

NORTHWESTERN UNIVERSITY

VACANCY CLUSTERING PHENOMENA IN SILICON CARBIDE REACTOR
COMPONENTS: A MULTIPLE TIME-SCALE ATOMISTIC PROBLEM

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Mechanical Engineering

By

David E. Farrell

EVANSTON, ILLINOIS

December 2008

© Copyright by David E. Farrell 2008

All Rights Reserved

ABSTRACT

Vacancy Clustering Phenomena in Silicon Carbide Reactor Components:
A Multiple Time-scale Atomistic Problem

David E. Farrell

The goal of this dissertation was to elucidate the coalescence of vacancies in the cubic polytype of silicon carbide in a nuclear reactor relevant system. Large-scale empirical potential atomistic simulations were employed and the exploration of long time-scale behavior was performed through a new application of the parallel replica dynamics method. To differentiate this dissertation from earlier work, emphasis has been placed on the determination of a predictive description of initial damage states, the study of the impact of pre-existing voids, and the exploration of long-time behavior of the damage configuration. Further, the methodology employed here is sufficiently general to be applied to any material.

From the simulations performed, several major conclusions were drawn. The simulations of 10 keV silicon primary knock-on atom cascades in a perfect silicon carbide crystal showed good qualitative agreement with previous simulations. No large clusters formed and most vacancies were isolated in the final state, in contrast to some similar studies in metals. Simulations at 1200 K yielded an approximate spatial distribution of vacancies, a completely

new result for silicon carbide. The temperature effect study indicated the existence of a high temperature defect relaxation regime at 2000 K that reduced the damage caused by the cascade. After the addition of the initial void structures to the 10 keV cascade studies, the results did not exhibit significant void growth or structural change. It is suspected that boundary conditions contributed to this effect. In both the perfect crystal and initial void studies, it was concluded that a dramatically larger timescale or multiple irradiation events would be needed to observe formation of large vacancy clusters. Parallel replica dynamics calculations were performed that demonstrated the major pathways for vacancy migration in cubic silicon carbide, in agreement with previous *ab initio* studies. This was the first application of parallel replica dynamics to a cubic silicon carbide system with a large number of migrating vacancies. It was concluded that the computational requirements for a long time calculation of the full cascade geometry would be prohibitive even on a large modern supercomputer.

ACKNOWLEDGMENTS

I wish to express my sincere gratitude to my advisor Professor Wing Kam Liu for his support, encouragement, and guidance throughout my graduate studies. I want to extend thanks to the members of my defense committee, Dr. Noam Bernstein, Prof. Greg Olson and Prof. Chris Wolverton, for taking the time to read this dissertation and provide their comments. I would also like to thank those who served as mentors to me through my undergraduate and graduate studies: Professor Kathleen Issen, Professor Harold Park, Professor Eduard Karpov, Dr. Patrick Klein and Dr. Noam Bernstein. Without these people, I would not have the background this work has required. Dr. Bernstein also was kind enough to provide the more than 27,000 CPU-hours of supercomputer time used for the work presented here.

I wish to thank my wife Melissa and our friends. These important people provided me much needed emotional support and relaxing outlets that while possibly reducing my time to work on this dissertation, assured that I remained nominally sane while doing so. Thanks also go to my lab-mates who were ‘in the trenches’ with me through my graduate studies, as well as the indispensable help of the Mechanical Engineering Department staff. Finally, I would like to extend a thank you to my parents, Carolyn and Joseph, and brothers, James and Sean, who have been a great source of encouragement and support through the years.

Contents

List of Figures	9
Nomenclature	12
1 Introduction	14
1.1 Radiation Basics	19
1.1.1 Useful Units and Terms	21
1.1.2 Types of Radiation	27
1.1.2.1 Electromagnetic Radiation	28
1.1.2.2 Particle Radiation	29
1.1.3 Particle Radiation Damage	30
1.2 Overview of Nuclear Power Generation	35
1.2.1 Fission Reactors	35
1.2.2 Fusion Reactors	37
1.3 Materials for Nuclear Applications	39
1.3.1 Metallic Materials	41
1.3.2 Ceramics and Fiber-reinforced Composites	42
1.3.3 Silicon Carbide (SiC) as an Engineering Material	45
1.3.4 SiC Manufacturing Techniques	51
1.4 Focus Application: Void Formation in SiC-based Reactor Components	53
2 Review of Relevant Short-Time Simulation Methods	57
2.1 The Binary Collision Approximation	58
2.2 Empirical Potential Molecular Dynamics (EP)	63
2.2.1 The Tersoff Potential	66
2.2.2 EP Initial/Boundary Conditions and Solution Procedures	68
2.3 Post-Processing and Data Mining EP Simulations for Defect Information	70
2.3.1 PKA Kinetic Energy and Displacement	72
2.3.2 Point Defect Identification	72
2.3.3 Damage Configuration Visualization	76

2.3.4	Defect Clustering and Distributions	77
2.4	Essential Density Functional Theory (DFT)	77
3	Review and Verification of Displacement Cascade Simulation Approach	84
3.1	Relevant Literature	85
3.2	Simulation Approaches and Verification Cases	94
3.2.1	Threshold Displacement Energy Simulations	95
3.2.1.1	System Setup	96
3.2.1.2	Results and Discussion	99
3.2.2	10 keV Si PKA Cascade Verification Case	102
3.2.2.1	System Setup	103
3.2.2.2	Results and Discussion	106
3.2.2.3	Point Defect Counts	107
3.3	Conclusions	113
4	Characterization of Initial Defect Distribution in a Neutron-Irradiated Perfect Crystal	114
4.1	10 keV Si PKA Cascade in 3C-SiC Perfect Crystal	115
4.1.1	PKA Displacement and Energy Evolution	117
4.1.2	Damage State Visualization	119
4.1.3	Defect Counts	125
4.1.4	Vacancy Distribution and Clustering	129
4.2	Temperature Effects in 10 keV Si PKA Cascade in 3C-SiC	135
4.2.1	PKA Displacement	136
4.2.2	Point Defect Counts	137
4.2.3	Vacancy Clusters and Distributions	145
4.3	Conclusions	152
5	Characterization of Initial Defect Distribution in a Neutron-Irradiated Crystal with Initial Void Structure	155
5.1	Simulation Configurations and Methodology	157
5.2	Effects of an Initial Void on Defect Configuration in 10 keV Si PKA Cascade in 3C-SiC	162
5.3	Effects of Initial Void Arrays on Defect Configuration	173
5.3.1	3 Void Linear Array	175
5.3.2	27 Void Cubic Array	180
5.4	Conclusions	186

6	Long-Time Behavior	190
6.1	Long-Time Simulation Techniques	191
6.1.1	Kinetic Monte Carlo (KMC)	191
6.1.2	Parallel Replica Dynamics (PRD)	194
6.2	Parallel Replica Dynamics Implementation	199
6.3	Parallel Replica Dynamics Example	201
6.4	Discussion of Parallel Replica Dynamics for Cascade Simulations	205
6.5	Conclusions	209
7	Summary and Conclusions	210
	Bibliography	218
A	Appendices	229
A.1	LAMMPS Tutorials	229
A.1.1	Threshold Displacement Energy Example Input	231
A.1.2	Cascade Simulation Example Input	232
A.1.3	Parallel Replica Dynamics Input Example	235
A.1.4	LAMMPS PRD Command Function	238
A.2	Post-processing Scripts, Codes	243
A.2.1	Point Defect Identification Scripts	243
A.2.2	LAMMPS *.dump-EnSight Conversion and Paraview Scripts	246

List of Figures

1.1	NRC map of reactor locations	15
1.2	NRC map of nuclear waste storage sites	16
1.3	DOE schematic of proposed Yucca Mountain Repository	17
1.4	Time and length-scales in nuclear engineering	18
1.5	Schematic definition of the steradian	23
1.6	Schematic definition of the threshold displacement energy	26
1.7	Displacement damage schematic	32
1.8	Uranium-235 fission reaction schematic	36
1.9	Tritium-Deuterium fusion reaction schematic	38
1.10	First wall and blanket placement in a Tokamak type fusion reactor	40
1.11	Micrograph of TRISO MPBR fuel pellet	44
1.12	Schematic of s, p electronic orbitals and sp^3 hybridization	47
1.13	Schematic of 3C-SiC unit cell	48
1.14	Schematic of 3C-SiC polytype stacking	49
1.15	Schematic of 2H-SiC polytype stacking	50
2.1	Schematic of classical binary collision	60
2.2	Schematic of point defect identification algorithms	74
3.1	Half-section of threshold displacement energy calculation geometry	97
3.2	Schematic of Devanathan et al cascade geometry	104
3.3	Devanathan et al reproduction PKA kinetic energy vs. time	108
3.4	Devanathan et al reproduction PKA displacement vs. time	108
3.5	Devanathan et al reproduction total vacancies & interstitials vs. time	109
3.6	Devanathan et al reproduction C & Si vacancies vs. time	110
3.7	Devanathan et al reproduction atoms with kinetic energy in 21-42 eV range	111
4.1	Schematic of perfect crystal cascade geometry	117
4.2	Perfect crystal PKA kinetic energy vs. time examples	120
4.3	Perfect crystal PKA displacement vs. time examples	121
4.4	Perfect crystal final PKA displacements	122

4.5	Perfect crystal final defect configuration examples	124
4.6	Perfect crystal run averaged defect counts	126
4.7	Perfect crystal run averaged antisites and replacements	127
4.8	Perfect crystal run averaged vacancy cluster size	130
4.9	Perfect crystal run averaged vacancy radial distribution	132
4.10	Perfect crystal run averaged vacancy spatial distribution	134
4.11	Percentage of PKA wrap-around vs. temperature	137
4.12	PKA displacement vs. temperature	138
4.13	Total vacancies vs. time and temperature	139
4.14	C and Si vacancies vs. time and temperature	141
4.15	Antisite defects vs. time and temperature	142
4.16	Atomic replacements vs. time and temperature	143
4.17	Vacancy cluster number vs. time and temperature	145
4.18	Vacancy cluster size vs. time and temperature	146
4.19	Vacancy clustering percentage vs. time and temperature	147
4.20	Vacancy RDF for each temperature, final	149
4.21	Vacancy RDF for each temperature, initial	150
4.22	Vacancy distribution for each temperature	151
5.1	Schematic of cascade base geometry	158
5.2	Void effect study configurations	160
5.3	Single void final PKA displacements	164
5.4	Single void final defect configuration examples	165
5.5	Single central void run averaged defect counts	167
5.6	Single near-surface void run averaged defect counts	168
5.7	Single central void cluster information	170
5.8	Single near-surface void cluster information	171
5.9	Single void crystal cluster percentages	172
5.10	Single void crystal phase I RDFs	174
5.11	3 void linear array PKA displacements	176
5.12	3 void linear array defect configuration examples	177
5.13	3 void linear array run averaged defect counts	179
5.14	3 void array RDFs	181
5.15	27 void cubic array PKA displacements	182
5.16	27 void cubic array defect configuration examples	183
5.17	27 void cubic array run averaged defect counts	185
5.18	27 void array RDFs	187
6.1	Schematic of the Parallel Replica Dynamics algorithm	196
6.2	PRD example case transitions	202

6.3 PRD example case final configuration 204

Nomenclature

3C-SiC cubic polytype of SiC, discussed in Section 1.3.3

Å Angstrom, 10^{-10} m

amu Atomic Mass Unit

as attosecond, 10^{-18} s

BCA Binary Collision Approximation

BCC Body-centered cubic crystal structure

C Carbon element symbol

Cu Copper element symbol

CVD Chemical Vapor Deposition

DFT Density Functional Theory

DPA Displacements Per Atom

EP Classical Empirical Potential Molecular Dynamics. Used to avoid ambiguity when discussing classical and quantum mechanical molecular dynamics methods.

eV Electron-volt, $1.60217653 \times 10^{-19}$ Joules

FCC Face-centered cubic crystal structure

Fe Iron element symbol

fs femtosecond, 10^{-15} s

KMC Kinetic Monte Carlo method

MD	Molecular Dynamics, most often used to denote classical empirical potential Molecular Dynamics.
MPBR	Modular Pebble Bed Reactor
Ni	Nickel element symbol
PBC	Periodic Boundary Condition
PKA	Primary Knock-on Atom
PRD	Parallel Replica Dynamics
ps	picosecond, 10^{-12} s
RDF	Radial Distribution Function
Si	Silicon element symbol
SiC	Silicon Carbide compound symbol
TB	Tight-Binding Molecular Dynamics
TDE	Threshold Displacement Energy
ZBL	Coulomb repulsive potential with Ziegler, Biersack and Littmark universal screening function. See Equation 2.9
Zr	Zirconium element symbol

Chapter 1

Introduction

Nuclear energy is not a new idea in the U.S., however, many of the nuclear power plants currently in operation were designed and built in the 1960s and 1970s. This is mainly due to the decreased cost of fossil-fuel based energy in the 1980s, which is purported to have resulted in a decrease in U.S. interest in nuclear energy. This decrease in interest was compounded by public fear caused by nuclear accidents, such as the Three-mile Island incident of March 1978 in Pennsylvania and the Chernobyl incident of April 1986 in the Ukraine. Finally, political support has been reduced because of the environmental and security issues related to the large number of nuclear reactor and waste storage sites currently licensed by the U.S. Nuclear Regulatory Commission (NRC). At present there are 33 test reactor and 104 commercial reactor sites as well as over 120 radioactive waste storage sites in 39 states across the U.S. (see Figures 1.1 and 1.2). Over the last few decades these factors,

coupled with decreased federal and commercial support, have led to decreased enrollments in university nuclear engineering programs and even department closures. This has resulted in a reduction in the pool of available expertise and facilities for materials design for nuclear reactor applications[1]. Despite this historical trend, nuclear energy and thus materials design for nuclear applications have recently experienced renewed interest in the United States. This can be attributed to concerns over other environmental problems, such as greenhouse gas emission, and the dependence of critical infrastructure on imported fossil fuels [2].

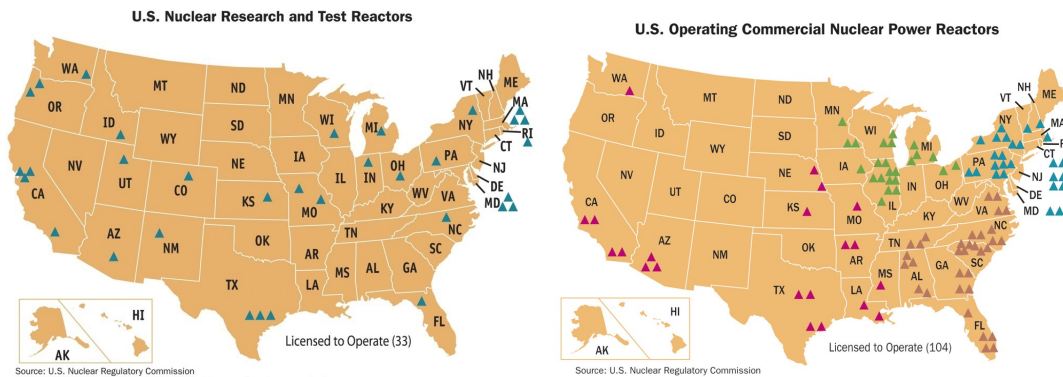


Figure 1.1: Nuclear Regulatory Commission map of U.S. test, research and commercial nuclear reactor sites. Images courtesy of the U.S. Nuclear Regulatory Commission.

The renewal of interest in nuclear power has spurred the development of a number of programs to encourage the development of new reactor designs[3]. Examples include the U.S. Department of Energy ‘Nuclear Power 2010 Program’, ‘Generation IV Initiative’ [4] and the multinational ‘Generation IV International Forum’[5, 6] which focus on further

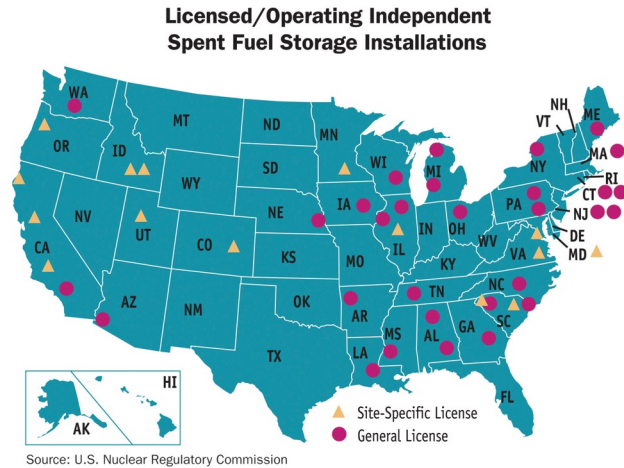


Figure 1.2: Nuclear Regulatory Commission map of U.S. nuclear waste storage sites. Image courtesy of the U.S. Nuclear Regulatory Commission.

development of fission-based energy generation capabilities. In addition to these reactor related initiatives, there is the proposed project to build a long-term radioactive waste storage facility at Yucca Mountain, Nevada (see Figure 1.3)[7], currently awaiting licensing approval. However, safety and the high costs (on the order of 10 billion dollars) associated with the licensing, construction and maintenance of a nuclear power plant remain major roadblocks to further development in the nuclear power sector [8]. One key area of research that will be necessary for nuclear power to be a competitive energy source is in basic materials science; in particular how defects are generated by particle irradiation and their subsequent behavior[1, 9].

From an engineering perspective, nuclear reactors and repositories are a good example of a complex multiscale system[2, 10]. The lifetime of a reactor facility is measured in decades and

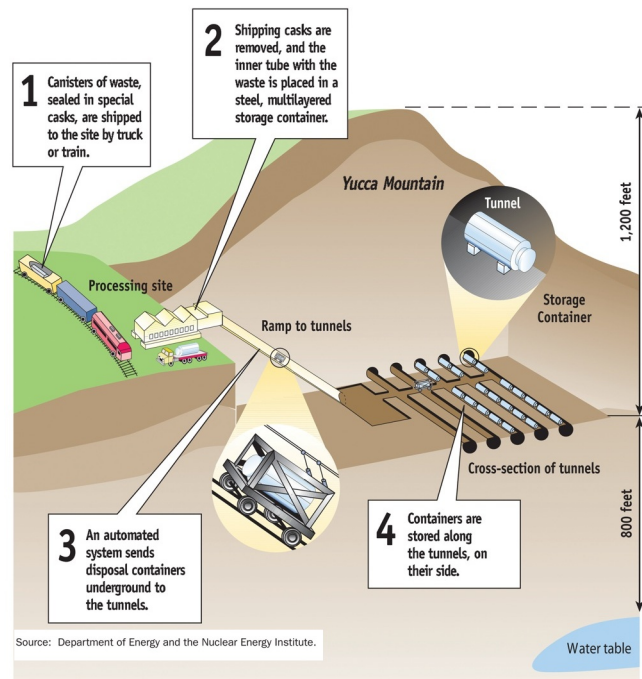


Figure 1.3: Department of Energy schematic of the proposed high-level waste repository at Yucca Mountain, Nevada. Image courtesy of the U.S. Nuclear Regulatory Commission.

centuries for a repository (for reference, a decade is on the order of 10^8 seconds). Meanwhile the atomic phenomena that occur in the reactor core or in the radioactive waste can be on the order of attoseconds (10^{-18} seconds); similarly for the lengthscales, the reactor and repository structure can be on the order of tens to hundreds of meters, while there are atomic and sub-atomic phenomena that occur over fractions of an Angstrom (10^{-10} meters). The massive disparity between these scales is illustrated in Figure 1.4. Finally, many different materials are involved from metals, ceramics and composites for structural components to water, liquid salts and molten metals for reactor coolants.

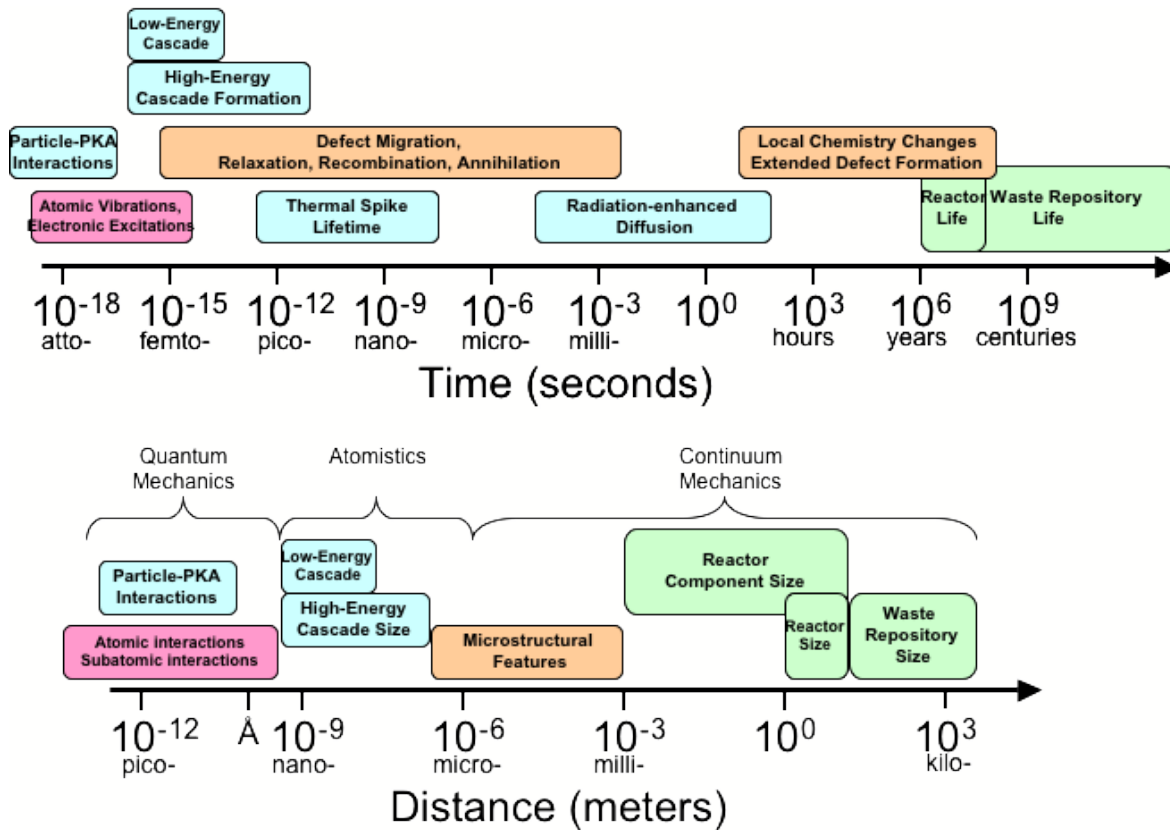


Figure 1.4: Schematic of the time and length-scales of relevance to nuclear engineering. Note the large span of the scales: 26 orders of magnitude difference in time, 13 orders of magnitude difference in distance.

Due to the materials, length and time-scales involved, experiments alone cannot properly describe the behavior of the materials in most nuclear applications. Further, the technological, environmental and safety issues are coupled such that great care must be taken to obtain a complete and reliable picture during the analysis and design phases[11]. Thus, new design paradigms are being explored around the world that make use of a combination of experimental, theoretical and computational approaches to obtain an accurate picture of

material properties and behavior in harsh environments. This dissertation will not explore the details of such a multiscale, multi-method design paradigm as that is best left for articles such as [12–14]. However, it is prudent to briefly review the application area that is the basis for this work. This will allow the reader to better understand the broad picture and specific material requirements for nuclear energy applications. To this end, this chapter aims to provide a broad picture of radiation science and nuclear power generation, then delve deeper into the specific topic area that is the backdrop for this work. This overview will lead into a discussion of the details of the relevant physics and materials in nuclear applications. From this discussion we will focus in on a single material to be studied further in the remaining chapters, namely silicon carbide (SiC). The motivation and introduction will then be concluded by the presentation of a specific focus application to give the broader impact of this work on materials design in the nuclear industry.

1.1 Radiation Basics

This section is intended to be a primer for those not familiar with some of the basic concepts of nuclear engineering or radiation science. This section will provide a general overview of these topics, such that it is equally applicable to problems in nuclear and aerospace engineering, micro and nano electronics as well as the biological sciences. The concepts presented will be important to understand as they will reappear throughout the the remainder of this dissertation.

This section assumes the reader is familiar with the concept of crystal lattices, crystallographic directions (Miller indices) and point defects (i.e. vacancy, interstitial, antisite defects, etc.) in a crystalline material. One can find good overviews in any basic materials science textbook, such as [15] or even online. Put simply, a point defect is the term for when an atom is absent from a lattice site (vacancy) or when an atom is in a place it should not be, such as in-between lattice sites (interstitial) or on a lattice site of a different atom type (antisite). Miller indices are used to indicate directions with respect to the basis vectors of a particular crystal structure. Thus a direction can be defined as follows, given a particular set of lattice basis vectors $\mathbf{a}_1, \mathbf{a}_2$ and \mathbf{a}_3 :

$$[l m n] \equiv l \mathbf{a}_1 + m \mathbf{a}_2 + n \mathbf{a}_3 \quad (1.1)$$

Where a negative value is indicated with an over-bar: $[\bar{l} m n] \equiv [-l m n]$. To denote the family of directions with the same symmetry as $[l m n]$, the notation $\langle l m n \rangle$ is used. One can also use this system to denote planes within the crystal. Generically, $(l m n)$ is used to denote a plane that intersects the basis vectors at some multiple of \mathbf{a}_1/l , \mathbf{a}_2/m and \mathbf{a}_3/n . Additionally, $\{l m n\}$ is used to denote all planes with the same symmetry as $(l m n)$. The Miller indices ideas are independent of the crystal type. There are two special cases that are useful to present briefly as they are commonly used in the materials science community,

- Cubic systems, the triplet of numbers is used (i.e. $[l m n]$) and the basis vectors gener-

ally correspond to a local cartesian coordinate system. For example, the direction of the body diagonal for a cube with sides oriented along the x,y and z directions would be denoted by $[1\ 1\ 1]$, and a plane with a normal vector aligned with this direction would be denoted as $(1\ 1\ 1)$. In cubic systems $[l\ m\ n]$ is the normal to $(l\ m\ n)$. This is not always true in other systems. This dissertation will only discuss cubic systems in detail.

- Hexagonal and rhombohedral systems are often treated using a slightly different system, called the Miller-Bravais indices. In this system a quadruplet of numbers is used (i.e. $[h\ k\ i\ l]$), where h, k, l are the Miller indices that correspond to the hexagonal or rhombohedral cell basis vectors (analogous to Equation 1.1), and i is an extra index whose purpose is to make certain symmetries more obvious. This extra index is defined such that $i = (-h - k)$.

1.1.1 Useful Units and Terms

Due to the particular length- and time-scales involved in nuclear applications, the nuclear physics and engineering community make use of units that are uncommon in traditional structural engineering communities, and even some which are not common in the nanomechanics or solid state physics communities. As some of these units will be used in the remainder of this dissertation, they are introduced here. For consistency, these units will be those which are accepted for use with the International System of Units (SI), but are not

necessarily part of the SI. Further, the definitions of many of these terms including notes on use in practice can be found in several standards[16, 17].

- Atomic Mass Unit (amu). $1 \text{ amu} = 1.66053886 \times 10^{-27} \text{ kg}$ or more precisely, exactly $1/12$ the mass of a carbon-12 atom.
- Angstrom (\AA). Atomic unit of length, $1 \text{ \AA} = 1 \times 10^{-10} \text{ meters}$. For comparison, the ‘diameter’ of a hydrogen atom in its lowest energy state is on the order of 1 Angstrom.
- Attosecond (as). A measure of time, $1 \text{ as} = 1 \times 10^{-18} \text{ s}$. An incredibly small unit of time, relevant in very high energy collisions like those that occur in particle irradiation. Physically, light travels only 3 \AA in this amount of time. In 2004, a group of experimentalists reported a measurement of a roughly 100 as interval[18], the smallest time interval measured as of 2007 [19].
- Electron-volt (eV). Atomic unit of energy, $1 \text{ eV} = 1.60217653 \times 10^{-19} \text{ Joules}$. For intuitive reference, the mean kinetic energy of a particle in an ideal gas at room temperature (300 K) is approximately .039 eV.
- Barn (unit of area). $1 \text{ barn} = 1 \times 10^{-28} \text{ m}^2$, or roughly the effective cross-sectional area of a uranium (U) nucleus.
- Steradian (sr). Measure of angular span in three dimensions (3D analog to the radian), defined by the angle whose projection on the surface of a sphere of radius r is an area of size r^2 . See Figure 1.5 for reference.




Figure 1.5: Schematic of the definition of the 3D angular measure referred to as the steradian. One steradian (sr) is defined as the angle whose projection on the surface of a sphere of radius r is an area of size r^2

- Integral flux measurements (particles/($\text{cm}^2 \cdot \text{sr} \cdot \text{day}$), for example). Integral flux measurements, such as neutron flux with units of neutrons/($\text{cm}^2 \cdot \text{s}$), are a measure of the number of incident particles per surface area, per unit time. Sometimes simply *flux* is used. The integral of the flux ($\phi(E, t)$) through a period of time is called the *fluence* ($\Phi(E)$), i.e. the number of impinging particles per surface area in a given amount of time at a given particle energy, E .
- Primary Knock-on Atom (PKA). The atom in the target that an incident high energy particle initially collides with and to which the particle transmits the bulk of its kinetic energy. For convenience, simulations often do not explicitly include the irradiating particle/PKA interaction and instead impart a known kinetic energy to an atom

designated as the PKA in the simulation domain. This is usually at the top or center of the cell, depending on the situation.

- Displacement cross-section (sometimes ‘DPA cross-section’ or just ‘cross-section’). It is often denoted by σ in the literature, but this work will use σ_{DPA} to avoid confusion with the Cauchy stress. The meaning of the subscript ‘DPA’ will be explained shortly. The displacement cross-section is a measure that relates the probability (P) of an interaction with an atom in a target irradiated at a given fluence ($\Phi(E)$) to that fluence. It is defined as:

$$\sigma_{DPA}(E) = \frac{P}{\Phi(E)} \quad (1.2)$$

As a result it has units of area, typically given in barns, though it is not actually a measure of an area in the geometrical sense. Heinisch and co-workers have calculated cross-section values averaged over the energy spectra for a number of fission test reactors and the ARIES-IV fusion reactor [20]. Typical values for Fe in a fission reactor run from 191 to 390 barns, while SiC varies from 158 to 423 barns. In the ARIES-IV (predecessor to the ARIES-AT that will be discussed later), the Fe cross-section is 762 barns while the SiC cross-section is 348 - thus far fewer neutron-atom collisions are expected in SiC as compared to Fe in the ARIES-IV.

- Threshold Displacement Energy (TDE). The amount of energy required to displace an atom from a stable lattice site into a stable vacancy-interstitial configuration (i.e.

does not recombine). This is typically taken as the smallest applied kinetic energy (i.e. a velocity in a given direction) required to displace an atom, and varies with crystallographic direction. see Figure 1.6.

- Displacements per atom (DPA). A bulk measure of particle radiation damage. Displacements per atom can be thought of as the maximum number of atoms that could be permanently displaced from stable lattice sites at a given fluence. Equivalently, it is the sample-averaged number of times the atoms in the sample have been displaced from stable lattice sites at a given fluence. This can be expressed in mathematical form as follows:

$$\text{DPA} = \frac{1}{N_{atoms}} \sum_i N_{displacements}^i \quad (1.3)$$

Where $N_{displacements}^i$ is the number of times atom i has been displaced and N_{atoms} is the total number of atoms in the sample. Thus, for a single PKA excitation event (as is often considered in a simulation) the larger the sample size, the lower the DPA. However in real applications this does not generally occur, thus the technique for determining DPA in experiments is rather different. Experimentally, it is calculated as follows for known a irradiation time, displacement cross-section and flux [17]:

$$\text{DPA} = \int_0^t \int_0^\infty \phi_{tot}(E, t) \sigma_{DPA}(E) dE dt \quad (1.4)$$

Displacements per atom is a very common measure of radiation damage in reactor

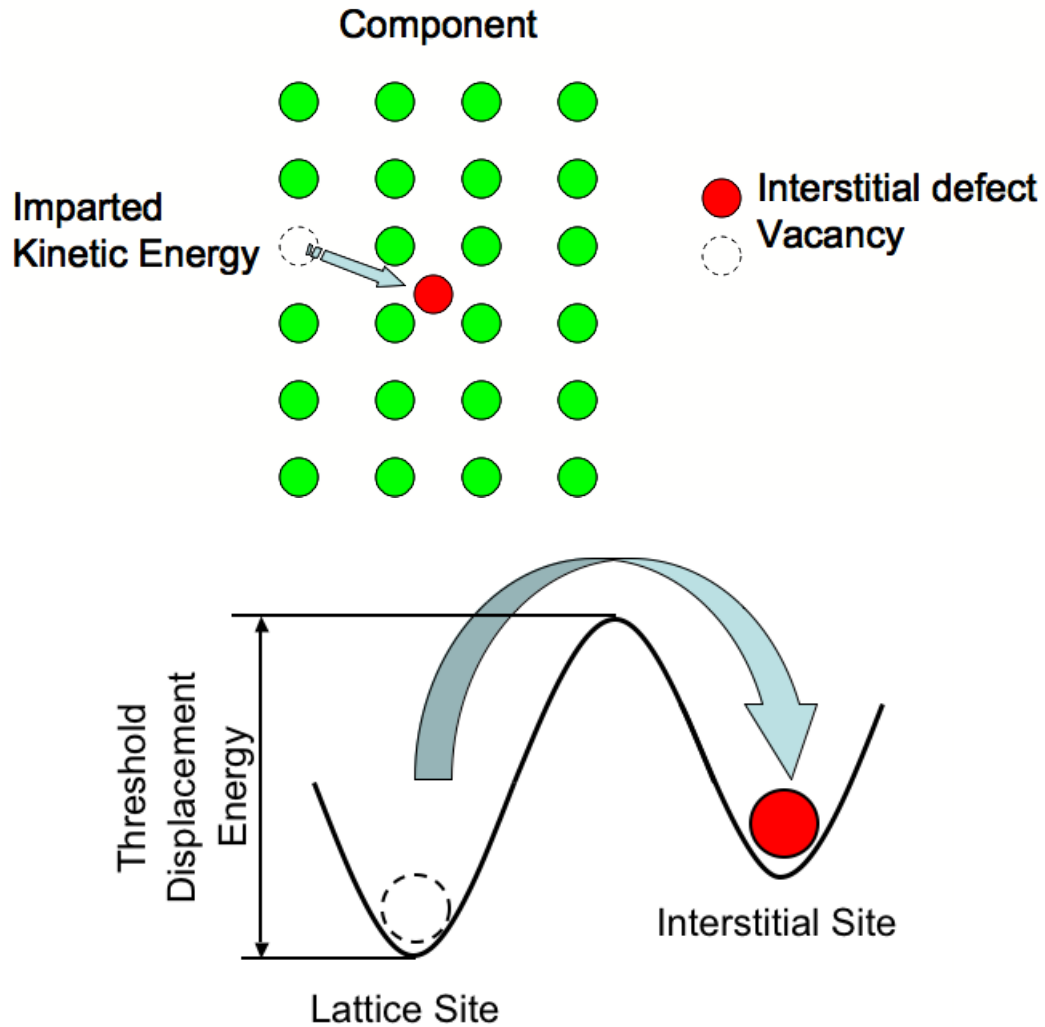


Figure 1.6: Schematic of the definition of the Threshold Displacement Energy(TDE). The TDE is typically taken as the smallest applied kinetic energy required to displace an atom, and varies with crystallographic direction. In other words, the TDE is the kinetic energy which corresponds to the minimum velocity in a given direction that needs to be applied to an atom for it to overcome the potential energy barrier between a lattice site and interstitial site.

applications. However, the drawback to DPA as a damage measure is that it does not provide any information on the distribution of defects within the sample. Because of this it will not be used in this dissertation.

- Gray (Gy) Unit of absorbed radiation dose. $1 \text{ Gy} = 1 \text{ J/kg}$ i.e., one Joule of energy from a radiation source absorbed by one kilogram of matter.
- Exposure. Measure of the strength of the radiation field at some point in a volume of air. Units include the Roentgen, where 1 R (Roentgen) is photon energy required to produce 1.610×10^{12} ion pairs in one cubic centimeter of dry air at 0°C .

1.1.2 Types of Radiation

There are two major types of radiation. The first is electromagnetic radiation (i.e., a stream of photons), such as X-rays, ultra-violet radiation and even visible light. The second is particle radiation, or bombardment by high energy charged (electrons, protons, ions, etc.) or uncharged (neutrons) particles. Both types can cause mechanical damage, microstructural modification and may introduce additional charge carriers into a material or cause ionization, thus are of interest for structural, biological, semiconductor and electrical applications. This section will discuss these two forms of radiation in more detail.

1.1.2.1 Electromagnetic Radiation

Electromagnetic radiation refers to a stream of photons with both energy and momentum. Electromagnetic radiation can occur on various wavelengths. Visible light, as well as infrared radiation and ultra-violet radiation are part of this type of radiation. However it is of greatest interest for nuclear and space applications to consider the very high energy, very low wavelength end of the spectrum such as gamma radiation, a high-energy emission on the order of 100 keV. Another form of importance is X-rays, commonly used in medical applications, with energies on the order of 10 keV.

In aerospace and nuclear applications, the generation of high-energy electromagnetic radiation as a result of interactions between particles (i.e. particle radiation interacting with a body) is of particular importance. This form of radiation is a form of secondary radiation known as *bremsstrahlung* (braking radiation, in German) and is a result of the stopping of incident charged particles by some material. The deceleration of the charged particle results in the release of electromagnetic radiation (a photon). If the energy of the primary particle is high enough, the resulting wavelength of the secondary radiation may be in the X-ray range. Thus high-energy secondary radiation is an important consideration for inhabited areas of spacecraft and for sensitive electronics such as navigation equipment. One easily identifiable instance of *bremsstrahlung* is the ‘Aurora Borealis’, the Northern Lights, visible light emitted as high-energy charged particles impinging on the earth are slowed down by the atmosphere.

1.1.2.2 Particle Radiation

Particle radiation refers to bombardment of an object by high-energy charged (electrons, protons, ions) or uncharged (neutrons) particles. These particles carry mass, energy and momentum. Particle radiation interacts with materials via collisions (momentum transfer) and chemical interactions (charge transfer, bonding, etc.). While particle radiation is a concern common to both nuclear and aerospace applications, there are important differences in the particular particles, energies and fluxes generally found in each case. However, the general principles remain the same.

Aerospace applications are often primarily concerned with the ionizing radiation that air and space-craft are exposed to when traveling at high altitudes or orbiting the Earth. The particle types typically encountered are protons, alpha particles (2 protons and 2 neutrons bound together), as well as electrons [21]. These particles can come from galactic cosmic rays (GCR), charged particles that originate in deep space, or from solar events such as coronal mass ejections or solar flares. There are also the ‘Van Allen Belts’, regions of trapped electrons and protons that encircle the earth. The energies of these particles typically vary from a few MeV up to 1 EeV (10^{18} eV), with a most probable energy of approximately 1 GeV (10^9 eV)[22]. The dosage rate due to GCR as measured near the earth (but still outside the atmosphere) has been observed to vary from roughly 40 μ Gy/day to almost 150 μ Gy/day through the solar cycles between 1989 and 1996 [22]. Luckily the earth’s magnetic field and atmosphere largely protects us from these dangerous particles.

Within nuclear applications, the radiation conditions vary greatly from those present in space. Further, in nuclear reactors the environment varies with the particular type of reactor. There are currently two major classes of nuclear reactor systems: fission and fusion. These two classes are named for the particular atomic reaction that is the basis for their operation. In both reactions, the bulk of the particles emitted are high energy neutrons (so-called ‘fast’ neutrons, at lower energies the term ‘thermal’ is often used), however electrons, protons and ions may also be present as reaction products. Because of the relevance of reactor applications to the present work, these two classes of reactor will be discussed individually in Section 1.2. But first it is prudent to review the fundamentals of particle radiation damage, which will apply to any environment where solids are bombarded by particle radiation.

1.1.3 Particle Radiation Damage

This section will focus on the forms of mechanical damage that result from bombardment of a crystalline solid by high energy particles. This is a topic that must be well understood in order to efficiently design materials and structures for use in environments where particle radiation is a significant concern. This can include applications in broad array of fields such as nuclear reactors, aircraft, spacecraft, and sensor design.

There are several modes of particle radiation damage, such as sputtering, ion entrapment and re-emission, as well as ablation and so-called ‘displacement damage’. Sputtering is a process by which material is ejected from a target body due to bombardment by high-energy

particles. Ablation on the other hand refers to thermally driven erosion processes where material is removed from a target through vaporization, or spallation due to local thermal gradients. Both of these have received much attention from the thin film community [23], where these processes are used to produce films from the eroded material. Ion entrapment (and to some extent re-emission) poses concerns in nuclear applications where helium nuclei and other ions are present as reaction products or fuel. In these cases, the entrapment of these ions may lead to gas bubble formation in structural materials and premature failure [24]. Finally, displacement damage is caused by the collision of a high energy particle with a near-surface atom of a solid target material (see Figure 1.7), and the resulting collisions within the body. It is a well known phenomenon with theoretical foundations that date back to studies of metals in the 1960s[25, 26] and have more recently been applied to crystalline non-metals [27]. Displacement damage can lead to localized microstructural modification and thus mechanical property changes [10]. Because of this, it is of great interest in the design of systems exposed to particle radiation such as nuclear reactors.

Displacement damage is initiated by the collision of a high energy particle with a near-surface atom of a solid target material. If the incident particle is of sufficient energy, the kinetic energy transfer from the incident particle to an atom at or near the surface of the target (the PKA) will result in the formation of a vacancy and interstitial defect pair[27]. The defect pair is considered ‘stable’ if the distance between the vacancy and interstitial is sufficient to avoid recombination due to the induced stress field in their vicinity. Further, in

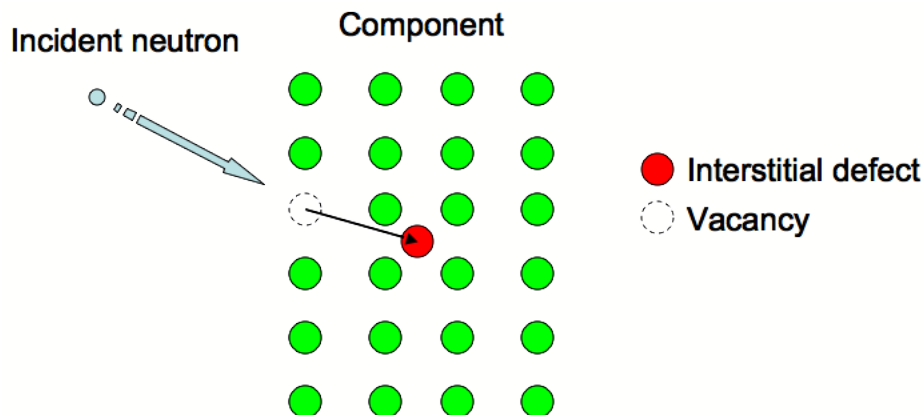


Figure 1.7: Simple schematic of displacement damage caused by an incident neutron. The incident high-energy neutron collides with a near-surface atom of the component. The near-surface atom, or ‘primary knock-on atom’ (PKA), is dislodged from its lattice site to an interstitial location, resulting in a Frenkel pair (vacancy and interstitial pair).

the case of non-zero temperatures it is possible that these defects may migrate via thermally activated diffusion.

If sufficient energy was transferred to the PKA by the initial collision, the PKA may collide with and displace neighboring atoms, who in turn collide with other atoms. This phenomenon is referred to as a *collision cascade* or *displacement cascade*. The cascade process in both metals and nonmetals can be broken down into three stages[10, 28]. The first stage, which occurs after the initial collision that set the PKA into motion, is the ‘collisional’ stage. In this stage, the PKA becomes involved in secondary collisions with surrounding atoms. After approximately .1 ps (for PKA kinetic energies less than about .1 MeV), this stage ends, with a majority of the initial PKA kinetic energy dispersed to other atoms via collisions. The defect structure at this point is a core region of vacancies (i.e.,

vacancy cluster), where the displaced atoms have moved to interstitial sites at the periphery of the cascade. Due to the large amount of kinetic energy left in the cascade region, local heating occurs. This begins the second stage of the cascade, sometimes referred to as the ‘thermal spike phase’. This phase has been estimated at roughly a picosecond, and for small defect regions (1-2 nm), little re-arrangement of the defect region occurs. However, in larger regions (5 - 10 nm), localized temperature gradients can result in void formation as vacancies at the periphery of the cascade diffuse toward the core. The relaxation stage is the final stage of cascade formation. This phase can last on the order of a microsecond or more, depending on the particle flux. Relaxation in the cascade may result in either defect diffusion out of the cascade region, or the transformation of defects into lower energy configurations[28].

The effects of displacement damage can vary with the incident particle type, irradiated material, and temperature. For different incident particle types, differing amounts of energy may be transmitted to the PKA, based on the interactions between the PKA and the incident particle. For example, Thompson [25] explains that a classical mechanics model for the PKA/particle collision is sufficient when the incident particle is heavier than a proton. However, for electrons, their mass and typical kinetic energy put them into the regime of relativistic quantum mechanics. Meanwhile, for neutrons, their short-range interactions generally require consideration via non-relativistic quantum mechanics. An exception is for neutrons and ions with kinetic energy less than 10 MeV, where a classical mechanics approximation of the collision is acceptable[28]. The expression for the energy transferred under

the classical collision approximation is as follows[29],

$$T = 4 \frac{m_1}{m_2} E_1 \quad (1.5)$$

Where T, m_1, m_2, E_1 are the energy transferred to the PKA, the masses of the neutron and PKA, and the total energy of the neutron, respectively. Consider a neutron-Si collision as an example. The mass of a neutron is approximately 1 amu and about 28.1 for Si, thus the fraction of energy transferred is approximately 1/7. Therefore, only a small fraction of the incident particle energy is transferred to the PKA. The equations for electrons and ions at higher energies have been given in [25], and are not repeated here for the sake of brevity.

The irradiated material also influences the effects of radiation damage. For pure, monolithic materials, the crystal structure and bonding between atoms affects the threshold displacement energy, as well as how cascades form due to the relative positions of lattice sites. Another effect of structure on displacement damage is *channeling*[25], where an atom that has been knocked out of its starting position travels between planes of atoms (possibly passing through many unit cells) before colliding with another atom. Another factor is temperature; it has a strong effect on displacement damage, due to the dependence of the mobility of defects on temperature and the thermal fluctuations of the atomic positions.

Now that the basics of radiation and particle radiation damage have been discussed, it is possible to more meaningfully discuss the environment of interest here: nuclear reactors. The

next section will provide a brief overview of nuclear power generation, including a discussion on different reactor types.

1.2 Overview of Nuclear Power Generation

The central idea of nuclear power generation is to make use of a nuclear reaction to release energy, which can then be used to create electricity. Usually this is done by extracting the energy from a nuclear reaction contained in a *reactor* using a liquid or gas coolant, this coolant is passed through heat exchangers and turbines to run generators that produce electricity. Currently there are two major classes of nuclear reactor, namely fission and fusion, and each operates based on different physical principles. The nuclear power plants that exist today are based on nuclear fission. There are currently no operational fusion-based power plants and it is speculated that it may be upwards of 50 years before one is built. However, experimental reactors do exist and a number of projects are underway around the world to create viable fusion power plants such as the ARIES-AT in the USA [30], the ITER project in the EU [31], the JET project in the UK [32] and the EAST project in China [33].

1.2.1 Fission Reactors

Fission reactors work on the principle of nuclear fission, breaking apart heavy nuclei, to release the energy contained in their nuclear bonds in the form of energetic neutrons and

'light' ions as in Figure 1.8. The reaction makes use of highly refined fuel, typically Uranium-235 (^{235}U) or Uranium-238 (^{238}U), to initiate and maintain the nuclear chain reaction. In order to facilitate fission, slow moving or thermal neutrons (kinetic energy less than 1 eV) are introduced into the reactor. The chain reaction releases energy in the form of heat which is used to generate steam (in water-cooled reactors) and run turbines to obtain electricity. However, after the fuel has been used it remains radioactive and must be carefully handled. This spent fuel is the majority of waste in the various waste storage sites mentioned earlier in this chapter.




Figure 1.8: Schematic of a typical fission reaction that uses slow moving or 'thermal' neutrons and Uranium-235 as fuel. The thermal neutron is absorbed by the ^{235}U nucleus, making it unstable. The nucleus then splits into Barium-142 (^{142}Ba), Krypton-92 (^{92}Kr), 2 neutrons and gamma radiation with a total kinetic energy of roughly 200 MeV.

The currently operating reactors are referred to as generation(Gen-) II or III reactors (depending on type), and the current trend is to develop so-called generation-IV reactors[1, 3, 6]. There are two major categories of fission reactor, thermal (slow) reactors, and fast reactors. Thermal reactors make use of moderating materials (often water or graphite [6, 34])

to slow neutrons to sustain the fission reaction. Most reactors currently in operation are of this type. In contrast, fast reactors do not use a moderating material and thus the neutrons remain highly energetic. These types of reactors require highly refined fuel, possibly even weapons-grade, a fact that leads to manufacturing difficulties, political and ethical concerns over nuclear weapons proliferation as a result of creating such fuel[34].

Examples of Gen-IV thermal reactors are molten salt reactor(MSR) and very-high temperature (VHTR) reactors. Gen-IV fast reactor types include the sodium-cooled fast reactor (SFR), lead-cooled fast reactor (LFR), and gas-cooled fast reactor (GFR)[6]. One reactor that has characteristics of both thermal and fast reactor types is the supercritical water-cooled reactor(SCWR). A notable example of a GFR is the Modular Pebble Bed Reactor (MPBR) [35–37]. These proposed designs for Gen-IV fission reactors call for a wide range of operating temperatures from less than 350 °C up to roughly 1000 °C and maximum expected neutron energies on the order of 1-5 MeV[3–6].

1.2.2 Fusion Reactors

Nuclear fusion seeks to combine light nuclei as occurs in the core of stars like the Sun. A fusion reaction requires a high energy ion collision in order for the nuclei to overcome their inherent nuclear repulsion and fuse together (see Figure 1.9). This requires the fuel (for example Tritium, ^3H and Deuterium, ^2H) to be in the form of a very high temperature plasma. The fuel must also be *confined* using a magnetic field to maximize the energy

available to cause fusion reactions and not destroy the reactor structure due to the extremely high temperature of the plasma. The required temperatures and the confinement problem are the main reasons that there are no operating fusion power plants at this time. However, experimental reactors do exist and there are a number of fusion power projects, as mentioned previously. As an interesting aside, the development of a viable first wall and blanket system has been described as one of the most challenging engineering problems in history [38].




Figure 1.9: Schematic of a typical fusion reaction that uses fast moving Tritium and Deuterium as fuel. Within the confined high temperature plasma containing the light fuel particles collide with very high kinetic energies, overcoming the inter-atomic repulsive forces to fuse together producing an unstable Helium-5 atom. The Helium-5 atom decays into a stable Helium-4 atom with roughly 3.5 MeV kinetic energy, and a neutron with roughly 14 MeV kinetic energy. However, additional tritium needs to be produced or *bred* from another reaction to allow the fusion reaction to be sustained.

Many of the proposed fusion power plant designs are based on the so-called Tokamak design, invented in Russia during the late 1950s. This design consists of a toroidal chamber surrounded by a series of magnets, inside a large vacuum vessel. The magnetic field serves to

confine the fusion plasma and reduce the plasma interactions with the chamber wall. The wall includes layers referred to as the ‘first walls’ and ‘blankets’, that absorb the electromagnetic radiation, heat and neutrons released from the fusion plasma. This absorption process heats the first wall and blanket layers and the resulting thermal energy is removed by the reactor coolant (see Figure 1.10). Currently, there are at least four fusion reactor designs that follow this design concept: the ARIES-AT (USA)[30], as well as TAURO [39] and the ITER project in the EU [31]. In the currently proposed reactor designs, particle energies of order 10 MeV are expected [1] with neutron wall loads of up to the order of 10 MW/m² corresponding to a neutron flux of approximately 10¹⁵ neutrons/s·cm² [30]; this is roughly one to two orders of magnitude more than is estimated for Gen-IV reactors. Estimated nominal operating temperatures at the chamber wall range from roughly 300 °C to 1500 °C, with possible spikes to even higher temperatures during pulsed operation.

1.3 Materials for Nuclear Applications

The previous sections have made it clear that the next generation of nuclear reactors will operate in a wide temperature range (roughly 300 °C to 1500 °C), and irradiating particle conditions in the confining structure (predominantly 1-5 MeV neutrons). Additionally, there is the desire to design systems that from the outset have a longer operational life than current systems [1, 3], which have mostly been upgraded through time to keep up with safety requirements. Thus, it is important to make thorough use of cutting edge materials




Figure 1.10: Cross-section of fusion power core illustrating the first wall (FW) and blanket placement (highlighted) in the ARIES-AT Tokamak type fusion reactor (ITER is also based on the Tokamak design). The blanket serves to absorb neutrons as well as transfer heat to the coolant and power generation systems. Scale at the bottom is for reference, in meters. Reprinted from *Fusion Engineering and Design*, 80(1), F. Najmabadi et. al., 'The ARIES-AT Advanced Tokamak, Advanced Technology Fusion Power Plant', pages 3-23, copyright (2006), with permission from Elsevier.

science and engineering to select the best materials for each design. This section seeks to provide an overview of the materials which are of interest in reactor systems, as well as the particular concerns for each material.

1.3.1 Metallic Materials

Traditional engineering structures such as skyscrapers, aircraft, ships, trains and automobiles make use of steel, aluminum alloys or perhaps some more exotic metallic alloys like titanium alloys. However, few of these structures have to operate in the same type of extreme environment as a nuclear reactor. Thus it is no surprise to find that the requirements and therefore materials are different. The candidate metals identified for the next generation of nuclear reactors include ferritic pressure vessel steels, Fe-based austenitic stainless steels, Ni-based stainless alloys and superalloys, zirconium alloys and ferritic/martensitic alloys [1]. These materials have a range of performance concerns in nuclear applications, such as radiation embrittlement and damage, corrosion due to coolant exposure, erosion, hydrogen embrittlement, creep, and swelling.

It has been found that the candidate metals will likely have their lower operational temperature limit set by the fracture toughness, which typically decreases with increasing amounts of radiation damage (radiation embrittlement)[40]. The reason for this forming the lower limit is that brittle failure modes are not desirable in most structures. This embrittlement phenomenon is most pronounced at temperatures below about 30% of the

material's melting temperature [1, 14, 40] (roughly 350 °C). Embrittlement due to radiation damage or helium bubble formation has been indicated as a major concern for all candidate metals through the entire operating temperature range proposed for Gen-IV and fusion reactors (below 350 °C to above 900 °C) [1].

At the higher end of the proposed operating temperature regime, creep behavior and corrosion due to chemical reactions with coolant materials become major concerns, in addition to embrittlement. This is particularly true in the case of the Zr and Ni-based alloys. Meanwhile austenitic stainless steel alloys and high-chromium steel alloys show a higher corrosion resistance in supercritical water (i.e. water with temperature and pressure above 647.3 K and 22.12 MPa, respectively) than ferritic-martensitic alloys [40]. These characteristics make ferritic-martensitic alloys and austenitic stainless steel alloys the primary choices for the majority of proposed Gen-IV fission reactor types. The only exception is the class of Very High Temperature fission Reactors (VHTR reactors) and fusion reactors where ceramics and ceramic composites are being considered and even recommended by some[40]. These materials will be covered in the next section.

1.3.2 Ceramics and Fiber-reinforced Composites

In recent decades, there has been a tendency to adopt ceramics and composite materials for use in structural applications where their tunable properties are thought to offer a greater return than traditional alloys. The Boeing 787 Dreamliner is one well publicized example of

composites being used to achieve cutting edge performance in a large-scale application. It is no surprise then that the proposed designs for the next generation of nuclear reactors follow suit.

As mentioned in the previous section, some fusion and Gen-IV fission reactors design call for structural components to be exposed to temperatures on the order of 1000 °C, thus many of the previously mentioned metallic alloys have to be ruled out due to the operational temperature being on the order of their melting temperature. Further, the designs require the consideration of components exposed to high energy particle bombardment as discussed in Section 1.1.3 and 1.2. In these regions, ceramic coatings may be used as shielding for underlying structural components or as a coating on fuel pellets for pebble bed reactors, like those shown in Figure 1.11[35, 40]. Additionally, ceramic composites (where a ceramic material is either the matrix material, fiber material or both) have become of considerable interest [1, 41, 42]. Currently, there are at least four fusion reactor designs that call for the use of ceramic composite blankets and first walls, namely the ARIES-AT (USA)[30], as well as TAURO [39] and the ITER project in the EU [31]. In particular, silicon carbide coatings and SiC-based composites, such as carbon fiber reinforced SiC matrix composites (SiC/C composites) are being considered for applications in the next generation of nuclear reactors. There has also been mention of carbon coatings for certain applications [40]. Both of these classes of materials bring some interesting properties to the table.

There is relatively little information available on SiC and SiC-based composites, as metals




Figure 1.11: A micrograph of a TRISO fuel pellet used in the MPBR [35] showing a cross section of the shell. Note the use of carbon and SiC as barrier coatings to contain the reacting fuel. This particle would be one of many contained in a single 60 mm fuel sphere. These fuel spheres consist of a 5mm graphite shell around an array of fuel pellets in a graphite matrix.

have a much longer history in nuclear engineering. Much of the information available is due to the interest in SiC within the fusion reactor design community [1]. In monolithic SiC and SiC composites, at temperatures lower than approximately 1000 °C, isotropic expansion (swelling) of the material is a relatively well known phenomenon [42–45]. Other phenomena include void formation and subsequent swelling at temperatures up to approximately 1500 °C [42, 46] and radiation-induced creep at temperatures below 1000 °C [47]. Despite a wide range of phenomenological studies in SiC, the microstructural events responsible for the swelling phenomena and radiation-induced creep are still not well understood[48].

In order to ensure that safety requirements are met, there is great interest in the deformation and failure of SiC. Other considerations are also of great importance in the selection

of materials for reactor applications, such as morphological changes due to irradiation, sputtering of coatings, and chemical interaction with coolants (such as H_2O , He, Li, Pb, etc). Specifically, it has been noted that it is necessary to further investigate radiation effects on defect production and evolution in ceramics[1, 3]. Thus, it is of significant technological interest to understand this material in further detail. The next section seeks to provide the reader with some background on this interesting material, which will serve as the focus of the remainder of this dissertation.

1.3.3 Silicon Carbide (SiC) as an Engineering Material

Up to this point, SiC has only been mentioned in general terms as a material of interest in nuclear applications. However, the remainder of this dissertation requires a much more detailed understanding of SiC. Therefore this section will introduce the structure and bonding characteristics of SiC, as they are responsible for its electronic, thermal and mechanical properties.

The history of artificially produced SiC can be traced as far back as the late nineteenth century, when a process patented in the U.S. by Edward G. Acheson (US patent number 615,648) produced SiC (sometimes referred to as carborundum) as a byproduct of heating clay and coke in a special furnace. At roughly the same time, the French scientist Henry Moissan produced a similar compound from a mixture of quartz and carbon (naturally occurring SiC is still referred to as moissanite). Acheson went on to establish the Carborundum

Company in 1894, which specialized in the production of industrial abrasives. Today, the use of SiC is widespread in a number of areas, well beyond its original use as an abrasive.

Both silicon (Si) and Carbon (C) are group 14 nonmetals, whose valence structure consists of 4 electrons split amongst an *s-orbital* and the 3 *p-orbitals* (p_x , p_y , p_z). It is typically accepted that when such atoms form bonds, it is energetically favorable to form an *sp hybrid* bond[49]. For SiC this means the Si-C bonds are so-called *sp³ hybrid* bonds, where the s and all 3 p orbitals of each involved atom hybridize to form orbital structures that have characteristics of each of the involved orbitals (see Figure 1.12). There are 4 sp³ hybrid orbitals, each made up of a large ‘positive’ lobe and one smaller ‘negative’ lobe. This follows from the unique combinations of the s and p orbitals as follows:

- $sp_i^3 \rightarrow s + p_x + p_y + p_z$
- $sp_j^3 \rightarrow s + p_x - p_y - p_z$
- $sp_k^3 \rightarrow s - p_x + p_y - p_z$
- $sp_l^3 \rightarrow s - p_x - p_y + p_z$

One of the unique aspects of SiC is that it is *polymorphic*[50], it possesses the ability to exist in different crystal structures. In particular SiC exhibits *polytypism*, where the different crystal structures (called *polytypes*) differ only in the stacking order of the Si-C planes (called ‘bi-layers’). Currently, over 200 polytypes of SiC have been determined[51]. Two major structures are zinc-blende with planar stacking like that of a face-centered cubic crystal (i.e.,

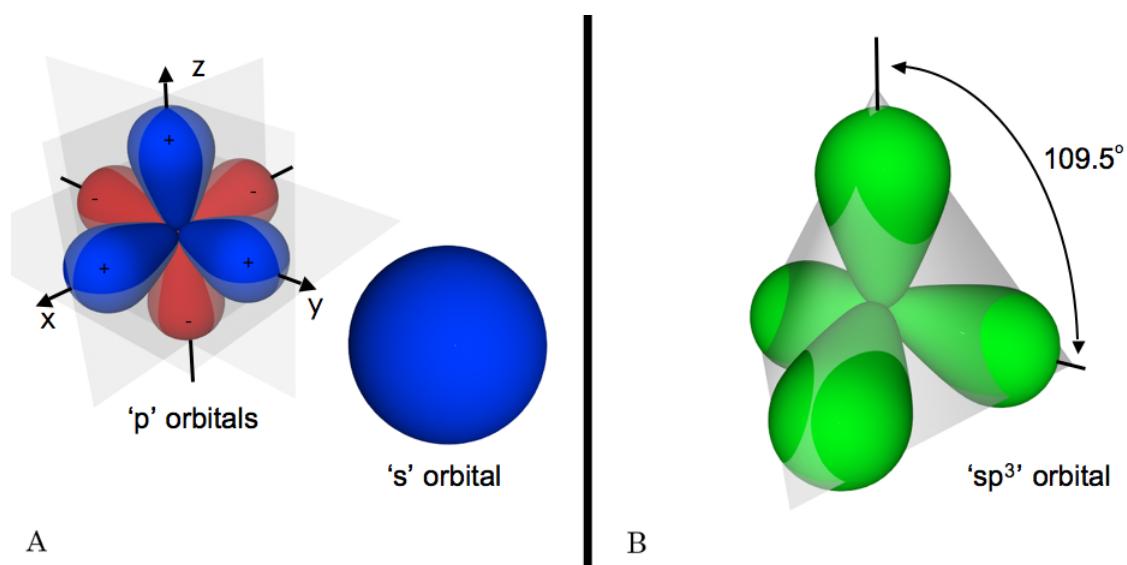


Figure 1.12: A. Schematic of the s and p_x , p_y and p_z orbitals (more correctly, the probability density of the orbital). The subscript of the p orbitals denotes the direction along which it lies. By convention, the positive lobe points in the positive direction. B. Schematic of the 4 sp^3 hybrid orbitals. The lobes of each of the sp^3 orbitals are aligned such that they form a regular tetrahedron, with a separation angle of 109.5 degrees. The smaller 'negative' lobes are not shown for clarity.

ABCABC stacking along $\langle 111 \rangle$) and wurtzite with hexagonal stacking (i.e., ABAB stacking along $\langle 0001 \rangle$). For reference see Figures 1.13, 1.14 and 1.15. The zinc-blende and wurtzite polytypes of SiC are also referred to as 3C and 2H where the number is the periodicity of the stacking (number of unique planes) and the C and H denote cubic and hexagonal planar stacking, respectively. The remaining polytypes of SiC are combinations of the hexagonal and rhombohedral (R) symmetry, with different stacking periodicities. Therefore a series of stacking faults or screw dislocations can lead to transformations from one polytype to another[52], because the polytypes only differ in the stacking direction.




Figure 1.13: Schematic of a 3C-SiC unit cell. This structure is known as the ‘zinc-blende’ structure. The smaller spheres are the carbon atoms, the larger ones the silicon. Dark lines indicate bonds between carbon and silicon atoms, light lines are for reference to the cubic cell. Note that the cell shown here is not a true repeating unit as it includes more carbon atoms than silicon, it is merely a convenient form for illustration.

Despite all of the polytypes having identical chemical compositions, they exhibit different




Figure 1.14: Schematic of the stacking order 3C-SiC (cubic stacking). 3C-SiC is one of the extremes of the SiC polytypes, and is the only cubic polytype. Miller indices are given to show the orientation, the large and small spheres in each bi-layer denote a silicon and carbon atom, respectively.




Figure 1.15: Schematic of the stacking order 2H-SiC (cubic stacking). 2H-SiC is one of the extremes of the SiC polytypes and has the smallest stacking periodicity. Miller indices are given to show the orientation, the large and small spheres in each bi-layer denote a silicon and carbon atom, respectively.

electronic, thermal and mechanical properties. For example 3C has the smallest band-gap (roughly 2.4 eV) while 2H has the largest (3.3 eV) [51], compared to 1.1 eV for Si [53]; 3C-SiC has a thermal conductivity of 5.0 W/cm-°C compared to 1.5 W/cm-°C for Si[54]. Finally, SiC-based composites where SiC is the base of both the matrix and fibers have been developed that demonstrate high tensile strength (up to 3 GPa) and high stiffness (up to 420 GPa), while having a low thermal expansion coefficient (roughly $3.5 (10^{-6}/K)$) [41]. These properties are some of the reasons SiC is of interest in engineering applications for extreme environments.

1.3.4 SiC Manufacturing Techniques

Besides the basic atomic structure aspects of SiC, it is important to review the sorts of processes used to manufacture SiC and SiC-based composite components. This section will review a few of the manufacturing techniques that are used for the creation of both large and small SiC-based components (as opposed to abrasives), with an emphasis on the introduction of defects. The reason for this focus is the hypothesis that initial defects act as ‘sinks’ for introduced defects, where they may recombine or cluster [55–58].

There are several major manufacturing techniques currently employed to manufacture SiC-based components. These are chemical vapor deposition (CVD), chemical vapor infiltration, vapor phase epitaxy (VPE), sintering and reaction bonding (including hot pressing). There have also been demonstrations of selective laser sintering (SLS) as a rapid-prototyping

method [59]; and with the appropriate raw materials, the manufacturing of layers of woven fabric composites has also been demonstrated [60].

Sintering and reaction bonding methods are generally used to create full components of SiC. CoorsTek [61] advertises this technique for producing plates and complex shapes for use in armor applications. They do not report the particular polytype generated (likely a polycrystalline 4H or 6H type due to the stability of these polytypes), but they report a grain size on the order of 5-15 microns.

Chemical vapor deposition is a widely used technique in the fabrication of thin films, and is employed by companies such as CoorsTek [61] and Dow-Corning [62] to produce wafers of SiC for semiconductor applications. Dow-Corning produces both 4H and 6H polytypes. The thickness is on the order of 400 microns, with a vertical orientation of $\langle 0001 \rangle$ and a tolerance of .5 degrees for 6H. For the n-type 4H, the orientation is up to 8 degrees off $\langle 1120 \rangle$, and contains roughly 15 micropipes per square centimeter. Micropipes are small, cylindrical defects that are common in CVD SiC. They are essentially large screw dislocation structures with hollow cores of roughly 1 micron size, and form during the deposition process[63]. Much work has been performed on the reduction of them in wafers due to the adverse effects on the electronic and mechanical properties of the finished wafer [64]. Coorstek advertises being able to produce wafers (of an unmentioned polytype) up to 20 inches (most likely diameter, no thickness measurements given), with a 99.9995% purity and an average grain size of 3-10 microns. Finally, a vacuum plasma spray process for deposited SiC films has been outlined

in United States Patent #5332601. The patent claims the process can produce a dense 3C or 4H film with .35-1.1% porosity and thicknesses over 200 microns. The porosity provided by this process is taken as the state of the art, and used as a basis for some of the work presented later in this dissertation.

1.4 Focus Application: Void Formation in SiC-based Reactor Components

This chapter has provided an overview of basic radiation science and nuclear engineering, described the basic physics of particle radiation damage as well as motivated its study in silicon carbide. This section will provide the finishing touches on the motivation by presenting further details of the general question that this dissertation seeks to answer: *How do voids form from distributed vacancies in neutron-irradiated 3C-SiC at temperatures and pressures relevant to nuclear reactors?*

This question has several parts, given below, some of which have already been addressed in the literature and will be discussed further in the appropriate chapter.

1. *What is a void?*
2. *What is a relevant temperature and pressure?*
3. *What is a relevant PKA initial kinetic energy?*

4. *What is the initial distribution of defects?*
5. *How do different temperature regimes affect the initial distribution?*
6. *How does the initial distribution evolve in time?*
7. *How does an initial void structure affect the initial distribution?*

The first part can be answered by the presentation of the definition of a void that will be used here. In this dissertation, a void is defined as a large, closely connected vacancy cluster. Closely-connected is taken to mean that the cluster is such that the typical distance between nearest neighbor vacancies is no more than the second nearest neighbor distance for the lattice structure. Large is taken to mean that it contains a sufficient number of vacancies to form a stable geometry and does not change shape or size drastically by short-term (picosecond or less) thermal fluctuations of the crystal structure. This will be on the order of 10 or more vacancies. This number is based on the literature, where different groups have defined ‘large’ as anywhere from about 5 to 100 vacancies[65–67].

The second and third parts have been briefly discussed in Section 1.2. For SiC-based components, the main interest is in the fusion reactor community and use in very high temperature fission reactors. These will both have similar operating temperatures and pressures, though the neutron energy spectra will likely vary. Thus, to obtain a rough estimate of relevant temperatures and pressures the ‘first wall’ structure of the ARIES-AT fusion reactor design [30] has been chosen. This design calls for operational temperatures near 1000 °C and

a pressure range of 50-175 MPa [30]. As for the PKA initial kinetic energy; we must consider three things. Firstly that the fusion reaction discussed in Section 1.2 will release energetic neutrons with maximum kinetic energies on the order of 10 MeV. Next, as mentioned in Section 1.1.3, for neutron-ion collisions with neutron kinetic energies greater than the order of 1 MeV, one needs to make use of quantum mechanics to determine the outcome of the collision. Finally, according to the classical collision model only a small fraction of the neutron energy will be transferred to the PKA. This has been shown to be a maximum of 1/7 of the neutron's initial kinetic energy in the case of a neutron-Si collision, as demonstrated in the example given with Equation 1.5. Thus, if we wish to stay in a classical regime throughout the cascade formation and assume an Si PKA, we should adapt a PKA kinetic energy that is less than about 140 keV. Following the radiation damage literature, and in particular [68], we adopt an Si atom as the PKA and a 10 keV PKA kinetic energy. This PKA kinetic energy is large enough to cause a sufficiently large number of defects but the energy is not so large that the early collisions require quantum mechanical considerations. An additional consequence of this choice is that the results are also relevant to high temperature fission reactors, where the irradiating particle energies are much lower than in the fusion case, as mentioned in Section 1.2.

The remaining parts of the question require a combination of short and long timescale atomistic calculations to answer. The reason for this has been discussed in Section 1.1.3: point defects are an atomistic phenomenon and thus their kinetics depends on the atomic

structure of the specimen. Additionally, the generation of point defects in displacement cascades occurs over fractions of a picosecond while the diffusion-dominated relaxation occurs over timescales that can reach to seconds, therefore one needs to study the physics through multiple time scales. Chapter 2 will describe the methodologies that have been employed in the literature for problems such as these. The final chapters of this dissertation will provide an outline of the problems studied to answer the final four parts of the question as well as a discussion of the results obtained with references to the literature.

Chapter 2

Review of Relevant Short-Time Simulation Methods

The previous chapter discussed the basic physical principles of radiation and radiation damage with a focus on nuclear engineering. From there, it motivated the study of an irradiated SiC component and focused in on the atomistic mechanisms responsible for the production and evolution of point defect clusters. A multi-part question was put forth, which briefly mentioned the need for atomistic-level simulations. Thus before beginning to answer the various parts of the question, it is prudent to review the types of simulation methods used in this dissertation and the radiation damage literature. This review is presented here in order to make the literature more accessible as well as establish some additional terminology not discussed in the introduction. Further, this chapter will provide a detailed discussion

on the post-processing methodology employed to mine useful information from the roughly one-hundred gigabytes of raw data produced by the over 27,000 CPU-hours of computations performed in the course of this dissertation.

There are five particular classes of atomic simulation methods of interest here, namely the Binary Collision Approximation (BCA), classical empirical potential Molecular Dynamics (EP to differentiate it from other atomistic dynamics methods, often just MD), quantum mechanical methods such as Density Functional Theory (DFT) , statistical methods such as Kinetic Monte Carlo (KMC) and long-time molecular dynamics methods such as Parallel Replica Dynamics (PRD). While it will be convenient to establish some basic information on BCA, DFT and KMC, these methods will not be used in this dissertation. The primary method will be EP, and for the study of long-time dynamics, PRD will be discussed. For clarity, this chapter will focus on the simulations that are either ‘quasi-static’ (i.e. energy minimization) or short-time dynamics. The long timescale methods such as KMC and PRD will be discussed in detail in Chapter 6.

2.1 The Binary Collision Approximation

Of great importance to the study of particle radiation damage is the concept of ‘stopping’ (particle’s motion halted) and ‘scattering’ (particle redirection) due to particle interactions with a body or target. This area received much attention in the past, particularly the scattering and implantation of ions in metallic solids (see for example [25, 28, 69]). From this

body of work there are several important concepts which will be relevant to the discussions later in this dissertation. Some of these concepts have already been presented in Section 1.1.3. This section will finish the overview of this area with the introduction of the ideas of the Binary Collision Approximation (BCA) [70], and the ‘universal screening function’ put forth by Ziegler, Biersack and Littmark [69].

The Binary Collision Approximation [70, 71] is a simplified physical model used to study high energy ion-atom and atom-atom collisions. It has been implemented in codes such as SRIM [72] and MARLOWE (now seemingly defunct) due to its relative simplicity over full molecular dynamics calculations. However, its nature dictates that it is most useful when studying collisions rather than the post-cascade formation defect migration phases that are of interest in this dissertation. On the other hand, it does yield some useful information and thus is presented briefly here.

The fundamental idea of the BCA is that one assumes that all collisions occur between pairs of classical particles (see Figure 2.1). Assuming that one is able to determine the scattering angles θ and ψ , the kinetic energy transfer from particle 1 to particle 2 can be written as follows[28],

$$T = T_{max} \cos^2(\psi) \tag{2.1}$$

$$T_{max} = \frac{4m_1m_2}{(m_1 + m_2)^2} E_1 \tag{2.2}$$




Figure 2.1: Schematic of the classical binary collision considered in the Binary Collision Approximation. Here, particle 1 has an initial kinetic energy E_1 and impacts particle 2 which was initially at rest. This is illustrated in the laboratory reference frame.

Where E_1 is the initial kinetic energy of particle 1, m_1 and m_2 are the masses of particles 1 and 2 respectively and T is the kinetic energy of particle 2 after the collision. The second equation is the classical expression for the maximum kinetic energy transfer. This expression is similar to that given in Equation 1.5 but without the assumption that the mass of particle 1 is much less than that of particle 2, as is the case for neutron/ion interactions. An important point to realize is that a collision event is far more likely to be ‘glancing’ than perfectly head-on (where $b = 0$), due to thermal vibrations, etc. Thus, a large fraction of the collisions in a cascade will result in the transferred energy being less than the maximum possible. This intuitive idea says nothing about the scattering angles, however. The scattering angles can be determined from the classical ‘scattering integral’ using center of mass coordinates and a series of coordinate transformations to get the result into the laboratory reference frame. These results are as follows[70],

$$\Theta = \pi - 2b \int_R^\infty \frac{1}{r^2 g(r)} dr \quad (2.3)$$

$$g(r) = \left[1 - \frac{b^2}{r^2} - \frac{U(r)}{T} \right]^{1/2} \quad (2.4)$$

$$\tan \theta = \frac{(m_1/m_2)f \sin \Theta}{[1 + (m_1/m_2)f \cos \Theta]} \quad (2.5)$$

$$\tan \psi = \frac{f \sin \Theta}{[1 + f \cos \Theta]} \quad (2.6)$$

$$f = \left[1 - \frac{Q}{T} \right]^{1/2} \quad (2.7)$$

Where b is the ‘impact parameter’ (i.e. distance between the original paths), r is the center-

to-center distance between the particles, $U(r)$ is the interatomic potential energy function that governs the interactions of the particles, T is the initial kinetic energy of particle 1 and Q is the energy lost to electron excitation.

Naturally, now that the BCA has provided a basic way to look at particle collisions from a classical mechanics point of view, it is necessary to think about how the particles in the collisions interact. In real solids, the particles involved are actually a set of nuclei and their associated electronic structure rather than hard spheres. During the sorts of collisions considered in the BCA, the particles interact over very small distances relative to the electronic orbital size, therefore the electronic interactions can be simplified or ignored. Thus, we are left with collisions between nuclei, possibly with some basic electronic contribution. It is well established that such interactions can be modeled by a Coulomb-like repulsive potential energy function[69, 73, 74]. The general form of this potential can be written as:

$$U^{Coul}(r_{ij}) = \frac{1}{4\pi\epsilon_0} \frac{Z_1 Z_2 e^2}{r_{ij}} \phi \quad (2.8)$$

Where r_{ij} is the distance between the particles' centers of mass, ϵ_0 is the permittivity of a vacuum ($.00552635 e/(eV \cdot \text{\AA})$), e is the electron charge (-1 'e' is usually used in atomic units), and Z_1 and Z_2 are the number of protons in each nucleus. In other words, the numerator of the second fraction is the product of the nuclear charges. The final term, ϕ , is the so-called *screening function* and accounts for the presence of the electronic structure of

the atoms at separations larger than the nuclear radius. One widely used screening function is the so-called ‘universal screening function’ put forth by Ziegler, Biersack and Littmark [69]. Ziegler and co-workers fitted a function to theoretically obtained potentials for a large number of atom pairs, thus the term ‘universal’. The form of the Coulomb-like potential described earlier then becomes,

$$U^{ZBL}(r_{ij}) = \frac{1}{4\pi\epsilon_0} \frac{Z_1 Z_2 e^2}{r_{ij}} \phi\left(\frac{r_{ij}}{a}\right) \quad (2.9)$$

$$a = \frac{0.8854a_0}{Z_1^{0.23} + Z_2^{0.23}} \quad (2.10)$$

$$\phi(x) = 0.1818e^{-3.2x} + 0.5099e^{-0.9423x} + 0.2802e^{-0.4029x} + 0.02817e^{-0.2016x} \quad (2.11)$$

This form will be referred to throughout the rest of this dissertation as the *ZBL potential*.

2.2 Empirical Potential Molecular Dynamics (EP)

Empirical potential molecular dynamics (EP) have come to be ubiquitous within the area of computational nanoscale mechanics. The basic theory was developed in the 1950s and has become increasingly popular as readily available computational power increases and improved interatomic interaction models are derived. This section will begin by discussing basic EP theory, and then discuss some particular aspects of EP simulations that are of interest in this dissertation. For a more complete discussion on EP, there are several classical texts on

the topic(for example [75, 76]) and some discussion put forth in a recent text by Liu and coworkers [77].

The principle of EP is simple: to obtain a numerical solution to the classical mechanics N -body problem. Or in less esoteric words: to explicitly integrate Newton's equations of motion in time for an assembly of interacting particles. However, the particulars of obtaining the solution, as well as the set-up of the problem itself present numerous complications.

An EP simulation consists of 4 major components:

- Geometry
- Interatomic potential energy function(s)
- Initial and boundary conditions
- Solution algorithm

The geometry is typically taken to be a cluster of point masses, representing individual atoms (i.e. the center of the nuclei) or groups of atoms. The particulars of the cluster depend on the simulation being performed. So, in the case of 3C-SiC, one would begin with a geometry representing the desired orientation of the zinc-blende structure (i.e. Figure 1.13) with some centers representing the position of Si atoms and others the C atoms.

The interactions between the point masses are determined by potential energy functions. A *potential energy function* (or *potential*) is an analytical or numerical approximation to the interaction between atoms or groups of atoms, and is often a parameterized function that

may be used for groups of similar materials. While the derivation of a potential is both a science and an art, the fundamental idea is to obtain a relatively simple function that reproduces the important physics of a particular material. This derivation is usually performed as a function-fit based on data obtained through a number of specialized experiments or first-principles calculations. Examples of fitting data include bulk elastic constants, lattice parameter, bulk modulus, cohesive energy, surface energy and defect formation energies. Another requirement for a potential to be useful is that it must be *transferable*, i.e., able to approximate the material response apart from those parameters explicitly accounted for in the fitting data. There are several common classes of potential, typically broken down by the type of interactions that are included:

1. *Pair (or two-body) potentials.* This class only accounts for interactions between pairs of particles (designated i and j). Pair potentials are typically a function of the interatomic separation distance, r_{ij} . For example, the simplest of all potentials, the harmonic potential (linear spring model). Others include the Morse[78] and Lennard-Jones[79] potentials.
2. *Three-body potentials.* This class accounts for both pairwise interactions as well as the interaction between triplets of particles, i, j, k . This allows for the inclusion of bond

angle dependence in the potential. These are typically of the form,

$$U_i = \sum_j U_2(r_{ij}) + \sum_j \sum_k U_3(r_{ij}, r_{ik}, \theta_{ijk}) \quad (2.12)$$

Here the potential is for the i^{th} particle, the first sum is the pairwise contribution from neighboring particles, and the second is the three-body term. θ_{ijk} is the angle between the interatomic position vectors \mathbf{r}_{ij} and \mathbf{r}_{ik} . An example of this type of potential is the Stillinger-Weber potential[80].

3. *Local-environment dependent potentials.* These potentials include some information about the environment around the atom of interest. The three-body and higher order classes of potential could be considered a sub-class of this category. Examples of this are the Embedded Atom Method (EAM) potential [81] and the Tersoff Potential[82].

2.2.1 The Tersoff Potential

The Tersoff potential[82] was put forth to describe two-component covalently bonded systems more accurately than previous potentials. In particular, it was geared towards the compounds SiC and Silicon-Germanium (SiGe). It has been widely used in EP simulations of SiC (see for example [68, 83–90]). The basic form of the Tersoff potential is that of a sum

of pair-like components,

$$E = \sum_i E_i = \frac{1}{2} \sum_{i \neq j} U_{ij}, \quad U_{ij} = f_C(r_{ij}) [f_R(r_{ij}) + b_{ij} f_A(r_{ij})] \quad (2.13)$$

$$f_R(r_{ij}) = A_{ij} e^{-\lambda_{ij} r_{ij}}, \quad f_A(r_{ij}) = B_{ij} e^{-\mu_{ij} r_{ij}} \quad (2.14)$$

$$f_C(r_{ij}) = \begin{cases} 1, & r_{ij} < R_{ij} \\ \frac{1}{2} + \frac{1}{2} \cos\left(\pi \frac{r_{ij} - S_{ij}}{S_{ij} - R_{ij}}\right), & R_{ij} \leq r_{ij} < S_{ij} \\ 0, & r_{ij} > S_{ij} \end{cases} \quad (2.15)$$

Where E is the total energy, f_C is an interaction cutoff function, f_R is the repulsive term, f_A is the attractive term, r_{ij} is the distance between atoms i and j , while the other terms are parameters that will be defined shortly. Unlike a pair-wise potential, however, the coefficient of the attractive term (b_{ij}) depends on the local environment and gives information related to the number of bonds as well as the bond angle (i.e. a so-called ‘bond order parameter’),

$$b_{ij} = \chi_{ij} (1 + (\beta_i \zeta_{ij})^{n_i})^{-\frac{1}{2n_i}}, \quad \zeta_{ij} = \sum_{k \neq i, j} f_C(r_{ik}) \omega_{ik} g(\theta_{ijk}) \quad (2.16)$$

$$g(\theta_{ijk}) = 1 + \frac{c_i^2}{d_i^2} - \frac{c_i^2}{[d_i^2 + (h_i - \cos(\theta_{ijk}))^2]} \quad (2.17)$$

Here, θ_{ijk} is the angle formed between the i, j -th and i, k -th bonds (i.e., the angle between the vectors \mathbf{r}_{ij} and \mathbf{r}_{ik}), and the term ω_{ik} is an additional fitting parameter that is available but seldom used (the cited works use $\omega_{ik} = 1$). In the preceding equations, the single

subscripts (such as β_i or c_i) indicate that the parameter depends only on a single atom type, while the double subscripts (such as λ_{ij} or χ_{ij}) indicate dependence on two atom types. The parameter χ_{ij} is determined for pairs of atom types (i.e. Si-C bonds), with $\chi_{ij} = \chi_{ji}$ and $\chi_{ii} = 1$. The remaining parameters that depend on multiple atom types can be determined from the single atom types as follows,

$$\lambda_{ij} = \frac{1}{2}(\lambda_i + \lambda_j), \mu_{ij} = \frac{1}{2}(\mu_i + \mu_j), A_{ij} = \sqrt{A_i A_j}, B_{ij} = \sqrt{B_i B_j} \quad (2.18)$$

$$R_{ij} = \sqrt{R_i R_j}, S_{ij} = \sqrt{S_i S_j} \quad (2.19)$$

Two parameter sets for Si and C are given in Table 2.1. There is, however, one other material dependent parameter required for EP simulations (and dynamics in general): the mass. The atomic masses for C and Si are also given in Table 2.1, based on the values from [91].

2.2.2 EP Initial/Boundary Conditions and Solution Procedures

Once the geometry, potentials and masses have been specified for a given system, one has the basic material information for a simulation. However, virtually all simulations seek to find the response of the system under non-homogeneous (i.e. non-zero) initial and boundary conditions. There are, in general, three major types of boundary conditions in EP simulations, namely: applied force, applied velocity and applied displacement. Similarly, the initial conditions consist of initial velocities and initial displacements. An additional consideration comes

Parameter	Tersoff		Devanathan et al	
	C	Si	C	Si
A (eV)	1393.6	1830.8	1544.8	1830.8
B (eV)	346.7	471.18	389.63	471.18
λ (\AA^{-1})	3.4879	2.4799	3.4653	2.4799
μ (\AA^{-1})	2.2119	1.7322	2.3064	1.7322
β	1.5724×10^{-7}	1.1000×10^{-6}	4.1612×10^{-6}	1.1000×10^{-6}
n	.72751	.78734	0.99054	.78734
c	38049	100390	19981	100390
d	4.384	16.217	7.034	16.217
h	-.57058	-.59825	-.33953	-.59825
R (\AA)	1.8	2.7	1.8	2.7
S (\AA)	2.1	3.0	2.1	3.0
χ_{C-Si}	.9776		1.0086	
mass, m_a (amu)	12.0107	28.0855	12.0107	28.0855

Table 2.1: Parameter set for Si and C, as provided by Tersoff in [82] and by Devanathan et al in [85]. Length parameters S and R were reported to not be optimized in both cases. Mass taken from [91]

into play when considering the evolution of systems at temperatures above absolute zero. At these so-called ‘finite’ temperatures, the atoms in the system vibrate randomly about their equilibrium positions with an amplitude that increases as the temperature increases. There exist numerous techniques to maintain a constant temperature within a simulation cell, such as velocity scaling based on a Maxwell-Boltzmann distribution, the Nose-Hoover thermostat [92], Berendsen thermostat [93] as well as the newly proposed phonon heat bath approach [94]. The final consideration for atomistic level simulations of solids is the representation of a system that is a small section of an infinite, periodically repetitive bulk. This form of boundary condition, called the *periodic boundary condition*, essentially imparts the state of the atoms on one side of the defined simulation domain to the atoms on the opposing side.

With the material information, the initial and boundary conditions, the remaining aspect of the simulation is the solution procedure. The general form of equations to be solved follow from Newton's Second Law for the position, \mathbf{r} , of an atom i . These equations are often cast in terms of displacements, $\mathbf{u}(t)$, from some defined starting position (i.e. $\mathbf{r}_i(t) = \mathbf{r}_i(0) + \mathbf{u}_i(t)$),

$$\mathbf{m}_{a,i} \ddot{\mathbf{u}}_i(t) = \mathbf{f}_i(t) = -\frac{\partial U(\mathbf{u})}{\partial \mathbf{u}_i} \quad (2.20)$$

This series of equations can be solved in numerous ways. For example, in an explicit, dynamic solution one can use any number of numerical time integration techniques. The simplest of these, while still being in widespread use, is the central difference method or Verlet algorithm[75]. For quasistatic simulations time is no longer considered, instead they use the idea of a 'load step'. At each load step, the system energy is minimized with respect to the positions of the atoms. This typically involves an iterative, non-linear solution method such as a conjugate gradient method or Newton iteration scheme [95].

2.3 Post-Processing and Data Mining EP Simulations for Defect Information

As mentioned in the previous section, EP simulations are based on the evolution of the positions, velocities and forces associated with points representing atomic nuclei. Therefore it

is usually necessary to perform some ‘post-processing’, calculations or data manipulations after the solution phase of the calculation, in order to obtain useful information from the simulation. In this dissertation and the cases most common in the radiation damage literature, the information sought will be related to point defects and the Primary Knock-on Atom (PKA). In particular, it is of interest obtain the following data:

- PKA kinetic energy as a function of time
- PKA displacement as a function of time
- Damage (point defects) configuration visualization
- Number of point defects as a function of time
- Number and size of point defect clusters as a function of time
- Radial distribution function (RDF) for vacancies
- Spatial distribution of vacancies in simulation cell

Additionally, due to the dependence of the EP results on the initial conditions it is necessary to consider not only information from individual simulations, but also this information averaged over a number of statistically identical initial conditions. In the result chapters, the bulk of the quantitative data presented will be averaged over some number of runs. For ease and clarity, the details of the averaging procedure used will be given with the results.

This section will discuss the procedures employed in this dissertation to obtain the above data. Examples of the scripts used have been given in Appendix A.2.

2.3.1 PKA Kinetic Energy and Displacement

In order to get a basic idea of the behavior of the PKA through the simulated time, it is useful to look at its kinetic energy and displacement. The raw data was obtained by outputting the position and kinetic energy of the PKA (a pre-defined atom) every 100 or 200 timesteps, depending on the particular simulation. The displacement was defined as simply the magnitude of the relative position vector relative to the initial position of the PKA (i.e. $|\mathbf{x}(t) - \mathbf{x}(0)|$, where $\mathbf{x}(t)$ indicates position at time t). This information was tabulated at each output step and plotted in the open-source graphing package gnuplot[96] as a basic X-Y plot.

2.3.2 Point Defect Identification

Prior to any analysis of point defect data based on an EP simulation, it is necessary to identify these defects as the simulation itself deals only with nuclear positions. Two techniques are used in the literature, both based off of the idea of finding the number of atoms associated with some group of reference points. The first method, used by Devanathan and co-workers [68] and further discussed in [97], is to create some reference geometry and compare the output geometry from each output step of the calculation to this reference geometry. Each

reference point is assigned a sphere of some small radius, generally just large enough to accommodate thermal fluctuations (Devanathan and co-workers use half the nearest neighbor distance, $\sim 0.9\text{\AA}$). If no atom in the output geometry is found within this sphere, the reference site associated with the sphere is labeled as a vacancy. The definition of interstitials in the scheme used by Devanathan and co-workers is to check if an atom has been displaced by at least half the nearest neighbor distance and remains at least half the nearest neighbor distance from all other atoms in the output geometry.

The technique used in this dissertation is based on a Wigner-Seitz (or Voronoi) cell analysis and is also mentioned in the literature[98–100]. This technique draws from the well-studied computational geometry problem of nearest neighbor searches. Similar to the sphere-based method, one begins with some reference geometry and a series of output steps. The reference geometry is used as the centers for a three-dimensional Voronoi diagram[95, 101], such that any given cell of the diagram contains the a volume closest to the center (reference point) associated with that cell. Each output step is then compared to this diagram. If an output point is found within a given cell, the reference point closest to it is the one that is the center for the given cell. Therefore, if one wishes to know how many output points are associated with each reference point, one need only determine the number of points within each Voronoi cell (i.e. the cell's *occupancy*). Highly optimized algorithms exist for just this purpose and have been implemented in codes such as Matlab. Within this method, point defects were identified based on the occupancy of the reference cell:

- Occupancy < 1 : vacancy
- Occupancy > 1 : interstitial site.
- Occupancy $= 1$: normally occupied site

A second check was performed on normally occupied sites to find antisite defects (atom type not the same as in the reference configuration) and atomic replacement events (atom type the same as in the reference configuration).

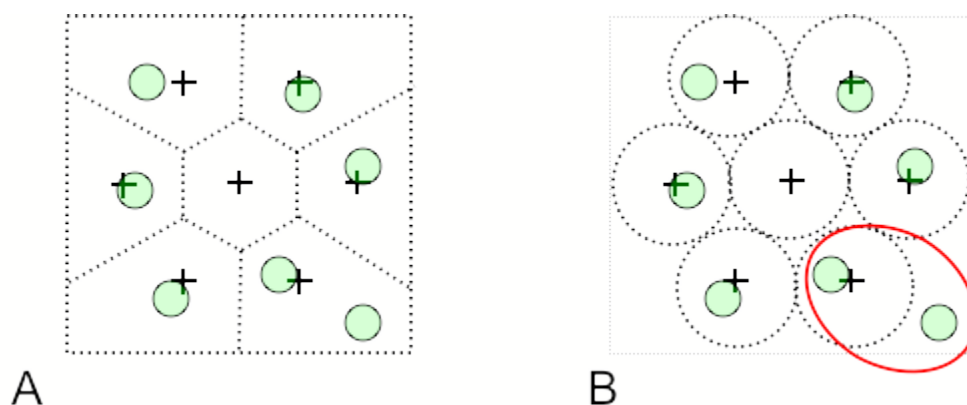


Figure 2.2: Schematic comparison of two point defect detection algorithms. The black crosses indicate the reference points, the circles the atoms in the output geometry. (A) illustrates the method based on Voronoi cells. (B) illustrates the method based on spherical boundaries. Notice that for the atom pair circled in (B), one may not be able to detect the interstitial site because the spheres are not space-filling.

These two methods are schematically represented in Figure 2.2. While the case represented in the figure is simple for the sake of illustration, it still allows for the discussion of the relative merits of each method. Both techniques rely on a reference geometry and implicitly

require that the bulk of the system not deviate significantly from a crystalline arrangement. In other words, neither of these approaches would be suited to an amorphous material, where the definition of vacancies and interstitials is nontrivial. The sphere-based technique is appealing as it is fairly simple to implement, and can build off of the ideas used in the neighbor searches that are generally part of an EP simulation already. The Voronoi-cell method is more daunting, as writing a code to construct the cells and perform the neighbor search efficiently in three dimensions is nontrivial; thankfully such codes are already available as mentioned before. However, ease of implementation is only one of the concerns in computational studies. One other concern is the ability for the technique to be able to accurately capture what is happening in the system. As illustrated in Figure 2.2, the major difference between the algorithms is that the Voronoi-cell method provides search volumes that are space-filling, i.e. there are no gaps or overlaps between adjacent domains. This means that for any given atom in the output configuration, it must necessarily be within one of the reference cells, and therefore associated with a reference point. Further, the Voronoi-cell method does not require the fitting of a search radius and ensures that a ‘true’ nearest neighbor is found. In contrast, with the sphere-based method it is possible for an atom to be outside of any of the reference domains or simultaneously associated with two or more, though the chances of either of these cases may be small in a well-behaved periodic system. As a compromise between the two techniques, one can take the sphere-based method, use larger domains (say about the size of the nearest neighbor distance) and an additional local

search to find the points that are closer to a given reference point than any other reference points. This method would produce the same result as the Voronoi-cell method and is easier to implement in a programming language such as C or C++. However, in the majority of work in this dissertation, the Matlab-based Voronoi cell method was used as it provided the fastest, most complete solution and easiest implementation. For the long-time studies, the C++-based hybrid method was used for ease, though it provided the same results as the Matlab code.

Two sets of files were generated by this algorithm. The first was a set of geometrical data for visualization of the position of the interstitial sites and vacancies, discussed in the next section. The second was a set of plain text files, with rows for each output step and columns that contained the number of vacancies of each type, interstitials, replacements of each type, C_{Si} antisites and Si_C antisites. These were then plotted for each run using the open source plotting program gnuplot [96] in a similar manner to the PKA data.

2.3.3 Damage Configuration Visualization

The geometrical output of the point-defect identification process were converted to EnSight format (a particular geometrical data file format) and visualized with the parallel visualization application ParaView [102]. One advantage of this application is that it is designed for large datasets and allows for the creation of scripts to facilitate batch processing of data. The scripts used in this work are given in Appendix A.2.2.

2.3.4 Defect Clustering and Distributions

To quantify the distribution of vacancies in the systems being studied, three approaches were used. The first method was to take the identified defect geometry, and given a search radius, identify vacancies that have neighboring vacancies within this radius. For vacancies with neighbors, the search is repeated on the neighbor and so on until no new vacancies are found. This group of associated vacancies were then considered a cluster. In order to quantify how isolated an ‘isolated’ vacancy was at the end of the simulations, a body-averaged vacancy radial distribution (histogram based on distance between vacancies, averaged over all vacancies in the body) was calculated. The data from these calculations was then plotted in the package gnuplot[96].

2.4 Essential Density Functional Theory (DFT)

Density Functional Theory (DFT) [103, 104] calculations are one common class of simulation methods that do not require the use of empirically determined interatomic interactions. In the past few decades DFT has been widely used in the solid-state physics and computational materials science communities to study a number of problems in solids such as electronic structures, defect formation and elastic properties. This section will briefly introduce the fundamental concepts of DFT. At this point, the reader unfamiliar with quantum mechanics may wish to reference a basic text such as the very accessible one by Griffiths [105].

Density Functional Theory begins with the Hohenberg-Kohn Density Functional Theorem [103]. This theorem states that for a system of N electrons moving in an external potential field, $V(\mathbf{r})$, described by a continuous, non-negative and normalized electron density field, $\rho(\mathbf{r})$, there exists an antisymmetric wavefunction $|\Psi\rangle$ that describes the system[73]. Thus, for an atomic system the total energy can be expressed as a functional of the position of the nuclei (contained in $V(\mathbf{r})$) and the electron density $\rho(\mathbf{r})$ [106]. The next major step comes from the Kohn-Sham *ansatz*[104] which assumes that the total ground state energy of the system can be written as a functional of the following form,

$$E[\{\mathbf{R}_i\}; \rho(\mathbf{r})] = T_0[\rho(\mathbf{r})] + E_{e,n}[\{\mathbf{R}_i\}; \rho(\mathbf{r})] + E_{xc}[\rho(\mathbf{r})] \quad (2.21)$$

Where $T_0[\rho(\mathbf{r})]$ is the kinetic energy of a system of ***non-interacting*** electrons at density $\rho(\mathbf{r})$, $E_{e,n}[\{\mathbf{R}_i\}; \rho(\mathbf{r})]$ contains the electrostatic energy of the nuclei and electrons, including electron self-interactions, while $E_{xc}[\rho(\mathbf{r})]$ is the so-called *exchange-correlation* energy. The exchange-correlation term includes the Hartree-Fock exchange energy (i.e. energy exchange between electrons), many-body corrections, terms that cancel the energy due to electron self-interactions in $E_{e,n}$ and any kinetic energy not captured by T_0 [106].

A difficulty arises in that the exact form of E_{xc} is in general unknown, therefore requiring assumptions to be made in order to determine the total energy in the DFT formalism. Turning back to $T_0[\rho(\mathbf{r})]$, Kohn and Sham noticed that if one defined an appropriate effective

external potential field for the system of non-interacting electrons, the solution for $\rho(\mathbf{r})$ in both the original and non-interacting systems would be the same and thus a compatible wavefunction for the original system can be found from the solution to the non-interacting system. The form of this effective potential is as follows,

$$V_{eff}(\mathbf{r}) = V_H(\mathbf{r}) + V_{xc}(\mathbf{r}) + V_{ext}(\mathbf{r}) \quad (2.22)$$

Where $V_H(\mathbf{r})$ is the *Hartree potential* field, the Coulomb potential resulting from the charge density of the electrons in the system. The term $V_{xc}(\mathbf{r})$ is defined as the functional derivative of E_{xc} with respect to $\rho(\mathbf{r})$, and V_{ext} is the external potential field due to the nuclei. The resulting Schrödinger equation for the n^{th} particle state, $\psi_n(\mathbf{r})$, with energy ϵ_n is of the form[105]:

$$\left(-\frac{\hbar^2}{2m} \nabla^2 + V_{eff}(\mathbf{r}) \right) \psi_n(\mathbf{r}) = \epsilon_n \psi_n(\mathbf{r}) \quad (2.23)$$

This can be expressed in matrix-vector form (using so-called *bra-ket* notation [105]) as

$$\hat{H} |n\rangle = (\hat{T} + \hat{V}_{eff}) |n\rangle = \epsilon_n |n\rangle \quad (2.24)$$

From the solutions to this equation, it is possible to construct the total kinetic energy of the system. This can be accomplished by first finding the expectation value of the kinetic energy operator, \hat{T} , for each state, and summing over all states that contain electrons (i.e.,

occupied states). Because the states are normalized by convention[105], the kinetic energy can be written cleanly as:

$$T_0 = \sum_n f_n \langle n | \hat{T} | n \rangle = \sum_n f_n \int \psi_n^*(\mathbf{r}) \left(-\frac{\hbar^2}{2m} \nabla^2 \right) \psi_n(\mathbf{r}) d\mathbf{r} \quad (2.25)$$

Where the occupancy of the n^{th} state is given by the scalar f_n . The occupancy is essentially a numerical representation of the number of electrons in each state.

Electrons are fermions, i.e. they follow Fermi statistics, and are subject to the Pauli exclusion principle that states two identical fermions cannot occupy the same state due to the antisymmetry of the wavefunction[105]. With this in mind, for the ground state system considered here, the lowest $N/2$ states are assumed to be occupied by 2 electrons (a valid assumption unless N is small[73]). The state occupancy can then be written as

$$f_n = \begin{cases} 2 & n \leq N/2 \\ 0 & n > N/2 \end{cases} \quad (2.26)$$

From Fermi statistics, this relationship can be extended to finite temperatures in the form of the Fermi distribution[73]. The occupancy expression then becomes

$$f_n = 2f_F(\epsilon_n) = \frac{2}{1 + \exp[\epsilon_n - E_F/k_B T]} \quad (2.27)$$

Where E_F is the so-called *Fermi energy*, the energy associated with the highest occupied

state at absolute zero. Additionally, the sum over the occupancy for all states gives the total number of electrons in the system, i.e.,

$$\sum_n f_n = N \quad (2.28)$$

Recalling the solutions of the single particle Schrödinger equation in (2.24), referred to as the *Kohn-Sham equations*, we can express the electron density as the following,

$$\rho(\mathbf{r}) = \sum_n f_n |\psi_n(\mathbf{r})|^2 = \sum_n f_n \langle n | n \rangle \quad (2.29)$$

Where $\langle n | n \rangle \equiv 1$, but it is kept to make the next manipulations more clear. The electron density also has the following property:

$$\int \rho(\mathbf{r}) \, d\mathbf{r} = N \quad (2.30)$$

Using this, the kinetic energy expression may be re-written in a more convenient form. First consider the expectation value of the effective potential operator,

$$\int \rho(\mathbf{r}) V_{eff}(\mathbf{r}) \, d\mathbf{r} = \sum_n f_n \langle n | \hat{V}_{eff} | n \rangle \quad (2.31)$$

This can then be re-arranged and added to the total kinetic energy from before,

$$\begin{aligned}
T_0 &= \sum_n f_n \langle n | \hat{T} | n \rangle \\
&= \sum_n f_n \langle n | \hat{T} | n \rangle + \sum_n f_n \langle n | \hat{V}_{eff} | n \rangle - \int \rho(\mathbf{r}) V_{eff}(\mathbf{r}) \, d\mathbf{r} \\
&= \sum_n f_n \langle n | \hat{T} + \hat{V}_{eff} | n \rangle - \int \rho(\mathbf{r}) V_{eff}(\mathbf{r}) \, d\mathbf{r}
\end{aligned} \tag{2.32}$$

This manipulation allows for the expression of the total kinetic energy of the interacting system of electrons in terms of the total energy of the non-interacting system, the effective potential and the electron density. Because the states $|n\rangle$ and therefore $\rho(\mathbf{r})$ satisfy (2.24), one may re-write the total kinetic energy expression (2.32) as,

$$\begin{aligned}
T_0 &= \sum_n f_n \langle n | \hat{T} + \hat{V}_{eff} | n \rangle - \int \rho(\mathbf{r}) V_{eff}(\mathbf{r}) \, d\mathbf{r} \\
&= \sum_n f_n \langle n | \hat{H} | n \rangle - \int \rho(\mathbf{r}) V_{eff}(\mathbf{r}) \, d\mathbf{r} \\
&= \sum_n f_n \epsilon_n - \int \rho(\mathbf{r}) V_{eff}(\mathbf{r}) \, d\mathbf{r}
\end{aligned} \tag{2.33}$$

After substitution into the Hohenburg-Kohn-Sham total energy functional (2.21) one obtains,

$$E[\{\mathbf{R}_i\}; \rho(\mathbf{r})] = \sum_n f_n \epsilon_n - \int \rho(\mathbf{r}) V_{eff}(\mathbf{r}) \, d\mathbf{r} + E_{e,n}[\{\mathbf{R}_i\}; \rho(\mathbf{r})] + E_{xc}[\rho(\mathbf{r})] \tag{2.34}$$

By collecting terms in the total energy functional, we obtain something similar to the

total energy form put forth by Chadi[107],

$$E [\{\mathbf{R}_i\}; \rho(\mathbf{r})] = \sum_n f_n \epsilon_n + F [\{\mathbf{R}_i\}; \rho(\mathbf{r})] = E_{band} + E_{res} \quad (2.35)$$

Where the contributions to the total energy not captured by the eigenvalue sum, i.e. band structure term E_{band} , are contained in the ‘residual’ energy term E_{res} :

$$E_{res} = F [\{\mathbf{R}_i\}; \rho(\mathbf{r})] = - \int \rho(\mathbf{r}) V_{eff}(\mathbf{r}) \, d\mathbf{r} + E_{e,n} [\{\mathbf{R}_i\}; \rho(\mathbf{r})] + E_{xc} [\rho(\mathbf{r})] \quad (2.36)$$

Chapter 3

Review and Verification of Displacement Cascade Simulation

Approach

This chapter is the first that will discuss the simulations performed for this dissertation. At the start of the project, radiation damage was a new topic to the author and the Liu group. Thus much preliminary work was required before carrying on to the heart of this dissertation, the question put forth in Section 1.4. In particular, it was necessary to review the radiation damage literature and understand the methodologies employed to simulate radiation damage at the atomic level. This chapter will focus on these two points. The first section will present a brief overview of the simulation literature most relevant to the

current work. Next come the simulation approaches used in this dissertation to determine the threshold displacement energy (TDE) and to create displacement cascades, based on work from the literature. The chapter will conclude with a set of verification simulations performed to understand the physics of displacement damage as well as ensure that the methodology used in the literature was fully understood and reasonable for the problems in this dissertation.

3.1 Relevant Literature

Atomic-scale simulations of particle radiation damage date back to the 1960's. Using rudimentary molecular statics and dynamics calculations, Gibson and co-workers examined radiation induced defects in Cu [98]. They represented the material as a cubic block of atoms, confined by a uniform pressure that balanced the forces caused by the purely repulsive, decaying exponential Born-Mayer potential[108] used to model the interatomic interactions. For their dynamic simulations of displacement cascade formation, they made use of a known kinetic energy (i.e. velocity specified) along a given direction applied to a chosen atom to simulate the energy transfer from the irradiating particle to the primary knock-on atom (PKA). They examined structures caused by PKA energies less than 100 eV, and for durations of roughly 300 femtoseconds in models of less than 1000 atoms. Despite the use of a very simple interatomic force model, they observed several defect configurations. One they termed a 'split' interstitial, where a 2 atom cluster occupied a lattice site. The other,

the ‘crowdion’, was where an interstitial atom was found between two adjacent lattice sites. They concluded that their predictions of the instability of Frenkel (interstitial/vacancy) pairs up to 4th neighbor separation distances was erroneous, possibly due to an oversimplified interatomic potential. Further, they indicated that their cascades produced a larger number of atomic replacements (atom knocked from one lattice site to another) rather than displacements (atoms knocked into interstitial sites) and that boundary effects may have had a significant impact on their reported values.

Jumping forward to the late 1990s, nearly forty years from the studies of Gibson and co-workers, computational power has increased to the point where larger and more sophisticated simulations than theirs are commonly carried out on single processor desktop computers. However, the approaches employed remain remarkably similar. Empirical potential molecular dynamics has found widespread use, with more ‘fine-grained’ calculations such as DFT being used to further elucidate interatomic and defect behavior. These advances, combined with additional work on radiation damage and particle ‘stopping’ and ‘scattering’ have led to significant improvements in the realism of the models used. One improvement of particular use to cascade simulations is the use of ‘hybrid’ potentials, so that more physically-realistic potentials may be constructed from simpler ones. One widely used technique for radiation damage simulations is to hybridize a potential particular to the material of interest with the

ZBL potential (Equation 2.9) as follows:

$$E = \frac{1}{2} \sum_i \sum_j U_{ij} \quad (3.1)$$

$$U_{ij} = [1 - f_F(r_{ij})] U_{ij}^{ZBL} + f_F(r_{ij}) U_{ij}^{LR} \quad (3.2)$$

$$f_F(r_{ij}) = \frac{1}{1 + e^{-A_F(r_{ij} - r_C)}} \quad (3.3)$$

Where f_F is the ‘Fermi Function’ with tunable parameters A_F and r_C . The first, A_F , controls the ‘sharpness’ of the transition from one potential to the other, and r_C is the parameter which controls the radius at which the transition occurs, typically taken to be about 1Å. The material dependent term, U_{ij}^{LR} , provides the longer-range (possibly multibody) behavior while the ZBL potential handles the short-range repulsion-dominated behavior.

Devanathan and co-workers made use of the hybrid potential idea in their studies of the threshold displacement energy (TDE) for 3C-SiC [85] as well as displacement cascade formation in 3C-SiC [68] and 6H-SiC [86]. To overcome issues with the representation of 3C-SiC at small interatomic separations due to errors in the ZBL potential at some ranges, they fitted a cubic spline to ab initio quantum mechanical results and connected this via a smoothing function to a Tersoff potential (described in detail in Chapter 2) for longer

separation distances. The resulting potential was of the following form,

$$E = \frac{1}{2} \sum_i \sum_j U_{ij} \quad (3.4)$$

$$U_{ij} = [1 - f_F(r_{ij})] U_{ij}^R + f_F(r_{ij}) U_{ij}^{Tersoff} \quad (3.5)$$

$$f_F(r_{ij}) = \frac{1}{1 + e^{-A_F(r_{ij}-r_C)}} \quad (3.6)$$

Where U_{ij}^R was the fitted repulsive spline function and $U_{ij}^{Tersoff}$ is the Tersoff potential described in Section 2.2.1. Additionally, [85] cited the use of a modified version of the Tersoff potential described by Tang and Yip [83]. This modified Tersoff potential made use of a cut-off function that scales with the volumetric changes to eliminate spurious bonds under large deformations. In this modified version, the scaled cutoffs take the following form,

$$R_{ij}^{scaled} = R_{ij} \left(\frac{V}{V_0} \right)^{1/3} \quad (3.7)$$

Where R_{ij}^{scaled} is the scaled cutoff parameter corresponding to R_{ij} , and V is the deformed simulation cell volume while V_0 is the undeformed simulation cell volume. Similarly for the remaining cutoff distance, S_{ij} . It is worth noting that the potential energy differences that motivated this change were of order 10% [69, 85] compared to more detailed calculations, a relatively small margin.

The TDE calculations of Devanathan and co-workers [85] made use of a cubic block of 10x10x10 3C-SiC unit cells (8000 atoms total, 4.359 Å lattice parameter), with periodic

boundary conditions applied to all sides. They used a constant system energy, volume and number of particles ('NVE' or 'microcanonical' ensemble), where the atoms had been given an initial velocity corresponding to a temperature of 150 K, sampled from an unnamed distribution (most likely a Maxwell-Boltzmann distribution). They allowed the system to remain unperturbed near 150 K for 2.1 ps with a 1 fs timestep (the thermal equilibration phase), after which they applied a chosen velocity to a chosen atom (the PKA) at the center of the simulation cell and allowed the system to evolve for 3.9 ps with a 1 fs timestep, during which the PKA would move and lose energy. To determine the C TDE, they chose the PKA to be a C atom and an Si atom for the Si TDE. In order to determine the TDE in a particular crystallographic direction, they chose the velocity such that the magnitude corresponded to the desired kinetic energy and the direction corresponded to the desired crystallographic direction. While they did not explicitly describe the technique used to vary the applied velocity magnitude, they do state that it was varied by decrementing its value from some energy where a permanent PKA displacement occurred until no such event occurred. In other words, they first applied a velocity high enough that the PKA left its original lattice site and did not return in the course of the evolution phase (determined by examining the position output). Thus, the energy applied was above the TDE for that crystallographic direction. They then re-ran the system, beginning from a perfect crystal, thermally equilibrated it and applied a smaller velocity corresponding to a decrease of 1 eV in the kinetic energy. The smallest energy at which a displacement event was observed was taken to be the TDE in

that particular direction for that particular atom type. In this manner they claimed to have determined the TDE within an error of 1 eV; i.e., the value they give is taken to be the ‘real’ TDE ± 1 eV, because they did not sample with increments smaller than 1 eV. The range for the C TDE was found to be 28 to 71 eV and for Si, 36 to 113 eV (with an estimated ± 1 eV error). The extremes of the TDE measurements were reported to be between the [111] (Si maximum, C minimum), $[\bar{1}\bar{1}\bar{1}]$ (C maximum) and [001] (Si minimum) directions, reflecting the lack of reflection symmetry along those directions and [110] was reported to have a high TDE due to the close packing of atomic layers (i.e. atoms are close together and have little room to move). The same authors later examined the TDE in 6H-SiC, and compared it to that of 3C-SiC [86]. They found that the TDE ranges were similar, as well as the anisotropy trends with varying crystallographic orientation.

A later, ab-initio DFT study of the displacement threshold energies for 3C-SiC was performed by Lucas and Pizzagalli [109]. They made use of a plane-wave pseudopotential method based on DFT to examine the details of displacements and compare with experimental and EP results from the literature. Their results for both the C and Si sub-lattices disagreed in general to those in [68, 86], with the lowest TDE for both Si and C being in the $[\bar{1}\bar{1}\bar{1}]$ direction (21 eV for Si and 16 eV for C), while the largest reported C TDE was in the [111] direction (38 eV), and the largest reported Si TDE was in the [100] direction (46 eV). Further, their initial temperature was 300K, so it is possible thermal fluctuations contributed to changes in the TDE. Additionally, the simulation cell was reported as being

‘very small’ (64 atoms), but was reported to produce results no different than the results produced with a larger, 216 atom cell (still smaller than those used in EP simulations). The authors noted that for SiC this small cell was appropriate as the PKA did not move far from its initial location. Others supported the small cell size [110], however [109] reported that this would be too small of a cell for pure Si to avoid interactions between the thermostat and the PKA. This was in disagreement with [110]. Further, it has been reported that the choice of interatomic potential can influence such quantities as defect formation energies and migration characteristics in metals (Fe in particular) [111]. This same article indicates that point defect formation energetics are strongly dependent on interatomic potential. While the specifics of this paper do not apply to SiC, as it is a very different material than bcc-Fe, it nonetheless presents the possibility that a more accurate interatomic interaction model may be required in areas containing defects or where defects may form. This is supported by the differences of up to nearly 65 % in the C TDE found by [110] for 3C-SiC represented by a Pearson potential, a Tersoff potential and a simplified local orbital/local density sp^3 basis Density Functional Theory (DFT) method (Fireball96[112, 113]). However, the Tersoff potential did repeatedly fare better in comparison to the DFT than the Pearson potential, with differences of roughly 22 %.

The simulation boundary and initial conditions in the cascade calculations in [68] were similar to the Gibson work [98], though [68] used a larger simulation cell, periodic boundaries, and provided a much larger initial kinetic energy to the PKA (10 keV to a Si PKA).

The simulation cell in [68] consisted of 20x20x60 3C-SiC unit cells (roughly 192,000 atoms, 4.359 Å lattice parameter), with {001} (i.e. (001)-like) crystal plane boundaries. They carried out their simulations keeping the number of atoms and system volume constant, and controlled the temperature of the system through a 4-atom layer velocity scaled thermostat region at 300 K. In this way, they were able to simulate the irradiation of a system initially at 300 K, but attached to a heat bath so that the system total energy did not remain constant. Physically, this could represent a localized irradiation event, such that the energy introduced by the PKA excitation would be carried away to the bulk of the body and allow for some local temperature changes. After the system was allowed to thermally equilibrate for 2 ps, the PKA (a chosen Si atom at the top center of the simulation cell) was given an initial velocity in a similar fashion to the TDE calculations in [85]. Here, however, the direction and magnitude were fixed. The 10 keV PKA kinetic energy was imparted in the form of an initial velocity (after thermal equilibration) of 2620 Å/ps along the $[4\ 11\ \bar{9}5]$ (just off the Z-direction) direction, reportedly to minimize channeling of the PKA. The total cascade evolution time they simulated was 10 picoseconds, however no explicit timestep size information is given. The authors found that after the PKA was excited, the number of point defects (Si or C interstitials or vacancies) would slowly increase, reaching a peak at roughly .1-.3 picoseconds, then decrease to roughly 1/2 of the maximum number of defects by .8 picoseconds. Prior to approximately .1 picosecond, the number of C and Si defects were roughly the same; however, as time carried on the number of C defects would hold steady at

roughly three times the number of Si defects. The authors attributed this behavior to the differences in TDE between C and Si. They also found that roughly 7% of the displacements led to anti-site defects (an Si settling in a C lattice site, or vice versa) and concluded that this may be an important energy storage mechanism, and may lead to irradiation induced amorphization. They further indicated that there was ongoing work to examine the energetics of particular defects in 3C-SiC.

Perlado and co-workers [114] carried out studies in SiC very similar to those in [68], with a PKA kinetic energy range of .5 to 8 keV using both C and Si PKAs, primarily at 300 K (a few runs at 1300 K were reported as well) . They varied their system size from 25x25x25 (150,000 atoms) to 40x40x60 (768,000 atoms) 3C-SiC unit cells, depending on the magnitude of the PKA energy, with a 2-4 atom layer thermostat region as in [68]. They found a final ratio of C to Si vacancies closer to 5:1 than the nearly 3:1 of Devanathan and co-workers. However, they also went on to study the size of clusters formed by vacancies in the system. To do this, they identified the number of vacancies within a given distance of another vacancy. They used the third-nearest-neighbor distance of 4.36 Å as the cutoff radius for the cluster search. They found that clusters of more than a few vacancies rarely occurred in their calculations and that the bulk of vacancies were isolated (not in a cluster). Meanwhile, they found less than 5 clusters of less than 5 vacancies formed in the energy range they studied. They also reported that these clusters were stable for timescales of a nanosecond even at 1300 K, and that the question that remains unanswered is if the clusters

are stable on longer timescales.

The work described in this section forms the basis for the work of this and subsequent chapters. All of the simulations performed for this dissertation were performed using the Sandia National Laboratories Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) [115] on a supercomputer at the US Army Research Laboratory Major Shared Resource Center High Performance Computing Center. Some features were added to LAMMPS by the author in order to better capture the physics of the problems studied. It is worth noting that LAMMPS is a different code from what has been used in the literature cited, most of which used now defunct empirical potential molecular dynamics (EP) codes called MOLDY and MDCASK. The average run-time varied from roughly an hour to 6 hours, and the number of processors used varied from 1 to 64 as needed for the particular run. For those not familiar with the software package, examples based on the TDE and cascade simulations are given in Appendix A.1.1 and A.1.2, respectively. The post-processing methods have been discussed in Section 2.3, and will not be discussed in detail here.

3.2 Simulation Approaches and Verification Cases

From the simulation work in the literature, there are two particular sets of EP calculations directly related to the goals of this dissertation. The first of these are the TDE calculations of Devanathan and Co-workers [85]. This particular calculation was important to ensure that the TDE in the present simulations compared well to what was observed for the nearly

identical interatomic potential in [85]. This also functioned as a basic check on the validity of the Tersoff/ZBL hybrid potential as implemented by the author in LAMMPS. The second set of calculations are the 10 keV Si primary knock-on atom (PKA) cascade formation calculations from [68]. This calculation was used as the basis for the majority of the current work, thus it could be used to provide a basic verification of the simulation and post-processing methodology employed in this work. For the sake of the interested reader, example input files for the simulations have been given in the brief LAMMPS tutorial in Appendix A.1.

3.2.1 Threshold Displacement Energy Simulations

The first verification case is the TDE calculations of Devanathan and co-workers [85]. A system as close as could be made to that used in [85] was created. A Tersoff/ZBL hybrid potential was used rather than the ab initio short range potential from [85], as there was not sufficient information provided in the paper to accurately recreate the potential. Further, based on a private communication with Bill Weber and Ram Devanathan, the ZBL short-range potential would provide very similar results. The Tersoff part of the hybrid potential, however, still made use of the parameter set provided in [85], as given in Table 2.1. For convenience, an example LAMMPS input file from this set of calculations is given in Appendix A.1.1.

3.2.1.1 System Setup

The geometry considered for the TDE calculations was a cube of 10x10x10 3C-SiC cubic unit cells (8000 atoms) with {001} boundaries (the orientation shown in Figure 1.13). To each boundary, periodic boundary conditions (PBCs) were applied. The lattice parameter used was 4.359 Å (roughly 52.2 Å per simulation cell side). An atom of the appropriate type (C or Si) in the center of the simulation cell was chosen as PKA for the C and Si TDE calculations. This geometry is represented schematically in Figure 3.1, and is the same as that used in [85].

The interatomic potential used in the calculation was a hybrid Tersoff/ZBL potential, similar to the hybrid potential mentioned in Section 3.1. The form of this potential can be expressed as in Equation 3.4, with U_{ij}^R given by Equation 2.9. The tunable parameters in the Fermi-like cutoff function were taken as $A_F = 14 \text{ Å}^{-1}$ and $r_c = .95 \text{ Å}$, based on the values from [85]. Despite this being different than the ab-initio spline function from [85], it is believed that this potential would still provide reasonable results as the reported margin of 10 % energy errors due to the ZBL potential is acceptably small.

The same methodology for finding the TDE value was employed here and in [85]. As in [85], all of the atoms in the system were given an initial velocity corresponding to 150 K based on a Boltzmann distribution, and allowed to equilibrate for 2.1 ps before the PKA excitation was applied. After the PKA excitation, the system was allowed to evolve undisturbed for 3.1 ps while position and kinetic energy data were periodically written to disk for later




Figure 3.1: Half-section along $[100]$ of the geometry used for the threshold displacement energy calculations. The circle indicates the position of the PKA. Small spheres are C atoms, larger ones are Si. The arrow is to simply indicate conceptually that an instantaneous velocity is applied to the PKA during the simulation.

inspection. The calculations were carried out with dynamics in a microcanonical ensemble (number of atoms, volume and total energy held constant across the system). The timestep used was between .1 and 1 fs, depending on the applied PKA kinetic energy, chosen such that the maximum displacement in the system in a single timestep was below 1/5 of the nearest neighbor distance of roughly 1.9 Å.

The PKA was excited by an applied instantaneous velocity in a known direction, corresponding to the desired kinetic energy and crystallographic direction. A displacement event was assumed if the PKA moved and stayed more than 1 Å from its initial position for the remainder of the simulation. This was done to ensure that thermal fluctuations would not be counted as displacement events. For each TDE investigated, an initial run was performed with the PKA velocity set to correspond to a kinetic energy value that was determined by taking that reported in [85] and rounding up to the next 10 eV (so 36 eV becomes 40 eV, for example). This was done to ensure that the initial guess for the TDE was not so high that a cascade formed nor too low to create a displacement. Several runs were necessary to find the TDE for a given atom type and direction. For each of these runs, the simulation was restarted with an identical starting configuration but a different PKA kinetic energy. The magnitude of the applied velocity (i.e. kinetic energy) was varied by decreasing the value from the initial guess until a value was found where an additional 1 eV decrease in the PKA kinetic energy did not result in a permanent displacement of the PKA from its original lattice site. This value was then recorded as the TDE for the PKA type and crystallographic

direction, and taken to be within a margin of error of ± 1 eV from the ‘actual’ TDE, just as in [85].

3.2.1.2 Results and Discussion

After the determination of the TDEs for the $\langle 001 \rangle$, $\langle 111 \rangle$ and $\langle 110 \rangle$ directions, the results of this study were compared to several previous studies. This comparison is given in Table 3.1. The present study and [85] used the same Tersoff potential parameter set, given in Table 2.1, but different short range potentials, as mentioned previously. Hensel et al [116] used Tersoff’s original parameter set [82] and the ZBL potential with an initial temperature of 0K. The work of Wong et al [117] used Tersoff’s parameters as well and an unreported short-range repulsive potential with an initial temperature of 10K. Reference [109] used DFT calculations at an unreported temperature. As is commonly done in the literature, a weighted average was calculated to get an idea of the TDE over the range of PKA directions. This average is can be considered a ‘material average’ TDE for each PKA type, taking into account the crystallographic symmetry of the material. This weighted average is calculated in the following manner,

$$\text{TDE}_{ave} = \left(\sum_{i=1}^{n_{dir}} \frac{1}{w_i} \right) \sum_{i=1}^{n_{dir}} w_i \text{TDE}_i \quad (3.8)$$

where the weight, w_i , is the number of directions that are equivalent to those explicitly simulated, n_{dir} is the number of directions explicitly simulated and the first term is the total

PKA Type	Direction	Present	Devanathan [85]	Hensel [116]	Wong [117]	Lucas [109]
Silicon	$\bar{1}\bar{1}\bar{1}$	43 ± 1	39 ± 1	46.5	40	22
	[001]	41 ± 1	36 ± 1	42.5	30-35	46
	[110]	64 ± 1	71 ± 1	65.5	85	45
TDE _{ave}		42	55	56	*	41
Carbon	[111]	20 ± 1	28 ± 1	21.5	25	16
	[001]	15 ± 1	31 ± 1	13.5	40	18
	[110]	27 ± 1	38 ± 1	17.5	30	14
TDE _{ave}		21	34	17	32	15

Table 3.1: Threshold Displacement Energy values for 3C-SiC obtained by different researchers (all values in eV). If error ranges are not given, they were not given in the reference. Note that in the present study and [85] there is a nearly 2:1 Si:C TDE ratio despite quantitative differences. The differences in the results can be attributed to differences in simulation methods and definitions of the TDE as discussed in the text, but the results of the present work are reasonable compared to the literature.

number of equivalent directions. For the 3C-SiC considered here, there are 6 $\langle 001 \rangle$ directions, 4 $\langle 111 \rangle$ directions and 12 $\langle 110 \rangle$ directions[109]. Thus $n_{dir} = 3$, $w_1 = 6$, $w_2 = 4$, $w_3 = 12$ and the total number of equivalent directions is 22.

From a glance at Table 3.1, it is possible to see that there is a clear difference (greater than about 10 %) between the present work and [85] in several of the directions determined, most evident in the C TDEs. However, it is also clear that there is a wide spread of predicted values for 3C-SiC in the literature, based on the other references. It is equally clear that all of the studies indicate a far lower average C TDE than that of Si, thus the most basic displacement behaviour (C being easier to displace than Si) should be seen in a cascade using any of these potentials. The differences between the various reported TDE results are attributed to differences in the definition of TDE (in particular what they deemed a

PKA Type	Direction	Malerba [118]		Present	Devanathan[85]
		Lower	Upper		
Silicon	[111]	38	75	43 ± 1	39 ± 1
	[001]	35	62	41 ± 1	36 ± 1
	[110]	73	93	64 ± 1	71 ± 1
Carbon	[111]	20	31	20 ± 1	28 ± 1
	[001]	30	30	15 ± 1	31 ± 1
	[110]	26	58	27 ± 1	38 ± 1

Table 3.2: Threshold Displacement Energy values for 3C-SiC obtained by Malerba and co-workers [118], compared with the present work and [85] (all values in eV). Notice that the present work falls in the range found by Malerba and co-workers as often as [85].

‘displacement event’), and the interatomic potentials. The work of Malerba and co-workers [118] illustrates with the same potential as [85] and a slightly smaller domain, however, that the reported results in the literature may not be as exact as the authors indicate. In particular, they cite the importance of considering an ‘uncertainty band’, bounded by an upper and lower TDE value. Within this band, metastable defects may be produced that appear to be displacement events but are not stable over longer timescales or initial condition perturbations. Malerba and co-workers published a set of upper and lower bounds for several crystallographic directions, as given in Table 3.2. From this comparison, it is likely that the differences between the present work and [85] are due to the criteria used to define a displacement event and possibly statistical fluctuations due to the initial conditions. It is likely if more than a single initial condition had been used and the results averaged, the present work would match [85] more closely.

The other major possibility for differences in the literature is the potential used. For

all of the directions simulated in the present work, the PKA did not collide directly with or displace one of its initial neighbors. Thus, due to the directions the PKA traveled, it is likely that long range (greater than 1 Å) interactions are more important in this particular calculation than the short-range and collision behavior. One interesting note on the data comes from [116] and [117]. Despite the use of the same long-range potential, they clearly differ in the estimates of most of the TDE values as well as the weighted average. One may suspect that this is due to the temperature difference (0K vs 10K), but more work is required to make this more conclusive. Finally, there is nearly a 2:1 ratio between the weighted average Si and C TDE, in both the present case and that reported in [85], despite some quantitative disagreement in the TDE for each direction. Thus, it is concluded that the variation between the present work and that in [85] is within the variation present in the literature due to subtle differences in methodology; therefore the potential implementation used in this work is expected to produce satisfactory displacement cascade behavior.

3.2.2 10 keV Si PKA Cascade Verification Case

Just as the TDE calculations of [85] were an important benchmark for the Tersoff/ZBL hybrid potential implementation, the cascade formation calculation presented in [68] was the starting point for much of the simulation work in this dissertation. Therefore, it is only fitting to use their results to verify the simulation approach used in this dissertation. A detailed verification effort was necessary due to the absence of some important simulation

details from [68] and to the different software package employed in this work. This verification case makes use of the Tersoff/ZBL hybrid potential, previously verified, as well as the same crystal structure from the previous section. To make this section more transparent, the LAMMPS input file used is given in Appendix A.1.2.

3.2.2.1 System Setup

The crystal geometry used was similar to that used in Section 3.2.1. It consisted of $28 \times 28 \times 64$ 3C-SiC cubic unit cells (401,408 atoms) with the long dimension ($[001]$) oriented along the Z-direction, $[100]$ along the X-direction and $[010]$ along the Y-direction. The system was subject to periodic boundary conditions on all sides to simulate being a small part of a infinite body. A 4-cell thick thermostat ‘skin’ was used on all sides of the box except the ‘top’ (topmost (001) atom planes). This was done to avoid any unphysical behavior due to the thermostat artificially reducing the energy of the the PKA and its neighbors at the beginning of the simulation. The domain size was chosen such that the non-thermostated region matched the $20 \times 20 \times 60$ cell (192,000 atom) ‘simulation domain’ described in [68]. The thermostat region was maintained at 300K via velocity scaling; in other words the velocities of all of the atoms in the thermostat region were scaled such that their temperature was equal to the target temperature. This allowed for a more natural behavior than a system-wide thermostat. The lattice parameter was taken as 4.359 \AA , the value given in [68]. The PKA was chosen as an Si atom 1 atom layer from the top-center of the simulation cell,

based on the description of how the PKA was chosen in [68]. As in [68], the 10 keV PKA kinetic energy was imparted after thermal equilibration, detailed in the next paragraph, in the form of an instantaneous velocity of $2620 \text{ \AA}/\text{ps}$ along the $[4\ 11\ \bar{9}5]$ direction (just off of the Z-direction), to minimize channeling. A schematic of the system is given in Figure 3.2.




Figure 3.2: Schematic of the geometry used for the reproduction of the 10 keV Si PKA cascade formation calculations of Devanathan et al [68]. The circle represents the initial position of the PKA.

A single simulation was carried out in several continuous stages (no restarting or re-initialization), based on the phases of cascade formation discussed in Section 1.1.3. These stages were as follows:

- Thermal equilibration, part 1. All atoms in the system have their velocity scaled every step to maintain an average temperature of 300K, for 1000 timesteps with a 1 fs

timestep (1 ps total time).

- Thermal equilibration, part 2. Interior region (everything except the 4 cell thermostat skin) thermostat removed. Thermostat skin maintained at 300 K. System allowed to equilibrate for 1000 timesteps with a 1 fs timestep (1 ps total time).
- Initial collision phase (labeled in later diagrams as I). PKA instantaneous velocity applied (but not maintained). System allowed to evolve for .2 ps (20,000 timesteps with a .01 fs timestep).
- Intermediate evolution phase (labeled in later diagrams as II). System allowed to evolve for 1 ps (10,000 timesteps with a .1 fs timestep).
- Final evolution phase (labeled in later diagrams as III). System allowed to evolve for 10 ps (10,000 timesteps with a 1 fs timestep). Total time from PKA excitation: 11.2 ps.

The first thermal equilibration stage was effectively an NVT ensemble. However, after the first thermal equilibration stage, the force calculations in the interior region were carried out in the microcanonical ensemble (number of particles, N , volume, V , and energy, E , constant), but the exterior was still subject to velocity scaling (a variation of the NVT ensemble, where the temperature, T , is nominally constant). Thus, the net effect is an ensemble where energy is allowed to leave the system through the interaction with the thermostat region (i.e. atoms with neighbors in the ‘skin’ are indirectly affected by the thermostat), but the

thermostat does not directly affect the energy in the cascade (interior) region. This is an important consideration as a system-wide thermostat would lead to an artificial reduction in the energies of the cascade atoms, and thus lead to unphysical cascade behavior. Physically, the thermostat and periodic boundary approach represents an isolated cascade forming in an infinite body at a given initial temperature. Thus, the kinetic energy that does not result in displacements would conduct through the body in the form of thermal waves, but would not be sufficient to increase the temperature of the infinite body. This provides a similar effect to that described in [68], where the authors cite the use of a similar periodic boundary condition and thermostat arrangement in addition to an undefined boundary condition that prevents energy that leaves the domain from re-entering it (presumably due to the PBCs).

3.2.2.2 Results and Discussion

The first step of the verification procedure was to examine the PKA evolution found in the present simulation and compare with that reported in [68] and the general physics discussed in Section 1.1.3. Of particular interest was the PKA displacement and kinetic energy as a function of time, as given in Figures 3.3 and 3.4 (see Section 2.3 for details). There are several important things that these plots imply. First, by the end of the initial collision phase (region I), the PKA kinetic energy has reached the thermal background energy (on the order of .1 eV). This indicates that the PKA has lost nearly all (over 99.99%) of its initial energy, likely through collisions with other atoms as it did not enter the thermostat

region. The sharp drops in the kinetic energy indicate that at least some of these collisions resulted in high energy transfers to single or multiple atoms in a short time. These other highly energetic atoms could go on to expand the cascade. Not surprisingly, the time at which a stable PKA displacement is reached late in phase I (see Figure 3.4) corresponds well to when the PKA has lost nearly all of its kinetic energy. Based on these plots, cascade formation is expected and the bulk of the PKA kinetic energy has been lost by the end of the initial collision phase (phase I). The final PKA displacement value is roughly 85 Å from the initial position, and compared very well to the 90 Å reported in [68]. Unfortunately, Figures 3.3 and 3.4 tell us nothing of the actual cascade formation or point defect (vacancy and interstitial) production in the system, nor were plots of this sort published by Devanathan and co-workers. Thus, an examination of point defect creation and evolution was necessary to provide additional verification data.

3.2.2.3 Point Defect Counts

From the raw data obtained from the simulation, point defects were identified as described in Section 2.3.2. The next logical step was to plot the number of defects as a function of time as done in [68]. Details on how these plots were obtained has been given in Section 2.3. The post-processed results from the present work, as well as the corresponding plot from [68] are shown in Figures 3.5 and 3.6.

Figure 3.5 is important in that it verifies that the number of interstitials and vacancies




Figure 3.3: PKA kinetic energy as a function of time from the reproduction of the work of Devanathan and co-workers [68]. The three major regimes of cascade formation are labeled by Roman numerals. Note that by the end of phase I, the PKA has lost over 99.9% of its initial kinetic energy.




Figure 3.4: PKA displacement as a function of time from the reproduction of the work of Devanathan and co-workers [68]. The three major regimes of cascade formation are labeled by Roman numerals. Note that a stable displacement is reached at roughly the same time that the PKA has lost nearly all of the applied kinetic energy.

is exactly the same at every step. This indicates that all of the displaced atoms have been accounted for in the the defect identification stage, as no atoms were gained or lost during the simulation. It also indicates the general trend that is expected based on Section 1.1.3, namely that the bulk of defects are created while the PKA is losing energy due to collisions (phase I). Once the defects have been created, the highly unstable defects relax (i.e. close vacancy and interstitial pairs recombine), until only defects that are stable for the simulated time-scale remain at the end of phase III.

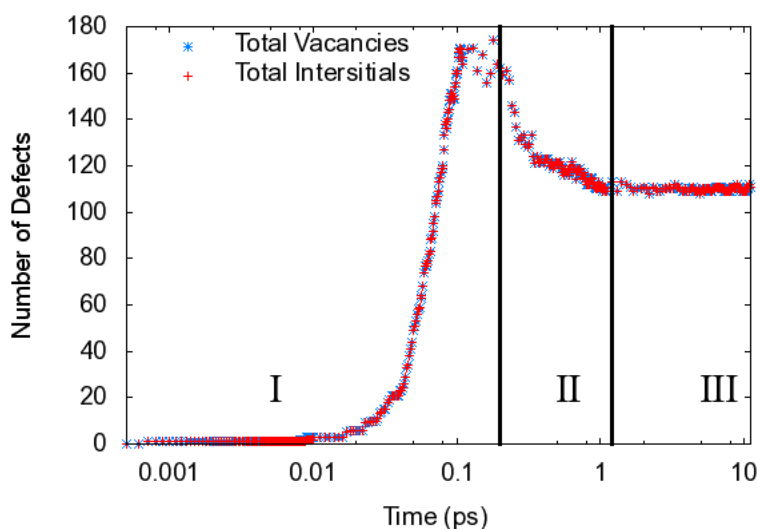


Figure 3.5: Plot of the total number of vacancies and interstitials (C defects plus Si defects) through the simulated time. Roman numerals indicate the phases of cascade formation.

Figure 3.6 supplies information that may be directly compared to [68]. These plots show the evolution of the number of C and Si defects as a function of time. For clarity, only the C and Si vacancies were plotted in the present study. From these plots two things are clear,

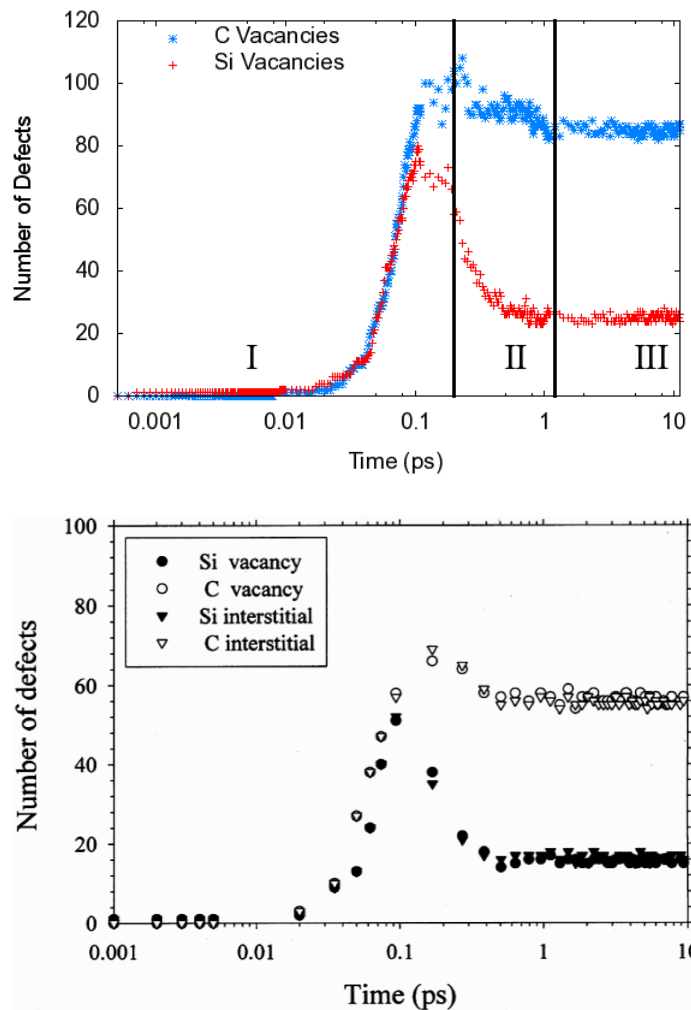


Figure 3.6: Plot of the number of C and Si vacancies through the simulated time. Roman numerals indicate the phases of cascade formation. Also included (on the bottom) is the corresponding plot from [68]. The second plot is reprinted from *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 141(1-4), Devanathan R., Weber W. J. and Diaz de la Rubia T., ‘Computer simulation of a 10 keV Si displacement cascade in SiC’, 118-122, Copyright (1998), with permission from Elsevier.

the first is that the qualitative shape of the evolution is the same, with the peak values at roughly the same time; the second is that the present study shows that the number of Si and C vacancies has nearly a 4:1 C:Si vacancy ratio, compared to the nearly 3:1 ratio reported by Devanathan and co-workers. The present study shows the formation of about 5-10 ($\sim 50\%$) more Si and 25 ($\sim 50\%$) more C ‘stable’ vacancies than in [68].

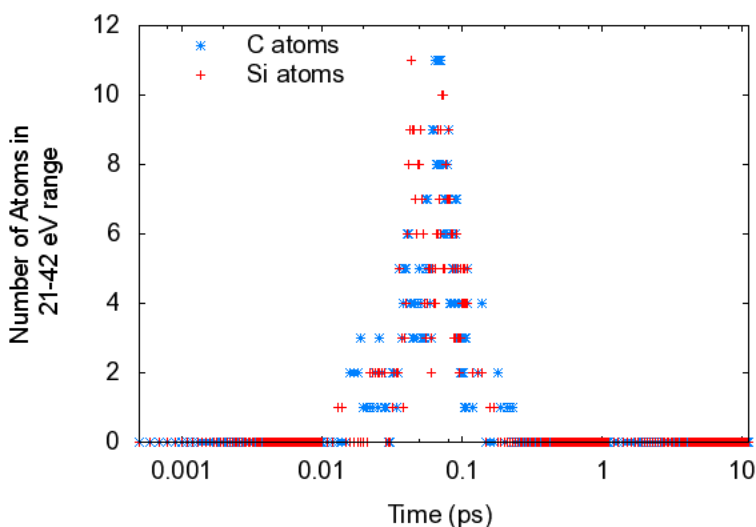


Figure 3.7: Plot of the number of atoms with kinetic energy in 21-42 eV range through the simulated time. This range was chosen as it would indicate the number of atoms that could cause C displacements but not Si displacements.

Devanathan et al stated in [68] that the difference in the number of Si and C defects was due to the nearly 2:1 Si:C average TDE ratio, a conclusion that makes sense intuitively (i.e. C atoms are easier to displace than Si). Table 3.1 indicates that in the present study there is a 2:1 ratio in TDE and Figure 3.7 shows that indeed there are C atoms with kinetic energies lower than the Si TDE, but higher than the C TDE for nearly the entire time the cascade is

forming. Thus this provides some additional evidence for the explanation given in [68] that the difference is due to the ratio of the TDEs. However, based on Figures 3.5 and 3.6, there is likely no quantitative link between the magnitude of the ratio of Si to C defects and the ratio of the respective average TDEs. This conclusion is supported physically as well, as the TDE difference merely indicates that a C atom would be easier (i.e. requires less energy) to displace than an Si atom, but says nothing of the probability of the displacement events in a physical system.

Having been satisfied that the current simulation parameters were as close as could be reasonably expected, a simple implementation of the sphere-based defect identification algorithm was performed based on the criteria for vacancies used by Devanathan et al [68]. Due to the large expense of the sphere-based calculation, only a single output step was used, the end of the initial collision phase I. The sphere based method was found to predict 253 C and 149 Si vacancies, in a similar spatial distribution to those obtained with the Voronoi-cell method but grossly disagreeing in the number of defects identified. Based on these results, it was determined that the differences between the present simulation and that of Devanathan et al were due to the method for identifying the point defects or some nuance in the simulations they performed but did not report. Thus, as the results matched within a reasonable margin and as mentioned before the Voronoi-cell method is a more rigorous method for detecting point defects, it was deemed that the methodology used here was of sufficient quality to continue the study at hand. It is also noteworthy that in a private

communication, Ram Devanathan agreed that the Voronoi-cell method described in Section 2.3.2 for defect identification was superior to the sphere-based method, though it can be more computationally intensive.

3.3 Conclusions

Based on the results obtained in the course of the verification study, it was determined that the simulation methodology was sufficiently close to that used in the literature. The point defect identification scheme agreed in the time behavior, but no comparison could be realistically made with regards to the spatial distribution, as such information is not available in the literature. Thus, now that the methodology has been verified, the next chapter will focus on answering the following questions from the problem posed in Section 1.4:

- *What is the initial distribution of defects?*
- *How do different temperature regimes affect the initial distribution?*

Chapter 4

Characterization of Initial Defect Distribution in a Neutron-Irradiated Perfect Crystal

This chapter looks to answer the fourth and fifth parts of the question put forth in Section 1.4, ‘What is the initial distribution of defects?’ and ‘How do different temperature regimes affect the initial distribution?’. This and subsequent chapters make use of the simulation approach reviewed and verified in Chapter 3. This chapter begins with a basic statistical analysis of point defect distribution after a single simulated neutron irradiation event. This was done to elucidate the ‘average’ damage state in a perfect 3C-SiC crystalline system. The final topic of discussion will be the simulations of cascade formation at different temperatures,

particularly how temperature affects the distribution of defects.

4.1 10 keV Si PKA Cascade in 3C-SiC Perfect Crystal

The simulations that this chapter will focus on were built from the validation work in Section 3.2.2. Nearly all of the details of the system configuration remained the same, with several exceptions. The simulated system is now a cube, 50 cubic cells to a side (1,000,000 atoms total), based on the results of a brief size effect study. The system was setup and equilibrated at a pressure of roughly 100 MPa (1000 bars) and a temperature of 1200 K. This resulted in a lattice parameter of 4.3765 Å and a cell side length of about 219 Å. The temperature and pressure were based on the expected operational values of the ARIES-AT first wall structure discussed briefly in Section 1.4, here they were also used to determine the lattice parameter of the system.

The multiple-phase timestep procedure from Section 3.2.2 was also followed:

- Thermal equilibration, part 1. All atoms in the system have their velocity scaled every step to maintain an average temperature of 300K, for 1000 timesteps with a 1 fs timestep (1 ps total time). This was done to speed the convergence of the thermal equilibration.
- Thermal equilibration, part 2. Interior region (everything except the 4 cell thermostat skin) thermostat removed. Thermostat skin maintained at 300 K. System allowed

to re-equilibrate after the removal of the thermostat on the interior region for 1000 timesteps with a 1 fs timestep (1 ps total time). Removing the interior thermostat was required to prevent non-physical quenching of the cascade in later phases.

- Initial collision phase (labeled in later diagrams as I). PKA instantaneous velocity applied (but not maintained). System allowed to evolve for .2 ps (20,000 timesteps with a .01 fs timestep).
- Intermediate evolution phase (labeled in later diagrams as II). System allowed to evolve for 1 ps (10,000 timesteps with a .1 fs timestep).
- Final evolution phase (labeled in later diagrams as III). System allowed to evolve for 10 ps (10,000 timesteps with a 1 fs timestep). Total time from PKA excitation: 11.2 ps.

Once the initial configurations were generated, a total of 20 runs were performed, each with different thermostat initial conditions but the same target temperature. This was done for each run in LAMMPS by providing different random number generator seeds to the velocity scaling thermostat function in the input files (similar to the file shown in Appendix A.1.2). This particular parameter variation was chosen as it maintained the temperature, but resulted in slight variations in initial positions and velocities of the atoms in the system while keeping the initial conditions of the systems statistically identical. Figure 4.1 shows a schematic of the system configuration. A similar post-processing procedure to that in

Section 3.2.2.2 was followed, with some additional analysis to extract basic information on vacancy cluster formation, the final spatial distribution of defects and the variations between runs, outlined in detail in Section 2.3.




Figure 4.1: Schematic of the geometry used for the simulations of 10 keV Si PKA cascades in a perfect 3C-SiC crystal. The circle represents the PKA initial position.

4.1.1 PKA Displacement and Energy Evolution

As before, the first step was to examine the behavior of the PKA as it provides a basic check of the system evolution. Figures 4.2 and 4.3 illustrate the PKA displacement and kinetic energy behavior from two particular runs of the system. The topmost plots in Figures 4.2 and 4.3 correspond to the same simulation (Run 1), and similarly for the bottom plots (Run 2). In Run 1, the PKA evolution matches the expected behavior based on the results

in Section 3.2. In particular, over 99.99% of the kinetic energy applied to the PKA was dissipated by the end of the first phase, presumably through interatomic collisions as there was no indication of thermostat effects on the PKA. Meanwhile, the PKA in Run 2 only lost 96.33% of its initial kinetic energy in the same time span.

The reason for this discrepancy can be inferred from Figure 4.3. It is clear that Run 2 has a very different PKA response than Run 1. In Run 2, what appears as a discontinuity in the PKA displacement can be explained as a PKA ‘wrap-around’ event, possibly caused by channeling of the PKA. More specifically, the PKA had sufficient energy to pass out of the simulation domain and wrap around to the other side due to the PBCs enforced on the domain. Further, the PKA had to pass through the thermostat region, where the PKA velocity would have been artificially decreased by the thermostat. Thus, the behavior after this wrap-around event is unphysical. Of the 20 cases run, only 2 exhibited this wrap-around behavior, and it is known from the literature that such events can occur (for example [68]). Because the results of these wrap-around runs were unphysical due to the thermostat interaction effects, they were excluded from the statistical analyses to be presented later. This is the approach taken in the literature, as a simulation domain where such events would not occur would be prohibitively large. The drawback to this approach is that this requires one to discard cascades that may have been physically meaningful in a larger simulation cell. However, the channeling and wrap-around events do illustrate the possibility of wide variations in cascade formation with only minor variations in initial conditions. Thus any


characterization of such systems must be done in a statistical sense as any one run may not be truly representative.

Figure 4.4 provides a clear picture of the PKA penetration variation for the runs that did not exhibit PKA wrap-around (18 of the 20 runs performed). To obtain this figure the final PKA displacement for each run was tabulated and plotted with each run represented by a bar. For the 18 runs tabulated, we see a range of roughly 60 Å to 200 Å, with a mean PKA displacement of approximately 121 Å and standard deviation of approximately 37 Å. This is inline with the PKA displacements seen by Devanathan and co-workers in [68]. Again, it is clear that there are variations in the system evolution and thus one must use basic statistical characterization methods for such systems to obtain quantitative information. Later sections will discuss the results of just such an analysis. However, prior to discussing the quantitative results, it is useful to visualize the defect states that resulted from the simulations performed.


4.1.2 Damage State Visualization

After the simulations were performed, the vacancies and interstitials were identified as discussed in Section 2.3.2. From the output generated, states of interest were visualized in Paraview [102]. The reference configuration used for the defect identification code was the atomic positions after thermal equilibration stages but before the PKA excitation was applied.

The particular state of interest here is the final state (end of phase III), after the cascade




(a) Run 1




(b) Run 2

Figure 4.2: Examples of the PKA kinetic energy through the simulation time for the perfect crystal case. These plots corresponds to the same simulations as those in Figure 4.3.



(a) Run 1



(b) Run 2

Figure 4.3: Examples of the PKA displacement through the simulation time for the perfect crystal case. The top figure, illustrates a run in which the PKA is steadily moving in the initial collision phase (I), until it reaches a stable defect site in the intermediate evolution phase (II). The bottom figure, on the other hand, illustrates ‘wrap-around’ of the PKA, where the PKA leaves the simulation domain and ‘wraps’ around to the other side due to the PBCs (hence the drop in the distance from the initial position).

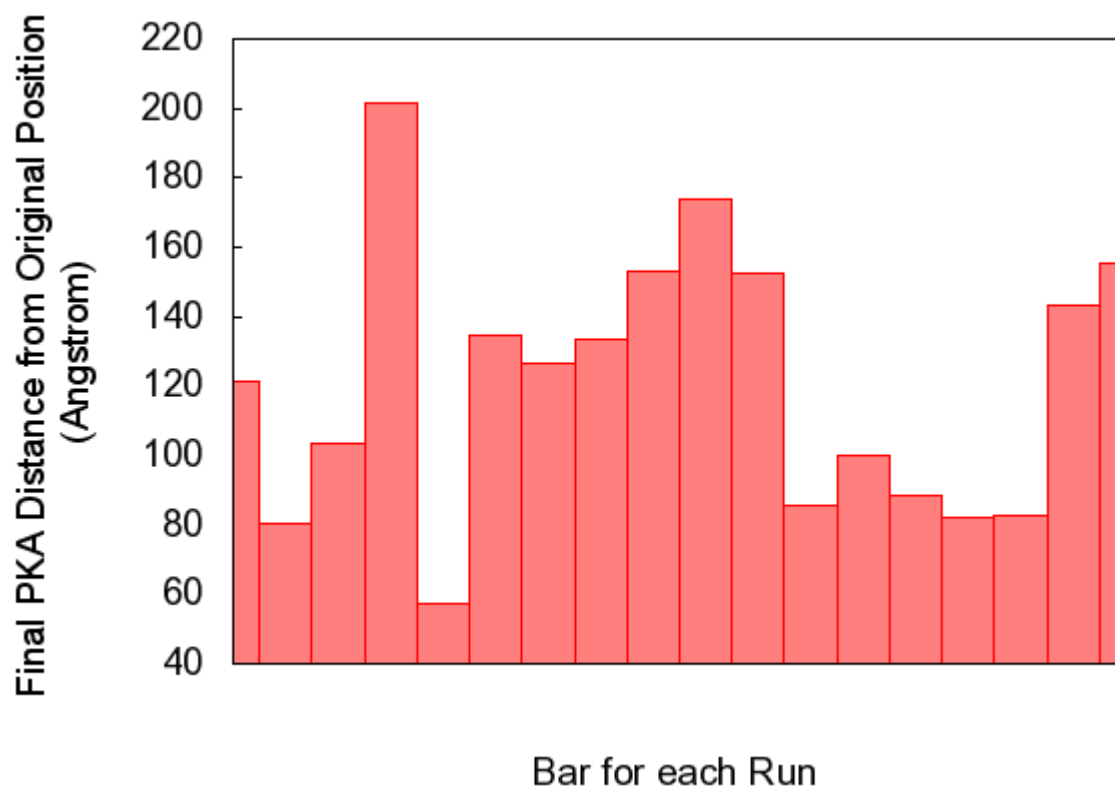


Figure 4.4: Plot showing the final PKA displacement for the 18 non-wrap-around runs. The average displacement was 120.923 \AA and the standard deviation was 37.3552 \AA .

has formed and the unstable defects have relaxed. For illustrative purposes and brevity, only 6 of the 20 runs are illustrated in Figure 4.5. A single example of a wrap-around case (Run 2) is given, in Figure 4.5(b). The views shown have the Z-direction upward, and the camera is looking back along the X-direction toward the origin. The Y-direction is the horizontal, with the origin at the back-bottom-left corner of the visualized domain. It is evident that given nearly identical initial conditions, it is possible to get wide variation in the final defect configuration, possibly due to large thermal vibration effects due to the high temperature. However, it is also clear that there are two trends that stand out - namely, damage clustering at the center of the body such as that in Figures 4.5(a), 4.5(c) & 4.5(e), and what appears to a more distinctly linear damage track such as in Figures 4.5(b), 4.5(d) & 4.5(f).

The remainder of this chapter aims to use several techniques to characterize these defect configurations, so that it is possible to better understand their evolution and general behavior. Before moving to the next section, however, one comparison to the literature is important to make. The simulations performed by Devanathan and co-workers discussed in Section 3.2.2.2 made use of a simulation domain that was nearly 1/4 the size of the one being used here (20x20x60 vs. 42x42x46 here). Yet the present simulations clearly demonstrate damage states that would have exceeded the boundaries of the domain used in their calculations. While it is possible that the differences between the results reported in [68] and the present work can be explained by changes in the lattice parameter and temperature, it seems more likely that the computational power available at the time (1997-2000) required them to

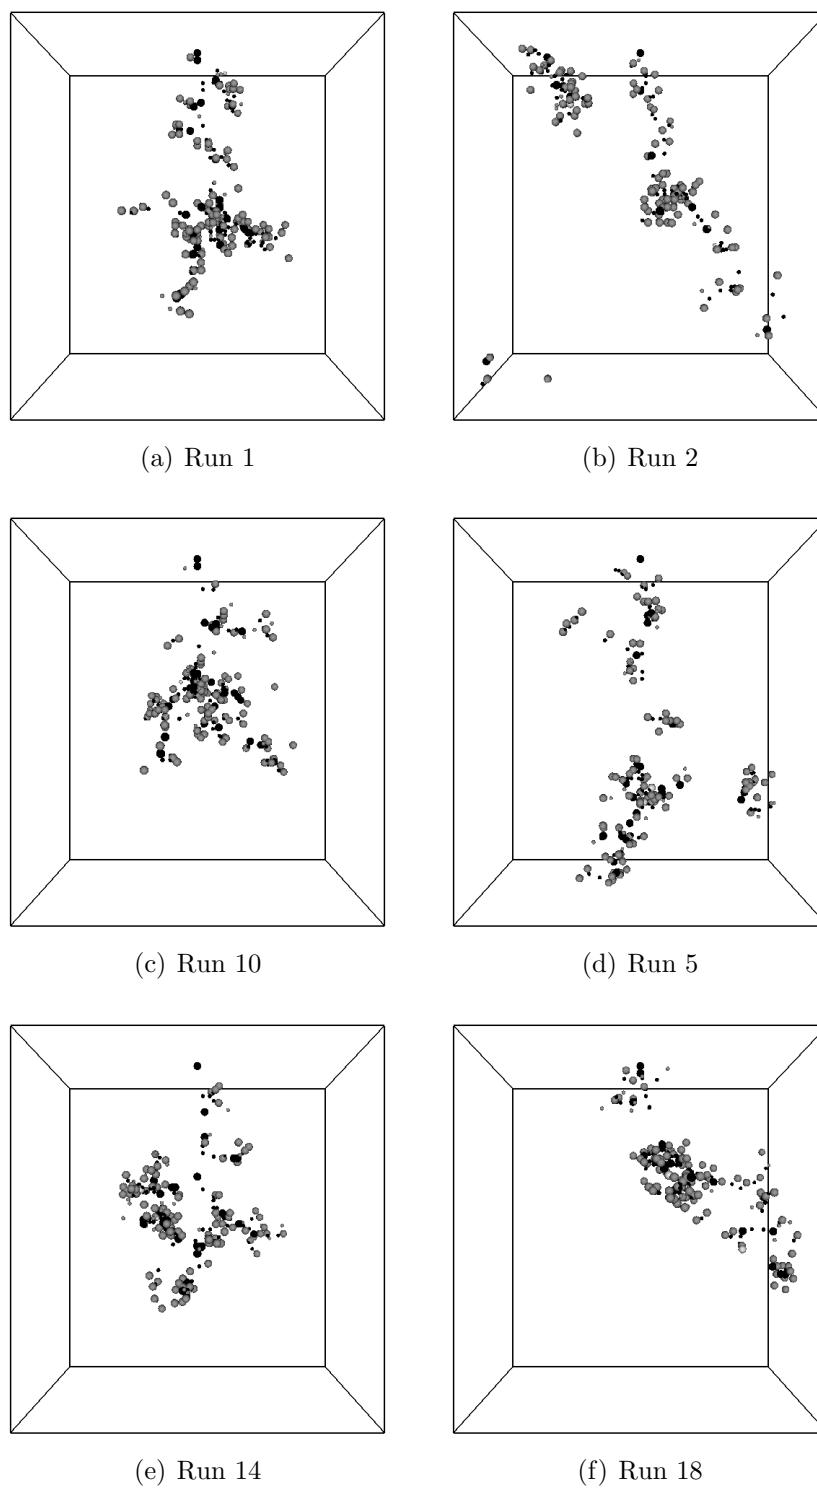


Figure 4.5: Plots of the final defect defect configuration for 6 of the 20 runs. The large spheres are the Si defect sites, and the smaller are the C sites. Dark sites are the vacancies, lighter ones are the interstitial sites. Run 2 demonstrates cascade wraparound.

use a much smaller cell than can be used today. Thus it is likely that they had to discard a large number of runs that would have been physically meaningful in a larger simulation cell. Nonetheless, the thermal vibration effects were studied for completeness and are presented in Section 4.2.

4.1.3 Defect Counts

In the course of the point defect identification, the number of each type of point defect at each output step was tabulated. The defect counts at each output step for all of the 18 non-wrap-around runs were then averaged, tabulated and plotted. The specific point defects tabulated were vacancies, interstitials, antisites and replacements for both Si and C.

The plots of the number of interstitials and vacancies are shown in Figure 4.6. Nearly the same trend can be seen here and in the simulations discussed in Section 3.2.2.2, as expected given the similarities of the systems. The maximum number of defects were created in the initial collision phase (I), and reached a final ‘stable’ number in phase III. In comparison to the results in Section 3.2.2.2, on average there was a greater number of total defects (roughly 130 vs. 120 for the validation calculation), though the individual runs varied by roughly ± 30 defects ($\sim 20\%$) from one-another. Further, more C vacancies formed in the present calculation leading to a roughly 5:1 C:Si vacancy ratio. Despite the increase in the number of C vacancies, the number of Si vacancies remained approximately the same as in the validation case.

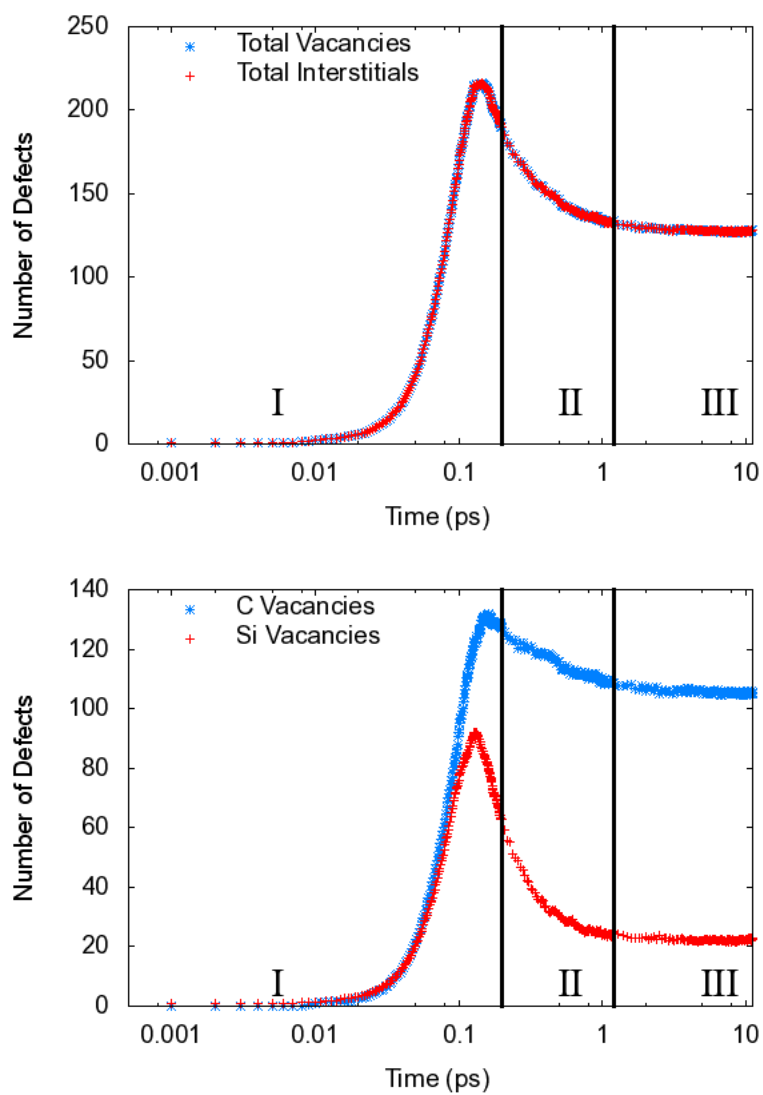


Figure 4.6: Plots of the defect counts averaged over the 18 non-wrap-around runs. Note the qualitative similarity to those discussed in Section 3.2.2.2.

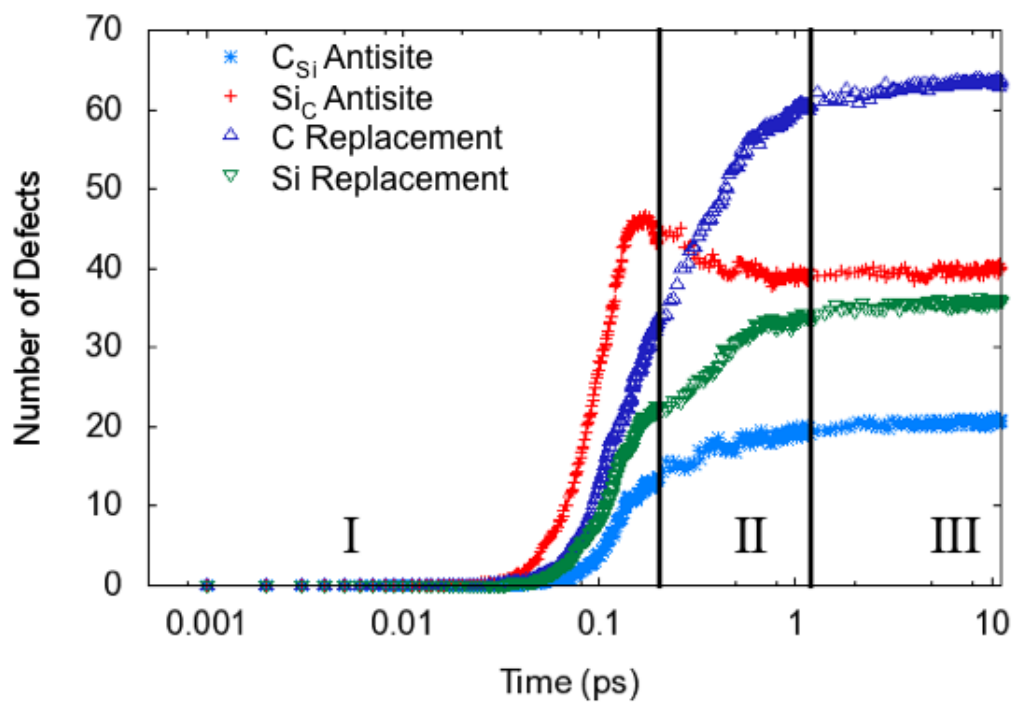


Figure 4.7: Plots of the number of Si and C antisite defects and replacements averaged over the 18 non-wrap-around runs. A replacement is when an atom is displaced from its original position and settles in another site, of the same type. An antisite is when the displaced atom settles in a site of the opposite type. A C_{Si} antisite defect is when a C atom fills what initially was an Si site and vice versa.

Figure 4.7 illustrates the number of replacements and antisite defects formed in the system. A replacement is defined as a displaced atom that settles into an empty lattice site (vacancy) of the same type, and an antisite is when a displaced atom settles in a site of the opposite type. A C_{Si} antisite defect is when a C atom fills what initially was an Si site and a Si_C antisite defect is when a Si atom fills what was a C site. Compared to Figure 4.6, it can be determined that the total number of ‘stable’ replacements and antisites (roughly 150) is almost 25 % higher than the number of interstitials (roughly 130). It is also evident that C replacements are more common than Si replacements. Additionally, Si_C antisite formation occurs more frequently than C_{Si} antisite defect formation and Si replacements. Further, it appears that the majority of the antisites formed are stable for the timescale studied in these calculations. The only decrease occurs in the Si_C antisites, where a 10% reduction after the initial collision phase is seen, possibly due to the instability of these defects when near an Si vacancy [119]. The existence of so many antisites also has the implication that one may find small (4 or more atoms) localized regions of Si or C that have not necessarily formed into distinct ‘inclusions’ nor become a locally amorphous structure. However, further investigation of the stability of Si_C antisites and the formations of local Si and C regions is left as future work.

Based on the total number of vacancies found, it was concluded that the formation of a large vacancy cluster (on the order of 100 vacancies) is possible. However, the plots shown thus far only offer a cell-total time history of the number of defects present in the system.

What is of greater interest in void formation is the distribution of these defects within the body. The spatial distribution of vacancies is of particular interest, as the hypothesis that this dissertation works from is that vacancies contribute to void formation in irradiated materials.

4.1.4 Vacancy Distribution and Clustering

Vacancy clusters were identified as discussed in Section 2.3, based on a tree-type neighbor search (find the neighbors of neighbors, etc. in a given radius). Here the search radius was chosen as 2.2 Å, the cutoff range for Si-Si interactions, and just larger than the roughly 1.9 Å nearest neighbor spacing, but smaller than the second neighbor spacing. Once this operation was performed, the number of clusters and the minimum, average and maximum cluster sizes were averaged and plotted (see Figure 4.8). On average, the runs tended toward roughly 3 (± 2) clusters of about 3 (± 2) vacancies. Compared to the number of vacancies at the end of phase III in Figure 4.6, it is apparent that the vast majority of vacancies at the end of the simulation are ‘isolated’ vacancies, i.e. they are not near other vacancies, by this metric. This is in general agreement with the conclusion reported in [114] for 5 keV cascades in 3C-SiC at 300 K - however, they used a third-nearest neighbor (roughly 4.36 Å) criterion. Thus it is possible that they detected more or larger clusters than is seen in this work, nevertheless, the conclusion is the same.

In order to get an idea of how isolated an ‘isolated’ vacancy is at the end of the simulations,

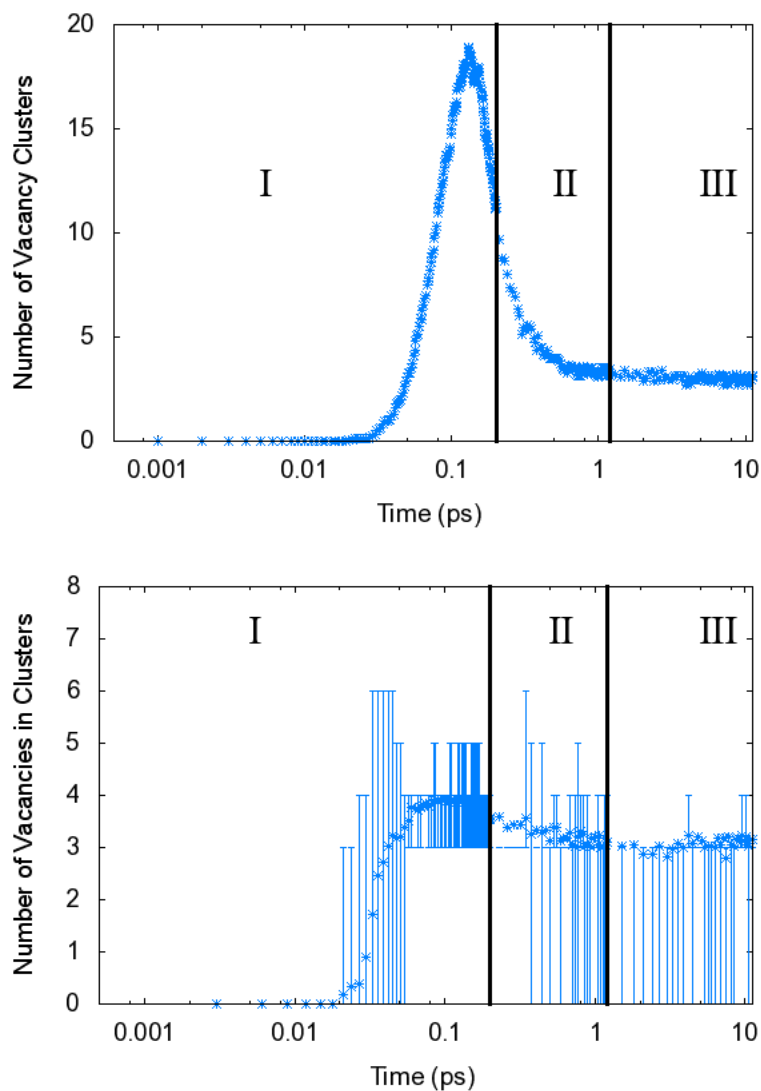


Figure 4.8: Plots of the number and size of vacancy clusters averaged over the 18 non-wrap-around runs. The errorbars in the second plot indicate the maximum and minimum cluster sizes in the runs. The majority of clusters consist of 3 vacancies, though a small minority of runs saw no stable clusters forming.

a body-averaged vacancy radial distribution was calculated for the last step of each run as described in Section 2.3. This gives a measure of the average distance between all of the vacancies in the system. An average over all of the non-wrap-around runs was then calculated to get an ‘average’ RDF for the final distribution. This distribution is given in Figure 4.9. From the figure it is clear that, on average, the majority of the vacancies are outside of the first and second nearest neighbor ranges (the first two boxes in Figure 4.9). This indicates that the majority of vacancies are indeed isolated, with average separations larger than the third nearest neighbor range. Finally, the average number of vacancies drops to zero above about 200 Å (the domain is only about 219 Å per side) and the form of the distribution resembles a log-normal distribution. Thus it is clear that these defects are not entirely closely clustered, nor are they uniformly distributed in the body. One possible explanation for this relates to the physics of cascade formation. In particular, when a high-energy transfer collision occurs, the interstitial created will ‘inherit’ a high kinetic energy. Therefore, it is possible that the defect pair created by such a collision may separate enough that recombination will not occur easily. Further, atoms with high kinetic energies may travel far from their original site before reaching the thermal background energy, colliding with other atoms along the way and may also move far enough to make stable defect pairs. In this manner, a series of dispersed vacancies could be generated initially, that could then migrate through the body.

With the findings from Figure 4.9 in mind, it was of interest to determine an averaged

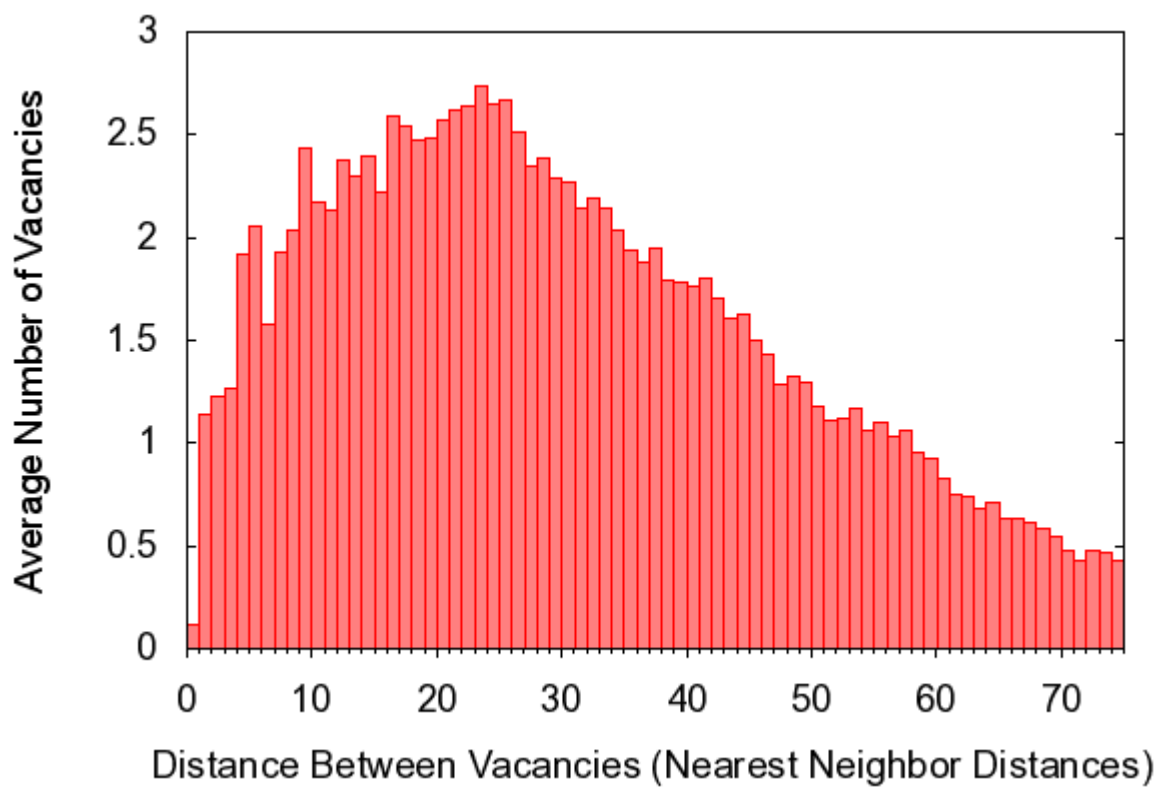


Figure 4.9: Plots of the body-averaged vacancy radial distribution averaged over the 18 non-wrap-around runs. Note that it is clear that the majority of vacancies are much further apart than the first neighbor distance (roughly 1.89 \AA - the first box on the left).

spatial distribution of the vacancies. This was done for two reasons, the first was an attempt to quantify the spatial distribution of the vacancies, something which when tracked through time would indicate clustering. The second reason was to make it possible to generate defect configurations for long-time simulations without resorting to EP simulations of cascade formation; cutting down on the overall cost of the simulation. Through a simple sort-and-bin operation of the vacancies as described in Section 2.3, it was possible to obtain a histogram in the X,Y and Z directions ([100],[010] and [001], respectively). The results, averaged over all 18 non-wrap-around runs are given in Figure 4.10. The error bars indicate the limits of the variation of the runs from the average.

What is clear from the figures is that there is wide variation in the distributions between runs. However, the shape of each is similar, therefore some trends can be extracted. The X ([100]) and Y ([010]) directions (essentially perpendicular to the direction of the PKA initial velocity) show a distribution not unlike a Gaussian distribution, an indication that the damage primarily occurs in the center of the simulation domain - not surprising as the general path of the PKA is in this direction. On the other hand, the Z-direction ([001]) has a maximum near the center of the cell, decays to zero near the bottom of the cell (0 on the horizontal axis), and is nonzero near the top of the cell. This was also expected, as the PKA and the atoms it collided with during the initial collision phase would likely have been energetic enough to travel some distance and cause the cascade to expand. This would result in a ‘damage track’ (see Figure 4.5) caused by the permanent displacement of the highly




Figure 4.10: Plots of the run-averages spatial distribution of vacancies for each orthogonal direction. The X and Y plots have a Gaussian fit overlaid on them for reference. The origin is a bottom corner of the simulation cell, so larger Z values indicate the top of the cell.

energetic atoms, despite the recovery of some of the cascade-induced damage. Nevertheless, the error bars in the figures make it evident that a very large number of simulations would likely be needed in order to get a spatial distribution that could be considered a good approximation to an ‘average’ cascade. However, even with the small number considered here we can begin to see the formation of a trend. Thus it is still possible to answer the question ‘What is the initial distribution of defects in a perfect crystal?’: there is a nearly Gaussian distribution in all 3 directions, as indicated by a larger amount of damage toward the center of the cell than at the periphery.

4.2 Temperature Effects in 10 keV Si PKA Cascade in 3C-SiC

The previous section addressed the question of the defect distribution in an irradiated perfect crystal. This section will address the effects of temperature, specifically thermal fluctuations of position, on the damage configuration in an identical perfect crystal system. To do this, a temperature study was done where the system was run at 5 additional temperatures. The system geometry, potential and timestep used in the study were identical to the system used in Section 4.1. While the temperature differed, the simulation volume was kept constant and thus the nominal lattice parameter would remain the same as before. This particular parameter variation allowed for the isolation of thermal effects on the initial damage state.

For each temperature, 10 runs were performed with different thermostat initial conditions (i.e. random number generator seeds). The particular target temperatures simulated were 0 K, 10 K, 100 K, 500 K and 2000 K. These results were then combined with the 1200 K results from Section 4.1.

4.2.1 PKA Displacement

Once the simulations were completed and post-processed as in Section 4.1, an examination of the PKA evolution was performed. At every temperature the decay of the PKA kinetic energy was nearly identical to that seen in the results shown in 4.1; thus it is of little value to show them here. Despite this, there are a few interesting aspects of the PKA evolution that merit discussion, such as the number of possible wrap-around events and PKA displacement.

Figure 4.11 shows the percentage of the 10 runs at each temperature that demonstrated PKA ‘wrap-around’ and Figure 4.12 presents the average and spread (maximum minus minimum) of PKA final displacements for each temperature. At 0 K, all of the PKA displacements are identical, approximately 100 Å and exhibit no wrap-around. This is not surprising given that the relative positions of the atoms in each run would change very little due to the small thermal energy. However, As the temperature (and therefore relative atomic displacements) increases to above 0 K, the development of a spread of final PKA displacements is seen (by 10 K). Figure 4.12 shows that the spreads remain similar for nearly all of the non-zero temperatures simulated. If either the 10 or 100 K spread are taken as an outlier, one can




Figure 4.11: The percentage of runs that exhibited wrap-around of the PKA (i.e. possible PKA channeling) at each temperature. Percentage used as the 1200 K case consisted of 20 runs while the others consisted of 10.

conclude that the PKA displacement spread shows no significant dependence on temperature. Another interesting conclusion from these plots is that for the particular PKA initial velocity direction used here, the average PKA displacement at 0 K may not be the minimum value. However, a detailed study with different PKA excitation directions would be needed to further support this conclusion and is left as future work.

4.2.2 Point Defect Counts

The next logical step was to examine the number of point defects through the simulated time and temperatures, as in the previous sections. An average over all of the runs at each temperature was computed and plotted as discussed in Section 2.3. Each of the specific

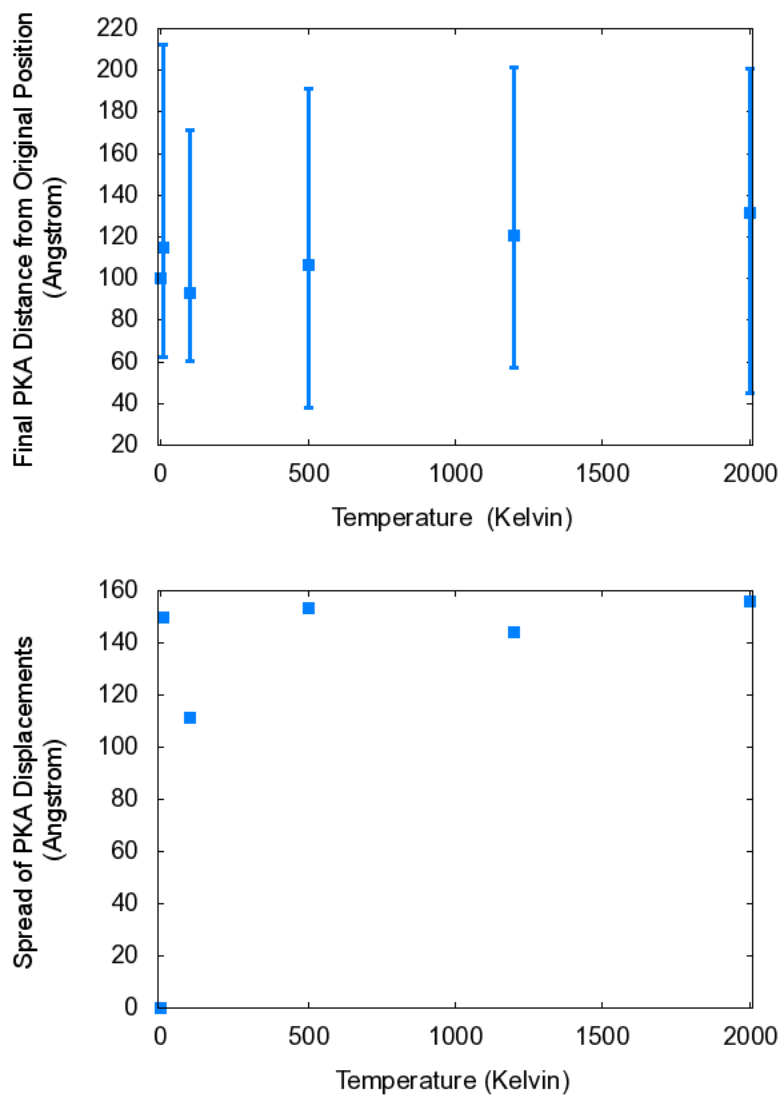


Figure 4.12: Plot of the final PKA displacement, averaged for all runs at each temperature. The errorbars show the spread of the data at each temperature. The 0 K case showed all PKA displacements to be equal. Note that the spreads for the non-zero temperatures are nearly the same though the average changes.

defect counts (total vacancies, Si vacancies, C vacancies, Si replacements, etc.) were then plotted together for each of the temperatures to elucidate any trends or lack thereof.




Figure 4.13: Plot of the run-averaged total vacancy counts for each temperature. There is little indication of temperature dependence on this measure, as the majority of the results do not vary significantly.

Figure 4.13 illustrates the run-averaged total number of vacancies at each temperature through the simulation time. From the figure, it is clear that while there is some small variation (roughly 20 vacancies) between the 0 K to 1200 K cases, the behavior is largely the same at each temperature. The 2000 K case is an exception, however it only varies by a maximum of 47 vacancies ($\sim 27\%$) at the peak and 16 ($\sim 10\%$) at the last step though the number of vacancies is clearly decreasing throughout phase III. By comparing each individual run at each temperature, it was found that the total defect variations in Figure 4.13 are within the variation seen between runs at the same temperature. The exception was the 0 K case,

where all runs yielded the same defect evolution. Despite this, there is a clear difference in the averages that suggests a slightly different behavior for the 2000 K case. Similarly, if one examines Figure 4.14, it is evident that there is no clear trend in the temperatures as the variations are within that seen in individual runs at a given temperature. Not surprisingly given the total vacancy plots, the 2000 K case exhibited different behavior in the individual counts than the other temperatures. Further, it is clear that the decrease in the total number of vacancies in phase III is due to a decrease in the number of C vacancies. This is likely due to C interstitial-vacancy recombination, as the thermal background energy at 2000 K is on the order of the migration barrier for C interstitials (roughly .75 eV [120]). Despite this difference at the highest temperature, it was concluded that there is no clear temperature dependence in the total number of vacancies nor the number of each type generated. The differences in the 2000 K case does indicate the possible existence of a high-temperature cascade regime where thermal effects actually decrease the effective damage due to the cascade. However, further study of this high temperature regime is left as future work.

As before, the number of vacancies only tells part of the story. Thus, it is useful to look at other defect types such as antisite defects and atomic replacements, shown in Figures 4.15 and 4.16 respectively. Unlike the vacancy counts, these plots do show trends that indicate a temperature dependence. The number of C_{Si} antisites generated in the intermediate and final evolution phases (II & III) clearly increases with temperature, and only in the 0 K case does the number clearly decrease from the value at the end of initial collision phase




Figure 4.14: Plot of the run-averaged vacancy counts for each temperature and atom type. There is little indication of temperature dependence on this measure as well.

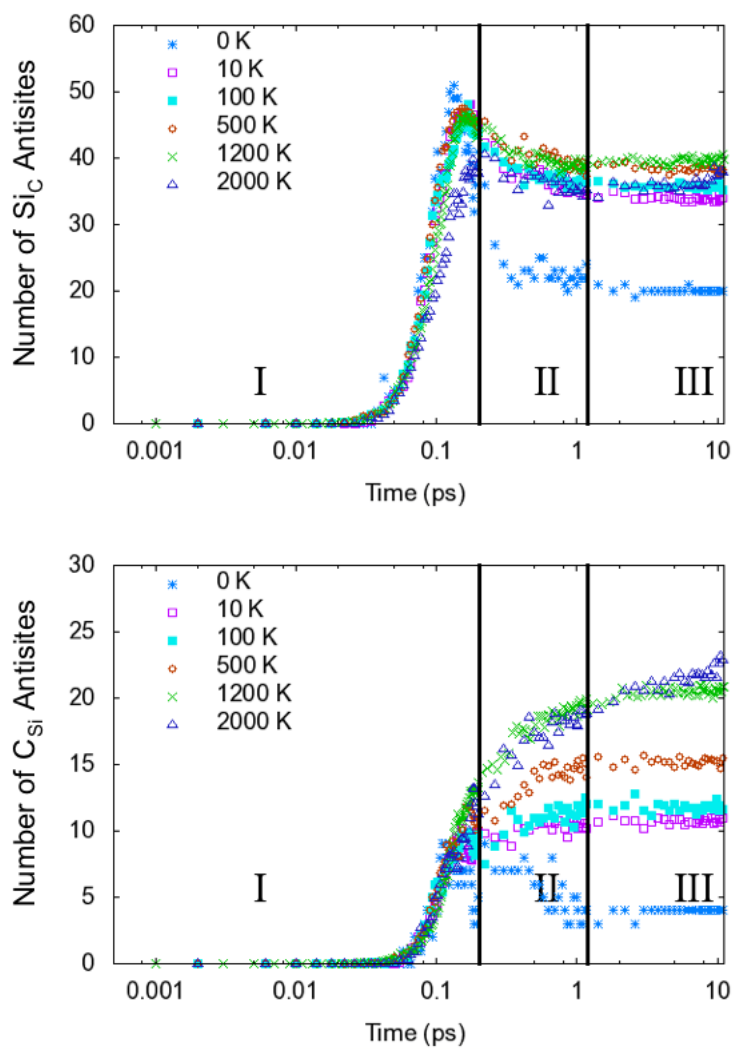


Figure 4.15: Plot of the run-averaged antisite defect counts for each temperature and defect type.

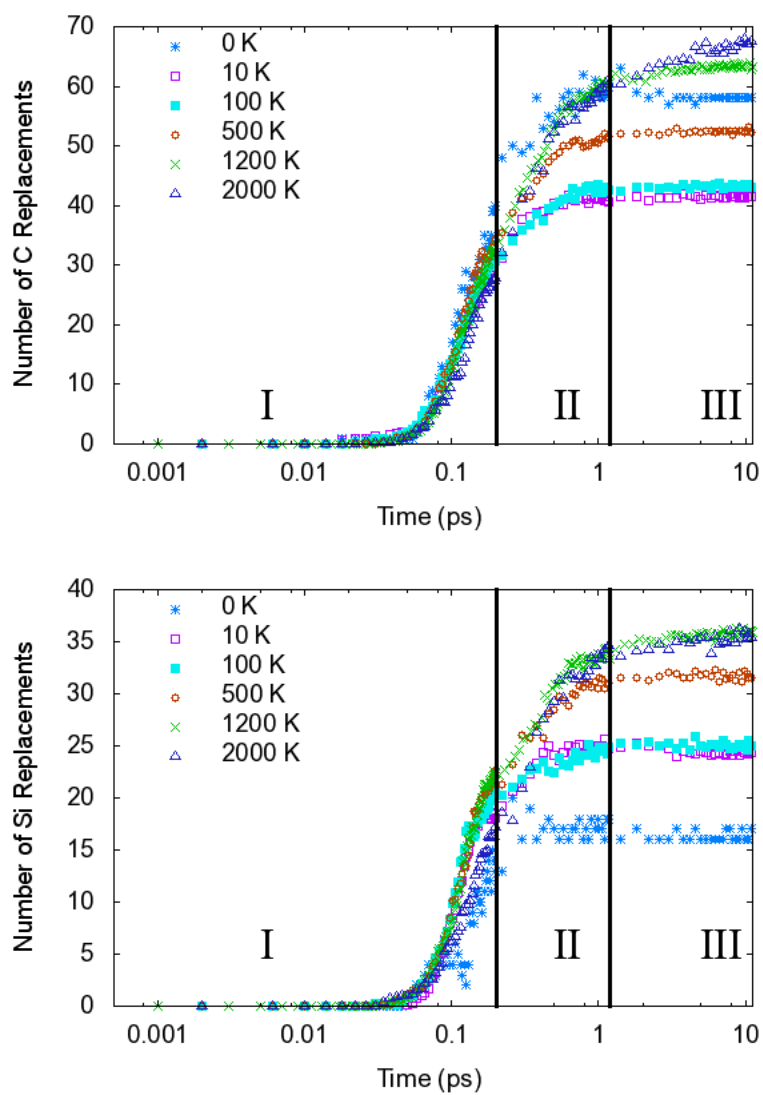


Figure 4.16: Plot of the run-averaged atomic replacement events for each temperature and defect type.

(I). Further, while the 0 K through 1200 K cases seem to have leveled out, the 2000 K case appears to increase throughout phase III. Meanwhile, there is less indication of temperature dependence on the formation of Si_C defects. In particular, the maximum number formed (in phase I) is nearly constant through all temperatures, though the final number increases with temperature except in the 2000 K case. In all cases, there is some amount of relaxation that occurs in phases II and III, though this effect is slightly reduced at higher temperatures. For nearly all the temperatures (0K C_{Si} defects are the exception), the relative stability of the antisites formed may be explained by the fact that forming an antisite defect from an existing interstitial-vacancy pair is more energetically favorable than the interstitial-vacancy pair[121]. However, this does not explain the decrease in the number of Si_C defects in phase II and increase in the number of C_{Si} defects throughout the simulations. The decrease in Si_C defects is possibly due to thermal vibrations kicking the Si atom out of the C site, allowing a nearby C atom to form a replacement. Similar trends appear in Figure 4.16, where the number of C and Si replacements after phase I increases with temperature, though more C replacements were seen at 0 K than at 10 K and 100 K. For all temperatures, more C replacements were found than C_{Si} antisites, though the number of Si replacements and Si_C antisites differed only by a margin of roughly 33 %. Finally, the increase in the number of C replacements in phase III is further evidence of the role of thermal vibrations in C vacancy-interstitial annihilation. Further study of this, along with the role of antisite formation and antisite-replacement swaps in cascade relaxation is left as future work.

4.2.3 Vacancy Clusters and Distributions

The final aspect of the simulations to examine is the vacancy clustering data as a function of time and temperature, based on the procedures used in Section 4.1. The average vacancy cluster number and size are given in Figures 4.17 and 4.18, respectively. The cluster definition is identical to that used in Section 4.1. An additional measure was introduced for this study by the author, the clustering percentage, to quantify the percentage of the total number of vacancies that are part of a cluster. This data is given in Figure 4.19.




Figure 4.17: Plot of the run-averaged number of vacancy clusters for each temperature.

From an examination of Figure 4.17, it can be seen that the final number of clusters increases as the temperature increases from 0 K to 500 K, after which the number drops. Thus, at the highest temperatures there are only a few clusters of vacancies (an average




Figure 4.18: Plot of the average vacancy cluster size (number of vacancies per cluster) for each temperature.

of about 3-5 clusters). Further, Figure 4.18 indicates that for all temperatures the average cluster size is between 2 and 4 vacancies. This indicates that at the higher temperatures, only a few stable clusters form and those are small clusters of only a few vacancies. A possible reason for fewer clusters at the highest temperatures is that the thermal energy allows for greater defect and atomic mobility, thus there is more energy available to breakup unfavorable defect configurations after they are formed by the cascade. This possibility is further evidenced by the fact that the number and average size of the vacancy clusters is nearly constant after phase II, when the cascade has dissipated most of its energy and the system approaches a stable configuration through thermally driven defect migration and annihilation.


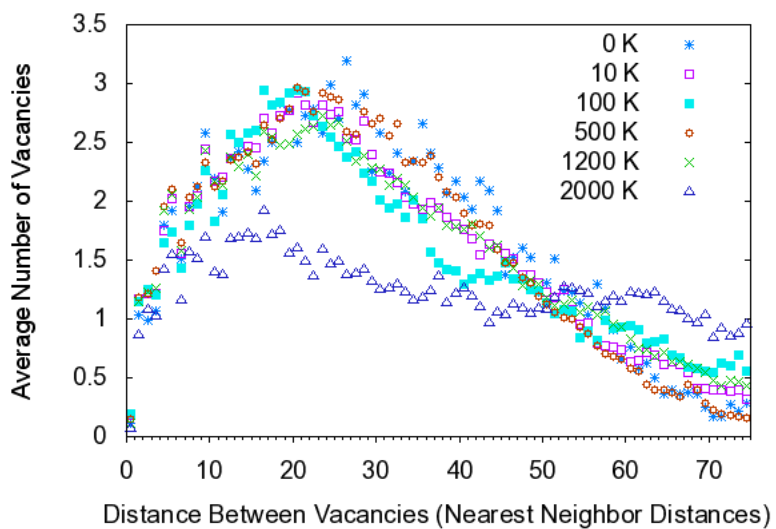


Figure 4.19: Plot of the average percentage of vacancies found in clusters at each temperature.


The results from the previous paragraph led to the question of what percentage of the vacancies in the system are actually in clusters as opposed to isolated or in Frenkel pairs (interstitial/vacancy pairs). Figure 4.19 answers this question. There are two points of note in this figure. Firstly, the peak percentage is roughly the same for all temperatures and occurs at very nearly the same time. This time corresponds to the maximum number of defects generated, i.e. the peak of the cascade formation from Figure 4.13. There is a similar trend to that seen in Figure 4.17; the final percentage of clustered vacancies increases from about 12 % at 0 K to about 30 % at 100K, when it steadily decreases until a minimum is reached at roughly 10 % at 2000 K. These results indicate the possibility of a moderate temperature at which a maximum amount of clustering occurs in the simulated timescale. Further, these

conclusions reinforce the findings of the previous section that the bulk (upwards of 70 %) of the vacancies in the final defect configuration are isolated from one another, and that this is the case regardless of the temperature. In fact, at the highest temperatures the effect is magnified to 80-90% isolated vacancies.

In order to further elucidate the thermal effects on the final spatial distribution of vacancies, the body-average vacancy radial distribution function was calculated for each run and has been averaged over the runs at each temperature (Figure 4.20). A similar trend to that seen in the clustering percentage is repeated in these figures, particularly at the small separations shown in Figure 4.20(b). While the distributions for the 0 K to 1200 K cases are similar, it is evident that the final relative distribution of vacancies in the 2000 K case is more uniform in space than at the lower temperatures. However, it is unlikely that thermal diffusion alone accounts for the difference, as even at 2000 K few atoms would have the kinetic energy to overcome the high potential energy barriers for bulk vacancy diffusion (roughly 4.1 eV and 2.3 eV for C and Si vacancies respectively[120]). This is reinforced by Figure 4.21, where it is evident that the vacancy distribution at the end of the initial collision phase (I) is more uniform than the other temperatures. Thus, it is more likely that recombination of interstitial/vacancy pairs during cascade formation plays a greater role as interstitials are known to be more mobile than vacancies in SiC[120]. Specifically, at 2000 K it is more likely that an atom would have a kinetic energy on the order of the interstitial migration barriers (roughly .75 eV and 1.5 eV for C and Si interstitial sites respectively)



(a)



(b)

Figure 4.20: Plot of the averaged radial distribution function for vacancies each temperature at the end of the simulation. (b) shows the detail of the RDF for separations less than 4 nearest-neighbor distances (roughly 7.58 \AA).






Figure 4.21: Plot of the averaged radial distribution function for vacancies each temperature at the end of the initial collision phase.

than at lower temperatures. This would also explain the decrease in the total number of vacancies and interstitials during phase III for only the 2000 K case shown in Figure 4.13.


The final point of comparison for the temperature study was the final vacancy spatial distributions. These results are illustrated in Figure 4.22. There is little clear indication of temperature dependence in the final configuration, and the distributions are largely similar. There is some indication of a trend towards more diffuse damage in the Z-direction as temperature increases, but the difference is only on the order of 1 percent between the low and high temperature cases.



(a)



(b)



(c)

Figure 4.22: Plots of the final vacancy spatial distribution for each temperature and direction.

4.3 Conclusions

This chapter began with the presentation of a set of 10 keV Si PKA cascade simulations in a system larger than that used previously in the literature to further reduce boundary effects. The only variation between the runs was the initial velocities of the atoms, though the velocities were chosen to maintain the same temperature. Once this set of simulations was performed, a temperature effect study was undertaken to explore 5 additional temperatures, from 0K to 2000K, using the same system setup as the 10 keV Si PKA cascades. These results were post-processed and visualized to identify key features and characterize the resultant damage state in an averaged sense.

From the data obtained from the large 10 keV Si PKA cascades, it was possible to draw several conclusions. The first is that despite a nearly 2:1 ratio in the Si:C TDE, a C:Si vacancy ratio of nearly 5:1 was found in all simulations. Thus, it is clear that the TDE alone does not tell the entire story of defect formation. It was concluded that lattice parameter and thermal fluctuations can change the TDE and cascade behavior, possibly dramatically. This was further evidenced by the variation among the 20 runs presented in Section 4.1. However, the time evolution of the number of defects was very similar between runs, regardless of the defect spatial configuration. From the study of vacancy clustering it was found that at the end of the simulations, the bulk of vacancies were isolated from one another with only a few small clusters formed. Thus, the likelihood of forming a large cluster that could be considered a ‘void’ based on the definition given in Section 1.4 is unlikely in the timescale

considered here. There is the possibility that multiple cascade events could lead to a greater amount of cluster formation, but simulating realistic neutron fluxes would require timescales many orders of magnitude greater than that which EP can achieve. The analysis did reveal a nearly Gaussian average defect distribution in the plane perpendicular to the initial PKA velocity (XY plane in this case), though in the depth direction (Z) it was not Gaussian. Thus, there was a tendency for defects to form near the center of the simulation cell - not entirely surprising given the PKA initial velocity vector and the direction of travel for the initial cascade atoms. Further, the damage distribution visualization revealed something of note in comparison to [68]. The simulations performed by Devanathan and co-workers discussed earlier made use of a simulation cell that was nearly 1/4 the size of that used here. Yet a number of the generated damage states would have approached or exceeded the boundaries of the domain used in their calculations. It was concluded that the computational expense of larger simulations at the time of [68] prevented the authors from using an appropriately large simulation cell. Thus, it is likely that size effects biased their results.

The temperature study performed yielded several interesting points. The first was that the PKA behavior was largely independent of temperature. At 0 K there was no spread in the response of the PKA and the spread at higher temperatures varied little between the non-zero temperatures. The point defect generation and evolution behavior indicated the possible role of enhanced C interstitial-vacancy recombination at 2000 K, a phenomenon that was not readily apparent at lower temperatures. Beyond this, there was little indication

of temperature dependence in the number of vacancies and Si_C antisites generated in the regimes studied. On the other hand, C and Si replacements and C_Si antisite formation did show a weak temperature dependence. The number of C and Si replacements increased with temperature, though non-linearly. The number of C_Si antisites was significantly higher at higher temperatures, likely due to enhanced atomic mobility. Analysis of the defect clusters indicated that regardless of temperature, the number of vacancies found in clusters was relatively small (10-30 % of the total number of vacancies). Additionally, at small ranges (1-4 nearest neighbor distances), there was little evidence of temperature dependence in the relative distribution of vacancies. Despite this, it was seen that the distribution of vacancies over long ranges (greater than 4 nearest neighbor distances) at 2000 K was significantly more uniform than at the lower temperatures. This is further evidence of the possible existence of a high-temperature regime where atomic and defect mobility effectively decreases the amount of concentrated damage in the system.

Chapter 5

Characterization of Initial Defect Distribution in a Neutron-Irradiated Crystal with Initial Void Structure

The bulk of Chapter 4 was concerned with the quantitative characterization of the ‘stable’ post-cascade defect structure in a neutron irradiated 3C-SiC perfect crystal. Particular attention was paid to the distribution and clustering of vacancies through time, to gain insight into the processes of vacancy clustering during displacement cascade formation and relaxation. It was concluded that there was relatively little evidence of large cluster formation in the perfect crystal on the simulated timescale, regardless of the irradiation temperature.

Due the manufacturing processes used to produce bulk SiC structures and surface coat-

ings (e.g. Chemical Vapor Deposition (CVD) or powder compaction) it is unlikely that perfect crystalline structures would actually be used in nuclear applications. Based on the experimental and manufacturing literature, SiC surface coatings and bulk structures can contain on the order of .25-1 % porosity (from United States Patent #5332601 for an example of a low-porosity CVD technique). In addition to process-induced defect considerations, there are defect behavior theories that indicate voids may act as vacancy ‘sinks’[55, 57, 122, 123]. Therefore, it is possible that vacancies would tend to migrate toward existing voids or vacancy clusters and slowly increase their size. Thus, it was of interest to examine the changes in the distribution of defects during cascade formation and relaxation due to the presence of an initial void structure. In other words, this chapter seeks to answer the question: ‘*How does an initial void structure affect the defect distribution?*’.

A pre-existing void can be considered a model for a structure with a manufacturing defect or one that contains a vacancy cluster from some previous irradiation event. Based on [55, 57, 122, 123] and the physics discussed in Section 1.1.3, there are several possible void structure roles in cascade formation and relaxation:

- Void acts as a vacancy sink. During the relaxation phases, vacancies may ‘attach’ to the void, decreasing the number that are free to migrate and increasing the size of the void.
- Void acts like a free surface. Acts to lengthen/redirect cascade when an atom travels through the void and strikes another on the other side.

- Cascade causes structure change of the void. Interstitials and vacancies may join with the void, but not in the same places, changing the shape of the void.

In order to investigate the different possible roles, this chapter will discuss several different system configurations. The first is a single large spherical void in the center of a cubic domain like that in discussed in detail in Section 4.1. In order to investigate how the position of a large void affects cascade formation and penetration, a second set of simulations were performed with a large spherical void near the top of the domain. These studies led to the the question of how defect arrays would affect cascade formation and relaxation, thus simulations were also performed with a 3-void linear array in the center of the domain and a 27-void cubic array, uniformly dispersed in the domain. Each section in this chapter will discuss one of these sets of simulations, with an emphasis on how they differ from the perfect crystal simulations of Section 4.1. To the authors knowledge, the work performed for this chapter was the first empirical potential molecular dynamics study of the impact of an initial void structure on cascade production and relaxation in neutron irradiated 3C-SiC.

5.1 Simulation Configurations and Methodology

The basic system configuration and simulation methodology used as a base for all of the ‘imperfect crystal’ simulations in this chapter was identical to that in Section 4.1 and is shown schematically in Figure 5.1:

- Tersoff/ZBL hybrid potential with Devanathan et al parameters (Table 2.1).
- Crystal structure: 3C-SiC (zinc-blende), 50x50x50 unit cells. (001)-type boundaries. Lattice parameter of 4.3765 Å.
- 4 cell thick thermostat ‘skin’ on all sides but the top. Target temperature of 1200 K.
- Primary knock-on atom (PKA) chosen to be a Si atom near the top, center of the cell.




Figure 5.1: Schematic of the geometry used as a basis for the simulations of 10 keV Si PKA cascades in imperfect 3C-SiC crystals. The circle represents the PKA initial position. The void structures were created in the ‘simulation domain’, as discussed in the text.

From the base geometry, the void geometry was extracted by deleting atoms from the base geometry using a built-in feature in LAMMPS. This was performed prior to thermal

equilibration, to ensure that the void structure was able to relax prior to the PKA excitation. 4 specific void configurations were investigated, with a void size determined based on the experimental and manufacturing literature. There, one can find estimates of porosity (i.e. volume of the voids divided by volume of the sample) of roughly .25-1 % (for example United States Patent #5332601). From this range and the simulation cell just described, the following void configurations were obtained using a constant void volume fraction of .2 %:

- 1 spherical void with a radius of roughly 1.6 nm (4 unit cells).
- 3 spherical voids in a linear array, each void of roughly 1.2 nm radius (2.7 unit cells).
- 27 spherical voids in a cubic array, each void roughly .6 nm radius (1.3 unit cells).

Two versions of the single void case were performed, one with the void in the center of the simulation cell, the other with it centered in the XY plane, but $3/4$ of the cell length from the bottom (i.e. closer to the PKA than the centered case). the 3 void linear array consisted of voids equally spaced in the Z-direction but centered in the XY plane. The 27 void cubic array consisted of evenly distributed voids arranged in a cubic lattice. These configurations are given in Figure 5.2.

For each void configuration, 10 runs with varied thermostat initial conditions and constant target temperature were performed. Each simulation was carried out in several phases, in the same manner as those in Section 4.1:

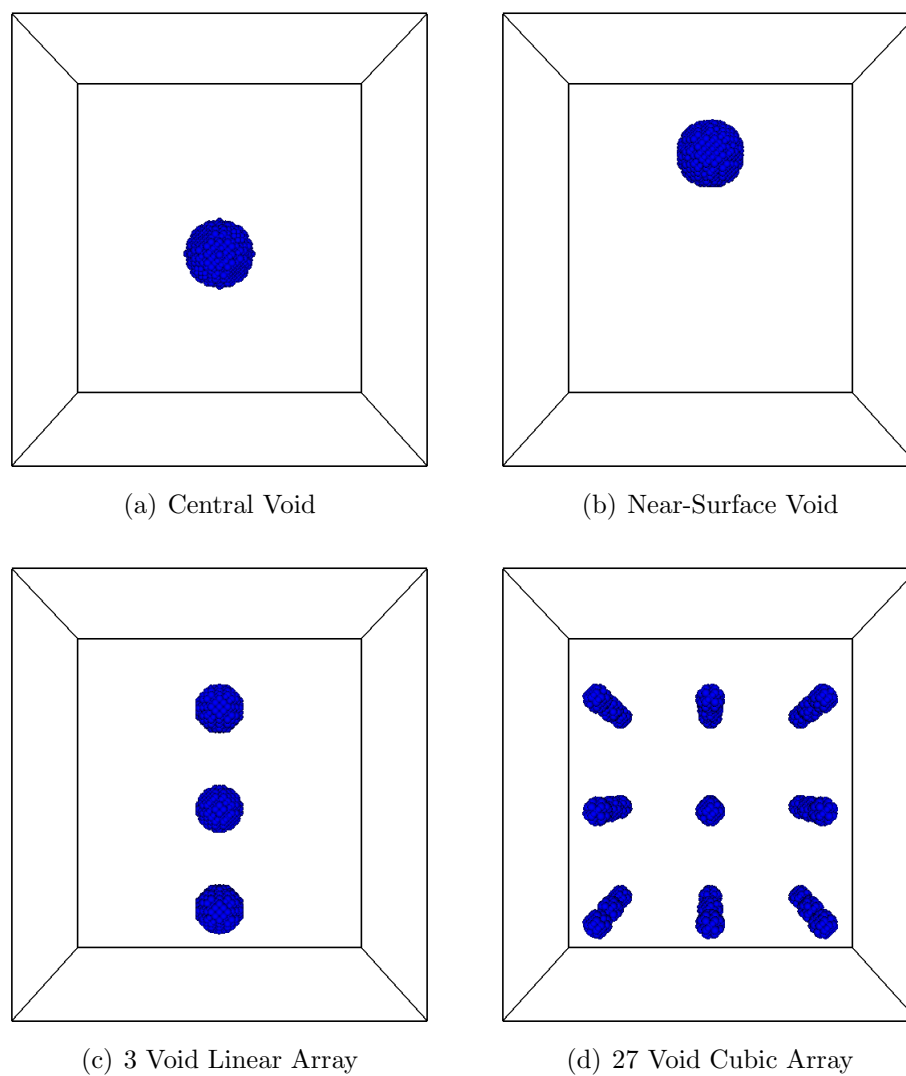


Figure 5.2: Initial void configurations used in this chapter. Only initial vacancies have been visualized here for clarity. Note that the boundaries shown do not include the thermostat region, and thus the void structures may not appear centered in the domain.

- Thermal equilibration, part 1. All atoms in the system have their velocity scaled every step to maintain an average temperature of 300K, for 1000 timesteps with a 1 fs timestep (1 ps total time).
- Thermal equilibration, part 2. Interior region (everything except the 4 cell thermostat skin) thermostat removed. Thermostat skin maintained at 300 K. System allowed to equilibrate for 1000 timesteps with a 1 fs timestep (1 ps total time).
- Initial collision phase (labeled in later diagrams as I). PKA instantaneous velocity applied (but not maintained). System allowed to evolve for .2 ps (20,000 timesteps with a .01 fs timestep).
- Intermediate evolution phase (labeled in later diagrams as II). System allowed to evolve for 1 ps (10,000 timesteps with a .1 fs timestep).
- Final evolution phase (labeled in later diagrams as III). System allowed to evolve for 10 ps (10,000 timesteps with a 1 fs timestep). Total time from PKA excitation: 11.2 ps.

Once the simulations were performed, the same Voronoi-cell based point defect identification methodology was employed as in Section 4.1. The perfect crystal reference geometry was used, such that the void structure appears in the initial and subsequent configurations. A very similar quantification procedure to that used in Section 4.1 was used for the clustering analysis. The specific differences will be discussed further in the relevant section.

	Ave. Displacement (Å)	Std. Deviation (Å)
Central Void	132.928	49.1819
Near-Surface Void	161.196	25.3887
Perfect Crystal	120.923	37.3552

Table 5.1: Table of PKA final displacements. Average and standard deviation given to accompany Figure 5.3. The text provides information on how many runs were used for each.

5.2 Effects of an Initial Void on Defect Configuration in 10 keV Si PKA Cascade in 3C-SiC


The first set of simulations to be performed were the set that considered two different positions of a single, large spherical void. The void remained centered in the XY plane, but one configuration had the void positioned much closer to the PKA along the depth (Z) direction.

From an examination of the PKA behavior, the decay in kinetic energy was much the same in both void configuration cases as one would expect from the perfect crystal case in Section 4.1. In the central void case, 3 out of the 10 runs exhibited wrap-around of the PKA. In the near-surface void case, only 2 out of 10 showed signs of this. In the runs that did not exhibit PKA wrap-around, the biggest difference in the PKA behavior was in the PKA final displacement (Figure 5.3 and Table 5.1). On average, the central void configuration resulted in a roughly 10 % increase in the PKA displacement over the perfect crystal case (approximately 133 Å compared to 120 Å), while the near-surface void configuration resulted in an increase closer to 20 % (approximately 161 Å versus 120 Å). This increase does fit with the qualitative difference expected from intuitive arguments, i.e., the gap in the material

would allow for greater penetration. Further, the larger penetration in the near-surface case agrees with intuition, i.e. the void being closer to the initial position of the PKA, the atoms which travel through it would maintain larger kinetic energies through a greater distance than in the central void case because of the reduction in the number of early collisions.


As with the perfect crystal, the next step was to visualize the final damage state, to obtain a better idea of how the cascades formed. Three examples from each initial void configuration are shown in Figure 5.4, for comparison to Figure 4.5. These figures show that there is a wide variation between the runs, much like Figure 4.5. That there are two dominant trends, one where the cascade intersects or approaches the void and one where it largely misses the void (such as that shown in Figure 5.4(c)). It was observed that the void did indeed affect the cascade form, most noticeably in the cases where the cascade intersected the void.

The first quantitative measure employed to examine the evolution of the damage structure was the evolution of the number of point defects through the simulation time. As in Section 4.1, the defect counts of the runs that did not exhibit PKA ‘wrap-around’ for each void configuration were averaged to obtain the general trends. These defect counts were then plotted, as given in Figures 5.5 and 5.6. In this set of calculations, only the net number of vacancies and interstitials were plotted. This was due to an algorithmic limitation in the point defect identification code that would not allow for the reliable definition of replacements or antisite defects in these cases. The code could have been re-written but due to this




Bar for each Run

(a) Central Void



Bar for each Run

(b) Near-Surface Void



Bar for each Run

(c) Perfect Crystal

Figure 5.3: Plot showing the final PKA displacement for the non-wrap-around runs of the single void cases. The perfect crystal results from Section 4.1 are shown for comparison.

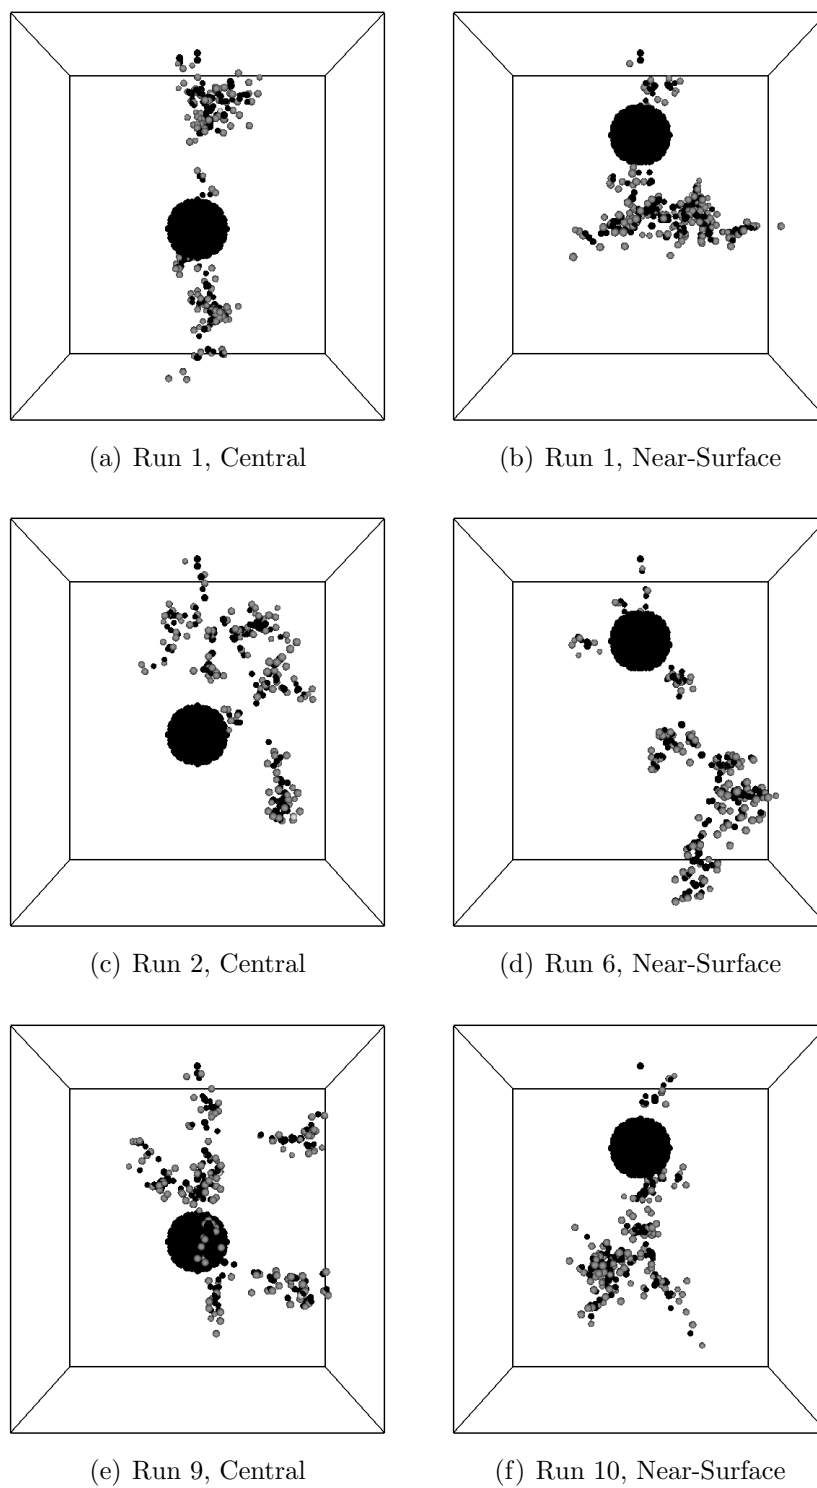


Figure 5.4: Plots of the final defect defect configuration for 6 of the single void runs. The large spheres are the Si defect sites, and the smaller are the C sites. Dark sites are the vacancies, lighter ones are the interstitial sites.

dissertation's focus on vacancies and vacancy clusters, it was deemed unnecessary. The net number of vacancies was determined by calculating the total number at each timestep and subtracting the number of vacancies initially in the void. This was done to obtain a measure of the change in the number of defects in the system due to the cascade, such that it can be compared to the perfect crystal case. The cluster analyses later will discuss the change in the size of the void and its effect on cluster formation.

These two sets of figures provide an interesting comparison to the perfect crystal case in Figure 4.6. The total number of defects at the peak and end of phase III as well as the general form of the curve are nearly identical to the perfect crystal case. On the other hand, there are fewer C vacancies formed in phase I in the imperfect crystal cases, and there is a small difference in final number as well (roughly 5% fewer C vacancies in the imperfect cases). The ratio of the types of surviving defects remains nearly 5:1 C:Si vacancies, just as in the perfect crystal case. Thus it is concluded that for this measure, there is no significant difference between the perfect and imperfect cases.

In order to quantify the impact of the void on clustering, the same cluster-identification and radial distribution scheme employed in Section 4.1 was applied to these results. From these results, several plots were then generated, one limited to the maximum cluster size at each output step (this would indicate a change in the void size), one that gives the average size of the smaller clusters (i.e. clusters of vacancies not associated with the void) and the total number of clusters in the system. These plots are shown in Figures 5.7 and 5.8. From




Figure 5.5: Plots of the defect counts averaged over the non-wrap-around runs for the single central void. Note the qualitative similarity to those discussed in Section 3.2.2.2 and 4.1 (i.e. Figure 4.6). The net number of defects is defined as the total number minus the initial number, as discussed in the text.




Figure 5.6: Plots of the defect counts averaged over the non-wrap-around runs for the single near-surface void. Note the qualitative similarity to those discussed in Section 3.2.2.2 and 4.1 (i.e. Figure 4.6). The net number of defects is defined as the total number minus the initial number, as discussed in the text.

these plots, it is evident that there is little stable growth of the void. In both cases, the variation of the void is only ± 2 vacancies, less than .1 %. Because of this, it is concluded that no significant void growth or shrinkage occurs due to the cascade on the timescale studied here. On the other hand, as in the perfect crystal case there are smaller clusters generated which remain in the system for the duration of the calculation. The average size of these clusters in the imperfect crystal cases are slightly smaller than in the perfect crystal case (~ 1 vacancy smaller - a $\sim 20\%$ difference). However, there is a nearly 3 fold increase (compared to the perfect crystal case) in the number of clusters present in the with-void cases. Despite this, based on Figure 5.9 the central void configuration has nearly the same percentage of clustered vacancies (not including those in the void) as the perfect crystal case. Meanwhile, the near-surface void configuration results in roughly 10% fewer vacancies clustering (again, of those not associated with the void) than the perfect crystal, though agrees well at the peak ($\sim 60\%$). It is suspected that the greater energy (relative to the same depth in the central void or perfect crystal case) of the atoms that formed of the cascade in the near-surface led to a more dispersed defect configuration being initially generated by the cascade. However, it seems clear that the presence of the void did not increase the amount of clustering significantly, thus the void does not act as a vacancy sink over the timescale simulated here.

In order to further explore the possibility that the near-surface void resulted in a more dispersed cascade than the central void case (and by extension, the perfect crystal case),

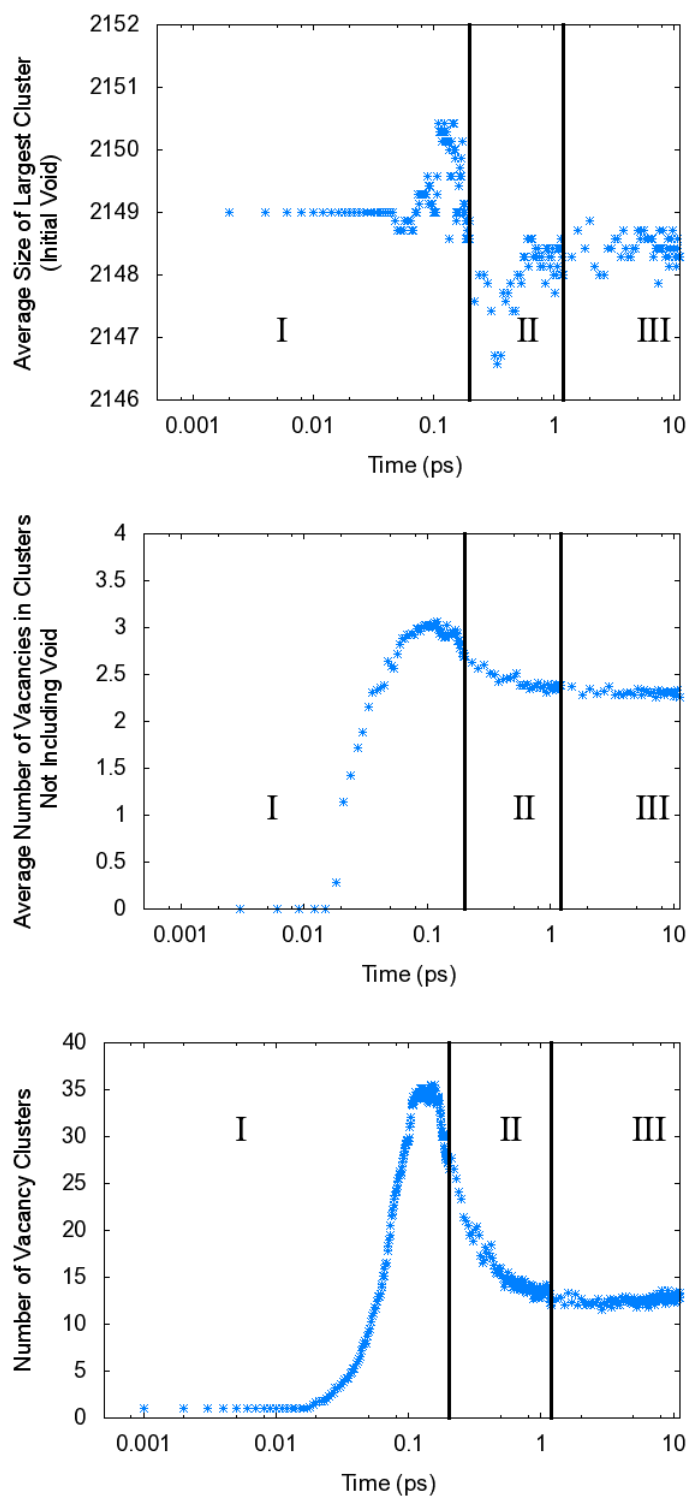


Figure 5.7: Plots of the number of vacancies in the largest cluster (i.e. the void), the average number of vacancies in the smaller clusters, and the number of clusters in the central void system.

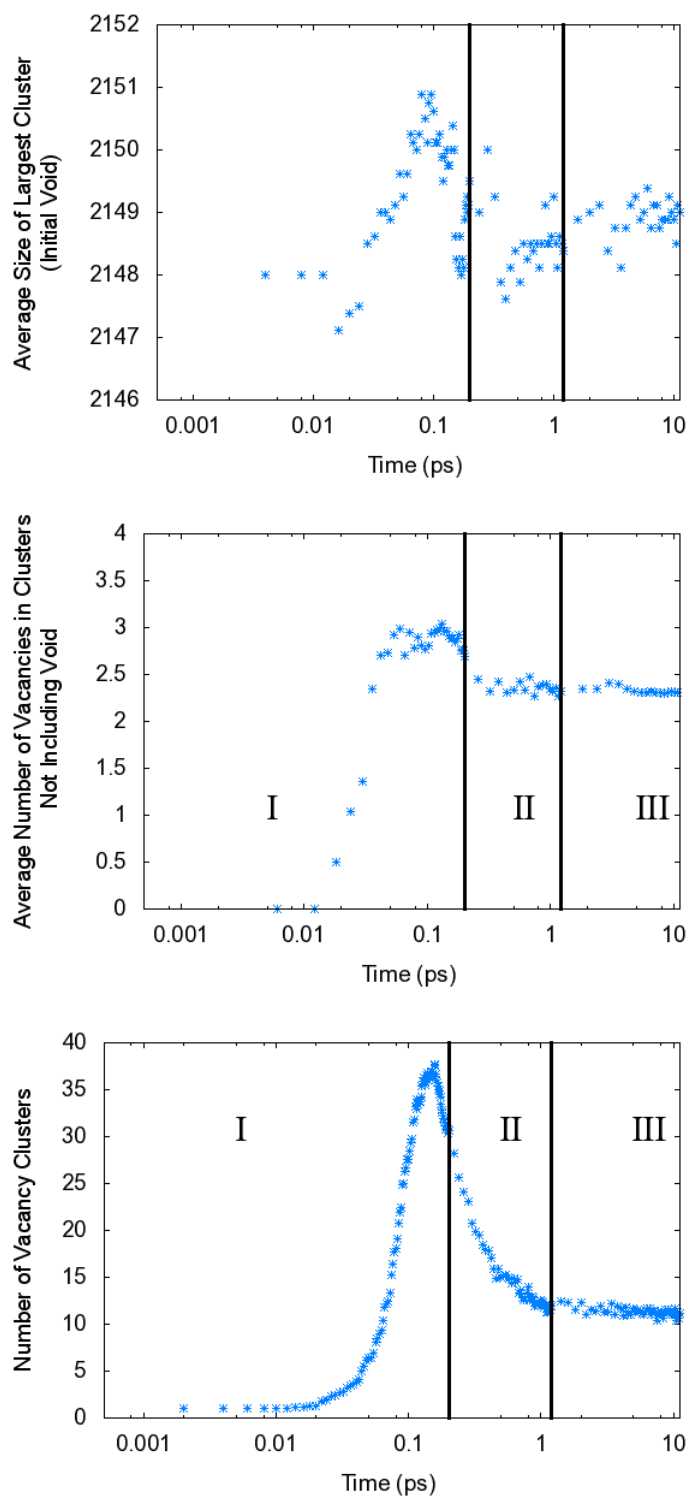




Figure 5.8: Plots of the number of vacancies in the largest cluster (i.e. the void), the average number of vacancies in the smaller clusters, and the number of clusters in the near-surface void system.



(a) Cental Void




(b) Near Surface Void

Figure 5.9: Plots of the percentage of the vacancies that are not associated with the void, but still in a cluster, for each void configuration.


a ‘net’ RDF was calculated by subtracting the RDF of the initial structure from the RDF calculated at the desired output step, to isolate the influence of the cascade from the initial void structure. The net RDFs from the end of phase I are given in Figure 5.10. From these figures, the near-surface void case clearly has a more dispersed vacancy configuration than the central void (i.e. more vacancies at further separations). For reference, the void radius corresponds to approximately 8 nearest neighbor distances.

5.3 Effects of Initial Void Arrays on Defect Configuration

The previous section briefly explored the impact of a single, large initial void on cascade formation, as well as how the position of the void impacted defect generation and the final configuration. In order to make the current study of void-structure effects more complete, it was desirable to study the impact of a more complex void structure on cascade formation. To do this, two additional defect structures were considered as mentioned earlier, a 3-void linear array and a 27-void cubic array as shown in Figure 5.2. These structures were chosen as they would allow for the exploration of void structure effects on the extent and distribution of damage due to irradiation. Based on intuition and the results of the previous study, it was expected that both of these structures would exhibit a clear increase in the penetration and dispersion of damage in the structure. Further, it was expected that these cases would



(a) Cental Void



(b) Near Surface Void

Figure 5.10: Plots of the net radial distribution function for vacancies in the single void configurations at the end of phase I.

demonstrate greater wrap-around frequency but the cell size was maintained to allow direct comparisons. As in the other cases, examining the effects of larger cell sizes than used here has been left as future work.

5.3.1 3 Void Linear Array

The first initial void structure studied was the 3 void linear array shown in Figure 5.2. The same simulation procedure was followed as in the single void cases, i.e. 10 runs were performed and those that exhibited PKA wrap-around were not considered in the analysis to follow. Of the 10 runs, only 4 *did not* exhibit PKA wrap-around. Of those that exhibited PKA wraparound, the PKA typically exited near the ‘bottom’ of the simulation cell, significantly deeper penetration than in the previous cases. The PKA final displacements for the non-wrap-around runs are given in Figure 5.11. The average penetration for the non-wraparound cases was approximately 131 Å, not significantly different than the central void case. However, if one includes the wraparound penetrations (i.e. assume the PKA stopped at 218 Å), the average would be closer to 175 Å, approximately a 10% increase over the near-surface void configuration.

The individual defect counts for the non-wrap-around cases are given in Figure 5.13 and the final defect state of 2 runs visualized in Figure 5.12. These plots did not differ significantly (less than 10% for the peak and final values) from the central void and perfect crystal cases. The only feature of the plots that is of note is the decrease in the number of Si vacancies






Figure 5.11: Final displacements for the PKA in the 4 non-wrap-around cases in the 3 void linear array configuration. Note that those that did exhibit wrap-around had the PKA leaving the bottom of the cell (greater than about 218 \AA displacement along $[00\bar{1}]$). For reference, the average displacement was 131.071 \AA and the standard deviation was 5.73316 \AA .



(a) Run 3



(b) Run 8

Figure 5.12: Visualization of the final damage state of two of the non-wrap-around runs. The large spheres are the Si defect sites, and the smaller are the C sites. Dark sites are the vacancies, lighter ones are the interstitial sites.

through phase III, however, the decrease is small enough to be considered insignificant, more so because of the relatively small number of runs. Compared to the perfect crystal case, the peak is much less defined, though that is attributed to noise due to the small number of non-wrap-around runs.

As in the single initial void cases, no significant growth or shrinkage of the initial voids was found based on the clustering data. A variation of 2-4 vacancies (less than 1% change) was observed, likely due to defects being created near the void, but these tended to relax or migrate away through the simulation. Further, the number of clusters formed was not significantly different than the previous cases. Because of this, these plots are not shown. Thus, an examination of the spatial distribution based on the vacancy radial distribution function is all that remains. The ‘net’ RDF was calculated by subtracting the initial RDF of the initial structure from the RDF calculated at the desired output step, to isolate the influence of the cascade. These ‘net’ RDFs are given in Figure 5.14 along with the initial RDF. Based on the results, it is evident that the cascade produces a relatively dispersed damage state (i.e. vacancies more uniformly distributed along a series of distances greater than a few nearest neighbor separations), and does slightly reduce the size of the voids (whose defect separations are noted by the 3 large bumps in the initial RDF plot). However, as mentioned earlier this reduction is relatively small throughout the simulation (less than 5% change in the RDF) and unstable, thus is concluded to be insignificant. Additionally, there seemed little indication of vacancies that connected the voids into a single large defect




Figure 5.13: Plots of the defect counts averaged over the non-wrap-around runs for the 3-void linear array configuration. Note the qualitative similarity to those discussed in Section 3.2.2.2 and 4.1 (i.e. Figure 4.6). The net number of defects is defined as the total number minus the initial number, as discussed in the text.


structure.

5.3.2 27 Void Cubic Array


The final initial void structure used was a cubic array of 27 small voids as shown in Figure 5.2. This structure had two aims in mind; the first was to see if the cascade extents were further increased, the second if the smaller voids were able to be changed by the cascade.

Based on the previous sections, it is not surprising that the PKA kinetic energy decay was very much the same as in previous cases. Of the 10 runs performed only 1 exhibited PKA wrap-around, with displacements given in Figure 5.15. It is expected that additional runs would yield additional wrap-around cases. Of the 9 non-wrap-around cases, the average displacement was only about 145 Å, due to 2 cases where the PKA displacement was relatively small (less than 100 Å). It is expected that additional runs would yield higher PKA penetrations that would increase this average displacement.


Three examples of the final defect configuration are given in Figure 5.16. These display 2 interesting trends, much like those seen in the single void cases. In particular, Run 8 (Figure 5.16(c)) displays a linear damage track that largely misses the void structure, and approaches the boundary. The second, illustrated by Runs 1 and 4 (Figures 5.16(a) and 5.16(b)), demonstrates the cascade interacting with at least 1 void clearly, though missing most of the others either due to a linear damage track or spreading out more in the horizontal direction (X-Y plane). The behavior seen in Figure 5.16(b), where the PKA has only traveled



(a) Initial RDF



(b) Net RDF, end of phase I



(c) Net RDF, end of phase III

Figure 5.14: Plots of the ‘net’ radial distribution function for vacancies in the 3 void linear array configurations at the end of two key phases of cascade formation. The function plotted is the RDF from the timestep (averaged over all non-wraparound runs) minus the initial RDF (prior to PKA excitation). The initial RDF is given for reference.




Figure 5.15: Final displacements for the PKA in the 9 non-wrap-around cases in the 27 void cubic array configuration. Note that the case that did exhibit wrap-around had the PKA leaving the bottom of the cell (greater than about 218 \AA displacement along $[00\bar{1}]$). The average displacement was 145.172 \AA and the standard deviation was 48.377 \AA .

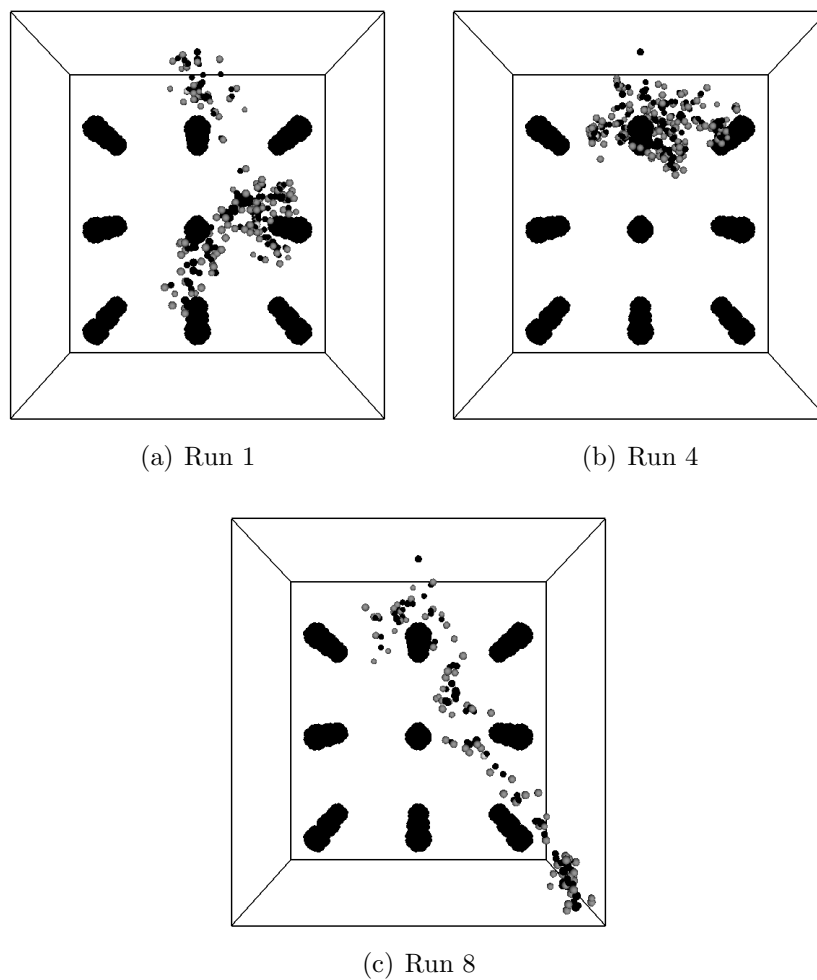


Figure 5.16: Visualization of the final damage state of three of the non-wrap-around runs for the cubic void array configuration. The large spheres are the Si defect sites, and the smaller are the C sites. Dark sites are the vacancies, lighter ones are the interstitial sites.

about 80 Å, is likely due to the direction of the energetic atoms that entered the void (or those involved in the cascade formed above the void). More specifically, if the thermal vibrations of the atoms resulted in an oblique impact, there could be a significant horizontal component to the atoms' velocities. Thus, rather than penetrating deeply, they cause the damage to be spread out closer to the PKA's original location. However, this applies equally to the other cases, so more work is required to determine if this is a significant trend.

Compared to the previous cases, the net vacancy and interstitial counts (given in Figure 5.17) are not significantly different in the peak or final values. In comparison to the 3-void run, there is a more gradual decrease in the net number of C vacancies in phases II and III, and the Si vacancy number remains largely stable through phase III. However, in comparison to the single void cases and the perfect crystal cases, the rate of reduction in the number of C vacancies is similar. Thus, it is concluded that the relatively small number of runs averaged to obtain the 3-void case plots are the reason for the difference. In other words, it is suspected that the plot of the 3-void case is not as representative as the plots in the other cases. Based on this, it is concluded that the presence of the void structure has relatively little effect on the net number of vacancies produced by the cascade. Further, as in the 3-void case, there was little indication of a change in clustering behavior due to the presence of the voids, nor was any significant void growth observed. The RDF, shown in Figure 5.18, also showed a similar trend to the 3-void case. In particular, the cascade created a dispersed set of vacancies, most noticeable in the range of about 10 to 80 nearest neighbor distances.





Figure 5.17: Plots of the defect counts averaged over the non-wrap-around runs for the 27-void cubic array configuration. Note the qualitative similarity to those discussed in Section 3.2.2.2 and 4.1 (i.e. Figure 4.6). The net number of defects is defined as the total number minus the initial number, as discussed in the text.

This fits with the expectation that the voids would allow for the formation of more dispersed cascade configuration. However, based on the visualized configurations, it was clear that not all of the voids played a role in this phenomenon. However, if multiple PKA excitation events occurred in the cell, it is possible that the other voids would become involved, but the exploration of that possibility is left for future work.


5.4 Conclusions

This chapter examined the role of initial void structures (i.e. large vacancy clusters) on the damage that resulted from a displacement cascade. Specifically, this chapter made use of four different geometries, all based on the perfect crystal simulation geometry, initial and boundary conditions of Section 4.1. The void structure configurations consisted of a single void placed at 2 different positions in the crystal, a 3-void linear array and a 27-void cubic array. Each of these configurations had the same nominal void volume fraction, thus the void size differed. The particular aspects of interest were the role of the voids on vacancy cluster formation and the extents of the cascade.


It was found that there was a link between the placement of a void and the PKA penetration, attributed to the tendency for a void to ‘offset’ part of the cascade. However, in the 3 and 27 void cases, the run-average penetration depth was closer to the central void or perfect crystal case than the near-surface void. This was attributed to the exclusion of those runs that exhibited wrap-around of the PKA, often due to the PKA exiting the simu-



(a) Initial RDF



(b) end of phase I



(c) end of phase III

Figure 5.18: Plots of the ‘net’ radial distribution function for vacancies in the 27 void cubic array configurations at the end of two key phases of cascade formation. The function plotted is the RDF from the timestep (averaged over all non-wraparound runs) minus the initial RDF (prior to PKA excitation). The initial RDF is given for reference.

lution cell at the ‘bottom’. In the 3-void case, this occurred in the majority of the 10 runs performed. Thus it is likely that the cell size impacted the solution. However, as the goal here was to compare to the results of Section 4.1, this chapter used the same simulation cell size. It is likely that a study that makes use of a larger cell would obtain results that show the expected increase in PKA penetration. On the other hand, the drawback would be that as the cell size increases, so too would the void size to maintain the void volume fraction. Thus, one would likely need to choose a smaller void volume fraction than that used here to reduce PKA wrap-around.

In an examination of the defect production, it was found that the net number of vacancies and interstitials introduced into the system by the cascade did not differ greatly (less than 10%) between the cases considered here and those of Section 4.1. This reflects the ideas presented in the literature that the amount of cascade damage depends in large part on the PKA kinetic energy and the TDE of the atoms in the crystal[20]. An examination of defect configurations and the RDF of the vacancies in system indicated a tendency for the simulations with initial voids to have a more dispersed damage configuration compared to the perfect crystal case. This was attributed to the voids allowing an increased probability that high-energy atoms could effectively ‘channel’ through them and collide with atoms on the other side. However, it was found in the 27 void case that the defect structure may not reach or interact with all of the voids, and may in some cases result in a relatively dense defect state due to the direction of travel of the high-energy atoms. Additional characterization

would be needed to make these results more conclusive.

From an examination of vacancy cluster formation, several things of note appeared. In the single void cases, where the size of the initial void and other clusters was tracked closely, it was clear that there was no stable void growth on the timescale of the simulations. However, a 3-fold increase in the number of clusters (apart from the void) was observed beyond that seen in the perfect crystal case. This indicated that the single void cases exhibited the possibility of enhanced clustering, an examination of the percentage of vacancies that were part of small clusters revealed that there was not a significant change from the perfect crystal case. Thus, it was concluded that the single voids did not greatly affect overall clustering, though they did result in smaller clusters than in the perfect crystal case. The cases with multiple vacancies did not exhibit clear changes in size, nor in the amount of clustering compared to the perfect crystal case. Thus, it was concluded that there was no evidence for the role of the voids as vacancy sinks on the timescale studied.

Chapter 6

Long-Time Behavior

The previous chapters focused on the picosecond scale behavior of 10 keV Si primary knock-on atom (PKA) displacement cascades in 3C-SiC, formed as a result of a simulated neutron bombardment. These studies provided important insights into the distribution of vacancies throughout the cascade formation and evolution. From the studies performed on a perfect crystal at 1200 K, it was demonstrated that one can extract quantitative information on vacancy clustering as well as an approximate spatial distribution of defects. However, on the timescale studied (~ 11 picoseconds), no large cluster formation or significant void growth was observed, even in the cases where an initial void structure was present. Further, it was concluded that an EP simulation that could capture a significant level of vacancy migration would likely be intractable due to the runtime and amount of data produced. As a result of these conclusions, it was decided that the application of a long-time simulation method

would extend the reach of the results presented here and provide a more complete answer to the question from the introduction, ‘*How does the initial distribution of defects evolve in time?*’.

This chapter will begin with a brief overview of two particular long-time simulation methods that have been used for atomic-level defect migration studies. The first to be discussed will be the widely used probabilistic method known as the Kinetic Monte Carlo Method (KMC) and the second will be the one applied in this chapter, Parallel Replica Dynamics (PRD). With the overview of these methods out of the way, the discussion will focus on PRD, its implementation, a relevant proof-of-concept application and close with a discussion of its advantages and disadvantages.

6.1 Long-Time Simulation Techniques

6.1.1 Kinetic Monte Carlo (KMC)

Kinetic Monte Carlo (KMC), non-equilibrium Monte Carlo or dynamic Monte Carlo methods [124–126] are a subgroup of Monte Carlo methods. Monte Carlo methods are mathematical simulation methods that rely on the use of random numbers in the solution process [127]. Monte Carlo methods have found widespread use in finance, computer science, physics and mathematics while KMC methods have widespread applications in solid-state physics and chemistry, particularly in the simulation of diffusion and reaction kinetics. It is of note that

papers prior to about the 1990s do not use the term kinetic Monte Carlo to refer to the method to be discussed[126] and instead use Monte Carlo. One early application [128] was the study of vacancy migration in binary alloys (SiC is an example of a binary alloy).

Rather than explicitly follow the trajectories of each particle in the system as done in methods such as EP, KMC takes a statistical approach by modeling the evolution of the ensemble of particles as a random walk between states via pre-identified transitions. An important aspect of KMC is that the timestep itself is a variable and is based on the transitions which occur. This idea allows for long time simulations of infrequent events as one need not select the timestep manually or proceed with a small (i.e. femtosecond scale) timestep throughout the course of the simulation in order to capture every transition event.

In a typical KMC algorithm, one must define a set of transition rates for each possible transition. These rates are usually obtained through some detailed deterministic calculations or experiments. These rates, as the name implies, are used to determine the probability that a certain transition will occur. These probabilities are then compared with a random number obtained from a random number generator. Based on this, a transition is chosen that is consistent with the desired transition probability distribution. After the transition, the simulation time is incremented based on the rate of the transition. The transition rates are then updated and the process repeated. A basic algorithm for this method is as follows, sometimes called the Bortz-Kalos-Liebowitz (BKL) algorithm[129]:

1. set time to zero ($t = 0$)

2. Form a list of all N possible transition rates, r_i
3. Calculate the cumulative function $R_i = \sum_{j=1}^i r_j$ for $i = 1, 2, \dots, N$
4. Get a random number, u' from a uniform random generator with limits 0 and 1
5. Find if a transition occurs by searching for an i such that $R_{i-1} < u'R_N \leq R_i$. If needed, allow transition i to occur and update rates
6. Get a random number, u from a uniform random generator with limits 0 and 1
7. Update time such that $t \leftarrow t - \frac{\log u}{R_N}$
8. Go to (3)

For the phenomena of interest here (i.e., point defect migration), the transitions could include a vacancy hopping from one lattice site to the next or a vacancy and interstitial recombination event. There are several key differences between an MD method and KMC calculations. The first is that in the case of point defect migration, one usually does not model the atoms but rather the defects themselves. Thus one must have a defined lattice upon which these defects can move, which may or may not be the same as the crystal structure of interest. For example, vacancies will move along the usual crystal structure but interstitials by definition do not. The second difference is that needs to be able to tabulate the possible transitions in the system, *prior* to running the simulation. Once the transitions are known, one would also need to characterize the rates at which these phenomena occur.

While the advantage of KMC methods is that one can reach very long times while retaining statistically meaningful physics, one needs to have a great deal of *a priori* knowledge about the physical system. In particular, generating the list of possible transition events is nontrivial. Similarly, the determination of the rates of transition and how they vary is a complex task. Finally, even though a KMC solution and a deterministic solution (of something such as Newton's equations of motion) lead to the same thermal equilibrium distribution, the KMC dynamics often have little resemblance to the deterministic dynamics [125]. For these reasons, it is desirable to have a method that combines the relative ease and intuitive interpretation of an EP calculation with the long timescale reachable by KMC. The next section will discuss one such technique.

6.1.2 Parallel Replica Dynamics (PRD)

The final simulation technique that will be discussed in this dissertation is an accelerated molecular dynamics method developed by Art Voter and co-workers [130, 131] referred to as Parallel Replica Dynamics (PRD). This technique has been successfully applied to the growth of interstitial clusters in Si [132], the transport of lithium in amorphous polyethylene oxide [133] and the stick-slip behavior of grain boundaries in copper [134]. The fundamental idea is simple: for events which are sufficiently rare that there is no strong dependence on previous events, one can utilize N statistically independent simulations (*replicas*) of the same system running in parallel to explore the solution space more rapidly. This can be done by

searching all of the replicas for the next event. In fact, most diffusive processes are such events [130].

Consider the thermally driven motion of a system of N atoms in a canonical ensemble (NVT), with an associated potential energy ‘landscape’ with a number of local minima and a global minimum. The classical dynamic exploration of this potential energy landscape can be done with methods such as EP, and for long enough times ends in the global minimum. If the trajectory during this long exploration is such that the behavior at any given time is statistically similar to what one would obtain from an ensemble of similar systems, each with slightly differing initial conditions, the ensemble of systems is called *ergodic*. Thus for ergodic systems, the average behavior through time of the single system is the same as the average behavior of the *ensemble of systems* at a single point in time[135]. This property is the basis for the fundamental idea of the acceleration of an EP simulation using PRD. From such systems one finds that there is an exponential decay in the probability of a transition from one state to the next as a function of time[131]:

$$P(t) = A \exp(-kt) \tag{6.1}$$

Thus the longer a system stays in a particular state, the less likely a transition will occur. As a result, the length of time between transitions can become large compared to the typical EP timescale.

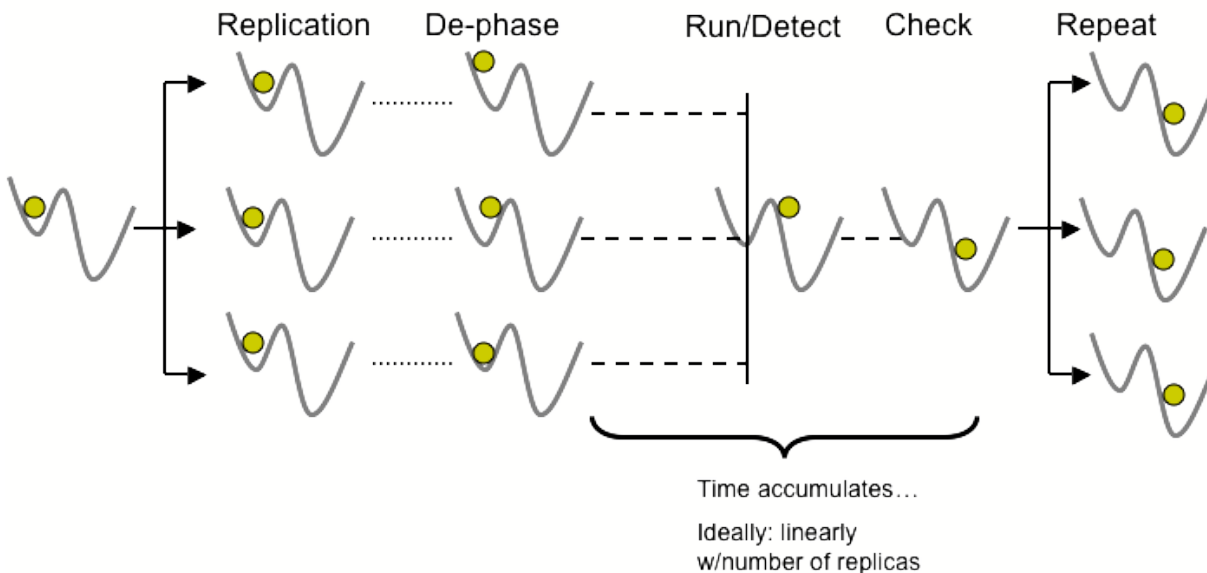


Figure 6.1: Schematic of the Parallel Replica Dynamics method. Note that the ideal scalability in the time-domain is linear with the number of replicas.

The PRD algorithm is fairly straightforward, and is shown schematically in Figure 6.1. Assuming that one has an initial state attached to a thermostat that requires a random number generator to assign velocities and a pre-defined transition criteria, the algorithm is as follows (based on [130, 131]):

1. set $t_{sys} = 0$
2. Generate N replicas of current state (i) across N sets of processors. Denote each individual replica by j .
3. De-phase each of the replicas by assigning each a different thermostat random number generator seed and running for a time $\Delta t_{dephase} \geq t_{corr}$. This runtime does not

contribute to the ‘simulated time’.

4. Continue to run each replica as a classical EP simulation and begin transition detection phase. Perform transition detection every $\Delta t_{check} \ll t_{corr}$. This runtime (Δt_j) contributes to overall ‘simulated time’ (i.e. $t_{sys} \leftarrow t_{sys} + \sum_{j=1}^N \Delta t_j$).
5. If a transition is detected in replica j , stop all other replicas. Continue to run replica j for additional $\Delta t_{cc} \geq t_{corr}$ checking for correlated transitions to previous state or new state. If stable transition detected, update state i to current state, $t_{sys} \leftarrow t_{sys} + \Delta t_{cc}$ and go to (2). Otherwise, go to (4) or exit.

Here, t_{corr} is the ‘correlation time’, the minimum time for which a subsequent transition would be independent of the current state. For typical bulk diffusion t_{corr} is on the order of 1 ps. During the transition detection phase, one can check at every integration step to ensure that nothing is missed. As long as the time between checks (Δt_{check}) is much smaller than the correlation time, one should not lose much accuracy [130] and will cut down on the expense of the calculation. In [130], the transition check requires an energy minimization at each check step and compares it to the results of an energy minimization of state i from (2). There is also the possibility of performing a more detailed search after a transition is detected to obtain a more accurate account of when the transition occurred.

In a statistical physics sense, the correlated transition check phase allows for short-time non-ergodicity of the system by allowing these events to occur within a single replica. After

this phase, when the algorithm is in steps (2)-(4), the system is presumed to be ergodic. This is why it is possible to say that the total simulated time is the sum of the times simulated in each of the replicas, plus the time required to run the non-ergodic events on a single processor. Ideally, the timescale of this simulation will scale linearly with the number of replicas, because of terms such as $\sum_{j=1}^N \Delta t_j$ in the time update. However, when the correlated transition check occurs, only a single processor advances the simulation time. Thus, the algorithm will be most efficient when the amount of time between transitions ($\Delta t_{transition}$) is much greater than the total time spent dephasing the replicas and searching for correlated events. Mathematically this can be expressed as

$$\Delta t_{transition} \gg N (\Delta t_{dephase} + \Delta t_{cc}) \quad (6.2)$$

Or, assuming one takes the minimum value for the dephasing and correlated search times,

$$\Delta t_{transition} \gg 2N t_{corr} \quad (6.3)$$

Further, if one takes t_{corr} to be on the order of 1 picosecond then one should expect good efficiency for transition timescales greater than $2N$ picoseconds.

The major advantage of PRD over other accelerated dynamics methods is that requires relatively little change from a typical EP simulation and no corrections to the system. Effectively, one could write a ‘wrapper’ script that would replicate the system and check each

replica for transitions, stop the non-transition replicas and run the correlated event checks. No major modifications to the EP solution procedure is necessary. However, implementing the PRD control-mechanism within the code does, in practice, make things easier. The implementation of PRD in LAMMPS will be discussed next.

6.2 Parallel Replica Dynamics Implementation

In order to further evaluate the applicability of PRD for use in cascade evolution studies, it was necessary to implement it. As this work had already made extensive use of LAMMPS, it was deemed most efficient to use it as the base for the implementation. The general algorithm was the same as that described in Section 6.1.2, though some additional assumptions were made. The code for the built-in LAMMPS ‘parallel tempering’ function formed the base of the implementation though only in as much as it helped set up the use of partitions and multiple, independent simulations. The implementation was created as a user package for the 21Apr08 release, current at the time of the work. The code should be easily updatable to later versions.

The first assumption is that the code would only be running on a parallel machine (i.e. more than 1 processor), such that the user could execute LAMMPS with the ‘-partition’ option. This option would allow LAMMPS to automatically set up to run multiple simulation instances from a single input file. The second is that one begins with an initial state in the form of a LAMMPS binary restart file, with all parameters, etc known. The third is that

one has some idea of the correlation time and an idea of the maximum number of steps required for a single transition. This information was used to ensure that the simulation did not enter an infinite loop and instead would exit, to avoid using large amounts of CPU time. The final assumption was that the system could be characterized by a maximum-displacement-based transition criteria. For the simulation type of interest here, this was a reasonable starting place (i.e. a vacancy migration would be the equivalent of an atomic migration in the opposite direction).

The implementation consisted of 2 parts, the input file and the source code. Thanks to the broad functionality built into LAMMPS, much of the low-level aspects (steps (2) and (3) in the algorithm) could be easily handled with standard LAMMPS input file syntax. The part that required a new function was the evolution and transition detection phase (steps (4) & (5) in the algorithm). By actually programming a new function, it was possible to get much more control than would be possible with the input file abilities. An example LAMMPS PRD input file is given in Appendix A.1.3. The new function, called ‘run_PRD’, was set up such that it would be started by each partition of processors, and a ‘master’ process would collect and disseminate comparison data at each detection step. The PRD master process was selected to be the same as the MPI master process for convenience. As of the writing of this dissertation, the command had not been submitted to the LAMMPS developers for inclusion in the code. The core code for the command has been given in Appendix A.1.4, with comments that describe in detail the algorithm as implemented.

6.3 Parallel Replica Dynamics Example

In order to demonstrate the basic functionality of the PRD implementation, a simple example simulation was performed. The example was related to cascades (in a simplified manner) and quite different from previous applications of PRD. The problem involved tracking the migration of a set of pre-existing vacancies in a small 3C-SiC simulation cell, similar to the cell used in the perfect crystal cascade simulations. The domain was 27,000 atoms (a 15x15x15 cell cube), with a lattice parameter of 4.3765 Å and a 2 cell thermostat skin on all sides but the top. From this geometry, 20 vacancies were created by deleting atoms according to a random distribution based on the Gaussian fit calculated for the perfect crystal cases. The input file for this simulation is given in Appendix A.1.3. For the PRD run, the thermostat temperature was kept at 2500 K via velocity scaling. This high temperature (about double the temperature used in the perfect crystal case) was chosen to enhance the migration of the vacancies in the system. Once the simulation was completed, a C++ version of the point defect identification code discussed in Section 2.3 was applied to obtain the positions of the vacancies.

The PRD example was run for 15 transition events (i.e. vacancy hops), using a 1.8 Å transition criteria with a 1 ps correlation time and a 1 fs timestep, checks for transitions were performed every 5 timesteps. The simulation ran in approximately 10 hours of wallclock time on 32 partitions of 1 processor each on the same ARL supercomputer used in the previous chapters. The estimated ‘simulated time’ was 742.52 picoseconds with a 1 fs timestep. The

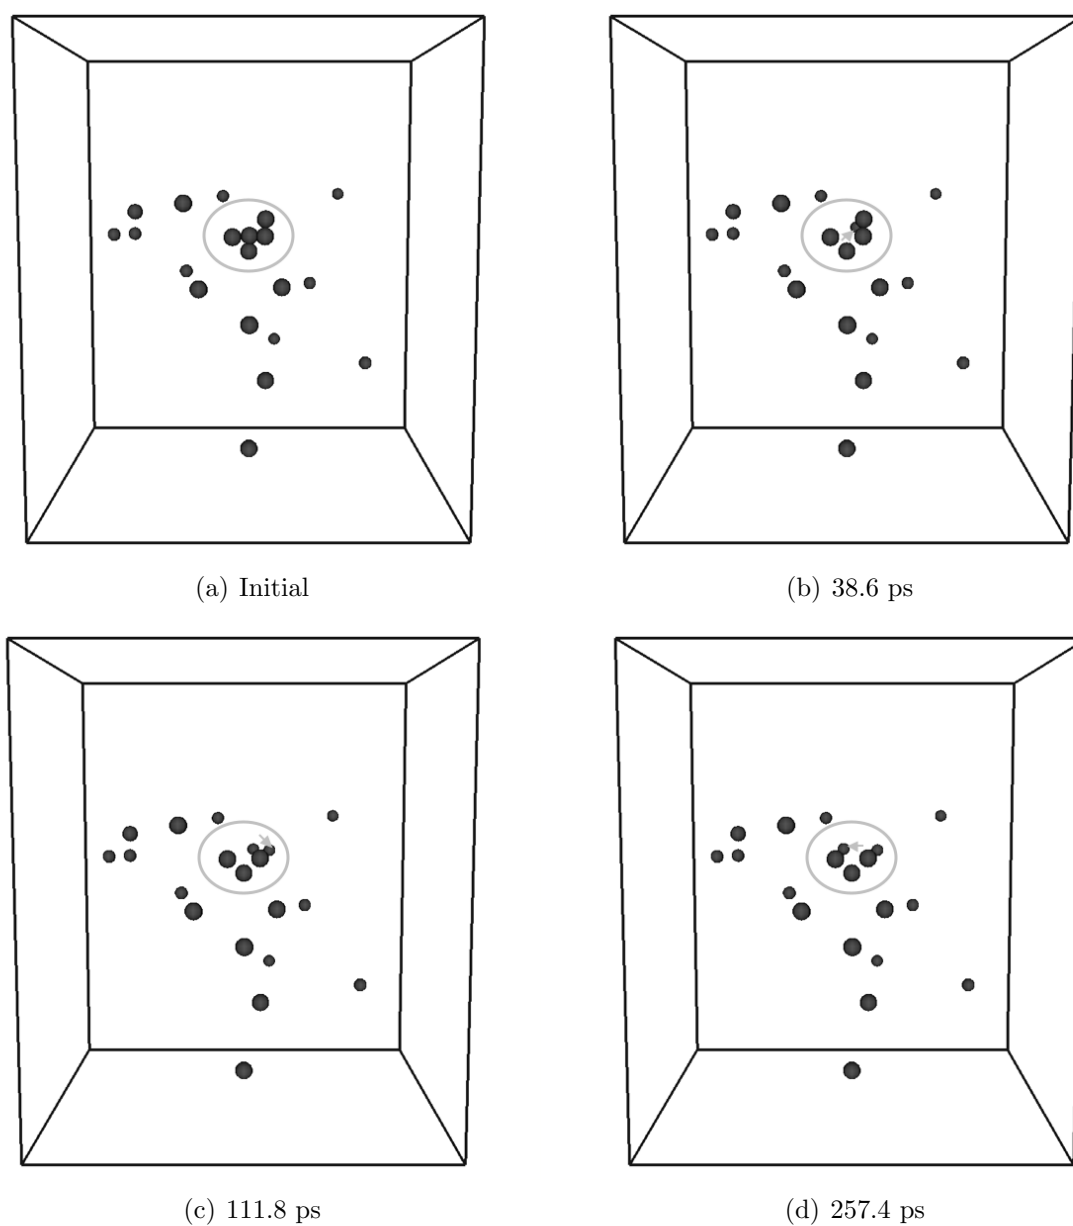


Figure 6.2: Visualization of initial configuration and the first 3 transitions in the PRD example calculation. Notice that the central initial small cluster is where the transitions primarily occur. The arrows are put in to act as visual guides showing the motion from the previous state to the current. The large spheres are Si vacancies and the smaller ones are C vacancies. Antisites and replacements have not been shown.

first 3 transitions are shown in Figure 6.2.

In the course of the simulations, stable transitions were found to occur between approximately 6 ps and 145 ps apart, with a mean of roughly 50 ps. Based on the standard EP single processor performance, the total of 742,520 timesteps simulated would have required on the order of 58 hours of CPU time to complete. Thus, PRD provided a wallclock speedup of roughly 5.8 times, nearly an order of magnitude lower than the ideal speedup. This is attributed to the relatively naive implementation and to the amount of time spent performing checks on possible transitions. The focus of the implementation provided here was to obtain a workable implementation, not the most efficient one. The method used to check for transitions and the communication of transitions between partitions is a likely culprit for inefficiency due to the amount of communication and waiting on each process. However, since the average time between transitions was about 50 ps, and each processor spent 1 ps de-phasing the replicas (for a total time across the replicas of 32 picoseconds), there is the possibility that the inefficiency is due to the timescale of the migration phenomena. It is possible that a decrease in the temperature to closer to 1000 K could lead to a more efficient computation, though it has been left as future work.

Despite these issues, some clear and physically interesting transitions were observed, though not enough data was generated to make any conclusive remarks. Figure 6.2 illustrates two of the major vacancy migration paths in 3C-SiC, a Si to C vacancy and C_{Si} antisite transition from the initial state to state 1 (Figure 6.2(b)) and a C second-neighbor vacancy




Figure 6.3: Visualization of final configuration in the PRD example calculation (742.52 ps). Notice that the central initial small cluster has grown from that seen in Figure 6.2. The large dark spheres are Si vacancies and the smaller dark ones are C vacancies. The large light sphere is an Si interstitial site. Antisites and replacements have not been shown.

hop from state 2 (Figure 6.2(c)) to state 3 (Figure 6.2(d)). These transitions compare well to the favorable transition paths observed in the *ab initio* study performed by Bockstedte and co-workers [136]. By the last few transitions, however, an interstitial has appeared (now there are 21 vacancies migrating) and the central cluster of vacancies has become more closely connected as seen in Figure 6.3. These results indicate that PRD is a viable method for looking at vacancy migration and clustering events in 3C-SiC. However, more work is needed before a meaningful vacancy cluster growth rate can be obtained in this manner. Such work and related issues are the topic of the next section.

6.4 Discussion of Parallel Replica Dynamics for Cascade Simulations

The proof-of-concept problem presented in the previous section demonstrated that the PRD technique could be applied to a problem that represented the basics of a cascade. While it neglected the presence of interstitials, the arrangement of vacancies was consistent with the cascade simulations performed in previous chapters. The PRD system did indeed show a non-negligible speedup (in wallclock time) compared to a single processor calculation, though the speedup was less than ideal. This non-ideal speedup was attributed mainly to lack of optimization in the implementation and to the timescale of the physical problem being studied. However, if one was to consider performing a PRD calculation on the final

configuration from the 1200 K perfect crystal (Section 4.1), there arises efficiency concerns from both spatial and temporal scales. This section will discuss these concerns and present some possible solutions.

Based on the proof-of-concept problem, the relatively short timescale of some transitions led to a reduction in efficiency. A problem that would likely only be exaggerated by the presence of a large number of interstitials as they tend to diffuse faster than vacancies. Further, the larger the number of defects, the more possible it becomes that one undergoes a transition at a relatively short time. Thus, one possible method to ensure that one is carrying out PRD calculations with maximum efficiency is to start with a small number processor partitions, run several transitions, and if the average time is long enough, increase the number of partitions and continue to run the simulation. This process can be repeated as needed and even automated, though it is possible that may become too cumbersome to maintain efficiency in the long run.

The second issue is due to the spatial scales of the types of simulations of interest in materials design. Truly, the cells used in Section 4.1 were found to be smaller than would be required to ensure that all of the cascades generated were contained within the simulation cell. The simulation domain used in Section 4.1 consisted of 1,000,000 atoms, and were run on 32 processors each. While PRD has no theoretical limitations on the size of the partitions, there is a practical hardware limitation: the number of CPUs available to the user. This is the major roadblock to PRD for large scale calculations and thus many materials design

applications.

For example, based on the simulations from Section 4.1, it was found that the simulation of the 10 ps final relaxation phase required roughly 4 hours on 32 processors of the Army Research Laboratorys MJM cluster. The nodes of the MJM cluster consist of 3.0 GHz Intel Woodcrest processors with Infiniband interconnects and operate at roughly 12.0 GFlops/processor (i.e. billion floating point operations per second per processor). Thus, the maximum time that could be simulated via EP would be approximately 240 ps in 96 hours (4 days) on 32 processors, the maximum allowable runtime for a standard user on MJM. Assuming linear scalability to 10,000 atoms per processor, a reasonable estimate based on experience and [115], this would result in roughly 960 ps in 96 hours on 128 processors. Based on estimates for runs of this size using PRD it would be possible to run a maximum of 8 128 processor partitions (1024 processors) for 96 hours, resulting in an ideal total simulation time of 7.7 ns, though to do this on MJM would have required special permissions. However, in order to truly say that we have approached a macroscopic timescale with such a simulation, say on the order of 1 microsecond, we would require a speedup factor of nearly 1000 over the single 32 processor simulation, which would translate to approximately 32,000 processors (i.e. 384 TFlops) working for 16 days, assuming ideal scalability. Further, in order to speak meaningfully of the statistics of such a system, a number of runs with different initial conditions would be needed, further inflating this amount and we would not have the luxury of simplified interatomic interactions to reduce expense due to the nature

of the physical system. To look at it another way, if one compares the scalability of an EP run with LAMMPS on 128 processors, one would expect to simulate 960 ps in 96 hours, meanwhile assuming perfect scalability in PRD, with the same number of processors divided into 4 partitions one would still only obtain 960 ps in 96 hours. Thus one would not expect to obtain a different timescale than EP unless a very large number of processors are available for an extended period of time, something unlikely on a machine with many users. Thus it is more likely that PRD would find use as a method to verify other simulations, such as KMC calculations, via smaller length scale problems. Further, PRD is a strong candidate to extend the timescale in simulations where a parallel EP code is not available or when the system size is too small to split efficiently among the processors. However, at best it scales linearly in the number of processors, identical to the scaling of an efficient parallel EP code. Thus, it can not be expected to ever reach the timescales reachable by methods like KMC.

There are other ‘hidden’ aspects of the PRD algorithm that can be applied to make EP simulations easier. One major aspect is the fact that PRD describes a technique that will only output data when there has been a clear state change. In EP calculations, the frequency at which output files are written is usually determined *a priori* in an *ad hoc* manner. Typically, this is set at some moderate number of timesteps (100 in the final evolution cases in Chapter 4) and iterated upon if more resolution is needed. Thus, one could write a function based on the transition criteria ideas used in PRD to control when data was output during an EP simulation.

6.5 Conclusions

This chapter began with a presentation of two long-time atomistic simulation methods, Kinetic Monte Carlo and Parallel Replica Dynamics. The implementation of PRD in LAMMPS was discussed. A simple proof-of-concept problem was presented and some aspects of its solution discussed, such as the C and Si vacancy diffusion pathways it demonstrated. The problem exhibited less than ideal scaling in time, which was attributed to the physics of the problem as well as the implementation. Finally, a discussion on the application of PRD to displacement cascade evolution was given. It was concluded that while PRD shows promise, it is only a viable method when a very large number of CPUs are available for an extended period of time. Despite this, PRD does have the benefit of providing a pseudo-deterministic method to verify KMC calculations and the transition criteria ideas are applicable to EP calculations to reduce the output requirements. Further, PRD can be used when a parallel EP code is not available or when the system size is too small to efficiently parallelize using typical means.

Chapter 7

Summary and Conclusions

This dissertation sought a quantitative description of vacancy cluster formation from distributed vacancies in neutron-irradiated 3C-SiC at temperatures and pressures relevant to nuclear reactors. This question was approached in several parts via empirical potential atomistic simulations. The first chapter of this dissertation provided a review of the literature regarding nuclear power applications, predicted operating conditions and materials selection issues. The relevant conditions were obtained based on the ARIES-AT fusion reactor design[30]. This design calls for operational temperatures near 1000 °C and a pressure range of 50-175 MPa [30]. A PKA kinetic energy of 10 keV was chosen for its relevance to both fusion and fission applications as well as being within the energy range where classical molecular dynamics simulations of atomic collisions are valid. Following [68], the PKA was chosen to be a Si atom. No simulations were performed with a C PKA, because previous

Units are eV	Empirical Potential	<i>Ab initio</i>
C interstitial migration barrier	.75 [120]	n/a
Si interstitial migration barrier	1.5 [120]	n/a
C vacancy migration barrier	4.1 [120]	4.8 [137]
Si vacancy migration barrier	2.3 [120]	1.7 [137]
C vacancy formation energy	5.15 [121]	5.48 [121]
Si vacancy formation energy	6.31 [121]	6.64 [121]
C _{Si} formation energy	0.35 [121]	1.32 [121]
Si _C formation energy	6.10 [121]	7.20 [121]
Average C TDE in 3C-SiC	21 (this work)	15 [109]
Average Si TDE in 3C-SiC	42 (this work)	41 [109]
Average Fe (BCC) TDE	36 [138]	n/a
Average Ni (FCC) TDE	38 [138]	n/a
Average Cu (FCC) TDE	44 [138]	n/a

Table 7.1: Characteristic point defect related energies for 3C-SiC, from several sources. Values of the average TDE for BCC Fe, FCC Ni and FCC Cu from [138] have been added for reference. Empirical potential values for SiC use either the potential used in this work or one that is similar.

simulations have shown cascade formation to be largely independent of the PKA type[114].

With the bulk of the relevant physics described in the first chapter, the second chapter focused on the theoretical and practical background of the methods employed to answer the question posed. The third chapter reviewed the relevant simulation literature and illustrated the 2-part verification procedure for the simulation approach. The first part of this procedure was a TDE calculation for the hybrid Tersoff/ZBL potential used in this dissertation, based off the work in [85] and [118]. In general, the TDEs calculated with the potential used here were within the bounds reported in the simulation literature despite the use of a slightly different small-separation potential function. The TDE variations found in this work were attributed to subtle differences in the definition of a ‘displacement event’ (compared to the

literature) and statistical fluctuations due to the initial conditions of the system. It was concluded that the potential performed sufficiently well to justify its use, but it was also concluded that averaging over several runs with varied initial conditions would likely have yielded results that match the literature. The second part of the verification procedure was the reproduction of the 10 keV Si PKA cascade simulation reported in [68]. The results showed good qualitative agreement with the published results, verifying the simulation and post-processing procedure. The $\sim 50\%$ difference in the number of stable vacancies was attributed to differences in identifying point defects and solution scheme not reported in [68]. Compared to other materials of interest for nuclear applications such as Fe (BCC structure), Ni (FCC structure) and Cu (FCC structure) the average Si TDE for SiC is higher while the average C TDE is lower, based on the values in [138]. The average TDEs for Fe (36 eV) and Ni (38 eV) fall almost half-way between the average Si and C TDE values in 3C-SiC, while Cu (44 eV) is much closer to the Si TDE in 3C-SiC. The average 3C-SiC TDEs and several additional material properties related to point defect production in 3C-SiC are given in Table 7.1.

The fourth chapter discussed the formation of a 10 keV silicon PKA displacement cascade in a perfect silicon carbide crystal and thermal effects on the damage state. The first part of the chapter focused on a quantitative description of the damage state after cascade formation. The key contribution of this was extracting vacancy spatial distributions and vacancy clustering trends at key steps in the cascade evolution. The results of this section

were the first unique contribution of this dissertation. From the final configurations, it was found that there was a gaussian distribution of vacancies in the [100], [010] directions and a nearly gaussian distribution in the [001] direction. This information was used to generate the vacancy distribution studied in the long-time calculations of the sixth chapter. It was concluded that further simulations and statistical analysis is required to improve the quality of the spatial distributions, though those calculated here provided a good initial approximation. Additionally, if the same procedure was applied to other point defects, it would be possible to obtain an approximate, partially relaxed, cascade without explicitly simulating its formation.

The second part of Chapter 4 investigated the role of thermal vibrations (i.e. temperature) on cascade formation and relaxation in a perfect crystal. The number of C_{Si} antisites was significantly higher at higher temperatures, likely due to enhanced atomic mobility. The most dramatic result of this section was the observation of enhanced cascade relaxation at 2000 K, which resulted in fewer point defects and more distributed vacancies than at lower temperatures. This led to the conclusion that there is a high-temperature cascade formation regime where atomic and defect mobility effectively decreases the amount of concentrated damage in the system. Further, it was found that the number of vacancies in clusters was relatively small (10-30 % of the total number of vacancies) regardless of temperature. In comparison to 20 keV cascades in BCC Fe [139], we observed the formation of nearly twice as many vacancy-interstitial pairs in the final state despite the lower PKA kinetic energy in

Fe. However the percentage of vacancies associated with clusters was smaller in 3C-SiC, by as much as 15% (at 600 K in Fe, compared to 500 K here) though the values at 100 K were similar for both materials. Compared to the FCC Cu and BCC Fe cascades in [138], we found nearly three times more vacancy-interstitial pairs in 3C-SiC. For 30 keV cascades at low temperatures (nominally less than 200 K) in Ni[140], some interatomic potentials predict a number of vacancies and interstitials similar to that seen in the 10 keV cascades here. In Ni, the formation of extended vacancy cluster structures (stacking fault tetrahedra, for example) was reported, unlike anything seen in the simulations here. Thus, there is a clear qualitative difference in the behavior of 3C-SiC compared to Fe, Cu and Ni. In particular, while the nominal system TDE in 3C-SiC (average of the Si and C TDEs) is similar to that of Fe and Ni, many more vacancies are generated in 3C-SiC at a lower PKA energy than in these two metals. This is most likely due to the relatively low TDE of C, which forms the vast majority of vacancies in 3C-SiC.

The fifth chapter extended the work of the fourth by considering the effect of an initial void structure. The void structures consisted of a single void placed at 2 different positions in the crystal, a 3-void linear array and a 27-void cubic array. Each of these configurations had the same nominal void volume fraction of .2 %. Each void was a sphere, thus the void radius in each case differed. The particular aspects of interest were the role of the voids on vacancy cluster formation and the extents of the cascade. It was found that there was a link between the placement of a void and the PKA penetration, attributed to the tendency for

a void to ‘offset’ part of the cascade. However, in the 3 and 27 void cases, the run-average penetration depth was closer to the central void or perfect crystal case than the near-surface void. This was attributed to the exclusion of those runs that exhibited wrap-around of the PKA, often due to the PKA exiting the simulation cell at the ‘bottom’. In the 3-void case, this occurred in the majority of the runs performed. Thus it is likely that the cell size impacted the solution. An examination of defect configurations and the RDF of the vacancies in the system indicated a tendency for the simulations with initial voids to have a more dispersed damage configuration compared to the perfect crystal case. In the single void cases, where the size of the initial void and other vacancy clusters was tracked closely, it was clear that there was no stable void growth on the timescale of the simulations. However, a 3-fold increase in the number of clusters (apart from the void) was observed beyond that seen in the perfect crystal case. This indicated that the single void cases exhibited enhanced clustering, though an examination of the percentage of clustered vacancies revealed that there was not a significant change from the perfect crystal case. Thus, it was concluded that the single voids did not greatly affect overall clustering, though they did result in smaller clusters than in the perfect crystal case. Further, it was concluded that there was no evidence for the role of the voids as vacancy sinks on the timescale studied, though the boundary conditions may have increased the stability of the void by not allowing it to collapse or re-configure due to the requirement that the number of atoms in the simulation cell remain constant. Because of this, it is possible that the addition of atoms in the void (such as hydrogen or

helium gas) could yield more dramatic structural changes.

The sixth chapter focused on long-time methods for atomistic simulations and a proof-of-concept parallel replica dynamics problem. This was the first application, to the author's knowledge, of PRD to a cubic silicon carbide system with a large number of migrating vacancies. The chapter closed with a presentation of the advantages and disadvantages of the use of parallel replica dynamics for cascade simulations. It was concluded that while PRD shows promise, it is only a viable method when a very large number of CPUs are available for an extended period of time. Despite this, PRD does have the benefit of providing a pseudo-deterministic method to verify KMC calculations and the transition criteria ideas are applicable to EP calculations to reduce data output. Most importantly, PRD can be used when a parallel EP code is not available or when the system size is too small to efficiently parallelize using typical means.

In closing, this work made another step toward a quantitative description of post-irradiation damage in 3C-SiC. The main contributions of this work were:

- Quantitative characterization of the spatial distribution of vacancies after simulated cascade formation in 3C-SiC.
- Evidence of enhanced defect migration leading to reduction in cascade-induced damage at 2000 K.
- First detailed computational study on the influence of temperature on defect distribu-

tion in 3C-SiC after simulated irradiation.

- First investigation of initial void structure influence on point defect generation and distribution in 3C-SiC after simulated irradiation.
- Discussion and demonstration of the applicability of parallel replica dynamics to post-cascade relaxation phase simulations in SiC.

Based on the experience gained during this work the following are recommended as avenues of future research that directly build on the work in this dissertation:

- Stability of point defects (antisites in particular) and point defect clusters over timescales much greater than 10 ps.
- High-temperature (2000 K and higher) cascade generation, relaxation mechanisms (role of antisites in particular) and timescales.
- Influence of void structure on defect structure in larger cells, longer times.
- Influence of hydrogen or helium gas in voids on void structure changes and vacancy clustering.
- Influence of multiple irradiation events and PKA directions on damage structure.

Bibliography

- [1] *Higher Temperature Reactors Materials Workshop*, 2002. Proceedings of DOE workshop, La Jolla CA, 18-21 March 2002.
- [2] *Basic Research Needs for Materials Under Extreme Environments*, 2007. Proceedings of the DOE Basic Energy Sciences Workshop, Bethesda, MD, 11-13 June 2007.
- [3] *Workshop on Advanced Computational Materials Science: Application to Fusion and Generation IV Fission Reactors*, 2004. Proceedings of DOE workshop, Washington DC, 31 March – 2 April 2004.
- [4] Generation IV Initiative website. <http://gen-iv.ne.doe.gov/>.
- [5] Generation IV International Forum website. <http://gif.inel.gov/>.
- [6] A technology roadmap for generation IV nuclear energy systems., 2002. U.S. DOE Nuclear Energy Research Advisory Committee and Generation IV International Forum Publication. GIF-002-00.
- [7] DOE Yucca Mountain Nuclear Waste Repository project overview website. http://www.ocrwm.doe.gov/ym_repository/index.shtml.
- [8] J. Boak. Cost slows nuclear plant drive. *Chicago Tribune*, 7 May 2008. Business Section pages 1 & 4.
- [9] B. Raj et al. Challenges in materials science research for sustainable nuclear energy. *MRS Bulletin*, 33:327–337, 2008. Special Issue on Harnessing Materials for Energy, April.
- [10] A.M. Stoneham, C.A. English, and W.J. Pythian. Challenges in radiation damage: From dynamics of atoms to nuclear plant life management. *Radiation Effects & Defects in Solids*, 144, 1998.

- [11] Yucca mountain, DOE's planned nuclear waste repository faces quality assurance and management challenges. statement of Jim Wells, director, natural resources and environment, 2006. US GAO report, GAO-06-550T.
- [12] W.K. Liu, E.G. Karpov, S. Zhang, and H.S. Park. An introduction to computational nanomechanics and materials. *Computer Methods in Applied Mechanics and Engineering*, 193(17-20):1529–1578, 2004.
- [13] B.D. Wirth et al. Multiscale modeling of radiation damage in Fe-based alloys in the fusion environment. *Journal of Nuclear Materials*, 329–333:103–111, 2004.
- [14] S. Zinkle. Fusion materials science: Overview of challenges and recent progress. *Physics of Plasmas*, 058101, 2005.
- [15] W.D Callister. *Materials Science and Engineering, an Introduction*. John Wiley & Sons, 5th edition, 2000.
- [16] *ASTM E170-08b, Standard Terminology Relating to Radiation Measurements and Dosimetry*. ASTM, 2008.
- [17] *ASTM E693-01, Standard Practice for Characterizing Neutron Exposures in Iron and Low Alloy Steels in Terms of Displacements Per Atom (DPA)*. ASTM, 2001. Reapproved 2007.
- [18] R. Kienberger and F. Krausz. Attosecond metrology comes of age. *Physica Scripta*, T110:32–38, 2004.
- [19] E. goulielmakis et al. Attosecond control and measurement: Lightwave electronics. *Science*, 10:769–775, 2007.
- [20] H.L. Heinisch, L.R. Greenwood, W.J. Weber, and R.E. Williford. Displacement damage cross sections for neutron-irradiated silicon carbide. *Journal of Nuclear Materials*, 307–311:895–899, 2002.
- [21] E.R. Benton and E.V. Benton. Space radiation dosimetry in low-earth orbit and beyond. *Nuclear Instruments and Methods in Physics Research B*, 184:255–294, 2001.
- [22] E.R. Benton and E.V. Benton. The radiation environment in low-earth orbit. *Radiation Research*, 148:S3–S10, 1997.
- [23] R. Kelly and A. Miotello. Comments on explosive mechanisms of laser sputtering. *Applied Surface Science*, 96–8:205–215, April 1996.

- [24] H. Trinkaus and B.N. Singh. Helium accumulation in metals during irradiation - where do we stand? *Journal of Nuclear Materials*, 323(2-3):229–242, 2003.
- [25] M.W. Thompson. *Defects and Radiation Damage in Metals*. Cambridge Monographs on Physics. Cambridge University Press, 1st edition, 1969.
- [26] P. Sigmund. Collision theory of displacement damage, ion ranges and sputtering. *Revue Roumaine de Physique*, 7, 1972.
- [27] S.J. Zinkle and C. Kinoshita. Defect production in ceramics. *Journal of Nuclear Materials*, 251, 1997.
- [28] E. Friedland. Radiation damage in metals. *Critical Reviews in Solid State and Materials Sciences*, 26(2), 2001.
- [29] R.S. Averback and T. Diaz de la Rubia. Displacement damage in irradiated metals and semiconductors. In H. Enrenreich and F. Spaepen, editors, *Solid State Physics, Advances in Research and Applications*, volume 51. Academic Press, 1997.
- [30] F. Najmabadi et. al. The ARIES-AT advanced tokamak, advanced technology fusion power plant. *Fusion Engineering and Design*, 80(1), 2006.
- [31] International Thermonuclear Experimental Reactor (ITER) project website. <http://www.iter.org/>.
- [32] Joint European Torus website. <http://www.jet.efda.org/index.html>.
- [33] Experimental Advanced Superconducting Tokamak website. <http://www.ipp.ac.cn/ENGLISH/research/EASTmain.htm>.
- [34] *Basic Research Needs for Advanced Nuclear Energy Systems*, 2006. Proceedings of the DOE Basic Energy Sciences Workshop, Bethesda, MD, 31 July - 3 August 2006.
- [35] Modular pebble bed reactor project, university research consortium annual report, 2000. MIT, INEEL collaborative report. INEEL/EXT-2000-01034, MIT-ANP-PR-075.
- [36] D.R. Nicholls. Status of the pebble bed modular reactor. *Nuclear Energy*, 39, 2000.
- [37] The South African Pebble Bed Modular Reactor company, PBMR (Pty) Ltd. website. <https://www.pbmr.com/>.
- [38] S.J. Zinkle, M. Victoria, and K. Abe. Scientific and engineering advances from fusion materials R&D. *Journal of Nuclear Materials*, 31–42:307–311, 2002.

- [39] A.S.P. Ramirez, A. Caso, L. Giancarli, N. LeBars, G. Chaumat, J.F. Salavy, and J. Szczepanski. TAURO: A ceramic composite structural material self-cooled Pb-17Li breeder blanket concept. *Journal of Nuclear Materials*, 233, 1996.
- [40] *Scientific Issues & Technological Challenges in the Development of Materials*, 2004. Proceedings of the 2004 Fredric Joliot & Otto Hahn Summer School, Cadarache, France, 25 August - 3 September 2004.
- [41] A. Hasegawa, A. Kohyama, R.H. Jones, L.L. Snead, B. Riccardi, and P. Fenici. Critical issues and current status of SiC/SiC composites for fusion. *Journal of Nuclear Materials*, 283, 2000.
- [42] B. Riccardi, L. Giancarli, A. Hasegawa, A. Katoh, A. Kohyama, R.H. Jones, and L.L. Snead. Issues and advances in SiCf/SiC composites development for fusion reactors. *Journal of Nuclear Materials*, 329–333, 2004.
- [43] R. Blackstone and E. H. Voice. The expansion of silicon carbide by neutron irradiation. *Journal of Nuclear Materials*, 39, 1971.
- [44] J.E. Palentine. The development of silicon carbide as a routine irradiation temperature monitor, and its calibration in a thermal reactor. *Journal of Nuclear Materials*, 61: 243–253, 1976.
- [45] A.I. Ryazanov, A.V. Klaptsov, A. Kohyama, and H. Kishimoto. Radiation swelling of SiC under neutron irradiation. *Journal of Nuclear Materials*, 307–311:1107–1111, 2002.
- [46] R.J. Price. Neutron irradiation induced voids in β -silicon carbide. *Journal of Nuclear Materials*, 48:47–57, 1973.
- [47] R. Scholz. Light ion irradiation creep of SiC fibers in torsion. *Journal of Nuclear Materials*, 258–263:1533–1539, 1998.
- [48] Y. Katoh, N. Hashimoto, S. Kondo, L.L. Snead, and A. Kohyama. Microstructural development in cubic silicon carbide during irradiation at elevated temperatures. *Journal of Nuclear Materials*, 351:228–240, 2006.
- [49] D. Pettifor. *Bonding and Structure of Molecules and Solids*. Oxford Science Publications. Oxford University Press, 2002.
- [50] A.R. Verma and P Krishna. *Polymorphism and Polytypism in Crystals*. John Wiley & Sons, 1st edition, 1966.

- [51] F. Bechstedt, P. Käckell, A. Zywietz, K. Karch, B. Adolph, K. Tenelsen, and J. Furthmüller. Polytypism and properties of silicon carbide. *Physica Status Solidi B*, 202(1), 1997.
- [52] P. Pirouz. Polytypic transformations in SiC. *Solid State Phenomena*, 56, 1997.
- [53] J.B. Casady and R.W. Johnson. Status of silicon carbide (SiC) as a wide-bandgap semiconductor for high temperature applications: A review. *Solid-State Electronics*, 39(10), 1996.
- [54] M. Mehregany, C.A. Zorman, N. Rajan, and C.H. Wu. Silicon carbide MEMS for harsh environments. *Proceedings of the IEEE*, 86(8), 1998.
- [55] J.L. Katz and H. Wiedersich. Nucleation of voids in materials supersaturated with vacancies and interstitials. *Journal of Chemical Physics*, 55(3):1414–1425, 1971.
- [56] L.K. Mansur. Mechanisms and kinetics of radiation effects in metals and alloys. In G.R. Freeman, editor, *Kinetics of Nonhomogeneous Processes*, chapter 8. Wiley-Interscience, 1987.
- [57] N.V. Doan and G. Martin. Elimination of irradiation point defects in crystalline solids: Sink strengths. *Physical Review B*, 67(134107), 2003.
- [58] B. O. Hall. Point defect clustering during irradiation. *Journal of Nuclear Materials*, 91:63–72, 1980.
- [59] R.S. Evans, D.L. Bourell, J.J. Beaman, and M.I. Campbell. Rapid manufacturing of silicon carbide composites. *Rapid Prototyping Journal*, 11(1):37–40, 2005.
- [60] N. Orlovskaya et. al. SiC/SiC_{wovenfabric} laminates: Design, manufacturing, mechanical properties. *Composites Part B*, 37:524–529, 2006.
- [61] CoorsTek Advanced Technical Ceramics website. <http://www.coorstek.com/>.
- [62] Dow Corning website. <http://www.dowcorning.com/>.
- [63] W.M. Vetter and M. Dudley. The character of micropipes in silicon carbide crystals. *Philosophical Magazine*, 86(9):1209–1225, 2006.
- [64] P.G. Neudeck and J.A. Powell. Performance limiting micropipe defects in silicon-carbide wafers. *IEEE Electron Device Letters*, 15(2):63–65, 1994.
- [65] M. Prasad and T. Sinno. Internally consistent approach for modeling solid-state aggregation. I Atomistic calculations of vacancy clustering in silicon. *Physical Review B*, 68(045206), 2003.

- [66] M. Prasad and T. Sinno. Internally consistent approach for modeling solid-state aggregation. II Mean-field representation of atomistic processes. *Physical Review B*, 68 (045207), 2003.
- [67] T. Frewen et al. A microscopically accurate continuum model for void formation during semiconductor silicon processing. *Journal of Crystal Growth*, 279:258–271, 2005.
- [68] R. Devanathan, W.J. Weber, and T. Diaz de la Rubia. Computer simulation of a 10-keV displacement cascade in SiC. *Nuclear Instruments and Methods in Physics Research B*, 141:118–122, 1998.
- [69] J.F. Ziegler, J.P. Biersack, and U. Littmark. *The Stopping and Ranges of Ions in Solids*, volume 1. Pergamon Press, 1985.
- [70] M.T. Robinson and I.M. Torrens. Computer simulation of atomic-displacement cascades in solids in the binary collision approximation. *Physical Review B*, 9(12):5008–5024, 1974.
- [71] M.T. Robinson. Basic physics of radiation damage production. *Journal of Nuclear Materials*, 216:1–28, 1994.
- [72] Stopping and Range of Ions in Matter code website. <http://www.srim.org/>.
- [73] M. Finnis. *Interatomic Forces in Condensed Matter*. Oxford University Press, 1st edition, 2003.
- [74] M.A. Ratner and G.C. Schatz. *Introduction to Quantum Mechanics in Chemistry*. Prentice Hall, 1st edition, 2001.
- [75] J. M. Haile. *Molecular Dynamics Simulations*. John Wiley & Sons, 1992.
- [76] M. P. Allen and D. J. Tildesley. *Computer Simulation of Liquids*. Oxford University Press, 1987.
- [77] W.K. Liu, E.G. Karpov, and H.S. Park. *Nano Mechanics and Materials : Theory, Multiple Scale Analysis, and Applications*. John Wiley & Sons, 1st edition, 2005.
- [78] P.M. Morse. Diatomic molecules according to the wave mechanics. II. Vibrational levels. *Physical Review*, 34, 1929.
- [79] J.E. Lennard-Jones. On the determination of molecular fields. I. From the variation of the viscosity of a gas with temperature. *Proceedings of the Royal Society of London*, 106:441–462, 1924.

- [80] F.H. Stillinger and T.A. Weber. Computer simulation of local order in condensed phases of silicon. *Physical Review B*, 31(4):5262–5271, 1985.
- [81] M.S. Daw, S.M. Foiles, and M.I. Baskes. The embedded-atom method – a review of theory and applications. *Materials Science Reports*, 9:251–310, 1993.
- [82] J. Tersoff. Modelling solid-state chemistry: Interatomic potentials for multicomponent systems. *Physical Review B*, 39, 1989.
- [83] M. Tang and S. Yip. Atomistic simulation of thermomechanical properties of β -SiC. *Physical Review B*, 52(21):15150–15159, 1995.
- [84] M. Tang. *Elastic Instabilities and Structural Responses of β -SiC Under Stress*. PhD thesis, Massachusetts Institute of Technology, 1995.
- [85] R. Devanathan, T. Diaz de la Rubia, and W.J. Weber. Displacement threshold energies in β -SiC. *Journal of Nuclear Materials*, 253:47–52, 1998.
- [86] R. Devanathan and W.J. Weber. Displacement energy surface in 3C and 6H SiC. *Journal of Nuclear Materials*, 278:258–265, 2000.
- [87] F. Gao, E.J. Bylaska, W.J. Weber, and L.R. Corrales. Native defect properties in β -SiC: Ab initio and empirical potential calculations. *Nuclear Instruments and Methods in Physics Research B*, 180:286–292, 2001.
- [88] C.L. Rountree et al. Atomistic aspects of crack propagation in brittle materials: Multimillion atom molecular dynamics simulations. *Annual Review of Materials Research*, 32:377–400, 2002.
- [89] M. Ippolito, A. Mattoni, L. Columbo, and F. Cleri. Fracture toughness of nanostructured silicon carbide. *Applied Physics letters*, 87, 2005.
- [90] A. Mattoni, L. Columbo, and F. Cleri. Atomic scale origin of crack resistance in brittle fracture. *Physical Review Letters*, 95, 2005.
- [91] WebElements online periodic table and atomic information. <http://www.webelements.com/>.
- [92] D.J. Evans and B. L. Holian. The Nose-Hoover thermostat. *Journal of Chemical Physics*, 83, 1985.
- [93] H.J.C. Berendsen et al. Molecular dynamics with coupling to an external heat bath. *Journal of Chemical Physics*, 81, 1984.

- [94] E.G. Karpov, H.S. Park, and W.K. Liu. A phonon heat bath approach for the atomistic and multiscale simulation of solids. *International Journal for Numerical Methods in Engineering*, 2005. preprint.
- [95] W. H. Press, S. A. Teukolsky, W.T. Vetterling, and B.P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing*. CRC press, 2nd edition, 1997.
- [96] gnuplot: open source graphing application website. <http://www.gnuplot.info/>.
- [97] I. Santos et al. Molecular dynamics characterization of as-implanted damage in silicon. *Materials Science and Engineering B*, 124–125:372–375, 2005.
- [98] J.B. Gibson, A. N. Goland, M. Milgram, and G.H. Vineyard. Dynamics of radiation damage. *Physical Review*, 120(4):1229–1253, 1960.
- [99] K. Nordlund et al. Defect production in collision cascades in elemental semiconductors and fcc metals. *Physical Review B*, 57(13):7556–7570, 1998.
- [100] S.M. Foiles. Detailed characterization of defect production in molecular dynamics simulations of cascades in Si. *Nuclear Instruments and Methods in Physics Research B*, 255:101–104, 2007.
- [101] J.E. Goodman and J. O’Rourke. *Handbook of discrete and computational geometry*. CRC Press series on discrete mathematics and its applications. Chapman & Hall, 2^{ed} edition, 2004.
- [102] ParaView: Parallel Visualization Application website. <http://www.paraview.org/>.
- [103] P. Hohenberg and W. Kohn. Inhomogeneous electron gas. *Physical Review*, 136(3B):864–871, 1964.
- [104] W. Kohn and L.J. Sham. Self-consistent equations including exchange and correlation effects. *Physical Review*, 140(4A):1133–1138, 1965.
- [105] D. J. Griffiths. *Introduction to Quantum Mechanics*. Prentice Hall, 1st edition, 1995.
- [106] M. J. Mehl and D. A. Papaconstantopoulos. Tight-binding parameterization of first-principles results. In C. Y. Fong, editor, *Topics in Computational Materials Science*, pages 169–213, Singapore, 1998. World Scientific.
- [107] D.J. Chadi. (100) surface atomic structures of covalent and ionic semiconductors. *Physical Review B*, 19(4):2074–2082, 1979.
- [108] A.A. Abrahamson. Born-mayer-type interatomic potential for neutral ground-state atoms with $z = 2$ to $z = 105$. *Physical Review*, 178(1):76–79, 1969.

- [109] G. Lucas and L. Pizzagalli. Ab initio molecular dynamics calculations of threshold displacement energies in silicon carbide. *Physical Review B*, 72, 2005.
- [110] W. Windl, T.J. Lenosky, J.D. Kress, and A.F. Voter. First-principles investigation of radiation induced defects in Si and SiC. *Nuclear Instruments and Methods in Physics Research B*, 141:61–65, 1998.
- [111] C.S. Becquart, C. Domain, A. Legris, and J.C. Van Duysen. Influence of the interatomic potentials on molecular dynamics simulations of displacement cascades. *Journal of Nuclear Materials*, 280:73–85, 2000.
- [112] A.A. Demkov, J. Ortega, O.F. Sankey, and M.P. Grumbach. Electronic structure approach for complex silicas. *Physical Review B*, 52(3):1618–1630, 1995.
- [113] O.F. Sankey et al. The application of approximate density functionals to complex systems. *International Journal of Quantum Mechanics*, 69(3):327–340, 1998.
- [114] J.M. Perlado, L. Malerba, A Sanchez-Rubio, and T. Diaz de la Rubia. Analysis of displacement cascades and threshold displacement energies in β SiC. *Journal of Nuclear Materials*, 276(1-3):235–242, 2000.
- [115] Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) website. <http://lammmps.sandia.gov/>.
- [116] H. Hensel and H.M. Urbassek. Preferential effects in low-energy si bombardment of SiC. *Nuclear Instruments and Methods in Physics Research B*, 142, 1998.
- [117] J. Wong et al. The threshold energy for defect production in SiC: A molecular dynamics study. *Journal of Nuclear Materials*, 212–215:143–147, 1994.
- [118] L. Malerba and J.M. Perlado. Basic mechanisms of atomic displacement production in cubic silicon carbide: A molecular dynamics study. *Physical Review B*, 65(045202), 2002.
- [119] R. Rurali et al. First-principles studies of the diffusion of B impurities and vacancies in SiC. *Physical Review B*, 69(125203), 2004.
- [120] F. Gao and W.J. Weber. Mechanical properties and elastic constants due to damage accumulation and amorphization in SiC. *Physical Review B*, 69(224108), 2004.
- [121] F. Gao, E.J. Bylaska, W.J. Weber, and L.R. Corrales. Ab initio and empirical-potential studies of defect properties in 3C-SiC. *Physical Review B*, 64(245208), 2001.

- [122] A.D. Brailsford and R. Bullough. The theory of sink strengths. *Philosophical Transactions of the Royal Society of London*, 302(1465):87–137, 1981.
- [123] A.A. Semenov and C.H. Woo. Classical nucleation theory of microstructure development under cascade-damage irradiation. *Journal of Nuclear Materials*, 323:192–204, 2003.
- [124] M.E.J. Newman and G.T. Barkema. *Monte Carlo Methods in Statistical Physics*. Oxford University Press, 1st edition, 1999.
- [125] K. Binder, editor. *Monte Carlo Methods in Statistical Physics*. Springer-Verlag, 1986.
- [126] A. F. Voter. Introduction to the kinetic Monte Carlo method. In K.E. Sickafus and E.A. Kotomin, editors, *Radiation Effects in Solids*. Springer NATO Publishing Unit, 2005.
- [127] I. Sobol. *A primer for the Monte Carlo method*. CRC press, 1994.
- [128] W.M. Young and E.W. Elcock. Monte Carlo studies of vacancy migration in binary ordered alloys: I. *Proceedings of the Physical Society of London*, 89(3):735–746, 1966.
- [129] A.B. Bortz, M.H. Kalos, and J.L. Lebowitz. New algorithm for Monte-Carlo simulation of Ising spin systems. *Journal of Computational Physics*, 17(1):10–18, 1975.
- [130] A. Voter. Parallel replica method for dynamics of infrequent events. *Physical Review B*, 57(22):R13985–R13988, 1998.
- [131] A.F. Voter, F. Montalenti, and T.C. Germann. Extending the time scale in atomistic simulation of materials. *Annual Review of Materials Research*, 32:321–346, 2002.
- [132] S. Birner et al. Accelerated dynamics simulations of interstitial-cluster growth. *Solid State Communications*, 120:279–282, 2001.
- [133] Y. Duan et al. Mechanisms of lithium transport in amorphous polyethylene oxide. *Journal of Chemical Physics*, 122(054702), 2005.
- [134] Y. Mishin et al. Stick-slip behavior of grain boundaries studied by accelerated molecular dynamics. *Physical Review B*, 75(224101), 2007.
- [135] D.J. Amit and Y. Verbin. *Statistical Physics: An Introductory Course*. World Scientific, 1999. Translated from Hebrew by R. Tzafriri.
- [136] M. Bockstedte, A. Mattausch, and O. Pankratov. Ab initio study of the migration of intrinsic defects in 3C-SiC. *Physical Review B*, 68(205201), 2003.

- [137] E. Rauls, Th. Frauenheim, A. Gali, and P. Déak. Theoretical study of vacancy diffusion and vacancy-assisted clustering of antisites in SiC. *Physical Review B*, (155208), 2003.
- [138] D.J. Bacon et al. Computer simulations of defect production by displacement cascades in metals. *Nuclear Instruments and Methods in Physics Research B*, 102:37–46, 1995.
- [139] A. Suidi et al. Dependence of radiation damage accumulation in iron on underlying models of displacement cascades and subsequent defect migration. *Journal of Nuclear Materials*, 355:89–103, 2006.
- [140] Z. Yao, M.J. Caturla, and R. Schäublin. Study of cascades damage in Ni by MD with different interatomic potentials. *Journal of Nuclear Materials*, 367–370:298–304, 2007.

Appendix A

Appendices

A.1 LAMMPS Tutorials

The two subsections given here provide two examples of the input scripts used to generate the raw data for the simulations discussed in Chapter 4. Each uses precisely the same code and the majority of the input script is the same, though there are important differences in the boundary and initial conditions, as well as the post-processing. The post-processing methods have been discussed in Section 2.3 and scripts are provided in Appendix A.2, therefore the details of the post-processing methodology will not be repeated here. Additionally, the version of LAMMPS quoted here is the latest version that was used in the course of the work. Later versions may work with or without adjustments, however, they have not been tested. At the time of writing, the USER-ZBLTERSOFF package had not yet been added

to the standard LAMMPS release (a copy was sent in to Steve Plimpton for evaluation and inclusion in June 2008), however an electronic copy has been given to Prof. Wing Kam Liu to keep with this dissertation so that it can be made available to the reader by someone other than the author.

This tutorial assumes a basic familiarity with Unix-like operating systems and command lines, thus syntax will not be explained. For those unfamiliar, there are a large number of good tutorials available online in addition to books for beginners. No implicit guarantees, etc. have been made regarding security, functionality or compatibility of this code. Use at your own risk.

Requirements:

- Unix-like OS environment (Mac OS X, Linux, Cygwin under Windows, etc)
- LAMMPS release 21Apr08 (see <http://lammps.sandia.gov/>)
- LAMMPS USER-ZBLTersoFF package (Tersoff/ZBL potential capability)

Once these components are installed and tested, the following input scripts can be called with the appropriate command for your system (see the LAMMPS online documentation for particulars). This command should be of the form:

```
$LAMMPSpath/src/lmp_-$BuildType < $InputFile
```

Where \$LAMMPSpath is the path to the LAMMPS release directory, \$BuildType is the ar-

chitecture (as determined by the build) and \$InputFile is the name of the input file (including path, if needed).

A.1.1 Threshold Displacement Energy Example Input

This is an example of the input script used in the TDE calculation validation case discussed in Section 3.2.1. This particular example is for a Si PKA with a 40 eV kinetic energy ‘kick’ along $\langle 001 \rangle$. The ‘kick’ provides a velocity to the atom, but this velocity is not kept constant throughout the simulation, and is affected by interactions with other atoms. A single run required less than 1 hour on 1 processor of an ARL supercomputer, and produced only a few megabytes of data. The ‘#’-sign at the beginning of a line indicates a comment, not a command.

```
# 3C SiC TDE test (Si PKA)
# Repro of Devanathan et al JNM 253 1998
#
# Description:
# 10 x 10 x 10 unit cells, cubic basis (should be 8,000 atoms)
# [100], [010], [001] periodic (w/fixed volume, not fixed atoms though)
# NVE simulation, T=150K initial temperature
# 1) set initial temp, let system equilibriate for 2.1 ps - 2100 timesteps
# 2) apply velocity corresponding to desired PKA KE
# 3) allow to evolve for 3.9 ps - 3900 timesteps

# Dave Farrell (Feb 8 2008)

units metal          # sets units to 'metal' units - Atomic units w/ ps timescale
dimension           3
boundary            p p p          # periodic BCs on x,y,z plane faces

atom_style          atomic          # default attribute style
neighbor            1.0 bin         # sets maximum neighbor search radius to cutoff+value, using bin-sort algorithm
neigh_modify        delay 5 check yes # checks if neighbor list should be rebuilt every 5 steps

# Create geometry using internal stuff
# uses diamond lattice with an 8 atom basis to construct the 3C-SiC structure
lattice              diamond 4.359
```

```

region          box block 0 10 0 10 0 10
create_box      2 box                # initialize box
create_atoms 2 box basis 1 2 basis 2 2 basis 3 2 basis 4 2 basis 5 1 basis 6 1 basis 7 1 basis 8 1 # create atoms

mass           1    12.0              # type 1 = C
mass           2    28.1              # type 2 = Si

# Define potential
pair_style zbl/tersoff .529 .00552635 1.0
pair_coeff * * SiC_Devanathan_JNM_1998_zbl.tersoff C Si

# Groups
region         rallatoms block INF INF INF INF INF INF
region         rnearPKA block 4 6 4 6 4 6
group PKA id id 4441 # Si - 4441, C- 4445
group nearPKA region rnearPKA

# Initialization
compute 1 all temp
compute 2 all coord/atom 2.2
compute 3 all ke/atom
velocity all create 150.0 12345 # initialize atom velocities to temperature, seed random num generator

# set up the ensemble for force calculations
fix 1 all nve

timestep 0.001
thermo 100
thermo_style custom step temp pe etotal press vol
run 2100 # run the equilibration phase

# set PKA velocity to correspond to desired energy
# 30 eV Si PKA KE... v = 143.5273
# 40 eV Si PKA KE... v = 165.7310

# 40 eV Si PKA <001>
set group PKA vx 0
set group PKA vy 0
set group PKA vz 165.7310

# run the evolution for a bit
# dump datafiles with position, kinetic energy, coordination number, etc.
dump 1 nearPKA custom 100 TDE_test.dump x y z c_3 tag type c_2
dump 2 PKA custom 100 PKA_traj_TDE_test.dump x y z c_3 tag type c_2
run 3900 # run the post-excitation phase for 3.9 ps

```

A.1.2 Cascade Simulation Example Input

This example is the input script used to generate the validation case discussed in Section 3.2.2. This ran in about 3 hours on 32 processors on an ARL supercomputer, and produced

several gigabytes of data. For testing purposes, run it for a smaller time period. The ‘#’-sign at the beginning of a line indicates a comment, not a command.

```

# Cascade formation in 3C SiC
# Reproduction of Devanathan et al NIMPRB 1998
#
# Description:
# 28 x 28 x 64 unit cells, cubic basis (should be 401,408 atoms)
# [100], [010], [001] periodic (w/fixed volume, not fixed atoms though)
# NVE simulation, T=300K rescaled boundaries
#(4 unit cells deep on each side of the simulation cell but top)
#
# 1) equilibriate for 2 ps (to ensure it is close to 300K starting)
# 2) apply velocity of 2620.549 ang/ps along [4 11 -95] direction
#    (corresponds to 10 keV kinetic energy) to Si atom at center, near top of box
# 3) for the first .2 ps, use a timestep of .00001 ps (.01 fs)
# 4) increase timestep to .0001 ps (.1 fs) and let evolve for 1 ps
# 5) in timestep to .001 ps (1 fs) and let evolve for 10 ps

# Dave Farrell (Feb 1 2008)

units metal                # sets units to 'metal' units - Atomic units w/ ps timescale
dimension      3
boundary       p p p      # periodic BCs on x,y,z plane faces
atom_style     atomic     # default attribute style
neighbor       1.0      bin # sets maximum neighbor search radius to cutoff+value, using bin-sort algorithm

# check if neighbor list should be rebuilt every 5 steps
neigh_modify delay 5 check yes

# Create geometry using internal stuff
# uses diamond lattice with an 8 atom basis to construct the 3C-SiC structure
lattice        diamond 4.359
region         box block 0 28 0 28 0 64
create_box     2 box      # initialize box
create_atoms   2 box basis 1 2 basis 2 2 basis 3 2 basis 4 2 basis 5 1 basis 6 1 basis 7 1 basis 8 1 # create atoms
region        rPKA      sphere 14.5 14 63.5 .1

mass          1      12.0      # type 1 = C
mass          2      28.1      # type 2 = Si

# Define potential
pair_style     zbl/tersoff .529 .00552635 1.0
pair_coeff     * * SiC_Devanathan_JNM_1998_zbl.tersoff C Si

# Groups
region        rallatoms block INF INF INF INF INF INF
region        rinterior block 4 24 4 24 4 INF
region        rexterior block 4 24 4 24 4 INF side out

group interior region rinterior
group exterior region rexterior
group PKA region rPKA

# Initialization

```

```

compute 1 all temp          # compute temperature
compute 2 interior coord/atom 2.2 # compute coordination number
compute 3 interior ke/atom    # compute per atom kinetic energy
velocity all create 300.0 34986 # initialize atom velocities to temperature, seed random number generator

# set up the ensemble for force calculations, thermostat
fix 1 all nve
fix 2 exterior temp/rescale 1 300.0 300.0 0.5 1.0 region rexterior
fix 3 interior temp/rescale 1 300.0 300.0 0.5 1.0 region rinterior

timestep 0.001
thermo 100
thermo_style custom step temp pe etotal press vol
run 1000 # run the equilibration phase

unfix 3 # remove the rescaling fix from the interior
run 1000 # continue to equilibrate

variable totvel atom "sqrt(vx[]^2 + vy[]^2 + vz[]^2)"
compute 4 all reduce max v_totvel

# set PKA velocity to correspond to ~ 10 keV - 2620.549 Ang/ps
set group PKA vx 109.5108
set group PKA vy 301.1547
set group PKA vz -2600.8815

# run initial collision phase with small timestep (.01 fs for .2 ps)
timestep 0.00001

# output step, temperature, system potential energy, total energy, pressure and max atom velocity
# every 100 timesteps to track overall progress
thermo 100
thermo_style custom step temp pe etotal press c_4

# dump datafiles with position, kinetic energy, coordination number, etc.
dump 1 interior custom 100 cascade_3_init_col_pt2.dump x y z c_3 tag type c_2
dump 2 PKA custom 100 PKA_traj_init_col_pt2.dump x y z c_3 tag type c_2
run 10000

# run intermediate phase with intermediate timestep (.1 fs for 1 ps)
timestep 0.0001
undump 1
undump 2
dump 3 interior custom 100 cascade_3_inter_evolve.dump x y z c_3 tag type c_2
dump 4 PKA custom 100 PKA_traj_inter_evolve.dump x y z c_3 tag type c_2
run 10000

# after cascade formed, let evolve for remainder of time using .001 ps timestep (1 fs for 10 ps)
undump 3 # close the previous dumps
undump 4
timestep .001
dump 5 interior custom 100 cascade_3_final_evolve.dump x y z c_3 tag type c_2
dump 6 PKA custom 100 PKA_traj_final_evolve.dump x y z c_3 tag type c_2
run 10000

```

A.1.3 Parallel Replica Dynamics Input Example

This is an example of the PRD capabilities added to LAMMPS. There are 2 input files given, the first generates the initial geometry and the second runs the PRD calculations discussed in Chapter 6. These input scripts require that LAMMPS be built with the USER-PRD and USER-ZBLTersoFF. The first should be run as a single processor simulation on the same architecture that the PRD system will be run on, to ensure that the restart file will be readable. This simulation will run fairly quickly as there is no integration, just initialization.

```
# Pre-generatd vacancies in 3C SiC
# PRD proof-of-concept cascade simulation
#
# Description:
# 15 x 15 x 15 unit cells, cubic basis (should be 27,000 atoms)
# [100], [010], [001] periodic (w/fixed volume, not fixed atoms though)
# Remove 20 vacancies placed according to perfect crystal spatial distribution

# Dave Farrell (Aug 4 2008)

units metal                # sets units to 'metal' units - Atomic units w/ ps timescale
dimension      3
boundary       p p p       # periodic BCs on x,y plane faces, z dimension fixed volume

atom_style     atomic      # default attribute style
neighbor       1.0 bin     # sets maximum neighbor search radius to cutoff+value, using bin-sort algorithm
neigh_modify   delay 5 check yes # checks if neighbor list should be rebuilt every 5 steps

# Create geometry using internal stuff

lattice        diamond 4.3765
region         box block 0 15 0 15 0 15
create_box     2 box      # initialize box
create_atoms   2 box basis 1 2 basis 2 2 basis 3 2 basis 4 2 basis 5 1 basis 6 1 basis 7 1 basis 8 1 # create atoms

mass          1      12.0    # type 1 = C
mass          2      28.1    # type 2 = Si

# Define potential
pair_style     zbl/tersoff .529 .00552635 1.0
pair_coeff     * * SiC_Devanathan_JNM_1998_zbl.tersoff C Si

# Groups
region        rallatoms block INF INF INF INF INF INF
region        rinterior block 2 13 2 13 2 INF
```

```

region          rexterior block 2 13 2 13 2 INF side out

group interior region rinterior
group exterior region rexterior

group vacancies id 4619 8226 9687 10030 11699 13398 13490 13627 15048 15539 16832 16957 17217 17220
17225 17226 17787 18896 19004 19165
delete_atoms group vacancies

run 0
write_restart PRD_test_initial_geom.restart

```

The second input file actually runs the PRD simulation. This requires that the ‘-partition’ flag be used and shows 32 partitions of 1 processor each being used. This simulation ran in about 10 hours on an ARL supercomputer. The line to execute it is

```
$mpirun$LAMMPSpath/src/lmp_$_BuildType - partition 32x1 - in $InputFile
```

Where *\$mpirun* is the appropriate MPI run command for your system. The input file is as follows:

```

# This script is intended to set up 32 different, parallel instances of LAMMPS
# Steps
# Read in restart file
# run dephasing period for 1 ps w/different seeds in each replica
# run standard evolution & check phase (check every .005 ps)
# if a transition is detected, run for an additional 1 ps to look for crossover events
# if stable event, output restart and dump file and exit

# Specify the loop variable, thermostat seeds (for dephasing)
variable LoopVar loop 15 # loop variable for number of transitions to track
variable therm_seed world 6260 17524 3575 10887 5544 749 26499 15848 3132 29529 20503 15809 17620 19736 28612 30931
8536 27279 20074 31591 5217 23502 12232 25470 11163 16601 14053 11691 18956 14859 24807 25216 # Thermostat seeds

label loopstart
# read restart file (final_evolve final config from the cascade)
if ${LoopVar} == 1 then "read_restart PRD_test_final_evolve.restart"
if ${LoopVar} > 1 then "clear"
if ${LoopVar} > 1 then "read_restart PRD.restart"

```

```

# re-specify everything that is needed
# Neighbor finder information
neighbor      1.0      bin          # sets maximum neighbor search radius to cutoff+value, using bin-sort algorithm
neigh_modify delay 5 check yes      # checks if neighbor list should be rebuilt every 5 steps

# Define potential
pair_style zbl/tersoff .529 .00552635 1.0
pair_coeff * * SiC_Devanathan_JNM_1998_zbl.tersoff C Si

# Regions
region                rallatoms block INF INF INF INF INF INF units box
region                rinterior block 8.753 56.8945 8.753 56.8945 8.753 INF units box
region                rexterior block 8.753 56.8945 8.753 56.8945 8.753 INF side out units box

# Computes
compute 1 all temp
compute 2 interior coord/atom 2.2
compute 3 interior ke/atom

# Simulation fixes
fix 1 all nve
fix 2 exterior temp/rescale 1 2500.0 2500.0 0.5 1.0 region rexterior

# basic run information
timestep 0.001
thermo 100
thermo_style custom step temp pe etotal press

#-----
# De-phase the replicas
#-----
velocity all create 2500.0 ${therm_seed} # initialize atom velocities to temperature, seed random num generator

# run the dephasing stage (1 ps)
run 1000

#-----
# now begin the evolution/check phase
#-----
# have to define some new computes
fix 3 all coord/original
compute 4 all displace/atom 3
fix 4 interior ave/atom 1 5 5 c_4[4]
compute 5 all reduce max f_4
thermo_style custom step temp pe etotal press c_5
thermo 100
dump 1 interior custom 1000000 PRD_test_*.dump x y z c_3 tag type c_2
restart 1000000 PRD.restart PRD.restart

run_PRD 10000 5 c_5 1.8 1000
# syntax here is max # of steps, check interval, criterion compute, criterion threshold,
# # steps to run for crossover check (should be same as dephasing)

next LoopVar
jump PRD_test.in loopstart

```

A.1.4 LAMMPS PRD Command Function

The following is the core code for the `run_PRD` function. This does not contain the information contained in the header files nor the code to print out the status of the run but should be sufficient to re-write the implementation if needed. Further, the required code has been given to Prof. Wing Kam Liu to keep with this dissertation so that it can be made available to the reader by someone other than the author.

```

/* -----
perform PRD evolution and checks with inter-world communications (for checks)

Syntax:

run_PRD N M compute trans_dist ncrosscheck

N - total number of simulation steps to run before exiting
M - check for a transition every M steps
compute - compute ID that determines maximum displacement
trans_dist - the displacement threshold criteria for a 'transition'
ncrosscheck - number of steps to run to check for followon events (at least correlation/dephase time)
----- */

void RunPRD::command(int narg, char **arg)
{
    if (universe->nworlds == 1)
        error->all("Must have more than one processor partition to run Parallel Replica Dynamics");
    if (domain->box_exist == 0)
        error->all("PRD command before simulation box is defined");
    if (narg != 5 ) error->universe_all("Illegal PRD command - check arguments");

    int nsteps = atoi(arg[0]);
    nevery = atoi(arg[1]);
    ncrosscheck = atoi(arg[4]);

    // get the compute ID
    char *id_comp_full = arg[2];
    char *word = strtok(id_comp_full, " \0");

    double my_maxdisp;
    class Compute *maxdispcompute;

    if (strncmp(word, "c_", 2) == 0)
    {
        int n = strlen(word);
        char *id = new char[n];
        strcpy(id, &word[2]);
    }
}

```

```

char copy[9];
strncpy(copy,id,8);
copy[8] = '\0';

n = modify->find_compute(id);
if (n < 0) error->all("Could not find maximum displacement compute ID");
if (modify->compute[n]->scalar_flag == 0)
    error->all("run_PRD compute ID does not compute scalar info");

maxdispcompute = modify->compute[n];
if (!maxdispcompute) error->all("Could not access compute");
my_maxdisp = 0.0;    // Initialize the max displacement to 0
}
else error->universe_all("Illegal compute ID in run_PRD command - check arguments");

// Distance criteria for transitions
double trans_dist = atof(arg[3]);

// Check frequency must evenly divide total # of timesteps
int nchecks = 0;
if (nevery == 0) error->universe_all("Invalid frequency in run_PRD command");
    nchecks = nsteps/nevery;
if (nchecks*nevery != nsteps)
    error->universe_all("Non integer # of checks in run_PRD command");

// setup for long PRD run
update->whichflag = 0;
update->nsteps = nsteps;
update->beginstep = update->firststep = update->ntimestep;
update->endstep = update->laststep = update->firststep + nsteps;

lmp->init();

// local storage
me_universe = universe->me;
MPI_Comm_rank(world,&me);
nworlds = universe->nworlds;
iworld = universe->iworld;

// create MPI communicator for root proc from each world
int color;
if (me == 0) color = 0;
else color = 1;
MPI_Comm_split(universe->uworld,color,0,&roots);

// world2root[i] = global proc that is root proc of world i
world2root = new int[nworlds];
if (me == 0)
    MPI_Allgather(&me_universe,1,MPI_INT,world2root,1,MPI_INT,roots);
MPI_Bcast(world2root,nworlds,MPI_INT,0,world);

// signal to update displacement fix, access my_maxdisp
// maxdisp[i] = maximum displacement of world i
modify->clearstep_compute();
my_maxdisp = maxdispcompute->compute_scalar();
maxdisp = new double[nworlds];
if (me == 0)
    MPI_Allgather(&my_maxdisp,1,MPI_DOUBLE,maxdisp,1,MPI_DOUBLE,roots);
MPI_Bcast(maxdisp,nworlds,MPI_DOUBLE,0,world);

```

```

// transition flag - initialize on all procs
int trans_world = -1;

// setup PRD runs
MPI_Status status;

if (me_universe == 0 && universe->uscreen)
    fprintf(universe->uscreen,"Setting up Parallel Replica Dynamics Evolve and Check run ...\\n");
if (me_universe == 0 && universe->ulogfile)
    fprintf(universe->ulogfile,"Setting up Parallel Replica Dynamics Evolve and Check run ...\\n");

update->integrate->setup();

MPI_Barrier(universe->uworld);          // wait for all processes to get here
timer->barrier_start(TIME_LOOP);

if (me_universe == 0)
{
    if (universe->uscreen) fprintf(universe->uscreen,"Step D1 D2 ...\\n");
    if (universe->ulogfile) fprintf(universe->ulogfile,"Step D1 D2 ...\\n");
    print_status();
}

// begin the loop over the checks
for (int ichecks = 0; ichecks < nchecks; ichecks++)
{
    // run for nevery timesteps before check
    update->integrate->iterate(nevery);

    // Compute the maximum displacement and communicate it
    // alert the compute to be re-computed
    modify->clearstep_compute();
    my_maxdisp = maxdispcompute->compute_scalar();
    if (me == 0)
        MPI_Allgather(&my_maxdisp,1,MPI_DOUBLE,maxdisp,1,MPI_DOUBLE,roots);
    MPI_Bcast(maxdisp,nworlds,MPI_DOUBLE,0,world);

    // status report (done at every check)
    if (me_universe == 0) print_status();

    // Perform check against transition criteria (each master process does the check)
    if (me == 0)
    {
        // loop through the worlds to see if which ones have exceeded the threshold,
        // take the maximum value as the one to continue with
        double maxdispcheck = 0.0;
        for (int i=0; i < nworlds; i++)
        {
            if ((maxdisp[i] >= trans_dist) && (maxdisp[i] >= maxdispcheck))
            {
                // if transition detected,
                // find world with highest displacement in case of multiple transitions
                maxdispcheck = maxdisp[i];
                trans_world = i;
            }
        }
    }
}
// broadcast the transworld to all procs in world (now all procs in all worlds should have same value)
MPI_Bcast(&trans_world,1,MPI_INT,0,world);

```



```

int trans_world_temp = trans_world;

// if a transition has been found, hold non-transitioned states.
// run 1 more check phase on the transitioned state to see if the transition is stable.
// If no transition detected, carry on as usual.
if ((me_universe == 0) && (trans_world > -1))
{
    if (universe->uscreen)
        fprintf(universe->uscreen,"Possible Transition Detected in World %d\n", trans_world);
    if (universe->ulogfile)
        fprintf(universe->ulogfile,"Possible Transition Detected in World %d\n", trans_world);
}

if ((trans_world > -1) && (iworld == trans_world))
{
    // run for ncrosscheck timesteps before check
    update->integrate->iterate(ncrosscheck);

    // Compute the maximum displacement and communicate it
    // alert the compute to be re-computed
    modify->clearstep_compute();
    my_maxdisp = maxdispcompute->compute_scalar();

    if (my_maxdisp >= trans_dist)
    {
        // Stable Transition
        if (universe->uscreen && me == 0)
            fprintf(universe->uscreen,"Stable Transition Detected in World %d\n", trans_world);
        if (universe->ulogfile && me == 0)
            fprintf(universe->ulogfile,"Stable Transition Detected in World %d\n", trans_world);
    }
    else
    {
        // Transition Unstable, reset the transition flag and continue
        if (universe->uscreen && me == 0)
            fprintf(universe->uscreen,"No Stable Transition Detected in World %d, continue search\n",
                trans_world);
        if (universe->ulogfile && me == 0)
            fprintf(universe->ulogfile,"No Stable Transition Detected in World %d, continue search\n",
                trans_world);
        trans_world = -1;
        MPI_Bcast(&trans_world,1,MPI_INT,0,world);
    }
}

MPI_Bcast(&trans_world,1,MPI_INT,world2root[trans_world_temp],universe->uworld);
MPI_Bcast(&trans_world,1,MPI_INT,0,world);
// Use a barrier to hold jobs - will wait for the check to be done if there was a transition,
// otherwise will just wait until all procs get here to begin the next check phase...
// easier than series of if statements
MPI_Barrier(universe->uworld);

// if there was a stable transition on any process, all processes on all worlds break out of the loop
if (trans_world > -1)
{
    if (universe->ulogfile && me_universe == 0)
        fprintf(universe->ulogfile,"Stable Transition Detected in World %d\n",
            trans_world);
    break;
}

```

```

    }
}

// convert simulation nsteps to the actual number of steps each world ran
update->nsteps = (update->ntimestep - update->beginstep);

// output the total number of steps run (and total time integrated) in this phase
if (me == 0)
{
    int totsteps = 0;
    MPI_Allreduce(&(update->nsteps), &totsteps, 1, MPI_INT, MPI_SUM, roots);

    if (universe->ulogfile && me_universe == 0)
    {
        double tottime = totsteps*(update->dt);
        fprintf(universe->ulogfile, "Total Simulation Steps: %d Simulated Time: %lf\n", totsteps, tottime);
    }
}

// Output the final thermo output, a restart and a dump file world with the stable transition state
if ((trans_world > -1) && (iworld == trans_world))
{
    if (domain->triclinic) domain->x2lamda(atom->nlocal);
    domain->pbcc();
    domain->reset_box();
    comm->setup();
    comm->exchange();
    if (domain->triclinic) domain->lamda2x(atom->nlocal);

    for (int idump = 0; idump < output->ndump; idump++)
        output->next_dump[idump] = update->ntimestep;
    output->next_dump_any = update->ntimestep;
    if (output->restart_every) output->next_restart = update->ntimestep;
    output->next_thermo = update->ntimestep;

    modify->addstep_compute_all(update->ntimestep);
    output->write(update->ntimestep);
}

MPI_Barrier(universe->uworld);
timer->barrier_stop(TIME_LOOP);

// finish up, cleanup
update->integrate->cleanup();
Finish finish(lmp);
finish.end(1);

update->whichflag = -1;
update->firststep = update->laststep = 0;
update->beginstep = update->endstep = 0;
}

```

A.2 Post-processing Scripts, Codes

A.2.1 Point Defect Identification Scripts

This section includes the source code for a Matlab script that can be used to identify point defects using the algorithm described in Section 2.3.2. Another, nearly identical, C-based version was used for the calculations in Section 4.2 and Chapter 5 due to concerns over speed and efficiency of the Matlab code. This script was tested and used under Matlab version 7.0. The ‘%’-sign denotes a comment, not a command.

```
% Script to find vacancies, interstitials by way of a Delaunay Triangulation

% takes input in LAMMPS *.dump custom format: x y z KE tag type CN
% Requires 1 file to act as the 'reference' geometry, which all the others will be compared against.
% The reference geometry should be set up such that all of the atoms are in the
% appropriate equilibrium locations

function occupancy_finder(init_geom_file, result_file, num_steps, output_file, defect_output_file)
    tic;
    init_geom_fid = fopen(init_geom_file,'r');
    if init_geom_fid == -1
        error('could not open reference geometry file')
    end
    % Read in the geometry, putting the header information (9 lines of data)
    % into another array, for later use
    header = cell(9,1);
    for i=1:9
        header_temp = fgetl(init_geom_fid);
        header{i} = header_temp;
    end
    num_atom = str2num(header{4});

    % Read in the atomic information
    ref_data = zeros(num_atom,7);
    ref_data = textread(init_geom_file, '', 'headerlines',9);
    fclose(init_geom_fid);
    dt=toc;
    disp(sprintf('Reading reference geometry:\t    DONE (dt = %f s)', dt))

    % Calculate the delaunay triangulation, store it for use through the run
    tic;
    TRI = delaunayn(ref_data(:,[1:3]));
    dt=toc;
    disp(sprintf('Calculating Delaunay Triangulation:\t    DONE (dt = %f s)', dt))
```

```

% Begin loop over number of comparison steps
result_file_fid=fopen(result_file,'r');
if result_file_fid == -1
    error('could not open result file')
end

output_file_fid=fopen(output_file,'w');

% format is [timestep, total # vacancies, # C vacancies, # Si Vacancies,
% total # interstitials, # C_Si antisites, # Si_C antisites, # C replacements, # Si replacements]
defect_time_array=zeros(num_steps,9);

for step=1:num_steps
    % read in the current step data
    tic;
    % Header first
    if (step==1)
        status = fseek(result_file_fid,0,'cof');
    else
        status = fseek(result_file_fid,2,'cof');
    end
    result_step_header = cell(9,1);
    for i=1:9
        header_temp = fgetl(result_file_fid);
        result_step_header{i} = header_temp;
    end
    num_result_atom = str2num(result_step_header{4});
    num_result_step = str2num(result_step_header{2});

    % Now the atomic data
    %result_data = zeros(num_result_atom,7);
    result_data = fscanf(result_file_fid,'%f %f %f %f %i %i %i',[7,num_result_atom]);
    result_data = result_data';
    dt=toc;
    pos = ftell(result_file_fid);
    disp(sprintf('Reading comparison geometry for comparison step %i:\t DONE (dt = %f s)',step,dt))

    % Perform the search for the current step
    % each row # in ind corresponds to the atom # in result_data,
    % each entry is the # of the nearest point in ref_data
    tic;
    ind = zeros(num_result_atom,1);
    ind = dsearchn(ref_data(:,[1:3]),TRI,result_data(:,[1:3]));
    dt=toc;
    disp(sprintf('Calculating nearest points:\t DONE (dt = %f s)',dt))

    tic;
    % now convert that to an occupancy
    % each row # in occup corresponds to the atom # in ref_data
    occup = zeros(size(ref_data,1),1);
    for i=1:length(ind)
        occup(ind(i)) = occup(ind(i)) + 1;
    end

    % get some data on number of point defects out
    vac_count = 0;
    vac_count_1 = 0;
    vac_count_2 = 0;
    inter_count = 0;

```

```

inter_site_count = 0;

for i=1:length(occup)
    if (occup(i) == 0)
        vac_count = vac_count+1;
        if (ref_data(i,6) == 1)
            vac_count_1 = vac_count_1 + 1;
        elseif (ref_data(i,6) == 2)
            vac_count_2 = vac_count_2 + 1;
        end
    elseif (occup(i) > 1)
        inter_site_count = inter_site_count+1;
        inter_count = inter_count + occup(i) - 1;
    end
end

antisite_count_1_2 = 0;
antisite_count_2_1 = 0;
replacement_count_1 = 0;
replacement_count_2 = 0;

for i=1:length(ind)
    if (occup(ind(i)) == 1)
        if (ref_data(ind(i),6) == 1 && result_data(i,6) == 2)
            antisite_count_2_1 = antisite_count_2_1 + 1;
        elseif (ref_data(ind(i),6) == 2 && result_data(i,6) == 1)
            antisite_count_1_2 = antisite_count_1_2 + 1;
        elseif ((ref_data(ind(i),6) == 1 && result_data(i,6) == 1) && ...
            (ref_data(ind(i),5) ~= result_data(i,5)))
            replacement_count_1 = replacement_count_1 + 1;
        elseif ((ref_data(ind(i),6) == 2 && result_data(i,6) == 2) && ...
            (ref_data(ind(i),5) ~= result_data(i,5)))
            replacement_count_2 = replacement_count_2 + 1;
        end
    end
end

% put the entries into the time-history array. Assumes occup > 1 is an interstitial
defect_time_array(step,:) = [num_result_step vac_count vac_count_1 vac_count_2 ...
    inter_count antisite_count_1_2 antisite_count_2_1 replacement_count_1 replacement_count_2];

dt=toc;
disp(sprintf('Converting neighbors to occupancy:\t DONE (dt = %f s)',dt))

% Print out some statistics from each comparison step
disp(sprintf('\n**** Step %i Summary ****',step))
disp(sprintf('Num Ref Atoms: %i \t Num Comp Atoms: %i \t Total Occupancy: %i',num_atom ...
,num_result_atom,sum(occup)))
disp(sprintf('Max Occupancy: %i \t Min Occupancy: %i',max(occup),min(occup)))
disp(sprintf('Num Over-Occupied Sites (Total Over-Occupancy): %i (%i) ...
\t Num Unoccupied Sites: %i', inter_site_count,inter_count, vac_count))
disp(sprintf('Num C_Si Antisites: %i \t Num Si_C Antisites: %i',antisite_count_1_2, ...
antisite_count_2_1))
disp(sprintf('Num C Replacements: %i \t Num Si Replacements: %i',replacement_count_1, ...
replacement_count_2))
disp(sprintf('*****\n'))
% Output the occupancy to a file, in LAMMPS *.dump custom format: x y z tag type OCCUP
output_data=cat(2,ref_data(:,[1:3 5 6]),occup);

% Output the header for the current step

```

```

fprintf(output_file_fid,'%s\n',header{1});
fprintf(output_file_fid,'%i\n',num_result_step);
fprintf(output_file_fid,'%s\n',header{3});
fprintf(output_file_fid,'%i\n',num_atom);
fprintf(output_file_fid,'%s\n',header{5});
fprintf(output_file_fid,'%s\n',header{6});
fprintf(output_file_fid,'%s\n',header{7});
fprintf(output_file_fid,'%s\n',header{8});
fprintf(output_file_fid,'%s\n',header{9});
% Now do the atomic data
fprintf(output_file_fid,'%f %f %f %i %i %i\n',output_data');
end
fclose(result_file_fid);
fclose(output_file_fid);

% output time history file, should be ready for plotting in gnuplot
t_hist_fid = fopen(defect_output_file,'w');
fprintf(t_hist_fid,'# Timestep\tTotal Vacancies\tC Vacancies\tSi Vacancies\tTotal Interstitials ...
\tC_Si Antisites\tSi_C Antisites\tC Replacements\tSi Replacements\n');
fprintf(t_hist_fid,'%i\t%i\t%i\t%i\t%i\t%i\t%i\t%i\t%i\n',defect_time_array');
fclose(t_hist_fid);

```

A.2.2 LAMMPS *.dump-EnSight Conversion and Paraview Scripts

These scripts make use of the Pizza.py python toolkit for interaction with LAMMPS. It is available from (<http://www.cs.sandia.gov/sjplimp/pizza.html>) and Paraview, the Parallel Visualization Application (<http://www.paraview.org>). The first script takes the last step of the final dump file that results from the point defect script being run on the output from the example input file given in Appendix A.1.2 and converts it to EnSight format to be read in Paraview by the second script.

The dump to EnSight script:

```

d = dump("last_step_final_evolve_occ.dump")
d.map(1,"x", 2,"y", 3,"z", 4, "id", 5, "type", 6, "occ")
e = ensight(d)
e.change = 1
e.maxtype = 2
e.one("last_step_final_evolve_occ_out", "type", "ATNUM", "occ", "OCC")
sys.exit()

```

The Paraview batch visualization script is below. This one isolates the defect by atom type, represents them with spheres of different sizes and colors and outputs an image to disk.

```
#!/usr/bin/python
#
#  ensight_render_test.py
#

from paraview import servermanager
import time
import sys

if (len(sys.argv) != 3):
    print "Usage: ensight_render_SiC_Cascade.py <case_file> <output_name>"
    sys.exit()

# Set up the connection to the Paraview Server
if not servermanager.ActiveConnection:
    connection = servermanager.Connect()
    if not connection:
        print "Connection to server failed."
        sys.exit()

# create an ensight source object
ensightSource = servermanager.sources.ensight()
ensightSource.CaseFileName = sys.argv[1]
ensightSource.UpdatePipelineInformation()

# Isolate the C, Si groups
C_atoms = servermanager.filters.ExtractLevel(Input=ensightSource,GroupRange=[0, 0])
Si_atoms = servermanager.filters.ExtractLevel(Input=ensightSource,GroupRange=[1, 1])

# Set up the Occupancy thresholds
C_vacancies = servermanager.filters.Threshold(Input=C_atoms,AllScalars=1,SelectInputScalars=['0','0','0','0','0CC'],ThresholdBetween=[0,0])
C_interstitials = servermanager.filters.Threshold(Input=C_atoms,AllScalars=1,SelectInputScalars=['0','0','0','0','0CC'],ThresholdBetween=[2,10])
Si_vacancies = servermanager.filters.Threshold(Input=Si_atoms,AllScalars=1,SelectInputScalars=['0','0','0','0','0CC'],ThresholdBetween=[0,0])
Si_interstitials = servermanager.filters.Threshold(Input=Si_atoms,AllScalars=1,SelectInputScalars=['0','0','0','0','0CC'],ThresholdBetween=[2,10])

# Set up the sphere glyphs for the defects
#C_sphere = servermanager.sources.SphereSource(Radius=1.2, ThetaResolution=32, PhiResolution=32) # basic sphere, radius 1.2
C_sphere = servermanager.sources.SphereSource(Radius=1.8, ThetaResolution=32, PhiResolution=32) # basic sphere, radius 1.8
C_vac_glyph = servermanager.filters.Glyph(Input=C_vacancies, Source=C_sphere, SetScaleMode=3, SetScaleFactor=1.0)
C_inter_glyph = servermanager.filters.Glyph(Input=C_interstitials, Source=C_sphere, SetScaleMode=3, SetScaleFactor=1.0)

#Si_sphere = servermanager.sources.SphereSource(Radius=2.4, ThetaResolution=32, PhiResolution=32) # basic sphere, radius 2.4
Si_sphere = servermanager.sources.SphereSource(Radius=2.2, ThetaResolution=32, PhiResolution=32) # basic sphere, radius 2.2
Si_vac_glyph = servermanager.filters.Glyph(Input=Si_vacancies, Source=Si_sphere, SetScaleMode=3, SetScaleFactor=1.0)
Si_inter_glyph = servermanager.filters.Glyph(Input=Si_interstitials, Source=Si_sphere, SetScaleMode=3, SetScaleFactor=1.0)

# Set up a bounding box... may make visualization easier/more accurate
box = servermanager.sources.CubeSource(Center=[109.4875, 109.4875, 117.6725], XLength=184.349, YLength=184.349, ZLength=200.719)

# Set up the view, render to a file
view = servermanager.CreateRenderView()
view.ViewSize[0:2] = [600,600] # Set up the rendering window resolution
view.Background[0:3] = [0.329412, 0.34902, 0.427451] # set background color to paraview default blue-grey
#view.Background[0:3] = [1.0, 1.0, 1.0] # set background color to white

# Set up representation of data to be visualized
rep1 = servermanager.CreateRepresentation(C_vac_glyph, view)
rep2 = servermanager.CreateRepresentation(C_inter_glyph, view)
rep3 = servermanager.CreateRepresentation(Si_vac_glyph, view)
rep4 = servermanager.CreateRepresentation(Si_inter_glyph, view)
rep5 = servermanager.CreateRepresentation(box,view)
rep1.Representation = 2 # set representation to points(0), wireframe(1), surface(2), outline(3), surface w/outline(4)
rep2.Representation = 2
rep3.Representation = 2
rep4.Representation = 2
rep5.Representation = 3 # set up the 'bounding box'
rep5.Opacity = 1.0
rep5.LineWidth = 2
rep5.Color[0:3] = [0.0, 0.0, 0.0] # set box color to black
```

```

# Set up the coloring for Atom Type (1 = C, 2 = Si) or Occupancy
lookup_table = servermanager.rendering.PVLookupTable()
rep1.LookupTable = lookup_table
rep1.ColorAttributeType = 0 # set up point coloring
rep1.ColorArrayName = 'OCC' # tell which coloring to use 'ATNUM' or 'OCC', next two are limits
minnum = 0
maxnum = 3
lookup_table.ColorSpace = 1
lookup_table.RGBPoints = [minnum, 0, 0, 1, maxnum, 1, 0, 0] # blue to red bar
#lookup_table.RGBPoints = [minnum, 0, 0, 0, maxnum, .9, .9, .9] # black to grey
# Accessing this to other representations which are to be colored
rep2.LookupTable = lookup_table
rep2.ColorAttributeType = 0
rep2.ColorArrayName = 'OCC'
rep3.LookupTable = lookup_table
rep3.ColorAttributeType = 0
rep3.ColorArrayName = 'OCC'
rep4.LookupTable = lookup_table
rep4.ColorAttributeType = 0
rep4.ColorArrayName = 'OCC'

# Do the rendering/viewport stuff
#view.StillRender() # first one needed to get camera position, etc.
# Or... set it by hand:
view.CameraPosition[0:3] = [600, 109.48730087280273, 120]
view.CameraFocalPoint[0:3] = [117.67165279388428, 109.48730087280273, 120]
view.CameraViewUp[0:3] = [0.0, 0.0, 1.0]
#view.ResetCamera() # Sets camera up to 'default' position based on data
#camera = view.GetActiveCamera()
#camera.Elevation(90)
view.StillRender()

view.WriteImage(sys.argv[2], "vtkPNGWriter", 1)

```