\cdot\text{cm}$, which would correspond to a doping of $2.8\times 10^{12}\text{cm}^{-3}$. The resistivity of the MCz is quoted as being higher than $1.8\text{k}\Omega\cdot\text{cm}$, which means that the doping is less than $7.7\times 10^{12}\text{cm}^{-3}$ [123].

There are five varieties of 3D-STC pad detector. All of the pad detectors have a 10×10 matrix of connected n-type readout columns, and all but one have a column spacing of $80\mu\text{m}$. Around this array, an additional set of linked columns form a guard ring. An image of one detector is shown in Fig. 6.2. All of these detectors have strips of p-stop outside the guard ring, and all but one have p-stops between the guard ring and the array of columns. The distinguishing features of each detector are listed below:

- *3D1*: A ring of p-stop isolation surrounds each individual column. The n-type readout columns are connected by metal tracks. Unlike the other detectors, there is no p-stop between the guard ring and the pad array.

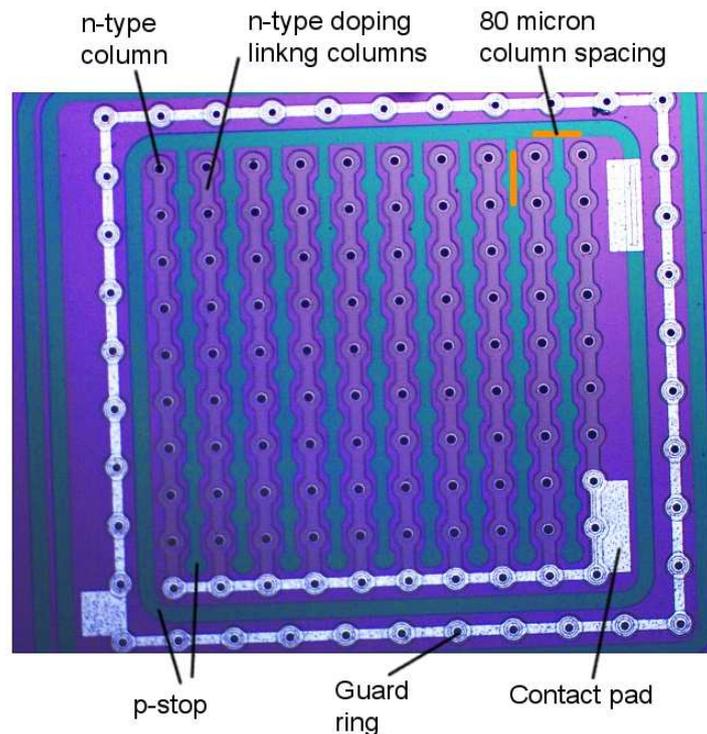


Figure 6.2: Image of single-type-column 3D pad detector, with a 10×10 matrix of connected n-type readout columns, surrounded by a guard ring. In this device (3D3) there are p-stops between the column matrix and the guard ring, and around the outside of the guard ring, to protect the pad against surface currents. There are also strips of p-stop between the columns within the matrix.

- *3D2*: There is no isolation between the columns within the pad detector. The n-type readout columns are connected by n-type implants.
- *3D3*: Shown in Fig 6.2. $20\mu\text{m}$ -wide strips of p-stop run between the n-type columns, like in a strip detector. The n-type readout columns are connected by n-type implants.
- *3D4*: Same as 3D2, but the electrode spacing is $100\mu\text{m}$, rather than $80\mu\text{m}$.
- *3D5*: Same as 3D3, but the p-stops are only $15\mu\text{m}$ wide.

Additionally, FBK-IRST fabricated planar pad detectors on the same wafers as the 3D-STC devices, as a basic test of the fabrication process. These planar pads are circular, with a relatively large diode area of 4mm^2 plus a guard ring.

The pad detectors have been electrically characterised, and their collection efficiency has been tested with the pad detector MIP setup described in section 3.3.

The 3D-STC strip detectors have a similar layout to the 3D-STC pad detectors. Each strip consists of a series of n-type readout columns connected by n-type doping and a metal layer. The spacing between columns along the strip is $100\mu\text{m}$, and the strips are 18.4mm long. Each detector has 64 strips, with $80\mu\text{m}$ pitch between the strips. At the edge of each detector there are guard rings and p-stops to prevent surface current reaching the strips. Since the readout columns are n-type, the devices have electrode isolation. One strip detector has strips of p-stop between the readout strips, and the other uses a p-spray.

Most of the tests on these strip detectors were done at the University of Freiburg. Glasgow collaborated with Freiburg to test two strip detectors in a beam test at CERN, using the LHCb-based readout electronics described in section 3.4. Some results from this test are presented in section 6.3.3.

6.2 Simulation

6.2.1 Simulation methods and device structure

The behaviour of this alternative 3D structure was investigated using Synopsys TCAD, as described in Chapter 2. Since all the single-type-column 3D detectors produced by FBK-IRST used n-type readout and p-type substrates, only this type of device was simulated. However, a variety of different structures were used.

Firstly, different structures were used for the pad detectors, which had $80\mu\text{m}$ spacing between columns, and for the strip detectors, which had $100\mu\text{m}$ spacing along the strips and $80\mu\text{m}$ between them.

Most of the simulations used $150\mu\text{m}$ -long columns with $10\mu\text{m}$ diameter and a $300\mu\text{m}$ substrate, to match the MCz devices. However, additional simulations were done with $250\mu\text{m}$ columns, in order to investigate how the column length affects the device's behaviour, and to provide a fairer comparison to the other 3D detectors simulated previously. Although the real devices have hollow columns, doped on the inside, the simulated devices had doped silicon columns to simplify the mesh.

For many of the simulations, the “quarter pixel” device structure was used, as before. However, some of the simulations used a larger region of the device, either to illustrate the electric field behaviour more clearly, or to investigate the effect of the weighting field on the charge collection. In particular, it was found that the weighting field was dramatically different between 3D-STC pad and strip detectors, unlike in full 3D and double-sided 3D devices.

Unless otherwise specified, the simulations used a substrate doping of $2.8 \times 10^{12} \text{cm}^{-3}$, matching the upper limit of the FZ wafer's quoted concentration.

6.2.2 Depletion and electric field behaviour

As a starting point, the depletion behaviour of the pad detector structure was simulated. Figure 6.3 shows the structure of the pad detector, and the hole concentration throughout the device at a range of bias voltages. (Only the upper region of the device is shown.) The depletion region has dramatically lower hole concentration than the undepleted p-type bulk. The depletion region initially appears around the n-type column where it meets the p-type substrate. As the bias is increased, the depletion region grows laterally outwards. At around 7.5V, the depletion regions from neighbouring n-type columns meet, and the region between the columns becomes fully depleted. As the bias is increased further, the depletion region continues to grow towards the bias contact on the back surface, much like in a planar device. Since the $150 \mu\text{m}$ columns only pass through half the thickness of the substrate, this second phase of planar-like depletion will have a significant effect on the device behaviour.

Next, the electric field pattern was simulated. Here, the results from a strip detector with p-stop isolation are shown. (The electric field in a pad device is very similar.) The detector structure is shown in Fig. 6.4. Four adjacent n-type columns were simulated, rather than just one, so that the same mesh could be used to calculate the strip detector's weighting field.

Figure 6.5 shows the electric field in a vertical cross-section of the 3D-STC strip detector, which passes diagonally across the device mesh between two of the n-type readout columns. The device is biased to 100V, and is fully depleted. The black streamtrace lines with arrows show the direction of the field, though their spacing is not proportional to the field strength. Due to the asymmetrical structure, the drift behaviour of electrons and holes will be different. When electron-hole pairs are generated between the readout columns, the electrons will travel a short distance horizontally to reach the columns, producing a fast signal. The vertical component of the field will also cause the electrons to move towards the front surface of the device. In contrast, the holes need to drift vertically across the substrate thickness to reach the back surface of the device, meaning that the hole drift distance will tend to be greater. In the process, the field will tend to channel the holes to the midpoint between columns, where the field is weaker. The field also becomes weaker towards the front surface, which means that holes generated here will take a particularly

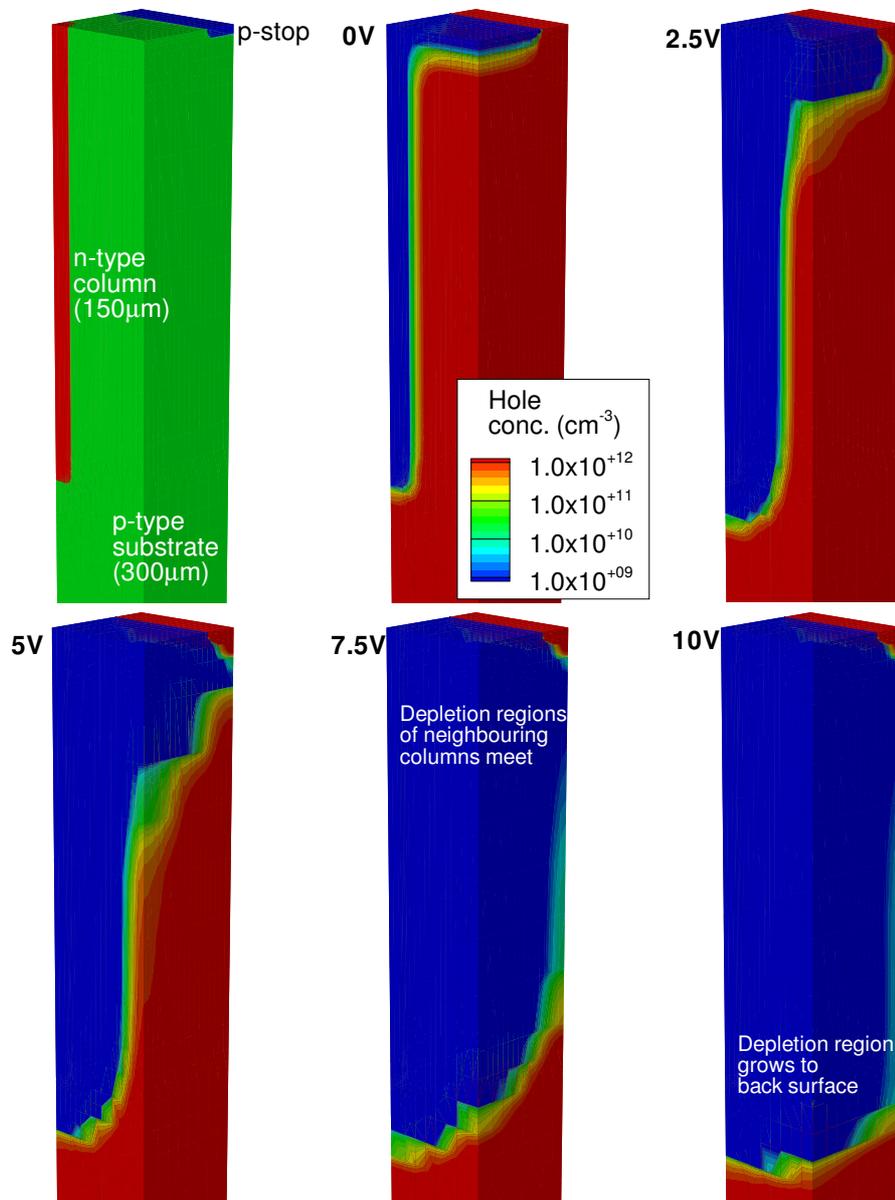


Figure 6.3: Depletion behaviour of a single-type-column 3D pad detector. The hole concentration is shown at 0V, 2.5V, 5V, 7.5V and 10V bias, and as the substrate depletes the concentration drops dramatically. The lower region of the device is not shown.

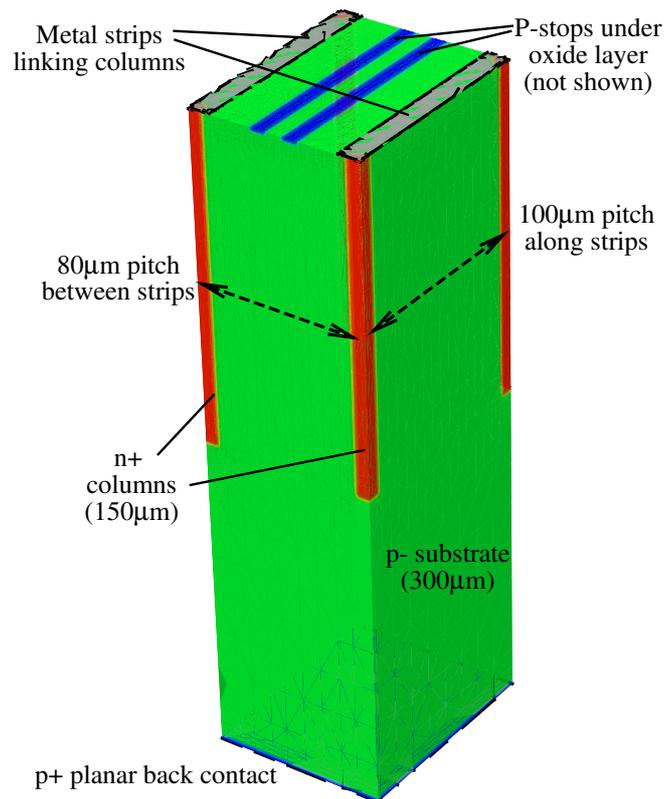


Figure 6.4: Structure of the simulated single-type-column 3D strip detector with p-stop isolation.

long time to cross the full substrate thickness. In the region beneath the columns, the field pattern and collection behaviour become more similar to a standard planar detector.

Fig. 6.6 shows the field when the bias is increased to 200V. It can be seen that increasing the bias increases the field in the region below the columns, but has relatively little effect in the region between the n-type columns. So, this low-field region is an inherent problem of the device structure. In a device with 250µm columns, these effects are exaggerated. The field becomes higher in the 50µm gap between the columns and the back surface, and the low-field region at the midpoint between the columns extends to a greater depth. So, increasing the column length won't necessarily improve the device's performance.

High field regions

It can clearly be seen from Fig. 6.6 that the highest-field regions occur around the tips of the n-type columns. Simulations were done to find how the maximum field within the 3D-STC structure varied with the applied bias. These simulations

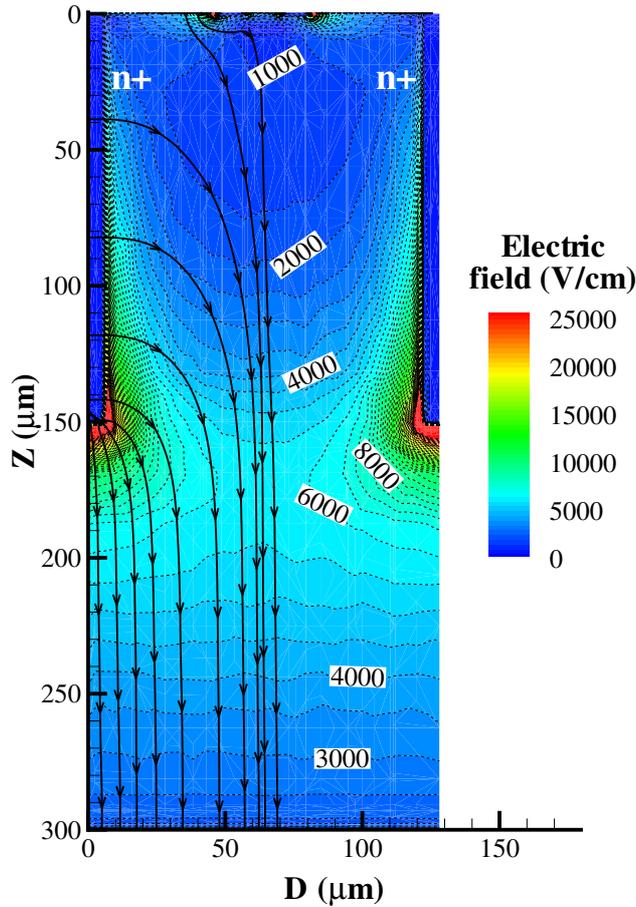


Figure 6.5: Electric field strength at 100V bias in a cross-section of a 3D-STC strip detector with p-stop isolation. The cross-section passes through two diagonally opposite n-type readout columns. The black streamtrace lines show the direction of the electric field.

used both the $150\mu\text{m}$ -column structure, and the $250\mu\text{m}$ -column structure. With $150\mu\text{m}$ -long columns, the maximum field reached $2.5 \times 10^5\text{V/cm}$ (a little below the breakdown field in silicon of $3 \times 10^5\text{V/cm}$) at 580V bias, which is much higher than the 170V obtained for full-3D and double-sided 3D detectors. However, with $250\mu\text{m}$ columns this fell to 245V. This is because the $250\mu\text{m}$ -column device has a smaller spacing between the column tip and the back contact, increasing the field.

6.2.3 Weighting field in pad and strip detectors

As shown back in section 2.2, the weighting field within each pixel in a full or double-sided 3D detector is relatively symmetric, meaning that electrons and holes make similar contributions to the total readout signal. The weighting field doesn't vary

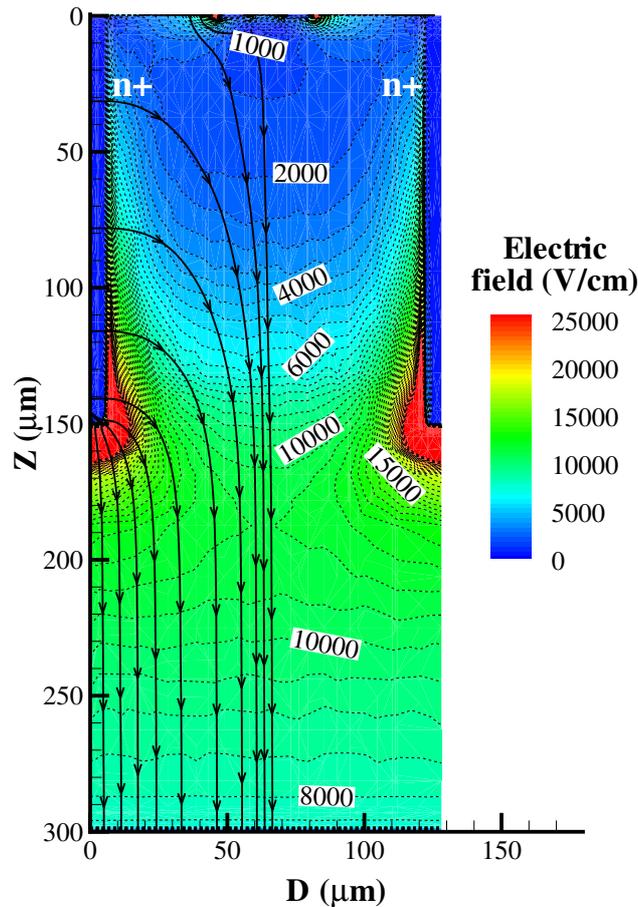


Figure 6.6: Electric field strength at 200V bias in a cross-section of a 3D-STC strip detector with p-stop isolation. The cross-section passes through two diagonally opposite n-type readout columns. The black streamtrace lines show the direction of the electric field. Compared to the simulation at 100V, the field strength has increased a lot in the region beneath the columns, but not between the columns.

much between pixel, strip and pad devices. Also, the weighting field of a particular readout electrode is generally low outside of the corresponding pixel cell, due to the presence of the bias electrodes—see Fig. 2.7.

When the weighting field of the fully-depleted 3D-STC detector was calculated, the results were substantially different, and varied a great deal depending on how the electrodes were connected. Figure 6.7 shows the weighting potential for one strip in the 3D-STC strip detector simulated in the previous section. Firstly, the weighting field of the readout strip extends well beyond the midpoint between the two strips, and has a relatively steep gradient even in the immediate vicinity of the neighbouring readout column. Suppose that an electron-hole pair is generated nearer to the neighbouring strip, say at position $D=85\mu\text{m}$, $Z=75\mu\text{m}$ as indicated by

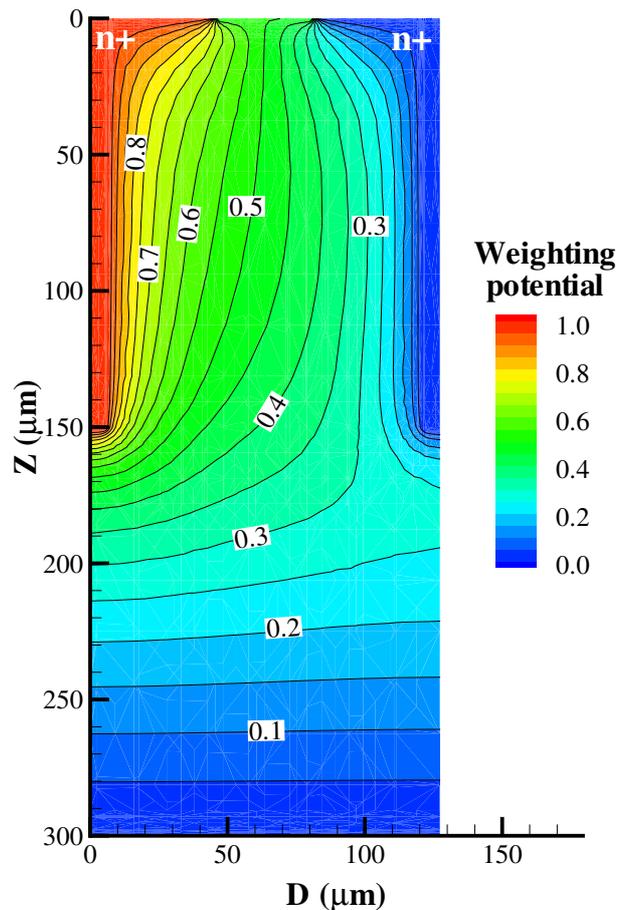


Figure 6.7: Weighting potential for one readout strip of a 3D-STC strip detector with p-stop isolation.

the axes. At this point the weighting potential is 0.4, so when the electron drifts to the right-hand readout electrode a fast negative-going signal of $-0.4e$ will be induced on the left-hand readout electrode. Although the hole drifting to the back surface will produce a signal of $+0.4e$ on the left-hand electrode, ultimately resulting in zero net charge, the difference in collection speed between electrons and holes will mean that the readout electronics will see a fast negative pulse followed by a much slower positive signal.

Figure 6.8 shows the recalculated weighting field when all the columns are connected together to emulate the behaviour of a pad detector. With this configuration, the entire region between the strips has a weighting potential ranging from just 0.8 to 1. This means that if an electron-hole pair is generated in this region, most of the resulting readout signal will be generated only when the hole reaches the region below the columns. So, the pad detector will produce a slower signal.

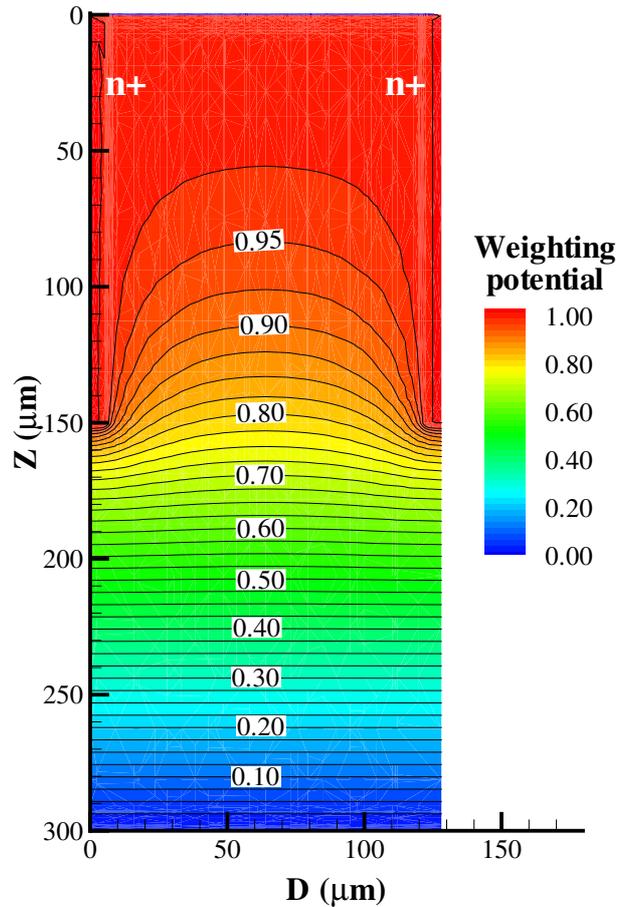


Figure 6.8: Weighting potential for a 3D-STC pad detector. This uses the same general structure as the strip detector simulated previously, but all the n-type columns are connected together.

6.2.4 Charge collection behaviour

The previous simulation results have shown that the 3D-STC detectors have distinctly different behaviour to full-3D and double-sided 3D detectors. Charge collection simulations were done to investigate this further. Since the previous section showed that the weighting field varied dramatically between pad and strip devices, both configurations were simulated. For each configuration, devices with $150\mu\text{m}$ and $250\mu\text{m}$ columns were used.

The pad detector simulation used the same mesh as the depletion simulation in Fig. 6.3. As with previous charge collection simulations, the entire device was flooded with a uniform concentration of excess electron-hole pairs, with the total number of carriers corresponding to the ionisation produced by a minimum ionising particle. The resulting readout current was simulated over a 200ns interval.

The strip simulation used a mesh similar to Fig. 6.4, except that the mesh was halved, leaving only one-type readout column per strip instead of two. This was done simply to speed up the simulation. When dealing with the strip detector, it was necessary to deposit charge only in the vicinity of one strip—if the entire simulated volume was flooded with uniform charge, then the result would be equivalent to using a pad detector structure. So, the region around one strip, extending from the centre of the strip (at one edge of the simulation mesh) to $5\mu\text{m}$ from the midpoint between the strips, was flooded with a uniform concentration of electron-hole pairs. The resulting readout current from both strips was simulated.

The fabricated 3D-STC devices used $150\mu\text{m}$ columns, but the other devices simulated in this thesis used $250\mu\text{m}$ columns. So, devices with both column lengths have been simulated to compare their behaviour.

Figure 6.9 shows the current signals obtained from these simulations, using a log-log scale. Initially, each 3D-STC detector has a fast signal pulse lasting for less than 1ns, much like the other 3D detectors simulated previously. The current pulse is larger in the devices with $250\mu\text{m}$ -long columns. However, after this fast pulse there is a long tail-off signal, decaying fairly smoothly for tens of nanoseconds, due to the slow drift of holes to the back surface. The tail-off signal is particularly long in the devices with $250\mu\text{m}$ columns, due to the low-field region between the columns. It can also be seen that the signals are faster in the devices which use a strip configuration, due to the more favourable weighting field.

The practical effects of this behaviour can be seen more clearly by looking at the total charge collected over time in each device, as shown in Fig. 6.10. For the strip detector simulations, the total charge collected on the neighbouring strip is also shown. The tail-off signal makes a large contribution to the total charge collection in each device. The longer tail-off signal in the $250\mu\text{m}$ -column devices means that, although they have a higher initial current pulse, they take over 50ns to approach 100% charge collection, whereas the $150\mu\text{m}$ -column devices achieve this in a few tens of nanoseconds. The behaviour of the $250\mu\text{m}$ -column pad detector is particularly poor. Additionally, in the strip detectors the neighbouring strips experience a fast negative current pulse as the electrons are collected, followed by a slow positive current from hole collection, as expected.

So, this device has relatively slow collection compared to a full or double-sided 3D detector. This will make it more prone to charge trapping following radiation damage. However, the 3D-STC structure still has the advantage of a low depletion voltage.

The radiation hardness of this structure was tested by running strip simulations,

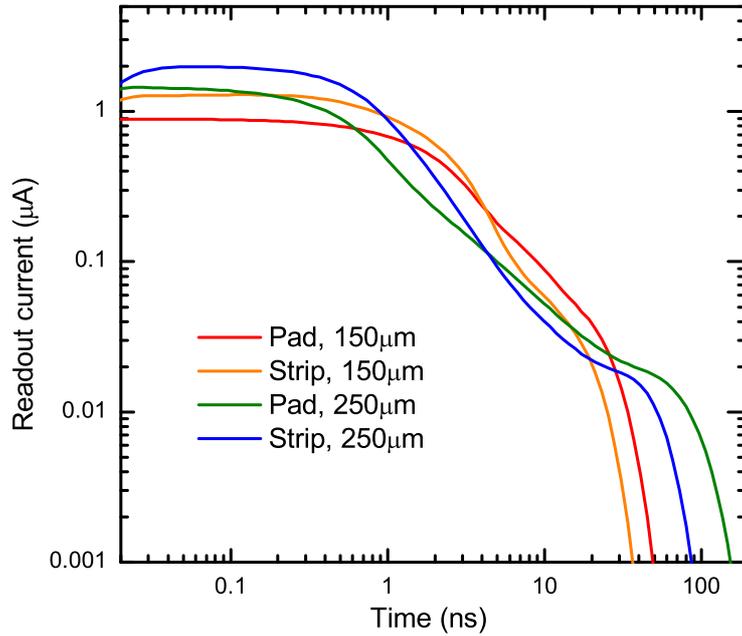


Figure 6.9: Simulated current signals in different 3D-STC devices. The devices use either a pad or strip structure, and either $150\mu\text{m}$ or $250\mu\text{m}$ -long columns.

using the structure and charge distribution described above, with different levels of radiation damage. Once again, the damage model from section 2.3 was used. The simulations were done with 200V bias, and the leakage current was subtracted from the results. At a damage level of $2 \times 10^{15} \text{ n}_{eq}/\text{cm}^2$, the charge signal was 8700 electrons, and at $5 \times 10^{15} \text{ n}_{eq}/\text{cm}^2$ the signal was 4700 electrons. In comparison, an n-on-p planar strip detector simulated at 900V gives 9000 electrons and 4000 electrons at these two fluences—see Fig. 2.13 earlier. So, the single-type-column detector can achieve the same collection signal as the planar detector at a much lower bias, but it doesn't match the extreme radiation hardness of a full or double-sided 3D sensor.

Additionally, the collection speed will affect the experimental results. The pad detector CCE setup uses an amplifier chain with a long peaking time of about $2\mu\text{s}$, so no ballistic deficit should occur. However, the strip detector CCE setup has a peaking time of about 30ns, as shown in Fig. 3.10. This means firstly that ballistic deficit may reduce the collection signal (particularly at lower voltages), and secondly that negative signals may be picked up on the neighbouring strips.

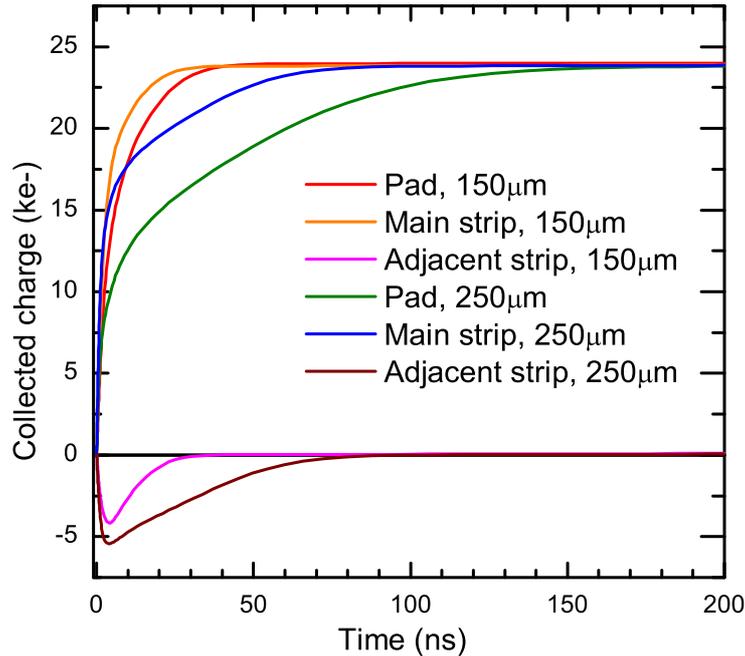


Figure 6.10: Simulated collected charge in different 3D-STC devices. The devices use either a pad or strip structure, and either $150\mu\text{m}$ or $250\mu\text{m}$ -long columns. In the case of the strip detectors, the collected charge is shown both for the main strip where the carriers were generated, and for the adjacent strip.

6.3 Experimental results

6.3.1 Pad detector I-V tests

The 3D-STC pad detectors were described in section 6.1.2. The five different pad detectors from the $300\mu\text{m}$ MCz substrate and the $500\mu\text{m}$ FZ substrate were I-V tested using the probe station, as described in section 3.1. During each test, probe needles at 0V made contact with the central pad and the guard ring, and the back contact of the detector was biased. The planar pad detectors were also tested for comparison.

Figures 6.11 and 6.12 show the I-V characteristics of the MCz and FZ pad detectors respectively. The figures give the current per cubic centimetre rather than the total detector current, to allow a fairer comparison between devices with different areas and substrate thicknesses.

The results with the two different substrate types and thicknesses are similar. All of the detectors except the 3D1 devices can be biased to 200V without breaking down. The MCz 3D1 device breaks down around 90V, and the FZ 3D1 shows a large

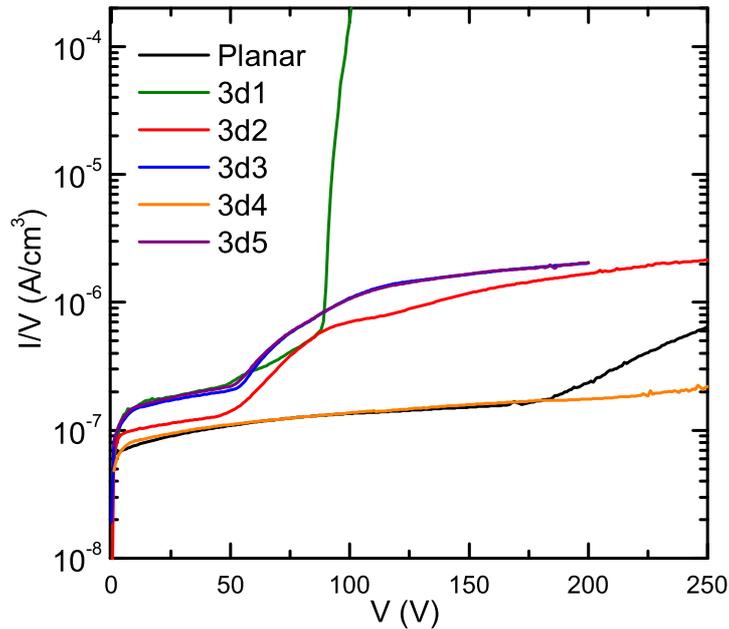


Figure 6.11: I-V characteristics of 300 μm -thick MCz 3D-STC pad detectors, and an equivalent planar detector. The current per cubic centimetre is given to allow fairer comparison between devices with different areas and substrate thicknesses.

increase in current above 125V. The 3D1 devices have a few unique features; they use rings of p-stop around each individual column, there is no p-stop between the pad and the guard ring, and metal tracks link the columns. Since the strip detectors work successfully with metal tracks, this probably means that the p-stop rings cause the avalanche breakdown, though it's also possible that the absence of the p-stop between the guard ring and pad might allow large surface currents to reach the pad.

In each substrate the 3D3 and 3D5 devices, which have identical structures apart from slightly different p-stop widths, have very closely-matched IV curves. They have slightly higher currents than 3D2, which is similar to 3D3 and 3D5 but has no p-stops between the readout columns within the pad. This is unsurprising, since these p-stops aren't actually necessary in a pad detector, and will create additional high-field regions. Still, these results show that the isolation doesn't cause any significant problems. The 3D4 detectors, which have a larger electrode spacing, have a lower current per unit volume. The planar detectors tested for comparison have similar currents per unit volume to the 3D4 devices.

During these tests, the guard ring currents were also measured. The results are shown in Fig. 6.13. The solid lines indicate the results from MCz devices, and

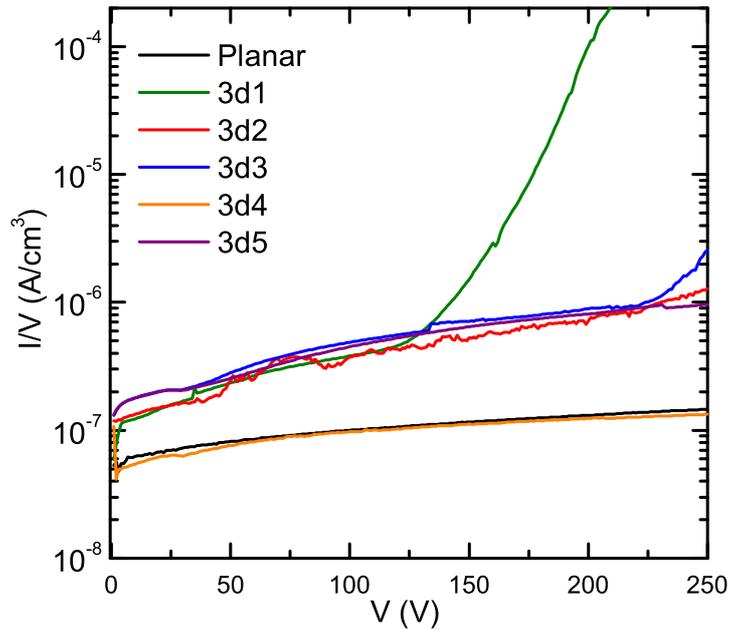


Figure 6.12: I-V characteristics of $500\mu\text{m}$ -thick FZ 3D-STC pad detectors, and an equivalent planar detector. The current per cubic centimetre is given to allow fairer comparison between devices with different areas and substrate thicknesses.

the dashed lines show the FZ results. Note that the plot shows the total guard ring current in each device, since there is no obvious way to normalise the results between devices with different areas and thicknesses. These results have a few interesting features. Firstly, most of the devices' guard currents become large at some point above 150V. This will lead to increased heating and power consumption at high biases. Secondly, the guard ring I-V curves show a step increase in current around 50V for MCz and 30V for FZ, but then level out again. This effect is larger in the MCz devices, which also have an increase in pad leakage current around 50V as shown in Fig 6.11.

As well as the guard ring of n-type columns, each device is enclosed by two p-stops, which will block surface currents by disrupting the layer of electrons at the oxide surface. As the device is biased, the depletion region around the n-type columns in the guard ring will grow outwards, and once the bias becomes high enough it's possible that the depletion region could extend beyond the p-stops. This would allow surface current to reach the guard ring more easily, explaining the kink in the guard ring current curve. Since the nominal doping concentration of the FZ substrate is lower than that of the MCz substrate, it should deplete more quickly,

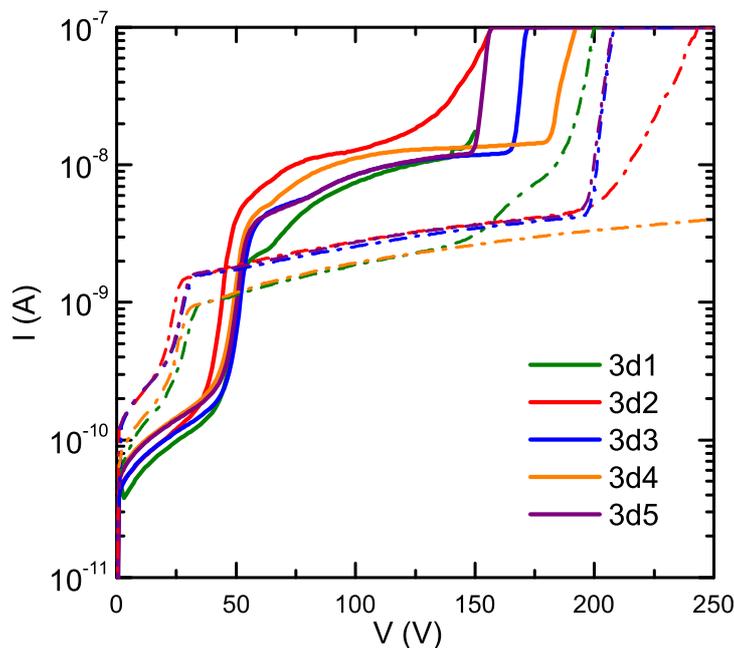


Figure 6.13: I-V behaviour of the guard ring of each 3D-STC pad detector. The results from $300\mu\text{m}$ MCz are indicated by solid lines, and the results from $500\mu\text{m}$ FZ are indicated by dashed lines. The different structures (3D1, 3D2...) are indicated by colour.

so this would explain why the effect occurs at a lower voltage in the FZ.

Pad detector C-V tests

The C-V tests on the pad detector structures were carried out using an older probe station than the one described in section 3.2. The results from the probe station tended to have poor repeatability, with the saturation capacitance varying from test to test. However, the C-V curves obtained from the MCz devices all levelled out at around 40–50V, which is in reasonable agreement with simulation. So, this can be taken as a measurement of the pad detectors' full depletion voltage.

6.3.2 Pad charge collection efficiency

The charge collection efficiency of the MCz 3D5 pad detector was tested using the pad detector beta source test setup described in section 3.3. The $300\mu\text{m}$ MCz substrate was used rather than the thicker FZ substrate so that the effect of the $150\mu\text{m}$ columns would be more distinct. The 3D5 structure was chosen because it

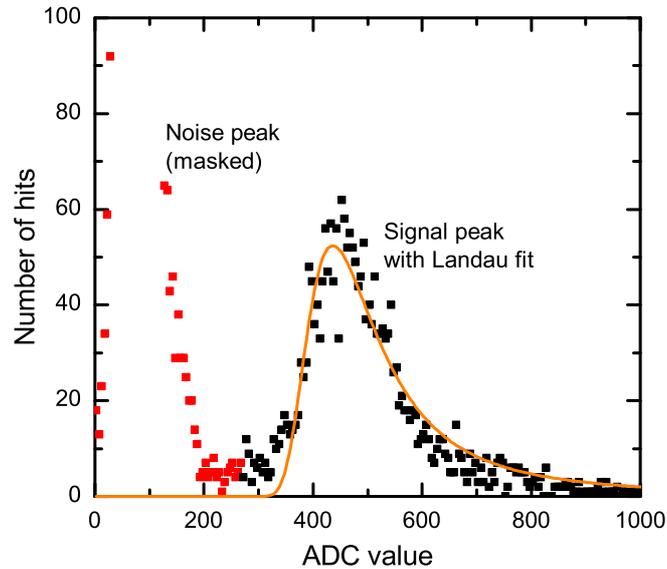


Figure 6.14: Signal spectrum obtained during a beta test of the MCz 3D5 pad detector at 30V bias. A Landau fit, shown in orange, has been applied to the spectrum to obtain the most probable value of the collected signal. The red data points have been masked during the fitting, to prevent the noise peak from disrupting the fit. An ADC level of 76 corresponds to zero signal.

gave satisfactory I-V test results, and because it used strips of p-stop between the readout electrodes, making it more similar to a strip or pixel detector than most of the other pad structures.

The test was done with the pad connected to the readout chain, the guard ring held at ground, and the back side of the detector biased. The device was tested at a range of different voltages. Figure 6.14 shows the signal spectrum obtained from the detector at 30V bias. As discussed previously, the energy deposited when a high-energy beta passes through the detector will vary from hit to hit, but the distribution of deposited energies will follow a Landau distribution. In the figure, a Landau fit has been made to the signal peak to find the most probable collection signal.

In addition to the signal peak, there is also a large noise peak around an ADC value of 76, which corresponds to zero signal. The area of the pad detector tested here was just 0.64mm^2 , whereas the area of the scintillator was much larger. So, although a collimator was used in the test setup, many betas reached the scintillator without passing through the active area of the detector. When these hits occurred, the ADC sampled the noise of the detector and readout chain, without any signal

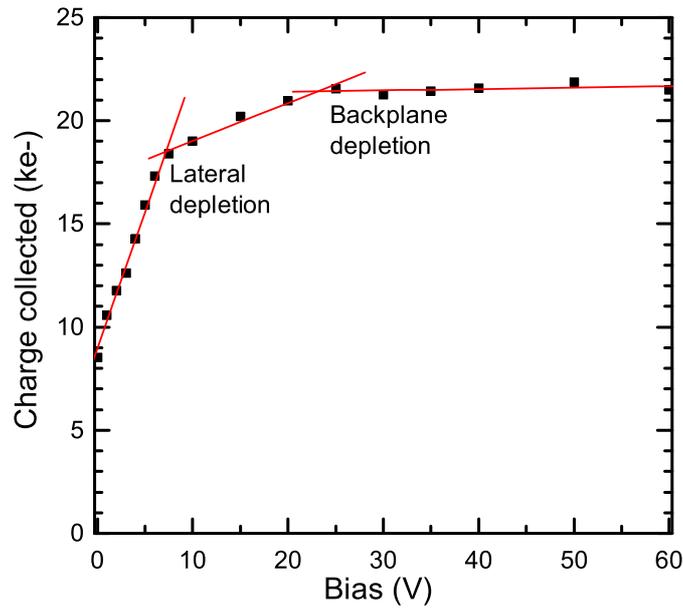


Figure 6.15: Charge collection results from the beta test of the MCz 3D5 pad detector. The fitted lines indicate the phases of the device’s depletion.

present. Comparing the number of counts in the signal and noise peaks in Fig. 6.15 shows that only 20–25% of the triggers coincided with a hit on the detector. To prevent the noise from interfering with the curve fit, the minimum of the signal spectrum between the Landau and noise peaks was found, and all the data points below this energy were masked. The masked points are shown in red in the figure.

After testing the detector at a range of voltages, the calibration obtained in section 3.3.3 was used to convert the most probable signal in ADC counts to electrons. The resulting charge collection curve is shown in Fig 6.15. Due to the long peaking time of the amplifier chain used in this setup, no ballistic deficit is expected. So, the charge collection should depend on the growth of the depletion region with bias, without being affected by the field strength within the depleted volume. Lines have been fitted to each section of the curve to indicate the corresponding stages of the depletion process.

As the bias is increased from zero to 7V, the collected charge increases rapidly. This will correspond to the initial phase where the depletion region grows laterally from each of the columns, rapidly depleting the volume between them. After lateral depletion is achieved, the collected charge continues to increase more slowly, as the depletion region grows downwards to the back surface of the detector, much like in a planar device. The curve levels out at around 25V, presumably when full

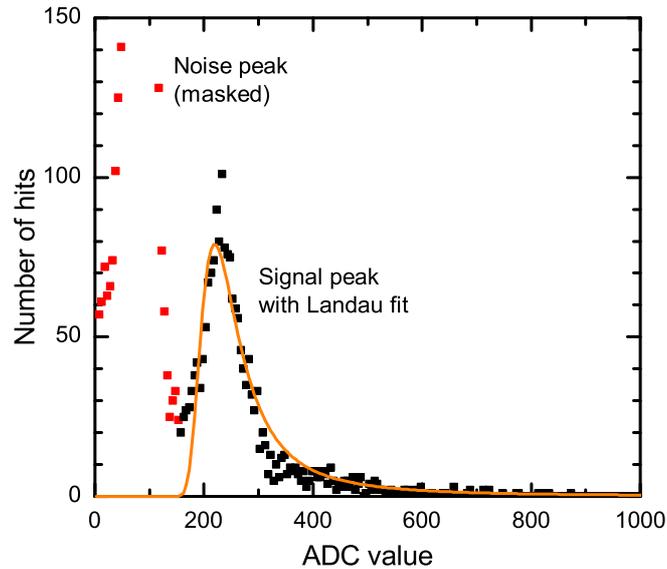


Figure 6.16: Signal spectrum obtained during beta test of the MCz 3D5 pad detector at 0V bias. A signal peak is clearly visible above the noise. A Landau fit, shown in orange, has been applied to the spectrum to obtain the most probable value of the collected signal. The red data points have been masked during the fitting, to prevent the noise peak from disrupting the fit.

depletion is reached. However, the C-V tests on the MCz pad detectors indicated that full depletion to the backplane was reached at a higher bias of 40–50V, as did the simulations. Since there is some variation in the data points, the levelling out of the CCE curve could perhaps be more gradual than these results suggest.

At lower biases, the charge collection also appears to be unexpectedly high. Even with zero applied bias, there is a collection signal of about 8000 electrons. Figure 6.16 shows the signal spectrum obtained at 0V bias. Although there isn't much of a gap between the signal peak and the noise peak, there is a clear minima between the two, so there is a genuine signal visible above the noise.

Even with zero bias applied, there will still be a depletion region where the n-type doping meets the p-type substrate. In a planar detector, this will consist of a narrow, flat layer, and each MIP passing through the pad will generate a small amount of charge in this layer. However, the depletion region around each n-type column in a 3D detector will be a long cylindrical column, as shown in Fig. 6.3. When betas pass through the detector, travelling parallel to the columns, many of them will miss the depletion region entirely and produce no signal. However, some will pass along the edge of the column and generate a substantial number of carriers

in the depletion region, producing a reasonably large signal. When the spectrum is plotted, the events with zero signal will be lost in the noise peak, and the larger hits will form a signal peak.

If this is the case, then at lower biases the number of hits seen in the Landau should decrease. At 0V, the number of hits seen in the Landau for a given number of triggers is 15% lower than at 20V. So, the hit efficiency is lower at 0V, but still surprisingly high. It's perhaps possible that the “zero” voltage setting in the test setup actually delivered a small voltage, increasing the depleted volume.

6.3.3 Pion beam test of a single-type-column 3D strip detector

In September 2007, a 3D-STC strip detector module was tested in a 120 GeV pion beam at CERN, using the LHCb-based strip readout system described in section 3.4. The basic principle of the beam test was the same as in the beta test setup described previously—the high-energy pions generated a signal as they passed through the detector, and by using a scintillator as a trigger each event could then be read out. However, the beam test also incorporated a “beam telescope” which could track the paths of the pions, making it possible to determine the location of each individual hit on the detector. So, this made it possible to map variations in the detector's response with hit position. The scintillator trigger system also recorded the time of each hit with sub-nanosecond resolution, which meant that the pulse shapes could be measured with greater accuracy.

The beam test was organised by members of the ATLAS 3D upgrade project, mainly from CERN and Manchester, who were testing another set of 3D detectors with ATLAS readout. The tests of the 3D-STC detectors with LHCb readout were done as a collaboration between Glasgow and the University of Freiburg. A preliminary data analysis was done at Glasgow, without the telescope data, and the full analysis was done at Freiburg, who have submitted the results for publication [124].

Detector and beam test setup

A photo of the test beam setup is shown in Fig. 6.17.

The two 3D-STC strip detectors used in the test were mounted on a detector module like the one shown in Fig. 3.9. As mentioned previously, the strip detectors had 64 strips, each 18.4mm long. The spacing between the n-type readout columns along each strip was $100\mu\text{m}$, and the spacing between strips was $80\mu\text{m}$. In one detector, the readout strips were isolated by strips of p-stop, and the other used a

uniform p-spray. The detectors had built-in AC coupling, so no RC network chips were required; the detectors were simply connected to the Beetle readout chips on the module by pitch adaptors.

Additionally, a module with $300\mu\text{m}$ -thick n-on-p planar strip detectors was also tested. This was used to obtain a signal calibration, and also acted as a test of the readout system. The same module was also used to calibrate the beta test setup. During the tests, the modules were held in the beam by two adjustable mounts. The mounts could be moved up and down to move the different detectors on each module in and out of the beam.

Particle tracking was done by the Bonn ATLAS Telescope (BAT [125]). This consisted of four telescope modules, each of which had two layers of crossed planar silicon strip detectors with a strip pitch of $50\mu\text{m}$. Due to the small pitch, each hit on a strip detector was shared across multiple strips, and by comparing the signal sizes on the strips the system was able to determine the hit position with greater precision. Later, the hit positions from each strip detector could be combined to find the path of the particle. The spatial resolution of the system as a whole was intended to be $5\mu\text{m}$, but since one of the modules wasn't functional during the tests the true value will be poorer.

The trigger was provided by two scintillators in the path of the beam, plus a third "veto" scintillator with a circular hole in its centre which was used to reject particles passing outside the area of the test devices. Both our LHCb readout system and the ATLAS system being used by members of the ATLAS 3D upgrade group were connected to the same 40MHz clock signal. A time-to-digital-converter (TDC) measured the time difference between the clock signal and the trigger.

The setup and operation of the LHCb-based readout electronics was much the same as for the beta tests that were done at Glasgow. The power and bias supplies for the detector modules and the repeater boards were placed next to the test bench, inside the beam enclosure. Long data and signal cables connected the repeater boards to the TELL1, which was inside the shielded hutch next to the enclosure. The trigger logic was set up so that five consecutive samples would be taken whenever a trigger signal occurred, covering a period of 125ns, so that the full shape of the signal pulse could be measured. Since the timing information was being recorded by the TDC, there was no need for the trigger logic to select hits with a particular timing. Although the telescope and the LHCb readout system both used the same trigger, they were not in direct communication with each other. So, data was taken in short runs of 50000 triggers, with both systems being started simultaneously to ensure synchronisation.

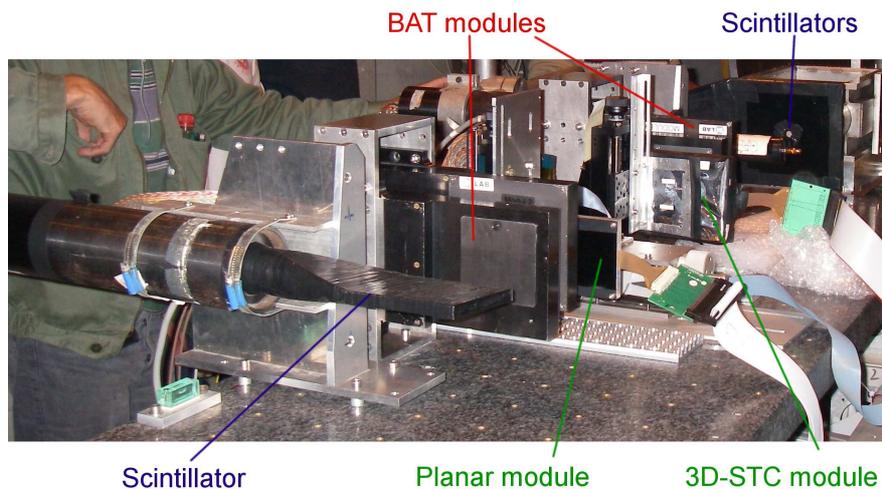


Figure 6.17: Photo of the 3D-STC test beam setup, looking along the path of the beam. The 3D-STC and planar detector modules are in the middle of the test bench, with two telescope modules in front of them and two behind them. The scintillators, which are used for triggering, are placed at the very front and back of the bench.

As discussed in section 1.1.3, if a singly-charged particle has sufficiently high energy, then the energy it will deposit in the detector will not strongly depend on its specific energy and mass. So, the number of carriers generated in the detector by 120 GeV pions will follow a Landau distribution with a most probable value of 80 electron-hole pairs per micron, just like when the detector is hit by high-energy betas. Hence, most of the preliminary analysis for the beam test was the same as for the beta tests, as described in section 3.4.4.

Basic analysis of results

The basic data analysis performed at Glasgow did not make use of the beam telescope data, and mainly involved checking that the data files could be processed properly, and that the results looked sensible. One set of detector data was analysed more fully. In this particular data set, the 3D-STC strip detector with p-spray isolation was in the beam, biased to 40V. The n-on-p planar detector in the beam was biased to 120V, above full depletion.

After processing the data with Vetra—see section 3.4.4—and adding the timing information from the TDC, the signal pulse shape from each detector was found. This was done using much the same method as for the lab beta tests. The TDC information meant that the pulse shape could be found with fine resolution; 1ns resolution was chosen, so that the number of events occurring in each time bin

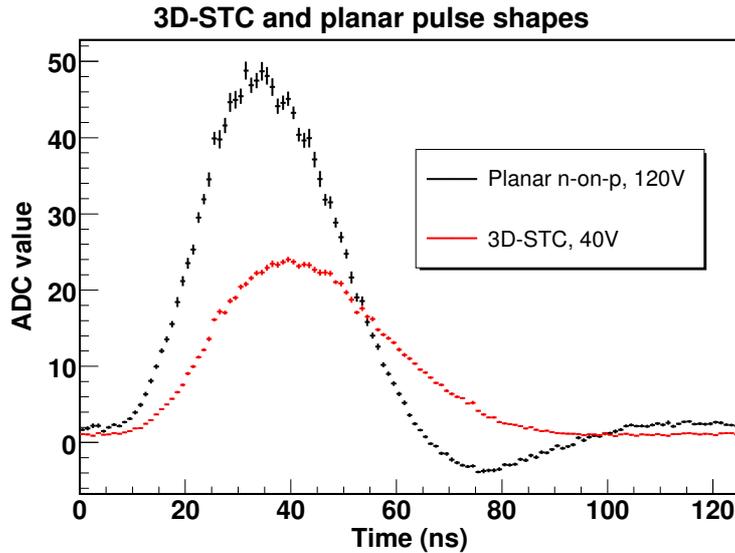


Figure 6.18: Pulse shapes obtained from beam tests on a 3D-STC strip detector with p-spray isolation at 40V bias, and an n-on-p planar strip detector at 120V. Both detectors are $300\mu\text{m}$ thick.

was reasonably high. The resulting pulse shapes are shown in Fig. 6.18. The 3D-STC detector's pulse has a smaller peak, which arrives 5ns later than the peak on the planar detector. While the signal on the planar sensor has a steep falling edge followed by an undershoot, due to the response of the shaping amplifier, the 3D-STC sensor has a long tail-off which cancels out the undershoot. The simulation results in section 6.2.4 show that the smaller peak and long tail-off signal are due to the slow vertical drift of holes in the region between the columns. Note that because the Beetle's amplifier chain has a 30ns peaking time it isn't possible to see any variations in the detector current occurring on shorter timescales. For example, the simulation results imply that the 3D-STC detector should show a short, high current signal immediately when the particle hit occurs, but this can't be distinguished.

Figure 6.19 then shows the signal spectrum from the 3D-STC detector. This was obtained by selecting only those events where the signal had been sampled in the time range of 37–46ns, around the peak of the amplifier pulse. A Landau convolved with a Gaussian was then fitted to the data, to find the most probable collection signal. The spectrum and most probable signal were also found for the planar sensor. Assuming that the planar detector gives 100% charge collection, the signal on the 3D-STC device was found to be 12400 electrons. Using the C-V tests as a guide, the detector should be approaching full depletion at this bias, so the charge loss from undepleted regions shouldn't be very high. However, the pulse shape results

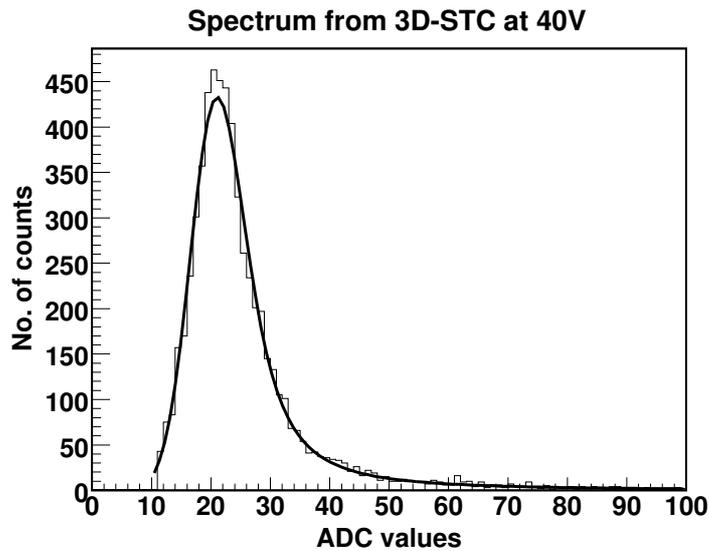


Figure 6.19: Signal spectrum obtained from the 3D-STC strip detector at 40V bias. The spectrum has been fitted with a Landau convolved with a Gaussian. The most probable signal was found to be 12400 electrons.

show that the charge collection process is slow compared to the peaking time of the Beetle's amplifier chain. So, the low signal here will largely be due to ballistic deficit. The noise level on the detector was found to be 1550 electrons, giving a signal-to-noise ratio of 8.

Key testbeam results

The full results of the analysis of the testbeam data are available in Ref. [124]. Here, the key results are summarised.

Firstly, during the testbeam the detector was tested at 60V and 80V bias, as well as the 40V seen above. As the bias was increased, the pulse shape became more similar to the pulse shape from the planar detector, with the peak becoming higher and earlier, and the tail-off signal becoming smaller. In particular, at 80V bias the tail-off became small enough for the undershoot from the main pulse to become visible. The charge collection obtained from the detector increased to 14700 electrons at 60V, and 17000e- at 80V bias.

In the simulation section, the weighting field and MIP simulations showed that when a particle hits one strip, the neighbouring strip on the side where the hit occurred should experience a fast negative signal as the electrons are collected, followed by a slower positive signal due to hole collection. The data from the strip detector at 40V was combined with the tracking data from the beam telescope. Using

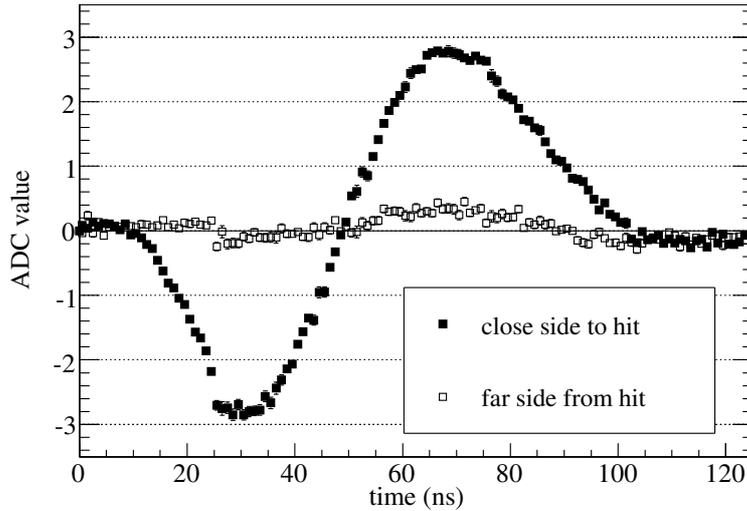


Figure 6.20: Average pulse shapes on the two neighbouring strips to the hit strip. Using the tracking data, the neighbouring strip on the same side as the hit has been distinguished from the neighbouring strip on the opposite side. These results are for the 3D-STC strip detector at 40V bias. The figure is taken from Ref. [124].

this additional hit position information, the average pulse shape was found for the neighbouring strip on the side where each hit occurred, and for the neighbouring strip on the opposite side. The pulse shapes are shown in Fig. 6.20. On the neighbouring strip on the same side as the hit, there is a negative-going signal, followed by a positive signal, as predicted. In the simulations, the negative-going signal is shorter and has a higher amplitude than the positive signal, but due to the pulse-shaping effects of the amplifier chain, the positive and negative parts of the output signal are similar in amplitude and duration. The zero-crossing occurs at 58ns, which is earlier than the undershoot signal on the main strip because the output signal is actively driven back to zero by the positive current flow, rather than by just the amplifier response. Also, the signal on the neighbouring strip on the opposite side to the hit is negligible, which confirms that the signal is due to weighting-field effects rather than any direct coupling between the readout electrodes.

The tracking information was then used to determine the variation in the detector's response with position. Since the detector consists of many copies of a basic unit cell (the n-type readout column and the surrounding detector volume), the hits from each cell in the device were superimposed to give a single cell with good statistics. The unit cell was divided into $5\mu\text{m}$ bins, and within each bin the hit efficiency was found. The efficiency was taken to be the proportion of hits with a signal

above 7000 electrons, which corresponds to a signal-to-noise ratio of at least 4. Note that the clustering algorithm was applied to each hit, which means that any signals shared across two strips have been combined before applying this signal-to-noise cut.

The results are shown in Fig. 6.21. Throughout most of the unit cell, the hit efficiency is close to the maximum value of 1. However, the hit efficiency is lower in the centre of the unit cell, due to signal loss in the n-type readout column. In this region, the efficiency drops to around 0.8. However, the column itself has only $5\mu\text{m}$ radius, and the efficiency map will be blurred by the telescope resolution, which will be poorer than $5\mu\text{m}$, and also any inaccuracies in the alignment process. So, the true efficiency will be lower than 0.8 in the column itself, but in the surrounding area it will be higher than the values given here.

There is also reduced efficiency along the midpoint between adjacent strips at $y=\pm 40\mu\text{m}$, where charge-sharing will occur. In contrast, along $x=0\mu\text{m}$ and $x=100\mu\text{m}$, which lie midway between readout columns belonging to the same strip, there is no drop in efficiency. This is probably because when carriers are shared between two strips, the chances of seeing a signal above the noise level will be reduced. In both of these regions, the charge collection will be dominated by the slow hole drift, and the resulting ballistic deficit will exacerbate this problem. Consider the weighting field in Fig. 6.7. Since the weighting potential is close to 0.5 at the midpoint between the columns, if electrons generated here are shared equally between the electrodes then the total contribution they make to the readout signal will be small.

6.4 Conclusions

The single-type-column 3D structure simplifies the 3D fabrication process substantially by only using electrode columns of one doping type, which are etched from the front surface and don't pass through the full detector thickness. The bias contact is provided by a planar implant on the back surface. FBK-IRST have produced pad and strip detectors using this structure, with $150\mu\text{m}$ -long n-type columns and p-type substrates.

Simulations of these detectors show that the region between the readout columns becomes depleted at a low voltage, much like in a full or double-sided 3D detector. However, since these devices have relatively short columns, a higher voltage is required to grow the depletion region downwards to the back contact. When electron-hole pairs are generated in the region between the columns, the electrons will drift a short distance horizontally to the columns, producing a fast signal. However, the

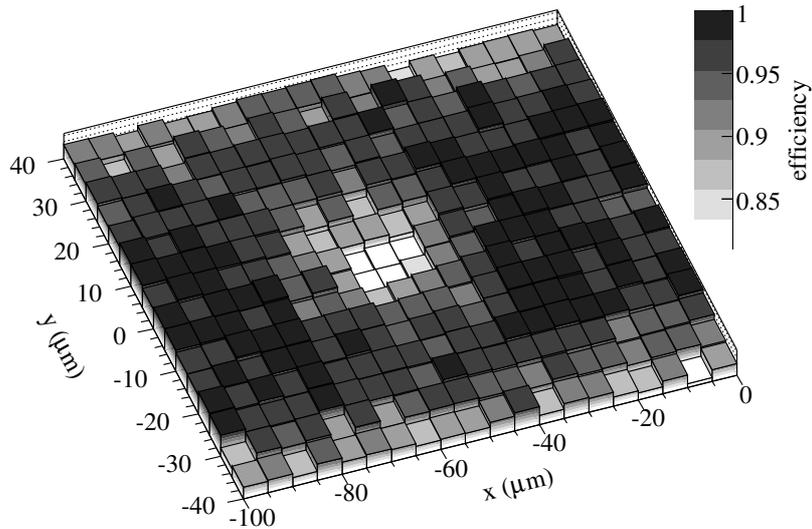


Figure 6.21: Hit efficiency map for a unit cell of the 3D-STC strip detector at 40V bias. To reduce statistical fluctuations, all the unit cells in the detector were superimposed. The hit efficiency is the proportion of hits where the signal-to-noise ratio exceeds 4. The figure is taken from [124].

holes need to cross the thickness of the substrate to reach the back contact, resulting in a slower signal. The vertical component of the field between the columns is relatively weak, particularly around the front surface, and the horizontal component of the field channels the holes to the midpoint between the columns where the field is weakest. So, the readout signal will have a long tail-off. Additionally, when the device is connected in a pad configuration, the majority of the signal will be produced by the slow hole drift. This means that these devices will be less radiation-hard than full and double-sided 3D detectors, since more charge trapping will occur.

I-V tests on the 3D-STC pad detectors show that most of the devices can be successfully biased up 200V. The only structure that breaks down at a lower voltage has individual rings of p-stop isolation around each n-type readout column, rather than strips of p-stop like in the other devices.

Charge collection efficiency tests were done on a MCz 3D-STC pad detector, using a beta source. Since the readout amplifier chain used a long peaking time, the variation in collection efficiency with bias was only dependent on the growth of the depletion region, and not the collection speed. The resulting CCE curve falls into distinct phases. Up to 7V, the collection efficiency increases rapidly with bias, corresponding to the phase of rapid depletion between the columns. Then, the charge collection increases more slowly as the depletion region grows to the back

surface.

A strip 3D-STC detector was also tested at CERN with 120 GeV pions, using LHC-speed readout electronics. Compared to a planar detector at 120V, the 3D-STC detector at 40V produced a signal pulse which peaked later and had a longer tail-off, due to the slow hole collection. This led to signal loss from ballistic deficit. Using tracking information from the beam telescope, it was shown that there was signal loss when particle hits occurred around the readout columns, as expected. Also, it was shown that negative-going signals were induced in adjacent strips, as predicted by the simulations.

Overall, these single-type-column detectors give inferior performance to full and double-sided 3D detectors. Aside from their low depletion voltages, they have no clear advantage over planar detectors, which are already widely used. So, although they are interesting test structures, they have little practical use.

Chapter 7

Conclusions

3D detectors are a novel variety of photodiode detector, with doped electrode columns passing through the thickness of a silicon substrate. These detectors have promising applications in future particle physics experiments at the Super-LHC, and as X-ray detectors in synchrotron light sources. However, before they can be used in practical applications, it will be necessary to optimise their design so that they give high performance and reliability, and can be produced on a reasonably large scale. The aim of the work in this thesis has been to characterise and understand the behaviour of various 3D detectors produced at other institutions, using a combination of lab tests and computer simulation. Using these results, future production runs can be redesigned to correct any flaws in the detectors and to improve their performance for different applications.

Firstly, the behaviour of 3D detectors was investigated in Chapter 2 using the “Synopsys TCAD” simulation package. This software can model the behaviour of a detector by approximating its structure with a “mesh” of discrete points and then applying appropriate semiconductor physics models to the mesh. As expected, the initial simulations of the device structure showed that it becomes fully depleted at a bias of just a few volts, and that the electrons and holes generated by ionising radiation will drift a short distance horizontally to reach the readout electrodes. The charge collection time was shown to be less than 1ns at a moderate bias of 100V, which is extremely fast for a photodiode.

Then, further simulations were used to find the best choice of electrode spacing for an ATLAS pixel detector that could be used at the Super-LHC. The main requirement for a Super-LHC pixel detector is radiation hardness up to a fluence of $10^{16}n_{eq}/\text{cm}^2$. The effects of this radiation damage were modelled by incorporating the carrier dynamics of radiation-induced defects into the simulation. As the electrode spacing of the 3D detector was made larger, its depletion voltage increased

quadratically, but there was little change in its breakdown voltage. Since radiation damage tends to increase a detector's depletion voltage, a smaller electrode spacing was required to achieve full depletion at high fluences. A smaller electrode spacing also gave faster charge collection, which led to less carrier trapping and a larger readout signal. However, as the electrode spacing was reduced, the readout capacitance increased rapidly, and at very small electrode spacings the signal-to-noise ratio was actually reduced. Additionally, as the electrode spacing was reduced, the total insensitive volume occupied by the columns increased. Overall, for detectors operating at $10^{16}n_{eq}/\text{cm}^2$ damage, the best trade-off was achieved with an electrode spacing of 40–55 μm . Given that the electrode spacing affects the device behaviour so strongly, future fabrication runs of 3D detectors should include devices with a range of electrode spacings centred around this optimum value, in order to compare their performance experimentally.

3D detectors designed in collaboration between CNM (an institute in Barcelona) and Glasgow, and produced by CNM, were simulated and tested in Chapter 4. These detectors had a “double sided 3D” structure, where the two sets of columns were fabricated from opposite sides of the substrate, and neither set passed through the full substrate thickness. This structure was intended to make the fabrication process easier; in particular, no support wafer was necessary, and the risk of cracking was reduced. However, the structure is less compatible with active-edge fabrication than full-3D. When this structure was simulated, most of the device volume behaved like a full-3D detector. However, around the front and back surfaces, the electric field was weaker, and a higher bias of about 8V was required to grow the depletion region to the back surface. Following radiation damage, the simulated collection efficiency was similar to that of a standard 3D detector. The double-sided 3D detector's breakdown behaviour was dominated by high-field regions around the column tips, rather than surface effects like in a full-3D detector. While the full-3D detector's p-spray isolation needed to be carefully matched to the oxide charge to ensure good breakdown behaviour, the double-sided detector achieved the same breakdown voltage without any fine-tuning, meaning that these devices are less likely to suffer from early breakdown.

Pad and strip detectors fabricated by CNM were electrically characterised. The C-V tests confirmed that they have a low lateral depletion voltage of about 3V. I-V tests showed some unreliability in the detectors' behaviour; the guard ring currents were large, and varied a great deal between devices. However, a second set of detectors produced by CNM, which used p-stop implants to reduce surface currents, were less prone to this problem. Charge collection tests on strip detectors with betas

produced mixed results. A heavily irradiated detector at a fluence of $5 \times 10^{15} n_{eq}/\text{cm}^2$ showed a relatively high charge collection signal of 12800 electrons at 200V bias, demonstrating the structure's radiation hardness. However, tests on an unirradiated detector showed unexpectedly low charge collection. For example, at 30V, above full depletion, the signal was just 13200 electrons, corresponding to 55% CCE. This was possibly due to poor coupling between the detector and the readout chip. Since CNM have recently produced a new set of detectors, this can be tested by measuring the collection efficiency of an AC-coupled unirradiated strip detector.

The collection signal seen on the irradiated double-sided 3D strip detector was similar to the values reported from tests on full-3D detectors produced at Stanford, as shown back in Fig. 4.28. At high fluences, these 3D detectors gave much higher collection efficiencies than planar detectors—for example, the double-sided 3D strip detector at $5 \times 10^{15} n_{eq}/\text{cm}^2$ had nearly twice the collection efficiency of a typical planar detector. Although the 3D detectors experienced higher noise levels, their signal-to-noise ratios were still superior to those of typical planar detectors at high fluences, as shown in Fig. 4.29. So, both double-sided and full-3D detectors are promising options for the inner pixel layer of an upgraded ATLAS detector at the Super-LHC. The next step in double-sided 3D research is to test the collection efficiency of the new devices from CNM across a wider range of radiation fluences. Bump-bonding the new double-sided 3D ATLAS sensors to readout chips and measuring their signal-to-noise behaviour would also make it possible to compare different detector technologies for the Super-LHC more directly.

CNM also produced Medipix2 double-sided 3D pixel detectors, which could potentially be used for experiments such as X-ray crystallography. The Medipix2 readout chip is designed for X-ray detection, and can count individual X-ray photons. This single-photon counting approach has a variety of benefits, such as a large dynamic range and reduced noise. Furthermore, by using 3D photodiodes, the image quality could be improved by reducing charge sharing. When three of the double-sided 3D detectors were bump-bonded to Medipix2 chips, all three worked successfully. However, two had some dead pixels along their edges, where the bonds did not make proper contact between the two chips. This was possibly due to wafer bending during fabrication, though damage during saw-cutting could also have this effect. So, CNM might need to alter their fabrication methods (such as the polysilicon deposition conditions) to reduce this problem.

Basic tests of the collection efficiency with an X-ray tube showed that most of the detector volume was depleted around 2V bias, and the full device was depleted at 9V. This agreed with the simulations and previous tests. The detectors were then

tested in an X-ray test beam at Diamond Light Source. In spectral response and edge-imaging tests, the 3D detectors showed lower charge sharing than standard planar detectors. For example, the spectral response showed that 24% of the hits on a double-sided 3D detector at 22V were shared, compared to 40% on a planar detector at 100V. Later tests with alpha particles also demonstrated reduced charge sharing. A 3D Medipix2 detector was also used to record a diffraction pattern from a powdered silicon sample, to demonstrate its use in a real experiment. Currently, Glasgow and Diamond are planning further tests where a focused X-ray beam will be scanned across a pixel to directly map the charge sharing and measure the signal loss inside the readout column. It would also be useful to make detective quantum efficiency measurements, to allow more direct comparison of image quality between these 3D detectors and other technologies.

Given the rapidly increasing intensity of new X-ray sources such as Diamond, it is almost inevitable that there will be a shift from using CCDs (which are relatively slow) to photodiode-based detectors. In addition to the reduced charge sharing demonstrated here, 3D detectors with active edges also offer a smaller dead area at the edge of each detector. With improved fabrication technology, it would be possible to use thicker 3D detectors to sense higher-energy X-rays. However, the usefulness of 3D detectors for X-ray detection will depend heavily on their cost, their reliability, and how easily they can be mass produced. In particular, there are alternative technologies that could improve X-ray detection. For example, new single-photon-counting readout chips (such as Medipix3) could compensate for charge-sharing effects by performing basic cluster analysis on each hit, and alternative semiconductor materials such as CdZnTe could be more effective for sensing high-energy X-rays. However, in some cases these technologies could be combined with the 3D structure.

Another set of detectors, fabricated at FBK-IRST in Trento, were also simulated and tested in Chapter 6. These had a simpler, single-type-column 3D structure, with a set of n-type readout columns etched from the front surface partway through the wafer, and a planar p-type contact on the back surface. In simulations, the region between the n-type columns depleted rapidly, but a higher bias was required to grow the depletion region to reach the back surface. The simulations also showed that electrons generated between the columns will drift a short distance horizontally to the columns and be collected, whereas holes must drift across the thickness of the substrate to reach the back surface. This meant that the simulated current signals consisted of a fast pulse from electron collection, followed by a long tail-off from hole collection. Beta tests with two different readout systems demonstrated these aspects

of the device's behaviour. A pad detector test setup, using an amplifier chain with a long peaking time, showed that the charge collection increased rapidly up to a bias of 7V, as the region between columns depleted, and the collection efficiency reached its maximum at 25V. However, when a strip detector was tested with LHC-speed readout electronics, there was substantial signal loss due to the slow charge collection. Overall, this showed that single-type column detectors are inferior to full and double-sided 3D, and have few advantages over standard planar detectors. So, although they are interesting test structures, they have little practical use.

The new, larger run of double-sided 3D detectors recently produced by CNM will be tested at Glasgow, the University of Freiburg, and Diamond Light Source. Other groups are also working to produce 3D devices in greater quantities and with optimised performance—for example, at Stanford Nanofabrication Centre (where the first 3D detectors were made) and at FBK-IRST. This research is primarily aimed at developing detectors for the Super-LHC in a decade's time. With this application driving the development of 3D detectors, and the necessary fabrication tools becoming more and more widely used, the possibility of using 3D detectors for other applications such as X-ray detection will also increase.

Bibliography

- [1] G. Lutz. *Semiconductor Radiation Detectors Device Physics*. Springer-Verlag, Berlin Heidelberg New York, first edition, 1999.
- [2] S. Sze. *Semiconductor Devices Physics and Technology*. John Wiley and Sons, New York, 1985.
- [3] E. S. Yang. *Microelectronic Devices*. McGraw-Hill International Editions, Electrical and Electronic Engineering Series. McGraw-Hill, Singapore, 1988.
- [4] S. Ramo. Currents induced by electron motion. *Proceedings of the IRE*, 27(9), 584–585, 1939.
- [5] W. R. Leo. *Techniques for Nuclear and Particle Physics Experiments: a how-to approach*. Springer-Verlag, Berlin Heidelberg New York, 2nd revised edition, 1994.
- [6] H. Sadrozinski. Tracking detectors for the LHC upgrade. *5th International Conference on Radiation Effects on Semiconductor Materials Detectors and Devices (RESMDD)*. Florence, Italy, 2004.
- [7] X. Llopart, M. Campbell, R. Dinapoli, D. S. Segundo, and E. Pernigotti. Medipix2: A 64-k pixel readout chip with $55\mu\text{m}$ square elements working in single photon counting mode. *IEEE Trans. Nucl. Sci.*, 49(5), 2279–2283, 2002.
- [8] S. L. Shapiro, W. M. Dunwoodie, J. F. Arens, J. G. Jernigan, and S. Gaalema. Silicon pin diode-array hybrids for charged-particle detection. *Nucl. Instr. and Meth. A*, 275(3), 580–586, 1989.
- [9] G. Lutz. Radiation damage in structured silicon semiconductor detectors. *IEEE Nuclear and Space Radiation Effects Conference*. Phoenix, Arizona, 2002.
- [10] T. P. Ma and P. V. Dressendorfer. *Ionizing radiation effects in MOS devices and circuits*. John Wiley and Sons, New York, 1989.

-
- [11] C. Piemonte. Device simulations of isolation techniques for silicon microstrip detectors made on p-type substrates. *IEEE Trans. Nucl. Sci.*, 53(3), 1694–1705, 2006.
- [12] L. Evensen, A. Hanneborg, B. S. Avset, and M. Nese. Guard ring design for high-voltage operation of silicon detectors. *Nucl. Instr. and Meth. A*, 337(1), 44–52, 1993.
- [13] B. Burke, P. Jorden, and P. Vu. CCD technology. *Experimental Astronomy*, 19(1-3), 69–102, 2005.
- [14] M. Bigas, E. Cabruja, J. Forest, and J. Salvi. Review of CMOS image sensors. *Microelectronics Journal*, 37(5), 433–451, 2006.
- [15] S. I. Parker, C. J. Kenney, and J. Segal. 3D - a proposed new architecture for solid-state radiation detectors. *Nucl. Instr. and Meth. A*, 395(3), 328–343, 1997.
- [16] V. A. Wright, W. D. Davidson, J. J. Melone, V. O’Shea, K. M. Smith, L. Donohue, et al. Three-dimensional Medipix - a new generation of X-ray detectors. *IEEE Trans. Nucl. Sci.*, 52(5), 1873–1876, 2005.
- [17] J. Morse, C. J. Kenney, E. M. Westbrook, I. Naday, and S. I. Parker. The spatial and energy response of a 3D architecture silicon detector measured with a synchrotron X-ray microbeam. *Nucl. Instr. and Meth. A*, 524(1-3), 236–244, 2004.
- [18] J. W. Judy. Microelectromechanical systems (MEMS): fabrication, design and applications. *Smart Materials and Structures*, 10(6), 1115–1134, 2001.
- [19] C. Kenney, S. Parker, J. Segal, and C. Storment. Silicon detectors with 3-D electrode arrays: Fabrication and initial test results. *IEEE Trans. Nucl. Sci.*, 46(4), 1224–1236, 1999.
- [20] G. Pellegrini, M. Lozano, M. Ullan, R. Bates, C. Fleta, and D. Pennicard. First double-sided 3-D detectors fabricated at CNM-IMB. *Nucl. Instr. and Meth. A*, 592(1-2), 38–43, 2008.
- [21] J. K. Bhardwaj and H. Ashraf. Advanced silicon etching using high-density plasmas. *Proc. SPIE*, volume 2639, pages 224–233. 1995.

-
- [22] C. J. Kenney, S. Parker, and E. Walckiers. Results from 3-D silicon sensors with wall electrodes: Near-cell-edge sensitivity measurements as a preview of active-edge sensors. *IEEE Trans. Nucl. Sci.*, 48(6), 2405–2410, 2001.
- [23] C. J. Kenney, J. D. Segal, E. Westbrook, S. Parker, J. Hasi, C. da Via, et al. Active-edge planar radiation sensors. *Nucl. Instr. and Meth. A*, 565(1), 272–277, 2006.
- [24] S. I. Parker, C. J. Kenney, D. Gnani, A. C. Thompson, E. Mandelli, G. Meddeler, et al. 3DX: An X-ray pixel array detector with active edges. *IEEE Trans. Nucl. Sci.*, 53(3), 1676–1688, 2006.
- [25] C. J. Kenney, J. Hasi, S. Parker, A. C. Thompson, and E. Westbrook. Use of active-edge silicon detectors as X-ray beam monitors. *Nucl. Instr. and Meth. A*, 582(1), 178–181, 2007.
- [26] A. Kok. *Signal formation and active edge studies of 3D silicon detector technology*. Thesis, Brunel University, 2005.
- [27] E. Wilson. *An introduction to particle accelerators*. Oxford University Press, Oxford, 2001.
- [28] K. Wille. Synchrotron radiation sources. *Reports on Progress in Physics*, 54(8), 1005–1068, 1991.
- [29] L. Evans and B. Philip. LHC machine. *Journal of Instrumentation*, 3(08), S08001, 2008.
- [30] O. Bruning and P. Collier. Building a behemoth. *Nature*, 448(7151), 285–289, 2007.
- [31] M. Jeitler. Trigger systems at LHC experiments. *Nucl. Instr. and Meth. A*, 598(1), 305–311, 2009.
- [32] P. W. Higgs. Broken symmetries and masses of gauge bosons. *Physical Review Letters*, 13(16), 508, 1964.
- [33] J. Ellis. Physics at the LHC. *European Physical Journal C*, 34(1), 51–56, 2004.
- [34] The ATLAS collaboration. The ATLAS experiment at the CERN Large Hadron Collider. *Journal of Instrumentation*, 3(08), S08003, 2008.

- [35] F. Gianotti, M. L. Mangano, T. Virdee, S. Abdullin, G. Azuelos, A. Ball, et al. Physics potential and experimental challenges of the LHC luminosity upgrade. *European Physical Journal C*, 39(3), 293–333, 2005.
- [36] F. Gianotti. Physics opportunities of the LHC luminosity upgrade. *Nuclear Physics B - Proceedings Supplements*, 147, 23–32, 2005.
- [37] Y. Unno. Silicon sensor development for the ATLAS upgrade for SLHC. *Nucl. Instr. and Meth. A*, 569(1), 41–47, 2006.
- [38] M. Moll. Development of radiation hard sensors for very high luminosity colliders - CERN-RD50 project. *Nucl. Instr. and Meth. A*, 511(1-2), 97–105, 2003.
- [39] C. da Via, J. Hasi, C. Kenney, V. Linhart, S. Parker, T. Slavicek, et al. Radiation hardness properties of full-3D active edge silicon sensors. *Nucl. Instr. and Meth. A*, 587(2-3), 243–249, 2008.
- [40] G. Greaves and I. Munro. *Synchrotron Radiation - Sources and Applications*. Scottish Universities Summer School in Physics, 1985.
- [41] N. Marks. Synchrotron-radiation sources. *Radiation Physics and Chemistry*, 45(3), 315–331, 1995.
- [42] Diamond Light Source. Diamond status report. Available: http://www.diamond.ac.uk/Publications/Status_Report_Aug05.htm, 2005.
- [43] Diamond Light Source. Diamond website. Available: <http://www.diamond.ac.uk>, 2008.
- [44] D. H. Bilderback, P. Elleaume, and E. Weckert. Review of third and next generation synchrotron light sources. *Journal of Physics B-Atomic Molecular and Optical Physics*, 38(9), S773–S797, 2005.
- [45] Synopsys Inc. Synopsys TCAD manuals. Available: <http://www.synopsys.com>, 2007.
- [46] W. Fichtner, D. J. Rose, and R. E. Bank. Semiconductor device simulation. *IEEE Trans. Electron Dev.*, 30(9), 1018–1030, 1983.
- [47] R. E. Bank, D. J. Rose, and W. Fichtner. Numerical methods for semiconductor-device simulation. *IEEE Trans. Electron Dev.*, 30(9), 1031–1041, 1983.

- [48] M. N. O. Sadiku. A simple introduction to finite-element analysis of electromagnetic problems. *IEEE Transactions on Education*, 32(2), 85–93, 1989.
- [49] R. E. Bank and D. J. Rose. Global approximate Newton methods. *Numerische Mathematik*, 37(2), 279–295, 1981.
- [50] W. Shockley and W. T. Read. Statistics of the recombinations of holes and electrons. *Physical Review*, 87(5), 835–842, 1952.
- [51] R. H. Richter, L. Andricek, T. Gebhart, D. Hauff, J. Kemmer, G. Lutz, et al. Strip detector design for ATLAS and HERA-B using two-dimensional device simulation. *Nucl. Instr. and Meth. A*, 377(2-3), 412–421, 1996.
- [52] R. E. Bank, W. M. Coughran, W. Fichtner, E. H. Grosse, D. J. Rose, and R. K. Smith. Transient simulation of silicon devices and circuits. *IEEE Trans. Electron Dev.*, 32(10), 1992–2007, 1985.
- [53] G. Pellegrini, C. Fleta, F. Campabadal, S. Diez, M. Lozano, J. M. Rafi, et al. Technology development of p-type microstrip detectors with radiation hard p-spray isolation. *Nucl. Instr. and Meth. A*, 566(2), 360–365, 2006.
- [54] H. F. W. Sadrozinski. Applications of silicon detectors. *IEEE Trans. Nucl. Sci.*, 48(4), 933–940, 2001.
- [55] V. A. Wright, A. Blue, M. Horn, I. Johnston, E. Koskiahde, L. Lea, et al. 3D detectors - a new direction for Medipix. *Nucl. Instr. and Meth. A*, 531(1-2), 56–61, 2004.
- [56] L. A. Hamel and M. Julien. Comments on Ramo’s theorem. *Hard X-Ray and Gamma-Ray Detector Physics III*, volume 4507, pages 255–263. Proc. SPIE, 2001.
- [57] V. van Lint. *Mechanisms of Radiation Effects in Electronic Materials*, volume 1. John Wiley and Sons, 1980.
- [58] V. Eremin, E. Verbitskaya, and Z. Li. Effect of radiation induced deep level traps on Si detector performance. *Nucl. Instr. and Meth. A*, 476(3), 537–549, 2002.
- [59] G. Lutz. Effects of deep level defects in semiconductor detectors. *Nucl. Instr. and Meth. A*, 377(2-3), 234–243, 1996.

- [60] M. Moll. *Radiation Damage in Silicon Particle Detectors - microscopic defects and macroscopic properties*. Thesis, Hamburg, 1999.
- [61] R. Wunstorf. *Systematische Untersuchungen zur Strahlenresistenz von Silizium-Detektoren für die Verwendung in Hochenergiephysik-Experimenten*. Thesis, Hamburg Univ., 1992.
- [62] G. Casse, P. P. Allport, and M. Hanlon. Improving the radiation hardness properties of silicon detectors using oxygenated n-type and p-type silicon. *IEEE Trans. Nucl. Sci.*, 47(3), 527–532, 2000.
- [63] V. Radicci, L. Borrello, M. Boscardin, M. Bruzzi, D. Creanza, G. F. Dalla Betta, et al. Study of radiation damage induced by 24 GeV/c and 26 MeV protons on heavily irradiated MCz and FZ silicon detectors. *Nucl. Instr. and Meth. A*, 570(2), 330–335, 2007.
- [64] P. P. Allport. ATLAS tracking at the Super-LHC. *Nucl. Instr. and Meth. A*, 579(2), 592–594, 2007.
- [65] V. Cindro, G. Kramberger, M. Lozano, I. Mandic, M. Mikuz, G. Pellegrini, et al. Trapping of electrons and holes in p-type silicon irradiated with neutrons. *IEEE Nuclear Science Symposium*, volume 1, pages 139–142. San Diego, 2006.
- [66] G. Kramberger, V. Cindro, I. Mandic, M. Mikuz, and M. Zavrtanik. Effective trapping time of electrons and holes in different silicon materials irradiated with neutrons, protons and pions. *Nucl. Instr. and Meth. A*, 481(1-3), 297–305, 2002.
- [67] M. Petasecca, F. Moscatelli, D. Passeri, and G. U. Pignatell. Numerical simulation of radiation damage effects in p-type and n-type FZ silicon detectors. *IEEE Trans. Nucl. Sci.*, 53(5), 2971–2976, 2006.
- [68] M. Petasecca, F. Moscatelli, D. Passeri, G. U. Pignatell, and C. Scarpello. Numerical simulation of radiation damage effects in p-type silicon detectors. *Nucl. Instr. and Meth. A*, 563(1), 192–195, 2006.
- [69] D. Passeri, P. Ciampolini, G. M. Bilei, and F. Moscatelli. Comprehensive modeling of bulk-damage effects in silicon radiation detectors. *IEEE Trans. Nucl. Sci.*, 48(5), 1688–1693, 2001.
- [70] M. Ahmed, S. J. Watts, J. Matheson, and A. Holmes-Siedle. Deep-level transient spectroscopy studies of silicon detectors after 24 GeV proton irradiation

- and 1 MeV neutron irradiation. *Nucl. Instr. and Meth. A*, 457(3), 588–594, 2001.
- [71] J. W. Slotboom and H. C. Degraaff. Measurements of bandgap narrowing in Si bipolar-transistors. *Solid-State Electronics*, 19(10), 857–862, 1976.
- [72] Michael Moll on behalf of the RD50 collaboration. Radiation tolerant semiconductor sensors for tracking detectors. *Nucl. Instr. and Meth. A*, 565(1), 202–211, 2006.
- [73] M. Lozano, G. Pellegrini, C. Fleta, C. Loderer, J. M. Rafi, M. Ullan, et al. Comparison of radiation hardness of p-in-n, n-in-n, and n-in-p silicon pad detectors. *IEEE Trans. Nucl. Sci.*, 52(5), 1468–1473, 2005.
- [74] P. P. Allport, G. Casse, M. Lozano, P. Sutcliffe, J. J. Velthuis, and J. Vossebeld. Performance of p-type micro-strip detectors after irradiation to 7.5×10^{15} p/cm². *IEEE Trans. Nucl. Sci.*, 52(5), 1903–1906, 2005.
- [75] G. Kramberger, V. Cindro, I. Mandic, and M. Mikuz. Impact of annealing of trapping times on charge collection in irradiated silicon detectors. *Nucl. Instr. and Meth. A*, 579(2), 762–765, 2007.
- [76] T. Lari and C. Troncon. Simulation of signals in ultra radiation-hard silicon pixel detectors. *IEEE Trans. Nucl. Sci.*, 53(5), 2923–2930, 2006.
- [77] I. Peric, L. Blanquart, G. Comes, P. Denes, K. Einsweller, P. Fischer, et al. The FEI3 readout chip for the ATLAS pixel detector. *Nucl. Instr. and Meth. A*, 565(1), 178–187, 2006.
- [78] G. Gagliardi. The ATLAS pixel detector: A hundred million channels vertex detector for LHC. *Nucl. Instr. and Meth. A*, 546(1-2), 67–71, 2005.
- [79] M. Kohler. *Studies of the Timing Behaviour of the ATLAS Pixel Detector*. Dissertation, University of Siegen, 2008.
- [80] O. Rohne and E. Bolle. CERN 3D testing: Summary and new test results. Slides available: <http://indico.cern.ch/materialDisplay.py?contribId=2&materialId=slides&%confId=21616>, 2007.
- [81] A. Andreazza. Progresses on the ATLAS pixel detector. *Nucl. Instr. and Meth. A*, 461(1-3), 168–171, 2001.

- [82] J. Grosse-Knetter. The ATLAS pixel detector. *Nucl. Instr. and Meth. A*, 568(1), 252–257, 2006.
- [83] D. Arutinov, M. Barbero, R. Beccherle, V. Buscher, G. Darbo, R. Ely, et al. Digital architecture and interface of the new ATLAS pixel front-end IC for upgraded LHC luminosity. *IEEE Nuclear Science Symposium*. Dresden, 2008.
- [84] A. Chilingarov. Recommendations towards a standardisation of the macroscopic parameter measurements part I: IV and CV measurements in Si diodes. RD50 Technical Note 2003/03, 2004.
- [85] H. Bichsel. Straggling in thin silicon detectors. *Rev. Mod. Phys.*, 60(3), 663, 1988.
- [86] R. Brun and F. Rademakers. The ROOT system homepage. Available: <http://root.cern.ch>, 2006.
- [87] The LHCb collaboration. The LHCb detector at the LHC. *Journal of Instrumentation*, 3(08), S08005, 2008.
- [88] M. Agari, N. van Bakel, C. Bauer, D. Baumeister, M. van Beuzekom, M. Feuerstack-Raible, et al. Beetle - a radiation hard readout chip for the LHCb experiment. *Nucl. Instr. and Meth. A*, 518(1-2), 468–469, 2004.
- [89] G. Haefeli, A. Bay, A. Gong, H. Gong, M. Muecke, N. Neufeld, et al. The LHCb DAQ interface board TELL1. *Nucl. Instr. and Meth. A*, 560(2), 494–502, 2006.
- [90] S. Lochner and M. Schmelling. The Beetle reference manual - chip version 1.3, 1.4 and 1.5. CERN report LHCb-2005-105, 2006.
- [91] L. Eklund, M. Pivk, J. Buytaert, P. Collins, D. Eckstein, and J. P. Palacios. Absolute calibration of Beetle 1.3 test-pulse, noise and header using MIP signals from an ATLAS reference detector. CERN report LHCb-2005-021, 2005.
- [92] G. Haefeli. *Contribution to the development of the acquisition electronics for the LHCb experiment*. Thesis, EPFL Lausanne, 2004.
- [93] Philips Semiconductors. The I²C bus specification. Available: http://www.nxp.com/acrobat_download/literature/9398/39340011.pdf, 2001.
- [94] T. Szumlak and C. Parkes. Description of the Vetra project and its application for the VELO detector. CERN report LHCb-2008-022, 2008.

- [95] B. Schmitt, C. Bronnimann, E. F. Eikenberry, G. Hulsen, H. Toyokawa, R. Horisberger, et al. Development of single photon counting detectors at the Swiss Light Source. *Nucl. Instr. and Meth. A*, 518(1-2), 436–439, 2004.
- [96] N. Demaria, S. Albergo, M. Angarano, P. Azzi, E. Babucci, N. Bacchetta, et al. New results on silicon microstrip detectors of CMS tracker. *Nucl. Instr. and Meth. A*, 447(1-2), 142–150, 2000.
- [97] C. da Via, E. Bolle, K. Einsweiler, M. Garcia-Sciveres, J. Hasi, C. Kenney, et al. 3D active edge silicon sensors with different electrode configurations: Radiation hardness and noise performance. *IEEE Nuclear Science Symposium*. Honolulu, 2007.
- [98] G. Casse, A. Affolder, and P. P. Allport. Charge collection efficiency measurements for segmented silicon detectors irradiated to $1 \times 10^{16} n.cm^{-2}$. *IEEE Trans. Nucl. Sci.*, 55(3), 1695–1699, 2008.
- [99] C. da Via and S. J. Watts. The geometrical dependence of radiation hardness in planar and 3D silicon detectors. *Nucl. Instr. and Meth. A*, In Press, Accepted Manuscript, 2009.
- [100] H. P. Myers. *Introductory Solid State Physics*. Taylor and Francis, London, 2nd edition, 1997.
- [101] C. Broennimann, E. F. Eikenberry, B. Henrich, R. Horisberger, G. Huelsen, E. Pohl, et al. The PILATUS 1M detector. *Journal of Synchrotron Radiation*, 13, 120–130, 2006.
- [102] RCSB. Protein data bank. Available: <http://www.rcsb.org>, 2008.
- [103] K. Moffat and Z. Ren. Synchrotron radiation applications to macromolecular crystallography. *Current Opinion in Structural Biology*, 7(5), 689–696, 1997.
- [104] I. A. Cunningham and R. Shaw. Signal-to-noise optimization of medical imaging systems. *J. Opt. Soc. Am. A*, 16(3), 621–632, 1999.
- [105] C. Bronnimann, E. F. Eikenberry, R. Horisberger, G. Hulsen, B. Schmitt, C. Schulze-Briese, et al. Continuous sample rotation data collection for protein crystallography with the PILATUS detector. *Nucl. Instr. and Meth. A*, 510(1-2), 24–28, 2003.
- [106] E. M. Westbrook. Performance characteristics needed for protein crystal diffraction X-ray detectors. *Proc. SPIE*, volume 3774, pages 2–16. 1999.

- [107] M. Chmeissani and B. Mikulec. Performance limits of a single photon counting pixel system. *Nucl. Instr. and Meth. A*, 460(1), 81–90, 2001.
- [108] M. G. Bisogni, P. Delogu, M. E. Fantacci, G. Mettivier, M. C. Montesi, M. Novelli, et al. A Medipix2-based imaging system for digital mammography with silicon pixel detectors. *IEEE Trans. Nucl. Sci.*, 51(6), 3081–3085, 2004.
- [109] X. Llopart. MPIX2MXR20 manual v2.3. Available: <http://medipix.web.cern.ch/MEDIPIX/Medipix2/indexMPIX2.html>, 2006.
- [110] E. R. Fossum. CMOS image sensors: electronic camera-on-a-chip. *IEEE Trans. Electron Dev.*, 44(10), 1689–1698, 1997.
- [111] S. Hejazi and D. P. Trauernicht. System considerations in CCD-based X-ray imaging for digital chest radiography and digital mammography. *Medical Physics*, 24(2), 287–297, 1997.
- [112] R. Ballabriga, M. Campbell, E. H. M. Heijne, X. Llopart, and L. Tlustos. The Medipix3 prototype, a pixel readout chip working in single photon counting mode with improved spectrometric performance. *IEEE Trans. Nucl. Sci.*, 54(5), 1824–1829, 2007.
- [113] A. Altarelli, R. Brinkmann, M. Chergui, W. Decking, B. Dobson, S. Dusterer, et al. The European X-ray Free-Electron Laser - Technical design report. DESY report 2006-097, available: http://xfel.desy.de/technical_information, 2007.
- [114] J. Moy. Signal-to-noise ratio and spatial resolution in X-ray electronic imagers: is the MTF a relevant parameter? *Med Phys.*, 27(1), 86–93, 2000.
- [115] C. Ponchut, F. Zontone, and H. Graafsma. Experimental comparison of pixel detector arrays and CCD-based systems for X-ray area detection on synchrotron beamlines. *IEEE Trans. Nucl. Sci.*, 52(5), 1760–1765, 2005.
- [116] L. Tlustos, J. Kalliopuska, R. Ballabriga, M. Campbell, S. Eranen, and X. Llopart. Characterisation of a semi 3-D sensor coupled to Medipix2. *Nucl. Instr. and Meth. A*, 580(2), 897–901, 2007.
- [117] Z. Vykydal, J. Jakubek, and S. Pospisil. USB interface for Medipix2 pixel device enabling energy and position-sensitive detection of heavy charged particles. *Nucl. Instr. and Meth. A*, 563(1), 112–115, 2006.

- [118] T. Holy and Z. Vykydal. Pixelman software. Available: <http://aladdin.utef.cvut.cz/ofat/others/Pixelman/Pixelman.html>, 2008.
- [119] H. Fujita, D. Y. Tsai, T. Itoh, K. Doi, J. Morishita, K. Ueda, et al. A simple method for determining the modulation transfer function in digital radiography. *IEEE Trans. Med. Imaging*, 11(1), 34–39, 1992.
- [120] NIST. Standard reference materials program. See <http://ts.nist.gov/measurementsservices/referencematerials/index.cfm>, 2000.
- [121] M. Campbell, E. Heijne, T. Holy, J. Idarraga, J. Jakubek, C. Lebel, et al. Study of the charge sharing in a silicon pixel detector by means of alpha-particles interacting with a Medipix2 device. *Nucl. Instr. and Meth. A*, 591(1), 38–41, 2008.
- [122] C. Piemonte, M. Boscardin, G. F. Dalla Betta, S. Ronchin, and N. Zorzi. Development of 3D detectors featuring columnar electrodes of the same doping type. *Nucl. Instr. and Meth. A*, 541(1-2), 441–448, 2005.
- [123] S. Ronchin, M. Boscardin, C. Piemonte, A. Pozza, N. Zorzi, G. F. Dalla Betta, et al. Fabrication of 3D detectors with columnar electrodes of the same doping type. *Nucl. Instr. and Meth. A*, 573(1-2), 224–227, 2007.
- [124] G. Pahn, R. Bates, M. Boscardin, G. F. Dalla Betta, S. Eckert, L. Eklund, et al. First beam test characterisation of a 3D-stc silicon short strip detector. *Submitted to IEEE Transactions on Nuclear Science*, 2009.
- [125] J. Treis, P. Fischer, H. Kruger, L. Klingbeil, T. Lari, and N. Wermes. A modular PC based silicon microstrip beam telescope with high speed data acquisition. *Nucl. Instr. and Meth. A*, 490(1-2), 112–123, 2002.