
Retention of hydrogen isotopes in diamond for applications in fusion energy

A combined computational and experimental study

By

JAMES ALEXANDER PITTARD



Department of Physics
UNIVERSITY OF BRISTOL

A dissertation submitted to the University of Bristol in accordance with the requirements of the degree of DOCTOR OF PHILOSOPHY in the Faculty of Science.

MARCH 2025

Word count: 58342

ABSTRACT

Diamond is a material of interest for various applications within fusion energy, including windows, diagnostics and even as a plasma facing material. However, as a carbon material, concerns remain over the retention of hydrogen isotopes and chemical etching of the surface. As such, this work looks to further develop understanding of the interaction of diamond with low energy hydrogen isotopes. Retention mechanisms were explored for various diamond samples post exposure to deuterium ions. To develop understanding of experimental results, a series of molecular dynamics simulations were carried out in which diamond surfaces and grain boundaries are bombarded with hydrogen isotopes. Diffusion calculations were also performed for both bulk diamond and grain boundaries. From both experimental and computational results, it was concluded that the interaction between hydrogen and diamond is likely to be limited to very shallow depths, leaving bulk properties largely unaffected. Some grain boundaries demonstrated channelling effects leading to anisotropic diffusion but, in contrast to other crystalline materials, presented slower diffusion rates than the bulk. The formation of a disordered top layer during continued bombardment means the presence of grain boundaries are unlikely to significantly impact total retention.

Lastly, a new, low energy ion source was developed. ExTEnD (Exposure to Low Energy Deuterium) uses an electrical discharge to create a plasma from which ions can be extracted via a biased sample stage. The bias applied to the stage dictates the ion energy, whilst the current required to maintain this bias can be used to estimate ion flux. The design, assembly, and testing of ExTEnD is presented, alongside a comparison with equivalent samples tested in an established facility. The increased flux of ExTEnD resulted in a notable increase in retention and loss of clear peaks in desorption spectra, possibly because of ion damage to the surface.

DEDICATION AND ACKNOWLEDGEMENTS

There are a lot of people I wish to thank for making my seven and a half years in Bristol so enjoyable and memorable. Firstly, I want to thank Prof. Neil Fox for his enthusiasm and support throughout this time, what started as a Summer internship progressed to a masters project and finally this PhD, and I am sure this would not have been possible without him. Dr. Mikhail Lavrentiev also supervised me throughout this time, who I would like to thank for his guidance in all things simulation and his patience at my relentless misspelling of ‘tokomak’. Dr. Ant Hollingsworth was also very helpful throughout my masters project and start of this PhD, and offered advice surrounding DELPHI and ExTEnD.

I have no doubt I would not have undertaken this PhD if my undergrad had not been so enjoyable. Callum, Carl and Connor, thank you for the years of meaningless chat and mothers meetings. Cycling has, undeniably, been a huge part of my life during this period, I want to thank the UoB cycling club for the racing, the trips and the mincing around the lanes. In particular, Charlie (despite the broken collarbone), Ted (no I’m not sure which crankset best represents my personality), Carl (for the mega rides, powered by garlic naan) and Joe (a fellow chainy addict).

I would like to thank the diamond group past and present for making the last few years fly by. In particular Ramiz, Liam, Catherine, Dr. James Smith for all his help with ExTEnD, as well as Prof. Paul May for gluing the diamond group together. I also want to thank Joe (different Joe) for fuelling the first two years of this work with his focaccia, Lottie and Ted (same Ted), I will always remember my time in ‘the commune’.

Again I am sure I wouldn’t be writing this if it wasn’t for the years of support and encouragement from my parents and sister. Finally, Liv, for the many happy years together, I couldn’t have done it without you.

AUTHOR'S DECLARATION

I declare that the work in this dissertation was carried out in accordance with the requirements of the University's Regulations and Code of Practice for Research Degree Programmes and that it has not been submitted for any other academic award. Except where indicated by specific reference in the text, the work is the candidate's own work. Work done in collaboration with, or with the assistance of, others, is indicated as such. Any views expressed in the dissertation are those of the author.

SIGNED: DATE:

PUBLICATIONS AND PRESENTATIONS

Publications

1. J.A. Pittard, N.A. Fox, A. Hollingsworth, M.Y. Lavrentiev, A. Wohlers and Y. Zayachuk, *Deuterium retention in CVD diamond: Combined experimental and computational study*, Fusion Engineering and Design, 188:113403, 2023.
Presented in Chapter 3 and 4.
2. J.A. Pittard, M.Y. Lavrentiev and N.A. Fox, *Simulated bombardment of diamond with hydrogen isotopes*, IEEE Transaction of Plasma Science, 52:3674-3679, 2024
Presented in Chapter 4.
3. J.A. Pittard, M.Y. Lavrentiev and N.A. Fox, *Simulated hydrogen diffusion in diamond grain boundaries*, Diamond and Related Materials 149:111665, 2024.
Presented in Chapter 4.
4. J.A. Pittard, J.A. Smith, A. Zafra and N.A. Fox, *Deuterium ion implantation of Eurofer with a new pulsed DC plasma source*, In preparation.
Presented in Chapter 5.
5. J.A. Pittard, M.Y. Lavrentiev and N.A. Fox, *Simulated bombardment of diamond grain boundaries with deuterium*, In preparation.
Presented in Chapter 4.

Conferences

1. 32nd Symposium on Fusion Technology (SOFT), Dubrovnik, Croatia, 2022.
Poster Presentation.
2. 30th IEEE Symposium on Fusion Engineering (SOFE), Oxford, UK, 2023.
Poster Presentation.
3. Universities' Nuclear Technology Forum (UNTF), Bristol, UK, 2023.
Oral Presentation.
4. 21st International Conference on Fusion Reactor Materials (ICFRM), Granada, Spain, 2023.
Poster Presentation.
5. 31st International Conference on Diamond and Carbon Materials (DIAM), Dresden, Germany, 2024.
Oral Presentation.

TABLE OF CONTENTS

	Page
List of Tables	xiii
List of Figures	xv
List of Abbreviations	xxxi
 1 Fusion Energy	 1
1.1 Process	1
1.2 Reactor Design	2
1.2.1 Magnetic Confinement Fusion	2
1.2.2 Inertial Confinement Fusion	4
1.2.3 Magnetised Target Fusion	4
1.3 Heating	5
1.4 Fuel and Tritium Breeding	5
1.5 Plasma Facing Material	6
1.5.1 Requirements	6
1.5.2 Graphite and Carbon Fiber Composites	8
1.5.3 Beryllium	9
1.5.4 Tungsten	9
1.5.5 Lithium	10
1.5.6 Other Considered Materials	10
1.6 A Role for Diamond in Fusion Reactors	10
1.6.1 Plasma Facing Material	11
1.6.2 Windows	15
1.6.3 Permeation barrier	15
1.6.4 Neutron Detectors	16
1.6.5 Tritium Detectors	17
1.7 Summary	18
 2 Methods	 19

TABLE OF CONTENTS

2.1	Diamond Growth Via Chemical Vapour Deposition	19
2.1.1	Growth Conditions	20
2.1.2	Types of Reactors	21
2.1.3	Dopants and Terminations	21
2.2	Deuterium Implantation - DELPHI	22
2.3	Analysis Methods	24
2.3.1	Thermal Desorption Spectroscopy	24
2.3.2	Nuclear Reaction Analysis	25
2.3.3	Elastic Recoil Detection Analysis	27
2.3.4	Raman Spectroscopy	27
2.3.5	Electron Microscopy	27
2.3.6	Secondary Ion Mass Spectroscopy	29
2.4	Molecular Dynamics	30
2.4.1	General Concept	30
2.4.2	Potentials (Force Fields)	30
2.4.3	Time Stepping	31
2.4.4	Thermostats	33
2.4.5	Boundary Conditions	34
3	Exposure of Diamond Samples to Deuterium Ions	35
3.1	Samples, Preparation and Methodology	35
3.1.1	Polycrystalline Diamond	37
3.1.2	Single Crystal Diamond	38
3.1.3	Boron Doped Diamond	38
3.1.4	Ion Damaged Diamond	40
3.2	Deuterium Retention in Diamond	46
3.2.1	Pristine Undoped Diamond	46
3.2.2	Pristine Boron Doped Diamond	52
3.3	Chemical Etching	55
3.4	Conclusions	60
3.5	Future Work	61
4	Simulation of Hydrogen's Interaction With Diamond	63
4.1	Testing	64
4.1.1	Temperature Control	65
4.1.2	Size	67
4.1.3	Flux	68
4.1.4	Timestep	68
4.1.5	Potentials	69

4.1.6	Orientated Diamond and Grain Boundaries	70
4.2	Continuous Bombardment of Diamond with Deuterium	74
4.2.1	Method	74
4.2.2	Results and Discussion	77
4.2.3	Summary	81
4.3	Repeated Single Bombardments of Diamond with Hydrogen Isotopes	81
4.3.1	Method	81
4.3.2	Results and Discussion	84
4.3.3	Summary	93
4.4	Hydrogen Diffusion within Diamond Grain boundaries	94
4.4.1	Method	95
4.4.2	Results and Discussion	97
4.4.3	Summary	106
4.5	Hydrogen Bombardment of Diamond Grain Boundaries	108
4.5.1	Method	108
4.5.2	Results and Discussion	110
4.5.3	Summary	125
4.6	Conclusions	127
4.7	Future Work	128
5	Design, Assembly and Testing of a New Low Energy Ion Source	131
5.1	DELPHI	132
5.2	Context and Aims	132
5.2.1	Design Considerations	134
5.3	Design	138
5.3.1	Electrodes	141
5.3.2	Sample Stage	143
5.3.3	Gas Supply	146
5.4	Assembly	146
5.4.1	Electrode Alignment	146
5.4.2	Initial Assembly and Pump Down	148
5.4.3	Hydrogen Fill Test	149
5.4.4	Power Supply and Interlocks	149
5.4.5	Sample Stage Mounting	150
5.4.6	Quartz Tube	151
5.5	Testing	152
5.5.1	Plasma Form	154
5.5.2	Power Supply Variables	154
5.5.3	Sample Stage Variables	161

TABLE OF CONTENTS

5.5.4	Fluence Estimation	170
5.5.5	Temperature Measurement	171
5.6	Preliminary Retention Study Using ExTEnD	172
5.6.1	Samples	172
5.6.2	Conditions	173
5.6.3	Exposure	173
5.6.4	TDS Results	176
5.7	Conclusions	179
5.8	Future Work	179
6	Conclusions	181
	Bibliography	185

LIST OF TABLES

TABLE	Page
3.1 Summary of samples discussed within this chapter. PC - Polycrystalline, E6 - Acquired from Element 6, TM100 - Element 6 product code ((100) orientation), BDD - Boron doped diamond either high, low or residual (Res.) boron concentration in gas phase. TDS - Thermal Desorption Spectroscopy, SEM - Scanning Electron Microscopy pre or post exposure in DELPHI, RS - Raman Spectroscopy. Anode voltage refers to the anode in DELPHI and corresponds to the energy of the incident deuterium ions in eV. Delay between exposure and retention measurement is indicated in brackets when applicable.	36
3.2 Summary of growth conditions of boron doped samples, grown in a hot filament CVD reactor. All samples were grown on a molybdenum substrate seeded via manual abrasion. Descriptions ('Lightly', 'Heavily' etc.) refer to the B:C gas phase ratio, whilst boron concentration [B] refers to SIMS measurement of samples. *Value from a different sample grown under equivalent conditions.	38
3.3 Parameters used for SRIM calculations.	43
3.4 Calibrated total counts from TDS data of six PC diamond samples implanted with deuterium at various energies. Note values marked * were heated to a lower max temperature, giving lower total count values.	49
3.5 NRA results for different diamond samples exposed to deuterium ions. DELPHI voltage refers to the potential difference used to accelerate deuterium ions. SC - single crystal, PC - polycrystalline.	50
3.6 Calibrated total counts from TDS data of PC diamond samples of different gas phase boron concentrations implanted with deuterium at 400 V.	52
3.7 Summary of deuterium retention of a variety of diamond samples exposed to deuterium ions accelerated across a given anode voltage in DELPHI. Retention values determined via TDS separates retention of HD and D ₂ , whereas retention determined via NRA only includes total deuterium. Samples included undoped polycrystalline (PC) and single crystal (SC) samples provided by Element 6 (E6), as well as boron doped diamond grown in the University of Bristol Diamond Lab. Face perpendicular to ion implantation is shown via Miller indices.	60

LIST OF TABLES

4.1	Activation energies and diffusion coefficients as determined in Fig. 4.25.	103
5.1	The range of variables used in testing of the new implanter. Combinations of the conditions below were selected to explore trends.	152
5.2	Standard operating conditions for ExTEnD.	169
5.3	Measured variables for three doped samples (2-5b, 2-6a and 2-17a) and an undoped (3-1b) sample exposed in ExTEnD. For the undoped sample, a gradual increase in stage bias was required to reach the -150 V that the majority of the exposure was performed at. Values in brackets indicate the value when the bias was at -150 V, whilst values not in brackets give values across the entire exposure.	175
5.4	Comparison of calibrated total counts from TDS data for diamond samples exposed to deuterium ions in two experimental setups (DELPHI and ExTEnD).	177

LIST OF FIGURES

FIGURE	Page
1.1 Main components of ITER.	3
2.1 Depiction of the chemical vapour deposition process for diamond synthesis.	20
2.2 A - Simplified cross section of apertures in a three grid ion optic setup. B - Example schematic of a three grid ion optic setup. This design is in reference for an ion thruster, which is used in small space craft.	23
2.3 Diagram of two common ion beam techniques used for materials composition and depth profiling. In Rutherford Backscattering (RBS), the back scattered primary ion (red arrow) is detected, whilst elastic recoil detection (ERDA) uses the energy of the recoiled target ion (blue circle/arrow). E_0 and m_0 indicate primary ion energy and mass respectively, RBS requires a m_0 less than the mass of the target atom m_t , whilst ERDA requires the opposite. Energy losses ΔE_1 and ΔE_2 correspond to ions moving through the material before and after the scattering event respectively. The changes in energy during scattering events, ΔE_S or ΔE_R , are a function of m_t , $f(m_t)$	26
2.4 Simplified diagram of the underlying mechanisms of Raman spectroscopy. Green arrows represent absorption of incident light whilst other arrows show the energy change during different scattering events.	28
2.5 Graphic depicting the different interactions used for various techniques in electron microscopy. Red and blue arrows represent electrons from the primary beam and the sample respectively.	29
2.6 Simplified diagram of periodic boundary conditions for a simulation cell that is too small (A) and sufficient size (B). Light blue atoms are exerting a force, and dark blue atoms are receiving a force. Black boxes indicate the boundaries of the simulation cell, with faded boxes showing the effect of periodic boundary conditions. When too small, two atoms can interact twice (green arrows), or an atom can interact with itself (red arrows). This does not occur with the larger system size.	34

LIST OF FIGURES

3.1	Offset Raman spectra for three CVD diamond films (samples 2-3a, 2-11a, 2-4b top to bottom) grown with various flow rates of diborane. Values in the legend refer to the B:C value of the growth gas mixture. Residual doping has no additional boron in the gas mixture, only the residual boron within the chamber. Sharp peak at 1332 cm^{-1} shows the presence of diamond whereas the broader peak around 1200 cm^{-1} is attributed to disorder as a result of boron dopants.	41
3.2	Secondary electron images of three boron doped diamond films. A - Residually doped (sample 2-8b), 15 kV, 5 nA probe current, 10.2 mm working distance, $\times 10000$. B - Lightly doped (sample 2-11a), 15 kV, 5 nA probe current, 9.3 mm working distance, $\times 10000$. C - Highly doped (sample 2-3a), 15 kV, 5 nA probe current, 10.1 mm working distance, $\times 10000$	41
3.3	Results from TRIM calculations for damage sustained to diamond along $\langle 110 \rangle$	42
3.4	Raman spectra taken with a green laser (514 nm) of single crystal (110) diamond samples exposed to 100 keV carbon ions at various doses.	44
3.5	Ion irradiated single crystal diamond samples. The four samples (3-1a, 3-2a, 3-3a and 3-4a left to right) were exposed to 100 keV carbon ions at various fluences to achieve the estimated doses of 1, 0.1, 0.01, 0 dpa respectively. Marks were used to distinguish between irradiated and unirradiated sides.	45
3.6	SEM images taken of a SC (110) diamond sample post irradiation to 100 keV carbon ions at a fluence of $5.2 \times 10^{16}\text{ C}^+\text{ cm}^{-2}$, achieving an estimated dose of 1 dpa within the top $0.2\text{ }\mu\text{m}$ (sample 3-1a). A - Raster marks on surface. 15 kV, 5 nA probe current, 19 mm working distance, $\times 20000$. B - 15 kV 30 nA probe current, 10 mm working distance, $\times 5000$. C - 15 kV, 30 nA probe current, 10 mm working distance, $\times 200$. D - Image C rotated 90° to ensure raster lines were physical. 15 kV, 30 nA probe current, 10 mm working distance $\times 2000$	45
3.7	Counts of masses 4, 3 and 2 (attributed to D_2 , HD, H_2 respectively), taken from six PC CVD diamond samples post deuterium implantation at various energies and one reference sample without implantation. The sample stage was heated at a ramp rate of 10 K min^{-1} from room temperature to 1273 K at which it was held for an hour. The 400 eV sample (blue) required a protective aluminium nitride layer, resulting in a reduced maximum temperature and the need for a temperature correction.	47
3.8	TDS spectra of mass 2, attributed to H_2 counts, for five PC CVD diamond samples post deuterium implantation at various energies (as indicated in the legend). Samples were heated at a ramp rate of 10 K min^{-1} from room temperature to 1273 K at which it was held for an hour. A typical temperature profile (brown dashed line) is also included.	48

- 3.9 Thermal desorption spectroscopy spectra for masses 2, 3 and 4 attributed to H₂, HD and D₂ respectively, for three diamond samples (2-3b, 2-11a and 2-4b are high [B], low [B] and residual [B] respectively) grown with different levels of gas phase boron concentrations and a undoped reference sample. A ramp rate of $\beta = 10 \text{ K min}^{-1}$ was used. Where there are large discrepancies in counts two axes have been used, coloured arrows indicate the axes each spectrum corresponds to. 53
- 3.10 SEM images taken of a PC (100) diamond sample post exposure to deuterium plasma (sample 1-1e). **A** - 15 kV secondary electron image of nanoparticles only present post exposure. 5 nA probe current, 9.8 mm working distance, $\times 5000$. **B** - 15 kV backscattered electron images of A. 30 nA probe current, 9.8 mm working distance, $\times 5000$. **C** - 20 kV secondary electron image of a nanoparticles at high magnification. 10 nA probe current, 7 mm working distance, $\times 100000$. **D** - 15 kV secondary electron image showing shadow of cross used to secure sample, nanoparticles were not present within the cross which gave decreased conductivity. 30 nA probe current, 9.8 mm working distance, $\times 250$ 56
- 3.11 SEM images taken of a lightly boron doped diamond film post exposure to deuterium plasma (sample 2-13a). **A** - Nanoparticles present outside of shadowed area post exposure, 15 kV, 5 nA probe current, 9.9 mm working distance, $\times 10000$. **B** - No nanoparticles present in shadowed area, 15 kV, 30 nA probe current, 9.9 mm working distance, $\times 5000$. **C** - Shadow of cross from wire used to secure sample, 15 kV, 30 nA probe current, 9.9 mm working distance, $\times 250$ 57
- 3.12 EM images taken of PC (100) diamond samples post exposure to deuterium plasma. **A** - Sample 1-1a, **B** - Sample 1-1b, **C** - Sample 1-1c, **D** - Sample 1-1d, **E** - Sample 1-1e, **F** - Sample 1-1g, **G** - Sample 1-2a, **H** - Sample 1-2b, **I** - Sample 1-2c. All images taken with a 15 keV beam, $\times 10000$ objective and a working distance of approximately 10 mm. All but subfigures H and I used a probe current of 5 nA, H and I used 30 nA. . . 58
- 4.1 Total number of carbon atoms and temperature from a test simulation of a $6 \times 6 \times 8$ unit cell block of diamond bombarded with 1000, 30 eV tritium atoms. Target temperature for this simulation was 300 K, but a combination of an overly large temperature damping parameter and additional energy in the system meant the temperature varied wildly, resulting in significant loss of carbon atoms. Carbon atoms below a cut off of 2 Å from the surface were counted, increases in the number of carbon atoms are a result of atoms relaxing back below this cut off. 66

- 4.2 Figure demonstrating two methods (**A** and **B**) of giving a bombarding atom (blue, of mass m_B) a velocity, v_B , corresponding to an energy of E_B . In **A**, a bombarding atom is created in a random position above the diamond (red atoms) and any atoms within this region are given a velocity. This can work as intended as seen in **Ai**. However, when reflected or etched atoms are also found within the box (as seen in **Aii**), this method results in an excessive change in system energy, ΔE . In **B**, the new atom is of a different atom type (green), allowing it to be grouped in isolation. The atom within this group can then be changed to the desired bombarding atom type and given the desired velocity. 67
- 4.3 Potential energy of, and force experienced by, a hydrogen placed at a varying separation, R , from a CH_3 molecules to create a methane molecule on minimisation using the AIREBO potential. 69
- 4.4 Simplified depictions of pure twist (left) and a pure tilt (right) grain boundaries. In a twist boundary, the surface vector of the grain boundary plane, \hat{n}_B , is parallel to the axis of rotation, \hat{t}_R , whereas, in a twist boundary, \hat{n}_B is perpendicular to \hat{t}_R . A tilt-twist boundary is a combination of the two rotations. 70
- 4.5 Snapshots taken from two simulations used to determine diffusion coefficients within 221 $\Sigma 9$ diamond grain boundaries. Red atoms are carbon and yellow atoms are hydrogen atoms which had been added to the grain boundary. Snapshots were taken after a run of 50000 steps (28.25 ps) at 3000 K. The top grain boundary (**A**) was produced by replicating out orientated unit cells, ensuring periodic boundary conditions could be applied effectively. The bottom grain boundary (**B**) was created by rotating and trimming a larger diamond block, causing issues with atom alignment across boundaries of the simulation box and resulting in a disordered structure. 71
- 4.6 Carbon-carbon bond length distributions for grain boundaries made by replicating a oriented unit cell (left, referred to as a ‘replicated’ datafile) and grain boundaries made by trimming rotated diamond blocks (right, referred to as a ‘trimmed’ datafile). ‘GB’ distributions groups atoms within $\pm 2.5 \text{ \AA}$ of the centre of a grain boundary, ‘bulk’ atoms are any atoms outside of this range. Blue line indicates the bond length determined for a perfect diamond structure of 1.545 \AA 72

- 4.7 Output file from partway through a LAMMPS simulation of diamond (red atoms) bombarded with deuterium atoms (blue atoms) with atomic density with respect to the depth overlaid on top. At this point in the simulation, the diamond has been bombarded with 480 out of a total 4000 deuterium atoms. Atom densities were calculated by counting atoms within a half unit cell thick slice taken in z (green shaded region) every 0.1 Å throughout the simulation system. Atoms in a unit cell thick slice of the non-bombarded surface were frozen in place (blue shaded region). Atom densities of carbon (black) and deuterium (blue) are presented. The surface was taken to be the solid red line and the dashed line was used as an ion cut-off, beyond which ions were not considered to be implanted. 75
- 4.8 Results from LAMMPS simulations of 30 eV and 100 eV deuterium bombardments on diamond. Central figure shows the heights of the diamond surfaces throughout the simulations. The six surrounding plots show atom density of carbon and deuterium atoms against the z axis at three points in the simulations. The two central density plots show the etch point – where the maximum swelling of the surface has occurred, and etching is about to begin. The definition of the surface (solid red vertical line) was based on the point where carbon atom density dropped below 50 % (solid green line) of pristine diamond. Dashed red line shows the ion cut-off, beyond which, ions were not considered to be implanted. 76
- 4.9 Heights of diamond surfaces bombarded with deuterium atoms at set energies simulated in LAMMPS. The definition of surface height was based on the point where the atom density with respect to the z axis dropped below 50 % of pristine diamond. Negative values show swelling of the surface from the initial height (0 Å) whereas positive gradients show etching. 77
- 4.10 Sputtering yield calculated by two methods from LAMMPS simulations of diamond bombarded with deuterium atoms at a set energy. The volume etched method was calculated by taking a linear fit of surface height with respect to fluence once the etching phase had begun and assumed an atom occupies an average volume of an eighth of a unit cell. Errors were taken to be the error on the linear fit. The final count method took the ratio of implanted to incident deuterium atoms. Linear fit of volume removed method: Gradient = $0.0083 \pm 0.0005 \text{ eV}^{-1}$, y-intercept = -0.07 ± 0.01 . Linear fit of final count method: Gradient = 0.0078 ± 0.0003 , y-intercept = -0.02 ± 0.02 78
- 4.11 Deuterium depth profiles from five LAMMPS simulations of perfect diamond bombarded with deuterium atoms at a set energy. The diamond had a 8×8 unit cell surface area and was bombarded with a fluence of $4.91 \times 10^{20} \text{ D m}^{-2}$ deuterium atoms. The original diamond surface was at 0 Å. 79

- 4.12 Path of incident atom (blue) at an incident angle, θ , towards a $5 \times 5 \times 4$ unit cell (UC) diamond block (red atoms). φ is the angle of implantation about the z -axis. A target location (x_t, y_t) is generated at random within a 3×3 central region (black dashed line). Initial position (x_i, y_i) is calculated so the hydrogen atom hits the target position when created at a height of 3 UCs above the surface. 82
- 4.13 Figure demonstrating the Wigner-Seitz vacancy analysis used to calculate vacancies and interstitials for a displaced configuration compared to a reference. Wigner-Seitz cells are constructed around each atomic point in a reference structure. The number of atoms from the displaced structure in each cell are counted (known as the occupancy). If no changes have occurred all cells will have an occupancy of one, cells with an occupancy of two or zero were classed as interstitials and vacancies respectively. . . . 83
- 4.14 Percentage of simulation repeats which resulted in the reflection of the incident atom. A single hydrogen atom was incident in each simulation, for a given set of conditions, a simulation was repeated 300 times. Incident angle, energy and mass of incident atoms were varied, as well as surface orientation and diamond temperature. Units of vertical momentum, P_z , have been selected for ease of conversion to energy, E , and mass, m , via $P_z = \sqrt{2Em} \cos(\theta)$ where θ is the incident angle. Data has been fitted with $y = y_0 + Ae^{-x/x_0}$, where $y_0 = 0.498 \pm 2.664 \%$, $A = 107.9 \pm 3.9 \%$, $x_0 = 9.4 \pm 0.8 (eV u)^{1/2}$. 84
- 4.15 Defect analysis of four datasets of 300 repeats of a (100) diamond block at 300 K being bombarded with single deuterium atoms of 30, 60, 80 and 100 eV normal to the surface. y -axes units are as a percentage of incident atoms (for example, 39 % of incident atoms caused a vacancy in the top atomic layer for 100 eV). Interstitials are the depth of any carbon atom with an occupancy of two, meaning each vacant site will result in two interstitials provided no carbon is lost. The percentage was halved to account for this. Dashed lines show the atomic layers of the diamond, with markers at every unit cell. Three depth categories ('Surface', 'Shallow', and 'Deep') were used for discussion and are marked on the deuterium density plot. 85
- 4.16 Fraction of hydrogen that caused one or more vacancies compared to all implanted hydrogen as a function of depth for all simulations of a (100) surface. Blue line shows the density of hydrogen depths, the red line also shows this but only for hydrogen that resulted in one or more vacancies forming. The black bars show the ratio between two densities for the hydrogen peaks. 86
- 4.17 Percentage of simulation repeats which resulted in the formation of one or more vacancies for different diamond temperatures (top) and incident hydrogen isotope masses (bottom). A single hydrogen isotope was incident in each simulation, for a given set of conditions, a simulation was repeated 300 times. All simulations performed on a (100) surface with a 0° incident angle. Deuterium (^2H) was used for temperature variation, whilst mass variation used a diamond temperature of 300 K. 87

-
- 4.18 Average depth for a series of repeated, single bombardment simulations of hydrogen isotopes of varying mass and energy at an incident angle of 0° on a (100) diamond surface at 300 K. Linear fits have been applied to the logarithmic plot giving gradients of ^1H : 2.08 ± 0.09 , ^2H : 2.15 ± 0.05 , ^3H : 1.95 ± 0.10 , ^{10}H : 1.55 ± 0.03 89
- 4.19 Results from repeated simulations of deuterium bombarding various diamond surfaces at 300 K. Where only a surface is specified in legend, bombarding atoms were at an incident angle of 0° . Directions indicated correspond to implantation angles upon a (100) surface. (110), (111), and (100) trimmed were all constructed by trimming a orientated diamond block, whereas (100) surfaces were made by replicating a diamond unit cell. The depth plotted takes into account the incident angle, θ , effectively plotting the distance travelled within the diamond rather than depth, z_{final} 90
- 4.20 Heat map of xy hydrogen positions with respect to position in the unit cell (UC) at three different depth regions following repeated simulations of hydrogen bombardment. Values on colour scale are arbitrary but indicate the density of hydrogen. Data was taken across all simulations of (100) surface orientations. Hydrogen with a final depth, z , of $-2 < z < 0$ Å was classified as surface hydrogen (**A**), $0 < z < 1.784$ Å as shallow (**B**) and $z > 1.784$ Å as deep (**C**). Red markers represent carbon atom positions. Where appropriate, different markers have been used to indicate the different atomic layers (AL) of the diamond ('X' - first AL, '+' - second AL, 'O' - third AL, 'Δ' - forth AL). 91
- 4.21 Atomic structures of four tilt grain boundaries – $111 \Sigma 3$, $221 \Sigma 9$, $114 \Sigma 9$ and $112 \Sigma 3$ (**A**, **B**, **C** and **D** respectively). All structures pictured contain no hydrogen atoms and are orientated with $\langle 110 \rangle$ into the page (y -axis). In the plots below, the red line shows the potential energy along the x -axis of hydrogen placed randomly within the structure and minimised. The black line shows the final positions of these hydrogen atoms, this has been normalised by the total hydrogen added across the whole structure (not just ± 20 Å from the grain boundary as shown here). The same scale has been used on all axes to aid comparison. Transmission electron microscopy (TEM) images of the replicated structure are also included 98
- 4.22 Raw mean square displacement (MSD) values for hydrogen diffusing in different diamond structures at a temperature of 2750 K. Miller indices in the legend indicate orientation of faces meeting in the grain boundary (GB) along x . The perfect diamond contained no grain boundary. All structures were orientated with $\langle 110 \rangle$ direction aligned with y . Diffusion coefficients were taken as the gradient of a linear fit on total MSD divided by 6 (Eq.4.5). 99

- 4.23 Heat maps of the potential energy of hydrogen placed in random positions and minimised for four diamond tilt grain boundaries. The top row of plots shows the grain boundary as if looking along the z direction, whereas the bottom along y . The same scale has been used on all the heat maps to ensure they are comparable to one another. The z direction corresponds to $\langle 112 \rangle$, $\langle 114 \rangle$, $\langle 221 \rangle$ and $\langle 111 \rangle$ for the 111, 221, 114 and 112 GBs respectively, whereas all y directions were aligned with $\langle 110 \rangle$ 100
- 4.24 sp^3 and sp^2 content across four tilt grain boundaries in diamond (111 $\Sigma 3$, 221 $\Sigma 9$, 114 $\Sigma 9$, 112 $\Sigma 3$) post energy minimisation. Coordination numbers (CN) of 3 and 4 were classed as sp^2 and sp^3 respectively and taken as a percentage of all atoms in that position. Gaps in data are a result of atomic layers leaving no carbon atoms in these positions. No sp^1 (CN 2) was found in any of the grain boundaries. 101
- 4.25 Diffusion coefficients as a function of inverse temperature for various diamond structures. Miller indices in the legend give the direction of the tilt grain boundary modelled, data from these simulations have been plotted alongside values for perfect diamond ('Perf.') which contained no grain boundary (simulating bulk diamond) and experimental values ('Exp.') of PC diamond. The log of the diffusion coefficients have been fitted with a linear fit to give activation energies, E_A , in Table 4.1. 102
- 4.26 Histograms comparing the bond lengths and angles for carbon carbon bonds within a supercell containing the specified grain boundary post energy minimisation. 'CN 3/4' refers to the coordination number of the carbon atom. Perfect diamond effectively gave single values (of 1.545 Å and 109.5°) as seen for 111 $\Sigma 3$ 105
- 4.27 Change in carbon bond length and angle from the inclusion of hydrogen atoms into the grain boundary post energy minimisation. Data shown only includes the carbon atoms closest to each hydrogen atom and only carbon atoms which went from a coordination number of 3 to 4 when the hydrogen was added. 106
- 4.28 Percentage of atoms deemed stable (not disordered) for the 221 GB bombarded with 30 eV deuterium along $\langle 110 \rangle$. The border of the disordered region and the point taken to be the surface are indicated by red and black dashed line respectively. The disorder border was defined as the highest z value where more than 50 % of atoms were disordered. As a result of how this was calculated, a random distribution of atoms would give approximately 40 % stable atoms. This has been overlayed on a snapshot of the simulation at the same timestep. 109

4.29	Defect analysis of 300 repeated simulations for five different diamond structures bombarded with single 60 eV deuterium atoms. All atoms were incident normal to the (110) surface, the bottom four structures contained a tilt grain boundary (GB) in the middle and at the edge of the simulation cell. The GB plane is indicated by the Miller indices in the legend. When a GB was present, only this region was bombarded. Carbon and deuterium atoms are given in red and green respectively, vacant sites (blue) and resulting carbon interstitials (black) were determined by a Wigner-Seitz analysis.	111
4.30	Defect analysis of four datasets of 300 repeats of a (110) diamond containing a grain boundary as indicated in the legend. The diamond was held at 300 K and bombarded with a single 30 eV deuterium atom normal to the surface. This process was repeated 300 times. y-axes units are as a percentage of incident atoms (for example, 40 % of incident atoms caused a vacancy in the second atomic layer for the 114 GB). Interstitials are the depth of any carbon atom with an occupancy of two, meaning each vacant site will result in two interstitials provided no carbon is lost. The percentage was halved to account for this. Dashed lines show the atomic layers of the diamond. Inset shows interstitials for the 111 GB.	112
4.31	Results from repeated simulations of single deuterium atoms bombarding various diamond structures at 300 K at an incident angle of 0° . Miller indices indicate the diamond surface, if present, the grain boundary plain is indicated, when not present legend states perfect ('Perf.'). 'Repl.' and 'Trim.' stand for replicated and trimmed which refers to the method used to create the structure.	113
4.32	Depth of 30 eV deuterium atoms incident normal to diamond (110) surfaces. Perfect diamond contained no grain boundary whilst the other structures contained a grain boundary as indicated in the legend. Perfect diamond (no grain boundary) Simulations consisted of single bombardments which were repeated 300 times.	114
4.33	Percentage of 300 repeated simulations of a deuterium atom of various energy incident normal to a (110) diamond surface containing various grain boundaries. Reflected atoms were classed as reflected if they were lost from the simulation box or more than 2 Å from the surface and atoms within 2 Å of the surface were classed as surface bound. Total is the sum of the two percentages and indicates the number of atoms which did not penetrate the surface.	115
4.34	Diamond surface height for various diamond structures bombarded with deuterium of different energies. Surface begins at 0 Å, negative values indicate swelling of the surface, whereas positive values of the surface being eroded. The surface was defined based on a fraction of initial atomic density.	117

4.35	Percentage of disordered atoms as a function of depth for the 221 Σ 9 grain boundary bombarded with a fixed fluence of deuterium at various energies. Atoms were classed as either bulk or GB atoms based of their position prior to implantation. Due to the method used to determine disorder, a completely random structure would result in 60.4 % disordered atoms, values consistently above this are likely to indicate a shift of atomic layers in z	118
4.36	Retention within different structures bombarded with deuterium of various energy. Retention values are given by the number of retained atoms divided by the exposed surface area, resulting in units of \AA^{-2} . Due to varying surface areas of structure cells fluence did vary small amounts between structures but was approximately 2.5 \AA^{-2} . The top plot shows retention of deuterium within $\pm 5 \text{ \AA}$ of a grain boundary centre, the second plot anything outside of this, and the final plot is the ratio between the two.	119
4.37	Deuterium depth profiles of various diamond structures exposed to deuterium of varying energy along the $\langle 110 \rangle$ direction. Retention values are given by the number of retained atoms within a slice in z (Depth) divided by the volume.	120
4.38	Comparison of two fluences for simulations of 30 eV deuterium bombardment of a (110) diamond surface containing a 221 Σ 9 grain boundary.	121
4.39	Deuterium density along the x axis for the 221 Σ 9 grain boundary bombarded by deuterium at a fluence of 2.47 \AA^{-2} and varying energy. Deuterium density values include all deuterium within the simulation box so includes both surface bound and retained deuterium.	122
4.40	The depth of the disordered region of diamond structures bombarded with deuterium at different energies. The disordered depth was defined as the highest depth at which more than 50 % of atoms were disordered (based off of a deviation of greater than 0.25 \AA from atomic planes in z).	123
4.41	Lost carbon atoms for simulations of diamond structures containing various grain boundaries bombarded with deuterium. Legend indicates the grain boundary present. GB/Bulk atoms lost per incident atom is similar to sputtering yield, but atoms can move into between regions as well as being sputtered so this is atoms lost from the region. Values for a prefect structure containing no grain boundary have been included for comparison for all plots.	125
5.1	Paschen curves for H_2 and D_2 obtained using voltage ramp (VR) and low current limit (LCL) methods.	135

5.2	An example of an output from an IBSimu simulation of an aperture which could be used in ion optics. Blue blocks represent the grids used in ion optics, cylindrical symmetry conditions have been applied, meaning these blocks effectively become the material surrounding an aperture. For this simulation, the screen grid (far left) has a voltage of 620 V, the extraction grid (middle) is at -410 V and the deceleration grid (right) is grounded. Equipotential lines (green) are shown, and colour bar gives plasma density.	137
5.3	Diagram of ExTEnD with key parts indicated, as viewed from the back of the implanter. A Penning gauge and Baratron were used for low pressure and operating pressure measurements respectively. The outlet valve separates the main chamber from turbo and scroll vacuum pumps. Additional valves and gauges not shown were present to control the flow of gas into the chamber. The high voltage supply for the powered electrode, was bolted directly onto an external face (labelled 'High Voltage Mount Point'), hence a PTFE cover was required for safety. The sample stage is biased, (supplied through the 'Sample Stage Feedthrough'). Sample stage is shown with grounded grid in place above the sample, this can be removed.	139
5.4	Cross section of mounted electrode assemblies.	141
5.5	A - Components for electrode mounting, including the two tungsten electrodes (i), collets and mounting screws (ii), the powered electrode mount (iii) and the grounded electrode mount (iv). B - Tungsten electrode within collet. C - Electrode and collet secured into the grounded electrode mount using the access hole and 5 mm hex key.	141
5.6	A - Powered electrode assembly. B - Grounded electrode assembly (i) with enlarged image showing the spacers used to maintain gas flow (ii).	142
5.7	A - The open end of the PTFE cover with securing bolts which overlap with the flange to prevent the cover from being removed unintentionally. B - The closed end of the cover with cable gland for high voltage supply to the electrode. The cable gland pictured was later replaced with a larger, M25, gland.	142
5.8	Cross section of sample stage assembly.	144
5.9	Assembly of the sample stage. A - Sample stage mount secured onto 2.75" blanking flange. B - Base plate with locking nut added. C - Insulating block, sample plate, and sample window added, full assembly without grounded grid. D - Full assembly including grounded grid.	145
5.10	Sample stage components. i - Insulating block, machined from MACOR. ii - Sample plate, 11×11 mm inset region for sample. Window plates for securing sample within sample plate, 8×8 mm ² exposure area with (iii) and without (iv) additional holes for grounded grid mounting. v - Insulating rods for grounded grid. vi - Grounded grid plate.	145

LIST OF FIGURES

5.11	A - Electrode misalignment due to the port used for the powered electrode (right) being off angle, view from top (i) and side (ii) . B - Realigned electrodes from top (i) and side (ii) . C - Shims used on the base of the powered electrode mount to correct port misalignment.	147
5.12	Piping and instrument diagram of the implanter setup. Valves are labelled 'V1' to 'V8', pressure gauges 'G1' to 'G4', and vacuum pumps 'P1' and 'P2'. V7 is a three-way valve which has two positions 'A' and 'B', gas flow of these positions are indicated below the V7 label. P1 was a turbo pump and P2 a dry scroll pump. G1 - Analogue gauge to measure gas cylinder pressure, G2 - Baratron (1 – 100 Torr) to monitor pressure during operation, G3 - Penning gauge for low pressure readings when pumping down, G4 - Pirani gauge to monitor scroll pump when pumping air.	148
5.13	High voltage power cable connected to the high voltage electrode through a cable gland in the PTFE cover. Blue wire is grounded and connected directly to the body of the chamber. The other end of the wire has the high current high voltage MIL-spec plug used to connect to the PDR power supply.	150
5.14	Sample stage assembly during mounting to 2.75 " tee. Kapton coated wire connected to the top of the sample stage supplied bias and measure stage current. Wire was connected to a vacuum feedthrough mounted to the port on the left of the tee.	151
5.15	Flared end of the quartz tube, required to prevent tube gradually sliding out of position.	152
5.16	Full assembly as used for first plasma. Valve and gauge labels (V1-V8 and G2-G4 respectively) refer to Fig. 5.12. Sample stage assembly had not yet been fitted.	153
5.17	Image of an early hydrogen test plasma in ExTEnd, the pink colour is typical of a low pressure hydrogen plasma and is a result of the red and blue emissions of the Balmer series. Pressure: 0.1 Torr, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Electrode-stage separation: 35 mm, Stage bias: 0 V.	155
5.18	ExTEnd hydrogen plasma at 0.1, 0.5, 0.75, 1, 1.5, 2, 3 Torr (A–G respectively). i - Primary discharge between electrodes, from the main viewport. ii - Secondary discharge within chamber, from main viewport. iii - View from secondary viewport. Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Electrode-stage separation: 35 mm, Stage bias: 0 V for A–F, -50 V for G.	156
5.19	Time averaged DC voltage supplied to the powered electrode to maintain hydrogen plasma at different pressures. Multiplying time averaged voltages by 1.43 gives the pulsed voltage. Electrode separation: 25 mm, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s.	157
5.20	Stage current for varying plasma power. The top plot shows various bias voltages for a fixed pressure of 1 Torr, the bottom plot shows various pressures for a fixed bias of -250 V. Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Electrode-stage separation: 35 mm.	158

- 5.21 Stage current as a function of percentage of pulse time period for which the pulse was on. The pulse off time was varied between 1.5 – 4.5 μs in 0.5 μs increments, whilst the 100 kHz pulse frequency gave a time period of 10 μs . A linear fit has been applied giving a gradient of $-0.072 \pm 0.002 \text{ mA}$, an intercept of $0.13 \pm 0.2 \text{ mA}$ and an adjusted R^2 of 0.98. Pressure: 0.1 Torr, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μs , Electrode-stage separation: 65 mm, Stage Bias: -300 V. 159
- 5.22 Stage current measured across a 25 hour period. Time of top plot is relative to the start of the first measurement, whereas bottom left and right plots are relative to measurements 3 and 4 respectively. Measurements 1, 2 and 5 were taken under equivalent conditions (50 W, -250 V). For measurement 3 bias was varied with a fixed 50 W power, whilst measurement 4 varied power for a fixed -250 V bias. Various gaps were left between measurements. Pressure: 1 Torr, Pulse frequency: 100 kHz, Pulse off time: 3 μs , Electrode-stage separation: 35 mm. 160
- 5.23 Stage current against pressure for varying stage biases. Values were collected for high pressures and 0 V, but many were positive and could not be shown on the log scale due to the change in polarity. Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μs , Electrode-stage separation: 35 mm. 162
- 5.24 Stage current for increasing stage bias. Includes curves for various different pressures as indicated in the legend. Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μs , Electrode-stage separation: 35 mm. 162
- 5.25 **A** - Desired operation, plasma discharge between electrodes (0 V stage bias), measured current corresponds to incident ions. **B** - Plasma discharge between electrodes and from the stage to the grounded electrode (-700 V stage bias). **C** - Plasma discharge from stage to grounded electrode with electrode power supply off (-700 V stage bias). Pressure: 1 Torr, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μs , Electrode-stage separation: 35 mm. 164
- 5.26 Simplified diagram of different sources of stage current, black arrows indicate the flow of electrons. **A** - Desired operation, discharge only between electrodes, measured current indicates incident ions. **B** - Discharge between electrodes and from the stage to the grounded electrode, measured current is the sum of ions incident on the surface and electrons leaving the stage. **C** - Discharge only present between grounded electrode and stage, measured current indicates electron emission from the stage. 165

5.27	Current measurements taken to measure breakdown voltage between the stage and grounded electrode. The figure on the left shows current with respect to time during this test. Initial spikes were a result of electron flow from the stage to the grounded electrode with the electrode power supply off. Following this, discrete steps in current are a result of changes in stage bias. Sudden drops in current are when the electrode power supply was turned off and on again. The figure on the right gives the measured stage current with the electrode power supply on and off, as well as the difference between them, for various stage biases. Pressure: 1 Torr, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Electrode-stage separation: 65 mm.	165
5.28	Breakdown voltage between the sample stage and grounded electrode for varying pressure (indicated by different symbols) and stage-electrode separation (indicated by different colours). For $p \leq 2$ Torr, the breakdown voltage, V_{BR} , was taken to be the highest bias at which the stage current returned to 0 mA when the plasma was turned off. For $p \geq 5$ Torr, no negative current reading was made prior to arcing, and V_{BR} was taken to be the lowest voltage required to trip the power supply (indicating a stage current of > 10 mA). Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s. . . .	167
5.29	The difference between stage current with the stage current with the plasma off (left) and the electrode plasma on and off (right) and for varying stage biases. The stage-electrode separation and resulting pL value is indicated in the legend. Pressure: 0.1 Torr, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s.	168
5.30	The shortest stage-electrode separation of 25 mm (A) distorts the form of the plasma, whereas a separation of 35 mm (B) does not. Pressure: 0.1 Torr, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Stage bias: -400 V.	168
5.31	A - Thermocouple isolated with Kapton tape attached to the sample stage prior to exposure. B - Temperature measurement during exposure. C - Thermocouple piercing the Kapton tape, taken post exposure.	171
5.32	Sample stage temperature during exposure to plasma and when cooling down. The plasma and stage bias were both on for the first 30 mins, before being turned off and left to cool for time > 30 mins.	172
5.33	Images of two samples during and post exposure. A - Undoped (i) and doped (ii) samples (3-1b and 2-6a respectively) post exposure with window plate removed. B - Undoped sample during exposure with no arc (i) and an arc between the sample and grounded electrode (ii). C - Optical images of the undoped sample surface post exposure, at $\times 5$ (i) and $\times 50$ (ii) magnification.	174

5.34 Comparison of TDS spectra for samples grown under equivalent conditions and exposed to deuterium ions in either DELPHI or ExTEnD. Counts have been scaled by both exposure area and a calibration factor to allow results from two different TDS facilities to be compared. Exposure conditions and retention values can be seen in Tables 5.3 and 5.4 respectively.	176
---	-----

LIST OF ABBREVIATIONS

AIREBO	Adaptive Intermolecular Reactive Empirical Bond Order
AL	Atomic Layer
BCA	Binary Collision Approximation
BDD	Boron Doped Diamond
BSE Imaging	BackScattered Electron Imaging
CN	Coordination Number
CVD	Chemical Vapour Deposition
DC	Direct Current
DELPHI	Device for Exposure to Low Energy Hydrogen Isotopes
DEMO	Demonstration Power Plant
dpa	Displacements per atom
E6	Element 6
EDX	Energy-dispersive X-ray spectroscopy
ELM	Edge-Localised Mode
ExTEnd	Exposure To low Energy Deuterium
GB	Grain Boundary
HF	Hot Filament
HTS	High Temperature Superconductor
HPHT	High Pressure High Temperature
IBA	Ion Beam Analysis
JET	Joint European Torus
LAMMPS	Large-scale Atomic/ Molecular Massively Parallel Simulator
MAST	Mega Ampere Spherical Tokamak
MCD	Microcrystalline Diamond
MCF	Magnetic Confinement Fusion
MD	Molecular Dynamics
MFC	Mass Flow Controller
MLG	Multi-Layer graphene
MS	Mass Spectrometry
MSD	Mean Square Displacement

LIST OF ABBREVIATIONS


MTF	Magnetised Target Fusion
NCD	Nanocrystalline Diamond
NEXAFS	Near Edge X-ray Absorption Fine Structures
NRA	Nuclear Reaction Analysis
NVT	Canonical ensemble
OVITO	Open Visualization Tool
PC	Polycrystalline
PDR	Pulsed Deposition Reactor
PFM	Plasma Facing Material
PTFE	Polytetrafluoroethylene
RBS	Rutherford Backscatter
REBO	Reactive Empirical Bond Order
RMSD	Root Mean Square Displacement
RS	Raman Spectroscopy
SC	Single Crystal
SEM	Scanning Electron Microscopy
SIMS	Secondary Ion Mass Spectrometry
SRIM	Stopping and Range of Ions in Matter
STEP	Spherical Tokamak for Energy Production
TDS	Thermal Desorption Spectroscopy
TEM	Transmission Electron Microscopy
TPD	Thermal Programmed Desorption
TRiCEM	Tritium Retention in Controlled and Evolving Microstructure
TRIM	Transport of Ions in Matter
UC	Unit Cell
XDS	X-ray Diffraction Spectroscopy
XPS	X-ray Photoelectron Spectroscopy
Z	Atomic Number

FUSION ENERGY


Fusion energy has long been promised as the ultimate source of clean energy, offering zero carbon emissions, no long-lived radioactive waste and no reliance on external factors such as weather. It is hoped this can be achieved by replicating the process that occurs within the sun - the nuclear fusion of hydrogen isotopes. Achieving this has proved incredibly challenging, with slow but steady progress over the last 40 years. This chapter will outline the general principles of fusion reactors, design considerations and progress to date with a particular emphasis on the use of diamond within these reactors.

1.1 Process

Generally, nuclear fusion is the combination of two light nuclei to form a heavier nucleus. If the resultant products have a lower mass than the two constituent nuclei, the mass difference is released as energy in accordance with Einstein's famous relation. Typically, lower Z nuclei are easier to fuse due to the smaller Coulombic repulsion, and additional neutrons of heavier isotopes make the nucleus less stable and again, easier to fuse. Because of this, most fusion reactors rely on the fusion of the hydrogen isotopes deuterium and tritium. Although other reactants have been explored [1], at conceivable energies, deuterium-tritium (DT) fusion has a cross section approximately 100 times bigger than other alternative reactions [2]. The underlying reaction is



Although the conditions used maximise the cross section of Eq.1.1, there is no way of restricting the reaction processes taking place to just Eq.1.1 and other reactions also take place.





The high energy neutrons produced in these reactions will result in the heating of surrounding materials which can be used to produce electricity.

In order to overcome the Coulombic repulsion between the two ions, incredibly high temperatures are required and the reaction occurs within a plasma - a partially ionised gas. Most of the design considerations of a reactor are related: achieving and maintaining this required temperature. A tokamak is a classic reactor design which uses magnetic fields to manipulate the plasma into a torus shape. An example of this design includes the Joint European Torus (JET) - a successful experimental reactor used for DD and DT fusion before beginning its decommissioning phase at the start of 2024. Despite the success of JET, larger tokamaks are required if positive energy output from a self-sufficient (or burning) plasma is to be achieved. ITER is an international, public funded, attempt to demonstrate both positive energy output and a burning plasma, alongside testing of other, vital, integrated technologies. Following this, there are plans to build a demonstration reactor (DEMO), which will produce power to the grid. The future of DEMO is under question due to notable delays and unforeseen costs surrounding ITER which is still under construction. With growing interest and need for clean energy sources, a number of private companies [3] have also been exploring alternative, cheaper designs in attempt to produce reactors quicker than public efforts. Discussion within this introduction will generally refer to DT fusion within tokamak designs.

1.2 Reactor Design

1.2.1 Magnetic Confinement Fusion

As the temperature of the plasma (of order 10^9 K) is far higher than any material can sustain, magnetic confinement is used to suspend the plasma and minimise energy transfer to the first wall. This confinement can be achieved with a toroidal magnetic field geometry in either tokamak or stellarator reactor designs [5]. Magnetic confinement fusion (MCF) contains a low density plasma in a relatively large volume and uses a variety of different heating methods to achieve the target temperature. These reactors have been the focus of most publicly funded research through JET and ITER, as well as some private companies.

The plasma is confined using large, coiled electromagnets. Although JET used conventional copper coil magnets, future reactors are to use superconductors to reach the very strong magnetic field required for larger tokamaks (ITER peak magnetic field is 11.8 T). Moving to superconducting magnets reduces the energy required to reach this field strength but presents its own challenges. Conventional superconductors (such as niobium-tin and niobium-titanium) require a temperature




Figure 1.1: Main components of ITER [4].

of around 4 K so must be cooled with liquid helium. This requirement means the reactor must be housed in an ultra-high vacuum cryostat as well as the need for an on-site cryoplant. A promising avenue is the development of high temperature superconductors (HTS). These materials offer a fairly modest but important easing of low temperature requirements (operating at around 24 K) and could help provide an increased strength of magnetic field. However, there are concerns over degradation of performance under neutron irradiation, with critical temperatures of HTS candidates decreasing with radiation dose [6].

A diagram highlighting the main components of ITER can be seen in Fig. 1.1. The central solenoid and poloidal field coils surround the vacuum vessel where the plasma is contained. Due to the low operating temperatures of these superconducting magnets, the reactor is housed within a cryostat. Lining the inside of the vacuum vessel is a protective layer known as the plasma facing material (PFM), which acts to protect other components from the incredibly harsh conditions within the reactor. The vessel consists of two main regions, the blanket and the divertor. The blanket is the upper portion of the first wall and faces less extreme conditions compared to the divertor at the bottom, which must withstand much higher thermal loads and ion fluxes. The blanket is also responsible for regenerating tritium through a reaction between lithium and neutrons (this is known as the breeding blanket and will be situated behind the PFM). Port plugs are used for plasma diagnostics and heating systems.

The success of traditional torus tokamaks (such as JET, ITER and DEMO) has long been a challenge of scale - for net energy output there is effectively a minimum reactor size. Building bigger reactors not only comes with considerable design challenges but massively increases the

cost. Some private fusion companies (such as Commonwealth Fusion Systems [7, 8] and Tokamak Energy [9, 10]) hope to exploit the higher fields possible with HTS to create a more compact reactor in order to reduce overall cost and increase commercial viability. Some public research bodies are also exploring more compact reactor designs using HTS. UKAEA's latest tokamak will be the Spherical Tokamak for Energy Production (STEP [11, 12]) and will use a spherical geometry in contrast to the D shaped coils of JET and ITER.

A stellarator is another type of magnetically confined fusion reactor. The shape of a tokamak means coils are more tightly packed on the inside of the ring than the outside, resulting in an uneven field and challenges surrounding plasma stability. The basic principle of a stellarator is to add twists to the magnetic field to force ions to alternate between the inside and outside of the plasma (between higher and lower coil density regions) in an attempt to maintain a consistent field. This leads to a far more complex magnetic field [13] which is challenging to create, but avoids the internal currents present in tokamak designs leading to a more stable plasma. Although internal currents lead to plasma instabilities, in a tokamak, they are used to heat the plasma. This is not an option in a stellarator making reaching required temperature more challenging.

1.2.2 Inertial Confinement Fusion

Inertial confinement fusion (ICF) relies on the rapid compression of a DT fuel pellet, which in turn results in densities and temperatures sufficient for self-heated fusion to occur. The compression can be done with lasers or a pulsed magnetic field, but its pulsed nature would mean many of these reactions would be required to approximate a continuous power output [15]. The National Ignition Facility (NIF) at the Lawrence Livermore National Laboratory [16] is perhaps the most well-known example of this type of reactor, and in December 2022 achieved a positive energy output (ignition). Although this is the first fusion technique to reach this milestone, there are concerns on how it would be possible to scale this up to a commercial reactor, alongside continuous scientific challenges surrounding the reliable ignition of the pellets [15].

1.2.3 Magnetised Target Fusion

Magnetised target fusion (MTF) is a combination of both MCF and ICF in which magnetic fields are used to contain a low density plasma which is then compressed (although to a lesser degree than in ICF). General Fusion is a private company attempting to create reactors using this method [17]. In this novel approach, a spinning cylindrical chamber is surrounded by liquid lithium held in place by the centrifugal force. Benefits of liquid lithium as a PFM are its ability to self-heal due to lack of rigid crystalline structure and breed tritium without the requirement for an additional breeding blanket. The DT plasma is confined magnetically in the chamber, and steam pistons are used to compress the liquid lithium wall and the plasma to reach the desired fusion conditions. In a similar vein to ICF, this has the potential to be significantly cheaper to

produce positive energy than MCF, but there are concerns over its pulsed nature and ability to scale up to reach commercially viable energy outputs.

1.3 Heating

In tokamak designs, the plasma itself has a finite resistance and the current induced by magnetic fields results in ohmic heating of the plasma [18]. Unlike most cases of ohmic heating, as the temperature of the plasma increases, the resistance decreases - limiting the temperatures that can be reached in this manner to around one tenth of what is required. Stronger magnetic fields from superconducting magnets can be used to increase temperature that can be achieved with ohmic heating, although concerns remain over operating temperatures and irradiation damage. Furthermore, in some reactors such as stellarators, the lack of internal plasma current means ohmic heating is not possible at all. Therefore, other external heating methods are also necessary.

Radio waves and microwaves are used to enhance the ohmic heating, adding velocity and chaotic motion to ions and electrons (respectively) within the plasma [19, 20]. Three different frequencies will be used in ITER to match the resonant frequencies of various charged species present within the plasma. Neutral beam heating is also used [21], in which a high energy neutral beam of deuterium atoms is injected into the plasma, transferring energy to ions within the plasma through collisions. Some heating is also provided by energetic ions produced in the fusion reaction. Neutrons produced will have limited heating effect on the plasma as, unlike ions, they are not confined by the magnetic field and are lost from the plasma. In future reactors, it is hoped that a large proportion of the plasma heating can be provided by the energy produced from fusion reactions, reducing reliance on external heating and reaching a burning plasma.

1.4 Fuel and Tritium Breeding

Deuterium is naturally abundant, making up approximately 1 in 5000 hydrogen nuclei, and as such is readily available. Tritium however, is radioactive and will decay via beta emission

$$(1.5) \quad {}^3_1\text{T} \rightarrow {}^3_2\text{He} + \beta^- + \bar{\nu}_e.$$

With a half-life of 12.32 years, this decay makes tritium scarce naturally (with an abundance of 10^{-18}), and challenging to source in suitable quantities, with the production and acquisition of tritium being highly monitored. The limited tritium inventory means tritium wastage must be minimised in order to maximise what is possible with the allocated amount. Components can absorb and retain tritium, so material selection is key to reduced wastage. Alongside reducing waste, DT reactors will be required to reproduce tritium during operation. Despite some tritium being produced through Eq.1.2, a dedicated breeding blanket is needed to produce the levels of tritium required (DEMO would be expected to use around 300 g of tritium per day of operation

[22]). The breeding blanket is typically situated behind the PFM, and uses the reaction between lithium and neutrons to produce tritium



or



Currently, lithium ceramics are favoured for tritium breeding in ITER, whilst other reactors are considering liquid lithium [23]. Beryllium is also included in most blanket designs to act as a neutron multiplier and increase the yield of tritium per incident neutron. No large scale breeding blanket has been attempted in current reactors, but ITER has set out to test different tritium breeding systems.

Despite the relatively small amounts of fuel required, one of the major challenges facing is sourcing tritium. Tritium is considered a necessity for most fusion reactor designs and is currently supplied as a by-product from CANDU-type fission reactors. These reactors do not produce sufficient quantities for commercial reactors such as DEMO and are due to be phased out over the next few decades. Despite plans for reactors to be self-sustaining in terms of tritium inventory, the global tritium inventory could be too low to meet the start-up inventories required [24, 25].

1.5 Plasma Facing Material

1.5.1 Requirements

The PFM is the material that lines the inside of the first wall and divertor of a reactor. A PFMs primary role is to protect more delicate components of the reactor from high thermal loads and low energy ion irradiation. Such materials will also need to endure high fluences of fast neutrons but typically will do little to protect other components from these.

There are several key considerations when selecting a PFM:

1. **Interaction with hydrogen isotopes** - How a PFM interacts with hydrogen ions will influence the retention and transport of tritium, cooling and contamination of the plasma, and changes to the material properties. The high temperature of a fusion plasma results in high ion energies within the plasma, but stray ions will lose much of this energy escaping the magnetic confinement, resulting in a low energy (10s of eV [26]) interaction with the PFM. Some ions will simply be reflected from the surface back into the plasma. Ions returning to the plasma after an interaction with the PFM (recycled ions) will have experienced an energy loss, and as such contribute to cooling of the plasma. Depending on the material, hydrogen ions may chemically etch the surface resulting in further cooling of the plasma as

well as contamination. Contaminants will influence the behaviour of the plasma and can be redeposited elsewhere in the reactor, typically exhibiting very high retention. The low ion energy means penetration depths are limited. The shallowness of this hydrogen results in a negligible contribution to the retained inventory. However, hydrogen is highly permeable in many materials, and the elevated temperature of the first wall means the diffusion and transport of tritium becomes a concern.

Generally, materials which are reactive with hydrogen (such as carbon materials) are more likely to exhibit chemical etching and resultant redeposition of etched atoms. In this case, reactions are likely to be restricted to the surface, and minimal changes to bulk properties would be expected. This is in contrast to most metals, which display a limited reactivity with hydrogen, but can result in higher levels of recycled hydrogen and increased diffusion. Diffusion of tritium is of particular concern, but collections of other gaseous molecules can still result in the formation of bubbles within the material changing its bulk properties [30].

2. **Thermal properties** - Quoting an operating temperature for a PFM is often not possible as it depends on many factors - the thermal properties of the materials, the position within the reactor, the inclusion of active cooling and the plasma form itself. Instead heat loads are typically discussed [31], and any PFM must be able to withstand the relevant heat loads. Typically, a high melting point is desirable to avoid dramatic changes in material properties. Temperature can also have more subtle effects on material physical properties such as its ductility and strength which must also be considered.

Alongside the high operating temperature, PFMs will also have to deal with sudden thermal shocks as a result of plasma instabilities such as edge localized modes (ELMs). Instabilities in the plasma can result in large amounts of energy being dumped on the first wall and a rapid increase in temperature. In order to manage these sudden changes, a high thermal conductivity is vital to transport heat away from a localised point. Irradiation damage sustained by a material will typically decrease its thermal conductivity, meaning a materials resistance to such damage must also be considered.

3. **Durability** - Any candidate PFM must demonstrate good resistance to ion and neutron damage to reduce the need for frequent maintenance. Furthermore, damaged materials commonly behave differently, so it must be ensured that any desirable properties are maintained in irradiated materials. Physical sputtering can also be a concern. Similarly to chemical etching, sputtering will result in contamination of the plasma. For metallic PFMs, a higher atomic number usually correlates with reduced sputtering and an increased melting temperature - both desirable properties. However, plasma contamination will result in radiative cooling of plasma, where energy is lost from the plasma ionising sputtered

atoms. Elements with higher atomic numbers require more energy to fully ionise and therefore result in a greater degree of radiative cooling.

4. **Fabrication and disposal** - A consideration often neglected when considering more exotic materials is whether it is possible to manufacture this material to sufficient quality and quantity for use within a fusion reactor. Cost of material and maintenance also comes into this, if fusion energy is ever going to be economically viable as an energy source long term.

Although fusion reactors will not produce the long lived highly active waste that fission reactors produce, it is expected to produce large amounts of intermediate level waste [29]. Tritium retention also presents challenges surrounding the disposal and decommissioning of reactors, and the removal of tritium from materials during decommissioning must also be considered [27, 28].

As fusion reactors continue to evolve into larger, more powerful, reactors, the conditions the PFM is exposed to become harsher. For example, one of the key goals for ITER, is to demonstrate continuous power output. This raises concerns over steady state erosion of the PFM [32–34] which has not needed to be considered for previous, smaller, reactors. Due to the constant development of different reactors, PFMs must and will continue to evolve. The following subsections summarise previously used and candidate PFMs.

1.5.2 Graphite and Carbon Fiber Composites

Graphite combined with carbon fibre composites (CFCs) in select areas were originally deployed in JET due to carbon’s low atomic number and exceptional thermal properties [62]. Graphite tiles were used in JET until 2011 when it was replaced with an ITER-like wall consisting of beryllium and tungsten. Graphite was considered to be quite forgiving as a PFM, as extreme thermal loads result in sublimation rather than melting which would occur with metallic PFMs, and the low atomic number resulted in minimal radiative cooling of the plasma. Until recently, the graphite configuration had yielded the highest energy returns for a fusion reactor (16.1 MW) [35].

When hydrogen ions are incident on a carbon material, a top layer which is completely saturated by hydrogen forms. Due to its thinness, this layer contributes a negligible amount to the retained inventory of deuterium/tritium but will dominate recycling processes [36]. Although the relatively strong sp^2 C-C bonding has a high threshold for physical sputtering, carbon materials are susceptible to chemical etching from hydrogenic species [37] as carbon readily forms bonds with hydrogen. In particular, the sp^2 hybridisation present in graphite displays very high chemical etching rates from hydrogen ions. Etched carbon contaminates the plasma, before being re-deposited elsewhere in the reactor as a carbon dust of unacceptably high retention [38, 63]. Doping can be used to alter graphites properties, typically reducing chemical erosion yield but at the cost of increased retention within the material [39]. Although there is some discussion over the use of graphite in future reactors [40], graphite has largely been discounted as a future

PFM due to the inadmissible retention in re-deposited carbon dust and has raised concerns over carbon materials in general.

1.5.3 Beryllium

Beryllium tiles were due to be deployed as the PFM for the first wall of ITER after testing in JET. Beryllium's low atomic number reduces the radiative cooling effect from sputtered ions and, as a metallic PFM, demonstrates minimal fuel retention and chemical etching from hydrogenic species. Beryllium has a low cohesive energy which makes it susceptible to melting. Although this may be acceptable for smaller test reactors, it remained a concern for future, larger, reactors. It was planned to use beryllium in some regions of the ITER first wall, leading to concerted testing in JET [41]. More recently however, it seems unlikely a beryllium PFM will be used in ITER due to concerns over its high toxicity - leading to additional challenges around assembly, maintenance and disposal. Beryllium is still expected to be present within the breeding blanket, where it will be used a neutron multiplier to increase the neutrons available to interact with the lithium and, therefore, the potential tritium yield.

1.5.4 Tungsten

Tungsten and tungsten alloys are considered the most promising plasma facing material for future and current reactors [42]. It has long been thought that tungsten was to be used in the divertor of current and future reactors leading to considerable testing, but it is also likely to be used in other areas of the first wall as well [43]. Tungsten exhibits a high thermal conductivity, limited physical sputtering and its lack of reactivity with hydrogen results in no chemical etching and reduces retention. The main concerns of tungsten surround fabrication challenges, its high atomic number (leading to increased radiative cooling), brittleness and a low recrystallisation temperature [44, 45]. Not only is tungsten brittle at room temperature but it will undergo irradiation hardening and embrittlement upon exposure to neutrons [46]. Brittleness is an obvious concern as can result in catastrophic material failure such as cracking. This brittleness is exacerbated by tungsten's high ductile brittle transition temperature. Lastly, tungsten's low recrystallisation temperature (the point where a new grain structure is formed, typically resulting in a decrease in strength and hardness of the material) makes it challenging to control grain structure and material properties.

Helium ions produced in the fusion reaction collect within tungsten PFMs forming helium bubbles. These bubbles collecting under the surface resulting in blistering and eventually the formation of tungsten 'fuzz' [47] - ultimately impacting thermal conductivity and mechanical properties of the tungsten [48]. Similarly, hydrogen isotopes have a high mobility in tungsten as with most metals, and can freely permeate to defects where they are trapped and collect to form bubbles.

1.5.5 Lithium

The melting point of lithium (454 K) means a solid lithium PFM would not be viable. Instead, liquid lithium is being considered as a PFM possibility [49]. Use of a liquid PFM has its practical challenges but can be held in place using a capillary system which increases surface tension. A major benefit of a liquid PFM is the surface can self-heal and restore its original form. Furthermore, lithium is required in essentially all first wall designs in order to breed tritium. Typically, it is thought this will be done with ceramic beads situated in the breeding blanket behind the PFM, but use of a lithium PFM could eliminate the need for this additional component. With an atomic number of just three, lithium contaminants in the plasma would result in very little radiative cooling. Lithium is also highly reactive with hydrogen, which effectively eliminates concerns over recycling which should result in a more stable plasma and reduce maintenance. Of course, having a PFM that is highly reactive with hydrogen raises concerns over tritium retention [50].

Although testing has been carried out in tokamaks [51], there are currently no plans to use a liquid lithium wall for ITER or DEMO largely due to the practical challenges of a liquid PFM. However, it is being considered for regions of the divertor in UKAEA's new spherical tokamak STEP and is a key design feature of General Fusion's MTF reactor.

1.5.6 Other Considered Materials

Molybdenum has been deployed in both the Alcator-C tokamak and as the first wall in EAST [52] and performs comparably to tungsten [53]. Additionally, this material is to be used as a mirror in diagnostics systems in ITER and future reactors [54]. Although molybdenum exhibits a better resistance to radiation-induced embrittlement, tungsten still has superior thermal properties, resistance to sputtering [55] and avoids concerns over long term activation present for molybdenum [56].

There are some ceramic alternatives which have also gained some interest. Silicon carbide displays good resistance to neutron damage, retains strength at high temperatures and a low hydrogenic diffusivity [57, 58]. This material has been largely untested to date, due to concerns of carbon redeposition and challenges in fabrication but remains a material of interest for future reactors. Other ceramics such as boron carbide have also been considered but mostly as a protective coating on other PFMs such as tungsten or graphite [59–61].

1.6 A Role for Diamond in Fusion Reactors

Diamond's intrinsic hardness, excellent thermal properties, resistance to radiation damage, semi-conducting properties and transparency make it a material of interest for numerous applications within fusion energy. These applications include plasma facing materials, windows, hydrogen permeation barriers as well as neutron and tritium detectors.

All of these potential applications involve the interaction with hydrogen isotopes on some level. As highlighted in previous sections, tritium retention is of particular concern regarding decommissioning and start-up inventories. As such, the main focus of the work is to explore diamonds interaction with hydrogen isotopes in the context of fusion reactors, with a particular emphasis on retention, diffusion and etching. The majority of techniques discussed in this section are explained in more detail in Chapter 2.

1.6.1 Plasma Facing Material

One of the early PFMs deployed in JET was graphite, which was selected for its excellent thermal shock resistance and thermal conductivity [62] (see 1.5.2). However, concerns over hydrogen retention of redeposited carbon limited this materials future for steady state reactors [63]. Despite being an allotrope of graphite, the sp^3 C-C bonding of diamond demonstrates a much greater resistance to hydrogen etching - this is exploited in chemical vapour deposition (CVD) growth of diamond, in which formation of graphitic phases is suppressed by hydrogen ions. Diamond would also be expected to show reduced fuel retention and similar, if not greater, resistance to physical sputtering. Additionally, diamond possesses many of the favourable attributes of graphite such as excellent thermal management, low Z and sublimation (rather than melting) under extreme thermal loads.

In order to assess different diamond materials suitability for this purpose, diamond's response to fusion relevant conditions needs to be understood. Porro *et al.* conducted a series of experiments [64–67] exploring surface modification and deuterium retention of CVD diamond materials exposed in two tokamak test facilities (MAST and TEXTOR) and a linear plasma device (Pilot-PSI). The work varied boron concentration and grain size in diamond materials, for use as protective coatings or standalone materials. Various methods were used to explore surface modifications, including Raman spectroscopy, scanning electron microscopy (SEM), X-ray photoelectron spectroscopy (XPS) and optical emission spectroscopy. Raman spectroscopy gave insight into sp^2 and sp^3 content of the diamond - finding increased sp^2 in diamond grain boundaries, and an increase in sp^2 content post exposure [64]. Shifts in the 1332 cm^{-1} diamond peak was used to indicate stress within the crystal from lattice mismatch with the substrate, the presence of dopants, and as a result of plasma exposure [65, 66]. Optical emission spectroscopy was used to analyse erosion rates post exposure in Pilot-PSI, with chemical erosion reported to be 50% less than graphite [64], although it was found the erosion rate increased with temperature [66]. When exploring nanocrystalline diamond (NCD) coatings, concerns were raised over the cohesion of the coating to metallic substrates [67], however, the other experiments of microcrystalline diamond (MCD) reported no delamination or dramatic failure [64, 66]. Some amorphous carbon was found on the surface of many of the samples, but graphitic phases were not found in XPS measurements or Raman spectra [64, 66], there was the suggestion that this could be redistributed from elsewhere within the reactor rather than from the samples itself [65]. Evidence of arcing was also

observed in undoped samples, with a notable reduction on the inclusion of boron dopants [65]. Deuterium retention was determined via nuclear reaction analysis (NRA) and increased with boron concentration but decreased with increasing grain size [64]. In both cases, this increase was concluded to be a result of a more disordered structure resulting in a more porous material as seen in boron doped graphite [68]. Etching mechanisms may also play a role in retention in these studies, as etching of the top surface would also result in a loss of retained deuterium in this region. This could contribute to why boron doped diamond presents higher retention but lower etching.

Porro *et al.* concluded that MCD would be more appropriate than NCD and diamond-like carbon (DLC) - the larger grains reducing sp^2 content and in turn reducing etching and deuterium retention. MCD also displayed preferable mechanical properties and generally a higher resistance to plasma exposure than the other diamond materials tested. Lightly doped boron was suggested as the most viable option, reducing damage from arcing through increased conductivity whilst keeping retention values to acceptable levels. Diamond materials showed a notable improvement to graphite.

Boron doped diamond has also been reported to have better thermal shock resistance than tungsten and other PFMs. This was tested by De Temmerman *et al.* [69], who exposed freestanding boron doped diamond to an electron gun to mimic extreme conditions expected from fusion plasma instabilities. Due to the use of an electron gun, boron doping was required to increase the conductivity of the samples and avoid charging. In conditions three times higher than tungsten's melting threshold and where significant erosion of CFC and graphite would be expected, no damage or surface erosion was exhibited. Raman spectroscopy was used prior and post exposure. For most conditions, the diamond peak was still clearly present, but for the more extreme tests the lack of diamond peak indicated the formation of amorphous carbon. In another study [70], De Temmerman *et al.* compared erosion effects in graphite, NCD, MCD and boron doped MCD. Some samples were exposed in the DIII-D tokamak, whilst sputtering yields were determined through experiments in Pilot-PSI and PISCES-B (both experimental plasma devices). Erosion rates of NCD were found to be similar to graphite whereas MCD exhibited about half this value. Erosion is highly dependent on the plasma mode used, in the DIII-D divertor, no measurable erosion was recorded for detached plasma conditions, but significant erosion was seen for an attached ELMy H-mode plasma.

The effect of ion irradiation on deuterium uptake on undoped diamond and graphite was explored by Deslandes *et al.* [71]. In this work, CVD diamond was irradiated with 5 MeV carbon ions to simulate the primary knock-on from 14 MeV neutrons. Ion irradiation was followed by exposure to a deuterium plasma. Post carbon ion irradiation, Raman spectroscopy and Near Edge X-ray Absorption Fine Structures (NEXAFS) results presented an increased sp^2 fraction, broadening of the diamond peak/excitation and additional peaks in the Raman spectra. Once exposed to deuterium plasma, many of the lost diamond features were recovered with a clear

reduction in sp^2 fraction, although the additional peaks seen in the Raman spectra remained. The recovery of the surface could indicate chemical etching of sp^2 phases has occurred. As Raman spectroscopy is a surface technique, it cannot be concluded that ion damage deeper within the sample isn't present. Surface graphitisation is more likely than graphitisation within the bulk due to the inherent instability of a sp^3 surface, so it is likely the increase in sp^2 is only present on the surface. This might explain why the sp^2 peaks were lost (from surface etching of the plasma) but other peaks remained as a result of ion damage deeper within the sample. Similarly to findings with graphite [72, 73], irradiation with carbon ions increased the plasma erosion exhibited by the diamond. When comparing the performance of graphite and diamond, Deslandes *et al.* found irradiated diamond displayed lower deuterium retention than irradiated graphite, although the opposite was true for unirradiated samples.

In a similar study, Deslandes *et al.* also explored the impact of 5 MeV carbon ions up to a dose of 1 dpa at varying diamond temperature [74]. The temperature range tested 300 – 873 K encompasses expected operational conditions of the ITER divertor, but below the point where graphitisation¹ of the surface might be expected (1000 K [75]). Surface modification was observed at all temperatures. Raman spectra from the highest temperature irradiation (873 K) was closest to that of the pristine sample - suggesting higher temperatures resulted in the removal of some unstable defects produced in the irradiation. Similarly, a decrease in sp^2 was also observed at higher temperatures. Shifts in the Raman diamond peak were used to conclude that the higher temperatures resulted in reduced stress in the lattice allowing the diamond structure to be recovered. Although, this reduction in stress could also be a result of the fewer stable defects at this temperature rather than it allowing the diamond surface recovery. Furthermore, the higher temperature may allow implanted carbon ions to diffuse to the surface rather than occupying interstitial positions which would further increase lattice stress. The study highlights the importance of considering probe depth of analysis technique. Both Raman spectroscopy and NEXAFS were used in this study, with the former having a shallower probe depth. At the surface electron stopping will dominate, resulting in reduced damage compared to deeper in the sample where nuclear stopping can occur. The shallower probe depth of Raman probed the lower damage electron stopping region, and therefore showed better diamond structure recovery than the NEXAFS would suggest. Deslandes *et al.* also exposed pristine diamond samples to deuterium ions of varying energy [76]. It was concluded that the exposure removed non-diamond phases from the surface, whilst increasing the energy resulted in an ion damaged layer alongside the removal of non-diamond phases. The energy dependence of etching highlights some of the limitations of experimental plasma exposure setups, which commonly use ion energies an order of magnitude higher than experienced at the first wall of a fusion reactor. In this study no notable deuterium retention was observed although it was acknowledged this could be a result of erosion

¹Graphitisation is a process in which a diamond's sp^3 structure can reconfigure to a graphitic sp^2 structure at high temperatures (above 1000 K). In the context of PFMs, this would be a major concern with such a transition negating any potential benefits of diamond over graphite.

of the region containing deuterium.

Guenette *et al.* [77] exposed CVD diamond samples to deuterium, helium and deuterium-helium plasmas, and evaluated the damage and deuterium retention to the samples. SEM images post exposure to plasmas containing deuterium showed a minimum fluence of around 10^{23} ions m^{-2} was required to exhibit surface modification. With the inclusion of helium, larger features were observed including some cracking of the surface. Small carbon protrusions/nanoparticles were seen on surfaces exposed to plasmas containing deuterium but not the pure helium plasma. This is likely to be evidence of redeposited carbon that had been etched from the surface. The fraction of sp^2 bonding seen on the surface was found using NEXAFS. For deuterium-helium plasma, the fraction of sp^2 increased with increasing fluence. Qualitatively, the number of protrusions seemed to scale with this fraction, suggesting the nature of the protrusions were in fact sp^2 . However, no clear trend between fluence and sp^2 fraction was seen for deuterium plasmas, with the lowest fluence exhibiting the largest sp^2 fraction of 9%. It is likely that the higher fluence resulted in greater etching of sp^2 bonds resulting in an unclear relationship between damage and sp^2 fraction. Overall, modifications to the diamond surface were concluded to be minor. Guenette *et al.* also observed saturation of deuterium in the subsurface of the diamond through Elastic Recoil Detection Analysis (ERDA) measurements. At 18 eV implantation energy used, saturation was seen at a fluence of 5×10^{22} m^{-2} and flux of 10^{21} m^{-2} . Averaged over the top 20 nm, the deuterium saturated at circa 5.4 at. %, a value lower than the predicted value of 26 at. % from modelling [78]. Different fluxes could account for discrepancies, as modelling used a flux eight orders of magnitude higher than Guenette *et al.* Again, it was concluded that the interaction was surface level [77, 79], with minimal changes seen to the bulk, in agreement with other work [71, 74, 76].

Although most research into this area has been experimental, some molecular dynamics (MD) studies performed using LAMMPS (Large-scale Atomic/Molecular Massively Parallel Simulator [80, 81]) have been carried out by Dunn *et al.* [75, 78, 82]. In simulations, diamond of various orientations and temperatures was bombarded with tritium atoms. Retention and erosion mechanisms were explored for various ion energies and fluences, with a particular emphasis was put on the graphitisation transition. The (111) surface is the closest match to the graphite structure, encouraging the graphitisation process and resulting in graphitisation at temperatures as low as 500 K [78] making it unsuitable for divertor PFM applications. For bombardment simulations, retention and etching values of the different surface orientations were indistinguishable [82]. A layer-by-layer erosion mechanism (in agreement with [77]) was observed resulting in minimal impact on the bulk structure, a linear etching rate, and a retention saturation point being reached. Some work exploring the presence of grain boundaries was carried out [78], which suggested transport through the grain boundaries could be possible and would enhance retention.

In summary, experimental work broadly concluded that only surface level damage was exhibited by diamond samples, with bulk properties left largely unchanged. Some etching was

observed from this surface layer, but generally erosion was reported to be significantly less than graphite. Higher quality and larger grain diamond contains less sp^2 carbon which reduced etch rates, suggesting MCD would be preferable over NCD or DLC. The increased conductivity of boron doped diamond reduced damaging from arcing, although the additional defects in the lattices resulted in an increased retention. Computational work carried out largely agrees with experimental work. More work is required to fully understand the interaction of hydrogen with diamond, particularly regarding grain boundaries and boron doping.

1.6.2 Windows

High intensity electromagnetic radiation at various frequencies is to be used to resonate plasma species and increase the plasma temperature [20] (see 1.3). Windows in the tokamak are required to apply such heating, whilst acting as a vacuum and tritium barrier [83]. Single crystal (SC) diamond is transparent to these frequencies, can survive the harsh conditions present in a fusion reactor and has an exceptional thermal conductivity which is beneficial for cooling, making it the material of choice for windows in ITER. Various diagnostics are also to be included in the window assembly to ensure the vacuum and tritium barrier is being maintained with both tritium and arc detection.

1.6.3 Permeation barrier

A key concern surrounding tritium is its ability to permeate through metals [84] making it challenging to contain tritium to one location, particularly at elevated temperatures. This can lead to contamination, or changes to physical properties such as hydrogen embrittlement. To try and counteract this, there has been work developing permeation barriers - a layer of a different material in which diffusion is limited, preventing the escape and aiding the containment of tritium. Typical candidate materials include metals, ceramics or glass [85], in particular ceramic metal oxide and rare earth oxide coatings [86].

Diamond has potential as a permeation barrier. Hydrogen diffuses slowly in diamond, and diamonds excellent thermal properties and resistance to radiation damage offer few limitations regarding location of the barrier. Although CVD is a well-established technique that can grow thin films over a variety of materials for use as a barrier, there are limitations. For example, diamond cannot be directly deposited onto steel without a buffer/intermediate layer (such as Cr-N [87]) due to irons strong affinity for carbon. Carbon materials have had some attention as a permeation barrier for hydrogen [87–89] but there is limited research surrounding diamond.

In order to assess diamonds suitability as a permeation barrier (alongside the other potential uses listed here) diffusion characteristics must be understood. Although there are studies [90–93] which have explored hydrogen diffusion in diamond, the impact of grain boundaries is relatively unexplored. Polycrystalline (PC) diamond is significantly cheaper and simpler to grow, with a

greater variety of viable substrate compared to SC diamond, making understanding the impact of grain boundaries important.

1.6.4 Neutron Detectors

Neutron detection is a key way of assessing the performance of a reactor. 2.5 and 14 MeV neutrons are a product of the DD and DT fusion reactions respectively. The detection of neutrons indicates the reaction is occurring successfully, whilst the neutron energy reflects the distribution of ion velocities within the plasma [94]. Selection of appropriate detection method is dependent on the energy of the neutron - as ultimately any detection method is relying on an interaction between the neutron and detector, and the cross section of this interaction will vary with energy. At low energies (thermal neutrons), neutrons can be absorbed by high cross section materials such as boron, resulting in emission of high energy ions which can then be detected. Moderators are often used to reduce neutron energy to detectable levels. Although detectors such as these are commonplace in the fission industry, detection of the fast neutrons produced in fusion reactors requires specific detectors.

Semiconductor based neutron detectors typically consist of a neutron conversion layer and a separate semiconductor layer [95]. As neutrons are not directly ionising, the neutron conversion layer consists of a thin layer of high cross section elements (such a lithium or boron) which will produce charged particles on the interaction with incident neutrons. In order to make a detection, charged particles will need to travel to the semiconductor layer, creating electron-hole pairs within the semiconductor, which drift towards the collecting electrode on the application of an appropriate bias. Neutron converter layers much be thin. Although thicker layers will increase the opportunity for neutron interaction, the short range of the particle means a thin layer increases the likelihood of detection.

As diamond reacts with neutrons (via $^{12}\text{C}(\text{n},\alpha)^9\text{Be}$) and is a wide band gap semiconductor, it can act as both the neutron converting and semiconducting layer eliminating concerns over charged particles lost in the converting layer. Furthermore, diamonds resistance to irradiation damage would allow it to be deployed in environments too extreme for typical semiconductors and its wide bandgap (5.5 eV compared to 1.14 eV for silicon) offers a very low dark current, allowing operation at high temperatures. However, the ionisation energy required to produce and electron-hole pair in diamond is over four times higher than silicon and produces one quarter of the primary electrons again compared to silicon [95]. As a result of this, the resultant signal in diamond very low compared to typically selected semiconductors.

Natural diamond has been used for these applications, however recent advances have opened opportunities for the use of SC, CVD diamond detectors [94]. Use of CVD also allows the possibility to include dopants or other carbon isotopes to alter the cross section for neutron interaction. Diamond based neutron detectors are due to be deployed in ITER and potentially future reactors as well [96, 97].

1.6.5 Tritium Detectors

Accurate and reliable tritium detection remains a significant challenge for safe tritium handling. As a hydrogen isotope, tritium can replace hydrogen in water and organic molecules which can be easily absorbed by the human body. This makes the detection and monitoring of tritium highly important. However, the energy and nature of tritium's decay (β^- of average energy 5.7 keV [98]) make it particularly challenging to detect with a maximum range of around 6 mm in air [99].

Methods for tritium detection typically rely on inferring its presence through the detection of decay products (namely β^- particles). Currently, there are no methods that can detect tritium in all phases (gas, liquid or trapped within solids). For liquids, liquid scintillation counters (LSC) are well established and accurate. In this method, the test liquid is mixed with a liquid scintillate, which will emit photons on the absorption of the β^- particle. Gas samples can be condensed or bubbled through water before being mixed with a liquid scintillator. This method is commonly used for detection in tritium facilities. Although accurate, there is a significant delay as samples are collected over the course of a week or so before testing.

Ionisation chambers are commonly used to measure activity within a gas. Incident radiation ionises an inert gas within a chamber. A voltage is applied across the gas resulting in a current as a result of gas ionisation. However, these detectors cannot operate under vacuum conditions and are highly sensitive to environmental changes. Moreover, the short penetrating distance of the β^- makes detection via this method challenging.

Solid crystal scintillators are another option. In a similar manner to a liquid scintillator, an appropriate scintillating material will generate photons as a result of incident β^- particles. Such scintillators require a moderate level of counts in order to make accurate measurements, again limiting their applications in fusion reactors, as many scenarios require low activity measurements. Furthermore, they often have a non-linear dependence with temperature, making them inappropriate in regions of significant temperature change.

Lastly, there are solid state detectors. These are semi-conducting materials in which incident radiation results in the promotion of electrons from the valance to conduction band - inducing a current. There are advantages to solid state detectors, as a solid, they can operate under vacuum conditions unlike gas counters, and their higher density allows for smaller volume detectors with the same sensitivity. Diamond has great potential solid state detector [83, 100]. Diamond possesses several advantages over more traditional semiconductors as discussed regarding neutron detectors - its resistance to harsh conditions and wide bandgap would allow it to operate at high temperatures without the need for cooling [83].

Retention and diffusion of tritium within a detecting material is an important consideration from multiple standpoints. As always, uptake is a concern from an inventory point of view - losses must be minimised to reduce startup inventory requirements and tritium retention is a concern regarding decommissioning. This is especially important for a detector, as tritium uptake may result in counts from retained tritium rather than a true measurement of tritium within the

environment.

1.7 Summary

The breadth and variety of different fusion reactor designs is significant. The vast majority of work to date has been surrounding D-shaped tokamaks with progressively larger tokamaks planned to reach positive energy output. Significant challenges and huge associated costs of these reactors have generated interest in smaller, compact reactors which make use of the high magnetic field possible with HTS - the route taken by many private companies due to the potential reduced cost and development time. Ultimately, there are very few things in the fusion landscape which are certain aside from the need for fusion reactors to supply carbon free energy without long-lived radioactive isotopes. As such, continuous research into different materials is required, and it is important to have as much information on as many potential materials as possible going forward.



There are various potential applications for diamond within fusion reactors, all of which involve an interaction with deuterium and tritium to some degree. This makes it important to understand diamonds interaction with hydrogenic species in terms of etching, retention and diffusion. Accepted applications of diamond include windows and neutron detectors. There is also a good amount of literature exploring diamond as a PFM, which demonstrates some of diamond's exceptional properties and highlights its potential. The work presented in this thesis will look to develop understanding of diamonds interaction with hydrogen isotopes through both experimental and computational work. Following this summary of the fusion landscape and literature, Chapter 2 will outline various relevant experimental methods referred to throughout this thesis, as well as an overview of general concepts of MD. Chapter 3 will present initial experimental work, exploring retention within undoped and boron doped diamond through TDS and NRA. Computational results are given in Chapter 4, which explored retention and diffusion of hydrogen isotopes within bulk diamond and physically accurate diamond grain boundaries. Lastly, Chapter 5 will outline the design, assembly and testing of a new setup for exposing samples to low energy deuterium ions. This includes results from a preliminary retention study, designed to compare the new setup to an established experimental facility used in Chapter 3.

This chapter outlines the operation and applications of a range of experimental techniques relevant to this research. This includes diamond growth, ion exposure, methods to determine hydrogen retention and methods to explore structural changes. Regarding computational work, a general discussion of MD methods, and underlying theory, is presented.

2.1 Diamond Growth Via Chemical Vapour Deposition

Initial attempts to form diamond synthetically looked to replicate the conditions in which natural diamond forms, known as HPHT diamond (high pressure high temperature). This method created small, high quality diamond but was not easily scaled up. Following this, CVD was developed, which allowed for larger areas of diamond to be grown making it ideal for coatings and films. CVD involves the atom-wise addition of carbon onto a substrate to gradually grow a film from a gas phase, the structure of which will be dictated by the growth conditions. To allow molecules to bind to the substrate, the precursor gas mixture must first be activated as shown in Fig.2.1. Varying the contents and ratios of the precursor gas mixture, as well as the pressure, temperature and substrate all impact the deposited film [102].

Only two gases are required for diamond growth: hydrogen and methane. These can be activated to give atomic hydrogen and methyl radicals



Carbon atoms are then able to bind onto the surface of an appropriate substrate, allowing the gradual growth of the film. Activation of precursor gases can be achieved in a variety of ways, including using hot filaments, microwaves or electrical discharges.




Figure 2.1: Depiction of the chemical vapour deposition process for diamond synthesis.

2.1.1 Growth Conditions

CVD reactors require a gas mixture of methane and hydrogen. Methane typically only makes up around 1% of the gas mixture, as the high hydrogen content is used to preferentially etch graphitic, sp^2 phases. Although diamond phases will also be etched by hydrogen, this happens at a slower rate than the etching of graphitic phases, permitting the slow growth (typically $0.1\text{--}10\text{ }\mu\text{m hr}^{-1}$ depending on the method [103, 104]) of a diamond film. Increases in methane concentration will result in faster film growth due to the greater availability of carbon atoms, but at the cost of film quality. High methane concentrations are often used for nanocrystalline and diamond-like carbon film growth [107]. Pressure also plays a key role in determining the quality of the deposited diamond. Typically, pressures of order 10 Torr are used [108], with pressures significantly higher than this increasing growth the rate but at the cost of diamond quality.

The substrate temperature must be between 1000 – 1400 K for diamond growth [105], with higher temperatures increasing the growth rate but, once again, decreasing the quality. This places restrictions on the substrates that can be used, as the melting point must be higher than this under operational pressures. The formation of a surface carbide is also useful to ensure adhesion to the substrate. Silicon, molybdenum and tungsten are all commonly used due to their ability to form a carbide whilst avoiding excessive carbon solubility [105]. Substrates must be seeded with diamond to act as the foundations for diamond growth. This can be done in a number of different ways each with its own benefits and limitations. For example, nanodiamonds can be ground into the surface (known as seeding by manual abrasion) which offers good adhesion to the substrate but scratches the substrate surface resulting in a rougher film. As an alternative, a

nanodiamond suspension can be accelerated across a potential difference onto the substrate (in a technique known as electrospray), offering a very even coverage but at the cost of adhesion. Seeding is required for non-diamond substrates, but is not for diamond substrates - commonly used for adding high quality doped layers for devices.

The length of the growth run has obvious implications to the thickness of the film. However, the film thickness can also impact the surface morphology. CVD has a columnar grow mechanism in which growth occurs from nucleation points, and diamond phases grow out as well as up. As a result of this, grains will coalesce during growth, resulting in larger grains present at the growth surface of thicker films.

2.1.2 Types of Reactors

Perhaps the simplest CVD reactor, hot filament (HF) reactors rely on thermal activation of the precursor gases. Filaments made from erosion resistive material such as tantalum are resistively heated by applying a large electrical current. These are positioned directly above the substrates and additional filaments can easily be added to grow films over large areas [109]. The growth rate is relatively slow, and the growth time is limited by the filaments which must be replaced regularly.

Plasma-assisted reactors use a plasma as the source of carbon and hydrogen radicals. This can be done with a direct current (DC) plasma, in which an electrical discharge is used to create free electrons and ions. In this technique, the substrate is placed on a grounded anode, whilst electrons are supplied via a powered cathode positioned directly above the substrate [110]. Electrons are accelerated between the electrodes and, with a sufficient potential difference, will create ions and more electrons through collisions (known as a Townsend avalanche). This will create a sustained plasma on top of the sample. One of the challenges facing these reactors is arcing¹. To avoid arcing, pulsed power supplies are used which skip pulses if the very high currents indicative of an arc are detected. Microwave sources can also be used to create a plasma, which rely on resonances of electrons to ionise gas molecules. Similarly to DC plasma reactors, the plasma ball is positioned on top of the substrate, which means high-power microwave systems are required for large area films [102]. Microwave reactors are the most common type of CVD reactor, and their ability to tolerate high power inputs and more varied gas pressures results in a faster growth rate than HF reactors.

2.1.3 Dopants and Terminations

Dopants are commonly used in diamond to alter electrical and mechanical properties, and are typically incorporated as a precursor gas. Due to the low concentration of methane in

¹Arcing produces sudden high currents from the thermionic emission of electrons, this differs to a Townsend discharge which is characterised by sustained low currents, and electrons produced via field emission.

the gas mixture, other gases are commonly diluted down in hydrogen to maintain reasonable carbon:dopant ratios whilst avoiding very slow flow rates. Common dopants include:

- **Boron** - Boron is largely used to increase charge carrier availability, giving p-type semi-conducting properties to the diamond [111] for sensor applications. By increasing the conductivity, diamond can also be used as corrosion resistant chemical electrodes [112]. Metal-like conductivity is achievable for very high boron concentrations ($[B] > 10^{21} \text{ cm}^{-3}$) but can also increase non-diamond/ sp^2 content within the film. Moderate levels of boron doping can improve physical properties and decrease the rate of erosion compared to undoped diamond [113]. Boron is typically incorporated into diamond by adding diborane, triethyl borane or trimethyl borane diluted in hydrogen into the gas phase.
- **Nitrogen** - Nitrogen impurities are present in all natural diamond and diamond synthesised via HPHT. It can be added to CVD diamond through the inclusion of ammonia into the gas mixture. Diamonds are commonly categorised based on their nitrogen content (type I - $[N] > 1 \text{ ppm}$, or type II - $[N] < 1 \text{ ppm}$). The defects created by the presence of nitrogen can be used to create fluorescent screens for use in synchrotrons [114].
- **Phosphorus** - Phosphorus is one of the few candidates for n-type doping in diamond [115]. n-type diamond is of particular interest for diamond diodes. Phosphine can be used as a gaseous precursor for microwave CVD.
- **Codoping** - Combinations of dopants can also be used to control electrical and chemical properties. For example, boron-nitrogen [116, 117], boron-oxygen [118] and boron-lithium [119] codoped diamond.

A large variety of different surface terminations are also possible and can also give rise to p-type surface conductivity, as well as insulating behaviour.

2.2 Deuterium Implantation - DELPHI

The Device for Exposure to Low Energy Hydrogen Isotopes (DELPHI, formally TRiCEM [120]) is a facility designed to assess the retention and release of hydrogen isotopes in potential first wall materials. Although many established ion beam facilities can implant hydrogen isotopes at high energies, fusion relevant tests require low energy ions (10^1 – 10^3 eV) at high fluence. Prior to the construction of this facility, there were limited setups capable of these conditions. Moreover, DELPHI was designed to be able to operate with both deuterium and tritium. For initial experiments, only deuterium was used, as this can be distinguished from the naturally abundant hydrogen and negates the necessary safety considerations when using tritium. An upgrade to allow the use of tritium was then attempted but, as a result of numerous issues with the ion source, has not been completed at the time of writing.




Figure 2.2: **A** - Simplified cross section of apertures in a three grid ion optic setup. **B** - Example schematic of a three grid ion optic setup. This design is in reference for an ion thruster, which is used in small space craft. Both figures taken from [121].

One of the main design considerations to allow for use of tritium is the gas circulation system. Typically, similar setups would use a ‘once through’ system, in which gas is continuously flowing in and out of the chamber. Although sufficient for deuterium, use of tritium requires the gas to be recirculated to avoid waste and unnecessary contamination. In ‘recycling’ mode, contaminants are removed from the gas and it is recirculated back into the chamber. Typically, experiments were run in recycling mode to try and ensure results were comparable between deuterium and tritium results.

Ions were generated by a SPECS ions source [122], which consisted of a microwave source, a small plasma chamber with a magnetic quadrupole, and ion optics. To maximise the ion flux incident on the samples a low pressure (around 0.7 mTorr) is used, minimising the scattering and neutralising of ions as they travel from the plasma to the sample. Due to the low pressures, a magnetic quadrupole surrounding the plasma chamber is used to increase plasma density and localise it to the plasma chamber. Ions are extracted via ion optics - a series of charged grids commonly used for initial ion extraction in ion accelerators and ion thrusters [121, 123]. Ion optics consist of 2-3 of biased grids (see Fig. 2.2). The screening grid (closest to the plasma) helps to contain the plasma, and will be positively biased for a plasma consisting of positively charged ions (such as a hydrogen plasma). The second grid (the acceleration/extraction grid) is negatively charged, and is responsible for the ion extraction. The aperture size of the extraction grid is smaller than the aperture of the screening grid (see Fig. 2.2A), ensuring some ions are still extracted from the plasma. A third deceleration grid can also be included to prevent damage from back streaming ions, this was grounded in the DELPHI system but can be positively biased. The energy of the extracted ions is dictated by the potential difference between the first and last grid. For DELPHI, this is the difference between the anode/screening grid and the grounded deceleration grid, meaning the ion energy is given by the anode bias. Typically, experiments in

DELPHI were carried out with an ion energy of 10^2 – 10^3 eV. The grid thicknesses, aperture size, separation and grid biases all impact the form, dispersion and ion current of the resulting ion beam.

The sample stage is located directly below the ion beam and is both thermally and electrically isolated from the chamber. A copper block directly connected below the sample stage contains a cooling channel as well as a 200 W heater, allowing the sample to be heated or actively cooled during operation. For exposures at ambient temperature and above, the cooling channel is filled with air, whilst a Hexid coolant is used for temperatures down to 4 °C. With no coolant and the heater on, a maximum temperature of 450 °C is achievable. Stage current is measured and linearly related to the sample current for fluence estimations. The stage current is measured in two places, the primary measurement is taken from an area approximately 1 cm outside of the perimeter of the sample and is used for fluence estimations. A secondary current measurement is taken for the area a further 1 cm outside of the primary current measurement area. This secondary region is used to give an indication of beam diversion.

Samples are loaded via a transfer arm, perpendicular to the direction of travel of the ion beam. Use of a transfer arm ensures the main chamber can remain under vacuum, reducing pump down time as well as minimising the potential risk of tritium contamination. The sample is secured onto a metal tab via a wire cross, this setup allows for some movement but is only intended to prevent the sample falling off during transfer. The metal tab clicks in place on top of the sample stage, and rotating the transfer arm detaches it from the arm.

2.3 Analysis Methods

2.3.1 Thermal Desorption Spectroscopy

Thermal desorption spectroscopy (TDS) or thermal programmed desorption (TPD) [124, 125] is a method of exploring the thermal desorption of molecules from a surface. A sample is heated at a linear rate commonly referred to as the ramp rate, β , so that the temperature, T , at a given time, t , can be given by

$$(2.1) \quad T(t) = T_0 + \beta t,$$

where T_0 is the initial temperature. Molecules desorbed from the surface are identified by their mass using a mass spectrometer (MS). A MS discriminates between different ions mass-to-charge ratio (m/z), allowing for the detected species to be inferred. Different types of MS are available, but the TDS used in this work featured a quadrupole mass analyser, which uses oscillating electric fields to determine an ions m/z . Other MS techniques use the speed of accelerating ions (time of flight MS) or the deflection of ions in a magnetic field to separate ions spatially (sector MS). The MS can be calibrated to give quantitative counts. This is typically done with a leak test calibration, in which a flow of known amount of the gas in question is put into the chamber.

Other, less accurate, calibrations are also possible, in which a sample of known concentration is measured.

Additional information on binding energy can be determined through repeating measurements with various β [126]. The desorption rate can be described by a version of the Arrhenius equation - the Polanyi-Wigner equation, which, when combined with Eq.2.1 gives

$$(2.2) \quad -\frac{d\Theta}{dT} = \frac{\Theta^n v_0}{\beta} e^{-E_{des}/k_B T}.$$

Here, Θ is the coverage of adsorbates in monolayer units and n is the reaction order of the process and contains information of the desorption mechanism. E_{des} and v_0 are the desorption barriers and attempt frequency respectively, which can be used to determine the desorption rate of each molecule. Although this higher level analysis can be used to gain quantitative description of desorption, it typically requires repeated measurements of equivalent samples and a good understanding of the adsorbent coverage, therefore limiting the number of different samples it would be possible to analyse.

2.3.2 Nuclear Reaction Analysis

NRA is an ion beam analysis (IBA) technique that can be used to infer the presence of hydrogen and deuterium [127]. As the name would suggest, the technique relies on a nuclear reaction occurring between an incident ion beam and the element to be measured. Ion beams of different energies and species can be used to allow detection and increase the sensitivity for a particular element. The technique is able to give concentration as a function of depth within the top few microns of the surface [128].

For hydrogen detection, NRA typically relies on the resonant reaction with nitrogen [129],



Both the γ -rays and α -particles can be detected to infer the concentration of hydrogen, although the former is typically preferred. A resonance refers to a sudden sharp peak in cross section for a reaction at a particular ion energy. For ion energies equal to the resonant energy, the cross section becomes much higher and a sudden increase in emitted γ -rays and α -particles is observed. Depending on the width of the resonance, a small increase or decrease in ion energy will see a drop in cross section, as well as the number of emitted γ -rays and α -particles. With the ion beam energy equal to the resonant energy, detected γ -rays will be emitted from the reaction with hydrogen on the surface. Increasing the ion beam energy beyond the resonance energy reduces the cross section, so surface hydrogen is no longer detected. However, incident ions will lose this excess energy as they travel through the sample until the ion energy is equal to the resonant energy, and the reaction can occur deeper in the material. Therefore, detected γ -rays for an ion energy greater than the resonant energy can be attributed to hydrogen deeper within the sample. The number of detected γ -rays will be proportional to the hydrogen concentration at that depth.

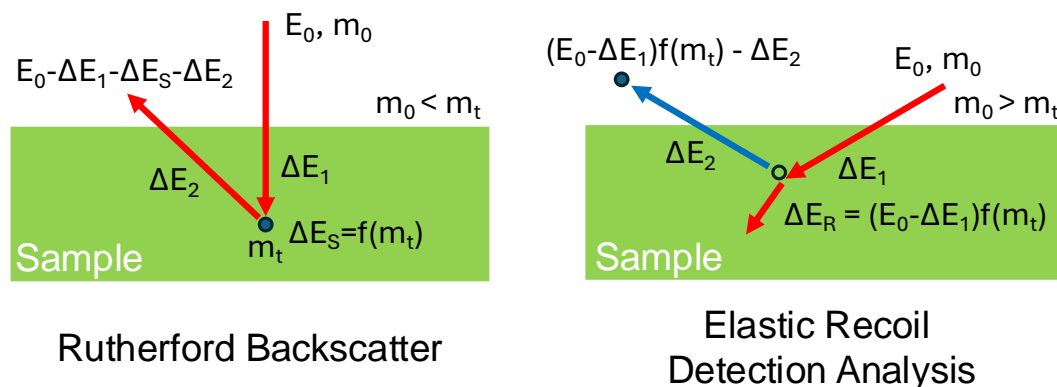


Figure 2.3: Diagram of two common ion beam techniques used for materials composition and depth profiling. In Rutherford Backscattering (RBS), the back scattered primary ion (red arrow) is detected, whilst elastic recoil detection (ERDA) uses the energy of the recoiled target ion (blue circle/arrow). E_0 and m_0 indicate primary ion energy and mass respectively, RBS requires a m_0 less than the mass of the target atom m_t , whilst ERDA requires the opposite. Energy losses ΔE_1 and ΔE_2 correspond to ions moving through the material before and after the scattering event respectively. The changes in energy during scattering events, ΔE_S or ΔE_R , are a function of m_t , $f(m_t)$.

A narrow resonance width helps ensure good depth resolution, as a smaller increase in ion energy can be used whilst avoiding overlap in γ emission peaks. Similarly, having a bigger difference between resonant and non-resonant cross sections will improve sensitivity in general.

NRA can also be used in a non-resonant manner, as is typically done for deuterium detections. Deuterium can be distinguished from hydrogen using the non-resonant, deuterium specific reaction [129]



Here, the detection of emitted protons can be used to calculate deuterium concentrations. As this is a non-resonant reaction, there is no sudden, significant increase in cross section that can be used for depth profiling. If only deuterium is present, or the ratio of deuterium and hydrogen concentrations is consistent throughout the sample, the reaction with ${}^{15}\text{N}$ (Eq.2.3) could be used in conjunction with ${}^2\text{D}({}^3\text{He}, \text{p}){}^4\text{He}$ to give an estimate of the deuterium concentration with depth. In this work, hydrogen is incorporated throughout the sample during growth whereas the implanted deuterium is restricted to the surface. Therefore, the ${}^1\text{H}({}^{15}\text{N}, \alpha\gamma){}^{12}\text{C}$ reaction which cannot distinguish between deuterium and hydrogen would not offer any insight into deuterium depth. Instead, the energy of the emitted proton can be used to estimate the depth at which the reaction occurred [130, 131]. For non-resonant NRA, the ion beam energy does not need to be varied for depth profiling, so it can be selected to maximise the signal to noise ratio based on the cross sections of the material and the element to be measured.

2.3.3 Elastic Recoil Detection Analysis

Other IBA techniques such as Rutherford Backscattering (RBS) and elastic recoil detection analysis (ERDA) are commonly used to determine the composition of materials [132], with the latter used for hydrogen detection [71]. In RBS, a low incident angle is used, and the incident primary ion is lighter than the target ion within the material. Primary ions backscatter from the heavier target ions, resulting in a loss of energy dependant on the mass of the target ion. Incident ions will also loose energy travelling through the material both to and from scattering events. Therefore, measuring the energy of the backscattered ion will give information on both the target ion mass and its depth. This is in contrast to the geometry used in ERDA (see Fig. 2.3), which uses a high incident angle and primary ion mass greater than the target ion. The energy of the primary ion is transferred to a recoiling target ion which is emitted from the surface and detected. Once again, the energy of this recoiled particle will contain information on its depth and mass. For hydrogen detection, a helium beam is commonly used as this will avoid recoil reactions with other, heavier elements.

2.3.4 Raman Spectroscopy

Raman spectroscopy [133] is an optical technique that relies on the interaction between a source of monochromatic light (typically a laser in the near-infrared to ultraviolet region) and phonons within the material. Reflected light will be elastically scattered, red or blue shifted (Rayleigh, Stokes or anti-Stokes scattering respectively, see Fig. 2.4). Both Stokes and anti-Stokes scattering can be used to probe the energetics of the materials phonons which are largely dependant on the elements present and the bonding between elements. Different wavelengths can be used to probe different phonon energies, therefore, certain wavelengths are more appropriate than others for identifying particular species and bonding. The probing depth is typically around 1 μm but is dependant on the sample and the laser wavelength. The resultant spectra is plotted as intensity against Raman shift (cm^{-1}) which indicates the difference in wave number of the incident beam compared to the excited beam reflected from the sample. Fluorescence (photon emission of an excited molecule returning to the ground state) will also be received by the detector and is typically removed from the spectra.

In the context of this work, Raman can be used to infer the presence of $\text{sp}^3\text{C-C}$ bonds present in diamond, as well as $\text{sp}^2\text{C-C}$ bonding and the presence of boron in doped films. For all spectra presented in this work, a 512 nm green laser was used, which gives a Raman shift of 1332 cm^{-1} for diamond bonding.

2.3.5 Electron Microscopy

Electron microscopy [134] is blanket term which includes a variety of different techniques that use the interaction of electrons with matter to create an image. A focused beam of primary




Figure 2.4: Simplified diagram of the underlying mechanisms of Raman spectroscopy. Green arrows represent absorption of incident light whilst other arrows show the energy change during different scattering events.

electrons is incident on a small area of the sample, the signal from which gives the data for a pixel of an image. The electron beam can raster across the sample with very high positional accuracy to build up an image of the surface. The electrons can interact with the sample in different ways as shown in Fig. 2.5, resulting in a variety of different techniques:

- **Secondary Electron Microscopy (SEM)** - Secondary electrons are produced via inelastic scattering mechanisms, where interactions with primary electrons result in electron emission from the sample. Many secondary electrons are emitted from a material's surface which can be easily detected, allowing the imaging of surfaces to a high resolution.
- **Backscattered Electron Imaging (BSE imaging)** - Primary electrons can also be backscattered from the sample through elastic (Rutherford) scattering. Heavier, higher Z elements produce a greater number of backscattered electrons and therefore can be used to contrast between different elements and phases. BSE that have undergone coherent Bragg scattering can be collected to create a diffraction or Kikuchi pattern (electron backscatter diffraction, EBSD [135]). This contains information of the materials grain structure, orientation and phase.
- **Transmission Electron Microscopy (TEM)** [136] - For thin (< 100 nm), electron transparent samples, primary electrons can be transmitted through the sample and detected the other side. There are different techniques that can be used to obtain a sample of this thickness: a pre-thinned sample can be polished to create holes with very thin regions surrounding the aperture, or a focused ion beam can be used to mill out a cross-sectional lamella from the sample. TEM is commonly used to study irradiation damage, as physical changes to the sample can be observed as a function of depth, even in-situ in some setups [137]. The contrast in a TEM image is given by electrons being scattered by different amounts as they pass through the sample, with brighter areas corresponding to reduced




Figure 2.5: Graphic depicting the different interactions used for various techniques in electron microscopy. Red and blue arrows represent electrons from the primary beam and the sample respectively.

scattering. Electron scattering can be impacted by changes in crystallographic orientation, defects and varying electron densities within the sample, allowing TEM to be used to image a variety of different structural features.

- **Energy-dispersive X-ray spectroscopy (EDX)** - When secondary electrons are emitted, electrons from higher energy levels drop down to the lower energy hole left behind. The transmission energy is released as a photon with a wavelength in the X-ray spectrum. Electron transitions emit discrete energies characteristic of that element. Analysing the energy spectrum of emitted X-rays can therefore be used to distinguish between different elements.
- **Auger Electron Spectroscopy (AES)** [138] - The transmission energy released from the ejection of secondary electrons can also be transmitted to other electrons, which are then emitted at low energy (an Auger electron). The energy of an Auger electron contains information of electron energy levels and therefore can be used to contrast different elements within a material. Although similar to EDX, the yield of Auger electrons is highest for low Z elements, allowing detection of elements as light as lithium. This is in contrast to X-ray yield which is greater for higher Z elements.

2.3.6 Secondary Ion Mass Spectroscopy

Secondary ion mass spectroscopy (SIMS) [139] is commonly used for determining the elemental composition of a material. The primary ion beam (commonly gallium or another heavy ion) rasters across on a small area of the sample, causing secondary ions to be sputtered from the material. Sputtered ions are accelerated via an electric field to a MS (the SIMS used in this work featured a sector MS) where ions are identified based on mass.

Ultimately, this is a destructive technique as part of the surface is etched away. The etch rate depends on the sample and the magnification of the primary ion beam. Determining the etch rate (by measuring the trench depth) allows for depth profiling of samples. Concentrations of elements can also be calculated, for example, the boron concentration of a boron doped diamond film. For this, the ratio of boron to carbon counts can be multiplied by a calibration factor. Calibrations are required to account for the sensitivities of different elements. This can be done with a reference sample of known concentration, typically from a high energy ion implantation where all incident ions are assumed to be implanted.

Although SIMS can be used to determine hydrogen/deuterium concentrations [140], it was not possible to make such measurements using the equipment available. It was attempted, however the limited depth of implanted deuterium meant most retained deuterium was lost almost instantly. SIMS was used for determining boron concentrations in boron doped diamond films.

2.4 Molecular Dynamics

2.4.1 General Concept

Generally, MD simulations numerically solve Newtons laws of motion for each atom in the system. The interatomic forces are dictated by a potential (also known as a force field), and the position, velocity and acceleration of each particle is updated at each step. Various conditions (such as temperature and pressure) can be applied to the simulation at each step to try and imitate a physical scenario, whilst calculations can be performed to estimate various properties (like total energy and volume) of the system as the simulation progresses. In this work, the computational facilities of the Advanced Computing Research Centre, University of Bristol and the Cambridge Service for Data Driven Discovery (CSD3) were used to carry out MD simulations using the LAMMPS code [80, 81].

2.4.2 Potentials (Force Fields)

The most computationally expensive part of a MD simulation is calculating the forces between particles. The total force experienced by a particle is the summation of forces exerted on it by its neighbours and can include both long-range and short-range forces. Potentials dictate the interatomic forces between atoms. Selection of an appropriate potential will dictate the accuracy and speed of the simulation, and only certain potentials will be appropriate for certain atom types. Basic potentials, such as the Lennard-Jones potential, make a two-body approximation, with forces that are solely dictated by the separation between pairs of atoms. In contrast, many-body potentials consider the impact of multiple interactions together, allowing for a more accurate representation of the local bonding environment. An example of this is the Reactive Empirical Bond Order (REBO) potential [141, 142]. This potential can accurately reproduce different carbon

bonding environments, by using a bond order function to account for the different bond orders (single bonds, double bonds etc.) possible between carbon atoms. This is in a similar manner to REBO's predecessor, the Tersoff potential [143, 144], but with the added capability of radical and conjugated carbon bonds. The Adaptive Intermolecular Reactive Bond Order (AIREBO) potential [145] is a further development of the REBO potential and is better suited to faster moving/higher energy particles due to the addition of two long range terms, but these come at a computational cost. Both REBO and AIREBO potentials are commonly used to study systems containing carbon and hydrogen and were both used in this work.

2.4.3 Time Stepping

To initiate the simulation, atom coordinates, types and masses are given. The symmetry of a crystalline solid means smaller unit cells can be replicated out in each direction to achieve the desired system size. Initial velocities are then given to the atoms, these are typically sampled from a velocity distribution that corresponds to the desired initial temperature.

Once these initial parameters have been established, a step forward in time can be taken by calculating the forces applied to the atoms and solving Newton's laws of motion. The force experienced by an atom will be a combination of interatomic forces dictated by the potential and any forces from additional conditions applied to the system (such as thermostats and barostats and boundary conditions). The acceleration of the atom can then be determined, from which the position of the atom at the next step can be determined with Verlet integration [146]. The new position at a time $t + \Delta t$ can be approximated by a Taylor expansion,

$$(2.5) \quad \mathbf{r}(t + \Delta t) = \mathbf{r}(t) + \mathbf{v}(t)\Delta t + \frac{1}{2}\mathbf{a}(t)\Delta t^2 + \frac{1}{6}\mathbf{b}(t)\Delta t^3 + O(\Delta t^4),$$

where \mathbf{r} , \mathbf{v} , \mathbf{a} and \mathbf{b} are position, velocity, acceleration and jerk respectively, and $O()$ indicates higher order terms. In Verlet integration, the position of the previous step is then considered,

$$(2.6) \quad \mathbf{r}(t - \Delta t) = \mathbf{r}(t) - \mathbf{v}(t)\Delta t + \frac{1}{2}\mathbf{a}(t)\Delta t^2 - \frac{1}{6}\mathbf{b}(t)\Delta t^3 + O(\Delta t^4).$$

Summing these equations gives

$$\mathbf{r}(t + \Delta t) + \mathbf{r}(t - \Delta t) = 2\mathbf{r}(t) + \mathbf{a}(t)\Delta t^2 + O(\Delta t^4)$$

$$(2.7) \quad \mathbf{r}(t + \Delta t) = 2\mathbf{r}(t) - \mathbf{r}(t - \Delta t) + \mathbf{a}(t)\Delta t^2 + O(\Delta t^4).$$

Now, the position at the next step can be calculated with the current position, the previous position and the acceleration - eliminating the jerk term and maintaining an error of $O(\Delta t^4)$. The acceleration of the particle is given by the force it experiences, \mathbf{F} , based on its position compared to other atoms and its mass, m ,

$$(2.8) \quad \mathbf{a}(t) = \frac{\mathbf{F}(\mathbf{r}(t))}{m}.$$

The velocity can be found using the Stormer-Verlet method. Subtracting Eq.2.6 from Eq.2.5 gives

$$\begin{aligned} \mathbf{r}(t + \Delta t) - \mathbf{r}(t - \Delta t) &= 2\mathbf{v}(t)\Delta t + \frac{1}{3}\mathbf{b}(t)\Delta t^3 \\ \mathbf{v}(t) &= \frac{\mathbf{r}(t + \Delta t) - \mathbf{r}(t - \Delta t)}{2\Delta t} + \frac{\mathbf{b}(t)\Delta t^3}{3\Delta t} \\ (2.9) \quad \mathbf{v}(t) &= \frac{\mathbf{r}(t + \Delta t) - \mathbf{r}(t - \Delta t)}{2\Delta t} + O(\Delta t^2). \end{aligned}$$

This is inconvenient, as information from the previous, current and next step is required to complete a step. Additionally, this equation relies on a fixed Δt . Use of a variable timestep is often preferable and means the timestep can be reduced if a certain criterion is fulfilled (such as an atom exceeding a certain displacement during a timestep) to maximise the default timestep whilst maintaining accuracy. The reliance on the previous step prevents the use of a variable timestep but can be avoided, by rewriting Eq.2.9 as

$$(2.10) \quad \mathbf{v}(t + \Delta t) = \frac{\mathbf{r}(t + \Delta t) - \mathbf{r}(t)}{\Delta t} + O(\Delta t).$$

However, this results in a large error of $O(\Delta t)$. The alternative is the velocity Verlet algorithm. As before, we can write the position, velocity and acceleration of the next step in as a Taylor expansion

$$(2.11) \quad \mathbf{r}(t + \Delta t) = \mathbf{r}(t) + \mathbf{v}(t)\Delta t + \frac{1}{2}\mathbf{a}(t)\Delta t^2 + O(\Delta t^3)$$

$$(2.12) \quad \mathbf{v}(t + \Delta t) = \mathbf{v}(t) + \mathbf{a}(t)\Delta t + \frac{1}{2}\mathbf{b}(t)\Delta t^2 + O(\Delta t^3)$$

$$(2.13) \quad \mathbf{a}(t + \Delta t) = \mathbf{a}(t) + \mathbf{b}(t)\Delta t + \frac{1}{2}\frac{d}{dt}\mathbf{b}(t)\Delta t^2 + O(\Delta t^3).$$

Rearranging Eq.2.13 for $\mathbf{b}(t)$ and substituting this into Eq.2.12 gives

$$\begin{aligned} \mathbf{v}(t + \Delta t) &= \mathbf{v}(t) + \mathbf{a}(t)\Delta t + \frac{\Delta t}{2}(\mathbf{a}(t + \Delta t) - \mathbf{a}(t)) + O(\Delta t^3) \\ (2.14) \quad \mathbf{v}(t + \Delta t) &= \mathbf{v}(t) + \frac{\Delta t}{2}(\mathbf{a}(t + \Delta t) + \mathbf{a}(t)) + O(\Delta t^3) \end{aligned}$$

This has a smaller error than both Eq.2.9 and Eq.2.10, and it does not require information from the previous step - permitting the use of variable timesteps. Eq.2.14 can also be written in terms of force with Eq.2.8 to give the steps of the algorithm

1. Calculate new positions with $\mathbf{r}(t + \Delta t) = \mathbf{r}(t) + \mathbf{v}(t)\Delta t + \frac{\mathbf{F}(\mathbf{r}(t))}{2m}\Delta t^2$.
2. Calculate forces for new positions ($\mathbf{F}(\mathbf{r}(t + \Delta t))$) using an appropriate potential.

3. Calculate new velocities with $\mathbf{v}(t + \Delta t) = \mathbf{v}(t) + \frac{\Delta t}{2m}(\mathbf{F}(\mathbf{r}(t)) + \mathbf{F}(\mathbf{r}(t + \Delta t)))$.

These steps can be repeated as required to run the simulation. If a variable timestep is used, a maximum displacement or maximum energy change of an atom during a simulation step is defined. For every atom at the end of each step, a conservative estimate is calculated of the largest timestep that could be used whilst preventing a displacement or increase in energy greater than the defined maximums. If this is smaller than the default timestep for any of the atoms in the system, the timestep is decreased for the next step, always ensuring no atom experiences a change in position or energy outside of the defined limits. To allow the use of a variable timestep, the velocity Verlet algorithm was used in this work.

2.4.4 Thermostats

In LAMMPS, a ‘fix’ can be applied to the system at the end of each step [147]. Fixes can be used to apply various conditions to the system, such as controlling temperature, pressure or applying boundary conditions.

Thermostat algorithms are used to maintain desired temperature of the system. This is useful for simulating a canonical NVT ensemble, where the amount of substance (N), volume (V) and temperature (T) are conserved. As temperature is a statistical quantity, and there are limited numbers of atoms present in a MD system, it can fluctuate significantly. Statistical temperature can be approximated with instantaneous temperature which is given by summing over the particles of the system [148]

$$(2.15) \quad T(t) = \sum_{i=1}^N \frac{m_i v_i^2(t)}{k_B N_f}$$

where m_i and v_i are the mass and velocity of a particle, and N_f is the number of degrees of freedom. As velocities are updated at each timestep, significant fluctuations in temperature can be expected, especially if energy is being added to the system (such as a new energetic atom). Temperature corrections are required to ensure instantaneous temperature fluctuates about the target temperature.

Velocity scaling is the simplest way in which temperature can be controlled. Here, velocities are scaled at each step to maintain a constant average kinetic energy per particle. However, this does not truly simulate a constant temperature ensemble, in which instantaneous temperature will fluctuate [148]. Thermostats are commonly used to avoid this which, instead of directly scaling velocities, act as a heat bath which the simulated system can transfer energy with. This is closer to how a physical system would behave if the simulation cell is considered to be a small volume of a larger material. Different thermostats approach this interaction between the system and heat bath in different ways. The Nosé-Hoover thermostat [149] was used in this work and works by including additional dynamic variables which are linked to particle velocities. The temperature bath offers resistance to temperature changes and pushes the velocities of




Figure 2.6: Simplified diagram of periodic boundary conditions for a simulation cell that is too small (**A**) and sufficient size (**B**). Light blue atoms are exerting a force, and dark blue atoms are receiving a force. Black boxes indicate the boundaries of the simulation cell, with faded boxes showing the effect of periodic boundary conditions. When too small, two atoms can interact twice (green arrows), or an atom can interact with itself (red arrows). This does not occur with the larger system size.

particles towards the target temperature. In LAMMPS, the rate at which these changes are made can be controlled by a time parameter t_{damp} . A shorter t_{damp} allows for faster return to target temperature, but can result in erratic temperature changes. A longer t_{damp} will offer more gradual, controlled temperature changes, but may not respond quick enough if energy is begin added to the system. Barostats operate in a similar manner to thermostats but are used to control pressure.

2.4.5 Boundary Conditions

The computational complexity of a MD simulation can result in more accurate results than other simulation types which fail to specifically model atomic structure. However, this puts constraints on the system size and timescales that are possible to simulate. As such, typical MD simulations consist of 1000–100000 atoms and a simulation volume on the nanometer scale. In order to mimic larger systems, the symmetry of a crystalline system can be exploited. Periodic boundary conditions allow atoms to interact with each other across the simulation boundary - effectively simulate another, identical system joined onto the periodic boundary. Although this is very effective at adding stability to the system, a minimum system size is still required. Smaller systems will exhibit more erratic temperature variations, especially if energetic atoms are to be added as these will make up a greater proportion of the total system energy compared to a larger system. Furthermore, long range interactions between atoms could result in self-interaction if the dimensions of the systems are too small, or the same atom producing opposing forces on another atom (see Fig. 2.6).

EXPOSURE OF DIAMOND SAMPLES TO DEUTERIUM IONS

As discussed in Chapter 1, diamond's interaction with hydrogen is one of the main concerns surrounding its use in fusion reactors. These concerns largely arose from the previously used allotrope of diamond - graphite, which displayed unacceptable levels of chemical etching and retention of hydrogen isotopes. In order to assess diamonds suitability for applications within fusion reactors, its interaction with hydrogen in fusion relevant conditions must be understood. This chapter presents results from experimental work focused on furthering understanding of hydrogen isotope retention within various diamond samples.

Some of the work presented in this chapter has been adapted from [150], namely, the TDS results of undoped PC samples presented in 3.2. Exposures in DELPHI were carried out by A. Wohlers (UKAEA, Abingdon), and the TDS measurements by Y. Zayachuck (UKAEA, Abingdon). All analysis and conclusions of results are the author's (J. Pittard) own, although discussed with the other authors of this paper. NRA measurements were performed by N. Catarino (IST, Lisbon) as part of a RADIATE grant. Selected samples were exposed to a C^+ ion beam to damage the surface. These ion irradiations were conducted by L. Antwis (UKNIBC, Surrey) and were funded through an EPSRC pump primer grant. Electron microscopy (SEM, BSE and EDX) results were acquired with the assistance of J.C. Eloi (CIF, Uni. of Bristol). SIMS measurements were carried out with the assistance of L. Cullingford (IAC, Uni. of Bristol).

3.1 Samples, Preparation and Methodology

UKAEA's DELPHI facility (described in 2.2) was used to expose diamond samples to deuterium ion clusters between 0.2 and 1.0 keV. Deuterium retained in diamond samples was determined using TDS and NRA. Alongside retention measurements, Raman spectra and electron microscopy measurements were taken pre and post exposure for selected samples. SIMS was used to deter-

Table 3.1: Summary of samples discussed within this chapter. PC - Polycrystalline, E6 - Acquired from Element 6, TM100 - Element 6 product code ((100) orientation), BDD - Boron doped diamond either high, low or residual (Res.) boron concentration in gas phase. TDS - Thermal Desorption Spectroscopy, SEM - Scanning Electron Microscopy pre or post exposure in DELPHI, RS - Raman Spectroscopy. Anode voltage refers to the anode in DELPHI and corresponds to the energy of the incident deuterium ions in eV. Delay between exposure and retention measurement is indicated in brackets when applicable.

Sample	Description	Anode Voltage	Retention Analysis	Other Analysis
1-1a	PC E6 TM100 10×10 mm ²	400	TDS (1 day)	SEM-Post
1-1e	PC E6 TM100 10×10 mm ²	600	TDS (7 days)	SEM-Post
1-1f	PC E6 TM100 10×10 mm ²	200	TDS (7 days)	-
1-1g	PC E6 TM100 10×10 mm ²	1000	TDS (7 days)	SEM-Post
1-2a	PC E6 TM100 10×10 mm ²	800	TDS (7 days)	SEM-Post
1-2b	PC E6 TM100 10×10 mm ²	300	TDS (7 days)	SEM-Post
1-3a	PC E6 TM100 10×10 mm ²	-	TDS	-
1-4b	SC E6 (110) 6.8×3.3 mm ²	400	NRA (1 month)	-
1-4c	SC E6 (110) 6.8×3.3 mm ²	800	NRA (1 month)	-
1-5b	PC E6 TM100 10×10 mm ²	400	NRA (1 month)	-
1-5c	PC E6 TM100 10×10 mm ²	800	NRA (1 month)	-
2-3a	High [B] BDD 10×10 mm ²	400	TDS (7 days)	SIMS, SEM-Pre/Post, RS
2-3b	High [B] BDD 10×10 mm ²	-	-	SEM-Pre, RS
2-4a	Res. [B] BDD 10×10 mm ²	-	-	SEM-Pre, RS
2-4b	Res. [B] BDD 10×10 mm ²	400	TDS (7 days)	SEM-Pre, SIMS, RS
2-5a	High [B] BDD 10×10 mm ²	-	-	SEM-Pre
2-5b	High [B] BDD 10×10 mm ²	-	-	SEM-Pre
2-6a	High [B] BDD 10×10 mm ²	-	-	SEM-Pre
2-6b	High [B] BDD 10×10 mm ²	-	-	SEM-Pre
2-8a	Res. [B] BDD 10×10 mm ²	-	-	SEM-Pre, RS
2-8b	Res. [B] BDD 10×10 mm ²	400	-	SEM-Pre/Post, RS
2-11a	Low [B] BDD 10×10 mm ²	400	TDS (7 days)	SEM-Pre, RS
2-11b	Low [B] BDD 10×10 mm ²	-	-	SEM-Pre, RS
2-13a	Low [B] BDD 10×10 mm ²	400	-	SEM-Post, RS
2-13b	Low [B] BDD 10×10 mm ²	-	-	RS
2-14a	Low [B] BDD 10×10 mm ²	-	-	SIMS, RS
2-14b	Low [B] BDD 10×10 mm ²	-	-	SIMS, RS
3-1a	SC E6 (110) 6.8×3.3 mm ² C ⁺ Irradiation to 1 dpa	-	-	SEM, RS
3-2a	SC E6 (110) 6.8×3.3 mm ² C ⁺ Irradiation to 0.1 dpa	-	-	SEM, RS
3-3a	SC E6 (110) 6.8×3.3 mm ² C ⁺ Irradiation to 0.01 dpa	-	-	SEM, RS
3-4a	SC E6 (110) 6.8×3.3 mm ²	-	-	SEM, RS

mine atomic composition of boron doped samples. A variety of different diamond samples were explored, including PC and SC diamond, boron doped diamond (BDD) and ion damaged samples. Although the intention was to expose all samples in DELPHI, delays surrounding this facility meant this was only possible for selected samples. A summary of all samples to be discussed in this section can be seen in Table 3.1.

3.1.1 Polycrystalline Diamond

PC samples acquired from Element 6 (product code 145-500-0536) were the first samples tested. Although SC diamond is required for window and sensor applications, PC diamond is significantly lower cost and easier to fabricate in large areas or as a coating on non-diamond substrates, making it more relevant for plasma facing applications. Furthermore, the moderate size ($10 \times 10 \text{ mm}^2$) of these samples acts to maximise total deuterium counts. This is particularly important for whole sample techniques such as TDS, as smaller samples will result in lower counts and increased signal to noise. Commercially grown samples were selected for initial testing to avoid concerns over consistency and delamination from a substrate, allowing other variables to be compared with greater confidence.

Eight of the PC diamond samples were exposed to deuterium ions in DELPHI with an anode voltage varying between 200 and 1000 V (corresponding to an ion cluster energy of 200–1000 eV). All samples were exposed for five hours at ambient temperature and a flux of $3.04 \times 10^{17} \text{ ions m}^{-2} \text{ s}^{-1}$ [120], reaching an estimated total fluence of at least $5.5 \times 10^{21} \text{ ions m}^{-2}$. An additional sample acted as a reference and underwent TDS without any exposure to deuterium ions beforehand.

Post exposure, retention in the samples was determined using either TDS or NRA. TDS measurements were taken 7 days after exposure for all samples except 1-1a (400 eV exposure of PC diamond) which was measured the following day. TDS was carried out using a Hiden Analytica Ltd Type 640100 TPD workstation [125]. Samples were placed on a hot plate which was heated from room temperature to 1273 K, at which it was held for an hour. A heating rate of 10 K min^{-1} was used. The presence of residue from an adhesive mounting disc meant one sample required a protective aluminium nitride (AlN) layer to protect the TDS from contamination. This protective layer reduced the temperature reached by the sample (resulting in lower total counts) and meant a temperature correction was required [151]. Signals of masses 2, 3 and 4 (corresponding to H_2 , HD and D_2 molecules) were detected and quantified using H_2 and D_2 calibrated leaks (with the calibration factor for HD being an average between the two).

NRA was carried out at IST Lisbon with support from the RADIATE project. Here, a helium-3 ion beam was incident on a $\varnothing 1 \text{ mm}$ circular area of the sample. Protons detected from the $^2\text{H}(^3\text{He},\text{p})^4\text{He}$ reaction were used to infer the presence of deuterium, the energy of which can be used to determine the depth profile (see 2.3.2). Initially, measurements were attempted with a 2.3 MeV ^3He beam (an energy commonly used when analysing fusion materials). However, deuterium proved to be in such low amounts, a 1 MeV beam was required (increasing the cross

Table 3.2: Summary of growth conditions of boron doped samples, grown in a hot filament CVD reactor. All samples were grown on a molybdenum substrate seeded via manual abrasion. Descriptions ('Lightly', 'Heavily' etc.) refer to the B:C gas phase ratio, whilst boron concentration [B] refers to SIMS measurement of samples. *Value from a different sample grown under equivalent conditions.

Sample	Description	Gas Phase B:C	CH ₄ (sccm)	H ₂ (sccm)	5% B ₂ H ₆ in H ₂ (sccm)	Growth Time (hrs)
2-4b	Residually Doped [B] = $8.34 \times 10^{20} \text{ cm}^{-3}$	Residual	200	2	0	4:15
2-11a	Lightly Doped, [B] = $*5.88 \times 10^{20} \text{ cm}^{-3}$	3500 ppm	200	2	0.07	4:00
2-3a	Heavily Doped [B] = $6.55 \times 10^{21} \text{ cm}^{-3}$	35000 ppm	200	2	0.70	4:10

section with deuterium whilst decreasing the cross section with carbon) to successfully detect deuterium within the samples.

3.1.2 Single Crystal Diamond

PC diamond is inappropriate for some of diamonds potential applications and SC must be used instead. For example, microwave transparency is needed for window applications, and reducing defects can improve carrier mobility and therefore the performance of detectors [152]. Additionally, PC diamond has a higher sp^2 content from the increased number of defects and the presence of grain boundaries. As sp^2 carbon exhibits greater chemical etching from hydrogen ions, SC diamond would be expected to exhibit increased resilience. The presence of grain boundaries could increase retention [64], and impact diffusion characteristics. The impact of grain boundaries has been further explored through simulation in Chapter 4.

SC diamond samples were acquired from Element 6 (product code 145-500-0573). Samples were $6.8 \times 3.3 \times 0.2$ mm in size and orientated with (110) faces aligned with the surface. Two of these were exposed in DELPHI with anode voltages of 400 and 800 V respectively. These voltages were selected as conclusions from PC samples suggested they might be of particular interest. Retention within these samples was evaluated using NRA.

3.1.3 Boron Doped Diamond

As discussed in 2.1.3, boron doping is commonly used in diamond to alter physical and electrical properties. Porro *et al.* [65] highlighted that the inclusion of boron into diamond could be of benefit for plasma facing applications as the increased conductivity reduces arcing effects. In graphite, boron doping was found to improve resistance to chemical etching [153] but at the cost of increased retention [39]. Furthermore, many sensor and diagnostic applications required boron doping, so it is important to consider the effect this has on retention.

Boron doped diamond films were grown using a hot filament CVD reactor on molybdenum substrates seeded via manual abrasion. Different flow rates of diborane gas were used to vary the boron concentration between different samples. Some samples were residually doped to try and achieve low levels of doping. To do this, no additional diborane is added, but boron desorbed from the chamber walls during growth is present. A summary of the growth conditions can be seen in Table 3.2. A relatively long growth time of at least 4 hrs was used for all samples to ensure reasonable film thickness was achieved. Provided the film thickness is greater than the penetration depth of incident deuterium (expected of order nm) exact thickness is not of concern for these experiments but is expected to be of the order of microns. However, film thickness often impacts other variables such as grain size, with thicker films typically exhibiting larger grains [102]. The relatively long growth period was used to try and minimise the likelihood of poor quality or nanocrystalline diamond phases.

SIMS was used to determine boron concentrations (see Table 3.2) for three samples, each grown with different diborane flow rate. To do this, the B:C ratio was averaged across an etch time of approximately 15 minutes. A correction factor had been previously determined¹ for this apparatus by measuring the B:C ratio of a reference sample of known concentration. Applying this correction gave the estimated boron concentration. The highly doped sample gave a concentration an order of magnitude higher than the other samples tested. Surprisingly, the residually doped sample gave a concentration slightly higher than the lightly doped sample, although both were within the same order of magnitude. The boron concentration of residually doped samples is largely dictated by the boron concentrations used in previous growths. Boron has a strong memory effect, and significant amounts will be taken up by the chamber walls from previous growths and released during later growths. If a particularly high boron concentration growth was carried out prior to the growth of the residual sample, a relatively high concentration may still be possible, despite no diborane being added to the gas mixture.

The presence of boron dopants can also be inferred using Raman spectroscopy. Fig. 3.1 shows three Raman spectra of boron doped diamond films taken with a green laser (512 nm) with the background removed. The sharp peak at 1332 cm^{-1} indicates the presence of the diamond structure and sp^3 bonds. The broader peak at 1200 cm^{-1} , and the diminishing of the diamond peak, indicates increased disorder in the lattice as a result of boron doping [154]. For the heavily doped diamond (35000 ppm B:C in gas phase) it can be seen that this merges with the diamond peak and is a classic indicator of highly boron doped diamond. A change in Raman shift of the 1332 cm^{-1} diamond peak indicates stress in the crystal and typically will redshift with boron doping [155, 156] as observed here. For the lightly doped sample, the two peaks are much more easily distinguished, whilst only a small peak at 1200 cm^{-1} is present for the residually doped sample.

Raman spectra of residual and lightly doped samples suggests different boron concentrations

¹By P. J. Heard (IAC, Uni. of Bristol, 27/10/14).

between the samples. This is in contrast to the SIMS results, which gave very similar concentrations for the two growth conditions. The most likely explanation is discrepancies between supposedly equivalent growth conditions, as the sample exposed in DELPHI and measured with the Raman was different to the sample measured in the SIMS. Despite the diborane flow rate being the same, it is possible there is significant variation between the concentrations in these samples because of differences in the boron content of the chamber. For the lightly doped sample, the boron contribution from the chamber may be significant in comparison to the amount from diborane in the gas mixture. If this were the case, the concentration within the film would also be dependent on the residual boron level rather than the just diborane flow rate. Alternatively, differences in concentration could be a result of the small probe area of both Raman and SIMS, meaning the measurement may not be representative of the whole sample. Additionally, all Raman spectra were taken shortly after the growth was complete, but this was not always possible with the SIMS measurements. For the residual and highly doped samples, SIMS measurements were taken after exposure in DELPHI and TDS, whereas the lightly doped sample was measured prior to exposure. Although timings of measurements were not considered to be important, it could be possible that the exposure to deuterium plasma and the effective annealing in the TDS resulted in some change in material properties that ultimately impacted the SIMS measurement.

The inclusion of boron into the diamond resulted in several morphology changes to the films as can be seen in Fig. 3.2. With increased boron concentration, a decrease in planar defects can be observed, with facets appearing larger and smoother for the highly doped sample. This is commonly observed [157], and a result of dopants lowering the barrier for vacancy diffusion and improving incorporation of carbon atoms to their correct lattice positions. The highly doped sample also exhibits more triangular facets typical of (111), rather than the rectangular (100) facets which can be seen, alongside (111) facets, in the lower [B] samples. Reduction of (100) facets is also common in boron doped diamond [158, 159]. Residually and lightly doped samples present with very similar morphologies, supporting the SIMS results suggesting these are of similar boron concentrations.

Boron doped samples were exposed in DELPHI with an anode voltage of 400 V. Standard conditions were used throughout and matched those used for undoped PC samples. Similarly, TDS was used to determine retention in the same manner as undoped samples.

3.1.4 Ion Damaged Diamond

For in-vessel applications, assessing the performance of damaged samples is key when considering the durability of materials. In a fusion reactor, materials will sustain significant damage as a result of high fluences of fast neutrons. However, outside of reactor conditions, it is challenging to achieve sufficient levels of damage via fusion relevant neutron irradiation². Instead, self-ion

²Although, the International Fusion Materials Irradiation Facility-DEMO Oriented Neutron Source (IFMIF-DONES) is under construction, and will achieve neutron doses of over 20 dpa/yr [160].

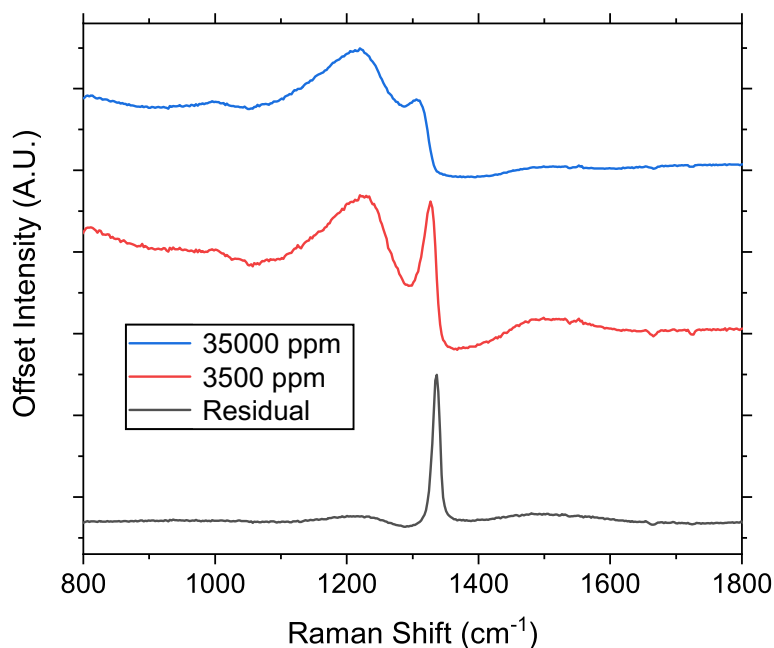


Figure 3.1: Offset Raman spectra for three CVD diamond films (samples 2-3a, 2-11a, 2-4b top to bottom) grown with various flow rates of diborane. Values in the legend refer to the B:C value of the growth gas mixture. Residual doping has no additional boron in the gas mixture, only the residual boron within the chamber. Sharp peak at 1332 cm^{-1} shows the presence of diamond whereas the broader peak around 1200 cm^{-1} is attributed to disorder as a result of boron dopants.




Figure 3.2: Secondary electron images of three boron doped diamond films. **A** - Residually doped (sample 2-8b), 15 kV, 5 nA probe current, 10.2 mm working distance, $\times 10000$. **B** - Lightly doped (sample 2-11a), 15 kV, 5 nA probe current, 9.3 mm working distance, $\times 10000$. **C** - Highly doped (sample 2-3a), 15 kV, 5 nA probe current, 10.1 mm working distance, $\times 10000$.




Figure 3.3: Results from TRIM calculations for damage sustained to diamond along $\langle 110 \rangle$.

irradiation is commonly used as a simple way to test the impact of damage to potential materials. Although a high level of damage can be reached in this manner, ion and neutron damage effects are not necessarily equivalent and care is required when comparing the two [161]. Ions will interact far more readily with a material, resulting in increased damage (vacancies, dislocations and other physical defects) and shallower interaction depths. Alongside physical damage, neutron irradiation has the potential for chemical changes. Neutron capture results in transmutation of atoms in the irradiated material, with the potential of creating unstable, radioactive, isotopes. This is of particular concern of heavy elements which will typically accept neutrons much more readily than lighter elements, although this is dictated by the cross section for neutron capture which is highly dependent on energy. Broadly, carbon materials present limited neutron activation compared to other materials, although formation of ^{14}C can be a concern under some conditions [162].

Both simulation and NRA results presented later suggested penetration depth of deuterium ions was limited (< 10 nm). As such, irradiation conditions were optimised for low depth damage, as having an increased damage levels deep within samples will be unlikely to impact retention which is limited to the surface. Irradiation conditions used were optimised for target displacements per atom (dpa) within the top $0.2\ \mu\text{m}$ of the surface.

Calculations to estimate dose for given irradiation conditions were carried out using SRIM (Stopping and Range of Ions in Matter [163]). This is a commonly used Binary Collision Approximation (BCA) code which simulates an amorphous material and probabilistically determines the interaction between a material and incident ions. By avoiding explicitly simulating atomic

Table 3.3: Parameters used for SRIM [163] calculations. Lattice Binding energy was set to 0 eV in agreement with [165].

SRIM Parameters	
Calculation Type	Quick: Kinchen-Pease
Atomic Density	3.53 g cm ⁻³
Displacement Energy	47.68 eV
Lattice Binding Energy	0 eV
Surface Energy	7.41 eV
Incident Ions	10000

structure as is done in MD, much greater system sizes (and higher ion energies) can be simulated. BCA codes become less accurate at low energies where atomic structure and channelling effects can be significant compared to higher energies where these can largely be neglected [164]. The parameters used in calculations can be seen in Table 3.3. Calculations were carried out following the method outlined in [165] in which vacancies are estimated by calculating the damage energy, T_{dam} , from the energy lost to target atoms and ionisation. From the vacancy distribution, the fluence required to reach a desired dpa at a certain depth could be determined.

At these energies, there are two ion stopping mechanisms to be considered: electronic and nuclear [166]. Electronic stopping is a result of inelastic interactions between electrons of the ion and within the target material. Typically, this results in reduced damage compared to nuclear stopping, which is a result of atomic collisions and as such gives much greater energy transfer. Initially, as an ion penetrates a material, ion energy is high and electron stopping will dominate resulting in gradual energy losses of the ion as it travels through the material. Once sufficient energy has been lost, the likelihood of nuclear collisions increases, and nuclear stopping will dominate deeper within the material. The sudden transfer of energy during nuclear stopping results in a peak damage at some distance beyond the surface, as well as a peak in implanted ion depth.

Initially, it was considered that the use of high energy protons would be most appropriate for simulating the neutron damage experienced by a first wall material. However, SRIM calculations revealed that electronic stopping is expected to dominate at the depth of interest for these high energies, meaning unfeasibly high fluence would be required to reach target dpas at the surface. Reducing the energy would move the nuclear stopping depth to the surface, but once again this would require unfeasibly high fluences to reach target damage levels from a lower energy irradiation. Instead, a higher mass was used to achieve the high target doses at a shallow depth (see Fig. 3.3). Use of self-ion irradiation also removes any impact from the chemical interactions between the implanted ions and the target material, or implanted ions and deuterium ions to be implanted later (which could influence TDS results).

At the UK National Ion Beam Centre (UKNIBC), 100 keV $^{12}\text{C}^+$ irradiations were carried out at three fluences ($5.2 \times 10^{18} \text{ C}^+ \text{ m}^{-2}$, $5.2 \times 10^{19} \text{ C}^+ \text{ m}^{-2}$ and $5.2 \times 10^{20} \text{ C}^+ \text{ m}^{-2}$) corresponding to an

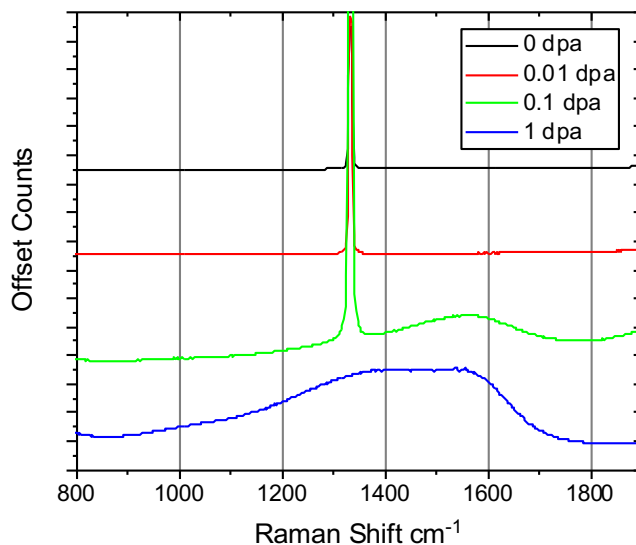


Figure 3.4: Raman spectra taken with a green laser (514 nm) of single crystal (110) diamond samples exposed to 100 keV carbon ions at various doses.

average dpa of 0.01, 0.1 and 1 respectively within the top 0.2 μm (as estimated with SRIM). All irradiations were carried out at room temperature. It was not possible to expose these samples in DELPHI as had been planned, and analysis was limited to SEM and Raman spectroscopy.

Raman spectra of ion irradiated SC diamond samples showed a stark contrast to pristine samples as seen in Fig. 3.4. Clearly, significant damage has occurred in the top micron of the diamond (approximate Raman probe depth), with the highest dose showing little evidence of the 1332 cm^{-1} diamond peak at all. For 0.1 and 1 dpa samples, a peak at 1580 cm^{-1} can be seen which could suggest some graphitisation of the surface has occurred. Deslandes *et al.* [71] performed 5 MeV C^+ irradiation and observed more subtle changes in Raman spectra, which was likely a result of the peak damage of the higher energy irradiation being beyond the probe depth of Raman spectroscopy. Visually, the colour of the samples has changed, with higher dose samples appearing darker as seen in Fig. 3.5.

SEM images of the 1 dpa sample can be seen in Fig. 3.6. Stripes across the surface are visible on the surface as seen in Fig. 3.6A. Rotating the sample stage 90° as was done in Fig. 3.6C to D confirmed these are genuine features of the surface not a result of the raster of the electron beam. Such marks are similar to ripples commonly observed on samples exposed to ion beams [167, 168]. Smaller circular marks are also present across the surface as well as larger particles. Lack of contrast in BSE images (not shown) suggested only carbon was present on the surface. It was apparent that the high dose samples exhibited less surface charging - suggestive of a higher electrical conductivity and, again, possible graphitisation. Few features of note were observed




Figure 3.5: Ion irradiated single crystal diamond samples. The four samples (3-1a, 3-2a, 3-3a and 3-4a left to right) were exposed to 100 keV carbon ions at various fluences to achieve the estimated doses of 1, 0.1, 0.01, 0 dpa respectively. Marks were used to distinguish between irradiated and unirradiated sides.




Figure 3.6: SEM images taken of a SC (110) diamond sample post irradiation to 100 keV carbon ions at a fluence of $5.2 \times 10^{16} \text{ C}^+ \text{ cm}^{-2}$, achieving an estimated dose of 1 dpa within the top 0.2 μm (sample 3-1a). **A** - Raster marks on surface. 15 kV, 5 nA probe current, 19 mm working distance, $\times 20000$. **B** - 15 kV 30 nA probe current, 10 mm working distance, $\times 5000$. **C** - 15 kV, 30 nA probe current, 10 mm working distance, $\times 200$. **D** - Image C rotated 90° to ensure raster lines were physical. 15 kV, 30 nA probe current, 10 mm working distance $\times 2000$

on the lower dose samples due to the polished surface, lack of grain boundaries and significant surface charging.

3.2 Deuterium Retention in Diamond

3.2.1 Pristine Undoped Diamond

Fig. 3.7 presents a comparison of TDS spectra for the six of the PC samples implanted at different energies and the reference PC sample which was not implanted. It can be seen in the D_2 spectra that, as the implantation energy of deuterium ions increased, an additional release peak in the 900–1100 K temperature range appears. A peak in this temperature range was not observed in HD or H_2 spectra - where implantation energy appears to have minimal effect.

Hydrogen and deuterium were incorporated into the diamond in different manners; the former was incorporated from the gas phase during growth, whilst the latter was energetically implanted. Differences in H_2 and D_2 spectra are likely to be a result of these differing processes. CVD diamond is grown in a hydrogen rich environment and exposed surfaces are hydrogen terminated. As grains coalesce during growth, hydrogen on the surface of the grains becomes trapped within grain boundaries, resulting in a high hydrogen content concentrated within the grain boundaries [90]. Therefore, peaks seen in H_2 spectra are likely to correspond to desorption from grain boundary binding sites. In contrast, deuterium was energetically implanted, meaning a wider variety of binding sites were available compared to the low energy hydrogen. Due to the similar mass and bonding mechanisms of hydrogen and deuterium, it would be expected that equivalent binding sites would desorb at the same temperature. As such, the first peak seen in the 600–800 K temperature range which is seen in some of the H_2 , HD and D_2 spectra is attributed to deuterium coming from grain boundary binding sites, as is the third peak seen above 1100 K. The second peak seen in the D_2 spectra (in the 800–1100 K range) becomes more prominent with increasing implantation energy. No significant peak in this temperature range can be seen in H_2 or HD spectra, suggesting the binding sites associated with this peak was not accessible to hydrogen bound within the diamond. Therefore, inter-grain binding sites are a possible candidate for this temperature range, as minimal hydrogen would be found within grains. It is also logical that increasing the implantation energy would result in a greater amount of deuterium successfully penetrating the grains.

Although it can be said with some confidence that hydrogen desorption will predominately occur from grain boundaries, the new desorption peak in D_2 spectra could also indicate ion damage of the surface has occurred. Damage sustained to the diamond could result in an increase in defects and potential trapping sites, offering a possible explanation for the broad peak observed only for ion energies above 600 eV. Su *et al.* [170] reviewed results from a number of TDS studies of (100) diamond. It was concluded the contributions from a high temperature main peak and a lower temperature shoulder were shown to vary with degradation of the surface. Across the studies




Figure 3.7: Counts of masses 4, 3 and 2 (attributed to D₂, HD, H₂ respectively), taken from six PC CVD diamond samples post deuterium implantation at various energies and one reference sample without implantation. The sample stage was heated at a ramp rate of 10 K min⁻¹ from room temperature to 1273 K at which it was held for an hour. The 400 eV sample (blue) required a protective aluminium nitride layer, resulting in a reduced maximum temperature and the need for a temperature correction.




Figure 3.8: TDS spectra of mass 2, attributed to H_2 counts, for five PC CVD diamond samples post deuterium implantation at various energies (as indicated in the legend). Samples were heated at a ramp rate of 10 K min^{-1} from room temperature to 1273 K at which it was held for an hour. A typical temperature profile (brown dashed line) is also included.

discussed in [170], activation energies for desorption were said to range from $1.6\text{--}3.4\text{ eV}$, with most values towards the top end of this range. As discussed by Su *et al.*, one study used electron irradiation to destroy the (100) surface [171], which lead to the lower temperature peak becoming the main peak with a higher temperature shoulder. In the results collected here, the new peak seen for higher energy implantation could indicate a similar effect. Furthermore, Deslandes *et al.* [76] observed notable increase in damage to diamond samples for 500 eV deuterium ions, and proposed a mechanism in which higher energy ions increase damage present in the subsurface of the diamond. Raman spectroscopy measurements taken post exposure were unable to highlight any notable differences between samples before and after exposure, however this is likely due to the probing depth being too high. Deslandes *et al.* used NEXAFS to evaluate the surface damage at a shallower depth and saw clear differences with increasing ion energy. Therefore, it seems likely that some surface damage has occurred for the higher energy implantations, which has resulted in significant trapping within diamond grains. This observation is in contrast to lower energy implantations in which trapping is thought to be dominated by grain boundaries. As grain boundaries make up a relatively small proportion of the overall surface, this suggests incident deuterium can diffuse into grain boundary trapping sites when stable binding sites in the grains are not available.

As the initial peak ($600\text{--}850\text{ K}$) is present in spectra of all masses (if not all samples) it is likely that desorption in this range is associated with surface level bound hydrogen/deuterium within grain boundaries. Simultaneous hydrogen and deuterium desorption is further supported

Table 3.4: Calibrated total counts from TDS data of six PC diamond samples implanted with deuterium at various energies. Note values marked * were heated to a lower max temperature, giving lower total count values.

Implantation Energy	TDS Total Counts		
	HD ($\times 10^{15} \text{cm}^{-2}$)	D ₂ ($\times 10^{15} \text{cm}^{-2}$)	Total D ($\times 10^{15} \text{cm}^{-2}$)
200 eV	3.62	0.925	5.47
300 eV	2.99	0.579	4.15
400 eV	1.72*	0.652*	3.02*
600 eV	3.21	0.933	5.08
800 eV	3.67	1.19	6.05
1000 eV	3.18	1.25	5.67
Reference	0.249	0.041	0.331

by the presence of a peak in HD spectra. No trends can be seen between implantation energy and the presence of a H₂ peak in this region with 300, 400 and 800 eV samples not presenting a H₂ peak in this temperature range. The absence of this peak in these samples is likely due to variation in growth conditions, as these samples were acquired together it is probable they were grown in the same batch. Similarly, the 200, 600 and 1000 eV samples were also acquired as a set, and all presented a peak in the 600–800 K range.

In a similar experiment, deuterium counts in the 700–1100 K range for diamond post 1000 eV deuterium implantation were attributed to sp³C-D defects which, upon desorption, resulted in sp²C-C bonding [172] (although no changes in Raman spectra were observed in this work). Overlap between the initial (600–850K) and central (850–1100 K) peaks can be seen in D₂ spectra, which suggests similar binding sites within the grain boundaries and the grains despite grain boundaries having a higher sp²C-C content [64]. This is because hydrogen would still be expected to bond to sp²C-C as sp³C-D defects [172]. The differences in these two temperature regions are a result of hydrogen/deuterium location rather than differences in bonding type. The absence of release peaks in the 800–1100 K range in HD and H₂ spectra shows a clear distinction between this and the 600–800 K region.

As the high temperature peak at >1100 K is prominent in HD and H₂ spectra, it is likely to correspond to grain boundary desorption. Significant hydrogen diffusion would be expected at these higher temperatures [173], allowing hydrogen from deeper within the crystal (potentially beyond the interaction volume of the deuterium) to outgas. This explanation is further supported by Fig. 3.8, which shows hydrogen counts against time, showing hydrogen desorption during the hour fixed at 1273 K (from approximately 5800 s onwards). Once held at 1273 K, H₂ desorption of deuterium implanted samples settles at a flat rate. This constant outgassing of H₂ would be expected for hydrogen diffusing out of the diamond, instead of a peak in desorption which would suggest a particular binding type that requires some threshold energy to desorb. The 400 eV sample was disregarded from this figure as the temperature correction made comparing desorption with respect to time against other samples invalid. It is unclear why the reference

Table 3.5: NRA results for different diamond samples exposed to deuterium ions. DELPHI voltage refers to the potential difference used to accelerate deuterium ions. SC - single crystal, PC - polycrystalline.

Sample	Description	DELPHI voltage (V)	Retention ($\times 10^{15} \text{cm}^{-2}$)
1-5b	PC E6 (100)	400	1.81 ± 0.13
1-5c	PC E6 (100)	800	2.97 ± 0.16
1-4b	SC E6 (110)	400	1.60 ± 0.12
1-4c	SC E6 (110)	800	3.18 ± 0.17

sample does not display the flat desorption rate of the deuterium implanted samples. The deuterium desorption in this fixed temperature region is predominately HD, as the amount of deuterium to successfully penetrate to these deeper depths is minimal and a greater amount of hydrogen relative to deuterium is present. Furthermore, the 800 eV sample showed the lowest H_2 counts with the highest D_2 counts at this peak, suggesting presence of hydrogen could be limiting deuterium uptake deeper within grain boundaries. Deuterium peaks in this temperature range in similar experiments have been attributed to D_2 bubbles produced during implantation [172]. This explanation seems unlikely from the results presented here, as this peak is most prominent in the spectra of grown-in hydrogen. Furthermore, bubbles typically require fast diffusion rates to allow the collections of hydrogen to form, and the slow diffusion of hydrogen in diamond at room temperature is unlikely to be sufficient for bubble formation.

As can be seen in Table 3.4, there is reasonable variation in total count values from the TDS. Conclusions drawn between total counts alone are limited. These differences are likely to be impacted by disparities in implantation fluence and between the samples themselves rather than solely correlated with implantation energy. No trends between total HD counts and implantation energy can be seen but the highest energies resulted in the highest D_2 counts. These observations support the inter-grain penetration theory. Deuterium and hydrogen desorption at the same temperature is likely to be from equivalent binding sites - with the presence of hydrogen suggesting grain boundary desorption. As such, peaks in HD correspond to easily accessible grain boundary desorption and would be expected to be largely uncorrelated with implantation energy. Whereas the amount of D_2 implanted in grains would increase with energy. Furthermore, it can be said that these values are comparable to values of metal PFM candidates tested in the same experimental setup [120, 174, 175]. It is worth noting the peak seen in the $>1100 \text{ K}$ range in Fig. 3.7 is at higher temperatures than peaks observed in these studies. Therefore, it is plausible that more deuterium remains within diamond samples post TDS than in the metal samples, which could contribute to comparably lower counts in diamond samples.

Retention values determined via NRA for four samples can be seen in Table 3.5. Deuterium within the samples was at such shallow depth a full depth profile was beyond the resolution of this method, and it could only be said that all deuterium is within the top 10 nm of the surface. This is in good agreement with computational work presented in Chapter 4 and implies there would

be no change to bulk properties as supported in other diamond-deuterium plasma experiments [71, 74, 76, 77, 79].

Anode voltages of 400 and 800 V were selected as TDS results suggested the higher voltage may result in the inter-grain penetration, whereas the lower energy may not be able to. It was thought this might present as a stark contrast between the two energies for SC samples, and less of a difference for PC samples.

For the PC samples, NRA retention values are in good agreement with TDS results, with values to the same order of magnitude. Results collected here show higher energy deuterium implantation (samples 1-5c and 1-4c) resulted in greater retention. TDS results suggested higher energy implantation allowed deuterium to access additional binding sites within the grains themselves. This, combined with the expected increase in penetration depth of higher energy ions, could result in greater retention³.

Trends between SC and PC samples are less obvious, with the SC samples giving higher retention for 800 eV implantation, but lower retention for 400 eV. Comparing the retention of the two energies, PC samples show a 1.64 factor increase from 400 eV to 800 eV, whereas SC samples give an increase of 1.99. The larger increase present for the SC samples could be a result of the inter-grain penetration suggested in previous TDS results, as a bigger difference is to be expected between the two energies when grain boundaries are not present. At 800 eV, the high retention values of the SC sample suggests surface orientation also impacts deuterium uptake. MD simulations have suggested that channels present in the $\langle 110 \rangle$ direction may result in increased penetration depth which in turn could result in increased retention (see 4.3). At the higher energy, where it is expected that deuterium can access the grains, the influence of grain boundaries appears to be outweighed by the impact of surface orientation.

Comparing TDS and NRA results can also give insight into the impact of the ion clusters (extracted ions are thought to be dominated by D_3^+ [120]) prevalent in DELPHI. TDS measurements did not record calibrated counts for masses this high, whilst NRA measurements rely on nuclear reactions with the deuterium rather than mass measurements, meaning deuterium in clusters of any size should be detectable. If there was significant deuterium content retained as D_3 clusters, retention determined via NRA would be notably higher, but this was not the case. Therefore, it seems likely that the weakly bound D^+ is separating from the D_2 on impact and TDS measurements of total inventory remains accurate despite the incident clusters being of higher mass than detected. The splitting of clusters would partly explain the large degree of HD present despite the lack of D^+ ions present in the plasma. Furthermore, as HD counts were consistently higher than D_2 counts, it seems the charge of the D^+ makes it more reactive and more likely to be retained within the material. Although there could be some contribution to HD counts from D_2 molecules splitting, trends in TDS data suggests this is not the dominant process. It would be expected that splitting of D_2 molecules would increase with energy, resulting

³TDS results did not present a clear increase in retention values for high energy ions, although this could be a result of deuterium remaining within the sample or varying hydrogen contents within the diamond impacting uptake.

in greater HD counts at higher energy, but the inverse of this is seen. D_2 counts increase with energy whilst HD counts show no clear trends, suggesting the higher energy is more effective at implanting the less reactive D_2 and D^+ is preferentially retained, particularly at low energy.

3.2.2 Pristine Boron Doped Diamond

Table 3.6: Calibrated total counts from TDS data of PC diamond samples of different gas phase boron concentrations implanted with deuterium at 400 V.

TDS Total Counts			
Sample	HD ($\times 10^{15} \text{cm}^{-2}$)	D_2 ($\times 10^{15} \text{cm}^{-2}$)	Total D ($\times 10^{15} \text{cm}^{-2}$)
2-3a, High [B]	5.12	1.29	7.70
2-3b, High [B]	6.56	1.53	9.61
2-11a, Low [B]	5.74	6.28	18.3
2-4b, Resid. [B]	5.47	7.11	19.7

TDS spectra for both the residually and lightly doped sample are remarkably similar as can be seen in Fig. 3.9. This aligns with SIMS measurements, that suggested these two samples may have very similar boron concentrations. As discussed in 3.1.3, Raman measurements suggested lightly and residually doped samples were of different boron concentrations, whilst SIMS measurements of comparable samples suggested very similar concentrations. The TDS results align better with SIMS results despite Raman spectra being of the actual sample (2-11a) rather than the equivalent sample tested in the SIMS. Either this suggests that the two concentrations produce similar TDS spectra and differences in concentration arose from varying levels of residual boron in the chamber. Or, that the timings of the measurements and the small probe areas of SIMS and Raman spectroscopy are impacting concentration values measured. A more systematic approach to determining boron concentrations would be required to understand this conclusively.

The residually and lightly doped samples were the only samples tested in which D_2 counts exceeded HD counts as shown in Table 3.6. The biggest contribution to D_2 was seen from the large peak at 450 K. In alignment with conclusions from undoped samples, the smaller HD peak and absence of H_2 peak suggests desorption from an area with minimal hydrogen. The relatively low temperature of this peak suggests it could be a result of D_2 weakly bound to the surface of the sample. It is thought that incident D_3^+ clusters will split upon impact into D_2 and D^+ , and that the single D^+ will be able to penetrate deeper into the surface and form more stable bonds resulting in the majority of HD being desorbed at higher temperatures than this initial D_2 peak. The 450 K D_2 peak was not observed from undoped samples. It is thought this is because the desorption from the surface will be significantly impacted by the topology of the sample and the rough surface of the unpolished boron doped films will offer more trapping sites for D_2 than the smooth, polished undoped samples. The highly doped sample would also present a similar degree of roughness, but this low temperature peak was not observed. TDS experiments of boron doped graphite concluded the inclusion of boron resulted in a more porous material [68],

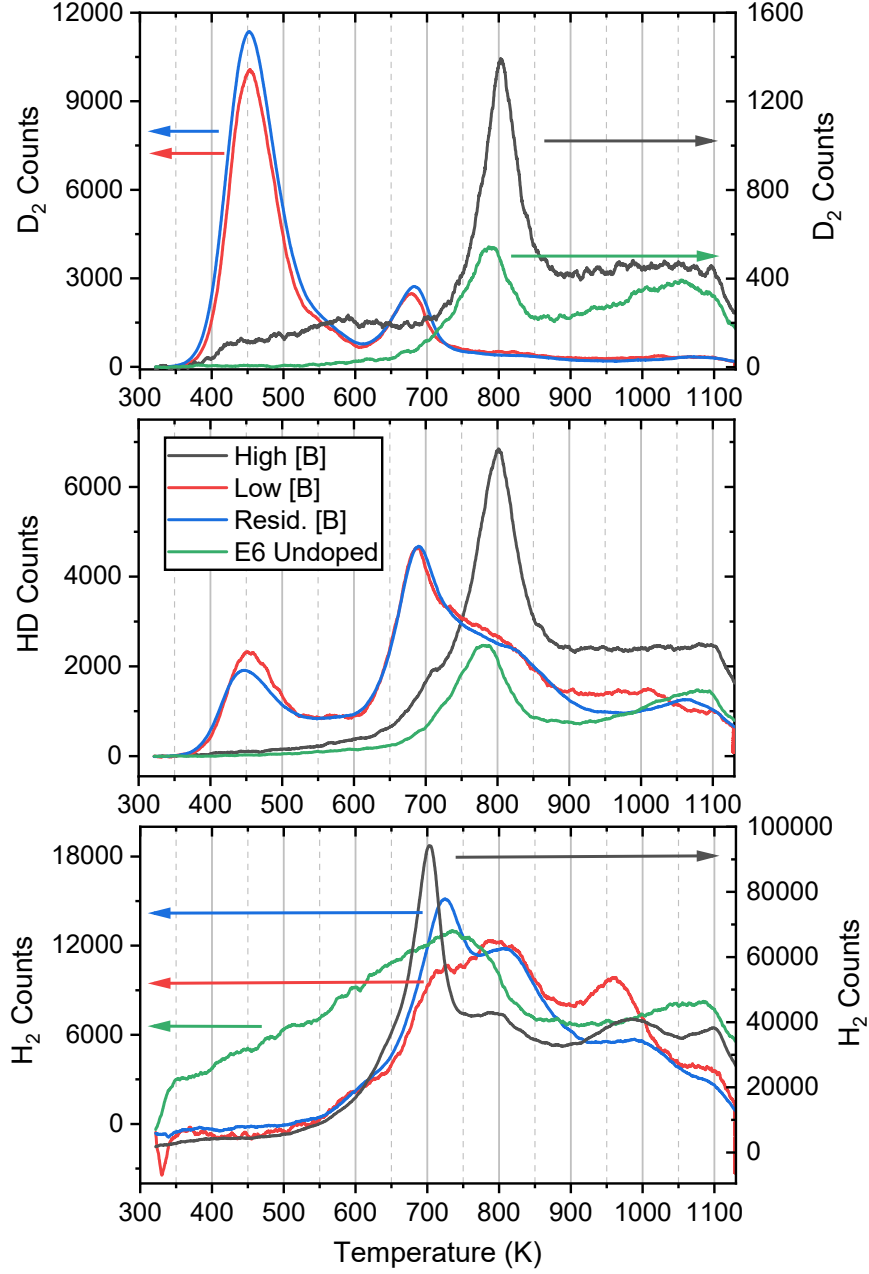


Figure 3.9: Thermal desorption spectroscopy spectra for masses 2, 3 and 4 attributed to H₂, HD and D₂ respectively, for three diamond samples (2-3b, 2-11a and 2-4b are high [B], low [B] and residual [B] respectively) grown with different levels of gas phase boron concentrations and a undoped reference sample. A ramp rate of $\beta = 10 \text{ K min}^{-1}$ was used. Where there are large discrepancies in counts two axes have been used, coloured arrows indicate the axes each spectrum corresponds to.

resulting in more deuterium located deeper within the material. Therefore, it is possible that D_2 was simply able to travel deeper into the material. Alternatively, the surface of the highly doped sample is dominated by (111) facets, whereas lower doped samples display a mix of (100) and (111). Applying the same logic as for the undoped samples, the lack of hydrogen peak at this temperature suggests this is corresponding to inter-grain desorption. Therefore, it could be a result of desorption from (100) facets, and the increase in counts compared to undoped (100) samples is simply a result of roughness. However, computational results presented in Chapter 4 suggest the formation of a disordered layer on the surface, reducing the impact of specific atomic structure and making this explanation less likely.

At intermediate temperatures (600–900 K), two distinct peaks at around 700 and 800 K can be seen across all the sample and species counts. The peak at 700 K was not observed for previously tested undoped samples and is most apparent in the H_2 counts of the highly doped sample. Whilst hydrogen concentrations are still thought to be highest within grain boundaries, for boron doped diamond, boron can also form hydrogen complexes [176], offering another potential trapping mechanism. It seems likely that this peak is a result of desorption from boron-hydrogen complexes formed during growth, explaining why it was not observed in undoped diamond, and it is most prominent in the sample with the highest degree of doping. A subtle peak in HD counts is seen at 700 K for the highly doped sample but is more prominent in the lightly and residually doped samples. This could be a result of the lower hydrogen content resulting in more available binding sites around boron atoms for incident deuterium, whilst the higher hydrogen content of the highly doped sample may offer less potential binding sites for incident deuterium. Similarly, a D_2 peak at this temperature is only seen for the lightly and residually doped samples, suggesting less of the binding sites from the boron are occupied by hydrogen in these samples.

The 800 K peak aligns well with desorption assigned from grain boundaries for undoped samples. The peak position is around 20 K higher compared to the undoped sample; this is likely due to the molybdenum substrate slightly restricting heat transfer compared to the free standing undoped diamond samples. A small increase in peak position would be expected, in a similar manner to the AlN protective coating [151]. For the highly doped sample, this peak is most prominent in HD and D_2 counts. Residual and lightly doped samples both show a shoulder across this temperature range in HD counts, but no notable D_2 counts. This indicates a wide variety of different binding sites and could suggest different grain boundaries are present from the mix of facets.

At temperatures above 900 K, a plateau in HD counts is observed, most prominently in the highly doped sample, which also shows a similar plateau in D_2 counts. In undoped samples, continuous outgassing of H_2 was observed when held at the max temperature (1273 K) which was thought to be a result of diffusion. This contrasts with what is observed here, where the plateau is present whilst the temperature is being increased, and counts drop off once held at a fixed temperature. The lack of continuous outgassing at fixed temperature suggests this is not a

diffusion effect. Instead, this is thought to indicate a wide variety of different binding sites available as a result of the inclusion of boron and the disruption this causes to the structure. Despite the thinking that boron doped diamond is more porous than undoped diamond [65, 68], the lack of diffusive behaviour and apparent availability of varying binding sites suggests boron dopants may act as trapping sites, reducing the rate of diffusion. Tailored diffusion experiments concluded hydrogen can become trapped in structural defects in boron doped diamond, particularly for low boron concentration [177].

Hydrogen contents also varied significantly between samples, with the highly doped sample giving significantly higher counts than any of the other samples. In part, this is thought to be a result of boron-hydrogen complexes, as well as the higher concentration of boron in this sample creating more disruption in the lattice. A consequence of this higher hydrogen content is a higher proportion of total deuterium is made up of HD counts rather than D₂. It is also of note that HD counts are comparable across all samples and show less variation than H₂ and D₂ counts between samples. This could suggest similar number of binding sites for hydrogen/deuterium are available in all samples, with samples of lower H₂ content will display higher D₂ counts but the number of HD counts remains consistent.

Overall, the boron doped sample gave higher retention values than undoped samples, however a large contribution to this for lightly and residual doped samples is thought to be from the roughness of the unpolished surface. A peak not observed in undoped samples was seen at 700 K, this is thought to be a result of desorption around boron impurities, either from disorder structure or boron-hydrogen complexes themselves. A peak at 800 K aligned with a peak seen in undoped diamond and is thought to be a result of grain boundary desorption. For the highest boron concentration in particular, a wide range of binding sites were observed at temperatures greater than 900 K which likely are result of structural traps from the presence of boron dopants.

In the context of fusion and the potential uses of diamond, it is unlikely the highly doped sample would be of interest. Lightly doping samples has been shown to help avoid arcing for plasma facing materials [65] and would be required for sensor applications. Although retention for lightly and residually doped samples is the highest seen, the biggest contribution was from a low temperature peak dominated by D₂ counts. As this is thought to be surface desorption (which will be impacted by surface topology) polishing samples could help reduce overall retention values significantly in this case.

3.3 Chemical Etching

SEM images were taken pre and post deuterium exposure for selected samples. Prior to exposure, the polished surfaces of undoped samples offered few features of interest. Post exposure, some of the samples presented with small nanoparticles on the surface as seen in Fig. 3.10. The absence of these particles in BSE images suggests they were likely to be carbon based, which




Figure 3.10: SEM images taken of a PC (100) diamond sample post exposure to deuterium plasma (sample 1-1e). **A** - 15 kV secondary electron image of nanoparticles only present post exposure. 5 nA probe current, 9.8 mm working distance, $\times 5000$. **B** - 15 kV backscattered electron images of A. 30 nA probe current, 9.8 mm working distance, $\times 5000$. **C** - 20 kV secondary electron image of a nanoparticles at high magnification. 10 nA probe current, 7 mm working distance, $\times 100000$. **D** - 15 kV secondary electron image showing shadow of cross used to secure sample, nanoparticles were not present within the cross which gave decreased conductivity. 30 nA probe current, 9.8 mm working distance, $\times 250$.




Figure 3.11: SEM images taken of a lightly boron doped diamond film post exposure to deuterium plasma (sample 2-13a). **A** - Nanoparticles present outside of shadowed area post exposure, 15 kV, 5 nA probe current, 9.9 mm working distance, $\times 10000$. **B** - No nanoparticles present in shadowed area, 15 kV, 30 nA probe current, 9.9 mm working distance, $\times 5000$. **C** - Shadow of cross from wire used to secure sample, 15 kV, 30 nA probe current, 9.9 mm working distance, $\times 250$.

was confirmed with XDS analysis, as regions with nanoparticles present were indistinguishable from pristine regions. Initially, it was unclear whether this was a result of the DELPHI exposure or an artifact from the deposition of the film. It seemed unlikely to be the latter, as the sample was polished to a high level and images of other equivalent samples showed no evidence of these. This was confirmed by the presence of a 'cross' on samples post exposure in DELPHI (Fig. 3.10D). A wire cross was used to secure samples to the stage during exposure and, when imaging, a faint outline of a cross could be observed on some of the samples. When present, nanoparticles were only observed in exposed areas that were unobscured by the wire cross, as can be seen in Fig. 3.11. One explanation could be that these nanoparticles are amorphous carbon which has been etched during the exposure and redeposited on the surface. The wire cross shielded a portion of the sample from this process leaving the surface clear of these particles and a shadow of the cross on the sample. At operating pressures (approximately 1 mTorr), it is unlikely that large numbers of individual atoms would aggregate together into a nanoparticles, as the expected mean free path at operating pressures is of order 10 cm. However, the details of the gas recirculation system are unknown and it is possible higher pressures could be used during this process. Alternatively, these particles could have been present below the surface, and etching of the material around them meant they became visible. Although the exact explanation is not known, it is thought that their presence indicates some degree of etching has occurred. These particles are not dissimilar to those observed by Guenette *et al.* [77], in a study where diamond was exposed to both deuterium and helium plasmas. Again, it was concluded that this was evidence on chemical erosion and subsequent redeposition, as these particles were only observed when hydrogenic species were within the plasma and not present for the pure helium plasma.

An increased conductivity was also observed in areas not obscured by the cross in Fig. 3.10D. Here, the cross region can be seen as much brighter than the rest of the sample suggestive of




Figure 3.12: EM images taken of PC (100) diamond samples post exposure to deuterium plasma. **A** - Sample 1-1a, **B** - Sample 1-1b, **C** - Sample 1-1c, **D** - Sample 1-1d, **E** - Sample 1-1e, **F** - Sample 1-1g, **G** - Sample 1-2a, **H** - Sample 1-2b, **I** - Sample 1-2c. All images taken with a 15 keV beam, $\times 10000$ objective and a working distance of approximately 10 mm. All but subfigures H and I used a probe current of 5 nA, H and I used 30 nA.

a lack of free electrons to interact with. Terminations can also impact surface charging effects, but this is unlikely to be the case here, as the TDS would likely have destroyed the hydrogen termination present post exposure. Instead, it is thought to be a result of the amorphous carbon nanoparticles present on the surface. These are likely to display a higher sp^2 content and therefore increase surface conductivity, helping minimise charging effects. Raman spectra taken post exposure indicated no increase in sp^2 , suggesting any changes are limited to the surface rather than the 1 μm probe depth of the Raman laser.

There are no obvious trends across the samples tested as to why these particles are only present some of the time. Fig. 3.12 shows all SEM images of undoped samples taken post exposure. Most samples pictured in this figure were equivalent, except for C (1-1c) which had a piece of multi layered graphene (MLG) on the surface from a previous experiment, and 1-1b which had been annealed preexposure. Implantation conditions also remained the same other than varying the ion energy. No nanoparticles were observed on samples 1-1g, 1-2a and 1-2c (Fig. 3.12C, F, G and I) which corresponded to implantations at 400, 1000, 800 and 400 eV respectively. The other images all show evidence of nanoparticles but in different forms. Sample 1-1a gives a uniform high density of small particles, in contrast to the collections of particles observed in 1-1b. Despite the MLG layer, sample 1-1c (Fig. 3.12C) featured very few particles although some are present. The presence of graphitic phases would typically offer an increase in chemical etching, but in this case, it has not resulted in an increase in redeposited carbon on the surface. Sample 1-1d (Fig. 3.12D) presents with similar size particles but a much sparser distribution with areas seemingly containing none at all. Samples 1-1e and 1-2b (Fig. 3.12E and H) present with similar size particles but the latter seemingly gave a higher concentration. The implantation energies were 400, 400, 400, 400, 600 and 300 eV for Fig. 3.12A, B, C, D, E and H respectively. There is no clear correlation between implantation energy and the presence of these nanoparticles.

The explanation behind these nanoparticles is currently unknown. The lack of particles below the shadowed area clearly indicates they are a result of the plasma exposure in DELPHI; however, it does not seem to be a result of the implantation energy used. Furthermore, samples 1-1a and 1-2c (Fig. 3.12A and Fig. 3.12I respectively) were equivalent samples implanted at the same energy, but one exhibited the particles and the other did not, further suggesting that ion energy is not the sole factor determining the presence of the particles. It seems likely that there were other differences in the exposure conditions of the samples, these could either be operational (variation in a controlled variable) or more random (variation in an uncontrolled variable). An example for the former could be the use of 'recycling' or 'once-through' modes. As described in 2.2, these are different gas circulation modes available in DELPHI that either recirculate previously used gas (designed for use with tritium) or remove it through an exhaust. Recirculating gas could result in an increase in redeposited carbon; however, it is believed that all samples ran in 'Recycling' mode. Other factors could also impact the presence of the particles such as contamination of the chamber from previous samples - information that was not possible to obtain.

Table 3.7: Summary of deuterium retention of a variety of diamond samples exposed to deuterium ions accelerated across a given anode voltage in DELPHI. Retention values determined via TDS separates retention of HD and D₂, whereas retention determined via NRA only includes total deuterium. Samples included undoped polycrystalline (PC) and single crystal (SC) samples provided by Element 6 (E6), as well as boron doped diamond grown in the University of Bristol Diamond Lab. Face perpendicular to ion implantation is shown via Miller indices.

Sample	Description	Retention Values			
		Anode Voltage (V)	HD ($\times 10^{15} \text{cm}^{-2}$)	D ₂ ($\times 10^{15} \text{cm}^{-2}$)	Total D ($\times 10^{15} \text{cm}^{-2}$)
1-1f	PC E6 Undoped (100)	200	3.62	0.925	5.47
1-2b	PC E6 Undoped (100)	300	2.99	0.579	4.15
1-1a	PC E6 Undoped (100)	400	1.72*	0.652*	3.02*
1-1e	PC E6 Undoped (100)	600	3.21	0.933	5.08
1-2a	PC E6 Undoped (100)	800	3.67	1.19	6.05
1-1g	PC E6 Undoped (100)	1000	3.18	1.25	5.67
1-3a	PC E6 Undoped (100)	-	0.249	0.041	0.331
1-5b	PC E6 Undoped (100)	400	-	-	1.81 ± 0.13
1-5c	PC E6 Undoped (100)	800	-	-	2.97 ± 0.16
1-4b	SC E6 Undoped (110)	400	-	-	1.60 ± 0.12
1-4c	SC E6 Undoped (110)	800	-	-	3.18 ± 0.17
2-3a	PC High [B] (100)	400	5.12	1.29	7.70
2-3b	PC High [B] (100)	400	6.56	1.53	9.61
2-11a	PC Low [B] (100)	400	5.74	6.28	18.3
2-4b	PC Resid. [B] (100)	400	5.47	7.11	19.7

3.4 Conclusions

Both undoped and boron doped diamond samples were exposed to deuterium ions in DELPHI. Deuterium retention was determined with either TDS or NRA, and values for all samples can be seen in Table 3.7. In undoped sample, peaks in different temperature regions of TDS spectra were attributed to differing deuterium locations, rather than bonding types, through comparison of hydrogen and deuterium spectra. Peaks seen in both hydrogen and deuterium spectra were attributed to grain boundary desorption and were observed at low temperatures (600–850 K) as well as high temperatures (>1100 K). It is believed the latter peak is a result of hydrogen diffusion through grain boundaries once a sufficient temperature was reached. A peak was observed between 850–1100 K in D₂ spectra which was assigned to deuterium from within the grains and was only present for ion energies ≥ 600 eV, suggesting some degree of ion damage may have occurred. Across all samples, total HD counts were more consistent than H₂ and D₂ counts. This could suggest a similar number of available hydrogen/deuterium binding sites in all samples, and the amount of D₂ desorbed is linked to the initial hydrogen content. Total retention of undoped samples was comparable, or lower, than retention values of other PFM candidates tested in the same experimental setup [120, 174, 175]. NRA results were in good agreement with

TDS results, with both giving retention values to the same order of magnitude. Although a full depth profile was not possible due to limits of resolution of the method, it can be said deuterium was within the top 10 nm of the diamond confirming this is ultimately a surface interaction and minimal bulk changes would be expected at these energies.

The inclusion of boron dopants increased retained deuterium as observed in the work of others [65]. The most significant increases are thought to be a result of an increase in surface roughness of the unpolished films. Another peak that is not present for undoped samples was observed at 700 K, this is thought to be a direct result of the inclusion of boron through either the formation of boron-hydrogen complexes or the inclusion of boron adding disorder to the atomic structure. At high temperatures, limited evidence of diffusion was observed in contrast to undoped samples.

Negligible changes were observed in Raman spectra taken prior to and post exposure, suggesting minimal changes in sp^2/sp^3 content within the top micron or so of the surface. However, the limited penetration depth of deuterium ions means any damage is likely to be constrained to very shallow depths, and is therefore unlikely to be detectable with this technique. SEM images did present some differences after exposure to the deuterium plasma, most notably the presence of carbon nanoparticles which are thought to be a result of etched and redeposited carbon. No clear trends were seen between which samples did or did not present these, highlighting the complex nature of chemical etching mechanisms. If a full understanding of these could be made in the context of fusion, conditions faced by the diamond could be controlled to minimise etching.

Raman spectra revealed significant changes to bonding on the surface of ion irradiated samples, with higher doses resulting in a higher degree of sp^2 carbon. Raman spectroscopy was also used to evaluate boron contents of boron doped diamond films grown under different conditions. Raman spectra suggested very different boron contents for each condition, whereas SIMS and TDS spectra implied that two of the samples were of very similar concentrations. The timeline of measurements meant it is possible that some changes to the material occurred during exposure and TDS measurements, but it is not possible to make firm conclusions regarding this.

3.5 Future Work

Despite preparing ion irradiated samples, it was not possible to expose these in DELPHI due to unforeseen delays during an upgrade to the setup. Results from these samples could indicate how retention would vary during operation of diamond as a PFM. Clear structural changes were observed in both SEM images and Raman spectra post irradiation, with the latter indicating an increase in sp^2 carbon for the higher dose samples. Others [71] have reported a decrease in sp^2 carbon on exposure to deuterium plasma as a result from the preferential etching compared to sp^3 . If such a decrease was observed, this would highlight chemical etching of the surface was occurring in DELPHI, further supporting the idea that carbon nanoparticles on the surface are likely to be redeposited etched carbon from the surface. For the ion irradiated samples, TEM

could be used to assess changes to the structure on an atomic scale would allow damage features as a function of depth and dose to be explored.

Chemical etching is one of the key concerns surrounding in-vessel use of carbon materials and, as the observation of carbon nanoparticles is thought to indicate chemical etching, a study devoted to exploring this would be of interest. Porro *et al.* observed a promising 50% reduction in etch rate of diamond compared to graphite [64], and it would be interesting to explore if this could be reduced further through alterations to the diamond (such as doping).

As a result of experimental delays, it was not possible to test as many samples as desired. Testing boron doped samples of lower, more relevant, concentrations as well as unpolished undoped samples should clarify the impact of boron dopants. With the high concentrations tested, as well as the uncertainty surrounding exact concentrations, it was challenging to pick out exact trends from the boron doped samples. Use of a microwave CVD reactor would allow for finer control over the boron content of these samples, whilst a more systematic approach to determining concentration with Raman and SIMS prior to exposure should clarify concentrations.

Peaks at certain temperatures were assigned to grain boundaries based on the presence of hydrogen peaks at the same temperature. This could be tested further by taking TDS measurements of SC samples. Preferably, these would need to be of a comparable size to PC samples ($10 \times 10 \text{ mm}^2$) to maintain a reasonable signal to noise ratio for the TDS counts.

SIMULATION OF HYDROGEN'S INTERACTION WITH DIAMOND

MD simulations were performed using the LAMMPS code [80, 81] to complement and develop the understanding of experimental results. A variety of different simulations were carried out to give a broad understanding of hydrogen isotope retention in diamond. Continuous bombardment simulations, in which a fluence of deuterium was incident upon a diamond surface, were performed to offer insight into retention and etching mechanisms that may occur during experimental work and within reactors. A series of single bombardments were also carried out, in which the impact of varying diamond orientation and temperature as well as incident hydrogen isotope mass, energy and angle was explored. Performing repeated single bombardments gave clarity when compared to continuous bombardment, revealing the impact of variables which can be hard to distinguish in a damaged material.

PC diamond shares many of its properties with SC diamond whilst being easier to produce on larger scales. It may prove a more viable option depending on the application, but the impact of grain boundaries will need to be assessed. To determine this, four common tilt grain boundaries were reconstructed from TEM images taken from the literature [178]. Both single and continuous bombardment simulations were performed on these grain boundaries and compared to bulk diamond.

Transport of retained hydrogen isotopes is also a key consideration for fusion applications of diamond. As such, a series of simulations were carried out to determine diffusion characteristics in both bulk and grain boundary structures. In diamond grown via CVD, the coalescing of hydrogen terminated surfaces during growth results in a large amount of hydrogen concentrated within the grain boundaries, meaning they could play an important role in determining overall diffusion characteristics. For implanted hydrogen isotopes, exploring diffusion characteristics may indicate whether this concentration within grain boundaries would still be expected and how retained hydrogen may diffuse through the material. This work combined with bombardment

simulations helped to build a picture of the hydrogen etching, retention and transport of PC diamond.

The work presented in this chapter has been adapted from the following papers [150, 179, 180]. A declaration of the contribution can be found at the start of each section when applicable. The author acknowledges the support of EPSRC, UKAEA and funding by the RCUK Energy Programme (Grant No. EP/W006839/1). This work was carried out using the computational facilities of the Advanced Computing Research Centre, University of Bristol - <http://www.bris.ac.uk/acrc> and the Cambridge Service for Data Driven Discovery (CSD3) - <https://www.csd3.cam.ac.uk>.

4.1 Testing

A series of tests were carried out to help determine parameters required to maintain accuracy whilst minimising computational load. This was particularly important for continuous bombardment simulations, which were both computationally expensive and inherently unstable as every incident atom added additional energy to the system. Generally, if a reasonable balance has been found between accuracy and computation time, moving towards more physical parameters (such as a larger system size or smaller timestep) should not significantly impact results. The following were parameters of focus:

- **Size** - As the size of the simulated system is far smaller than the physical material being simulated, it must be ensured that results do not vary significantly if the size of the system increases. Using periodic boundary conditions and frozen regions can help to mimic larger systems without the computational cost of additional atoms.
- **Flux** - Limitations on the physical time it is feasible to simulate means very high fluxes are typically used to achieve the moderate fluence. A flux must be selected at which no significant changes are observed when dropping to lower fluxes.
- **Timestep** - The largest possible timestep whilst maintaining accuracy is desirable, so it must be ensured that using a smaller timestep does not impact results. Variable timesteps can also be used will reduce the timestep from some default value if a condition is fulfilled (such as an atom moving more than a certain distance in a single step).
- **Fixes** - A fix is an operation applied to the system during a timestep [181], typically used to control temperature (thermostatting), pressure, position etc. of the system during the simulation. Appropriate selection of these alongside the parameters that control the fixes (such as the T_{damp} parameter used for thermostatting, see 2.4.4) are key to achieving physical accuracy. This is particularly important for continuous bombardment simulations, as each bombardment adds energy to the system and consequently a temperature rise. T_{damp} must be sufficiently fast to allow the system to return to the target temperature

between bombardments, but slow enough to avoid unnecessary, large, rapid fluctuations in temperature.

- **Potentials** - The potential or force field selected will dictate the forces experienced between atoms. Selection of a particular potential is largely dictated by the type of atoms included in the simulation, as typically a potential will only be accurate for certain atom types. The Reactive Empirical Bond Order (REBO [141, 142]) and Adaptive Intermolecular Reactive Empirical Bond order (AIREBO [145]) potentials are commonly used for modelling systems containing carbon and hydrogen. The latter is more accurate, particularly at higher energies, but is computationally slower.

Analysis of test simulations used the same code as used for bombardment simulations. Deuterium retention, depth profiles and sputtered carbon were all calculated, whilst other variables such as temperature were monitored throughout simulations. For testing, standard parameters consisted of a $6 \times 6 \times 8$ unit cell (UC), (100) orientated, block of diamond, maintained at a temperature of 300 K by simulating a canonical ensemble via a NVT fix with a T_{damp} of 0.1 ps. Periodic boundary conditions were used in x and y dimensions and a finite boundary condition was applied in z . A fixed timestep of 0.1 fs was used. In testing, the diamond was bombarded with 10–100 eV deuterium atoms at a flux of $7.28 \times 10^{26} \text{ m}^{-2} \text{ s}^{-1}$, achieving a fluence of $2.18 \times 10^{20} \text{ m}^{-2}$. Throughout testing, the number of bombarding atoms was scaled to maintain fluence and flux for different surface areas. The REBO potential was used for carbon-carbon and carbon-hydrogen interactions. Initial simulations were carried out using the above parameters, and were compared to simulations in which a parameters relating to the system size, flux, thermostat and timestep had been adjusted. If changing a parameter towards a more physical value made a notable difference to the simulation results, the new parameter was used going forward.

Datafiles for structures containing grain boundaries and different orientations of diamond also were tested. Two methods were used to produce these ('Replicated' and 'Trimmed'). The main focus surrounding testing the new structures was stability, particularly at high temperatures, as incorrect atomic positions and misalignment at the simulation box boundary resulted in unstable, disordered structures.

4.1.1 Temperature Control

In the first set of testing, it was observed that high flux simulations were very unstable, resulting in large volumes of the diamond etching away rapidly. The loss of carbon corresponded closely with long periods of elevated temperatures as seen in Fig. 4.1. The temperature spikes were typically observed in the latter parts of simulations. A faster T_{damp} had minimal effect in reducing the spikes. It transpired there was an error in the input file, namely, the manner in which the bombarding atoms were created and given a velocity. The method used is depicted in Fig. 4.2A, where a region was defined above the sample in which the bombarding atom was




Figure 4.1: Total number of carbon atoms and temperature from a test simulation of a $6 \times 6 \times 8$ unit cell block of diamond bombarded with 1000, 30 eV tritium atoms. Target temperature for this simulation was 300 K, but a combination of an overly large temperature damping parameter and additional energy in the system meant the temperature varied wildly, resulting in significant loss of carbon atoms. Carbon atoms below a cut off of 2 Å from the surface were counted, increases in the number of carbon atoms are a result of atoms relaxing back below this cut off.

created with a random position. Atoms within the region (initially just the bombarding atom as in Fig. 4.2A1) were then given a downward velocity corresponding to the target energy. However, any reflected bombarding atoms and etched carbon atoms within this region would also be given the same velocity, resulting in numerous atoms being propelled towards the diamond.

When working correctly, the energy increase of the system, ΔE , as a result of the bombarding atom should simply be energy of the bombarding atom, E_B . However, if N_C carbon atoms and N_R reflected bombarding atoms are given the same velocity as the bombarding atom, the resultant energy increase is given by

$$(4.1) \quad \Delta E = E_B(1 + N_R + N_C \cdot \frac{m_C}{m_B})$$

where m_C and m_B are the masses of carbon and the bombarding atom respectively. Taking deuterium as the bombarding atom, if just ten carbon atoms are given the same velocity, ΔE would be 61 times greater than the expected energy change. The system is simply not able to relax fast enough from such a big energy increase between bombardments, resulting in a series of temperature spikes and dramatic losses of carbon atoms as seen in Fig. 4.1.

The solution to this issue can be seen in Fig. 4.2B, where new bombarding atoms are created as a different atom type in order to distinguish them from all other atoms within the simulation. This atom can then be grouped in isolation based on its atom type, then changed to the bombarding

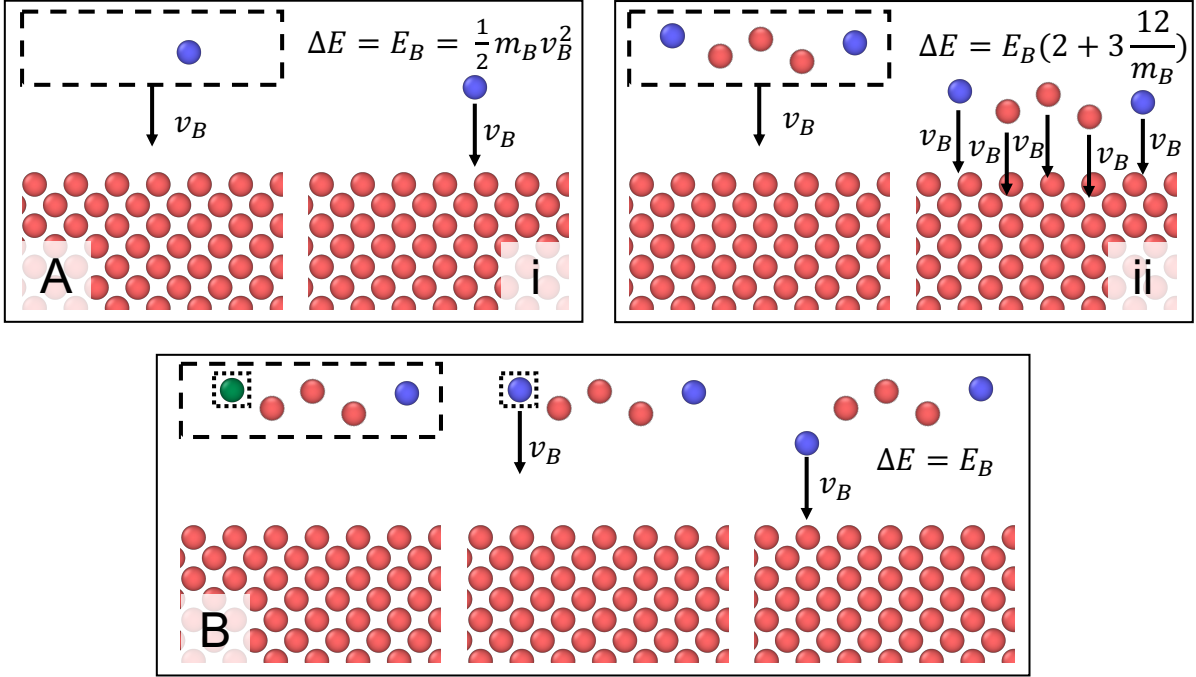


Figure 4.2: Figure demonstrating two methods (**A** and **B**) of giving a bombarding atom (blue, of mass m_B) a velocity, v_B , corresponding to an energy of E_B . In **A**, a bombarding atom is created in a random position above the diamond (red atoms) and any atoms within this region are given a velocity. This can work as intended as seen in **Ai**. However, when reflected or etched atoms are also found within the box (as seen in **Aii**), this method results in an excessive change in system energy, ΔE . In **B**, the new atom is of a different atom type (green), allowing it to be grouped in isolation. The atom within this group can then be changed to the desired bombarding atom type and given the desired velocity. Base diagram rendered in OVITO [182].

atom type and given its desired velocity. This process can be repeated for each bombarding atom, allowing the velocity to be given to only one atom each time. Deleting any atoms within this box prior to creation of the next atom also avoids this error, with the additional benefit of clearing the path of the new atom. Operating at a lower, more physical flux, would allow these atoms to leave the system anyway, so deleting atoms makes the high fluxes required in MD more physically accurate. Both of these methods were used in future bombardment simulations.

4.1.2 Size

Increasing system size comes at a significant computational cost, not only from the increase in atoms, but a larger surface area requires additional bombardments to reach an equivalent fluence. The main concern of a system that is insufficient in size is its stability and ability to moderate temperature - as each bombardment will add a proportionally higher energy for smaller

systems. To increase stability of the system, a 1 UC¹ thick slice at the bottom of the diamond was frozen in place, preventing any movement of the system. Maximum penetration depth was monitored closely to ensure no bombarding atoms came within 1 UC of the frozen region. The frozen layer, along with periodic boundary conditions in x and y , meant a stable system was achieved for typical surface areas of 8×8 and 5×5 UCs for multiple and single bombardments respectively. Therefore, these were taken to be the minimum system sizes for future simulations.

4.1.3 Flux

To test the impact of flux for continuous bombardment simulations, the time between bombardments was varied on the same $8 \times 8 \times 8$ UC, (100), diamond block. When increasing flux, the main consideration is the simulation's ability to return to the target temperature in between bombardments. At high fluxes, the energy/temperature in the system gradually creeps up, resulting in an unstable system and increased etching. A shorter T_{damp} helps to bring the temperature back to the target temperature faster, however this also results in rapid temperature fluctuations which are not physically accurate and could result in increased etching. Higher flux also can become a concern when sputtered atoms don't have time to leave the simulation box, potentially resulting in obstruction of the next bombarding atom. This effect can be minimised by deleting etched atoms within the creation box before the next incident atom is created.

Fluxes of order 10^{29} Atoms $\text{m}^{-2} \text{s}^{-1}$ (0.3–0.5 ps between bombardments) paired with a T_{damp} of 0.1 ps was selected to effectively maintain target temperature and allow a reasonable fluence to be reached without unfeasibly long simulations. Flux incident on different structures did vary slightly due to small differences in surface area. Due to limitations of computational modelling, this flux is far higher than seen in DELPHI or in reactors [183]. However, testing showed no significant changes were observed from decreasing flux by a factor of four (for fixed fluence). Lower fluxes than this would be counterproductive, as this would limit the fluence that would be possible to simulate.

4.1.4 Timestep

Increasing the timestep did not have an incremental effect on the results, rather, the simulation simply failed if the timestep became too large. When the timestep is too large, the numerical integration underpinning the simulation becomes unstable leading to sudden failure. It is likely there are some subtle changes before this point, particularly for fast moving atoms. Therefore, a variable timestep was selected which ensured no atoms moved more than 0.1 Å in a single timestep, with a default value of 0.1 fs.

¹The notation 'UC' is used to refer to the length of diamond's conventional unit cell (or lattice constant) of 3.567 Å.




Figure 4.3: Potential energy of, and force experienced by, a hydrogen placed at a varying separation, R , from a CH_3 molecules to create a methane molecule on minimisation using the AIREBO potential.

4.1.5 Potentials

As discussed in 2.4.2, REBO and AIREBO potentials were used for this work. These potentials are commonly selected for modelling interactions between carbon and hydrogen and are capable of reproducing the different structures possible with carbon. They have been selected as a balance between physical accuracy and computational speed. Faster, pair-potentials neglect multiatomic interactions meaning they struggle to recreate structures typical of carbon materials (diamond, graphene, nanotubes etc.) accurately. ReaxFF [184] is a more computationally heavy potential which can offer increased accuracy compared to AIREBO particularly for unbound atoms. But, as only carbon and hydrogen (bound atoms) are included in these simulations this was not deemed necessary, and the use of more computationally heavy methods would limit the scale and number of simulations that could be performed. REBO was used for bombardment simulations as no differences could be seen when comparing to equivalent simulations run with AIREBO. The decreased computational complexity of the REBO potential took priority over the negligible improvements in accuracy observed for AIREBO, as it was deemed more important to carry out longer simulations with a lower flux.

Diffusion calculations were performed with AIREBO as diffusion rates were noted to be faster with the supposedly more accurate potential. However, Tolladay *et al.* [185] highlighted some possible issues exploring interatomic forces when breaking carbon-carbon bonds and using these




Figure 4.4: Simplified depictions of pure twist (left) and a pure tilt (right) grain boundaries. In a twist boundary, the surface vector of the grain boundary plane, \hat{n}_B , is parallel to the axis of rotation, \hat{t}_R , whereas, in a twist boundary, \hat{n}_B is perpendicular to \hat{t}_R . A tilt-twist boundary is a combination of the two rotations. Recreated from [186].

potentials. For this work, these issues are of particular concern regarding the hydrogen diffusion simulations, which are effectively modelling carbon-hydrogen bond breaking. To test this, a series of very simple simulations were carried out. Simulations consisted of a methane molecule, where one of the hydrogen atoms was progressively moved further away. The whole system was then minimised. The potential energy of, and the force experienced by, the displaced hydrogen was outputted at every timestep in order to collect data points for as many separations as possible. Fig. 4.3 shows results from these simulations. Both energy and force show the expected general shape for bond breaking. Although some variation is seen in force around a separation of 1 Å (roughly the expected C-H bond length), the potential energy was well behaved and the AIREBO potential was deemed acceptable for diffusion simulations.

4.1.6 Orientated Diamond and Grain Boundaries

In order to simulate different surfaces and grain boundaries, diamond blocks of various orientations were created. Generally, a grain boundary has five degrees of freedom, making simulating such a large number of potential structures effectively impossible. Instead, the focus was put on grain boundaries common to CVD diamond. Grain boundaries broadly can be classed as tilt, twist or tilt-twist as shown in Fig. 4.4. In a tilt grain boundary, the axis of rotation is perpendicular to the direction incident to the boundary plane, whereas in a twist grain boundary these directions are parallel. A tilt-twist grain boundary is a combination of both of these transformations. Grain boundaries are commonly described by the orientation of the surfaces meeting, and a coincident site lattice (CSL, Σ) number [186]. Σ is a measure of the relative degree of misorientation between




Figure 4.5: Snapshots taken from two simulations used to determine diffusion coefficients within 221 Σ 9 diamond grain boundaries. Red atoms are carbon and yellow atoms are hydrogen atoms which had been added to the grain boundary. Snapshots were taken after a run of 50000 steps (28.25 ps) at 3000 K. The top grain boundary (**A**) was produced by replicating out orientated unit cells, ensuring periodic boundary conditions could be applied effectively. The bottom grain boundary (**B**) was created by rotating and trimming a larger diamond block, causing issues with atom alignment across boundaries of the simulation box and resulting in a disordered structure.

the two grains at a boundary and is given by the size of a repeating unit cell at the grain boundary compared to the normal unit cell size. A larger Σ corresponds to larger degree of misorientation and typically a smaller rotation angle.

Two methods were used to create grain boundaries and different surface orientations data files, referred to as ‘trimmed’ and ‘replicated’. Trimmed datafiles were made by rotating and trimming a larger diamond block. To form a grain boundary, a second block was also produced using this method of the desired orientation and placed next to the first block. For both grain boundaries and orientated diamond blocks an energy minimisation was performed at the end, and the final positions from this minimisation was used as atom coordinates.

Creating replicated datafiles involved making new unit cells orientated in different ways from which a larger block could then be replicated from. To do this, once again a large diamond block was created and rotated to the desired orientation. Approximate dimensions of the new unit cells were determined by measuring the distance where atomic positions repeated themselves. All atoms were then reduced into the new unit cell by plotting the remainder of their coordinates divided by the unit cell dimensions. The length of lattice vectors could then be adjusted to ensure atoms of equivalent positions lay in the correct coordinates. Duplicate atoms were then removed, leaving a new unit cell which could be replicated to give a block of the desired orientation. To create a twinned grain boundary, the same process was repeated for a second block which was then mirrored.

Although trimmed data files achieved the desired surface, trimming the edges of the large

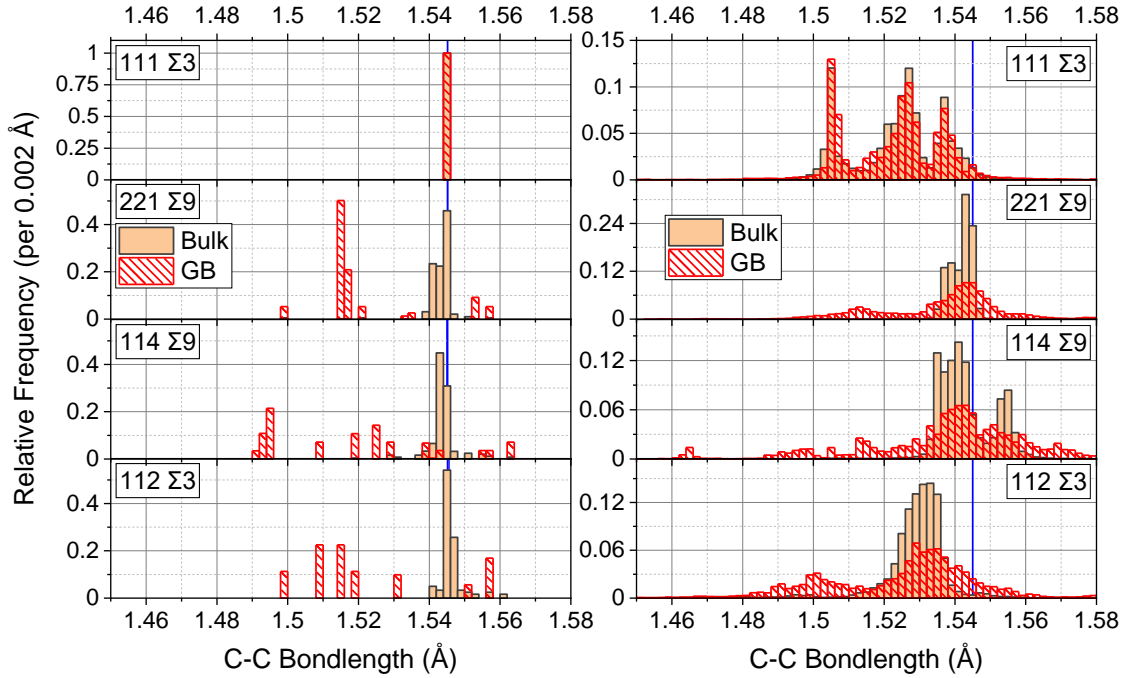


Figure 4.6: Carbon-carbon bond length distributions for grain boundaries made by replicating a oriented unit cell (left, referred to as a ‘replicated’ datafile) and grain boundaries made by trimming rotated diamond blocks (right, referred to as a ‘trimmed’ datafile). ‘GB’ distributions groups atoms within ± 2.5 Å of the centre of a grain boundary, ‘bulk’ atoms are any atoms outside of this range. Blue line indicates the bond length determined for a perfect diamond structure of 1.545 Å.

block resulted in imperfect atom alignment across the periodic boundary conditions despite efforts to correct for this. For localised testing such as the single bombardment simulations, the effect of this could be minimised by increasing the surface area to 8×8 UC surface area and restricting the area where atoms were incident on the surface to a central 3×3 UC area. However, for whole system tests such as the continuous bombardment and diffusion simulations, atomic alignment across the system boundaries of the simulation box needed to be correct. The result of imperfect alignment can be seen in Fig. 4.5, which compares the 221 GB constructed via trimmed and replicated methods at the end of a diffusion simulation at 3000 K. The trimmed datafile shows a disordered structure at the edges of the simulation box, particularly at the corners, whereas the replicated file maintains atomic alignment across the boundary even at these high temperatures. For diffusion simulations, these disordered regions resulted in trapping and restricted diffusion. Whereas, for bombardments, the disorder decreased channelling effects and increased etching.

Even at zero temperature, issues with the trimmed structures could be observed. Fig. 4.6 compares the bond length distributions for the replicated and trimmed datafiles (on the left and right respectively). For bulk atoms in trimmed datafiles, there is a significant deviation from the

expected 1.545 Å bond length found for perfect diamond. For the 112, 114 and 221 GBs² this was originally dismissed as the presence of the grain boundary creating additional strains and, in turn, distorting the bond length distribution. Furthermore, the peak around 1.51 Å has been attributed to sp²-sp³ bonds (averaged between planar and non-planar sp² geometries [187]) where the bulk atoms may be bound to atoms in the grain boundary. However, the bulk distributions for the 111 GB could not be explained by this. As discussed later, the 111 GB behaved largely the same as perfect diamond, and yet this bond length distribution did not align with the perfect diamond value. This finding acted as another indication that there was an error in the trimmed datafiles. The improvement of the replicated data files is clear to see, showing a much tighter distribution about the perfect diamond value of 1.545 Å for bulk atoms and the 111 GB. Although not presented here, similar improvements were observed in bond angle distributions.

For grain boundary data files, two methods were used to establish the optimal distance between the grains. The first method simply selected the gap that resulted in the lowest total system energy upon minimisation. The second method considered atom displacements along x (across the grain boundary) during the minimisation to see whether the periodic boundary conditions exhibited a compressive or tensile force on the system - a compressive force may indicate a larger gap was required and visa versa. The gap at which the force moved from compressive to tensile was selected, this generally corresponded closely to the minimum energy and at a distance similar to the C-C bond length in diamond (1.54 Å). For diffusion calculations, the difference this gap makes is expected to be minimal. As the temperature was controlled by simulating an isothermal-isobaric (NpT) ensemble the system volume, and therefore the grain boundary gap, could vary. For bombardment simulations use of a NVT ensemble means the volume is fixed and makes selection of this gap more important. Bond length distributions for replicated structures at zero temperature (see Fig. 4.6) gave expected values indicating it is unlikely there is excessive stress in the system as a result of an incorrect grain boundary gap.

It has been suggested that removing random atoms from within the grain boundary can achieve lower system energies [188, 189]. Although, others have reported this made minimal difference [190], suggesting any benefit may be grain boundary specific. It was not deemed necessary to deploy this method as significant potential wells were already observed across the grain boundaries and structures were in good agreement with the TEM images they are reproducing.

²The acronym 'GBs' (grain boundaries) will typically be used in reference to the different structures (for example, 111 GB denotes the 111 Σ3 grain boundary structure), whilst 'grain boundary' will be used to describe the spatial regions ± 2.5 Å from the grain boundary centres. Similarly, 'perfect' diamond will be used to describe a structure without a grain boundary, and 'bulk' will be used to describe the spatial region that is not classified as a grain boundary.

4.2 Continuous Bombardment of Diamond with Deuterium

The work presented in this section has been adapted from computational work in [150]. The author (J. Pittard) carried out simulations, analysis and wrote the original draft. M.Y Lavrentiev also carried out simulations and provided comments on the original draft and throughout the review process, as well as providing supervision over work. N.A. Fox provided supervision as well as the original proposal of the work.

4.2.1 Method

Generally, simulations consisted of a (100) orientated block of diamond bombarded from below by deuterium atoms. Parameters used in these simulations were dictated by the previous tests. As such, an 8×8 UC surface was used, with a depth ranging from 6–18 UCs depending on deuterium energy. Again, periodic boundary conditions in x and y and a finite boundary condition in z were used. To prevent movement of the block during bombardment, atoms within the top UC of the block were fixed in position. A variable timestep was used with a default of 0.1 fs which, if required, was reduced to maintain a maximum atom movement of 0.1 Å. Simulations were performed in a NVT ensemble at a temperature of 300 K. This temperature was selected to allow comparison to experimental results for room temperature implantation presented in 3.2. A temperature damping parameter of 100 fs was used in line with testing and recommendations in the LAMMPS manual [191].

Deuterium atoms were created at random positions within a box 25 Å below the diamond and given positive z velocity corresponding to the desired implantation energy. Any etched carbon atoms within this box were deleted prior to creation of the next deuterium atom. The REBO potential [141, 142] was used to describe carbon-deuterium and carbon-carbon interactions. 4000 deuterium atoms were incident on the simulated diamond surfaces, with 0.5 ps between each bombardment corresponding to a fluence of 4.91×10^{20} D m⁻² and a flux of 2.46×10^{29} D m⁻² s⁻¹.

Simulations explored the 10–140 eV energy range. The lower end of this range represents typical energies expected to be experienced by PFMs in ITER [183], whilst the top end is closer to the hundreds of eV used in DELPHI. It proved challenging to simulate higher energies that directly overlap with experimental work. Higher energies would require much larger system sizes and increased relaxation time between bombardments to cope with the greater penetration and larger temperature fluctuations respectively. Additionally, the increased atom velocities would require a shorter timestep and, perhaps, use of a more computationally heavy potential such as AIREBO. As the average cluster size of incident species in DELPHI is 2.96 [120, 192], the energy range explored in 3.2 of 200–1000 eV becomes an average of 68–337 eV per atom. Simulations exploring the impact of different masses and higher energies are presented in 4.3.

Carbon and deuterium atoms were counted to calculate sputtering yields and deuterium retention. In order to do this, a surface must be determined to distinguish between atoms within




Figure 4.7: Output file from partway through a LAMMPS simulation of diamond (red atoms) bombarded with deuterium atoms (blue atoms) with atomic density with respect to the depth overlaid on top. At this point in the simulation, the diamond has been bombarded with 480 out of a total 4000 deuterium atoms. Atom densities were calculated by counting atoms within a half unit cell thick slice taken in z (green shaded region) every 0.1 \AA throughout the simulation system. Atoms in a unit cell thick slice of the non-bombarded surface were frozen in place (blue shaded region). Atom densities of carbon (black) and deuterium (blue) are presented. The surface was taken to be the solid red line and the dashed line was used as an ion cut-off, beyond which ions were not considered to be implanted. Rendered in OVITO [182].

the diamond and those outside. It was observed that the system would swell upon deuterium implantation, making surface determination challenging and the use of a fixed surface height inappropriate. Instead, the atomic density of carbon along the z axis was considered. Atomic density plots were made by taking a slice in z , 0.5 UCs in thickness (see Fig. 4.7), counting the atoms within this slice and then dividing by the slice volume. The slice began in the frozen region, then the z limits of the slice moved 0.1 \AA along the simulation box and the atoms were recounted. This process was repeated until the length of the simulation box had been considered, and an atom density value for each z value had been calculated. The surface height was taken to be the first z value that the atomic density values of the previous 2 \AA averaged below half the density of pristine diamond. Carbon outside of the surface line was considered to be etched, whilst monitoring of surface height as the simulation progressed gave an insight into swelling and etching of the diamond. To allow for surface bonding, a secondary line 2 \AA from the surface was considered as the cut-off for deuterium. Any deuterium beyond this cutoff would not be considered as implanted. Atom density against z for deuterium atoms was also determined to give depth profiles.

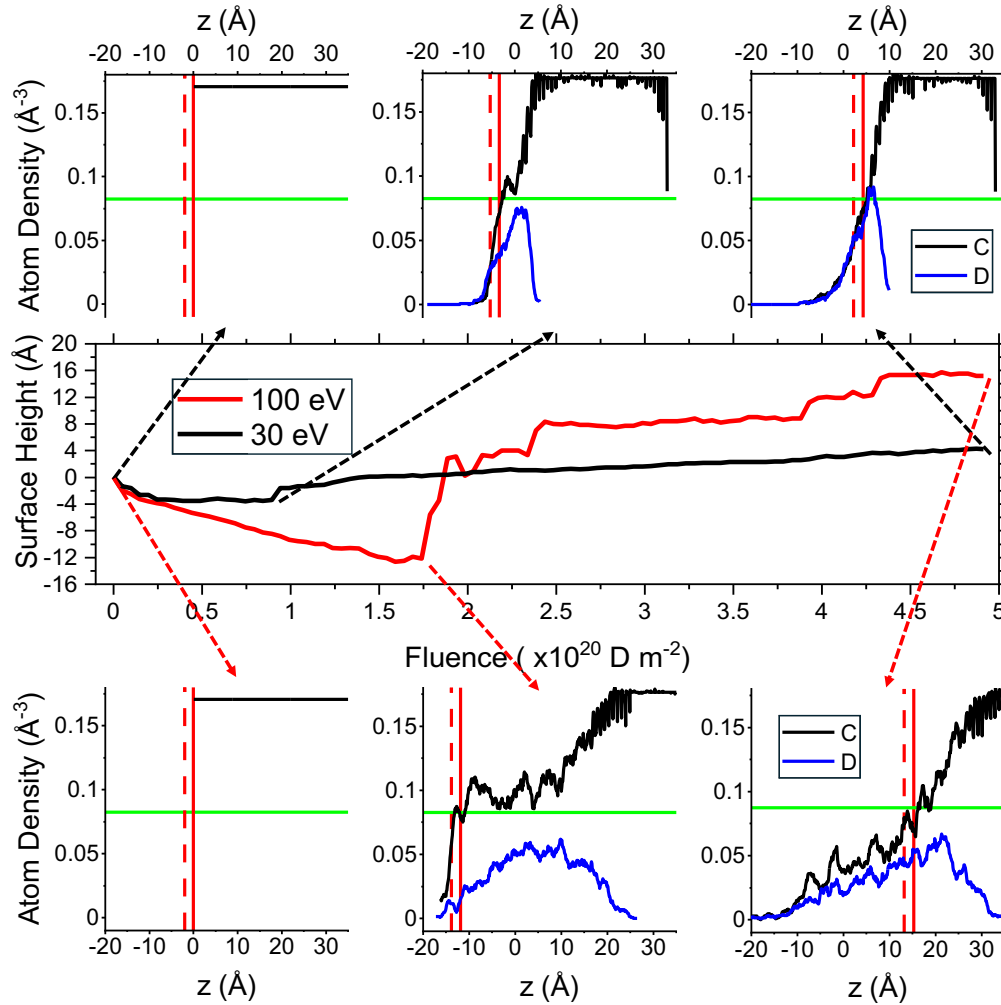


Figure 4.8: Results from LAMMPS simulations of 30 eV and 100 eV deuterium bombardments on diamond. Central figure shows the heights of the diamond surfaces throughout the simulations. The six surrounding plots show atom density of carbon and deuterium atoms against the z axis at three points in the simulations. The two central density plots show the etch point – where the maximum swelling of the surface has occurred, and etching is about to begin. The definition of the surface (solid red vertical line) was based on the point where carbon atom density dropped below 50 % (solid green line) of pristine diamond. Dashed red line shows the ion cut-off, beyond which, ions were not considered to be implanted.




Figure 4.9: Heights of diamond surfaces bombarded with deuterium atoms at set energies simulated in LAMMPS. The definition of surface height was based on the point where the atom density with respect to the z axis dropped below 50 % of pristine diamond. Negative values show swelling of the surface from the initial height (0 Å) whereas positive gradients show etching.

4.2.2 Results and Discussion

Results suggest a two-step etching mechanism, in which incident deuterium atoms are, initially, unable to sputter carbon atoms in the energy range studied here but are able to penetrate small distances (of order nm) into the surface. With increasing fluence, more deuterium atoms penetrate the diamond, resulting in a disordered region and swelling as the diamond loses its close packed structure. Eventually a point is reached where incident atoms are able to remove carbon from the disordered region; this is referred to as the etch point.

The swelling and etching phases can be seen in Fig. 4.8, which shows the surface height as the simulation progresses, alongside atom density plots from key points of the simulations. Negative surface height values indicate swelling of the surface before the etch point is reached. Density plots at this point (the two central plots) show an expanded sub surface region of lower density. In the etching phase, the higher energy implantation results in carbon being removed in clusters as seen in the steps in surface height for 100 eV implantation. The lower energy displays a much more controlled, steady and consistent etching phase due to the smaller disordered region, giving less opportunity for large clumps of carbon to be removed. A shorter swelling phase and slower etching is observed for lower energies (see Fig. 4.9).

Fig. 4.10 shows sputtering yields calculated using two methods. The final count method is calculated in the standard manner, where sputtering yield, S_{count} , is given by taking the ratio of




Figure 4.10: Sputtering yield calculated by two methods from LAMMPS simulations of diamond bombarded with deuterium atoms at a set energy. The volume etched method was calculated by taking a linear fit of surface height with respect to fluence once the etching phase had begun and assumed an atom occupies an average volume of an eighth of a unit cell. Errors were taken to be the error on the linear fit. The final count method took the ratio of implanted to incident deuterium atoms. Linear fit of volume removed method: Gradient = $0.0083 \pm 0.0005 \text{ eV}^{-1}$, y-intercept = -0.07 ± 0.01 . Linear fit of final count method: Gradient = 0.0078 ± 0.0003 , y-intercept = -0.02 ± 0.02

etched carbon, N_C over incident deuterium, N_D .

$$(4.2) \quad S_{count} = \frac{N_C}{N_D}.$$

Whereas the volume removed method used the volume etched per deuterium atom, V_{etch} , for various energies once the etching phase had been reached. V_{etch} values were found by taking linear fits of surface height (on equivalent plots to those seen in Fig. 4.9) after the etch point had been reached. Once the etch point is reached, a roughly linear etch rate would be expected from the regular atom bombardments. The average volume occupied by a single carbon atom was taken to be an eighth of a unit cell's volume (given by the cube of the lattice constant, α), allowing V_{etch} to be converted to a sputtering yield, S_{vol} , via

$$(4.3) \quad S_{vol} = \frac{8V_{etch}}{\alpha^3}.$$

Values calculated in this manner are approximately equivalent to values calculated via the final count method, but the volume removed method allows an average over the etching phase to be taken. For the 10 eV ion energy, the volume method could not be applied as a clear etching




Figure 4.11: Deuterium depth profiles from five LAMMPS simulations of perfect diamond bombarded with deuterium atoms at a set energy. The diamond had a 8×8 unit cell surface area and was bombarded with a fluence of 4.91×10^{20} D m $^{-2}$ deuterium atoms. The original diamond surface was at 0 Å.

phase was not present. S is approximately linear with energy for both methods, with the volume removed method giving a gradient of $(8.3 \pm 0.5) \times 10^{-3}$ eV $^{-1}$. This linear dependence is typical of the knock-on sputtering regime [193], where incident ions are of greater energy than the surface binding energy and are able to dislodge bound particles. As this point is only reached after the swelling, this would suggest a decrease in binding energy in the disordered carbon region as would be expected.

With increasing deuterium energy, more deuterium is retained, with the peak of the depth profile located deeper within the material as seen in Fig. 4.11. Generally, penetration depths were low, with the highest energy, 140 eV, showing a peak in the depth profile around 2.5 nm from the initial surface and max penetration of around 5 nm. This is in reasonable agreement with NRA results presented in 3.2, which found deuterium depths were beyond the resolution of the technique, meaning no deuterium was detected at depths greater than 10 nm. The broader deuterium profiles observed for higher energy implantation also resulted in more deuterium located at negative z values (outside the initial material). This is a consequence of the thicker disordered region resulting in greater swelling and trapped deuterium a further from the initial surface, which can be seen in the final density plots in Fig. 4.8. Total retention for experimental work presented in 3.2 were to the same order (10^{19} m $^{-2}$) as results from these simulations. Deslandes *et al.* [194], experimentally obtained a higher retention value for an ion energy comparable to the computational work presented here. This is likely to be a result of additional

factors, such as simulations using a perfect crystalline structure and perfectly flat surface (reducing surface area) which would be expected to reduce deuterium retention for a given energy.

One of the key differences between simulations and the experimental work in 3.2 is the single deuterons modelled here compared to the clusters (average size of 2.96 atoms [120]) incident in DELPHI. A larger mass would be expected to reduce penetration depth for the same energy, which can be shown by considering a head on collision. For the incident particle of mass m , directly colliding with a stationary particle of mass M , the ratio of final kinetic energy, ϵ_f , to initial kinetic energy, ϵ_i , for the incoming particle is given by

$$(4.4) \quad \frac{\epsilon_f}{\epsilon_i} = \left(\frac{m - M}{m + M} \right)^2.$$

Hence, going from the simulated mass of 2.01 u, to the average experimental mass of 5.96 u (2.96 deuterons), would give a drop in the energy carried through the collision from 51 % to 11 % and a greater degree of energy transferred to the stationary carbon atom. Treating these clusters as a single particle, the greater energy transfer compared to single atoms would be expected to reduce penetration depths but increase damage. It is possible the smaller interaction volume might behave similarly to the lower energy implantations, with reduced swelling and a more consistent etching rate. However, clusters are not single particles and it seems likely to break apart into a D^+ ion and a D_2 molecule on impact. How the cluster energy is distributed between deuterons in the cluster is unclear, making it challenging to draw definitive conclusions on the impact clusters would have.

Similar simulations were conducted by Dunn *et al.* [82], in which a variety of diamond surface orientations were bombarded with 15 eV tritium atoms. Comparing these results to results obtained for the 15 eV deuterium implantation presented here, shows very similar retention values (approximately $4.5 \times 10^{19} \text{ D m}^{-2}$ compared to $4.3 \times 10^{19} \text{ D m}^{-2}$ respectively). The etch rate of 0.073 obtained by Dunn *et al.* deviated slightly from S_{vol} and S_{count} values of 0.055 and 0.045 respectively. This could be a result of the larger mass of tritium compared to deuterium. As seen in Eq.4.4, for a fixed ϵ_i , a larger mass would result in greater energy transfer to the stationary particle, increasing the likelihood of sputtering. For an equal ϵ_f , the initial energy of a deuterium would need to be around 19.6 eV, or for an equal final momentum, an initial energy of 22.5 eV. This small increase in energy would result in a slightly higher sputtering yield as seen in Fig. 4.10. Other differences between simulations are also likely to play a role, such as the lower flux used in this study and subtle differences in thermostating. Lastly, Dunn *et al.* used a fixed cut off to define sputtered atoms unlike the variable surface heights used in this work. Although not discussed by Dunn *et al.* in depth, an initial period of reduced etching can also be seen in the results presented, similar to the two-step etching mechanism described here.

Hydrogen present in physical samples but not in simulation could act to reduce the uptake of deuterium. Although CVD diamond typically has a significant hydrogen content, this is concentrated within the grain boundaries [90], meaning the absence of simulated hydrogen in

bulk regions is unlikely to be drastically different to physical bulk regions. Simulations presented in 4.5 explore the impact of grain boundaries on retention.

4.2.3 Summary

LAMMPS simulations were carried out in which a fluence of deuterium atoms were incident upon a (100) diamond surface held at 300 K. Deuterium energy was varied between 10 – 140 eV and results were compared to experimental work presented in 3.2. Simulations revealed a two-step etching mechanism in which retained deuterium resulted in surface swelling and the formation of a disordered structure prior to carbon removal. Total retention for simulations and experimental work was to the same order, although the comparisons that can be made between experimental work and simulations are reduced due to the different energy regime and deuterium cluster size.

4.3 Repeated Single Bombardments of Diamond with Hydrogen Isotopes

In order to develop understanding of results in the previous section, this section explores how different variables might impact hydrogen³ retention and the damage sustained by the diamond. Results from continuous bombardments left questions regarding the influence of isotope mass and hydrogen energy in particular. To explore the impact of these and other variables in isolation, a series of repeated single bombardment simulations were carried out with LAMMPS. Each simulation consisted of a block of diamond hit with a single hydrogen atom which was then repeated a minimum of 300 times for each set of conditions. The mass, energy and incident angle of bombarding hydrogen, as well as diamond orientation and temperature were all varied. The impact of such variables was evaluated through measurements of vacant sites, interstitial carbon, sputtered carbon, reflected hydrogen and hydrogen positions.

The work presented in this section has been adapted from [179]. The author (J. Pittard) carried out simulations and analysis, as well as writing the original draft. M.Y Lavrentiev provided comments on the original draft and throughout the review process, as well as providing supervision and guidance throughout the work which he initially proposed. N.A. Fox provided supervision as well as the original proposal of the broader research.

4.3.1 Method

Simulations consisted of a diamond block of 5×5 UCs in area and a depth that was varied based on the expected implantation depth. Periodic boundary conditions in x and y dimensions and a finite boundary condition in z helped mimic a larger diamond block whilst allowing atoms to

³Throughout this section, ‘hydrogen’ will be used to refer to hydrogen isotopes collectively. Protium, deuterium and tritium will describe isotopes of simulated mass 1, 2 and 3 respectively.




Figure 4.12: Path of incident atom (blue) at an incident angle, θ , towards a $5 \times 5 \times 4$ unit cell (UC) diamond block (red atoms). ϕ is the angle of implantation about the z -axis. A target location (x_t, y_t) is generated at random within a 3×3 central region (black dashed line). Initial position (x_i, y_i) is calculated so the hydrogen atom hits the target position when created at a height of 3 UCs above the surface. Base diagram rendered in OVITO [182].

leave out the top of the simulation box. The total z length of the simulation box was typically 180 Å, with an empty region above the diamond of 80 Å. A 1 UC thick frozen layer was once again used on the non-implantation side to prevent shifting during implantation.

For each repeat, a hydrogen isotope was created 3 UCs above the surface in a position that ensured it was incident on the randomly generated target position by consideration of incident and planar angles (θ and ϕ respectively, see Fig. 4.12). For non-zero θ , ϕ is randomly generated unless specific directions were used (such as $\langle 110 \rangle$ or $\langle 111 \rangle$). Random seeds in the input file were also regenerated for each repeat.

(100) diamond blocks were constructed by replicating out an eight-atom unit cell. To create other orientations, trimmed datafiles were used as discussed in 4.1.6. The concerns surrounding trimmed datafiles were largely mitigated by ensuring bombarding atoms were only incident on a central 3×3 UC region on the surface. Nevertheless, a (100) trimmed structure was also produced for comparison. As part of the grain boundary bombardments presented in 4.5, replicated datafiles of different orientations were tested and gave very similar results to those presented here, suggesting the different approach had minimal impact on these results. However, it could be possible that additional strains are placed on the diamond by imperfect periodicities, which could have some impact on these results. Such strains are unlikely to be present for the (100) trimmed




Figure 4.13: Figure demonstrating the Wigner-Seitz vacancy analysis used to calculate vacancies and interstitials for a displaced configuration compared to a reference. Wigner-Seitz cells are constructed around each atomic point in a reference structure. The number of atoms from the displaced structure in each cell are counted (known as the occupancy). If no changes have occurred all cells will have an occupancy of one, cells with an occupancy of two or zero were classed as interstitials and vacancies respectively.

sample as the length of the unit cell is known so periodicity can be maintained.

Three different diamond temperatures (300, 600 and 1000 K) and four different diamond surfaces ((100), (100) trimmed, (110) and (111)) were tested. Hydrogen masses of one (protium), two (deuterium), three (tritium) and an unphysical mass ten were all tested. For every set of conditions, hydrogen energies of 10, 30, 60 and 100 eV were used, however additional energies were also explored for some sets of parameters with an overall range of 5 – 500 eV. Incident angle, θ , was varied between 0 – 85° with a random ϕ , as well as $\langle 110 \rangle$ and $\langle 111 \rangle$ directions.

The REBO potential was used to dictate carbon-carbon and carbon-hydrogen interactions [141, 142]. All bombarding atoms were treated as hydrogen regardless of mass. Atoms outside of the frozen layer were given initial velocities sampled from a Gaussian distribution corresponding to the target temperature. Target temperature was maintained throughout simulations by simulating an NVT ensemble. A short temperature damping parameter of 100 fs was used to allow for a rapid return to target temperature post implantation. A variable timestep ensured atoms moved no more than 0.1 Å per timestep, with a default value of 0.5 fs. An initial run of 3 ps allowed the target temperature to stabilise before the bombarding atom hit the surface. Once the incident hydrogen had been created, the simulation was run for 10 ps, giving sufficient time for the atom to hit the surface and come to rest in a stable position. Following this, the temperature was reduced to 0 K to remove thermal movement to improve vacancy analysis accuracy.

The analysis performed looked to quantify the impact of variables concerning both the hydrogen and the diamond. Hydrogen was classified as either implanted or reflected. If there was no hydrogen in the simulation box, or it was at a height greater than 2 Å from the surface, it was considered to be reflected. Hydrogen within 2 Å of the surface was taken to be implanted.

A Wigner-Seitz analysis (outlined in Fig. 4.13) was implemented via the OVITO Python package to determine vacancies and interstitials [182, 195]. The final output file (post bombardment, at a temperature of 0 K with hydrogen removed) was compared to the initial file which acted as a




Figure 4.14: Percentage of simulation repeats which resulted in the reflection of the incident atom. A single hydrogen atom was incident in each simulation, for a given set of conditions, a simulation was repeated 300 times. Incident angle, energy and mass of incident atoms were varied, as well as surface orientation and diamond temperature. Units of vertical momentum, P_z , have been selected for ease of conversion to energy, E , and mass, m , via $P_z = \sqrt{2Em} \cos(\theta)$ where θ is the incident angle. Data has been fitted with $y = y_0 + Ae^{-x/x_0}$, where $y_0 = 0.498 \pm 2.664$ %, $A = 107.9 \pm 3.9$ %, $x_0 = 9.4 \pm 0.8$ (eV u) $^{1/2}$.

reference. A Wigner-Seitz (or Voronoi) cell was constructed about each of the atomic sites (dotted lines in Fig. 4.13). Using atom coordinates from the final output file, the occupancy of carbon atoms for each cell was determined. For cells with an occupancy of one, it was deemed no point defect has occurred, whereas cells with occupancies of two or zero corresponded to interstitials and vacancies respectively. Typically, zero or one vacancy formed, as such, the percentage of simulations containing one or more vacancies was considered (referred to as ‘vacancy percentage’). Although sputtered atoms were counted (taken to be the difference between the number of interstitials and vacancies) values were very low and no meaningful conclusions could be drawn from this data.

4.3.2 Results and Discussion

4.3.2.1 Reflected Atoms

As can be seen in Fig. 4.14, percentage of reflected hydrogen was largely dictated by vertical angular momentum. Low vertical moment can be a consequence of lower isotope mass, lower isotope energy or an increased incident grazing angle. Increased vertical momentum increases the likelihood that an incident atom will continue into the diamond after collisions with carbon atoms and reduces the likelihood of it being reflected. Atoms can either be reflected directly from




Figure 4.15: Defect analysis of four datasets of 300 repeats of a (100) diamond block at 300 K being bombarded with single deuterium atoms of 30, 60, 80 and 100 eV normal to the surface. y -axes units are as a percentage of incident atoms (for example, 39 % of incident atoms caused a vacancy in the top atomic layer for 100 eV). Interstitials are the depth of any carbon atom with an occupancy of two, meaning each vacant site will result in two interstitials provided no carbon is lost. The percentage was halved to account for this. Dashed lines show the atomic layers of the diamond, with markers at every unit cell. Three depth categories ('Surface', 'Shallow', and 'Deep') were used for discussion and are marked on the deuterium density plot.




Figure 4.16: Fraction of hydrogen that caused one or more vacancies compared to all implanted hydrogen as a function of depth for all simulations of a (100) surface. Blue line shows the density of hydrogen depths, the red line also shows this but only for hydrogen that resulted in one or more vacancies forming. The black bars show the ratio between two densities for the hydrogen peaks.

the surface or travel some distance into the diamond and out again after a series of collisions. At the lower energies and vertical momentums explored in this study reflection is more prominent, and it would be expected that the former reflection mechanism would dominate.

As Fig. 4.14 shows data across all variables tested (different temperatures, incident angles, surfaces and isotope masses), it suggests diamond temperature and orientation have minimal impact on reflection. However, this is not to say these factors would not impact retention. Although hydrogen retention is likely to be impacted by reflection, it would be expected that penetration depth and vacancy formation would also be significant, particularly for an incident fluence of atoms. Higher penetration depths result in a larger interaction volume, and more hydrogen can be taken up before saturation of the top layer and the etch point is reached (see 4.2). Additionally, damage sustained to the diamond could create new binding sites that would not be accessible in pristine diamond. As discussed later, both diamond temperature and orientation saw changes in vacancy formation and minor changes in average depth, suggesting an impact in retention could still be expected even if changes to reflection are negligible.

4.3.2.2 Vacancy Formation

Fig. 4.15 shows depths of carbon interstitials, vacant sites and implanted deuterium for four datasets of different energies. Regardless of energy, it can be seen that vacant sites are restricted




Figure 4.17: Percentage of simulation repeats which resulted in the formation of one or more vacancies for different diamond temperatures (top) and incident hydrogen isotope masses (bottom). A single hydrogen isotope was incident in each simulation, for a given set of conditions, a simulation was repeated 300 times. All simulations performed on a (100) surface with a 0° incident angle. Deuterium (^2H) was used for temperature variation, whilst mass variation used a diamond temperature of 300 K.

to the top unit cell of the diamond block, despite the majority of higher energy deuterium residing beyond this region. Vacancies readily form in the top few atomic layers where carbon atoms can be pushed out of the surface - as seen from the carbon interstitial positions. The deuterium does not have sufficient energy to produce vacancies deeper within the diamond where interstitial positions for the carbon are less readily available.

Fig. 4.16 shows the correlation between hydrogen depth and vacancy formation. As with Fig. 4.15, peaks in hydrogen density correspond to regions occupied by hydrogen within the diamond. The bars on Fig. 4.16 show the fraction of hydrogen in these locations that resulted in one of more vacancies forming ('vacancy fraction'). It can be seen that hydrogen occupying positions

deeper within the diamond are less likely to have formed a vacancy. This is thought to be a result of vacancies only forming within the top atomic layers of the diamond as shown in Fig. 4.15. Hydrogen retained within the surface layers can push carbon out of the surface, as suggested by the depth of interstitials in Fig. 4.15, whereas hydrogen retained deeper within the diamond is unable to significantly displace carbon deeper than the surface. In some cases, positions occupied by hydrogen appear to only become available when a vacancy has formed, this seems likely for the peak at a depth of 0.2 Å which gives a vacancy fraction of nearly 1 meaning practically no hydrogen occupied this depth without forming a vacancy. These new sites are only forming, and therefore occupied, near the surface, resulting in an increase in vacancy fraction. Lower vacancy fractions (~ 0.6) are also observed at depths at which vacancies can form, suggesting the lateral position of the hydrogen is impacting vacancy formation alongside depth. However, as some hydrogen is able to form a vacancy in the surface whilst coming to rest at a point beyond vacancy forming depths, vacancy formation cannot be solely dictated by exact position and the amount of energy transferred to surface atom in initial collisions must be of some importance.

At low energies, small increases in energy result in a significant increase in vacancy formation as seen in Fig. 4.17. After peak vacancy formation is reached, a more gradual decrease is observed with increasing energy. Although the energy range of the initial increase in vacancy formation corresponds to the drop off in reflected atoms seen in Fig. 4.14, it is not a result of decreased reflection. Reflected hydrogen contributes a significant amount to total vacancies formed, with 52 % of reflected hydrogen resulting in a vacancy (across all simulations). This percentage was typically highest with higher energies, as reflection will more commonly occur from a head on collision which is more likely to create a vacancy due to greater energy transfer. Instead, this increase is simply a result of some minimum energy required to create a vacancy, with the energy of vacancy formation reported at around 9 – 15 eV [196]. As discussed regarding Fig. 4.15, the gradual decrease is thought to be a result of the reduction in hydrogen retained within the first few atomic layers.

In Fig. 4.17, it can be seen that the energy at which peak vacancy formation occurs also increases with mass and shows a longer tail afterwards. Reduced energy transfer between collisions (see Eq.4.4 in 4.2) results in lighter masses travelling deeper into the diamond. This can be seen in Fig. 4.18, where the linear fit in the log-log plot is suggestive of a power law dependence between energy and depth. Deeper penetration depths correlate with reduced vacancy formation (see Fig. 4.16), so the drop off in vacancies is faster for lower masses. These observations are in contrast to what is observed for temperature variation, where no change in the energy for peak vacancy formation was observed. However, an increase in temperature did result in an increase in vacancies. The higher average energy of carbon atoms meant they were more susceptible to displacement and vacancy formation.

A significant difference between the three surface orientations was observed in Fig. 4.19, with both (110) and (111) surfaces showing a notable reduction in vacancies formed compared to all




Figure 4.18: Average depth for a series of repeated, single bombardment simulations of hydrogen isotopes of varying mass and energy at an incident angle of 0° on a (100) diamond surface at 300 K. Linear fits have been applied to the logarithmic plot giving gradients of ^1H : 2.08 ± 0.09 , ^2H : 2.15 ± 0.05 , ^3H : 1.95 ± 0.10 , ^{10}H : 1.55 ± 0.03 .

(100) surfaces including (100) trimmed. It is thought there are two reasons why these surfaces displayed lower vacancy formation. (111) has a significantly lower surface energy than (100) [198] and a higher displacement energy in this direction [199] giving a reduction in vacancies compared to (100). However, (110) has a surface energy very similar to (100) and a lower displacement energy and yet still displayed lower vacancy formation. Instead, the reduction is thought to be a result of atom alignment. In $\langle 110 \rangle$, clear channels can be observed which reduce the chance of collision with carbon atoms in the top atomic layers and, therefore, vacancies. The same reduction was not observed for (100) in the $\langle 110 \rangle$ direction as the 45° incident angle meant more of the collisions that occurred would have been in the top atomic layers.

4.3.2.3 Hydrogen Positions

Changes to the temperature and orientation of the diamond resulted in minimal changes in average depth, although small increases were observed for higher temperatures and the (110)




Figure 4.19: Results from repeated simulations of deuterium bombarding various diamond surfaces at 300 K. Where only a surface is specified in legend, bombarding atoms were at an incident angle of 0° . Directions indicated correspond to implantation angles upon a (100) surface. (110), (111), and (100) trimmed were all constructed by trimming a orientated diamond block, whereas (100) surfaces were made by replicating a diamond unit cell. The depth plotted takes into account the incident angle, θ , effectively plotting the distance travelled within the diamond rather than depth, z_{final} .




Figure 4.20: Heat map of xy hydrogen positions with respect to position in the unit cell (UC) at three different depth regions following repeated simulations of hydrogen bombardment. Values on colour scale are arbitrary but indicate the density of hydrogen. Data was taken across all simulations of $\langle 100 \rangle$ surface orientations. Hydrogen with a final depth, z , of $-2 < z < 0$ Å was classified as surface hydrogen (**A**), $0 < z < 1.784$ Å as shallow (**B**) and $z > 1.784$ Å as deep (**C**). Red markers represent carbon atom positions. Where appropriate, different markers have been used to indicate the different atomic layers (AL) of the diamond ('X' - first AL, '+' - second AL, 'O' - third AL, 'Δ' - forth AL).

surface. The higher temperature encourages movement of the carbon atoms, resulting in slightly increased penetration depths. As previously discussed, channels can be seen in the $\langle 110 \rangle$ direction which could increase penetration depths. In fact, hydrogen incident in $\langle 110 \rangle$ on a $\langle 100 \rangle$ surface resulted in a significant increase in average penetration depth when incident angle was accounted for. This is likely to be impacted by the real depth being shallower and shallow hydrogen occupying different interstitial positions (see Fig. 4.15). However, it is still a notable increase on the $\langle 111 \rangle$ direction on a $\langle 100 \rangle$ surface which had the same incident angle but different φ – suggesting these channels are significant.

Parameters surrounding the incident hydrogen had a much more significant effect than parameters relating to the diamond. Increasing energy of hydrogen resulted in a non-linear increase in penetration depth. This relation is dominated by 'sticking' to the surface at low energies (< 30 eV) and then transitions to an approximately quadratic relationship for isotope masses much less than the mass of carbon, approaching $z \propto E^{3/2}$ for mass ten (see Fig. 4.18). This superlinear behaviour has been observed in simulations of other materials [164] and was attributed to channelling effects which require modelling atomic structure explicitly, as is done in MD. Other techniques which effectively treat structures as amorphous and uniform, such as BCA codes, would typically predict a more linear relation. Higher mass and energy atoms, which are more effective at moving carbon atoms, are less influenced by these channels and display a relationship more similar to the expected output from BCA codes. It is possible the formation of the amorphous disordered layer might also result in better agreement between MD and BCA.

Hydrogen that has sufficient energy to penetrate the first few atomic layers will travel deeper into the crystal through atomic channels within the crystal structure. For continuous

bombardments like those presented in 4.2, the greater than linear relation between depth and energy is expected to be exaggerated as lower energy hydrogen creates more vacancies and trapping sites near the surface - further reducing penetration at low energies. The formation of a disordered region in seen in continuous bombardments would negate the impact of atomic channelling and generally reduce penetration depths. For higher energy implantation, it takes longer for this region to form allowing channelling to occur at higher fluence, which would exaggerate the penetration depth compared to lower energies further. Channelling along grain boundaries and the impact of disorder is explored in 4.5.

In contrast to the depth profiles presented here, a fluence of incident deuterium at equivalent energies resulted in a much wider range of occupied depths (see Fig. 4.11 in 4.2). For example, 60 eV implantation of single deuterium atoms gave no significant amounts beyond depths of 2 UCs (7.134 Å, see Fig. 4.15), whereas for continuous bombardment of 60 eV deuterium gave depths up to 30 Å. This observation is likely to be a result of the formation of the disordered region during continuous bombardments that is not present for single implantations. The disordered region allows etching to occur, reducing the surface height of the diamond for subsequent deuterium resulting in deuterium at higher depths. This region is also saturated with deuterium, which can be pushed further into the diamond from collisions with incident atoms. The reduction in density in this layer results in carbon and deuterium atoms spilling out beyond the original surface. This is similar to what is observed in the interstitial positions in Fig. 4.15 as a result of hydrogen retained within the top atomic layers pushing carbon out. However, significant retention above the surface, other than the surface positions occupied for low energies, is only present in continuous bombardment simulations.

A more detailed look into interstitial positions occupied by the hydrogen was also carried out. Hydrogen depths seen in Fig. 4.15 were considered in three regions: surface hydrogen within 2 Å of the surface; shallow hydrogen at a depth between 0 and 1.784 Å; and deep hydrogen at depths greater than this. xy positions of hydrogen with respect to the unit cell were plotted for each depth selection as seen in Fig. 4.20. Surface hydrogen is positioned directly above atoms on the top atomic layer as seen in Fig. 4.20A. The shallow hydrogen in Fig. 4.20B mostly occupies positions in midpoints between atomic sites in the third atomic layer and the tetrahedral positions on the first/surface atomic layer (a C-site [90]). This occupation is only observed for shallow hydrogen, at a depth where there are no carbon atoms directly above this position. For deep hydrogen, there are carbon atoms above this position which seems to make its occupation energetically unfavourable as different positions are occupied. Other positions, more similar to those observed for deep hydrogen are also present at shallow depths but to a lesser extent. The dominant peak at around 0.892 Å (see Fig. 4.15) corresponds to these C-sites, the occupation of C-sites resulted in less carbon displacement than the other shallow positions, which gave much higher vacancy fractions. Deeper into the crystal, two peaks either side each atomic layer are present. Looking at both Fig. 4.15 and Fig. 4.20C, it can be seen that these positions are along the diagonal between

nearest neighbour carbon atoms on two atomic layers but not at the midpoint. This is not a recognised interstitial position in diamond and might be an artefact of the REBO potential used in these simulations. These simulations were repeated with the AIREBO potential, but both exhibited the same positions.

4.3.3 Summary

Repeated, single bombardment simulations gave insight into the impact of a wide variety of variables. At the lower energies expected within a fusion reactor's first wall (around 15 eV within the divertor of ITER [82]), results would suggest significant reflection, with approximately 55 % of incident deuterium and 48 % of tritium being reflected at a 0° incident angle. These percentages would be expected to increase with higher incident angles. As reflection is largely dictated by atomic structure in the top few atomic layers, the presence of a disordered region as observed in other simulations would impact reflection significantly for moderate fluences. As saturation of the disordered region is observed in continuous bombardment simulations, a much higher reflection coefficient must be present (or at least significant degrees of substitution and etching).

Vacancy formation was seen to peak around 20 eV which could be a concern surrounding etching for continuous bombardment. Using different surface orientations had a notable impact on reducing vacancy formation, so could be considered a way to reduce damage sustained by the diamond. Vacancy formation was very localised to the surface level, regardless of hydrogen energy for the energy range tested here. This, the low penetration depths, and the limited diffusion of hydrogen in diamond suggests minimal changes to bulk properties would be expected as a result of hydrogen interaction. A deeper understanding of diffusion characteristics is required to conclude this and is explored in the next section (4.4).

Small changes in penetration depths were observed when changing diamond orientation and temperature. An increase in penetration depth would result in a larger interaction volume and therefore increased retention. Although the (110) surface potentially saw a small increase in penetration depth, its resistance to vacancy formation/structural changes is likely to outweigh small increases in retention. The (111) surface showed a similar reduction; however, this surface is the most susceptible to graphitisation [200] which could make it unsuitable regardless.

Interstitial positions of incident hydrogen were also explored. C-site occupation was observed for shallow hydrogen, but at higher depths only undefined positions were observed. Some positions were only available when a vacancy had formed, whereas others caused less disruption and could be occupied with no damage.

Results from mass variation did confirm the expected trend discussed in 4.2, with higher masses giving lower penetration depths. The relationship between average depth and energy was non-linear and also varied between isotope mass. Average protium depth varied with $E^{2.08}$ compared to $E^{1.55}$ for the mass ten isotope. The non-linear relation is thought to be a result of channelling [164]. As expected, higher mass atoms also gave an increase in vacancy formation.

Despite these confirmations, the impact of different deuterium cluster sizes remains unclear. Although the behaviour of a single particle of increased mass has been established, clusters will likely break apart on impact and it is unclear how the energy will be redistributed when this happens. Specific modelling exploring this is required to fully understand the effect of clusters, and how this may impact comparing experimental test facilities with low temperature plasmas to the high temperature plasma present in tokamaks.

Depth profiles were typically narrower than those presented in 4.2 for continuous bombardment. It is thought that hydrogen implanted within the top layers of the diamond can be pushed in further upon additional incident atoms. Etching is likely to have played a role here as well, reducing the surface height and the formation of an amorphous layer with deuterium distributed throughout. This disordered layer will remove any channelling effects present due to a lack of repeating atomic structure.

The sputtering yield for multiple bombardments gave a linear relation with energy in 4.2, whilst a clear peak was seen at lower energies in vacancy formation in this section. This may go towards explaining why higher energies required higher fluences to reach the etch point in continuous bombardments - as the etching is reliant on the formation of the amorphous layer, and higher energies appear less effective at creating vacancies. In this scenario, hydrogen must steadily build up within the top layer disrupting the structure rather than relying on individual atoms to create vacancies and interstitials.

4.4 Hydrogen Diffusion within Diamond Grain boundaries

Up to this point, computational work has only focused on perfect diamond as, even in PC diamond, there are proportionally many more carbon atoms within bulk diamond than grain boundaries. However, the vast majority of hydrogen⁴ impurities can be found within the grain boundaries of CVD diamond [90], meaning they may play a disproportionate role in hydrogen isotope retention. Furthermore, PC diamond may prove more appropriate for some applications due to the additional cost associated with SC diamond, making the understanding of both grain boundaries and bulk regions key.

When considering retention, diffusion is a key factor in determining at what point the material will saturate as well as addressing concerns surrounding transport of restricted substances such as tritium. Previous conclusions have stated minimal impact on bulk regions beyond the surface but only considered a very short timescale (nanoseconds), diffusion characteristics must be determined to understand how retention may develop on longer timescales. Typically, crystalline materials demonstrate accelerated diffusion within grain boundaries. As hydrogen collects within the grain boundaries, it is important to assess whether this is the case for diamond as well. Continuous desorption of H₂ in TDS results (see 3.2) suggested evidence of diffusion of hydrogen

⁴Once again, hydrogen will be used in reference only to protium, rather than hydrogen isotopes collectively.

through the grain boundaries - where H_2 is most prevalent. A deeper level of understanding is required to assess whether diffusion is likely at those temperatures from either the grain boundaries or the bulk.

CVD growth conditions will dictate the grain size of the resulting diamond. Longer growths typically result in thicker films and larger grain sizes, decreasing the proportion of carbon atoms in grain boundaries and also changing the type of grain boundaries present. For example, in nanocrystalline diamond (NCD), grain boundary carbon atoms can make up to 10 % of the total atoms [187] with both tilt and twist boundaries present [201], of which the (100) $\Sigma 5$ GB [188, 202] is particularly prevalent. The increased number of grain boundaries in NCD results in an increased sp^2 content, and etching of sp^2 remains a key concern surrounding the selection of diamond for many potential fusion applications meaning larger grains are thought to be preferable. The columnar growth of CVD diamond results in thicker, larger grained films having a greater proportion of tilt grain boundaries [203], making these the focus of this work. Sawada *et al.* [178] identified atomic structures of four common tilt grain boundaries (111 $\Sigma 3$, 112 $\Sigma 3$, 114 $\Sigma 9$ and 221 $\Sigma 9$) from TEM images of CVD diamond - meaning these real world structures could be recreated in simulation. Other work [203–208] has also reported the presence of these grain boundaries with the 111 $\Sigma 3$ appearing most prevalent. Higher order $\Sigma 27$ and $\Sigma 81$ boundaries have also been observed but these form from the meeting of two lower order boundaries and are therefore less common. The same grain boundaries were also seen in boron doped diamond [209].

The work presented in this section has been adapted from [180]. The author (J. Pittard) carried out simulations, prepared input files and performed analysis (specifically surrounding calculations of coordination numbers, hydrogen energetics, bond length/angle changes and interpretation of the results) and wrote the original draft. M.Y Lavrentiev proposed the work, carried out simulations to determine mean-square displacement values and ultimately diffusion coefficients, provided comments on the original draft and throughout the review process, as well as providing guidance and supervision regarding the scope of the work. N.A. Fox provided supervision as well as the original proposal of the broader work.

4.4.1 Method

MD simulations were performed in LAMMPS [80, 81] to study the energetics and atomic structure of diamond grain boundaries, as well as the diffusion of hydrogen within. For all simulations and energy minimisations, the AIREBO [141, 142, 145] potential was used to dictate carbon-carbon and carbon-hydrogen interactions. For systems containing a grain boundary, it was orientated perpendicular to x (parallel with the yz plane). Periodic boundary conditions were used in all dimensions - simulating larger grain boundary planes in y and z dimensions, and a second grain boundary at the edge of the simulation cell in x . A variable timestep was used that ensured a single atom could not change kinetic energy by more than 100 eV or move more than 0.1 Å in a single timestep. The default timestep was 1 fs. The temperature was controlled by simulating

an isothermal-isobaric (NpT) ensemble with a temperature damping parameter of 1 ps and no external pressure applied.

Tilt grain boundaries from TEM images taken from [178] were used to recreate replicated data files as outlined in 4.1.6. Total system size was approximately $57 \times 57 \times 57 \text{ \AA}$ (30000 atoms) but did vary between the different structures due to the discrete size of the UCs used to produce them.

A series of simulations were carried out to map the potential energy of hydrogen within the diamond. To do this, hydrogen atoms were created in random positions within the simulation box prior to performing an energy minimisation to allow atoms to reach stable positions. The number of atoms added was varied to correspond to 1 at.%, this was to reduce hydrogen-hydrogen interactions and distortion of the structure whilst minimising the number of simulations required to collect sufficient data. This process was repeated 500–800 times, to give a total of around 138000–264800 hydrogen atoms in total depending on the grain boundary. To give the hydrogen positions and potential energy along a certain direction, hydrogen atoms within a 0.5 \AA slice were counted and the potential energy of these atoms averaged. The bounds of this slice moved along x in 0.05 \AA increments until the entire cell had been covered, in a similar manner to depth profiles in previous sections. Heat maps of the potential energy in two dimensions were also produced.

In order to measure diffusion, 25 hydrogen atoms were created in both grain boundaries (50 in total) in each simulation cell. Hydrogen atoms were added to the grain boundary randomly within boxes of $2 \times 2 \text{ \AA}$ in size. Boxes were separated from one another by 10 \AA (which gives the minimum hydrogen separation). A similar process was used for evaluating bulk diffusion coefficients within a block of perfect diamond. An energy minimisation was then performed, before a run of approximately 4 ns at an elevated temperature (ranging between 1500 – 3125 K). Diffusion coefficients were estimated from hydrogen mean-square displacement (MSD) values. Plotting MSD as a function of time, t , gives a gradient of six times the diffusion coefficient, D , as

$$(4.5) \quad \text{MSD} = 6Dt.$$

D varies as a function of temperature, T , according to the Arrhenius law

$$(4.6) \quad D(T) = D_0 e^{-E_A/RT}.$$

Therefore, plotting the natural log of D as a function of reciprocal temperature should give a linear gradient from which the activation energy, E_A , can be estimated.

Various methods were used to explore carbon hybridisations present within the structures. Hybridisation refers to the different types of covalent bonding that can occur through the mixing of atomic orbitals. Depending on the local environment, carbon-carbon bonds can form as sp (linear, 180° bond angle), sp^2 (trigonal planar, 120° bond angle) or sp^3 (tetrahedral, 109.5° bond angle). Different hybridisations result in different properties, graphitic carbon is dominated by sp^2 C-C bonds, resulting in planar sheets (graphene) of carbon atoms held together by delocalised

electrons. In contrast, diamond is characterised by sp^3 C-C bonds, resulting in a very hard and resilient material. Coordination numbers (CN) are commonly used [210] to distinguish between different carbon hybridisations and were calculated using a 1.73 Å cutoff distance. Atoms within this distance were classed as nearest neighbours and were used to calculate bond lengths and angles. Although bonds were not explicitly modelled, bond lengths were taken to be the distance between two nearest neighbours, and bond angles the angle between the displacement vectors of two pairs of nearest neighbours.

Changes in hybridisation on the addition of hydrogen to the grain boundaries were also explored. To do this, 100 hydrogen atoms were added randomly between both grain boundaries (up to 5 at.% within the grain boundaries) in each simulation box and the system was minimised. For the 114 GB, only 50 hydrogen atoms were added in an attempt to minimise disorder. Depending on the grain boundary, these minimisations were repeated 10 – 50 times, with hydrogen in different positions for each repeat. As coordination number will typically increase on the inclusion of additional atoms, changes in bond length and angle were also evaluated. For this, the carbon atom closest to each hydrogen atom was considered. The carbon-carbon bond lengths and angles for this atom was compared to the structure without the addition of hydrogen.

4.4.2 Results and Discussion

Fig. 4.21 shows the atomic structures of the grain boundaries, as well as the potential energy and final position of hydrogen placed within the simulation box in random positions and minimised. There are clear potential wells across the grain boundaries for the 112, 114, and 221 GBs, resulting in hydrogen trapping within the grain boundary. This is in good agreement with physical expectations, as hydrogen within CVD diamond is typically concentrated to the grain boundaries [90]. The exception to this is the 111 GB, which shows no significant increase in hydrogen across the grain boundary, but more variation across the bulk. Both of these observations are a result of atomic alignment. (111) is a stable surface in diamond with atomic planes parallel to the grain boundary plane, resulting in a coherent grain boundary without any significant deviation from the bulk structure. As the atomic planes are aligned with x across the grain boundary, collections of hydrogen in between atomic planes are observed at set x values, unlike in the other grain boundaries where atomic layers do not align with x and the hydrogen trapped between layers is averaged across x . This can be seen in Fig. 4.21A where the peaks in hydrogen correspond to the gaps between atomic layers. The average potential energy of hydrogen within the bulk regions was very consistent between the different orientations - all giving values of -1.08 eV to two decimal places with an average across the four structures of -1.0819 ± 0.0008 eV. Therefore, the hydrogen potential energy within bulk regions is largely independent of orientation as should be the case.

A consequence of this hydrogen trapping is the observation of anisotropic diffusion within the grain boundary, where hydrogen within grain boundaries preferentially diffuses within the grain

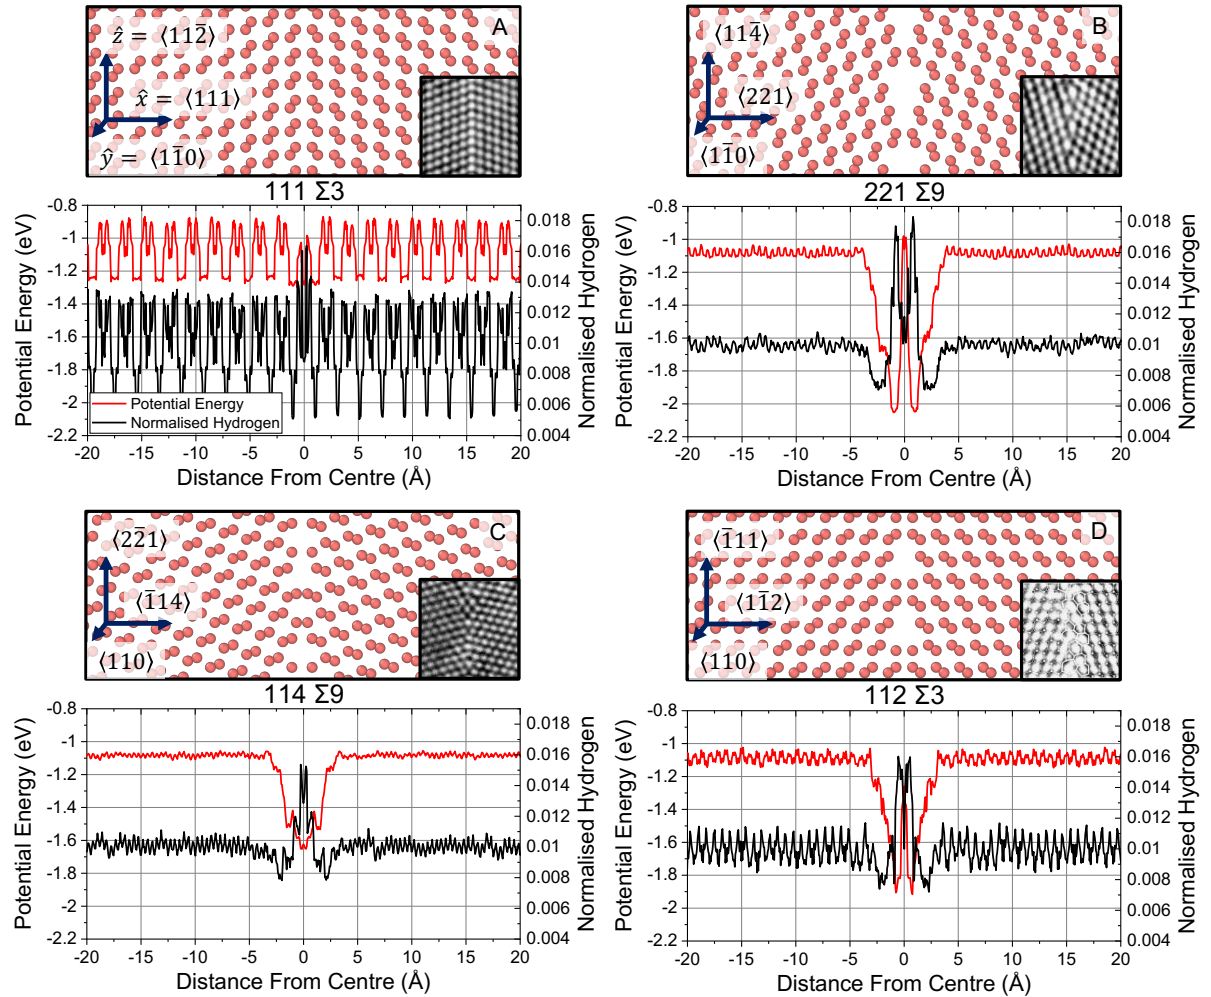


Figure 4.21: Atomic structures of four tilt grain boundaries – 111 $\Sigma 3$, 221 $\Sigma 9$, 114 $\Sigma 9$ and 112 $\Sigma 3$ (A, B, C and D respectively). All structures pictured contain no hydrogen atoms and are orientated with $\langle 110 \rangle$ into the page (y -axis). In the plots below, the red line shows the potential energy along the x -axis of hydrogen placed randomly within the structure and minimised. The black line shows the final positions of these hydrogen atoms, this has been normalised by the total hydrogen added across the whole structure (not just ± 20 Å from the grain boundary as shown here). The same scale has been used on all axes to aid comparison. Transmission electron microscopy (TEM) images of the replicated structure are also included (taken from [178]). Rendered in OVITO [182].




Figure 4.22: Raw mean square displacement (MSD) values for hydrogen diffusing in different diamond structures at a temperature of 2750 K. Miller indices in the legend indicate orientation of faces meeting in the grain boundary (GB) along x . The perfect diamond contained no grain boundary. All structures were orientated with $\langle 110 \rangle$ direction aligned with y . Diffusion coefficients were taken as the gradient of a linear fit on total MSD divided by 6 (Eq.4.5).

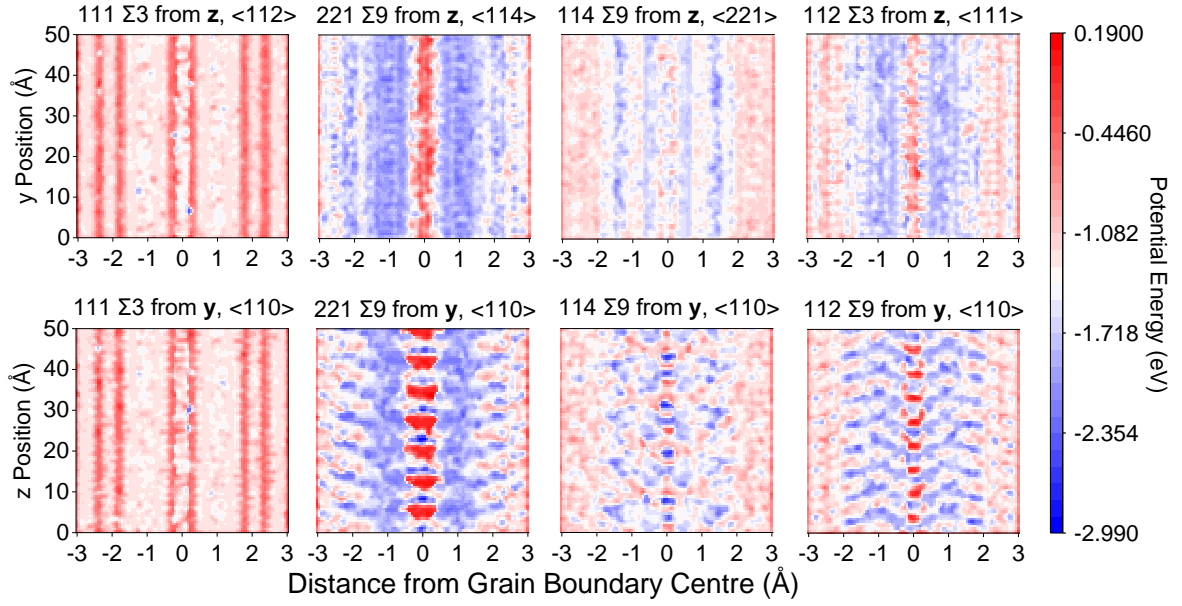


Figure 4.23: Heat maps of the potential energy of hydrogen placed in random positions and minimised for four diamond tilt grain boundaries. The top row of plots shows the grain boundary as if looking along the z direction, whereas the bottom along y . The same scale has been used on all the heat maps to ensure they are comparable to one another. The z direction corresponds to $\langle 112 \rangle$, $\langle 114 \rangle$, $\langle 221 \rangle$ and $\langle 111 \rangle$ for the 111, 221, 114 and 112 GBs respectively, whereas all y directions were aligned with $\langle 110 \rangle$.

boundary plane as opposed to drifting into the bulk. This behaviour is similar to observations of hydrogen diffusion in grain boundaries of different materials [211]. Anisotropic diffusion can be seen in Fig. 4.22, which presents the mean square displacement (MSD) in each direction. For perfect diamond and the 111 GB, x , y and z components are approximately equal, suggesting no preferential direction of travel and isotropic diffusion. The same cannot be said for the three other grain boundaries, which present negligible diffusion in the x direction (across the grain boundary). In the 221 GB, the y component dominates. This direction aligns with $\langle 110 \rangle$, where clear channels can be seen in the grain boundary (see Fig. 4.21), offering preferential diffusion in comparison to the other directions. The isotropic diffusion observed for perfect diamond suggests no preferential diffusion in $\langle 110 \rangle$, meaning the larger channels present in the grain boundary are required. Although the 114 GB does show some displacement in the z direction (also within the grain boundary plane) this does not display a steady increase, suggesting atoms may have moved in this direction into channels which they then diffuse along in y . The 112 GB shows a gradual increase in displacement for both y and z directions, suggesting hydrogen is trapped within the grain boundary but able to move in both directions within.

Channels were explored further by plotting Fig. 4.23, which maps out the potential energy in




Figure 4.24: sp^3 and sp^2 content across four tilt grain boundaries in diamond (111 Σ 3, 221 Σ 9, 114 Σ 9, 112 Σ 3) post energy minimisation. Coordination numbers (CN) of 3 and 4 were classed as sp^2 and sp^3 respectively and taken as a percentage of all atoms in that position. Gaps in data are a result of atomic layers leaving no carbon atoms in these positions. No sp^1 (CN 2) was found in any of the grain boundaries.

two dimensions across the grain boundaries using the same data as presented in Fig. 4.21. It would be expected that regions of high contrast (indicating a significant difference in potential energy) would result in restricted diffusion. At a distance of ± 1 Å from the centre of the 221 GB, there are seemingly unobstructed low energy regions along the z direction. As very little diffusion was observed in this direction, it would suggest the fast diffusion in y is instead a result of areas of low energy along the centre of grain boundary. These are separated by high potential regions preventing diffusion in z and offer very effective channels for hydrogen transport. Similarly, the 114 GB also has its lowest energy points along the centre of the grain boundary. However, as other channels are present throughout this grain boundary, it is hard to determine whether these central channels are dominating. The regions of higher energy separating regions of lower energy along z result in more dominant diffusion along y . The contrast between these regions is less distinct than those in the 221 GB and some movement in y was initially observed of hydrogen moving into these channels. The 112 GB gave almost equal diffusion within the grain boundary plane. Although channels are visible in the y direction, the regions of high energy that separate them along z are less distinct than the 221 and 114 GBs, allowing some diffusion to occur in this direction. The 111 GB presents with a much more homogeneous potential energy with the only variation corresponding to atomic layers, resulting in the observed isotropic diffusion.

Hydrogen trapping within the grain boundary is thought to be a result of the increased sp^2 nature compared to the bulk. This can be seen in Fig. 4.24 which presents coordination number as a percentage across the grain boundaries. Within the 112, 114 and 221 GBs, a clear increase




Figure 4.25: Diffusion coefficients as a function of inverse temperature for various diamond structures. Miller indices in the legend give the direction of the tilt grain boundary modelled, data from these simulations have been plotted alongside values for perfect diamond ('Perf.') which contained no grain boundary (simulating bulk diamond) and experimental values ('Exp.') of PC diamond taken from [91]. The log of the diffusion coefficients have been fitted with a linear fit to give activation energies, E_A , in Table 4.1.

in sp^2 and decrease in sp^3 can be observed. The same cannot be said for the coherent 111 GB, in which only sp^3 is present throughout which is in good agreement with physical expectations [205, 212]. Hydrogen position peaks in Fig. 4.21 broadly align with peaks in sp^2 seen in Fig. 4.24. The double peak in hydrogen position in the 112 GB aligns with the double peak in sp^2 carbon within the grain boundary, the single peak in the 114 GB is comparable to its sp^2 content and the increase in sp^3 at the centre of the 221 GB gives an increase in potential energy and reduction in hydrogen occupation. This could simply be a result of the lower atom density offering more positions for the hydrogen to occupy and giving lower coordination numbers. However, in contrast to what is typically observed in crystalline solids, most grain boundaries presented higher activation energies than the bulk (or perfect diamond) as can be seen in Fig. 4.25 and Table 4.1. This behaviour would not be expected if diffusion was solely dictated by atomic density. Lower atomic density in grain boundaries would be expected to enhance diffusion compared to the bulk as seen in other crystalline solids. The difference is the type of bonding present in the covalently bonded carbon. On the inclusion of hydrogen into the grain boundary, hydrogen can form sp^3C-H bonds with the sp^2 carbon present [213–215], resulting in restricted diffusion compared to the

Table 4.1: Activation energies and diffusion coefficients as determined in Fig. 4.25. Experimental values taken from [91].

	D_0 (m^2s^{-1})	E_A (eV)
Experimental	2.6×10^{-4}	2.8 ± 0.2
Perfect (Bulk)	4.3×10^{-9}	0.70 ± 0.06
111 $\Sigma 3$	3.7×10^{-9}	0.72 ± 0.03
221 $\Sigma 9$	3.0×10^{-8}	1.1 ± 0.1
114 $\Sigma 9$	3.0×10^{-8}	1.9 ± 0.1
112 $\Sigma 3$	2.5×10^{-7}	2.1 ± 0.2

bulk where sp^3 carbon dominates and hydrogen occupies interstitial positions [90]. Formation of these more stable bonds will also contribute to the hydrogen trapping. For other diffusing species which cannot form covalent bonds with carbon, diamond grain boundaries have demonstrated preferential diffusion compared to the bulk [216], indicating the type of bonding is playing a role.

Experimental results for PC CVD diamond (taken from [91]) align best with 112 and 114 GBs both in terms of activation energy and diffusion coefficients at lower temperatures. As the experimental values are a combination of bulk and grain boundary contributions, these results suggest the higher concentration of hydrogen within grain boundaries and the restricted diffusion of this hydrogen is dominating, despite bulk regions making up a far greater proportion of PC diamond. As observed throughout the results presented here, the lack of a distinct grain boundary for the 111 GB resulted in very similar behaviour to perfect diamond. Although the 221 GB gave similar diffusion coefficients to perfect diamond and the 111 GB, the activation energy was still notably higher. As discussed, the diffusion within this grain boundary is thought to be dominated by hydrogen moving through low energy channels along the centre of the grain boundary (as shown in Fig. 4.23). This channelling effect, as well as the presence of sp^3 carbon at the centre of this grain boundary (see Fig. 4.24) explains why the 221 GB presents with faster diffusion compared to the 114 and 112 GBs.

Bond length and angle distributions presented in Fig. 4.26 highlight the lack of deviation of the 111 GB from the perfect diamond structure, with both bond length and angle distributions presenting single values at 1.545 \AA and 109.5° respectively. CN 4 carbon in the other grain boundaries showed some variation but distributions were centred about perfect diamond values. Where CN 3 carbon is present (for 221, 114 and 112 GBs), bond lengths are mostly shorter than CN 4 values. All three grain boundaries gave bond lengths in the $1.49\text{--}1.53 \text{ \AA}$ range, which align well with the range given for planar $\text{sp}^3\text{--sp}^2$ bonding in [187] and is likely a result of bulk sp^3 carbon bonding to sp^2 carbon within the grain boundary. Bond lengths around $1.40\text{--}1.42 \text{ \AA}$ present for the 114 and 112 GBs could indicate the presence of some pure $\text{sp}^2\text{C--C}$ bonds. The extreme high and low values present for the 114 GB are likely to be a result of the cut off used when determining coordination number. The bond lengths at 1.7 \AA are unlikely to be physical and instead could indicate some presence of sp^1 carbon as, if a shorter cutoff was used, these

would not be counted, and the carbon atom would be classed as CN 2/sp¹. The presence of sp¹ would also explain the bond length values of < 1.4 Å

Similarly, bond angles present consistent distributions around 109.5° for CN 4 carbon, and generally larger angles for CN 3 carbon. Here, only the 114 GB presents evidence of pure sp²C-C bonding, with bond angles around 120°. Although CN 3 carbon did give some bond angles < 120°, generally, bond angles were higher than CN 4 carbon as would be expected for carbon with fewer nearest neighbours.

Upon adding hydrogen to the grain boundary many atoms presented an increase in coordination number as to be expected. The exception to this is when the presence of the hydrogen has pushed a carbon atom beyond the cutoff distance, resulting in no increase in coordination number. This scenario was more common for CN 4 carbon. It can be challenging to conclude whether an increase in coordination number is a true indicator of the sp² going to sp³C-H bonding as suggested in this work, as such, changes in bond length and angle were also considered.

Changes in bond length and angle on the inclusion of hydrogen can be seen in Fig. 4.27. Here, only data from the carbon atom closest to each hydrogen atom have been considered and only carbon that displayed an increase in coordination number from 3 to 4 (sp² to sp³). As perfect diamond and the 111 GB contained no sp², these structures contained no carbon atoms that fulfilled these criteria. For all the grain boundaries with sp², there is a reduction in the short bond lengths associated with sp² and sp¹ carbon on the inclusion of hydrogen. The 221 GB gives a clear peak at bond lengths that align well with sp³ bonding, offering a good indication that sp³C-H may be forming. Similarly, adding hydrogen to the 112 GB removed the peak at 1.43 Å and shifted the most predominant peak to around 1.52 Å. Although this length is 0.025 Å shorter than seen for pure sp³, this increase still could indicate the presence of sp³C-H. The 114 GB gave a much broader distribution of bond lengths once the hydrogen was added, but again there was a general increase in bond length with the shortest values being lost. For this grain boundary in particular, the inclusion of hydrogen seems to result in a much more significant deviation from the initial structure and a general increase in disorder. This is despite the number of hydrogen atoms added being halved and the number of repeated calculations being increased to try and improve the clarity of the data. The 114 GB is the only grain boundary with no sp³ at the centre of the grain boundary, and the only one whose bond lengths suggested the presence of sp¹. This indicates more space is present within the grain boundary, allowing movement of carbon atoms and an increase in disorder on the inclusion of hydrogen.

There are more similarities between the grain boundaries when looking at the bond angles. Once hydrogen was added, the bond angle distributions reduced to a peak centred about 110° - in good agreement with what would be expected for the formation of sp³C-H. The 114 GB in particular showed a notable reduction in 120° bond angles which are commonly associated with sp² carbon.

Overall, in grain boundaries containing sp² carbon, changes in bond length and angle on the




Figure 4.26: Histograms comparing the bond lengths and angles for carbon carbon bonds within a supercell containing the specified grain boundary post energy minimisation. ‘CN 3/4’ refers to the coordination number of the carbon atom. Perfect diamond effectively gave single values (of 1.545 Å and 109.5°) as seen for 111 Σ3.

inclusion of hydrogen are present which could be associated with a transition from sp^3 - sp^2 to sp^3 C-H. The 221 GB gave the clearest evidence of this, whereas the 114 GB had a more general increase in disorder. The formation of sp^3 C-H is thought to contribute to restricted diffusion observed in these grain boundaries compared to perfect diamond or the 111 GB. Despite the 221 GB showing the clearest indication of sp^3 C-H forming, it gave the fastest diffusion and lowest activation energy of the grain boundaries containing sp^2 carbon. This is thought to be a result of the low energy channels present in this grain boundary, as well as the sp^3 carbon aligned with these channels. However, the formation of sp^3 C-H with sp^2 outside of these channels, still resulted in a higher activation energy than perfect diamond. The 112 GB showed evidence of sp^3 C-H forming and minimal channelling effects were present, resulting in slow overall diffusion. It cannot be said definitively that sp^3 C-H formed in the 114 GB due to the broad bond length distribution on the addition of hydrogen, however, some evidence of this was present. There were some channels observed for this grain boundary, but the changes in bond length and angle indicate notable deviation from the initial structure and such disorder is likely to outweigh the impact of these and restrict the overall diffusion rates.




Figure 4.27: Change in carbon bond length and angle from the inclusion of hydrogen atoms into the grain boundary post energy minimisation. Data shown only includes the carbon atoms closest to each hydrogen atom and only carbon atoms which went from a coordination number of 3 to 4 when the hydrogen was added.

4.4.3 Summary

MD simulations were performed to explore hydrogen diffusion in diamond grain boundaries. Four common grain boundaries were replicated from TEM images from [178], and it was found that all but the 111 GB had a significant potential well across the grain boundary. This resulted in hydrogen trapping, in good agreement with expected physical behaviour. Anisotropic diffusion was also observed in these grain boundaries with hydrogen preferentially travelling along channels in $\langle 110 \rangle$ but only within grain boundaries. The exception to this was the 111 GB, which gave isotropic diffusion and generally behaved much more similarly to the perfect diamond.

Overall diffusion rates are thought to be influenced by sp^2 content and channels present within the grain boundaries. When present, sp^2 carbon resulted in restricted diffusion from the formation of sp^3C-H bonds rather than the interstitial positions occupied in the bulk [90, 213–215]. In contrast, the presence of channels and sp^3 carbon in some grain boundaries acted to increase diffusion rates. Evidence of the formation of sp^3C-H with sp^2 carbon was observed through consideration of bond length and angle distributions.

Diffusion characteristics will dictate the impact annealing and other similar treatments will have on the hydrogen content of diamond. The conclusions that can be drawn regarding this from the results presented here are limited without an understanding of the distribution of different

grain boundaries present. For boron doped diamond, Mora *et al.* [209] reported 65 % of grain boundaries were identified as $\Sigma 3$, with $\Sigma 9$ and $\Sigma 27$ grain boundaries making up 30 % and <5 % respectively. Generally, it would be expected the lower order $\Sigma 3$ boundaries would dominate, as higher orders are typically formed from the meeting of two lower order boundaries. Results presented here suggest that diffusion characteristics are dictated by more than just the CSL (Σ) value, as 111 $\Sigma 3$ and 112 $\Sigma 3$ presented very different results. Therefore, a more detailed grain boundary distribution of exact grain boundaries would be required to fully understand the significance of each of these grain boundaries diffusion effects.

Nevertheless, there are some broader conclusions which can be made. For tilt grain boundaries rotated about $\langle 110 \rangle$, these results would suggest hydrogen would remain within grain boundaries during heat treatments, whereas hydrogen in bulk regions is more likely to be removed or become trapped in grain boundaries. Different grain boundaries clearly present different degrees of hydrogen trapping and diffusion rates, meaning the hydrogen content is likely to vary significantly between grain boundaries especially post heat treatment. Atomic alignment in $\langle 110 \rangle$ results in significant channels within the tilt grain boundaries. This may not be the case for $\langle 100 \rangle$ and $\langle 111 \rangle$, so more work would be needed to draw conclusions for other diamond orientations. Diffusion was shown to be highly directional and dependant on the grain boundary, with the 221 GB only giving significant diffusion in one direction. This behaviour could potentially be exploited to have greater control on transport/extraction of retained hydrogen isotopes. However, it is challenging to assess the potential of this without having a thorough understanding of the distributions of grain boundaries present.

Experimental work presented in 3.2 concluded sustained hydrogen counts suggested hydrogen diffusion from within the grain boundaries. Diffusion coefficients calculated here can be used to evaluate this conclusion. The root mean square displacement, $RMSD$, is given by

$$(4.7) \quad RMSD = \sqrt{6Dt}.$$

Taking the D to be the slower grain boundary diffusion coefficients, at a temperature of 1273 K for an hour gives

$$(4.8) \quad RMSD = \sqrt{6 \times 10^{-14} \times 3600} \approx 10 \text{ } \mu\text{m}.$$

These results suggest moderate diffusion within grain boundaries could occur at these temperatures, with hydrogen being able to travel from 10s of microns within the surface. Boron doped samples showed limited evidence of diffusion, suggesting structural defects caused by the inclusion of boron may act as trapping sites, limiting diffusion in a similar manner to the grain boundary diffusion observed computationally.

4.5 Hydrogen Bombardment of Diamond Grain Boundaries

As part of the work presented in 4.2 ([150]), data files containing grain boundaries were continuously bombarded with deuterium to evaluate differences that could be expected between SC and PC materials. Trimmed data files were tested with arbitrary rotations (15° , 30° and 45°) applied about the z axis to one of the grains to tilt grain boundaries with a (100) surface. These are not presented or discussed, but no discernible differences between retention or etching could be seen when compared to perfect diamond across the energies tested.

However, as highlighted in the 4.1, there were found to be numerous flaws with the trimmed style data files, and the use of arbitrary rotations may mean structures are physically irrelevant. As the replicated data files produced for diffusion calculations were stable and reproduced from real grain boundaries, bombardment of these should give a more accurate insight to any potential differences for PC diamond. Both repeated single bombardments and continuous bombardments were carried out on previously constructed replicated grain boundaries. Single bombardments were carried out to offer a clear comparison between perfect diamond and grain boundaries, whilst continuous bombardments give a better insight into how these grain boundaries may evolve under implantation conditions.

4.5.1 Method

The same replicated grain boundary structures produced for the diffusion work were used for these simulations. Again, these consisted of the 111 Σ 3, 221 Σ 9, 114 Σ 9 and 112 Σ 3, all orientated with $\langle 110 \rangle$ aligned with the x axis. As bombardments occur along the z direction, if the original orientation of the structures was used, only the 112 GB have a surface of physical relevance (the (111) surface in this case) along z . To avoid this, and ensure results between structures were comparable, the structures were rotated to align the common (110) face with the z axis, allowing this face to be bombarded. For comparison, bombardments of a replicated perfect (110) surface were also carried out.

For both continuous and repeated single bombardments of grain boundaries, deuterium atoms of various energies were incident at an angle of 0° . For continuous bombardment, deuterium was incident on the full surface (approximately $60 \times 53 \text{ \AA}^2$ in size), whereas repeated single bombardments were only carried out on the grain boundary itself. To help mimic a larger material, periodic boundary conditions were used in x and y and the bottom 3.567 \AA (1 UC) of the block was frozen in place to prevent movement during bombardment. A finite boundary condition was applied in z . The REBO potential [141, 142] was used throughout simulations to dictate carbon-carbon and carbon-hydrogen interactions. Temperature of the diamond was maintained at 300 K by simulating an *NVT* ensemble with a temperature damping parameter of 0.1 ps as before. Once again, a variable timestep was used that ensured no atoms moved more than 0.1 \AA in a single timestep, with a default value of 0.5 fs. Prior to bombardment, a 3 ps run was carried




Figure 4.28: Percentage of atoms deemed stable (not disordered) for the 221 GB bombarded with 30 eV deuterium along $\langle 110 \rangle$. The border of the disordered region and the point taken to be the surface are indicated by red and black dashed line respectively. The disorder border was defined as the highest z value where more than 50 % of atoms were disordered. As a result of how this was calculated, a random distribution of atoms would give approximately 40 % stable atoms. This has been overlayed on a snapshot of the simulation at the same timestep which was rendered in OVITO [182].

out to allow the temperature to stabilise.

For the repeated single bombardments, deuterium atoms were created in random positions in a region 3 UCs above the surface and within ± 3 Å of the centre of one of the two grain boundaries. Random seeds in the input file were regenerated for each repeat. Any deuterium that left the simulation box or was at a height greater than 2 Å from the surface was classed as reflected. Vacancies and interstitials were determined using the Wigner-Seitz analysis utilised in 4.3. This was done by comparing the occupation of Voronoi cells in the initial structure to the final structure which had been cooled to 0 K to remove thermal fluctuations (see Fig. 4.13).

In continuous bombardment simulations, 0.2 ps runs were carried out between atom creations. This is slightly shorter than the value for previous simulations (in 4.2) to ensure a comparable flux on the larger area of these structures. Fluxes and fluences ranged between $1.23\text{--}1.94 \times 10^{29}$ D m $^{-2}$ s $^{-1}$ and $2.47\text{--}2.80 \times 10^{20}$ D m $^{-2}$ respectively due to slightly different surface areas of grain boundary cells.

Similar evaluation methods were carried out to those used in 4.2 and 4.3 offering a direct comparison. To allow for swelling and etching phases present in continuous bombardments, the surface was determined through consideration of atomic density as was done before (see 4.2.1). Due to the presence of the grain boundary the atomic density was lower than that of perfect

diamond, so the reference atomic density was calculated for the structure prior to bombardment. The surface was once again taken to be the height where the average density of the 2 Å region below it dropped below half the reference atomic density. If at no point this condition was fulfilled the surface was taken as the initial height of 0 Å. Once the surface had been determined, sputtered carbon and retained deuterium could be found. Carbon outside of the surface was classed as sputtered and hydrogen below 2 Å above the surface was classed as retained. Depths were all measured from the calculated surface height.

In addition to the analysis carried out in 4.2, the position and size of the disordered region was measured for continuous bombardment. This was determined by consideration of the periodicity of atomic layers. Atomic layers are separated by 1.261 Å along $\langle 110 \rangle$. Looking at the remainder of an atoms z coordinate divided by this value gives the vertical distance from the perfect atomic layer position the atom sits. Any atoms greater than 0.25 Å from an ideal atomic layer was considered 'disordered'. Although atoms could be displaced along the xy components as well, the majority of displacements would be expected to align with the direction of bombardment (z). Only considering z displacement offered a simple and consistent way to estimate disorder as all files were orientated with $\langle 110 \rangle$ along z . The percentage of disordered atoms as a function of z could then be plotted. The highest depth at which more than 50 % of atoms were disordered was classed as the disorder boundary (see Fig. 4.28), which could then be tracked as a function of time or fluence. It is worth noting that, as atomic layers are separated by 1.261 Å and the region to be classified as not disordered is 0.5 Å thick, a completely random distribution of atoms would result in a disorder percentage of approximately 60.2 %.

As was carried out in diffusion simulations, atoms within the grain boundary could be distinguished to atoms outside of it. Atoms were classed as within the grain boundary if their x position was within ± 2.5 Å of a grain boundary centre - sharing the same definition as before. Atoms were classified as within the grain boundary or bulk prior to bombardment (as undamaged structure), meaning if a bulk atom was displaced into a grain boundary during a simulation, it would still be classified as a bulk atom during analysis. The grain boundary regions and atoms were used within analysis to separate out the bulk and grain boundary effects.

4.5.2 Results and Discussion

4.5.2.1 Repeated Single Bombardments

Fig. 4.29 shows the positions of defects for repeated single bombardment of the (110) surface of the different diamond structures with 60 eV deuterium. Both perfect (110) diamond, the 111 GB and previous perfect (100) diamond simulations (see 4.3) present similar behaviour. For these structures, deuterium is able to penetrate beyond the surface, but vacancy formation is restricted to the surface. Previously, it was concluded that there is not sufficient space in perfect diamond for carbon interstitials, so the vacant sites could not form without significant deviation from the perfect structure which was not possible to achieve at the energies tested. Therefore, interstitials

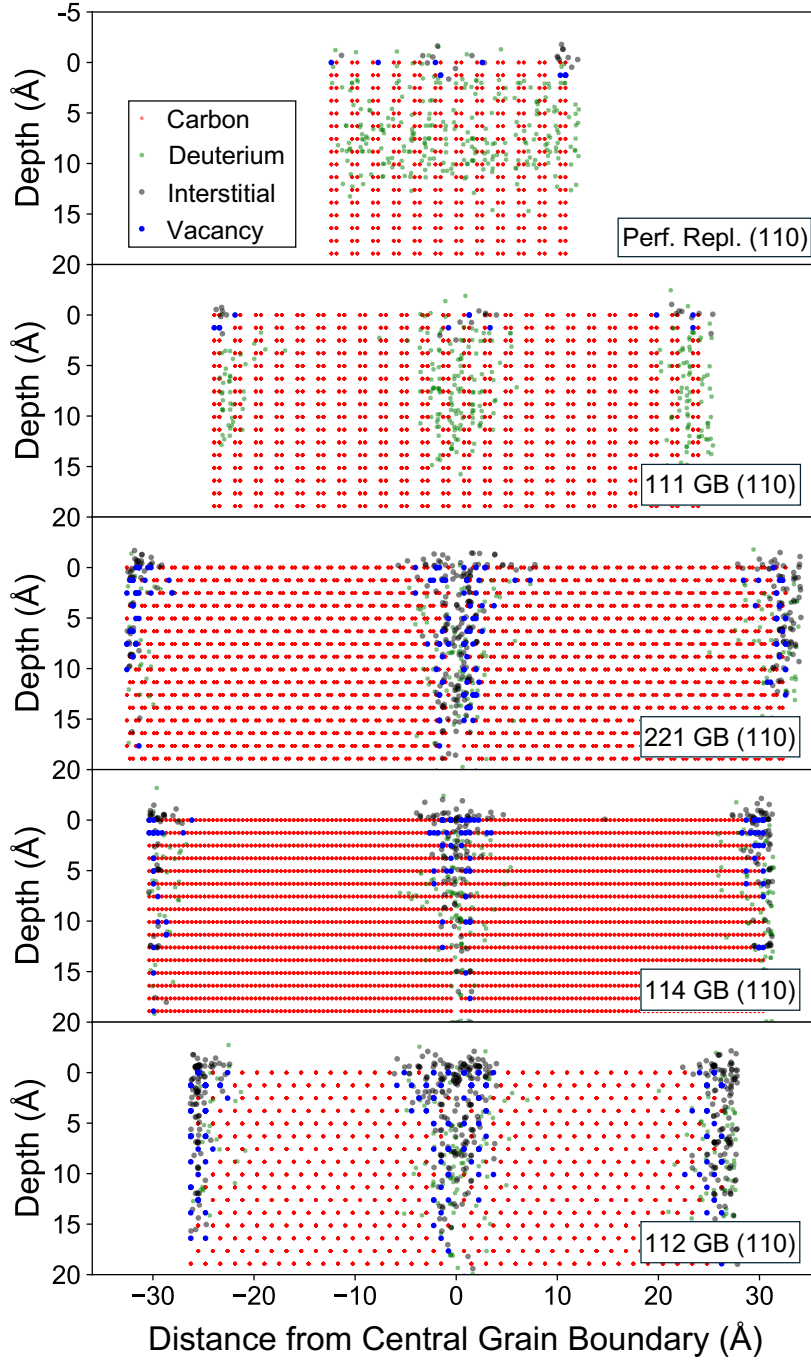


Figure 4.29: Defect analysis of 300 repeated simulations for five different diamond structures bombarded with single 60 eV deuterium atoms. All atoms were incident normal to the (110) surface, the bottom four structures contained a tilt grain boundary (GB) in the middle and at the edge of the simulation cell. The GB plane is indicated by the Miller indices in the legend. When a GB was present, only this region was bombarded. Carbon and deuterium atoms are given in red and green respectively, vacant sites (blue) and resulting carbon interstitials (black) were determined by a Wigner-Seitz analysis [195].




Figure 4.30: Defect analysis of four datasets of 300 repeats of a (110) diamond containing a grain boundary as indicated in the legend. The diamond was held at 300 K and bombarded with a single 30 eV deuterium atom normal to the surface. This process was repeated 300 times. y -axes units are as a percentage of incident atoms (for example, 40 % of incident atoms caused a vacancy in the second atomic layer for the 114 GB). Interstitials are the depth of any carbon atom with an occupancy of two, meaning each vacant site will result in two interstitials provided no carbon is lost. The percentage was halved to account for this. Dashed lines show the atomic layers of the diamond. Inset shows interstitials for the 111 GB.




Figure 4.31: Results from repeated simulations of single deuterium atoms bombarding various diamond structures at 300 K at an incident angle of 0° . Miller indices indicate the diamond surface, if present, the grain boundary plane is indicated, when not present legend states perfect ('Perf.'). 'Repl.' and 'Trim.' stand for replicated and trimmed which refers to the method used to create the structure.

could only form at the surface, where carbon could be pushed out and vacant sites were only found within the top few atomic layers. This resulted in a peak in vacancy formation at low energies when hydrogen isotopes would occupy shallow positions near to the surface. These observations are in contrast to what is seen when bombarding 221, 114 and 112 GBs. The additional space present in these grain boundaries allows interstitials and vacancies to form at much greater depths (comparable to the depth of the implanted deuterium) than seen in perfect diamond as seen in Fig. 4.29, where interstitials can be shown throughout the grain boundary as well as at the surface.

These observations are further supported in Fig. 4.30, which presents data for 30 eV deuterium incident on the (110) surface of the four grain boundaries⁵. With the 221, 114 or 112 GB

⁵Perfect diamond was also tested, but this gave very similar results to the 111 GB and as such was omitted from




Figure 4.32: Depth of 30 eV deuterium atoms incident normal to diamond (110) surfaces. Perfect diamond contained no grain boundary whilst the other structures contained a grain boundary as indicated in the legend. Perfect diamond (no grain boundary) Simulations consisted of single bombardments which were repeated 300 times.

present, a notable increase in vacancies can be seen, with peak vacancy depth beyond the surface for some. Once again this is a result of the depth at which there is space available for interstitials to occupy, which now can be seen within the diamond rather than just above the surface.

Fig. 4.31 also shows the increase in vacancy percentage for 221, 114 and 112 GBs. Perfect diamond gave very low vacancy formation in this direction for both replicated and trimmed data files and, once again, the 111 GB behaved very similarly to the perfect diamond produced by the same method (Perf. Repl. (110) in Fig. 4.31). As discussed in 4.3, the lower vacancy percentage of the perfect (110) surface compared to the perfect (100) surface is thought to be because of the link between penetration depth and vacancy formation in perfect diamond, and $\langle 110 \rangle$ channels and increasing average depth. The additional space present for the 221, 114 and 112 GBs resulted in a general increase in vacancies, as well as the formation of vacancies at depths beyond the top unit cell. The ability for vacancies to form deeper within the material removes the link between penetration depth and vacancy formation. In contrast to the 30 eV peak seen for perfect (100), the 221, 114 and 112 GBs, present no low energy peak in vacancy percentage, it simply increases in an almost linear fashion with energy. These grain boundaries also demonstrated an increased sp^2 content (see 4.3), but it is thought that the increase in vacancies is a result of additional

the plot to aid clarity.




Figure 4.33: Percentage of 300 repeated simulations of a deuterium atom of various energy incident normal to a (110) diamond surface containing various grain boundaries. Reflected atoms were classed as reflected if they were lost from the simulation box or more than 2 \AA from the surface and atoms within 2 \AA of the surface were classed as surface bound. Total is the sum of the two percentages and indicates the number of atoms which did not penetrate the surface.

space rather than the hybridisation. This is because sites of sp^3 carbon can still become vacant as seen in the results for perfect diamond in positions deeper than the top atomic layer.

Small amounts of deuterium at high depths in Fig. 4.30, the increase in average depth in Fig. 4.31 and the long tail present in Fig. 4.32 indicate channelling of deuterium along the 221, 114, and 112 GBs. Even at lower energies, small amounts of deuterium can be seen at very high depths, far greater than the average values and depths of perfect diamond and the 111 GB. The additional space in these grain boundaries offers a clear path for incident atoms, even if diffusion was slower within them compared to perfect diamond in 4.4. At these energies, penetration depth is dictated by collisions with atoms rather than diffusion characteristics.

Bombarding the 221, 114 and 112 GBs resulted in lower amounts of reflected atoms compared to perfect diamond and the 111 GB as shown in Fig. 4.33. As momentum/energy of incident atoms increased, any differences become negligible as the total percentage drops and the majority of atoms are able to penetrate through the top atomic layers. This difference was most obvious at a vertical momentum of $4.7 \text{ eV}^{0.5} \text{ u}^{0.5}$ (corresponding to 10 eV deuterium). At higher momentum than this, a non-linear decrease in reflected atoms similar to that presented for the (100) surface (Fig. 4.14) can be seen. However, at lower momentum, a levelling off or even decrease is present in the number of reflected atoms. Fig. 4.33 shows this decrease in reflection is a result of an increase in surface bound atoms rather than an increase in atoms penetrating the surface – which is at its lowest for this momentum. This can be seen clearest for the 111 and 114 GBs, where almost all atoms were either reflected or bound to the surface at low momentum, so the decrease in reflection at very low momentum must be a result of an increase in surface bound atoms. The 221 and 112 GBs gave the lowest total values, indicating that even at low momentum/energy deuterium was able to penetrate beyond the surface. Looking at the structures of these grain boundaries along $\langle 110 \rangle$, clear channels can be seen as discussed in 4.3 (see Fig. 4.21 and Fig. 4.23). If an atom is incident on one of these channels it is able to travel into the material even at low energy. Whereas, for the tighter atomic structures of the other surfaces, low energy incident atoms are more likely to collide with surface atoms and become reflected or bind to the surface. Additionally, the 221 and 112 GBs are also the grain boundaries which presented the highest average penetration as a result of these channels.

4.5.2.2 Continuous Bombardment

When bombarding grain boundaries with a fluence, much of the same behaviour presented in 4.2 for perfect (100) diamond was seen. Fig. 4.34 shows the surface height as a function of fluence for the different structures and energies tested, equivalent to Fig. 4.9 in 4.2. Similarly to these previous results, the two-step etching process can be seen, where an initial swelling period is required prior to the removal of carbon atoms. The etch point is present for all but the 10 eV simulations suggesting, at these energies, incident atoms are of insufficient energy to form a significant disordered region and remove carbon atoms. Once again, higher energy deuterium




Figure 4.34: Diamond surface height for various diamond structures bombarded with deuterium of different energies. Surface begins at 0 Å, negative values indicate swelling of the surface, whereas positive values of the surface being eroded. The surface was defined based on a fraction of initial atomic density.

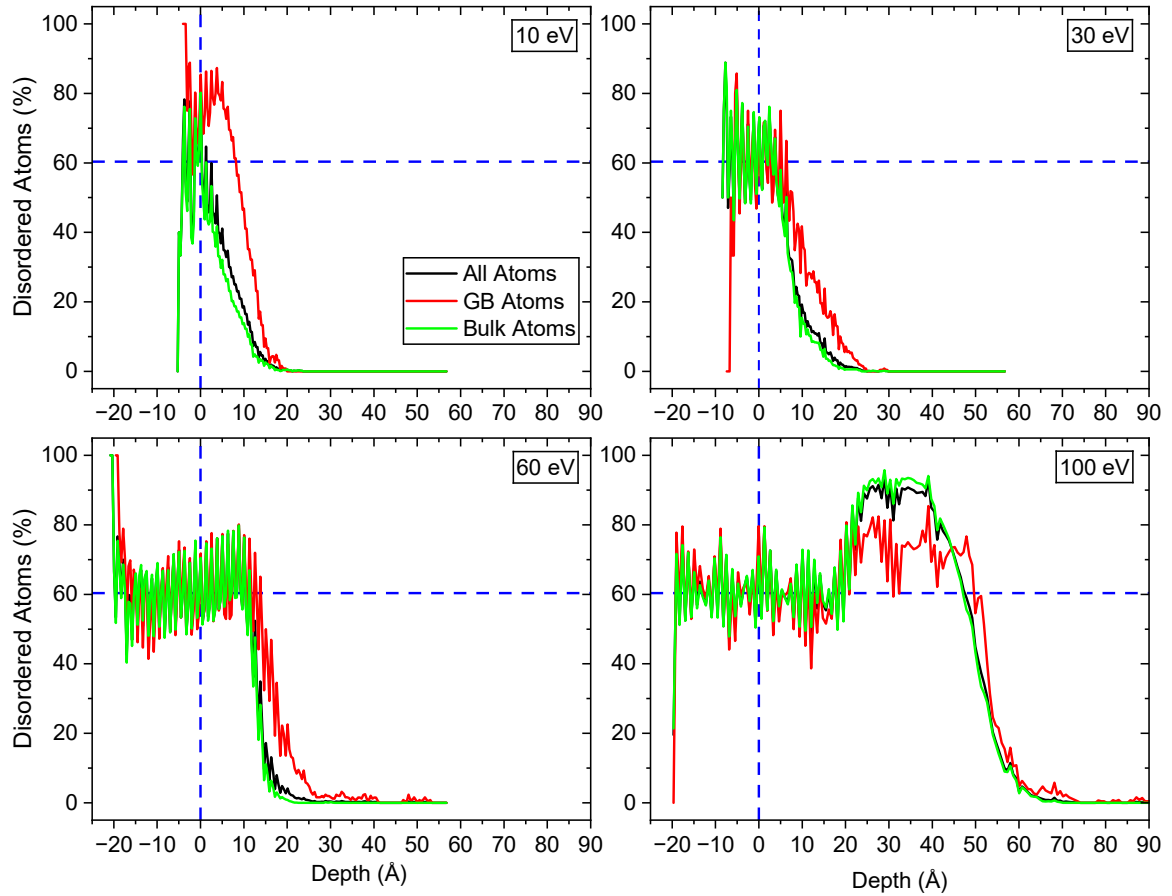


Figure 4.35: Percentage of disordered atoms as a function of depth for the 221 Σ 9 grain boundary bombarded with a fixed fluence of deuterium at various energies. Atoms were classed as either bulk or GB atoms based on their position prior to implantation. Due to the method used to determine disorder, a completely random structure would result in 60.4 % disordered atoms, values consistently above this are likely to indicate a shift of atomic layers in z .

resulted in an increased rate of swelling, a later etch point and an increased rate of etching. Differences between the different structures are subtle with all structures displaying similar gradients (swelling and etch rates). The 221 and 112 GBs seem to reach the etch point slightly later for 30, 60 and 100 eV. This is unlikely to be significant but could be a result of increased space within these GBs meaning an increased fluence is required to reach equivalent disorder.

Fig. 4.35 shows the degree of disorder at varying energy for the 221 GB. The disorder percentage can be seen to fluctuate about of 60 % (indicating random atomic positions) for all but the 10 eV simulation. At 10 eV, disordered atoms are present deeper within the grain boundary compared to the bulk. The peak percentage of disordered atoms is also greater than the peak of bulk atoms but at 80 % is likely to indicate a shift in atomic layers. 30 and 60 eV simulations also show the same increased depth of disorder within the grain boundary. The 100 eV shows very

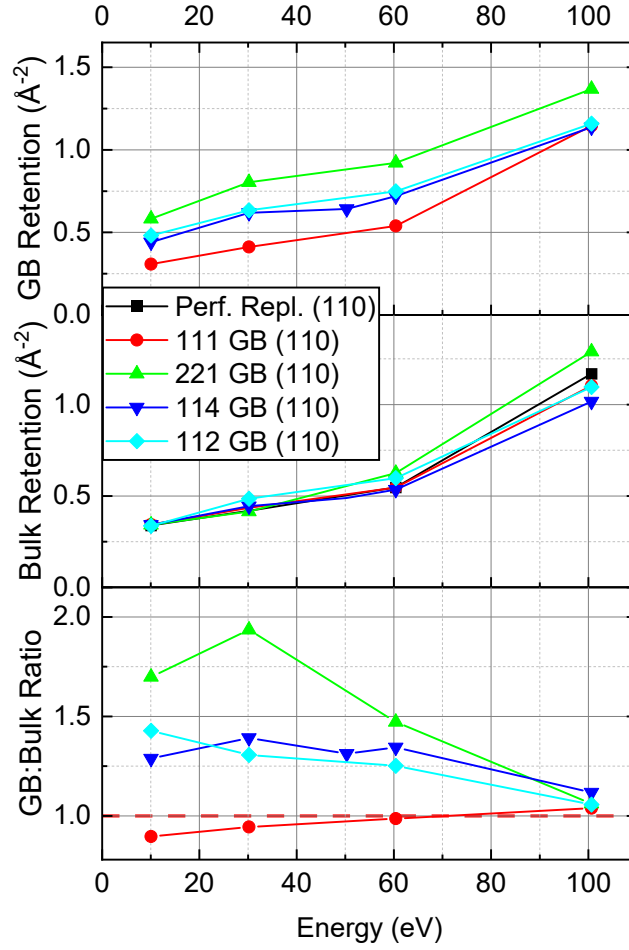


Figure 4.36: Retention within different structures bombarded with deuterium of various energy. Retention values are given by the number of retained atoms divided by the exposed surface area, resulting in units of \AA^{-2} . Due to varying surface areas of structure cells fluence did vary small amounts between structures but was approximately 2.5 \AA^{-2} . The top plot shows retention of deuterium within $\pm 5 \text{ \AA}$ of a grain boundary centre, the second plot anything outside of this, and the final plot is the ratio between the two.

little difference between GB and bulk disorder. A peak is seen at greater than the 60 % value for a random distribution, again this could be indicating a shift of atomic layers. There is some evidence of disorder at high depth within the grain boundary, where early on in the simulations incident atoms have travelled some depth into the grain boundary.

Generally, disorder within the grain boundary could be seen at higher depths than disordered atoms within the bulk. The exception to this was the 100 eV simulations and the 111 GB, where there were negligible differences between datasets. At this higher energy, incident atoms are able to create disorder in bulk regions just as effectively as grain boundaries so minimal difference is




Figure 4.37: Deuterium depth profiles of various diamond structures exposed to deuterium of varying energy along the $\langle 110 \rangle$ direction. Retention values are given by the number of retained atoms within a slice in z (Depth) divided by the volume.

observed.

The 221, 114 and 112 GBs exhibited higher deuterium retention than both bulk regions within the same simulations, as well as a perfect (110) structure. This can be seen in Fig. 4.36, where the ratio of deuterium retained in the GB:Bulk is greater than 1 for all but the 111 GB at energies less than 100 eV. The 221 GB in particular showed the highest ratio, with nearly twice as much deuterium retained within the grain boundary compared to the bulk. Broadly, a gradual decrease in ratio is seen with increasing energy, with 100 eV showing no preferential retention within the grain boundary. Across all energies, bulk retention values are approximately the same. There is some spread at 100 eV, but this could be a result of variations in determining the surface from the greater disorder present at this energy. However, as the surface height is the same for both GB and bulk regions, this should have minimal impact on the retention ratios, which clearly




Figure 4.38: Comparison of two fluences for simulations of 30 eV deuterium bombardment of a (110) diamond surface containing a 221 Σ 9 grain boundary.

show increased GB retention at low energies.

Depth profiles for the different structures can be seen in Fig. 4.37. For 10, 30 and 60 eV, the grain boundaries depth profiles show a slight broadening, indicating deuterium retained at slightly increased depths compared perfect diamond and the 111 GB. As the main peak is of similar height across the structures, the increased retention of grain boundaries is due to these slightly higher depths. Particularly for energies ≥ 30 eV, depth profiles of 221, 114 and 112 GBs show a long tail of deuterium reaching very high penetration depths. At 100 eV, deuterium profiles are more or less indistinguishable from one another, a result of the increased disorder at this higher energy removing any impact the presence of a grain boundary makes. This is in good agreement with the observations for the GB: Bulk retention ratio previously, which went to 1 for 100 eV.

The observation of deuterium at very high depths is similar to results seen for single bombardments (see Fig. 4.32), suggesting incident atoms are being channelled along grain boundaries. This can only happen for initial bombardments, prior to the formation of the disordered region, as channels will no longer be accessible after this point. Therefore, the impact such channels have on retention is thought to be minimal for any moderate fluence. Similarly, single bombardment results also showed lower reflection for 221, 114 and 112 GBs, but as reflection is largely dictated by atomic structure on the surface, these differences are unlikely to be present once the disordered




Figure 4.39: Deuterium density along the x axis for the $221 \Sigma 9$ grain boundary bombarded by deuterium at a fluence of 2.47 \AA^{-2} and varying energy. Deuterium density values include all deuterium within the simulation box so includes both surface bound and retained deuterium.

region has formed. Instead, the increase in retention is thought to be a result of the decreased atomic density in these grain boundaries, allowing deuterium in the disordered region to be pushed deeper into the grain boundary compared to the bulk, as can be seen in the wider depth profiles in Fig. 4.37. A high fluence simulation of 30 eV deuterium incident on the 221 GB was also carried out (see Fig. 4.38). The ratio of GB:Bulk retention remained consistent even during etching. This further supports the idea that the increase in retention is not a result of channelled deuterium at very high depths, as the ratio remains constant even as the material is etched away. Results from a shorter fluence simulation were also included in this figure which demonstrated good reproducibility.

The energy of the incident atoms dictates the thickness and time taken for the disordered region to form, as discussed in 4.3. At high energies, atoms are more capable at displacing carbon, penetration depth is increased, and a higher proportion of atoms are able to penetrate the surface, resulting in a thicker disordered layer. The greater disorder means the presence of a grain




Figure 4.40: The depth of the disordered region of diamond structures bombarded with deuterium at different energies. The disordered depth was defined as the highest depth at which more than 50 % of atoms were disordered (based off of a deviation of greater than 0.25 Å from atomic planes in z).

boundary has minimal impact at 100 eV but is more significant for the shallower disorder at lower energies. Again, this can be seen in Fig. 4.39, which shows an increased deuterium density across the grain boundary for 10, 30 and 60 eV but not 100 eV. Similar results were collected for other grain boundaries with the exception of the 111 GB which gave no notable increase across the grain boundary and very similar retention to bulk regions. This increase in retention is in good agreement with similar simulations of other diamond grain boundaries [78].

Fig. 4.40 shows the maximum depth at which more than 50 % of atoms were classed as disordered as a function of fluence. With increasing energy, the disordered region reaches deeper

into the material as a result of the increased penetration depth. There are also subtle differences between the different structures tested. The results for perfect diamond closely tracked those of 111 GB as has come to be expected. Both of these consistently gave the lowest depths across the four energies tested. The inclusion of 221, 114 and 112 GBs increased the depths the disordered region reached, with the 221 GB presenting the most notable increase and the other two displaying similar values. Again, this is thought to be a result of the increase in penetration depth and space within these grain boundaries. Beyond initial peaks (likely a result of the method used rather than a point of physical importance) less variation is seen in this data than the surface height data. The swelling phase of the surface seemingly has minimal impact on the depth of disorder despite the surface being of greater height. A gradual increase in depth, presumably as a result of surface etching can be seen towards the end of some of the simulations. Although not shown, a gradual increase in the disorder depth was seen for the high fluence, 30 eV, 221 GB simulation, confirming the disordered region moves with the etching of the surface and has a consistent thickness.

Depth profiles presented in Fig. 4.37 highlight a potential limitation of tracking the depth of the disordered region in this manner. For 10, 30 and 60 eV, the depth of the disordered region shown in Fig. 4.40 broadly corresponds to upper end of the depth profile presented here. 100 eV simulations are the exception to this. Looking at the depth profile, it might be expected that the depth of disorder may reach around 20 Å, however the value is closer to 40 Å. The reason behind this can be seen in Fig. 4.35. For 100 eV, the randomly disordered region (corresponding to 60 % disordered atoms) does end at around 20 Å as would be expected from the depth profile. However, the border was taken to be at around 50 Å, where values first rise to close to 100 % which, as previously discussed, is more likely to correspond to a shift in atomic layer rather than true disorder. This needs to be considered for higher energy implantations, but the others follow the expected trend.

Data presented in Fig. 4.41 shows the number of atoms lost from the GB, bulk and full surface respectively. As carbon can move between GB and bulk regions, these plots have been labelled as 'lost', whereas any carbon lost from the full surface must have been sputtered. For low energies, negative values can be seen for GB atoms, indicating an increase in carbon atoms within this region. This is a result of carbon atoms from the bulk being pushed into the grain boundary at a greater rate than carbon was being sputtered from the surface. As the energy increases, this effect is reduced as the rate of sputtering dominates. Similarly, lost bulk atoms will be a combination of sputtered atoms and atoms moving into the grain boundary and as such these are consistently higher than GB lost atoms for an equivalent energy and structure. The sputtering yield for structures containing a grain boundary are generally slightly higher than perfect diamond. This indicates that, overall, the presence of most grain boundaries increases sputtering yield despite the GB lost atoms values being lower than bulk regions due to the movement of carbon between the regions. The 111 GB gave a sputtering yield closest to that of perfect diamond, although




Figure 4.41: Lost carbon atoms for simulations of diamond structures containing various grain boundaries bombarded with deuterium. Legend indicates the grain boundary present. GB/Bulk atoms lost per incident atom is similar to sputtering yield, but atoms can move into between regions as well as being sputtered so this is atoms lost from the region. Values for a perfect structure containing no grain boundary have been included for comparison for all plots.

values lost for the different regions suggest some carbon is still moving into the grain boundary region.

4.5.3 Summary

The clearest differences between grain boundaries and perfect diamond were seen for repeated single bombardments on a pristine surface. In these simulations, the 221, 114 and 112 GBs showed channelling along the grain boundary plane which resulted in high penetration depths and reduced reflection from the surface. The presence of these grain boundaries also gave space for interstitial carbon, allow vacancies to form at much higher depths than for perfect diamond.

There were minimal differences between the 111 GB and perfect diamond due to similarities in structure.

Channelling effects had limited impact on results from continuous bombardment simulations due to the formation of a disordered layer, particularly for higher energy bombardment. The disordered layer blocked channels reducing their effect on both reflection and penetration depth. For the thinner disordered regions present for lower energy (10 – 30 eV) bombardments, some effects of the grain boundary were still prominent. Depth profiles showed a small increase in penetration at the grain boundary, alongside a handful of atoms at very high depths which travelled down grain boundary channels prior to the disordered region forming. The small increase in depth at the grain boundary is thought to be a result of the increased space, allowing more deuterium atoms to be pushed in. As a result of this, greater retention was observed in 221, 114 and 112 GBs at low energies, whilst negligible difference was seen at 100 eV. These grain boundaries also seemed slightly more susceptible to damage, with disordered regions reaching deeper into the material at low energy, and the sputtering yield of full surfaces being slightly higher than that of perfect diamond. Evidence of carbon moving from bulk regions into grain boundaries was also observed as the disordered region forms.

Overall, it appears that channelling effects arising from a distinct atomic structure dictates results for a pristine surface, whereas the additional space present in grain boundaries gives rise to the observations made for continuous bombardment simulations. Different grain boundaries presented different behaviours, so an understanding of grain boundary distribution would be required to fully understand their impact on retention and etching. In 4.4, it was noted that grain boundaries could play a disproportionate role in determining overall diffusion characteristics, as hydrogen is typically concentrated within grain boundaries. In contrast, bombardments are equally distributed across an entire surface, so grain boundaries would have to present very dominant differences to bulk regions to make a significant impact to overall effects. At lower energies and lower fluences, the distinct effects that arose from channelling are still present (if diminished) making the presence of grain boundaries more significant. At energies greater than 100 eV, it seems unlikely grain boundaries would play a significant role in determining both retention and etching. Although differences were more distinct at low energies, they remain reasonably subtle and as bulk diamond would make up the vast majority of PC diamond the overall impact is thought to be low. These observations support the idea that ion damage of the surface may be the cause of the additional desorption peak observed for PC diamond exposed to higher energy deuterium ions (see 3.2). Despite simulations suggesting negligible impact from the presence of grain boundaries, trapping of deuterium for low energy implantation in 3.2 was attributed to grain boundaries. This could be explained by deuterium incident on the grains diffusing into grain boundaries during implantation - an effect that would not be observed on the short timescale of simulation but could occur during the 5 hr implantations of physical sample. Therefore, the preferential trapping of hydrogen in grain boundaries observed in diffusion

simulations may be important for physical implantations.

4.6 Conclusions

Simulations were carried out in LAMMPS to develop understanding of hydrogen isotopes interactions with diamond. These simulations explored continuous bombardments, repeated single bombardments and diffusion calculations for diamond of different orientations, structures and temperatures, and hydrogen isotopes of different masses, energies and incident angles. Results from simulations offered insight into experimental findings presented in Chapter 3.

Continuous bombardment simulations gave retention results to the same order as experimental work and the computational work of others [82]. The maximum deuterium depths were well within the 10 nm limit given by NRA measurements (see 3.2). Limited sputtering was observed initially, but with increasing fluence a disordered surface region formed resulting in steady state etching. The thickness of this region was dependant on incident atom energies. At energies relevant for ITER and fusion reactors, a thickness of order nm could be expected, leaving the vast majority of the bulk unaffected and the amount of etched material low.

Repeated single bombardments suggested different orientations played a significant role in reflection, penetration depth, and vacancy formation. These simulations also showed that, at these low energies, vacancies could only form where there was space for interstitial carbon and were only observed on the surface or along grain boundaries.

Diffusion characteristics in PC diamond are thought to be significantly impacted by the presence of grain boundaries, which showed hydrogen trapping and restricted diffusion as a result of the formation of $\text{sp}^3\text{C-H}$ bonds. Diffusion rates support the conclusion that hydrogen diffusion was observed at the higher temperatures in experimental work. Despite grain boundaries making up a relatively small degree of PC diamond compared to the bulk, hydrogen trapping within them is thought to mean they play a disproportion role in the overall diffusion effect. The same cannot be said for continuous bombardments of grain boundaries. Although a small increase in retention and penetration depth was observed, the impact of this is thought to be negligible when considering the proportion of grain boundary to bulk carbon within a microcrystalline diamond - particularly for bombardments at ≥ 100 eV. These subtle changes are thought to be a result of the increased space within grain boundaries, in contrast to the more obvious channelling effects seen when the disordered region is not present. For the single bombardments of pristine diamond, the distinct atomic structure present gave clearer differences between perfect diamond of different orientations as well as grain boundaries. The formation of the disordered layer in continuous bombardment simulations dampened the impact of variables relating to the atomic structure.

4.7 Future Work

The work presented in this chapter has aimed to assess a variety of different variables to gain a broad understanding of diamonds interaction with hydrogen. This approach has offered valuable insight and developed understanding of experimental work, however, carrying out simulations that matched experimental work closer could help unpick results further.

The most notable difference between simulation and experimental work is the clusters which are typically incident in DELPHI and other low temperature plasma sources but have not been simulated here. An attempt at this was made, but clusters of three deuterium atoms were very unstable in simulation and would simply break apart prior to impacting the surface during bombardments. It may be possible, perhaps with use of a different or new potential, to accurately model D_3^+ clusters to explore how the energy of the cluster is distributed between the three atoms. In 3.2, it was speculated that clusters may split into a single deuterium atom and a pair, explaining the amount of HD present in TDS data despite the lack of D^+ extracted from the plasma. Further simulation work may be able to support or contradict this theory.

Although cluster simulations are of interest from an experimental aspect, the incredibly high temperature of a fusion plasma means it will consist almost entirely of single ions rather than clusters. This increases the relevance of this work, even if simulating clusters could help bridge the gap between experimental setups and fusion reactors. That being said, simulation parameters could be adjusted to closer replicate fusion relevant conditions. Performing continuous bombardment simulations at higher temperatures, as well as simulating dramatic changes in temperature to mimic plasma instabilities, would be one way to tailor simulations towards more fusion relevant scenarios. Additionally, the impact of neutron irradiation has not been explored. Limitations on system size and inability to reproduce transmutation or the quantum interactions of neutrons means MD is not appropriate for simulating 14 MeV neutron irradiation. However, there are other ways of mimicking neutron damage in classical MD simulations. The primary knock-on atoms (PKAs) from neutron irradiation can be simulated by imparting a certain amount of energy to an atom within the system [219], or damage could be introduced by bombarding with another species prior to retention simulations. However, the energy of PKAs/incident atoms would still be restricted by system size. Instead, physical defects present in neutron irradiated diamond could be recreated to allow the impact of neutron damage to be explored without the unfeasible system sizes required for high energy bombardment.

Additional steps could also be made to improve the physical accuracy of the diamond structures. Although grain boundaries were explored, this was limited to four structures due to a lack of atomic resolution images available in the literature. Expanding this work to additional structures would give a broader understanding of the expected behaviours of grain boundaries and PC diamond. As reference throughout this work, CVD diamond is expected to contain high concentrations of hydrogen within the grain boundaries [90]. Placing hydrogen within grain boundaries before bombardments may prevent the small increase in retention seen for continuous

bombardment of grain boundaries. Incident atoms may also be able to remove loaded hydrogen, which could act to contaminate/cool the plasma in the context of fusion reactors so would need to be assessed. Lastly, it would be interesting to explore diffusion characteristics of the structures post bombardment, as retained deuterium may diffuse deeper into the material or simply be lost from the surface.

DESIGN, ASSEMBLY AND TESTING OF A NEW LOW ENERGY ION SOURCE

Up to this point, experimental work had revolved around exposing diamond samples to low energy deuterium ions in order to evaluate retention. These experiments had been carried out in DELPHI (Device for Exposure to Low energy Hydrogen Isotopes) at the Culham Centre for Fusion Energy (CCFE) - a setup designed to perform such tests. Results from this facility can be seen in Chapter 3, and an overview of the setup in Chapter 2. DELPHI underwent an upgrade process to allow tritium implantations to be carried out along with the previously used deuterium, leading to a shutdown period of over two years and severely impacting the experimental work possible. Driven by this, and the desire to be able to carry out retention experiments on site, the decision was made to develop a new experimental setup to carry out similar experiments in the University of Bristol CVD Diamond Lab. This chapter outlines the design, assembly and testing of ExTEnD (EXposure To low ENergy Deuterium).

The author would like to acknowledge and thank the following people for their contribution to this work. James Smith (Uni. of Bristol) offered his guidance, technical knowledge and practical assistance throughout every aspect of this build. Dan Faircloth (UKRI), Olli Tarvainen (UKRI) and Marco Schippers (PSI) provided advice surrounding ion optics and biased sample stages. Technical information regarding the design and operation of DELPHI was supplied by Ant Hollingsworth (CFS) and James O'Callaghan (UKAEA). The mechanical and glass workshops made the custom parts required for this build. TDS measurements were performed by A. Zafra (Uni. of Oxford), who also discussed results with the author.

5.1 DELPHI

DELPHI (see 2.2) was constructed with the aim of exploring retention mechanisms in potential PFMs and other fusion relevant materials. As discussed in Chapter 1, retention of hydrogen isotopes (particularly tritium) is a huge challenge facing the success of commercial fusion. This facility is one of a broad variety of setups used to explore retention in fusion materials, including ions beams [220] and linear plasma devices [221], plasma ion extraction [120, 192, 222], plasma submersion [224] and gas permeation at elevated temperature/pressure. Ion extraction setups (such as DELPHI) typically offer the greatest control over implantation conditions, allowing low energy implantations at reasonable fluences without submersion in the plasma. From the outset, the goal of this facility was to perform tritium implantations, which significantly complicated the setup due to the notable challenges surrounding safe tritium handling. From an operational standpoint, the biggest challenge this introduces is the requirement for the entire setup to be housed in a glove box. Using tritium also places restrictions on many aspects of the design, from the materials that can be used to the gas flow systems. As ExTEnD will exclusively be used with deuterium, many of these features can be negated, significantly simplifying this new setup. The backbone of DELPHI consists of three main components: a plasma, an ion extraction system and a sample stage. Any equivalent system must include versions of each of these components.

5.2 Context and Aims

Once it became clear that completing the planned experiments in DELPHI was unlikely, alternatives were explored. Initially, established facilities were considered. This included both plasma-based sources such as DELPHI and more traditional ion beams. Although ion beams can carry out deuterium implantations, typically, they operate at ion energies too high to be relevant for plasma facing component applications. Setups more similar to DELPHI were contacted, but access to these was not possible.

It was then considered that, by replacing the H_2/CH_4 gas mixture used for diamond growth with D_2 , many of the CVD diamond reactors in the lab could be used to strike a deuterium plasma. If the ion optics and sample stage could be housed within the chamber, one of these could make a viable alternative to DELPHI. The reactor of particular interest was the largely unused pulsed deposition reactor (PDR, see 2.1.2), in which a plasma is formed via an electrical discharge between two molybdenum electrodes. Compared to the other reactors, the large size of the PDR's chamber offered the best potential to house the additional components.

Ultimately it was deemed unfeasible to adapt the PDR in this manner. Despite the chamber being one of the largest in the lab, the main problem was trying to fit the ion optics and sample stage into the chamber. The PDR was designed to be run at relatively high powers (up to 3.5 kW) and as such required additional considerations such as water cooling, which limits space within the chamber. The deuterium plasma in DELPHI is a far lower power (around 50 W), making

many of the PDR's design features excessive for the planned experiments. The PDR's electrodes are also mounted vertically, meaning the ion optics and sample stage would have to be mounted on the side of the chamber rather than more logical position at the bottom. Furthermore, the chamber only has one 6" mounting flange, currently used as a viewport. To adapt this reactor, this one port would need to be used to mount all the additional components (viewport, sample stage, ion extraction system and required feedthroughs) which would have been very challenging. Lastly, the chamber was covered in carbon residue and would need to be thoroughly cleaned to avoid contamination.

Instead, it was decided the best approach was to design and build a dedicated ion implanter. The new setup was to operate in a similar manner to the PDR (creating a plasma via electrical discharge) even sharing the same power supply, but with the addition of a sample stage and ion extraction system. This maintains the simplicity of a discharge plasma, whilst the separation between plasma and sample prevents heating, and allows for better control over ion energy and more accurate fluence estimation. Several iterations and designs were considered prior to the final design being reached. The following sections will outline the design considerations which led to this final design. There were three main targets of this setup:

1. **Maximise ion current** - In order to ensure a good fluence can be reached in reasonable a time, the ion current incident on the sample must be maximised. Ideally, an exposure time of six hours or less would result in detectable levels of implanted deuterium. Under standard operating conditions, DELPHI could achieve a flux of approximately $3.0 \times 10^{17} \text{ D m}^{-2} \text{ s}^{-1}$, resulting in fluence of at least $5.5 \times 10^{21} \text{ D m}^{-2}$ [120] during a five-hour exposure possible in working day. Ideally, this setup would be able to reach similar fluences to DELPHI.
2. **Energy selection** - Results presented in Chapter 3 showed ion energy plays an important role in retention mechanisms. As such, having some control over ion energy is important to maintain in this new setup, especially if results are to be compared to DELPHI. For fusion applications, low energy ions are of particular interest, DELPHI typically operated at an ion energy of a few hundred eV to ensure measurable amounts of deuterium could still be implanted in a reasonable time.
3. **Current measurement** - An accurate estimation of fluence hitting the sample will be key to evaluating retention. In DELPHI, ion current on the area surrounding the sample combined with knowledge of the beam divergence was used to estimate the fluence incident on the sample.

DELPHI had some additional features which were not possible to include in the new setup. Firstly, DELPHI's sample stage could be cooled or heated to vary the stage temperature from 4 to 450°C [120]. Although heated sample stage would be fairly straight forward to include, this was not deemed necessary at this point due to limited time and the fact that all previous implantations

were carried out at room temperature. Similarly, continuous temperature measurement was not thought to be required, as large increases in temperature were not expected from the low power plasma.

The new setup will also neglect accurate characterisation of extracted ions, which are likely to be a combination of D_3^+ , D_2^+ and D^+ . For a setup similar to DELPHI [192], the former dominated, giving an average of 2.96 deuterium atoms per incident ion (as determined using an energy resolving mass spectrometer). Due to the similarities in the setup, the same value was used for DELPHI without additional measurement [120]. This value impacts both the current to fluence conversion and the estimated average energy per deuteron. Manhard *et al.* [192] also demonstrated variation in extracted ion composition as a result of microwave power and deuterium gas pressure, suggesting there could be a small deviation from the 2.96 value for this new setup, as both the operating pressure and plasma striking method differ. Nevertheless, it is still likely that D_3^+ would dominate, as this is formed via the ion-molecule reaction



which is particularly prevalent for low temperature plasmas and higher pressures [192] as would be the case for the new setup. As such, in the absence of a method for accurate characterisation, it seems reasonable to apply the same 2.96 value assumed for DELPHI. Ultimately, this will be a limitation on any conclusions that can be made regarding total retention and exact ion energy. However, using TDS to analyse samples gave total retention as well as insight into retention mechanisms, allowing this setup to provide information of interest. Furthermore, for this and similar setups, fluence is ultimately an approximation based on ion current. As such, small variations in the cluster size should not impact the ability to compare inventory values to other setups with similar conditions.

5.2.1 Design Considerations

5.2.1.1 Striking Plasma

For simplicity, the new setup still relies on electrical discharge to strike a plasma within the chamber in a similar manner to the PDR. A microwave source would be preferable from an operational point of view and was used in DELPHI as part of a SPECS ion source. The ion source is a commercially available component, consisting of a microwave source, a small plasma chamber with a magnetic quadrupole and ion optics (see 2.2). This ion source allowed for low operational pressures (of order 10^{-3} mbar) - decreasing the likelihood of ions being neutralised or scattered as they travel to the sample.

Use of microwaves was deemed unfeasible due to the significant cost and complexity such a setup would entail. Using an electrical discharge is a much simpler method, and meant the power source from the PDR could be used. This power source is not only capable of the high voltages required to strike a plasma but, as a pulsed power supply, offers arcing suppression. If the power




Figure 5.1: Figure taken from [225]. Paschen curves for H_2 and D_2 obtained using voltage ramp (VR) and low current limit (LCL) methods. Data from other studies [226–230] is also presented.

supply detects the high currents characteristic of an arc, it will skip subsequent pulses until the lower current, typical of a plasma discharge, resumes.

5.2.1.2 Arcing and Operational Pressure

As the plasma is formed via an electrical discharge between electrodes, the challenge is avoiding secondary discharges in other regions of the chamber. The breakdown voltage of a gas, V_{BR} , is well described by its Paschen curve, which shows how V_{BR} varies as a function of pressure, p , and distance, L . Generally, Paschen curves are plotted on a log-log plot, and consist of a linear region at high pL , a minimum V_{BR} typically found at a pL of 1–10 Torr cm and a steep region where low pL results in very high breakdown voltages as shown in Fig. 5.1. This behaviour can be explained by considering the avalanche breakdown that is required for a discharge plasma to form. Also known as a Townsend discharge [231], free electrons in the gas accelerated by the electric field collide with other gas molecules resulting in the production of more free electrons. This results in an avalanche multiplication leading to more and more free electrons and ions, ultimately resulting in a glow discharge. No discharge will occur if the electric field is insufficient to cause the ionisation of other atoms from these collisions. For pL greater than the Paschen minimum, increasing pL results in a reduction in electron's mean free path - making it small in comparison to the electrode gap increasing the number of collisions it experiences. Each collision will randomise the direction of travel and result in an energy loss of the electron, making it less likely it will have sufficient energy to ionise other atoms and reducing the likelihood of

the avalanche. Therefore, a smaller pL results in a lower V_{BR} . This decrease is only true up to the Paschen minimum, below which the electron's mean free path is long in comparison to the electrode gap and is unlikely undergo a sufficient number of collisions to give the avalanche, resulting in a sharp increase in V_{BR} at low pL .

The pressure used in the PDR is dictated by both the pressure required for diamond growth and the pressure which results in an achievable V_{BR} across a gap of order a few centimetres. Typically, around 10 Torr would be used with an electrode gap of 5 cm, a striking voltage of 15 kV is more than sufficient. In order to prevent a plasma discharge to the grounded chamber, a quartz tube is used to insulate the high voltage electrode to try and ensure only discharge occurs to the grounded electrode (the anode). At these pressures and distances, the pL value is above the Paschen minimum and at the bottom of the linear region. This ensures a plasma discharge is achievable whilst remaining predictable as the discharge should occur at the shortest distance to ground. At low pL , this is not the case, and plasma discharges will preferentially occur at longer distances making it very challenging to predict where it will happen.

The low operational pressure of the microwave plasma in DELPHI not only increases ion current, it also effectively eliminates concerns over arcing. Although DELPHI had no high voltage electrodes, there is still a significant potential difference present between screen and extraction grids in the ion optics (used for ion extraction) across a small gap of approximately 1 kV across a 1 mm gap. However, concerns of arcing are negligible, as the V_{BR} value would be huge for any reasonable distance at this low pressure. Under similar conditions, only low pressures could be used to avoid arcing between the grids. Low pressures are also preferable in the new setup in order to maximise ion current incident on the sample. However, as gas ionisation between the electrodes is still required to strike the plasma, a pL must be selected at which this is still feasible.

5.2.1.3 Ion Optics vs biased stage

Ion optics, as used in DELPHI, are a series of charged grids used to extract ions from a plasma with energy selection (see 2.2) and are used for a variety of applications from synchrotrons to thrusters for small space craft. From the outset, design and successful implementation of the ion optics was known to be a significant challenge. In part, this was due to the large number of parameters that would need to be determined to ensure the ion optics performed as desired. Distances from the plasma to the ion optics and from the ion optics to the sample would need to be decided, as well as grid separation, aperture size and operating voltage for both screening and extraction grids. Careful selection of all of these parameters is required to maximise ion current, have effective energy selection of ions, and avoid arcing between grids, from the electrodes to the grids or from the grids to the chamber. Optimisation of ion optics is commonly done with the aid of simulation packages such as IBSimu [232] or SIMION [233]. The former was used to gain an understanding of how varying different parameters impacted the resulting ion beam (an example




Figure 5.2: An example of an output from an IBSimu [232] simulation of an aperture which could be used in ion optics. Blue blocks represent the grids used in ion optics, cylindrical symmetry conditions have been applied, meaning these blocks effectively become the material surrounding an aperture. For this simulation, the screen grid (far left) has a voltage of 620 V, the extraction grid (middle) is at -410 V and the deceleration grid (right) is grounded. Equipotential lines (green) are shown, and colour bar gives plasma density.

plot can be seen in Fig. 5.2). Although this was helpful for developing understanding of ion optics, it was very challenging to find a starting point from which to work from with so many potential degrees of freedom.

Ion optics are typically used when the form of the ion beam (divergence, ion energy etc.) is of particular importance, or if total separation of plasma and ion is required. As this setup was simply designed to perform low energy implantation of deuterium ions, no such limitations are in place, and the aim was to reach a moderate fluence of deuterium at a low energy. Furthermore, at the moderate pressures used in this setup, the large potential differences and small separation of the grids raise significant concerns of arcing.

Instead, a far simpler approach was used in the final design. Applying a negative bias to sample stage would extract the positive deuterium ion clusters from the plasma at an energy dictated by the magnitude of the bias. Additionally, this design removed any major concerns over arcing and was far more feasible in the time frame available. A grounded grid above the sample could also be included to act as a reference potential (improving the energy selection) with ion energy given by the potential difference between the grounded grid and the sample stage. Use of a biased stage removes many of the optimisation and arcing concerns surrounding the use of ion optics, and means only the stage bias, stage height and gas pressure need to be considered. The stage bias will largely be dictated by target ion energy, whilst stage height and pressure can be adjusted to maximise ion current whilst avoiding arcing and excessive heating. Other glow discharge setups have also used a biased stage for ion extraction [222].

As described, the biased sample stage behaves similarly to a Langmuir probe - a plasma diagnostics instrument in which a biased metal probe is inserted into the plasma and the current response is used to determine electron temperature and other plasma characteristics. The polarity of the applied bias influences the species collected, with a positive bias increasing incident electrons and a negative bias increasing incident ions. The current required to maintain a bias on the stage can be measured with an ammeter in series between the stage and DC voltage

supply. With a negative bias, positive ions will be preferentially attracted to the stage resulting in a build-up of positive charge. In order to maintain the negative setpoint bias, a negative current is required from the power supply. Therefore, a negative current reading is a measurement of ion current incident on the stage and can be used to estimate the ion fluence on the sample.

Arcing between the powered cathode and the sample stage is unlikely as both are negatively charged, and the distance is greater than the electrode separation. Again, this is in contrast to an ion optics setup in which the grid closest to the electrode is the positively biased screen grid, greatly increasing risk of electron flow from the electrode to the grid. Although the inclusion of a grounded grid improves ion energy selection, the additional height and lack of negative bias would also increase the likelihood of electron flow and arcing. Decreasing the height of the sample stage could overcome this but at the cost of ion current, meaning the potential benefits of the grid could be outweighed by the disadvantages.

5.2.1.4 Gas Flow

DELPHI had two gas flow modes: once through and recycling. In the former, gas is simply passed through and pumped out, whilst the latter the gas is recirculated and has impurities removed - this is required for operation with tritium. As tritium is never to be used in this setup, there is no need to recycle the gas and flowing gas once through the system is the simpler choice. A once through method must match the gas flow in and out to maintain a constant pressure, usually with a mass-flow controller (MFC) and needle valve as is used in the hot filament CVD reactor. The main advantage of this method is consistent operating conditions and removal of contaminants. The pressure will remain constant throughout the exposure with no drop off with time. A static flow or fixed volume method can also be used for a once through system as is commonly done to reduce wastage of high value gases. In this simple method, gas fills the chamber to the required pressure and is sealed off, resulting in a fixed amount of gas within the reactor. Using a static flow will reduce gas wastage but could result in a build up of contaminants if significant sputtering/etching of the sample or components occurs for sustained runs. Due to the large volume of the chamber and expected low temperatures, no significant drop off in pressure would be expected in this setup, so it was designed to operate as a fixed volume for simplicity and to reduce deuterium wastage. If required, this could be updated at a later date with the inclusion of a MFC and a needle valve.

5.3 Design

The main chamber in the final design consisted of a large ConFlat DN160CF (8") tube with two DN40CF (2.75") ports and one DN63CF (4.5") as seen in Fig. 5.3. The 8" top port is used for a large Kodial window to ensure plasma sample alignment and monitor the plasma. The 4.5" port and the 2.75" port aligned opposite to one another were used to mount grounded and powered




Figure 5.3: Diagram of ExTEnd with key parts indicated, as viewed from the back of the implanter. A Penning gauge and Baratron were used for low pressure and operating pressure measurements respectively. The outlet valve separates the main chamber from turbo and scroll vacuum pumps. Additional valves and gauges not shown were present to control the flow of gas into the chamber. The high voltage supply for the powered electrode, was bolted directly onto an external face (labelled 'High Voltage Mount Point'), hence a PTFE cover was required for safety. The sample stage is biased, (supplied through the 'Sample Stage Feedthrough'). Sample stage is shown with grounded grid in place above the sample, this can be removed.

electrodes respectively. The port for the powered electrode used a 30 kV ceramic break to isolate the high voltage flange from the grounded chamber. The powered electrode mount is attached to a 2.75" blanking flange. A high voltage is applied to the external face of this blanking flange making this end live, as such, a large PTFE cover is used to cover the blanking flange and the ceramic break. The grounded electrode goes through the gas inlet flange, which allows gases into the chamber for filling and venting. A three-way Swagelok valve (not pictured) was mounted onto the gas inlet tube and ensures there is no direct line from the gas cylinder to atmosphere.

Three pressure gauges were used in conjunction to accurately measure different pressure regimes. During operation, a Baratron gauge accurate from 0.1 – 100 Torr was used to monitor chamber pressure (referred to as the ‘operating pressure gauge’ in Fig. 5.3). In these gauges, a metal diaphragm is adjacent to a fixed dual electrode [234]. The diaphragm deforms from changes in pressure, resulting in a change in capacitance between this and the electrode which can be measured to determine the pressure. The major advantage of this style of pressure gauge is its accuracy is independent of the type of gas being measured. This is in contrast to Pirani gauges, which rely on the variation thermal conductance of the gas at different pressures [235]. As such, scaling factors must be applied to ensure accuracy for different gases. A Pirani gauge was used downstream of the turbo to monitor the backing pump as, typically, this will be used to measure air when pumping down. Finally, a Penning gauge was used to measure high vacuum (down to 10^{-8} Torr) to ensure adequate vacuum has been reached between runs. These gauges rely on ionisation of gas molecules to induce a current, the measurement of which can be used to infer the pressure [235]. This was attached via a 2.75" cross to the last 2.75" port (‘low pressure gauge’ in Fig. 5.3), situated perpendicular to the electrode ports. The other two ports on this cross are connected to a manual inline valve (going to the vacuum pumps) and a secondary viewport.

The final 8" port at the bottom is used to mount the sample stage. Here, an 8" to 2.75" adaptor was attached to a 2.75" tee which the sample stage assembly protrudes through and into the main chamber. A feedthrough was mounted to the side port of the tee and was used for current measurement and biasing of the sample stage. The sample stage can be accessed by removing the tee.

Tabs welded on the bottom 8" flange were used to mount the chamber to a trolley. The height of the trolley is approximately 120 cm, making it easy to look through the main viewport at the top whilst still possible to remove the bottom tee for sample mounting and stage adjustment. Both scroll and turbo vacuum pumps were mounted to the trolley below the implanter. The former was simply mounted directly onto aluminium struts at the bottom of the trolley via rubber mounts in an attempt to minimise vibrations. The turbo vacuum pump was mounted just above this on an aluminium plate cut to size with drill holes to allow for direct mounting.




Figure 5.4: Cross section of mounted electrode assemblies.




Figure 5.5: **A** - Components for electrode mounting, including the two tungsten electrodes (i), collets and mounting screws (ii), the powered electrode mount (iii) and the grounded electrode mount (iv). **B** - Tungsten electrode within collet. **C** - Electrode and collet secured into the grounded electrode mount using the access hole and 5 mm hex key.

5.3.1 Electrodes

Electrodes consisted of two $\varnothing 10 \times 72$ mm tungsten rods. Tungsten is commonly used as an electrode material due to its strength and resilience to sputtering. Electrode mounts were machined from a single piece of stainless steel (see Fig. 5.4) and consisted of a mounting plate at the base and a hollow tube with a lip near the top. Above the lip is an 18 mm inset, the diameter of which was a slip fit for the collet used to affix the electrode to the mount (see Fig. 5.5). Tightening the screw into the collet causes it to compress around the electrode offering a secure fit, whilst maintaining parallel electrode faces and avoiding the need alter the brittle tungsten. For this to work, the internal surface at the tip of the electrode mount was angled to 14° to match the angle of the collet. The length of the electrode mounts was selected to give a 25 mm electrode




Figure 5.6: **A** - Powered electrode assembly. **B** - Grounded electrode assembly (i) with enlarged image showing the spacers used to maintain gas flow (ii).




Figure 5.7: **A** - The open end of the PTFE cover with securing bolts which overlap with the flange to prevent the cover from being removed unintentionally. **B** - The closed end of the cover with cable gland for high voltage supply to the electrode. The cable gland pictured was later replaced with a larger, M25, gland.

gap. Spacers can also be used to decrease this gap to 15 mm if desired.

The hollow tubes of the electrode mounts have various venting holes to promote gas flow during pump down. The high voltage electrode was mounted directly onto a tapped, 2.75" blanking flange with six vented M4 bolts (shown in Fig. 5.6A). As seen in Fig. 5.6B, the grounded electrode used spacers to add separation from the base plate and the flange. This, as well as additional venting holes in the base plate, acts to promote gas flow particularly around the pressure gauge and gas inlet. Again, this was mounted with six M4 bolts threaded into the 4.5" to 2.75" adaptor flange.

As the blanking flange used to mount the powered electrode would be live during operation, a cover is required for safety. For this, a $\varnothing 100$ mm PTFE tube with a $\varnothing 72$ mm cavity was used to cover the 2.75" ($\varnothing 69.9$ mm) blanking flange and the ceramic break. Screws going through tapped holes in the cover protruded the far side of the grounded flange of the ceramic break and prevented it from coming off unintentionally (see Fig. 5.7A). A M25 tapped hole in the covered end allows a cable gland to be threaded straight into the cover (see Fig. 5.7B) and a secure connection for the high voltage power supply to the electrode.

5.3.2 Sample Stage

There were several key considerations surrounding the sample stage beyond offering a secure mounting for the sample. The separation between plasma and sample stage impacts incident ion current, the likelihood of a secondary plasma discharge forming on the stage, and the temperature of the sample. As such, the sample stage assembly was designed to allow the height to be adjusted. A bias was also required on the top of the sample stage to extract ions from the plasma, as well as a way to measure current for fluence estimation. Therefore, the top of the sample stage needed to be electrically isolated from the sample stage mount. Lastly, the option for the inclusion of a grounded grid was incorporated into the design for better ion energy selectivity. All of these points were addressed with the following design.

The different components of the sample stage can be seen in Fig. 5.8 and Fig. 5.9. The baseplate was machined from one piece of stainless steel and consisted of a $\varnothing 28 \times 5$ mm disc with a 100 mm M10 threaded rod protruding out of the bottom. This threads into the sample stage mount, which takes a similar form to the high voltage electrode mount - a long stainless steel tube mounted to a 2.75" blanking flange by six M4 bolts. Again, venting holes were added along the length of the tube to avoid restricted gas flow. The top 50 mm of the internal of this tube is tapped. To adjust the height, the sample stage can be screwed in or out and secured in place with the locking nut. This design was simple and effective, but meant it was not possible to adjust the height during operation or under vacuum, and the sample stage must be removed to do this. This compromise was deemed acceptable as height adjustment was not expected to be required once standard operating conditions had been established.

Four M4 clearance holes were added to the base plate to allow the mounting of the insulating




Figure 5.8: Cross section of sample stage assembly.

block. The block was made from MACOR - a machinable ceramic material, which isolates the biased sample stage from the grounded base plate. Tapped M4 holes in the base of the block allow it to be mounted to the base plate with ceramic bolts¹. The block is $\varnothing 30 \times 20$ mm, this thickness alongside the slight overhang compared to the base plate and sample stage was designed to prevent arcing. Additionally, a $\varnothing 28 \times 2.5$ mm recessed region was made for the sample stage to sit in, again reducing the likelihood of arcing from the edge of the biased stage.

The stage itself consisted of a sample plate and window plate (see Fig. 5.10ii and iii/iv respectively), both of $\varnothing 27$ mm and a combined thickness that ensured a good fit in the recessed region of the block. The sample plate has a thickness of 2 mm, with a square 11×11 mm², 1 mm recessed region in the centre to accommodate 10×10 mm² samples. The window plate is used to secure the sample to the stage and give a defined implantation area. In the centre there is a

¹The design was initially for these to be vented stainless steel bolts, but during machining of the block, these holes met with the holes to secure the window plate above (see Fig. 5.8). Therefore, ceramic bolts were required but vented was not, as the bolts from above were vented stainless steel.




Figure 5.9: Assembly of the sample stage. **A** - Sample stage mount secured onto 2.75" blanking flange. **B** - Base plate with locking nut added. **C** - Insulating block, sample plate, and sample window added, full assembly without grounded grid. **D** - Full assembly including grounded grid.




Figure 5.10: Sample stage components. **i** - Insulating block, machined from MACOR. **ii** - Sample plate, 11×11 mm inset region for sample. Window plates for securing sample within sample plate, 8×8 mm² exposure area with (**iii**) and without (**iv**) additional holes for grounded grid mounting. **v** - Insulating rods for grounded grid. **vi** - Grounded grid plate.

square 8×8 mm window, ensuring a 10×10 mm² sample will always have the same exposure area even if there is some lateral movement in the inset region. A laser cutter was used to cut out window plates from a 0.2 mm thick tantalum foil and meant different window sizes could easily be made for different samples. The window plate was secured to the block via four vented M3 bolts. To load and unload samples, the 2.75" tee at the bottom of the setup is removed and the sample stage assembly is withdrawn from the base. The window plate can then be removed, the sample placed in the inset of the sample plate, and the window plate reattached to secure the sample.

The optional grounded grid consists of 19 $\varnothing 1.95$ mm holes separated by 2.5 mm in a $\varnothing 28 \times 0.2$ mm tantalum plate, with four M3 clearance holes around the edge. Again, this was made using a laser cutter. The grid also must be electrically isolated from the sample plate, which is achieved via four $\varnothing 6 \times 25$ mm MACOR rods. These go through the sample and window plates and slot into $\varnothing 6 \times 5$ mm recesses in the block. M3 \times 30 mm bolts go through holes in the grounded grid and the centre of these rods into tapped holes in the insulating block below the recess. Additional venting holes were added perpendicular to these as vented M3 bolts of this length are not readily available.

5.3.3 Gas Supply

For retention experiments, ExTEnd is run with deuterium, as this can be distinguished from hydrogen impurities common within diamond. To supply the deuterium, a small, double ended, 500 ml cylinder was mounted next to the implanter and refilled as required. One end was connected to the chamber, whilst the other is used for refilling. Despite the modest volume, this cylinder was more than sufficient to perform repeated runs without the need to refill, as the canister can be filled up to 1-3 bar and operating pressure is roughly three orders of magnitude lower.

Initial tests were performed with hydrogen gas as this is more readily available, this was also supplied via the 500 ml cannister. The cannister has its own pressure gauge to give an indication of the amount of gas remaining. A series of valves separate the cylinder and the chamber. This allowed incremental filling of sections of tubing to add small amounts to the chamber at a time due to the lack of accurate flow control.

5.4 Assembly

5.4.1 Electrode Alignment

Upon putting together the electrode assembly and mounting them onto the main chamber, significant misalignment between the two electrodes was observed as shown in Fig. 5.11A. It was not the electrode mount or the ceramic break causing this off angle, but the port of the chamber itself. As with some of the other components used in this build, the main chamber was




Figure 5.11: **A** - Electrode misalignment due to the port used for the powered electrode (right) being off angle, view from top (i) and side (ii). **B** - Realigned electrodes from top (i) and side (ii). **C** - Shims used on the base of the powered electrode mount to correct port misalignment.

not new, and its history and previous uses were largely unknown². On closer inspection, it seems ports had been added to the chamber, as there was evidence of a sketched-out position for an additional port on one of the sides. The port used for the powered electrode was, deliberately or otherwise, at a slight off angle of approximately 1° . Such a discrepancy was challenging to see with the naked eye but, due to the length of the powered electrode, became apparent on assembly with a misalignment of approximately 3–4 mm. This degree of misalignment was unacceptable for the intended use. Misaligned electrodes would result in the shortest distance between the electrodes being at an edge, resulting in a smaller plasma and increasing the risk of sputtering of the electrode. To correct for this, shims were placed between the blanking flange and the electrode mount. Shimming material of different thicknesses were tested to determine

²Although some components were not new, all were in good condition and were sandblasted and chemically cleaned before use to minimise potential contaminants.




Figure 5.12: Piping and instrument diagram of the implanter setup. Valves are labelled 'V1' to 'V8', pressure gauges 'G1' to 'G4', and vacuum pumps 'P1' and 'P2'. V7 is a three-way valve which has two positions 'A' and 'B', gas flow of these positions are indicated below the V7 label. P1 was a turbo pump and P2 a dry scroll pump. G1 - Analogue gauge to measure gas cylinder pressure, G2 - Baratron (1 – 100 Torr) to monitor pressure during operation, G3 - Penning gauge for low pressure readings when pumping down, G4 - Pirani gauge to monitor scroll pump when pumping air.

the desired thickness of 25 μm , which was in agreement with estimations made beforehand. A more permanent solution (as shown in Fig. 5.11C) was made using the shim material and adding a bolt hole to ensure the shim remained in place. The realigned electrode can be seen in Fig. 5.11B.

5.4.2 Initial Assembly and Pump Down

Once fully assembled, an initial pump down was carried out. Initially, the gas cylinder was disconnected and only the main chamber was evacuated. Left overnight with the turbo running, a pressure of order 10^{-7} mbar was achieved, indicating no significant leaks were present and the pumps were operating as expected.

The system was then vented and the gas supply assembly connected. The complete system was then pumped down to 10^{-2} mbar using the scroll pump. All valves between the main chamber and gas cylinder were then closed. The turbo was turned on, and once a sufficiently low pressure (10^{-6} mbar) was reached, valves were opened one at a time to incrementally reintroduced sections and check for leaks. Due to the small diameters of some of the gas connections this took longer to pump down, but over a weekend achieved a base pressure of 10^{-8} mbar. Although the operating pressure is a lot higher than this, being able to reach this low base pressure indicates a relatively clean and well assembled setup devoid of leaks.

5.4.3 Hydrogen Fill Test

Prior to the first plasma strike, a fill test was performed to find the best method to reach target pressure within the chamber. Fig. 5.12 shows a simplified piping and instrument diagram of the setup and will be referred to throughout this chapter. Firstly, V1 was attached to an existing hydrogen gas line and the vent on the scroll pump was attached to the ventilation within the lab. Opening the valves to and from the gas cylinder (V1 – V6 open, V7 position A) the setup and gas line was pumped down. Due to the long length of the hydrogen gas line and the minimal conductance of the small diameter tube, after two hours with the turbo pump on, a pressure of 10^{-3} mbar was achieved. Closing the valve between the gas cylinder and hydrogen line (V1) saw the pressure rapidly drop, indicating there was no leak in the main system.

Once pumped down, the gas cylinder was closed off from the main chamber by closing V2 – V4. The gas line and cylinder were then filled with hydrogen gas to a pressure of 1.3 bar (the pressure of the gas line). V1 was then closed and hydrogen that remained in the gas line was pumped out using a different pump attached to another part of the gas line (in this case, the undoped microwave CVD reactor). V5 was then closed.

Due to the lack of MFCs or needle valves, care is required when filling the chamber to desired pressure. The series of valves between the cylinder and the chamber (V2-V4) allow for small volumes to be filled one at a time in an attempt to make this process easier. For example, with V3 and V4 closed, V2 can be opened, allowing hydrogen at close to 1.3 bar to fill the piping between V2 and V3. V2 can then be closed, opening and closing V3 then fills a small volume of tubing between V3 and V4. Opening V4 and allowing this gas volume into the chamber corresponded to an increase in chamber pressure of around 0.5 Torr allowing for incremental increases. Alternatively, a more gradual increase is possible by partially opening V4.

The chamber was filled to approximately 2 Torr as, at the time, this was thought to be around the expected operating pressure. With V3, V4, V5 and V8 closed, the chamber was then left as a leak test. Over three weeks, the chamber showed no measurable change in pressure indicating a sufficiently good seal had been achieved.

5.4.4 Power Supply and Interlocks

As the power supply for this setup need to be shared with the PDR, a solution was required to allow the connections to be swapped easily and safely. To do this, a metal box housed a socket to which the output of the power supply was connected. High voltage cable could then be directly connected to the powered electrode of both ExTEnD and the PDR. A plug on the other end of these cables could then be connect to the power supply when in use. Fig. 5.13 shows the cable with plug connected to the setup.

Various interlocks were required to minimise the risks surrounding the high voltage power supply. The first interlock was on the door of the cabinet that housed the power supply, which turns off the power supply if the door is open. As the socket box is to be mounted within the




Figure 5.13: High voltage power cable connected to the high voltage electrode through a cable gland in the PTFE cover. Blue wire is grounded and connected directly to the body of the chamber. The other end of the wire has the high current high voltage MIL-spec plug used to connect to the PDR power supply.

cabinet and the power supply cannot be on whilst the door is open, it is not possible to tamper with the plugs or socket with the power supply on.

Another consideration was the grounding of the setup. As there are multiple power supplies and pieces of equipment connected to the mains, there is the potential for numerous different grounds. This is in part a safety concern, as considerable potential differences can be present between two grounded items but also could impact ion current measurements. As such, a common ground was implemented between all equipment across this setup and the power supply.

5.4.5 Sample Stage Mounting

The sample stage assembly was mounted through the straight section of a 2.75" tee as shown in Fig. 5.14. A high-voltage vacuum feedthrough was mounted to the side port of the tee and connected to the top of the sample stage via a Kapton coated wire. Additional Kapton tape was also used around the crimp mount on the wire to prevent any possible arcing. Excess cable was put within the tee to allow for removal of the sample stage as well as height adjustment. For initial measurements, the height was adjusted to be around 4 cm below the centreline of the electrodes and mounted to the chamber (3.5 cm shortest distance). To apply a stage bias and take



Figure 5.14: Sample stage assembly during mounting to 2.75 " tee. Kapton coated wire connected to the top of the sample stage supplied bias and measure stage current. Wire was connected to a vacuum feedthrough mounted to the port on the left of the tee.

stage current readings, a 250 V DC voltage supply was connected to the feedthrough, with an ammeter in series between the power supply and stage.

5.4.6 Quartz Tube

The $\varnothing 30$ mm quartz tube that surrounded the powered electrode was a snug fit for the $\varnothing 31.75$ mm internal of the ceramic break. It was hoped the tight fit would be sufficient to keep the tube in place, and no additional mounting would be required. However, due to the length of the quartz tube, the tube tilted downwards a very small amount towards the inside of the chamber. Although barely visible, the vibration of the scroll pump resulted in the gradual shifting of the tube towards the centre of the chamber. This was a very slow process, but after a week the tube was obscuring the end of the mounted electrode. To avoid this problem, a $\varnothing 33$ mm flare was added to one end of the tube (see Fig. 5.15). At this diameter, the end of the tube becomes caught on the ceramic - preventing it from sliding forward.




Figure 5.15: Flared end of the quartz tube, required to prevent tube gradually sliding out of position.

Table 5.1: The range of variables used in testing of the new implanter. Combinations of the conditions below were selected to explore trends.

Testing Parameters	
Pressure (Torr)	0.1, 0.2, 0.5, 1, 1.5, 2, 3, 5, 10, 20
Electrode Power (W)	40–80, 10 W increments
Stage Bias (-V)	0–800, 10–50 V increments
Stage-Electrode Separation (mm)	25, 35, 50, 65
Pulse t_{off} (μ s)	0.1–0.45, 0.5 increments

5.5 Testing

A large variety of different conditions were explored in order to improve the general understanding of the setup and optimise conditions for implantation. The range of parameters tested can be seen in Table 5.1. These included, stage bias, stage height, chamber pressure, electrode power and pulse off time. This section presents a summary of conclusions from across the testing phase.

After initial tests of the power supply the sample stage was fitted, and preliminary current measurements were made with no stage bias by connecting the feedthrough directly to an ammeter (a Keithley 2000 series multimeter was used for stage current measurements). Following this, a 250 V DC supply was connected, with the ammeter in series between the negative terminal of the DC supply and the sample stage. After initial tests up to a -250 V stage bias, this supply was replaced with a 3 kV DC supply in order to extend IV curves and reach the standard operating anode voltage of DELPHI (-400 V).

Stage current was measured for each of the conditions tested. Initially, the stage current was simply read from the ammeter display after the current reading had stabilised (10–30 s). When the 3 kV supply was installed, a laptop was also connected to the ammeter to record current measurements over time. This was done with Python script utilising the PyVISA package [236], which collected current measurements every second. As discussed, it was hoped the sample stage




Figure 5.16: Full assembly as used for first plasma. Valve and gauge labels (V1-V8 and G2-G4 respectively) refer to Fig. 5.12. Sample stage assembly had not yet been fitted.

would behave similarly to a Langmuir probe, where the current required to maintain a setpoint bias indicates the ion current incident on the sample stage. With no bias on the stage, the current reading gives the balance between electrons and ions hitting the stage, with a positive current reading indicating a greater number of electrons. Under certain conditions, the stage can act like the grounded electrode, and electrons can stream from the negatively biased powered electrode directly to the stage resulting in high positive current readings. This is unlikely to occur with a moderate negative stage bias, as the potential difference will be greater between the electrodes than the powered electrode to the stage. Conversely, with a large stage bias that exceeds the stage-electrode breakdown voltage, electrons are emitted from the stage to the grounded electrode resulting in a large negative current measurement and the formation of a plasma discharge. Positive current measurements are a result of incident electrons, either from the plasma or from the powered electrode, whereas negative current readings are a result of incident positive ions or electrons emitted from the stage. For accurate fluence estimates, it is important that arcing and stage discharges are avoided so the current measurement can be related to the number of incident ions.

5.5.1 Plasma Form

An early test plasma in ExTEnD can be seen in Fig. 5.17. At 0.1 Torr, a hemispherical plasma on the end of the powered electrode can be seen, as well as a secondary discharge along the internal of the 2.75" port that connects the main chamber to the Penning gauge and vacuum pumps. When increasing the pressure, the plasma is seen to increase in density, and condense around the electrodes, with the plasma sheath extending down the length of the powered cathode. This can be seen in Fig. 5.18, which presents hydrogen plasmas at various pressures. As the pressure increases, the secondary plasma fades away until it is no longer seen for pressures above 1 Torr (Fig. 5.18A–Gii). The faint circular plasma visible in the centre of Fig. 5.18A–Eiii is the same plasma observed in Fig. 5.18A–Eii but viewed along the port. Striations are ionisation standing waves that are commonly observed in plasmas [237, 238] and can be seen in the quartz tube at pressures of 1 Torr and above. With increasing pressure, the number of striations decreases, and they appeared to move along the quartz tube away from the centre of the chamber. In Fig. 5.18Fi, plasma discharge on the sample stage is visible, indicating a flow of electrons from the powered electrode and was observed for pressures of 2 Torr and above with a 0 V bias. Applying a negative bias to the stage prior to striking the plasma decreased the potential difference between the cathode and the stage to below the breakdown voltage and prevented the discharge. This was done for higher pressures (Fig. 5.18Gi).

5.5.2 Power Supply Variables

The electrode power supply was run in constant power mode. In this setting, a voltage above the breakdown voltage is used and the supplied current is varied to meet the setpoint power. Fig.




Figure 5.17: Image of an early hydrogen test plasma in ExTEnD, the pink colour is typical of a low pressure hydrogen plasma and is a result of the red and blue emissions of the Balmer series. Pressure: 0.1 Torr, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Electrode-stage separation: 35 mm, Stage bias: 0 V.

5.19 shows the electrode voltage with increasing pressure. The voltage reading is time averaged across the pulses. Therefore, the pulse voltage, V_P , is given by

$$(5.2) \quad V_P = \frac{V_{avg}}{1 - f t_{off}},$$

where V_{avg} is the voltage reading, f is the pulse frequency and t_{off} is the off time of the pulse. Standard operation used a f of 100 kHz and a t_{off} of 3 μ s, giving a pulsed voltage 1.43 times greater than the voltage reading and the values on Fig. 5.19. With increasing pressure, it can be seen the electrode voltage decreases sharply before displaying a more gradual increase. This behaviour suggests the implanter is operating at pL values around the Paschen minimum. Despite the minimum in breakdown voltage being at a pL of 1 Torr cm [225], the minimum electrode voltage corresponds to a pL of 10 Torr cm (5 Torr, 2.5 cm separation). Discrepancies




Figure 5.18: ExTEnD hydrogen plasma at 0.1, 0.5, 0.75, 1, 1.5, 2, 3 Torr (A–G respectively). **i** - Primary discharge between electrodes, from the main viewport. **ii** - Secondary discharge within chamber, from main viewport. **iii** - View from secondary viewport. Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Electrode-stage separation: 35 mm, Stage bias: 0 V for A–F, -50 V for G.




Figure 5.19: Time averaged DC voltage supplied to the powered electrode to maintain hydrogen plasma at different pressures. Multiplying time averaged voltages by 1.43 gives the pulsed voltage. Electrode separation: 25 mm, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s.

between these values are likely a result of differences in experimental setup. For example, the setup in [225] was designed to determine the minimum breakdown voltage for a given pL value. This is in contrast to the electrode power supply, which operates above the breakdown voltage to ensure consistent striking. Furthermore, it is unclear the impact the pulsed power supply would have. Regardless of these differences, a minimum in electrode voltage is seen at a pL value close to the expected minimum in breakdown voltage.

Alongside electrode voltage measurements, electrode current was also recorded to calculate power. It was noted there is significant variation between setpoint power and calculated power, with calculated power being within ± 5 W of the setpoint power. If the power was varying by this amount, it could result in inconsistent operating conditions. Conversely, precision of the current reading was low (0.2 \AA increments), so it is possible the power is actually closer to the setpoint value, and the current reading is not of sufficient precision. Although consistent conditions for a given power and pressure would be preferable, measuring stage current over the entire exposure still allows for an estimate of total fluence, regardless of variations.

When measuring stage current, it was observed increasing power resulted in greater stage current as shown in Fig. 5.20, as well as a visibly brighter plasma. For a pressure of 1 Torr, a plateau can be seen from 70 W and above for the three stage biases tested, this is comparable to ion flux measurements made with a Langmuir probe in other setups [239]. Current values scale




Figure 5.20: Stage current for varying plasma power. The top plot shows various bias voltages for a fixed pressure of 1 Torr, the bottom plot shows various pressures for a fixed bias of -250 V. Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Electrode-stage separation: 35 mm.

with increasing stage bias. With a fixed bias (-250 V), the power curves vary form and magnitude with different pressures, with lower pressures appearing more linear and showing limited evidence of a plateau. At 5.2 Torr, the pressure is too high for notable current measurements, with the power supplied to the plasma having no impact on the 0 mA current reading. A key concern with operating at a higher plasma power is excessive sample heating. Running at 50 or 60 W should allow for reasonable stage current whilst reducing sample heating.

Adjusting the amount of time the power supply is on for (by controlling the pulse timings) also impacts the stage current. Fig. 5.21 shows a linear relation between the time the power supply is on for and the stage current, suggesting t_{off} could be used to adjust stage current. Reducing t_{off} too much effectively makes the pulsed output continuous and removes the power supplies arcing suppression. It is unclear whether changes in pulse timings are truly reducing the ion flux or simply decreasing the time that the same high flux is being applied for (reducing the average flux). If the latter were true, the linear fit of Fig. 5.21 would be expected to pass through the origin (directly proportional), meaning halving the pulse length would results in half the stage current. However, the intercept of 0.13 mA means halving pulse length gives a stage current of lower than half the original value. For example, going from a time on of 80%




Figure 5.21: Stage current as a function of percentage of pulse time period for which the pulse was on. The pulse off time was varied between $1.5 - 4.5 \mu\text{s}$ in $0.5 \mu\text{s}$ increments, whilst the 100 kHz pulse frequency gave a time period of $10 \mu\text{s}$. A linear fit has been applied giving a gradient of $-0.072 \pm 0.002 \text{ mA}$, an intercept of $0.13 \pm 0.2 \text{ mA}$ and an adjusted R^2 of 0.98. Pressure: 0.1 Torr, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: $3 \mu\text{s}$, Electrode-stage separation: 65 mm, Stage Bias: -300 V.

to 60% is a 25% reduction in pulse length but results in a 30% reduction in stage current. This lack of direct proportionality might suggest that longer pulse times increases the number of ions available during periods with the pulse off, meaning increasing or decreasing pulse length could be used to impact ion flux.

In order to assess the consistency of the setup, stage current was measured with respect to time for longer periods. Fig. 5.22 shows data for three runs at fixed conditions (1 Torr, 50 W, -250 V bias). The first (black line) was run for approximately two hours before the plasma was turned off overnight. The following morning, the plasma was turned on with no changes made to the setup or conditions, this can be seen in the red line. Following this, different conditions were tested, in which the stage bias (blue line) and electrode power (green line) were varied. The plasma was then turned off again for around 2 hours. A final run was then performed (pink line) towards the end of which the plasma was turned on and off whilst still recording.

Across the three runs, the stage current gradually levelled off at -0.7 mA, after dropping from the larger current readings at the start of each run. A warmup period would be expected when turning on the plasma, in which the power supply and plasma become thermally stable. This could explain the initial decrease at the start of each run, during which the system is reaching thermal stability. However, this does not explain why each run gave different initial current

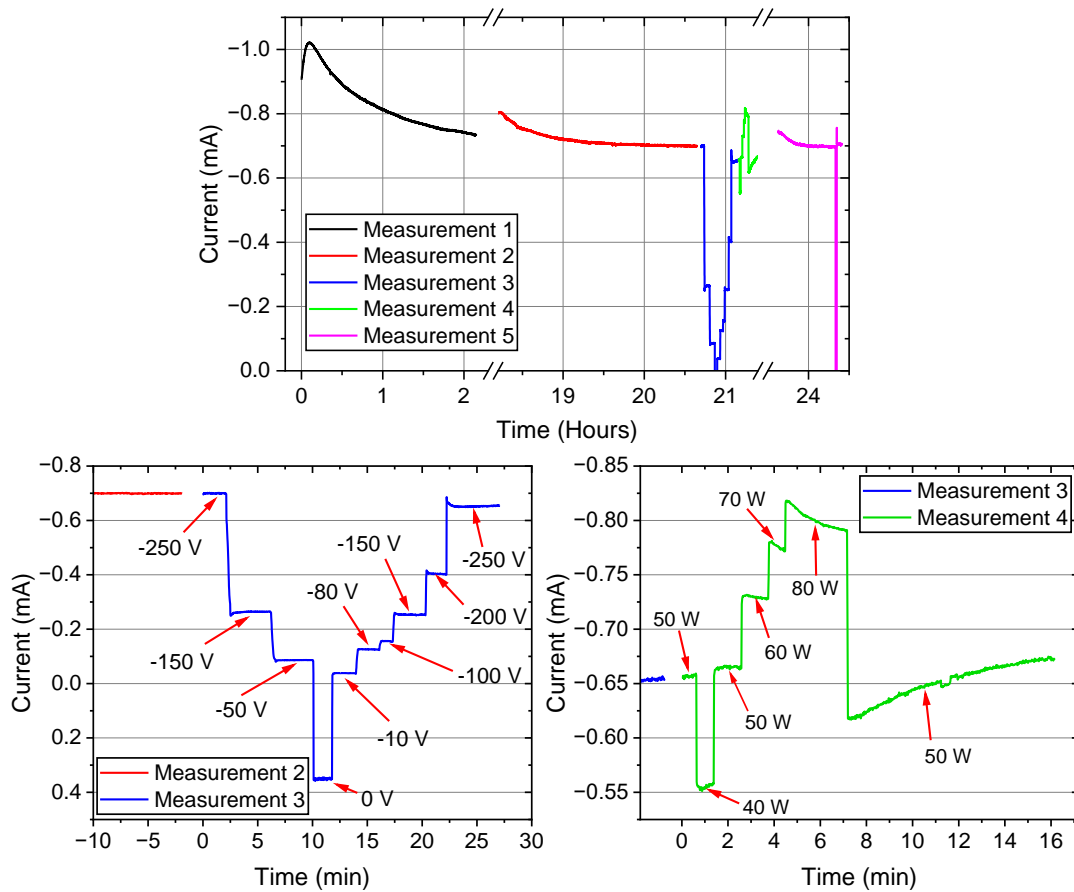


Figure 5.22: Stage current measured across a 25 hour period. Time of top plot is relative to the start of the first measurement, whereas bottom left and right plots are relative to measurements 3 and 4 respectively. Measurements 1, 2 and 5 were taken under equivalent conditions (50 W, -250 V). For measurement 3 bias was varied with a fixed 50 W power, whilst measurement 4 varied power for a fixed -250 V bias. Various gaps were left between measurements. Pressure: 1 Torr, Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Electrode-stage separation: 35 mm.

readings. 15 hours passed between the first and second runs, offering ample time for the setup to return to ambient temperature. If current variation was solely dictated by the temperature of the setup and power supply, the curve seen for the first run should be replicated the following day, which is not the case. This suggests there is another factor dictating the longer-range variation that settles over the course of 5-6 hours of the plasma being on. Although it is not clear why, it seems the longer-range variation must be a result of the same gas being used throughout runs. With new gas, more variation is observed, with a peak in stage current seen after around 10 minutes. The same peak is not present for subsequent runs with the same gas, and a quicker return to the stable current is observed. Sputtered positive metal ions would also be attracted to the sample stage and recorded as negative current, so it could be linked to the production of such ions. However, the sputtering rate of the tungsten electrodes and tantalum sample window are low, so ion contamination from these surfaces is not expected. Furthermore, these would act to increase detected current, rather than the reduction observed over long ranges here.

Varying bias resulted in effectively instant changes in current, which remained stable after the changes. Biases of -150 and -250 V were repeated, and second measurements were slightly smaller than the first measurements. This is likely a result of the gradual levelling off observed across the runs. Some variation in current was seen when changing power values, this is likely to be from the power supply adjusting to the new conditions. Increases in power resulted in a spike in current which gradually reduced, whereas decreases in power gave a gradually increasing current.

These current measurements highlighted the importance of measuring current throughout implantations. To calculate total charge for fluence estimation, it is not sufficient to use a single current value or an average, and integration of the current-time plots should be carried out. These results also highlight limitations of stage current measurements in the next section, which took an initial current reading. These values will not only be impacted by the long-range current variation, but the order in which variables were tested could also impact results (as increasing power resulted in an artificial increase in initial current). As such, initial current measurements should be used qualitatively - showing general trends rather than exact expected values.

5.5.3 Sample Stage Variables

As expected, pressure had a significant impact on stage current. No notable negative current (indicating either incident ions or stage electron emission) could be measured for pressures above 3 Torr for biases up to -250 V as shown in Fig. 5.23. As previously discussed, high pressures are likely to decrease ion current from a reduction in mean free path of the ions which are scattered or neutralised before reaching the stage. However, the highest stage currents were not measured at the lowest pressures but around 0.75 – 2 Torr. This could be explained by a Paschen curve and suggests that electron flow between the stage and the grounded electrode could be occurring. A 35 mm stage-electrode separation gives pL values in the range of 2.6 – 7.0 Torr cm. From this data




Figure 5.23: Stage current against pressure for varying stage biases. Values were collected for high pressures and 0 V, but many were positive and could not be shown on the log scale due to the change in polarity. Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Electrode-stage separation: 35 mm.




Figure 5.24: Stage current for increasing stage bias. Includes curves for various different pressures as indicated in the legend. Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Electrode-stage separation: 35 mm.

alone, it is challenging to conclude whether increases in stage current are a result of incident ions, or electrons emitted from the stage and accelerated to the grounded electrode. These pL values are close to the expected Paschen minimum (pL of 1 Torr cm) and the shape of the Paschen curve means an increase in stage-electrode current might be observed for intermediate pressures but not the highest and lowest pressures. Equally, the higher plasma density could contribute to increased incident ions, although generally it would be expected for extracted ions to drop off with pressure. Higher biases and different stage heights were required to try and distinguish these effects.

The impact of the bias voltage was explored further as seen in Fig. 5.24. At 0 V bias, positive stage currents (indicative of electrons incident on the sample stage) were seen. Electrons are either sourced from the plasma, or directly stream from the powered electrode to the stage. The latter scenario was clearly occurring for 0 V stage bias and pressures higher than 1.5 Torr, where a plasma discharge would be visible on the stage (as seen in Fig. 5.18Fi). This would give large positive current readings as the stage is acting as a drain for the electrons in a similar manner to the grounded electrode.

With a larger negative bias, the negative stage current increased, indicating a higher flux of positive ions incident on the stage, or a greater rate of electrons lost to the grounded electrode. The increase is non-linear, and the shape of the curve is linked to the pressure. At low pressures, a gradual increase can be seen. The curve becomes steeper with increasing pressure, and the bias required for the curve to begin moves to higher pressures. At 3 Torr, the stage current was very unstable and fluctuated between 0 and -2 mA in a step like manner. This sudden increase is more indicative of the stage bias exceeding the breakdown voltage than a sudden increase in ion extraction. No significant current was measured at 5.2 Torr.

Once the 3 kV supply had been installed, IV curves presented in Fig. 5.24 could be extended to higher biases. It was hoped that increasing the stage bias would eventually result in reaching the ion saturation current typical of a Langmuir probe. However, no evidence of a plateau was seen with increasing bias and the current measurements continued to increase. Furthermore, the formation of a plasma on top of the stage was observed under certain conditions (see Fig. 5.25). At low bias, turning off the electrode power supply resulted in a 0 mA stage current, as no ions or electrons are incident on the stage with no plasma present. However, this was no longer the case when a stage plasma was present - current measurements remained high, and the plasma remained present on the stage as shown in Fig. 5.25C.

The consequence of this behaviour on stage current is presented in Fig. 5.26. During desired operation (Fig. 5.26A), there is only electron flow between the electrodes, and a negative bias on the sample stage is used to extract positive ions from the plasma. In this case, the ion energy is dictated by the stage bias and the measured current corresponds to the flux of incident ions on the surface. However, when the bias applied to the stage exceeds the breakdown voltage for the conditions, electrons stream from the negative stage to the grounded electrode, creating a




Figure 5.25: **A** - Desired operation, plasma discharge between electrodes (0 V stage bias), measured current corresponds to incident ions. **B** - Plasma discharge between electrodes and from the stage to the grounded electrode (-700 V stage bias). **C** - Plasma discharge from stage to grounded electrode with electrode power supply off (-700 V stage bias). Pressure: 1 Torr, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Electrode-stage separation: 35 mm.

Townsend avalanche and ionising hydrogen to form the plasma on the stage (as shown in Fig. 5.26B). Here, the stage is acting in a similar manner to the negatively biased powered electrode. When this occurs, the current reading no longer corresponds to the flux of ions hitting the stage but is a combination of electrons leaving and ions incident on the stage, and both ion flux and energy becomes unclear. With the electrode power supply off and a sufficient stage bias, only the stage discharge is present, and current measurements correspond to electrons leaving the stage (Fig. 5.26C). As the powered electrode is either negatively biased or isolated, electrons will always preferentially flow to the grounded electrode. Evidence of this can be seen in Fig. 5.25C, where plasma can only be seen on the stage and around the underside of the grounded electrode.

When performing a bias ramp, the resulting *IV* curve is a combination of ion flux on the stage, and loss of electrons from the stage. In order to perform and measure ion extraction at a set energy, the stage-electrode discharge must be avoided, and conditions must be determined in which the electron flow between the stage and electrodes is negligible. The variables which influence the stage-electrode discharge the most are chamber pressure and stage-electrode separation. As




Figure 5.26: Simplified diagram of different sources of stage current, black arrows indicate the flow of electrons. **A** - Desired operation, discharge only between electrodes, measured current indicates incident ions. **B** - Discharge between electrodes and from the stage to the grounded electrode, measured current is the sum of ions incident on the surface and electrons leaving the stage. **C** - Discharge only present between grounded electrode and stage, measured current indicates electron emission from the stage.

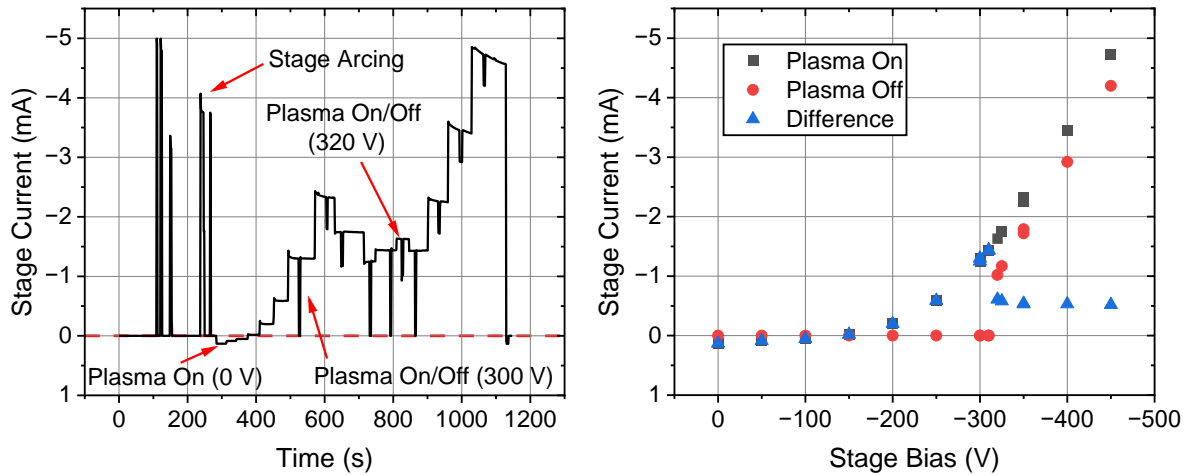


Figure 5.27: Current measurements taken to measure breakdown voltage between the stage and grounded electrode. The figure on the left shows current with respect to time during this test. Initial spikes were a result of electron flow from the stage to the grounded electrode with the electrode power supply off. Following this, discrete steps in current are a result of changes in stage bias. Sudden drops in current are when the electrode power supply was turned off and on again. The figure on the right gives the measured stage current with the electrode power supply on and off, as well as the difference between them, for various stage biases. Pressure: 1 Torr, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Electrode-stage separation: 65 mm.

such, bias ramps were performed at four different stage heights (corresponding to stage-electrode separation of 25, 35, 50 and 65 mm) at six pressures (0.1, 0.5, 1, 2, 5, 10 Torr). In an attempt to determine at what stage bias the stage-electrode discharge forms, the power supply to the electrodes was turned off and on at selected biases. With a negative stage bias and the electrode power supply off, any stage current must be a result of electrons flowing from the stage to the grounded electrode.

An example of these tests can be seen in Fig. 5.27. Current measurement with respect to time during the tests can be seen on the left, whilst the current measurement at each stage bias are presented on the right. The initial peaks on the current-time plot were a result the stage-electrode discharge with the electrode power supply off. This was done to get an estimate of the breakdown voltage prior to the bias ramp. The first off/on of the electrode power supply can be seen at around 500 s (-300 V bias). Here, the stage current returns to 0 mA when the plasma is off, indicating no electron flow is present with the plasma off. The following bias step (560 s, -350 V) shows a decrease to a non-zero current measurement, indicating electron flow from the stage to the grounded electrode with the plasma off. The lowest voltage that this was true for is -320 V. For voltages higher than this, the reduction in current is consistently around 0.5 mA, as shown in the difference between plasma on and off in Fig. 5.27.

As the stage current, I_{Stage} , is thought to be a combination of contributions from electron flow to the electrode and incident ions, it can be written

$$(5.3) \quad I_{Stage} = I_{Discharge} + I_{Ion},$$

where $I_{Discharge}$ is estimated by the current with the plasma off, $I_{PlasmaOff}$. Therefore, it would follow that ion current, I_{Ion} , can be approximated with

$$(5.4) \quad I_{Ion} = I_{Stage} - I_{PlasmaOff}.$$

However, there are clearly limitations to this simplification, as it relies on ion and discharge currents being independent of one another. Although it can be said with some certainty that the plasma off current is solely a result of electron flow to the grounded electrode, it cannot be said that this contribution remains the same with the plasma on. There is evidence of this at biases just below the stage-electrode breakdown voltage. At these biases for some conditions, the plasma off current was 0 mA but a stage discharge was present with the electrode plasma on. This indicates that, under these conditions, the current contribution from the discharge was not equivalent with the plasma on and off. Similarly, at this bias, the difference in current was greater than the constant value observed for biases beyond the breakdown voltage. These observations suggest the presence of free ions and electrons are enhancing the stage-electrode discharge, allowing it to occur despite the stage bias being below the breakdown voltage.

The constant value observed in the current difference of Fig. 5.27 could be an indication of an ion saturation current as observed in IV curves of Langmuir probes. However, this constant




Figure 5.28: Breakdown voltage between the sample stage and grounded electrode for varying pressure (indicated by different symbols) and stage-electrode separation (indicated by different colours). For $p \leq 2$ Torr, the breakdown voltage, V_{BR} , was taken to be the highest bias at which the stage current returned to 0 mA when the plasma was turned off. For $p \geq 5$ Torr, no negative current reading was made prior to arcing, and V_{BR} was taken to be the lowest voltage required to trip the power supply (indicating a stage current of > 10 mA). Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s.

value did not follow trends that would be expected for ion current measurements (decreasing with increasing pressure or separation). Furthermore, as the presence of the electrode plasma enhances the stage-electrode discharge, it is challenging to conclude with any confidence what this is a result of. Beyond the breakdown voltage the contribution from electron flow is effectively unknown, as it was concluded that discharge current is unlikely to be equivalent with the plasma on and off. Therefore, for effective ion extraction of known energy and measured fluence, a bias beyond the breakdown voltage cannot be used despite Eq.5.4 suggesting a correction might be possible. Instead, conditions must be selected that ensure the stage-electrode electron flow is not present with the plasma on, and $I_{Stage} = I_{Ion}$.

Fig. 5.28 shows the highest voltage at which the stage current returned to 0 mA for pressures up to and including 2 Torr. For pressures of 5 Torr or more, the bias supply would suddenly trip at high bias - indicating a stage current in excess of -10 mA and significant electron flow between the stage and grounded electrode. As no negative current readings were observed before this point, for $p \geq 5$ Torr, Fig. 5.28 presents the highest stage bias at which no trip occurred. Although not a typical measure of breakdown voltage, the shape of this curve bears a resemblance

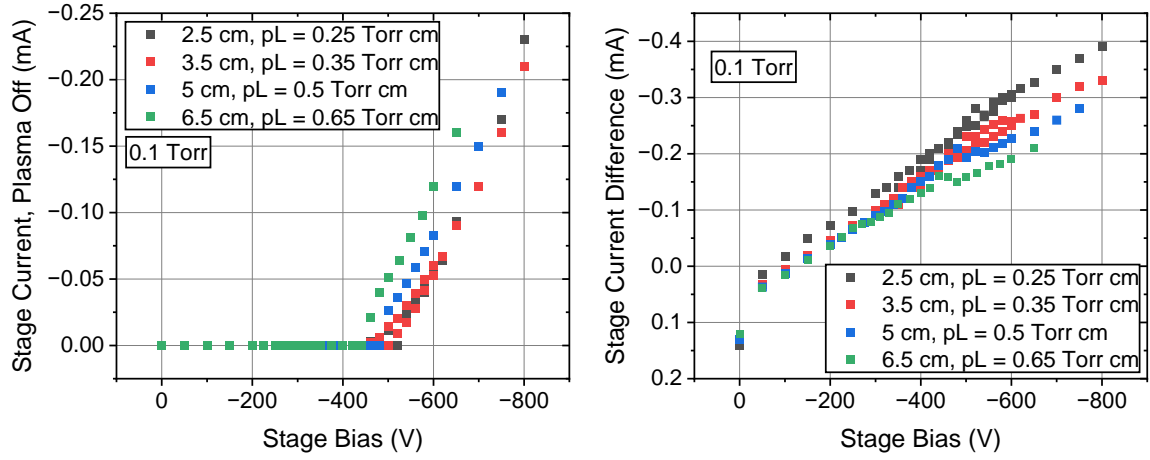


Figure 5.29: The difference between stage current with the stage current with the plasma off (left) and the electrode plasma on and off (right) and for varying stage biases. The stage-electrode separation and resulting pL value is indicated in the legend. Pressure: 0.1 Torr, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s.




Figure 5.30: The shortest stage-electrode separation of 25 mm (A) distorts the form of the plasma, whereas a separation of 35 mm (B) does not. Pressure: 0.1 Torr, Power: 50 W, Pulse frequency: 100 kHz, Pulse off time: 3 μ s, Stage bias: -400 V.

to a Paschen curve, with a minimum in breakdown voltage occurring at a pL of 5–10 Torr cm. Measurements at 0.1 Torr are of particular interest. As can be seen in Fig. 5.28, this pressure gives pL values below the Paschen minimum and therefore minimises the likelihood of the stage-electrode discharge occurring. Furthermore, ion current would be expected to drop off with pressure and stage-electrode separation, so minimising pressure should maximise the extracted ion current. Varying pL allowed the two current sources to be separated, as incident ions would be expected to drop off with separation and pressure, whilst discharge current increases with pL provided it is below the Paschen minimum.

Fig. 5.29 shows the stage current with the plasma off, and the difference in stage current with the plasma on and off, for a pressure of 0.1 Torr. As discussed, any measured stage current

Table 5.2: Standard operating conditions for ExTEnD.

ExTEnD Exposure Conditions	
Pressure	0.1 Torr
Stage-Electrode Separation	35 mm
Stage Bias	-400 V
Electrode Power	50 W
Pulse t_{off}	0.3 μ s
Exposure Time	Approx. 30 minutes
Gas	Deuterium

with the plasma off must be a result of stage-electrode electron flow. As the pL values for these datasets are all below the Paschen minimum, increasing pL results in a lower breakdown voltage and a larger current. In contrast, the difference in stage current decreases with increasing pL . The stage current with the plasma on is thought to be a combination of charge contributions from incident ions and electrons flowing from the stage to the grounded electrode. Although there are limitations, the difference between plasma on current and plasma off current could be used to indicate the ion current contribution³, as presented in Eq.5.4. The difference in current behaves in the same expected manner as ion current and decreases with L , whilst the plasma off current, which must be a measure of lost electrons, increases with L . Therefore, it was concluded that below the breakdown voltage at these low pL values, the rate of electron emission from the stage is minimal and stage current can be used to measure incident ions on the stage. At higher pressures, the lower breakdown voltage creates more restrictions on possible operating conditions, and it becomes more challenging to differentiate between stage discharge and ion extraction regimes. At pressures of 5 Torr and above, the applied bias had a minimal impact on the measured current (which consistently read low positive values) until a sudden arc would form at the breakdown voltage. This behaviour suggests no ion current could be measured at higher pressures due to the increased scattering effects, and any negative stage current measured is solely a result of electron emission from the stage. These results emphasise the challenges of relating stage current to ion current in setups such as ExTEnD, which rely on a discharge plasma as a source of ions. Other glow discharge setups [222] also related ion current but it is unclear whether this is a true representation of the incident ion flux or whether this has been impacted by electron flow.

Based on these conclusions, a stage-electrode separation of 35 mm was selected. Although a shorter separation could help minimise the risk of stage discharge further, at 25 mm, the stage began to distort the form of the plasma as shown in Fig. 5.30. Standard operational conditions for ExTEnD can be seen in Table 5.2.

³Under these conditions, the current contribution from stage-electrode discharge when present appears to be minimal, and below the breakdown voltage the difference in current is simply the measured stage current.

5.5.4 Fluence Estimation

Assuming stage current is a result of ions incident on the sample stage, it can be used to estimate the total fluence of deuterium incident on a sample during an exposure. The total accumulated charge, Q , on the stage is given by

$$(5.5) \quad Q = \int I(t)dt,$$

where $I(t)$ is the stage current as a function of time, t . Assuming a uniform beam distribution, this can be scaled by the ratio of the exposure area, A_0 , to stage area, A_1 , to give the charge accumulated on the sample

$$(5.6) \quad Q_{Sample} = \frac{A_0}{A_1} \int I(t)dt.$$

Multiplying this by the average ion cluster size, α , and accounting for the charge on each ion cluster, e , gives an estimation for the number of incident deuterium ions on the sample in ExTEnD,

$$(5.7) \quad N_{Sample} = \frac{\alpha A_0}{e A_1} \int I(t)dt,$$

the fluence,

$$(5.8) \quad f_{Ex} = \frac{\alpha}{e A_1} \int I(t)dt.$$

and the flux,

$$(5.9) \quad F_{Ex} = \frac{\alpha I}{e A_1}.$$

Therefore, to match the average flux of DELPHI [120], F_D , the ion current required would be

$$(5.10) \quad I_D = \frac{e F_D A_i}{\alpha} = \frac{-1.6 \times 10^{-19} \cdot 3.04 \times 10^{17} \cdot \pi \cdot (27 \times 10^{-3}/2)^2}{2.96} = -9.4 \times 10^{-6} \text{ A}$$

For the 0.1 Torr pressure used in ExTEnD, current measurements were of order 0.1 mA, resulting in a flux 10.64 times greater than DELPHI and exposure times than can be scaled down accordingly to reach equivalent fluence. For the standard 5 hr exposure time used in DELPHI, a run of approximately 30 minutes in ExTEnD should result in an equivalent fluences. Adjusting plasma power and the pulse timings can be used to further control ion current, as these effect ion current whilst having minimal impact on the breakdown voltage. This approximation does rely on a uniform beam across the sample stage, as well as the stage current being a result of positive ions incident on the stage.




Figure 5.31: **A** - Thermocouple isolated with Kapton tape attached to the sample stage prior to exposure. **B** - Temperature measurement during exposure. **C** - Thermocouple piercing the Kapton tape, taken post exposure.

5.5.5 Temperature Measurement

Temperature has a significant impact on the uptake of molecules into a material. Generally, the greater atomic movement in higher temperature materials increased uptake. Unlike DELPHI, ExTEnd has no temperature controlled stage, and the temperature of the sample and stage will be a result of room temperature and exposure conditions. No significant temperature increases were expected, but a one-time temperature measurement was taken to ensure results from ExTEnd could be compared to previously collected room temperature results from DELPHI.

To perform the measurement, the secondary viewport was replaced with a thermocouple feedthrough, and a k-type (positive and negative leads made of Ni-Cr and Ni-Al respectively) thermocouple connected. In an attempt to electrically isolate the thermocouple from the biased stage, the plasma and the chamber, the bare wire was covered in Kapton tape as shown in Fig. 5.31A. This had the added benefit of ensuring the two leads remained separated. The probe was secured to the sample stage by lightly pinching it under one of the M3 bolts that secure the sample and window plates to the insulating block. The temperature measurement was carried out under standard operational conditions (see Table 5.2).

Upon striking the plasma, an instantaneous increase to around 50 °C was observed. The rate of increase clearly indicated this was not a true measurement of temperature and that the




Figure 5.32: Sample stage temperature during exposure to plasma and when cooling down. The plasma and stage bias were both on for the first 30 mins, before being turned off and left to cool for time > 30 mins.

attempt to electrically isolate the thermocouple from the stage had not been sufficient. Turning off the plasma and bias gave an instantaneous drop to slightly above room temperature. Clearly some degree of sample heating was occurring, but measurement would only be possible with the plasma and bias off. As such, the stage bias and electrode power supplies were briefly turned off to make a temperature measurement. Negligible heat loss would occur during the measurement as the plasma was only off for a few seconds. Fig. 5.32 shows the stage temperature during a 30 minute run followed by a cool down period. The maximum temperature measured was 23.5 °C. Judging by the plateau in temperature, it seems reasonable to claim the sample is not expected to exceed a temperature of 25 °C, even for longer runs. Temperature decrease is as expected and indicates that the maximum temperature measured is a true reading, unaffected by the plasma and stage bias.

The thermocouple was removed, and the secondary viewport was replaced for future implantations. When inspecting the thermocouple probe post measurement, it was clear the probe had pierced through the Kapton tape (Fig. 5.31 C), resulting in an electrical connection to the stage and explaining the erroneous 50°C measurements.

5.6 Preliminary Retention Study Using ExTEnD

5.6.1 Samples

There were concerns over how effective ExTEnD would be for a non-conductive sample (such as undoped diamond) due to its reliance on a biased stage. For a conductive sample, the bias is present across the entire sample stage and sample, and the connection to the voltage supply removes surface charging effects. With a non-conductive sample this is not the case, and the

setup relies on the general direction of electric field to accelerate ions into the sample. However, divergence of the beam could lead to ions avoiding the sample but still being recorded as ion current on the sample stage, whilst surface charging on a non-conductive sample could further repel incident ions. To understand the capabilities of the current setup, both conductive and non-conductive samples were tested.

Three highly boron doped sample of close to metallic conductivity were selected for the initial retention experiment. Equivalent samples have been tested previously in DELPHI, allowing total retention and TDS spectra to be compared. The samples (2-5b, 2-6a and 2-17a) were equivalent to 2-3a in Table 3.1 - a BDD film grown on a molybdenum substrate in the hot filament reactor. The samples were grown for 4 hours, with gas flow rates of 200, 2 and 0.7 sccm for H_2 , CH_4 and B_2H_6 in H_2 (5%) respectively. These conditions gave a gas phase B:C of 35000 ppm and an expected [B] of order 10^{21} cm^{-3} based on SIMS results of an equivalent sample. For one of the samples (2-5b), silver DAG was used to ensure an electrical connection from the sample plate to the substrate, with a small dap on the corner of the diamond surface as well. No silver DAG was used for 2-6a, due to concerns that TDS measurement may not be possible with the DAG present. Both 2-5b and 2-6a were grown over 3 years prior to testing. A new sample (2-17a) was also grown under the same conditions and tested.

Additionally, an undoped, non-conductive, E6 sample (equivalent to all other undoped PC samples tested in 3.2), was exposed. A gold mesh was wrapped around the sample in an attempt to mitigate surface charging effects. Due to the relatively large gaps in the mesh (20 lines per inch, corresponding to approximately $1.27 \times 1.27 \text{ mm}^2$ holes), it was wrapped around the sample a few times to decrease the size of individual exposed areas. Again, no silver DAG was used for this sample.

5.6.2 Conditions

Exposure conditions were selected based on results from the testing phase presented in 5.5. The standard operating conditions (see Table 5.2) should allow for a stage bias of -400 V with accurate energy selection and fluence estimates. In testing, these conditions gave stage current values of approximately 0.1 mA, of which the contribution from stage-electrode discharge was thought to be minimal. As shown in Eq.5.5-5.10, this is more than sufficient to reach fluences equivalent to DELPHI. The accumulated charge on the stage was monitored during the implantation to try and match the target fluence of previous samples ($5.5 \times 10^{21} \text{ D cm}^{-2}$).

5.6.3 Exposure

The 2-5b and 2-6a were implanted under the standard operating conditions. The accumulated charge was calculated via Eq. 5.5 throughout the exposure, although an error in this calculation meant the final fluence was slightly higher than intended. Small differences in fluence are not expected to have a significant impact, due to the saturation of the top surface of the diamond




Figure 5.33: Images of two samples during and post exposure. **A** - Undoped (i) and doped (ii) samples (3-1b and 2-6a respectively) post exposure with window plate removed. **B** - Undoped sample during exposure with no arc (i) and an arc between the sample and grounded electrode (ii). **C** - Optical images of the undoped sample surface post exposure, at $\times 5$ (i) and $\times 50$ (ii) magnification.

observed in computational results. A summary of exposure conditions can be seen in Table 5.3. Examining the samples post exposure, the exposed region appears slightly lighter than the unexposed border of the sample as seen in Fig. 5.33Aii. A similar effect had previously been observed in DELPHI where the unexposed area below a wire cross appeared darker than the exposed area.

The stage current for 2-17b was much higher than other samples under standard conditions. In an attempt to reduce the ion flux, the electrode power was dropped from 50 W to 30 W, giving a reduction in stage current from 0.27 to 0.17 mA. Despite this, the stage current was still higher than other samples resulting in a shorter exposure to reach the desired fluence.

The undoped sample was more challenging to implant. Initial attempts to use a -400 V stage resulted in arcing and the stage supply tripping - indicating a current exceeding 10 mA. To avoid this, the stage bias was gradually increased from 0 to -200 V. Although this bias was reached without tripping, regular arcs from the stage to the grounded electrode (see Fig. 5.33B) were forming every 10 s or so. These arcs must be a result of surface charging, as were originating from

Table 5.3: Measured variables for three doped samples (2-5b, 2-6a and 2-17a) and an undoped (3-1b) sample exposed in ExTEnD. For the undoped sample, a gradual increase in stage bias was required to reach the -150 V that the majority of the exposure was performed at. Values in brackets indicate the value when the bias was at -150 V, whilst values not in brackets give values across the entire exposure.

Sample	2-6a	2-5b	2-17a	3-1b
Stage Bias (V)	-400	-400	-400	Various (-150)
Electrode Power (W)	50	50	30	50
Exposure time (minutes:seconds)	28:51	22:07	17:37	35:34 (30:17)
Average Stage Current (mA)	-0.115 ± 0.004	-0.159 ± 0.009	-0.19 ± 0.03	-0.100 ± 0.007 (-0.0980 ± 0.0007)
Accumulated Sample Charge (C)	-0.0223	-0.0230	-0.0223	-0.0228 (-0.0188)
Estimated Fluence ($\times 10^{21} \text{ D m}^{-2}$)	6.44	6.64	6.40	6.60 (5.44)
Estimated Average Flux ($\times 10^{18} \text{ D m}^{-2}\text{s}^{-1}$)	3.72	5.18	6.10	3.09 (3.00)

the sample itself rather than the stage. As the arc is going to the grounded electrode, rather than the negatively biased powered electrode, it must be a result of a build up of negative charge on the surface. For a conductive sample, electrons supplied by the bias supply matches the fluence of positive ions on the stage - resulting in no charging effects. The build up of negative charge here suggests that, although the number of electrons supplied to the stage is equivalent to the number of ions incident on the stage, a reduced number of ions are incident on the sample, resulting in a build up of negative charge from the continuous supply of electrons. Therefore, the stage current is unlikely to be an accurate representation of the fluence on the sample when stage arcing is occurring.

To reduce arcing, the bias was dropped to -150 V where occasional arcs did persist but much less frequently. The lower voltage also resulted in a lower average current and a slightly longer exposure time was required to reach the target fluence. The explanation above would suggest the rate of arcing might decrease linearly with stage current, as lower stage current would result in a slower build up of surface charge. Qualitatively, this was not observed, and dropping the stage bias to -150 V resulted in a notable reduction in arcing rate with minimal changes to stage current. This could suggest the stronger electric field of the higher stage bias was diverting more ions away from the sample, and that sample current may be a reasonable approximation for fluence provided arcing is not occurring. Examining the surface of this sample in an optical microscope, marks on the surface from the arcing were observed, alongside spattering of gold where the arcing had occurred (see Fig. 5.33C). Arcing would need to be addressed if non-conductive samples were to be run regularly to avoid contamination and damage of the sample. The gold mesh also left a shadow on the sample, in a similar manner to the wire cross used in DELPHI.

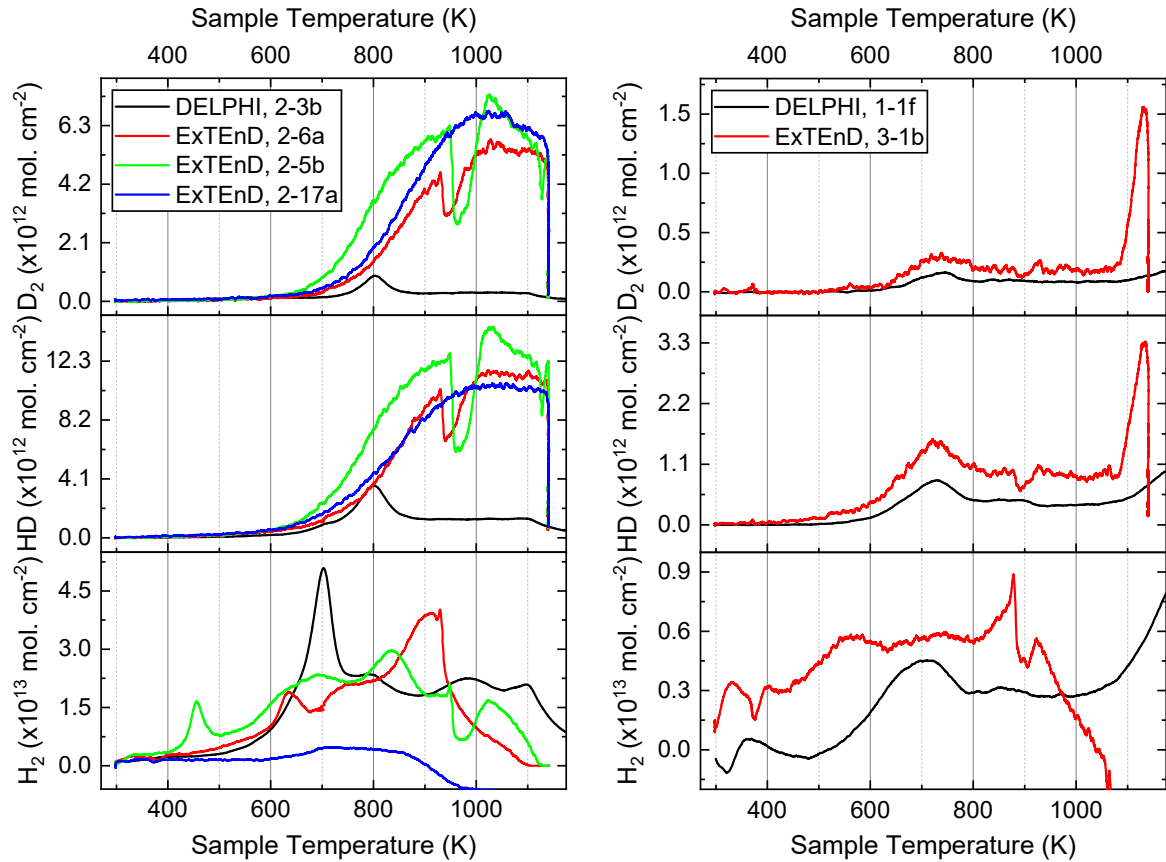


Figure 5.34: Comparison of TDS spectra for samples grown under equivalent conditions and exposed to deuterium ions in either DELPHI or ExTEnD. Counts have been scaled by both exposure area and a calibration factor to allow results from two different TDS facilities to be compared. Exposure conditions and retention values can be seen in Tables 5.3 and 5.4 respectively.

5.6.4 TDS Results

TDS measurements were taken using an equivalent (but different) Hiden Analytica Ltd. Type 640100 TPD workstation to that used for previous measurements [125]. As before, the sample stage was heated from room temperature to 1273 K at a ramp rate of 10 K min^{-1} and held at the maximum temperature for one-hour. Again, an AlN layer was placed between the heater and the sample, and a temperature correction to account for the slower heat transfer to the sample was applied. The temperature correction was determined by measuring argon desorption from silicon, which displays a narrow peak at a known temperature, and was comparable to the temperature calibration value determined for the other TDS (outlined in [151]). Background counts were measured for this temperature profile and removed from final results. Leak calibration tests were performed with both H_2 and D_2 to obtain calibration factors, with the calibration factor for HD taken to be the average between the two.

Table 5.4: Comparison of calibrated total counts from TDS data for diamond samples exposed to deuterium ions in two experimental setups (DELPHI and ExTEnD).

Sample	Exposure setup	TDS Total Counts			
		Energy (eV)	HD ($\times 10^{15} \text{cm}^{-2}$)	D ₂ ($\times 10^{15} \text{cm}^{-2}$)	Total D ($\times 10^{15} \text{cm}^{-2}$)
E6 undoped 1-1f	DELPHI	200	3.62	0.925	5.47
E6 undoped 3-1b	ExTEnD	150	5.07	1.90	8.88
Boron doped 2-3b	DELPHI	400	6.56	1.53	9.61
Boron doped 2-6a	ExTEnD	400	25.20	16.69	58.88
Boron doped 2-5b	ExTEnD	400	31.80	21.40	74.61
Boron doped 2-17a	ExTEnD	400	25.91	21.46	68.83

Spectra of the four samples, alongside previously collected spectra of equivalent samples exposed in DELPHI, can be seen in Fig. 5.34. The undoped, non-conductive sample, is in good agreement with data collected in DELPHI (sample 1-1f in Chapter 3), with desorption peaks seen in both spectra at both 740 and 1125 K. Overall counts and retention values are higher than the DELPHI sample (see Table 5.4), with the high temperature peak in particular being notably larger. As counts remained high at the end of the temperature ramp, it is likely that this peak would continue to grow and suggests significant amounts of deuterium occupying very stable binding sites.

In contrast, TDS spectra of the boron doped samples are notably different to the equivalent sample tested in DELPHI. For samples 2-6a and 2-5b, there is a sudden drop seen in all counts around 950 K. Discontinuities such as this can be a result of fluctuations in stage temperature, or desorption peaks in background counts. However, this is not the case here, as background counts were consistently very low for both HD and D₂, and the temperature profile in this region was smooth and linear. Instead, this is believed to be an interface effect, where trapped volumes between the substrate and the film can result in sudden changes in outgassing. Both of these samples exhibited small degrees of delamination from the substrate. In contrast, sample 2-17a showed no delamination and did not present this dip in TDS spectra, supporting the idea that it was a result of an interface effect. Discounting this region, the desorption spectra all take a similar form, and gradually rise towards a very broad peak, with negligible counts detected below temperatures of 600 K. Overall, deuterium retention was an order of magnitude higher than the equivalent sample (2-3b) previously tested in DELPHI. As discussed in 3.2, sample 2-3b presented with two peaks at 700 and 800 K in HD and D₂ spectra, followed by a plateau at higher temperatures. These features were previously attributed to the presence of boron dopants increasing disorder and potentially desorption from boron-hydrogen clusters. However, none of this finer detail is present here, and it appears that the single large peak observed in ExTEnD samples has simply eclipsed these peaks. The width and size of the peak could suggest damage to the diamond has created a large variety of potential binding sites and increased total retention.

Similarly, the undoped sample (3-1b) displayed increased counts at high temperatures, which could also be explained by damage to the lattice creating new, stable, binding sites.

Assuming the samples tested in ExTEnD and DELPHI are comparable, this increase in damage must be a result of exposure conditions. Temperature measurements confirmed no significant increase in sample temperature is expected and, as DELPHI samples were also exposed at room temperature, this is unlikely to be the cause of any differences present. An increased ion energy could also result in increased damage but, again, this is unlikely. The ion energy is dictated by the potential difference between the sample stage and the plasma. The potential of the stage is known to reasonable accuracy, so a significantly incorrect ion energy would require a plasma potential of order 10^2 – 10^3 V which is far higher than the few volts expected. Furthermore, previous exposures of undoped diamond at higher energies (see 3.2) presented a energy dependant peak at 900 K which was not observed in sample 3-1b, suggesting this and sample 1-1f (exposed in DELPHI) used a comparable ion energy. Instead, differences in spectra are thought to be a result of the higher flux of ExTEnD, which is approximately an order of magnitude greater than DELPHI. The impact of flux has been explored in similar studies of other materials. Poon *et al.* [240] compared the impact on retention of varying deuterium ion flux and fluence on single crystal tungsten, and observed a sharp decrease in retention below fluxes of 10^{18} D m⁻² s⁻¹ and fluences of 10^{21} D m⁻². As both flux and fluence in ExTEnD is at the threshold value, it is possible decreasing to the 10^{17} D m⁻² s⁻¹ flux of DELPHI could have a significant effect on retention mechanisms. That being said, Poon *et al.* used a higher ion energy (1.5 keV) and the impact of flux on retention will be material dependent, so although it seems likely varying the flux would impact retention, comparisons to this study are limited. The material dependence of the impact of flux was highlighted in the work of Zayachuck *et al.* [242], who observed higher retention for high flux exposure of tungsten, whereas the inverse was true for a tungsten-tantalum alloys. In other studies of tungsten [241, 243], variation in retained deuterium was linked to structural changes to the material which are dependent on fluence and flux. These findings support the idea that the higher flux of ExTEnD could increase damage sustained to the diamond. t' Hoen *et al.* [243] observed deeper deuterium for higher fluences of ions, and small differences in fluence between samples exposed in ExTEnD and DELPHI maybe present as a result of non-uniform beam divergence. However, t' Hoen *et al.* varied fluences by orders of magnitude whereas subtle differences in fluence from beam divergence are unlikely to result in obvious changes to desorption spectra observed here.

H₂ spectra take a notably different form to those of the samples exposed in DELPHI, particularly at high temperatures. Conclusions that can be made from these spectra relating to ExTEnD are limited, as H₂ was not implanted but incorporated into the diamond during growth. Furthermore, the TDS used for these measurements is not typically run at such high temperatures resulting in large amount of background hydrogen outgassing from the chamber which, when removed from the spectra, resulted in negative hydrogen counts for some samples.

Overall, the dependence of deuterium retention on flux and fluence appears to be linked to physical changes to the material which will vary with flux. From this work, it seems likely the higher flux has resulted in increased damage to the diamond, creating a large variety of stable binding sites and greater retention. Differences in fluence are not thought to play a significant role in this work, as any deviations from measured flux as a result of a non-uniform ion beam are thought to be small and computational work suggested saturation of the subsurface would be expected.

5.7 Conclusions

A new setup capable of performing low energy hydrogen/deuterium implantations with known ion energy and fluence has successfully been assembled and tested. Under standard operating conditions, ExTEnD can perform ion exposures across an $8 \times 8 \text{ mm}^2$ area at room temperature, with ion energies up to 400 eV and a flux of approximately $4 \times 10^{18} \text{ D m}^{-2} \text{ s}^{-1}$. Use of a pulsed DC plasma aided the simplicity and cost of the setup, although added restrictions on the exposure conditions possible. Through careful selection of operational conditions, it was possible to find settings which meant stage current could be used as an approximation for flux, and stage bias could be used to dictate ion energy. Preliminary retention results were compared to results for equivalent samples tested in an established facility. Although similarities were present, samples exposed in ExTEnD presented with notably higher deuterium retention, thought to be a result of the higher flux of ExTEnD leading to damage to the sample.

5.8 Future Work

There were clear issues surrounding surface charging and arcing when exposing the non-conductive sample. Replacing the window plate with a grid of regular holes could offer a known exposure area whilst preventing surface charging more reliably than the gold mesh. Further adaptations to the sample stage would be required to ensure a good electrical connection to the surface of the sample. The non-linear relation between arcing frequency and stage current suggested some beam divergence around non-conductive samples may be occurring. This could be explored further via a series of masking experiments. One method to do this would involve creating non-conductive masks (using MACOR or a similar ceramic) and measuring ion current whilst masking different regions of the sample stage. If the ion current scales with the amount of conductive area exposed, this would suggest incident ions remain uniform across the surface of the sample stage even with non-conductive samples. If there are less significant changes in measured ion current, this would indicate ions are avoiding the non-conductive areas and are predominantly incident on the conductive regions - suggesting more work would be required to avoid surface charging effects.

Scaling the current for area as suggested above relies on the distribution of ions on the sample stage being uniform. If the ion current is not uniform across the surface, additional scaling factors may be required when comparing areas of masked and unmasked regions, as well as relating measured current to estimated sample fluence. The uniformity of the beam can be assessed through the etching of an amorphous carbon film as was carried out for DELPHI [120]. Amorphous carbon is readily etched by hydrogen ions, if an area is receiving a higher fluence, a greater amount of material will be etched away in this region. Measuring the topology of the film post exposure using atomic force microscopy or ellipsometry (as was used for DELPHI) would effectively map out the ion current across the sample stage.

There may also be a more quantitative way to evaluate the stage-electrode discharge. Measuring the current flowing through the grounded electrode to ground would give the total current as a result of electrons flowing from the powered electrode and the stage. Subtracting the current supplied to the powered electrode would then give the current from the stage. Conditions could then be selected to ensure stage-electrode electron flow was zero. In practice, this would be challenging to implement, as this would require the grounded electrode to be electrically isolated from the chamber to allow an ammeter to be added between this and the grounding point. Additional current measurements would also be required for the powered electrode, as the current output on the power supply is not to the required precision.

The flux in ExTEnD is an order of magnitude higher than that of DELPHI, potentially leading to differences in TDS spectra. It is unlikely this could be explored computationally, as simulations already rely on fluxes many orders of magnitude higher than experiments, so a dedicated set of experiments would be required to explore this further. It is possible to change the flux in ExTEnD by adjusting either the electrode power or the pulse off time as discussed in 5.5.2. Although adjusting the pulse off time can be used to decrease stage current and average flux, it is not clear whether this is truly reducing the flux or simply maintaining the same high flux and leaving longer gaps in between pulses. To explore this further, current measurements for various pulse frequencies and pulse lengths would need to be made. Comparing these to different electrode powers (thought to give a true change in flux) and retention results could indicate how these parameters impact flux.

CONCLUSIONS

Both computational and experimental studies helped develop an understanding of diamond's interactions with low energy hydrogen isotopes. In MD simulations, different diamond surfaces were bombarded with hydrogen isotopes of varying energies, allowing an atomistic understanding of retention mechanisms. Whilst experimentally, various diamond samples were exposed to deuterium ions in two setups: the established DELPHI facility, and a new setup (ExTEnD) that was shown to be capable of performing equivalent experiments.

Desorption peaks in H_2 TDS spectra were used to separate deuterium desorption from grain boundaries and grains, with significant desorption from the latter only observed for ion energies greater than 600 eV (3.2.1) suggestive of ion damage to the surface. Boron doped samples displayed higher retention than undoped samples (3.2.2), as a result of increased defect concentration leading to more potential binding sites, the formation of boron-hydrogen complexes and the additional surface area from unpolished samples. More testing of diamond with lower boron concentrations would be of interest and more applicable to fusion applications.

Computational work helped to develop an atomistic understanding of experimental results and bridge the gap between test facilities and fusion reactors. Low energy ions incident on diamond surfaces can be reflected, bind to the surface or penetrate small distances into the diamond. Occupation of interstitial positions in the top atomic layers pushes carbon atoms out of the surface to create vacant sites, but these are unable to form beyond the surface (4.3) and significant physical sputtering was not observed under any of the conditions tested. As the bombardment continues, the concentration of hydrogen gradually builds up within the top nanometers of the sample and a disordered layer begins to form. Increasing the ion energy results in deeper penetration depths and a thicker disordered region, and a greater fluence is required to saturate this volume and create the disordered region. Only once this layer has

formed is steady etching of the surface observed, with the disordered region maintaining a roughly constant thickness (4.2). Channels present in grain boundaries act to funnel ions to much higher penetration depths than seen for bulk diamond, whilst the additional space present allows for vacancies to form along the grain boundary deeper within the material (4.5.2.1). Although these effects are significant for perfect diamond structures, the formation of the disordered layer during continuous bombardments obscures these channels leading to grain boundaries having minimal impact on uptake of ions or the etching of the surface (4.5.2.2). However, experimental results suggested that at low ion energies grain boundary trapping sites dominate hydrogen uptake, and that damage to the surface was required for significant grain trapping. This contradiction could be explained by deuterium diffusing across the grains into grain boundaries during the exposure - a phenomenon not observed on the short simulation timescales.

In diffusion simulations, grain boundaries were shown to dictate diffusion characteristics (4.4), meaning they could impact retention over time. Potential wells across most of the grain boundaries tested led to hydrogen trapping - in good agreement with physical expectations. Hydrogen trapping lead to anisotropic diffusion, where diffusion was restricted to one or two dimensions within the grain boundary plane. Despite this, diffusion coefficients within all grain boundaries tested were less than or equal to bulk coefficients. This observation was concluded to be a result of sp^2 hybridised carbon within grain boundaries forming sp^3C-H bonds, which are much stronger than the interstitial positions occupied by hydrogen in the bulk.

Overall, it was concluded that interactions were limited to the surface of the diamond, leaving bulk properties largely intact. Saturation of the surface and low penetration depths resulted in limited retention and etching that was restricted by the formation of a disordered top layer. Retention values from experimental and computational results were of the same order of magnitude and were slightly lower than other materials tested following an equivalent experimental procedure. These results offer promising indications for the use of diamond under fusion relevant conditions. However, carbon nanoparticles were observed on some samples (3.3) and were thought to be a result of surface etching but no trends could be found to suggest the conditions at which this occurred.

A new experimental setup was designed, assembled and used to implant four diamond samples to validate the setup (5.3, 5.4 and 5.6 respectively). Testing prior to exposure indicated that ion energies up to 400 eV were possible and typically resulted in an ion flux of order $10^{18} \text{ m}^{-2}\text{s}^{-1}$ (5.5). Due to the use of a biased stage to extract ions, there were challenges surrounding non-conductive samples. Despite attempts to mitigate surface charging, arcing was present, and a lower ion energy of 150 eV had to be used. Nevertheless, TDS results for this sample presented peaks at temperatures in agreement with previous experiments. Conductive, highly boron doped samples were also tested. These did not exhibit the arcing observed with the conductive sample, but spectra differed to past results. For both doped and undoped samples, deuterium retention was notably higher than equivalent samples exposed in DELPHI. Differences in retention and

TDS spectra are thought to be a result of the increased flux of ExTEnD leading to damage in the samples, creating a variety of stable binding sites.

The breadth of this work has highlighted several areas that could be developed with future work. TDS results of undoped PC samples suggested some ion damage may have occurred resulting in new binding sites. The probe depth of Raman spectroscopy meant it was not possible to detect structural changes at the surface and a different, more surface sensitive method, would be required to confirm this. Similarly, carbon nanoparticles were observed on the surface of some samples which, although it was unclear whether these were revealed or redeposited, were thought to indicate etching of the surface had occurred. Future work could look to determine the cause of these nanoparticles, the etch rates of different diamond materials, and the conditions (temperature, ion energy etc.) in which etching can be minimised. The increased conductivity of boron doped diamond could be used to reduce arcing damage for PFMs, and is required for sensor applications. Retention results here suggested a significant increase in retention, but at boron concentrations greater than useful for these applications. Furthermore, there were caveats surrounding the unpolished surfaces and uncertainty surrounding the boron concentration of the samples. As such, a more systematic approach of testing these samples, including testing samples at more relevant concentrations could more conclusively find the impact of boron dopants.

Although computational work offered valuable insight into the retention mechanisms behind experimental results, simulating deuterium clusters could help bridge the understanding between experimental setups and fusion reactors further. D_3^+ clusters are the most common species extracted from the low temperature plasmas which are common in experimental setups. This is unlike fusion reactors which will contain exclusively single ions. Understanding how these clusters behave when incident on a surface could be important in understanding how results from the lab could translate into real world fusion reactors. The AIREBO potential was not able to simulate a stable D_3^+ cluster, but could be used to test D_2 and D bombardments separately. Alternatively, a different potential may have more success in recreating clusters.

The majority of bombardment simulations were performed at room temperature to try and make them comparable to experimental work. However, in fusion reactors, significant thermal loads would be present. Performing simulations at elevated temperatures could give insight into how temperature effects etching of the surface and permeation of hydrogen in a damaged diamond. Furthermore, damage to the diamond could be simulated to create additional trapping sites which would presumably impact both retention and diffusion within the diamond.

Simulation of grain boundaries presented interesting behaviour for both bombardments and diffusion simulations. However, with the exception of continuous bombardments at high energy, this behaviour was shown to be dependant on the specific grain boundary. The large variety of grain boundaries present in a PC diamond means testing many different grain boundaries would be required to apply broader conclusions to polycrystalline diamond in general.

Regarding ExTEnD, more work is needed to allow for non-conductive samples to be tested at

higher energies. Replacing the window plate with a grid of holes could help minimise surface charging, whilst still giving a known exposure area and reduced sputtering compared to the previously used gold mesh. There were also concerns over changes to beam divergence for non-conductive samples - where a charged surface could divert ions away from the sample. In fact, for all samples, a measurement of beam divergence (using the etching of an amorphous carbon film) would allow for more accurate sample flux and fluence estimates.

Compared to DELPHI, the increased flux of ExTEnd resulted in notably different desorption spectra, suggestive of ion damage to the surface of the sample. As with previous samples, it would be interesting if features in the desorption spectra could be associated with physical changes to the sample as a result of ion damage. As ion flux can be varied in this setup, exploring how this impacts desorption spectra, as well as damage to the surface, could offer valuable insight. It is thought that both pulse timings and plasma power can be used to adjust flux, but this would need to be confirmed through testing. Exploring a greater variety of samples and comparing them to results from DELPHI would allow for a broader understanding of the setup. For example, there is data available for fusion relevant metals such as Eurofer, tungsten and beryllium, and these would not present the challenges of non-conductive samples.

BIBLIOGRAPHY

- [1] R. M. Magee, K. Ogawa, T. Tajima, I. Allfrey, H. Gota *et al.*
First measurements of $p^{11}\text{B}$ fusion in a magnetically confined plasma.
Nat. Commun., 14:955, 2023.
- [2] C. L. Smith and S. Cowley.
The path to fusion power.
Philos. T. Roy. Soc. A., 368:1091–1108, 2010.
- [3] T. Griffiths, R. Pearson, M. Bluck and S. Takeda.
The commercialisation of fusion for the energy market: a review of socio-economic studies.
Progress in Energy, 4:042008, 2022.
- [4] ITER, ITER Organization.
<https://www.iter.org/>, Accessed 2025-02-07.
- [5] H. Zohm.
Magnetic Fusion Energy - Plasma performance, burn and sustainment.
Woodhead Publishing, pages 7–30, 1st Edt., 2016.
- [6] R. J. Nicholls, S. Diaz-Moreno, W. Iliffe, Y. Linden, T. Mousavi *et al.*
Understanding irradiation damage in high-temperature superconductors for fusion reactors
using high resolution X-ray absorption spectroscopy.
Commun. Mater., 3:52, 2022.
- [7] A. J. Creely, D. Brunner, R. T. Mumgaard, M. L. Reinke, M. Segal *et al.*
SPARC as a platform to advance tokamak science.
Phys. Plasmas, 30:090601, 2023.
- [8] Commonwealth Fusion Systems, CFS.
<https://cfs.energy/>, Accessed 2024-05-07.
- [9] Tokamak Energy, Tokamak Energy Ltd.
<https://tokamakenergy.com/>, Accessed 2024-05-07.

- [10] A. E. Costley.
Towards a compact spherical tokamak fusion pilot plant.
Philos. T. Roy. Soc. A., 377:20170439, 2019.
- [11] I. T. Chapman, S. C. Cowley and H. R. Wilson.
The Spherical Tokamak for Energy Production: Theme issue introduction.
Philos. T. Roy. Soc. A., 382:20230416, 2024.
- [12] H. Anand, O. Bardsley, D. Humphreys, M. Lennholm, A. Welander *et al.*
Modelling, design and simulation of plasma magnetic control for the Spherical Tokamak
for Energy Production (STEP).
Fusion Eng. Des., 194:113724, 2023.
- [13] K. Risse, T. Rummel, H. Bosch, V. Bykov, A. Carls *et al.*
The magnet system of Wendelstein 7-X stellarator in operation.
Fusion Eng. Des., 136:12–16, 2018.
- [14] P. Magaud and G. Marbach.
Encyclopedia of Energy - Nuclear Fusion Reactors.
Elsevier Science, pages 365–381, 2004.
- [15] F. B. Marcus.
Systems Approaches to Nuclear Fusion Reactors - Helical, Linear and Inertial Fusion
Reactor Concepts: Inertial Fusion and Magnetic Fast Pulsed Systems.
Springer, pages 401–435, 1st Edt., 2022.
- [16] Lawrence Livermore National Laboratory, National Ignition Facility and Photon Science,
LLNL.
<https://lasers.llnl.gov/>, Accessed 2024-05-08.
- [17] M. Laberge, S. Howard, D. Richardson, A. Froese, V. Saponitsky *et al.*
Acoustically driven Magnetized Target Fusion.
Proceedings of IEEE 25th Symposium on Fusion Engineering 2013, 42, 2013.
- [18] W. Jiang, Y. Peng, Y. Zhang and G. Lapenta.
Numerical modeling of tokamak breakdown phase driven by pure Ohmic heating under
ideal conditions.
Nucl. Fusion, 56:126017, 2016.
- [19] D.W. Swain and R.H. Goulding.
ITER ion cyclotron system: Overview and plans.
Fusion Eng. Des., 82:603–609, 2007.

- [20] T. Omori, F. Albajar, T. Bonicelli, G. Carannante, M. Cavinato *et al.*
Progress in the ITER electron cyclotron heating and current drive system design.
Fusion Eng. Des., 96-97:547–552, 2015.
- [21] R. S. Hemsworth, D. Boilson, P. Blatchford, M. D. Palma, G. Chitarin *et al.*
Overview of the design of the ITER heating neutral beam injectors.
New J. Phys., 19:025005, 2017.
- [22] Tritium Breeding, ITER Organization.
<https://www.iter.org/mach/TritiumBreeding>, Accessed 2022-11-15.
- [23] M. Abdou, N. B. Morley, S. Smolentsev, A. Ying, S. Malang *et al.*
Blanket/first wall challenges and required R and D on the pathway to DEMO.
Fusion Eng. Des., 100:2–43, 2015.
- [24] C. R. Malone, H. B. Flynn, A. D. Somers, P. A. Rowell and G. K. Larsen.
Approach to startup inventory for viable commercial fusion power plant.
Fusion Eng. Des., 206:114563, 2024.
- [25] R. J. Pearson, A. B. Antoniazzi and W. J. Nuttall.
Tritium supply and use: a key issue for the development of nuclear fusion energy.
Fusion Eng. Des., 136:1140–1148, 2018.
- [26] A. Loarte, B. Lipschultz, A. Kukushkin, G. F. Matthews, P. C. Stangeby *et al.*
Power and particle control.
Nucl. Fusion, 47:S203, 2007.
- [27] M. R. Gilbert, Ž. Zacharauskas, P. Almond, N. Scott-Mearns, S. Reynolds *et al.*
Fusion waste requirements for tritium control: Perspectives and current research.
Fusion Eng. Des., 202:114296, 2024.
- [28] D. Matveev, D. Douai, T. Wauters, A. Widdowson, I. Jezu *et al.*
Tritium removal from JET-ILW after T and D–T experimental campaigns.
Nucl. Fusion, 63:112014, 2023.
- [29] S. M. Gonzalez de Vicente, N. A. Smith, L. El-Guebaly, S. Ciattaglia, L. Di Pace *et al.*
Overview on the management of radioactive waste from fusion facilities: ITER, demonstration machines and power plants.
Nucl. Fusion, 62:085001, 2022.
- [30] K. O. E. Henriksson, K. Nordlund, A. Krashenninnikov and J. Keinonen.
The depths of hydrogen and helium bubbles in tungsten: A comparison.
Fusion Sci. Technol., 50:43–57, 2006.

- [31] J. Linke, J. Du, T. Loewenhoff, G. Pintsuk, B. Spilker *et al.*
Challenges for plasma-facing components in nuclear fusion.
Matter. Radiat. Extrem., 4:056201, 2019.
- [32] B. Kalin, A.N. Suchkov, V.T. Fedotov, O. Sevryukov, A. Ivannikov *et al.*
Brazing of Be with CuCrZr-bronze using copper-based filler metal STEMET.
Nucl. Mater. Energy, 9:388–393, 2016.
- [33] T. Hirai, S. Panayotis, V. Barabash, C. Amzallag, F. Escourbiac *et al.*
Use of tungsten material for the ITER divertor.
Nucl. Mater. Energy, 9:616–622, 2016.
- [34] M. Ulrickson, V. Barabash, S. Chiochio, G. Federici, G. Janeschitz *et al.*
Selection of plasma facing materials for ITER.
Proceedings of 16th International Symposium on Fusion Engineering, 1:394–398, 1995.
- [35] Joint European Torus, UKAEA.
<https://ccfe.ukaea.uk/research/joint:european:torus/>, Accessed 2022-11-17.
- [36] R. A. Causey.
Hydrogen isotope retention and recycling in fusion reactor plasma-facing components.
J. Nucl. Mater., 300:91–117, 2002.
- [37] J. Kiippers.
The hydrogen surface chemistry of carbon as a plasma facing material.
Surf. Sci. Rep., 22:249–321, 1995.
- [38] Y. Nobuta, Y. Yamauchi, T. Hino, S. Akamaru, Y. Hatano *et al.*
Tritium absorption of co-deposited carbon films.
Fusion Eng. Des., 87:1070–1073, 2012.
- [39] M. Mayer, M. Balden and R. Behrisch.
Deuterium retention in carbides and doped graphites.
J. Nucl. Mater., 252:55–62, 1998.
- [40] L. L. Snead and M. Ferraris.
Comprehensive Nuclear Materials - Graphite and Carbon Fiber Composite for Fusion.
Elsevier, 6:54–92, 2nd Edt., 2020.
- [41] F. Romanelli and J. E. Contributors.
Overview of the JET results with the ITER-like wall.
Nucl. Fusion, 53:104002, 2013.

- [42] C. Luo, L. Xu, L. Zong, H. Shen and S. Wei.
Research status of tungsten-based plasma-facing materials: A review.
Fusion Eng. Des., 190:113487, 2023.
- [43] J. Bucalossi, J. Achard, O. Agullo, T. Alarcon, L. Allegretti *et al.*
Operating a full tungsten actively cooled tokamak: overview of WEST first phase of operation.
Nucl. Fusion, 62:042007, 2022.
- [44] V. Philipps.
Tungsten as material for plasma-facing components in fusion devices.
J. Nucl. Mater., 415:S2–S9, 2011.
- [45] R.G. Abernethy.
Predicting the performance of tungsten in a fusion environment: a literature review.
Mater. Sci. Tech. Ser., 33:388–399, 2017.
- [46] M. Rieth, S.L. Dudarev, S.M. Gonzalez de Vicente, J. Aktaa, T. Ahlgren *et al.*
Recent progress in research on tungsten materials for nuclear fusion applications in Europe.
J. Nucl. Mater., 432:482–500, 2013.
- [47] L. Pentecoste, P. Brault, A. L. Thomann, P. Desgardin, T. Lecas *et al.*
Low energy and low fluence helium implantations in tungsten: Molecular dynamics simulations and experiments.
J. Nucl. Mater., 470:44–54, 2016.
- [48] V. Kh Alimov and J. Roth.
Hydrogen isotope retention in plasma-facing materials: Review of recent experimental results.
Phys. Scripta, T128:6–13, 2007.
- [49] Z. Ye, X. Ma, P. He, Z. Wang, Q. Liu *et al.*
The investigation of plasma-induced wettability of liquid tin-capillary porous system.
Nucl. Mater. Energy, 20:1001–1008, 2019.
- [50] M. J. Baldwin, R. P. Doerner, S. C. Luckhardt and R.W. Conn.
Deuterium retention in liquid lithium.
Nucl. Fusion, 42:1318–1323, 2002.
- [51] A. N. Shcherbak, S. V. Mirnov, A. S. Dzhurik, V. B. Lazarev, Y. A. Vasina *et al.*
Experiments on the capture of Li, H, and D by lithium collectors at different surface temperatures at the T-11M tokamak.

- Plasma Phys. Rep.*, 44:1001–1008, 2018.
- [52] Z. X. Liu, X. Gao, S. C. Liu, S. Y. Ding, T. Y. Xia *et al.*
H-mode power threshold and confinement in a molybdenum wall with different magnetic configurations on the EAST tokamak.
Nucl. Fusion, 53:073041, 2013.
- [53] J. N. Brooks, L. El-Guebaly, A. Hassanein and T. Sizyuk.
Plasma-facing material alternatives to tungsten.
Nucl. Fusion, 55:043002, 2015.
- [54] M. Yu Lavrentiev, A. Hollingsworth, J. Hess, S. Davies, A. Wohlers *et al.*
Effects of self-irradiation on deuterium retention and reflectivity of molybdenum, fusion plasma-facing material: Combined experimental and modeling study.
J. Appl. Phys., 132:125902, 2022.
- [55] R. A. Causey, C. L. Kunz and D. F. Cowgill.
Deuterium retention and release from molybdenum exposed to a Penning discharge.
J. Nucl. Mater., 337-339:600–603, 2005.
- [56] M.R. Gilbert, L.W. Packer and T. Stainer.
Experimental validation of inventory simulations on molybdenum and its isotopes for fusion applications.
Nucl. Fusion, 60:106022, 2020.
- [57] G. Sinclair, T. Abrams, S. Bringuier, D. M. Thomas, L. Holland *et al.*
Quantifying erosion and retention of silicon carbide due to D plasma irradiation in a high-flux linear plasma device.
Nucl. Mater. Energy, 26:100939, 2021.
- [58] T. Abrams, S. Bringuier, D. M. Thomas, G. Sinclair, S. Gonderman *et al.*
Evaluation of silicon carbide as a divertor armor material in DIII-D H-mode discharges.
Nucl. Fusion, 61:066005, 2021.
- [59] L. B. Begrambekov and O. I. Buzhinskij.
Features and advantages of boron carbide as a protective coating of the tokamak first wall.
Plasma Devices Oper., 15:193–199, 2007.
- [60] L. B. Begrambekov, A. V. Grunin and Y. A. Sadovskiy.
Behavior of B₄C coating on tungsten under exposition on T-10 tokamak plasma.
J. Phys. Conf. Ser., 1281:193–199, 2019.
- [61] J. Linke, H. Bolt, R. Doerner, H. Grübmeier, Y. Hirooka *et al.*

- Performance of boron/carbon first wall materials under fusion relevant conditions.
J. Nucl. Mater., 176-177:856–863, 1990.
- [62] Y. Gotoh, T. Yamaki, T. Ando, R. Jimbou, N. Ogiwara *et al.*
Sputtering characteristics of B₄C-overlaid graphite for keV energy deuterium ion irradiation.
J. Nucl. Mater., 196-198:708–712, 1992.
- [63] J. Roth, E. Tsitrone, A. Loarte, Th. Loarer, G. Counsell *et al.*
Recent analysis of key plasma wall interactions issues for ITER.
J. Nucl. Mater., 390-391:1–9, 2009.
- [64] S. Porro, G. De Temmerman, P. John, S. Lisgo, I. Villalpando *et al.*
Effects in CVD diamond exposed to fusion plasmas.
Phys. Status Solidi A, 206:2028–2032, 2009.
- [65] S. Porro, D. T. G., S. Lisgo, D.L. Rudakov, A. Litnovsky *et al.*
Diamond coatings exposure to fusion-relevant plasma conditions.
J. Nucl. Mater., 415:S161–S164, 2011.
- [66] S. Porro, G. De Temmerman, D. A. MacLaren, S. Lisgo, D. L. Rudakov *et al.*
Surface analysis of CVD diamond exposed to fusion plasma.
Diam. Relat. Mater., 19:818–823, 2010.
- [67] S. Porro, G. D. Temmerman, S. Lisgo, P. John, I. Villalpando *et al.*
Nanocrystalline diamond coating of fusion plasma facing components.
Diam. Relat. Mater., 18:740–744, 2009.
- [68] M. Mayer, M. Balden and R. Behrisch.
Deuterium retention in carbides and doped graphites.
J. Nucl. Mater., 252:55–62, 1998.
- [69] G. De Temmerman, J. Dodson, J. Linke, S. Lisgo, G. Pintsuk *et al.*
Thermal shock resistance of thick boron-doped diamond under extreme heat loads.
Nucl. Fusion, 51:052001, 2011.
- [70] G. De Temmerman, R. P. Doerner, P. John, S. Lisgo, A. Litnovsky *et al.*
Interactions of diamond surfaces with fusion relevant plasmas.
Phys. Scripta, T138:014013, 2009.
- [71] A. Deslandes, M. C. Guenette, C. S. Corr, I. Karatchevtseva, L. Thomsen *et al.*
Deuterium retention and near-surface modification of ion-irradiated diamond exposed to fusion-relevant plasma.
Nucl. Fusion, 54:073003, 2014.

- [72] B.I. Khripunov, A.N. Brukhanov, O.K. Chugunov, V.M. Gureev, V.S. Koidan *et al.*
Evidence of radiation damage impact on material erosion in plasma environment.
J. Nucl. Mater., 390-391:921–924, 2009.
- [73] B. Khripunov, V. Gureev, V. S. Koidan, S. Latushkin, V. Petrov *et al.*
Plasma impact on materials damaged by high-energy ions.
Phys. Scripta, 2011:014052, 2011.
- [74] A. Deslandes, M. C. Guenette, K. Belay, R. G. Elliman, I. Karatchevtseva *et al.*
Diamond structure recovery during ion irradiation at elevated temperatures.
Nucl. Instrum. Meth. B., 365:331–335, 2015.
- [75] A. R. Dunn and D. M. Duffy.
A molecular dynamics study of diamond and graphite under tritium bombardment.
J. Appl. Phys., 110:104307, 2011.
- [76] A. Deslandes, M. C. Guenette, C. M. Samuell, I. Karatchevtseva, M. Ionescu *et al.*
Initial damage processes for diamond film exposure to hydrogen plasma.
Fusion Eng. Des., 88:3101–3107, 2013.
- [77] M. C. Guenette, A. Deslandes, L. Thomsen, C. S. Corr and D. P. Riley.
D and D/He plasma interactions with diamond: Surface modification and D retention.
Diam. Relat. Mater., 49:103–110, 2014.
- [78] A. Dunn.
A molecular dynamics study of diamond as a plasma facing material for fusion - PhD
thesis.
University College London Department of Physics and Astronomy, 2011.
- [79] M. C. Guenette, A. Deslandes, C. M. Samuell, A. Tadich, L. Thomsen *et al.*
NEXAFS spectroscopy of CVD diamond films exposed to fusion relevant hydrogen plasma.
Diam. Relat. Mater., 34:45–49, 2013.
- [80] P. J. in't Veld, S. J. Plimpton and G. S. Grest.
Accurate and efficient methods for modeling colloidal mixtures in an explicit solvent using
molecular dynamics.
Comput. Phys. Commun., 179:320–329, 2008.
- [81] A. P. Thompson, H. Metin Aktulga, R. Berger, D. S. Bolintineanu, W. Michael Brown *et al.*
LAMMPS - a flexible simulation tool for particle-based materials modeling at the atomic,
meso, and continuum scales.
Comput. Phys. Commun., 271:108171, 2022.

- [82] A. R. Dunn, D. M. Duffy and A. M. Stoneham.
A molecular dynamics study of diamond exposed to tritium bombardment for fusion applications.
Nucl. Instrum. Meth. B., 269:1724–1726, 2011.
- [83] F. Mazzocchi, S. Schreck, D. Strauss, G. Aiello, A. Meier *et al.*
Diamond windows diagnostics for fusion reactors—Updates of the design.
Fusion Eng. Des., 123:820–824, 2017.
- [84] B. Weaver.
Tritium handling and safe storage.
Department of Energy Handbook, DOE-STD-1129-2015, 2015.
- [85] V. Nemanič.
Hydrogen permeation barriers: Basic requirements, materials selection, deposition methods, and quality evaluation.
Nucl. Mater. Energy, 19:451–457, 2019.
- [86] L. M. Luo, Y. L. Liu, D. G. Liu, L. Zheng and Y. C. Wu.
Preparation technologies and performance studies of tritium permeation barriers for future nuclear fusion reactors.
Surf. Coat. Tech., 403:126301, 2020.
- [87] M. Tamura and T. Kumagai.
Hydrogen permeability of diamondlike amorphous carbons.
J. Vac. Sci. Technol. A., 35:04D101, 2017.
- [88] A. Hatta, S. Kaneko and K. Hassan.
Hydrogen permeation through diamond-like carbon thin films coated on PET sheet.
Plasma Process. Polym., 4:S241–S244, 2007.
- [89] J. M. Lackner, W. Waldhauserm and M. Kahn.
Study of gas permeation through thin ta-C:H films.
Acta Phys. Pol. A, 127:1236–1239, 2015.
- [90] J. P. Goss, R. Jones, M. I. Heggie, C. P. Ewels, P. R. Briddon *et al.*
Theory of hydrogen in diamond.
Phys. Rev. B, 65:1–13, 2002.
- [91] D. J. Cherniak, E. B. Watson, V. Meunier and N. Kharche.
Diffusion of helium, hydrogen and deuterium in diamond: Experiment, theory and geochemical applications.
Geochim. Cosmochim. Ac., 232:206–224, 2018.

- [92] D. Ballutaud, F. Jomard, B. Theys, C. Mer, D. Tromson *et al.*
Hydrogen diffusion and stability in polycrystalline CVD undoped diamond.
Diam. Relat. Mater., 10:405–410, 2001.
- [93] I. Z. Machi, J. E. Butler, S. H. Connell, B. P. Doyle, R. D. Maclear *et al.*
Diffusion characteristics of hydrogen in diamond.
Diam. Relat. Mater., 8:1611–1614, 1999.
- [94] K. Ogawa, M. Isobe, C. Weiss, E. Griesmayer, S. Sangaroon *et al.*
Fusion product diagnostics based on commercially available chemical vapor deposition
diamond detector in large helical device.
J. Instrum., 18:P01022, 2023.
- [95] Y. Diawara and R. Riedel.
Particle Acceleration and Detection, Neutron Detectors for Scattering Applications - Other
Detectors.
Springer, pages 215–237, 1st Edt. 2023.
- [96] M. Osipenko, M. Ripani, G. Ricco, B. Caiffi, F. Pompili *et al.*
Response of a diamond detector sandwich to 14 MeV neutrons.
Nucl. Instrum. Meth. A., 817:19–25, 2016.
- [97] J. Liu, H. Jiang, Z. Cui, Y. Hu, H. Bai *et al.*
Simultaneous measurement of energy spectrum and fluence of neutrons using a diamond
detector.
Sci. Rep., 12:12022, 2022.
- [98] P. Fichet, A. Bultel, S. Markelj and C. Moreno.
Review of the different techniques to analyse tritium.
TRANSAT Research an Innovation Action, 2018.
- [99] T. Tanabe.
Tritium: Fuel of Fusion Reactors - Characteristics of Tritium.
Springer, pages 27–48, 1st Edt. 2017.
- [100] K. Nakamura and (Particle Data Group).
Review of particle physics.
J. Phys. G, 37:075021, 2010.
- [101] B. Nie.
Fusion reactors.
Elsevier, pages 391–427, 2023.

- [102] P. W. May.
Diamond thin films: a 21st-century material.
Trans. R. Soc. Lond. A, 358:473–495, 2000.
- [103] S. Ohmagari.
Single-crystal diamond growth by hot-filament CVD: a recent advances for doping, growth rate and defect controls.
Func. Diam., 3:2259941, 2023.
- [104] A. P. Bolshakov, V. G. Ralchenko, V. Y. Yurov, G. Shu, E. V. Bushuev *et al.*
Enhanced deposition rate of polycrystalline CVD diamond at high microwave power densities.
Diam. Relat. Mater., 97:107466, 2019.
- [105] P. W. May.
CVD diamond: a new technology for the future?
Endeavour, 19:101–106, 1995.
- [106] M. J. Verstraete and J. Charlier.
Cite as.
Appl. Phys. Lett, 86:191917, 2005.
- [107] P. W. May, M. N. R. Ashfold and Yu. A. Mankelevich.
Microcrystalline, nanocrystalline, and ultrananocrystalline diamond chemical vapor deposition: Experiment and modeling of the factors controlling growth rate, nucleation, and crystal size.
J. Appl. Phys., 101:053115, 2007.
- [108] C. A. Wolden and K. K. Gleason.
On the pressure limits of diamond chemical vapor deposition.
Diam. Relat. Mater., 5:1503–1508, 1996.
- [109] A. Amorim, P. A. P. Nascente, V. J. Trava-Airoldi, E. J. Corat, A. R. Alves *et al.*
Two-step growth of HFCVD diamond films over large areas.
Vacuum, 83:1054–1056, 2009.
- [110] P. Hartmann, R. Haubner and B. Lux.
Characteristics of a pulsed DC-glow discharge CVD reactor for deposition of thick diamond films.
Int. J. Refract. Met. H., 16:207–216, 1998.
- [111] K. Muzyka, J. Sun, T. H. Fereja, Y. Lan, W. Zhang *et al.*

- Boron-doped diamond: current progress and challenges in view of electroanalytical applications.
Anal. Methods, 11:397, 2019.
- [112] Y. Einaga, J. Foord and G. Swain.
Diamond electrodes: Diversity and maturity.
MRS Bull., 39:525–532, 2014.
- [113] X. Wang, X. Shen, F. Sun and B. Shen.
Influence of boron doping level on the basic mechanical properties and erosion behavior of boron-doped micro-crystalline diamond (BDMCD) film.
Diam. Relat. Mater., 73:218–231, 2017.
- [114] M. Degenhardt, G. Aprigliano, H. Schulte-Schrepping, U. Hahn, H. Grabosch *et al.*
CVD diamond screens for photon beam imaging at PETRA III.
J. Phys. Conf. Ser., 425:192022, 2013.
- [115] M.-A. Pinault, J. Barjon, T. Kociniewski, F. Jomard and J. Chevallier.
The n-type doping of diamond: Present status and pending questions.
Physica B., 401-402:51–56, 2007.
- [116] S. Kunuku, M. Ficek, A. Wieloszynska, M. Tamulewicz-Szwajkowska, K. Gajewski *et al.*
Influence of B/N co-doping on electrical and photoluminescence properties of CVD grown homoepitaxial diamond films.
Acs. Sym. Ser., 33:125603, 2021.
- [117] Y. Miyake, T. Kondo, A. Otake, Y. Einaga, T. Tojo *et al.*
Boron- and nitrogen-codoped diamond electrodes for the improved reactivity of electrochemical CO₂ reduction reaction.
ACS Sustain. Chem. Eng., 11 (23):8495–8502, 2023.
- [118] X. J. Hu, Y. G. Shen, X. P. Hao and B. Y. Wang.
The structural properties of B–O codoped diamond films.
Diam. Relat. Mater., 18:210–212, 2009.
- [119] D. Zhang, X. Sun, Y. Zhang, W. Shen, F. Dong *et al.*
Synthetic pathway of shallow n-type donor: Theoretical study of Li and B co-doped diamonds.
Diam. Relat. Mater., 141:110599, 2024.
- [120] A. Hollingsworth, M.Yu. Lavrentiev, R. Watkins, A.C. Davies, S. Davies *et al.*
Comparative study of deuterium retention in irradiated Eurofer and Fe–Cr from a new ion implantation materials facility.

- Nucl. Fusion*, 60:016024, 2019.
- [121] D. Goebel and I. Katz.
Fundamentals of Electric Propulsion: Ion and Hall Thrusters - Ion Thruster Accelerator
Grids.
JPL Space Science and Technology Series, pages 189–241, 2008.
- [122] SPECS Surface Nano Analysis GmbH.
PCS-ECR 2.0, Plasma Cracker Source.
SPECS User Manual, 2013.
- [123] A. Goncharov.
The electrostatic plasma lens.
Rev. Sci. Instrum., 84:021101, 2013.
- [124] S. Ogura and K. Fukutani.
Compendium of Surface and Interface Analysis - Thermal Desorption Spectroscopy.
Springer, pages 719–724, 1st Edt., 2018.
- [125] A. Baron-Wiechec, K. Heinola, J. Likonen, E. Alves, N. Catarino *et al.*
Thermal desorption spectrometry of beryllium plasma facing tiles exposed in the JET
tokamak.
Fusion Eng. Des., 133:135–141, 2018.
- [126] A. M. de Jong and J. W. Niemantsverdriet.
Thermal desorption analysis: Comparative test of ten commonly applied procedures.
Surf. Sci., 233:355–365, 1990.
- [127] A. Oudriss, F. Martin and X. Feaugas.
Mechanics - Microstructure - Corrosion Coupling - Experimental Techniques for Dosage
and Detection of Hydrogen.
Elsevier, pages 245–268, 1st Edt., 2019.
- [128] M. Wilde, S. Ohno, S. Ogura, K. Fukutani and H. Matsuzaki.
Quantification of hydrogen concentrations in surface and interface layers and bulk materi-
als through depth profiling with nuclear reaction analysis.
J. Vis. Exp., 109:53452, 2016.
- [129] M. Wilde and K. Fukutani.
Hydrogen detection near surfaces and shallow interfaces with resonant nuclear reaction
analysis.
Surf. Sci. Rep., 69:196–295, 2014.

- [130] D. Dieumegard, D. Dubreuil and G. Amsel.
Analysis and depth profiling of deuterium with the $D(^3\text{He}, p)^4\text{He}$ reaction by detecting the protons at backward angles.
Nucl. Instrum. Methods, 166:431–445, 1979.
- [131] I. Takagi, T. Kobayashi, Y. Ueyama, H. Moriyama, M. Nakamichi *et al.*
Deuterium diffusion in a chemical densified coating observed by NRA.
J. Nucl. Mater., 386-388:682–684, 2009.
- [132] G. W. Grime.
Encyclopedia of Spectroscopy and Spectrometry - High-Energy Ion Beam Analysis.
Academic Press, pages 103–121, 3rd Edt., 2017.
- [133] A. Orlando, F. Franceschini, C. Muscas, S. Pidkova, M. Bartoli *et al.*
A comprehensive review on raman spectroscopy applications.
Chemosensors, 9:262, 2021.
- [134] B. J. Inkson.
Materials Characterization Using Nondestructive Evaluation (NDE) Methods - Scanning electron microscopy (SEM) and transmission electron microscopy (TEM) for materials characterization.
Woodhead Publishing, pages 17–43, 1st Edt., 2016.
- [135] D. J. Dingley, G. Meaden, D. J. Dingley and A. P. Day.
A review of EBSD: from rudimentary on line orientation measurements to high resolution elastic strain measurements over the past 30 years.
IOP Conf. Ser.: Mater. Sci. Eng., 375:012003, 2018.
- [136] L. E. Franken, K. Grünewald, E. J. Boekema and M. C. A. Stuart.
A technical introduction to Transmission Electron Microscopy for soft-matter: Imaging, possibilities, choices, and technical developments.
Small, 16:1906198, 2020.
- [137] J. A. Hinks.
A review of transmission electron microscopes with in situ ion irradiation.
Nucl. Instrum. Meth. B., 267:3652–3662, 2009.
- [138] K. Bhavyasri, M. Sreshta and R. Swethasri.
Auger Electron Spectroscopy-A Review.
Asian J. Pharm. Res., 7:100–103, 2019.
- [139] N. P. Lockyer, S. Aoyagi, J. S. Fletcher, I. S. Gilmore, P. A. W. van der Heide *et al.*
Secondary Ion Mass Spectrometry.

- Nat. Rev. Methods Primers* , 4:32, 2024.
- [140] Y. Aboura and K. L. Moore.
Nanosims analysis of hydrogen and deuterium in metallic alloys: Artefacts and best practice.
Appl. Surf. Sci., 557:149736, 2021.
- [141] D. W. Brenner, O. A. Shenderova, J. A. Harrison, S. J. Stuart, B. Ni *et al.*
A second-generation reactive empirical bond order (REBO) potential energy expression for hydrocarbons.
J. Phys-condens. Mat., 14:783–802, 2002.
- [142] D. W. Brenner.
Empirical potential for hydrocarbons for use in simulating the chemical vapor deposition of diamond films.
Phys. Rev. B, 42:9458–9471, 1990.
- [143] J. Tersoff.
Empirical interatomic potential for silicon with improved elastic properties.
Phys. Rev. B, 38:9902–9905, 1988.
- [144] J. Tersoff.
Modeling solid-state chemistry: Interatomic potentials for multicomponent systems.
Phys. Rev. B, 39:5566–5568, 1989.
- [145] S. J. Stuart, A. B. Tutein and J. A. Harrison.
A reactive potential for hydrocarbons with intermolecular interactions.
J. Chem. Phys., 112:6472–6486, 2000.
- [146] R. Santamaria.
Molecular Dynamics - Integrating the Equations of Motion.
Springer, pages 359–377, 1st Edt., 2023.
- [147] Copyright 2003-2024 Sandia Corporation.
LAMMPS Manual - Commands - fix command.
<https://docs.lammps.org/fix.html>, Accessed 6/11/24.
- [148] D. Frenkel and B. Smit.
Understanding Molecular Simulation - Molecular Dynamics in Various Ensembles.
Academic Press, pages 139–163, 2nd Edt., 2002.
- [149] D. J. Evans and B. L. Holian.
The Nose–Hoover thermostat.
J. Chem. Phys., 83:4069–4074, 1985.

- [150] J. A. Pittard, N. A. Fox, A. Hollingsworth, M. Y. Lavrentiev, A. Wohlers *et al.*
Deuterium retention in CVD diamond: Combined experimental and computational study.
Fusion Eng. Des., 188:113403, 2023.
- [151] Y. Zayachuk, I. Jecu, M. Zlobinski, C. Porosnicu, N. Catarino *et al.*
Fuel desorption from JET-ILW materials: assessment of analytical approach and identification of sources of uncertainty and discrepancy.
Nucl. Fusion, 63:096010, 2023.
- [152] P.J. Sellin, A. Lohstroh, A.W. Davies, A. Galbiati, J. Parkin *et al.*
Charge transport in polycrystalline and single crystal synthetic diamond using ion beam induced charge imaging.
Nucl. Instrum. Meth. B., 260:293–298, 2007.
- [153] C. G. Ña Rosales and M. Balden.
Chemical erosion of doped graphites for fusion devices.
J. Nucl. Mater., 290-293:173–179, 2001.
- [154] V. Mortet, Z. Vlčková Živcová, A. Taylor, O. Frank, P. Hubík *et al.*
Insight into boron-doped diamond Raman spectra characteristic features.
Carbon, 115:279–284, 2017.
- [155] V. Mortet, Z. Vlčková Živcová, A. Taylor, O. Frank, P. Hubík *et al.*
Insight into boron-doped diamond Raman spectra characteristic features.
Carbon, 115:279–284, 2017.
- [156] D. L. Emiliano, F. Álvarez, E. Burgos, E. Halac, H. Huck *et al.*
Stress analysis on single-crystal diamonds by Raman spectroscopy 3D mapping.
Mater. Sci. Appl., 4:191–197, 2013.
- [157] X.H. Wang, G.-H.M. Ma, W. Zhu, J.T. Glass, L. Bergman *et al.*
Effects of boron doping on the surface morphology and structural imperfections of diamond films.
Diam. Relat. Mater., 1:828–835, 1992.
- [158] J. Cifre, J. Puigdollers, M.C. Polo and J. Esteve.
Trimethylboron doping of CVD diamond thin films.
Diam. Relat. Mater., 3:628–631, 1994.
- [159] P. Hartmann, S. Bohr, R. Haubner, B. Lux, P. Wurzinger *et al.*
Diamond growth with boron addition.
Int. J. Refract. Met. H., 16:223–232, 1998.

- [160] C. Torregrosa-Martin, A. Ibarra, J. Aguilar, F. Ambi, F. Arranz *et al.*
Overview of IFMIF-DONES diagnostics: Requirements and techniques.
Fusion Eng. Des., 191:113556, 2023.
- [161] S.J. Jepeal, L. Snead and Z.S. Hartwig.
Intermediate energy proton irradiation: Rapid, high-fidelity materials testing for fusion
and fission energy systems.
Mater. Design, 200:109445, 2021.
- [162] C. W. Lee, Y. O. Lee, D. W. Lee, S. Cho and M. Y. Ahn.
Effect of graphite reflector on activation of fusion breeding blanket.
Fus. Eng. Des., 109–111:503–507, 2016.
- [163] J. F. Ziegler, M.D. Ziegler and J.P. Biersack.
SRIM – The stopping and range of ions in matter (2010).
Nucl. Instrum. Meth. B., 268:1818–1823, 2010.
- [164] Z. Yang, Q. Xu, R. Hong, Q. Li and G. N. Luo.
Molecular dynamics simulation of low-energy atomic hydrogen on tungsten surface.
Fusion Eng. Des., 85:1517–1520, 2010.
- [165] R. E. Stoller, M. B. Toloczko, G. S. Was, A. G. Certain, S. Dwaraknath *et al.*
On the use of SRIM for computing radiation damage exposure.
Nucl. Instrum. Meth. B., 310:75–80, 2013.
- [166] R.L. Shepherd, D.R. Kania, L.A. Jones, D.H. Schneider and R.E. Stewart.
Strongly Coupled Plasma Physics - The Measurement of Transport Properties in Strongly
Coupled Plasmas.
Elsevier, :433–437, 1st Edt., 1990.
- [167] G. Carter, M. J. Nobes, F. Paton, J. S. Williams and J. L. Whitton.
Ion bombardment induced ripple topography on amorphous solids.
Radiat. Eff., 33:65–73, 1977.
- [168] A. Datta, Y. R. Wu and Y. L. Wang.
Real-time observation of ripple structure formation on a diamond surface under focused
ion-beam bombardment.
Phys. Rev. B, 63:125407, 2001.
- [169] M. Schwander and K. Partes.
A review of diamond synthesis by CVD processes.
Diam. Relat. Mater., 20:1287–1301, 2011.

- [170] C. Su and J.-C. Lin.
Thermal desorption of hydrogen from the diamond C(100) surface.
Surf. Sci., 406:149–166, 1998.
- [171] A. Hoffman, K. Bobrov, B. Fisgeer, H. Shechter and M. Folman.
Effects of deuterium adsorption-desorption on the state of diamond: Surface degradation and stabilization of sp^3 bonded carbon.
Diam. Relat. Mater., 5:977–983, 1996.
- [172] H. Kimura, M. Sasaki, Y. Morimoto, T. Takeda, H. Kodama *et al.*
Thermal desorption behavior of deuterium implanted into polycrystalline diamond.
J. Nucl. Mater., 337-339:614–618, 2005.
- [173] D. Ballutaud, F. Jomard, B. Theys, C. Mer, D. Tromson *et al.*
Hydrogen diffusion and stability in polycrystalline CVD undoped diamond.
Diam. Relat. Mater., 10:405–410, 2001.
- [174] A. Hollingsworth, M.-F. Barthe, M. Yu Lavrentiev, P.M. Derlet, S.L. Dudarev *et al.*
Comparative study of deuterium retention and vacancy content of self-ion irradiated tungsten.
J. Nucl. Mater., 558:153373, 2022.
- [175] M. Yu. Lavrentiev, A. Hollingsworth, J. Hess, S. Davies, A. Wohlers *et al.*
Effects of self-irradiation on deuterium retention and reflectivity of molybdenum, fusion plasma-facing material: Combined experimental and modeling study.
J. Appl. Phys., 132:125902, 2022.
- [176] E.B. Lombardi.
Boron–hydrogen complexes in diamond and their vibrational properties.
Diam. Relat. Mater., 18:835–838, 2009.
- [177] Z. Teukam, D. Ballutaud, F. Jomard, J. Chevallier, M. Bernard *et al.*
Trap limited diffusion of hydrogen in boron-doped diamond.
Diam. Relat. Mater., 12:647–651, 2003.
- [178] H. Sawada, H. Ichinose and M. Kohyama.
Atomic structure of the σ_3 and σ_9 grain boundaries in CVD diamond film.
Scripta Mater., 51:689–692, 2004.
- [179] J.A. Pittard, M.Y. Lavrentiev and N.A. Fox.
Simulated bombardment of diamond with hydrogen isotopes.
IEEE T. Plasma Sci., 52:3674–3679, 2024.

- [180] J. A. Pittard, M. Y. Lavrentiev and N. A. Fox.
Simulated hydrogen diffusion in diamond grain boundaries.
Diam. Relat. Mater., 149:111665, 2024.
- [181] Copyright 2003-2024 Sandia Corporation.
LAMMPS Manual - Commands - fix command.
<https://docs.lammps.org/fix.html>, Accessed 6/11/24.
- [182] A. Stukowski.
Visualization and analysis of atomistic simulation data with OVITO - the Open Visualization Tool.
Model. Simul. Mater. Sc., 18:015012, 2009.
- [183] G. Federici, J. N. Brooks, M. Iseli and C. H. Wu.
In-vessel tritium retention and removal in ITER-FEAT.
Phys. Scripta, T91:76, 2001.
- [184] T. Senftle, S. Hong, M. M. Islam, S. B. Kylasa, Y. Zheng *et al.*
The reaxff reactive force-field: development, applications and future directions.
npj Comput. Mater., 2:15011, 2016.
- [185] M. Tolladay, F. Scarpa and N. L. Allan.
Interatomic forces breaking carbon-carbon bonds.
Carbon, 175:420–428, 2021.
- [186] W. Cai and W. D. Nix.
Imperfections in Crystalline Solids - Grain Boundary Geometry.
Cambridge University Press, pages 433–454, 1st Edt., 2016.
- [187] P. Zapol, M. Sternberg, L. A. Curtiss, T. Frauenheim and D. M. Gruen.
Tight-binding molecular-dynamics simulation of impurities in ultrananocrystalline diamond grain boundaries.
Phys. Rev. B, 65:454031–4540311, 2002.
- [188] P. Steneteg, V. Chirita, N. Dubrovinskaia, L. Dubrovinsky and I. A. Abrikosov.
Missing-atom structure of diamond $\Sigma 5$ (001) twist grain boundary.
Phys. Rev. B, 84:144112, 2011.
- [189] S. Von Althan, P. D. Haynes, K. Kaski and A. P. Sutton.
Are the structures of twist grain boundaries in silicon ordered at 0 K?
Phys. Rev. Lett., 96:055505, 2006.
- [190] T. Watanabe, B. Ni, S. R. Phillpot, P. K. Schelling and P. Keblinski.

- Thermal conductance across grain boundaries in diamond from molecular dynamics simulation.
J. Appl. Phys., 102:063503, 2007.
- [191] Copyright 2003-2024 Sandia Corporation.
 LAMMPS Manual.
<https://docs.lammps.org/>, Accessed 13/08/24.
- [192] A. Manhard, T. Schwarz-Selinger and W. Jacob.
 Quantification of the deuterium ion fluxes from a plasma source.
Plasma Sources Sci. T., 20:015010, 2011.
- [193] A. H. Simon.
 Handbook of Thin Film Deposition - Sputtering Process.
 William Andrew Publishing, pages 55–88, 3rd Edt., 2012.
- [194] A. Deslandes, M. C. Guenette, C. S. Corr, I. Karatchevtseva, L. Thomsen *et al.*
 Deuterium retention and near-surface modification of ion-irradiated diamond exposed to fusion-relevant plasma.
Nucl. Fusion, 54:073003, 2014.
- [195] OVITO User Manual 3.8.4 documentation 2023.
 Wigner-Seitz defect analysis.
 OVITO GmbH, Germany, 2023, accessed 15/05/23.
- [196] J. C. Bourgoin.
 An experimental estimation of the vacancy formation energy in diamond.
Radiat. Eff., 79:235–239, 1983.
- [197] T. Halicioglu.
 Calculation of surface energies for low index planes of diamond.
Surface Science Letters, 259:L714–L718, 1991.
- [198] J. Zhang F. Ma K. Xu X. Xin.
 Anisotropy analysis of the surface energy of diamond cubic crystals .
Surf. Interface Anal., 35:805–809, 2003.
- [199] D. Delgado and R. Vila.
 Statistical Molecular Dynamics study of displacement energies in diamond .
J. Nucl. Mater. , 419:32–38, 2011.
- [200] B. B. Bokhonov, D. V. Dudina and M. R. Sharafutdinov.
 Graphitization of synthetic diamond crystals: A morphological study.
Diam. Relat. Mater., 118:108563, 2021.

- [201] C. Baruffi and C. Brandl.
On the structure of (111) twist grain boundaries in diamond: atomistic simulations with Tersoff-type interatomic potentials.
Acta Mater., 215:117055, 2021.
- [202] Pooja, R. Mucherla and R. Pawar.
Envisioning the hydrogen dissociation in $\Sigma 5$ (100) grain boundary in diamond.
Diam. Relat. Mater., 138:110222, 2023.
- [203] O. A. Shenderova, D. W. Brenner and L. H. Yang.
Atomistic simulations of structures and mechanical properties of polycrystalline diamond: Symmetrical $\langle 001 \rangle$ tilt grain boundaries.
Phys. Rev. B, 60:7043–7052, 1999.
- [204] W. Luyten, G. Van Tendeloo, S. Amelinckx and J. L. Collins.
Electron microscopy study of defects in synthetic diamond layers.
Philos. Mag. A, 66:899–915, 1992.
- [205] O. R. Monteiro.
Structural defects and sp^2 localization in CVD diamond.
J. Mater. Sci., 54:2300–2306, 2019.
- [206] A. H. King, F.-R. Chen, L. Chang and J. J. Kai.
Toward understanding polycrystalline aggregate structure: Analysis of a twin intersection and the interactions between interfaces in diamond.
Interface Sci., 5:287–303, 1997.
- [207] O. A. Shenderova and D. W. Brenner.
Atomistic simulations of structures and mechanical properties of $\langle 011 \rangle$ tilt grain boundaries and their triple junctions in diamond.
Phys. Rev. B, 60:7053–7061, 1999.
- [208] D. Shechtman, J.L. Hutchison, L.H. Robins, E.N. Farabaugh and A. Feldman.
Growth defects in diamond films.
J. Mater. Res., 8:473–479, 1993.
- [209] A. E. Mora, J. W. Steeds and J. E. Butler.
Grain boundaries in boron-doped CVD diamond films.
Diam. Relat. Mater., 11:697–702, 2002.
- [210] L. Li, M. Xu, W. Song, A. Ovcharenko, G. Zhang *et al.*
The effect of empirical potential functions on modeling of amorphous carbon using molecular dynamics method.

- Appl. Surf. Sci.*, 286:287–297, 2013.
- [211] C. Onwudinanti, M. Pols, G. Brocks, V. Koelman, A. C. T. van Duin *et al.*
A ReaxFF Molecular Dynamics study of hydrogen diffusion in ruthenium - the role of grain boundaries.
J. Phys. Chem. C, 126:5950–5959, 2022.
- [212] H. Ichinose and M. Nakanose.
Atomic and electronic structure of diamond grain boundaries analyzed by HRTEM and EELS.
Thin Solid Films, 319:87–91, 1998.
- [213] M. Sasaki, Y. Morimoto, H. Kimura, K. Takahashi, K. Sakamoto *et al.*
Energetic deuterium and helium irradiation effects on chemical structure of CVD diamond.
J. Nucl. Mater., 329-333:899–903, 2004.
- [214] H. Kimura, M. Sasaki, Y. Morimoto, T. Takeda, H. Kodama *et al.*
Thermal desorption behavior of deuterium implanted into polycrystalline diamond.
J. Nucl. Mater., 337-339:614–618, 2005.
- [215] N. Nishikawa.
Evidence of hydrogen termination at grain boundaries in ultrananocrystalline diamond/hydrogenated amorphous carbon composite thin films synthesized via coaxial arc plasma.
J. Vac. Sci. Technol. B, 38:062803, 2020.
- [216] D. Berman, S. A. Deshmukh, B. Narayanan, S. K. R. S. Sankaranarayanan, Z. Yan *et al.*
Metal-induced rapid transformation of diamond into single and multilayer graphene on wafer scale.
Nat. Commun., 7:12099, 2016.
- [217] P. J. Fallon and L. M. Brown.
Analysis of chemical-vapour-deposited diamond grain boundaries using transmission electron microscopy and parallel electron energy loss spectroscopy in a scanning transmission electron microscope.
Diam. Relat. Mater., 2:1004–1014, 1993.
- [218] S. Rubanov, B. A. Fairchild, A. Suvorova, P. Olivero and S. Prawer.
Structural transformation of implanted diamond layers during high temperature annealing.
Nucl. Instrum. Meth. B, 365:50–54, 2015.

- [219] T. Liu, T. Shoa, F. Lyu, X. Lai and A. H. Shen.
Molecular dynamics simulations to assess the radiation resistance of different crystal orientations of diamond under neutron irradiation.
Model. Simul. Mater. Sci. Eng., 30:035005, 2022.
- [220] Z. Harutyunyan, Y. Gasparyan, V. S. Efimov, S. Ryabtsev and A. Pisarev.
Retention of deuterium in the surface layers of tungsten preliminarily irradiated with helium ions.
Bull. Russ. Acad. Sci. Phys., 84:727–731, 2020.
- [221] A. Kreter, C. Brandt, A. Huber, S. Kraus, S. M. O. Ller *et al.*
Linear plasma device PSI-2 for plasma-material interaction studies.
Fusion Sci. Technol., 68:8–14, 2015.
- [222] E. Malitckii, Y. Yagodzinskyy and H. Hänninen.
Hydrogen charging process instrument.
Fusion Eng. Des., 100:142–145, 2015.
- [223] T. Hino, Y. Yamauchi, Y. Kimura, K. Nishimura and Y. Ueda.
Fuel hydrogen retention of tungsten and the reduction by inert gas glow discharges.
Fusion Eng. Des., 87:876–879, 2012.
- [224] S. Ishikawa, K. Katayama, Y. Ohnishi, S. Fukada and M. Nishikawa.
Sorption and desorption behavior of hydrogen isotopes from tungsten deposits caused by deuterium gas or deuterium plasma exposure.
Fusion Eng. Des., 87(7):1390–1394, 2012.
- [225] I. Korolov and Z. Donkó.
Breakdown in hydrogen and deuterium gases in static and radio-frequency fields.
Phys. Plasmas, 22:093501, 2015.
- [226] J. M. Meek and J. D. Craggs.
Electrical breakdown of gases.
Oxford At the Clarendon Press, 1978.
- [227] W. R. Carr.
On the laws governing electric discharges in gases at low pressures.
Philos. Trans. R. Soc. London Ser. A, 201:403, 1903.
- [228] V. Stojanović, J. Boin, Z. Lj. Petrović and B. M. Jelenković.
Excitation of the $C^3\Pi_u$ state of N_2 and $B^2\Sigma_g^+$ state of N_2^+ by electrons at moderate and high values of E/N.
Phys. Rev. A, 42:4983–4992, 1990.

- [229] S. Githens.
The influence of discharge chamber structure upon the initiating mechanism the high frequency discharge.
Phys. Rev., 57:822–828, 1940.
- [230] R. J. Armstrong and T. K. Bennett.
Electric breakdown in deuterium and hydrogen at low pressures.
J. Appl. Phys., 82:2147–2149, 1997.
- [231] D. Xiao.
Gas Discharge and Gas Insulation - Fundamental Theory of Townsend Discharge.
Springer, pages 47–88, 1st Edt., 2016.
- [232] T. Kalvas, O. Tarvainen, T. Ropponen, O. Steczkiewicz, J. Rje *et al.*
IBSIMU: A three-dimensional simulation software for charged particle optics.
Rev. Sci. Instrum., 81:02B703, 2010.
- [233] D. A. Dahl.
SIMION for the personal computer in reflection.
Int. J. Mass Spectrom., 200:3–25, 2000.
- [234] E. H. Higham and J. M. Paros.
Instrumentation Reference Book - Measurement of Pressure.
Butterworth-Heinemann, pages 145–163, 3rd Edt., 2010.
- [235] A. S. Morris and R. Langari.
Measurement and Instrumentation - Pressure Measurement.
Butterworth-Heinemann, pages 463–491, 2nd Edt., 2016.
- [236] PyVISA Documentation.
<https://pyvisa.readthedocs.io/en/latest/>.
Copyright 2024 PyVISA Authors, Accessed 2024-12-05.
- [237] V. I. Kolobov.
Striations in rare gas plasmas.
J. Phys. D. Appl. Phys., 39:R487, 2006.
- [238] M. M. Tahiyat, J. C. Stephens, V. I. Kolobov and T. I. Farouk.
Striations in moderate pressure DC driven nitrogen glow discharge.
J. Phys. D. Appl. Phys., 55:085201, 2022.
- [239] D. Gahan, S. Daniels, C. Hayden, P. Scullin, D. O’Sullivan *et al.*
Ion energy distribution measurements in RF and pulsed DC plasma discharges.
Plasma Sources Sci. Technol., 21:024004, 2012.

- [240] M. Poon, R. G. Macaulay-Newcombe, J. W. Davis and A. A. Haasz.
Flux dependence of deuterium retention in single crystal tungsten.
J. Nucl. Mater., 307–311:723–728, 2002.
- [241] K. Tokunaga, M. J. Baldwin, R. P. Doerner, N. Noda, Y. Kubota *et al.*
Blister formation and deuterium retention on tungsten exposed to low energy and high
flux deuterium plasma.
J. Nucl. Mater., 337–339:887–891, 2005.
- [242] Y. Zayachuk, A. Manhard, M. H. J. 't Hoen, W. Jacob, P. A. Zeijlmans van Emmichoven *et al.*
The effect of ion flux on plasma-induced modification and deuterium retention in tungsten
and tungsten–tantalum alloys.
J. Nucl. Mater., 464:69–72, 2015.
- [243] M. H. J. 't Hoen, M. Balden, A. Manhard, M. Mayer, S. Elgeti *et al.*
Surface morphology and deuterium retention of tungsten after low- and high-flux deu-
terium plasma exposure.
Nucl. Fusion, 54:083014, 2014.
- [244] L. Buzi, G. De Temmerman, B. Unterberg, M. Reinhart, A. Litnovsky *et al.*
Influence of particle flux density and temperature on surface modifications of tungsten and
deuterium retention.
J. Nucl. Mater., 455:316–319, 2014.

