

Breaking the bottleneck in radiation materials science with transient grating spectroscopy

by
Sara Elizabeth Ferry

Dual B. S., Nuclear Science and Engineering and Physics (2011)
Massachusetts Institute of Technology, Cambridge
S. M., Nuclear Science and Engineering (2016)
Massachusetts Institute of Technology, Cambridge

SUBMITTED TO THE DEPARTMENT OF NUCLEAR SCIENCE AND ENGINEERING
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY IN NUCLEAR SCIENCE AND ENGINEERING
AT THE
MASSACHUSETTS INSTITUTE OF TECHNOLOGY

JUNE 2018

©Massachusetts Institute of Technology
All rights reserved

Signature of author: _____
Sara Elizabeth Ferry
Department of Nuclear Science and Engineering
May 17, 2018

Certified by: _____
Michael Short
Norman Rasmussen Assistant Professor of Nuclear Science and Engineering
Thesis Supervisor

Certified by: _____
Ju Li
Battelle Energy Alliance Professor of Nuclear Science and Engineering
Professor of Materials Science and Engineering
Thesis Reader

Certified by: _____
Michael Demkowicz
Associate Professor of Materials Science and Engineering, Texas A&M University
Thesis Committee Member

Accepted by: _____
Ju Li
Chair, Nuclear Science & Engineering Department Committee on Graduate Students

Breaking the bottleneck in radiation materials science with transient grating spectroscopy

Sara E. Ferry

Submitted to the Department of Nuclear Science and Engineering on May 17, 2018 in Partial Fulfillment of the Requirements for the Degree of Doctor of Science in Nuclear Science and Engineering

Nuclear power applications are characterized by harsh mechanical, chemical, thermal, and irradiation environments that present a challenge for the materials engineer. Nuclear materials research and development is a subject of managing constraints: a component must be proven to retain its integrity in the reactor environment for the entirety its operating lifetime, and the material must not impede the delicate neutronics balance that makes a reactor work.

It is not surprising, then, that materials often represent the major engineering hurdle in moving a new reactor concept closer to reality, especially since many advanced reactor concepts utilize higher temperature regimes, larger radiation fluxes, and more corrosive coolants. However, if nuclear materials research is the bridge between academic concept and commercial reality, it is frequently a long and expensive bridge to cross. In order to validate a new material for use in a specific reactor environment, one must test the material in representative conditions, or test the material in a sufficient number of conditions that the material's response to an arbitrary reactor environment can be accurately predicted.

Transient grating spectroscopy (TGS), long used in the materials science field to characterize the properties of thin films, is adapted for use as a method of characterizing radiation-damaged samples. TGS has the ability to simultaneously measure elastic, thermal, and acoustic material properties. It is also non-contact and non-destructive, and relatively inexpensive to build and adapt for different uses. This means it is an ideal candidate for moving the field of nuclear materials closer to the goal of having the ability to fully characterize the radiation-induced property changes in samples in situ and in real-time while they are irradiated. This thesis demonstrates, via a TGS setup built in the MIT Mesoscale Nuclear Materials laboratory, that TGS will be a valid method for quantifying radiation damage by using it to characterize (1) cold-worked irradiated samples, (2) samples with high concentrations of constitutional vacancies, and (3) samples irradiated for 14 years in the EBR-II reactor.

In (1), it is shown that TGS is a viable method for measuring thermal diffusivity changes due to radiation damage at low doses in cold-worked single crystal niobium. In particular, an initial decrease in thermal diffusivity at very low doses is measured, which is attributed to electron scattering by point defects, followed by an increase and saturation of thermal diffusivity as dose increases, which is attributed to less efficient electron scattering as point defects cluster into mesoscale defects. In (2), the impact of vacancies on the TGS signal is considered by using a material with a high concentration of constitutional vacancies that are stable at room temperature. Molecular dynamics simulations showed that increasing vacancies led to a softening material, but the opposite effect was observed in experiments. This study underlines the importance of having better methods of measuring radiation damage in situ, in real time, because ex situ experiments are not capable of capturing defect populations that are produced during irradiation but which anneal out when the irradiation source is removed. In (3), we observe a similar increase in thermal diffusivity with irradiation as was observed in (1), but in this case, the effect is due to radiation-induced segregation removing minor alloying elements. Study (3) also demonstrates the utility of using TGS on real nuclear materials, as the TGS results are consistent with the extensive characterization carried out on these samples by previous

researchers. These three studies illustrate the utility of TGS for characterizing radiation damage in nuclear materials in a cost-effective, time-efficient manner.

Thesis Supervisor: Michael Short

Title: Norman Rasmussen Assistant Professor of Nuclear Science and Engineering

Thesis Committee Member: Michael Demkowicz

Title: Associate Professor of Materials Science and Engineering, Texas A&M University

Thesis Reader: Ju Li

Title: Battelle Energy Alliance Professor of Nuclear Science and Engineering

Professor of Materials Science and Engineering

Chair, Nuclear Science & Engineering Department Committee on Graduate Students

for Clara

Contents

1	<i>Prologue:</i>	
	Breaking the bottleneck in nuclear materials research	31
2	Introduction to radiation damage in metals and alloys	37
2.1	Introduction	37
2.2	The types of radiation	38
2.2.1	Alpha radiation	38
2.2.2	Beta radiation	39
2.2.3	Gamma radiation	39
2.2.4	Neutron radiation	39
2.2.5	Ion radiation	40
2.3	Quantifying the interaction of radiation with materials	40
2.3.1	Stopping power and range	40
2.3.2	Quantifying damage in terms of atomic movement	43
2.4	Mechanisms and effects of radiation damage	45
2.4.1	Radiation damage mechanisms	46
2.4.2	Damage cascades associated with different types of radiation	47
2.4.3	Radiation-induced defects and resultant property changes	51
2.4.3.1	Point defects	51
2.4.3.2	Dislocations	52
2.4.3.3	Defect clusters	54
2.4.3.4	Voids and swelling	54
2.4.3.5	Irradiation-enhanced creep	57
2.4.3.6	Radiation hardening	57
2.4.3.7	Blistering and fuzz on plasma-facing surfaces	57
2.4.3.8	Phase instability and radiation-induced segregation	58
2.4.4	Neutron irradiation versus ion irradiation	60
3	Detecting and measuring radiation damage with existing methods and with transient grating spectroscopy	63
3.1	Detecting and measuring radiation damage	63
3.1.1	A brief overview of destructive methods of studying radiation damage	65
3.1.2	Nondestructive methods	67
3.2	Transient grating spectroscopy: the basic idea	68

3.3	The transient grating spectroscopy experimental setup in the MIT Mesocale Nuclear Materials group	69
3.4	TGS signal analysis	76
4	Transient grating spectroscopy characterization of cold-worked and ion-irradiated single-crystal niobium	79
4.1	Introduction	79
4.2	Original motivation for cold working and irradiating the samples	80
4.3	Preparation and characterization of the single-crystal niobium	83
4.4	Cold-working the single-crystal niobium	83
4.4.1	Polishing the single crystal cold-worked niobium samples	85
4.5	X-ray diffraction (XRD) characterization of the cold-worked niobium samples	87
4.5.1	Basic principles of X-ray diffraction	87
4.5.2	XRD characterization of the cold-worked niobium samples	88
4.5.2.1	Control sample	89
4.5.2.2	1000 lb sample	90
4.5.2.3	2000 lb sample	92
4.5.2.4	2500 lb sample	92
4.5.2.5	3000 lb sample	93
4.5.2.6	Phi scans of all six niobium samples	96
4.5.2.7	Summary of XRD characterization of the niobium sample orientations	97
4.5.3	Pole figures of the cold-worked niobium	97
4.5.4	Characterizing dislocation density via HRXRD rocking curves	100
4.5.5	Conclusions of XRD characterization campaign	104
4.6	Transient grating spectroscopy of the cold-worked single-crystal niobium samples: $v_{SAW}(\theta)$	104
4.6.1	Making the TGS measurements	104
4.6.2	TGS results: $v_{SAW}(\theta)$ of cold-worked single crystal niobium	106
4.6.3	Comparison with predicted TGS response	111
4.7	Analysis of $v_{SAW}(\theta)$ measurements for cold-worked niobium	114
4.8	Transient grating spectroscopy of irradiated cold-worked niobium	117
4.8.1	$v_{SAW}(\theta)$ results	118
4.8.2	Thermal diffusivity and acoustic damping results	123
4.9	Analysis of the TGS results for irradiated and cold-worked niobium	127
4.9.1	$v_{SAW}(\theta)$ results	127
4.9.2	Thermal diffusivity results	129
4.9.3	Acoustic damping results	130
4.10	Conclusions of the irradiated, cold-worked niobium study	132
5	Transient grating spectroscopy of intermetallic Nickel Aluminum	134
5.1	Introduction to B2-NiAl	134
5.1.1	Motivation for using NiAl in this study	135
5.1.2	The structure and lattice parameter of B2-phase NiAl	137
5.2	The NiAl samples used in the TGS experiments	142
5.2.1	NiAl sample fabrication: Batch I (Fall 2015)	142
5.2.1.1	Fabrication of B2-phase NiAl in the literature	142
5.2.1.2	Raw materials	142

5.2.1.3	The Busso samples	142
5.2.1.4	Preparation for arc melting	144
5.2.1.5	Arc melting and post-melt heat treatment	145
5.2.1.6	Metallographic examination of the as-arc-melted NiAl buttons	146
5.2.1.7	Growth of single crystals at Los Alamos National Laboratory	147
5.2.1.8	Preparation of single crystal NiAl samples	148
5.2.1.9	Mixing of samples and loss of identification (Batch I)	149
5.2.1.10	Orienting the samples using X-ray diffraction.	149
5.2.1.11	Composition analysis of single crystal NiAl samples (Batch I)	153
5.2.1.12	Making a second set of samples from Batch I (Batch 1B)	153
5.2.2	NiAl sample fabrication: Batch II (Summer 2017)	154
5.2.2.1	Motivation for making a second batch of NiAl samples	154
5.2.2.2	Making the arc melted buttons for Batch II	154
5.2.2.3	Compositional analysis of Batch II samples	155
5.2.2.4	Metallurgical examination of the post-heat-treatment Batch II samples	156
5.3	TGS of intermetallic NiAl	163
5.3.1	Initial $v_{SAW}(\theta)$ measurements of Batch I and lessons learned	163
5.3.2	TGS measurements of v_{SAW} as a function of NiAl composition	166
5.3.2.1	Results	166
5.3.2.2	Analysis	167
5.3.3	TGS measurements of thermal diffusivity as a function of composition	171
5.3.3.1	Results	171
5.3.3.2	Analysis	171
5.3.4	TGS measurements of acoustic damping in NiAl	173
5.3.4.1	Results	173
5.3.4.2	Analysis	175
5.4	Simulating TGS experiments on NiAl in LAMMPS	176
5.4.1	Building the NiAl test structures	177
5.4.1.1	Initial tests	177
5.4.1.2	Convergence tests	178
5.4.1.3	Making a tool to build NiAl test structures with appropriate concentrations of Ni, Al, vacancies, and anti-site defects	178
5.4.2	LAMMPS lattice parameter tests	180
5.4.3	Results and analysis of TGS simulations on B2-NiAl in LAMMPS	182
5.5	Conclusions of the B2-NiAl study	186
6	TGS examination of stainless steel samples irradiated in EBR-II	188
6.1	EBR-II hex block section samples: original locations	188
6.2	Radiation damage in the EBR-II samples	193
6.2.1	Radiation and temperature history	193
6.2.2	The EBR-II study control sample	195
6.2.3	Hex block swelling and density changes due to radiation damage	195
6.2.4	TEM microscopy of the EBR-II samples and quantification of defect populations	199
6.3	EBR-II TGS Experiments	203
6.3.1	Results	203
6.4	Analysis of TGS measurements of EBR-II samples	206

6.5	Conclusions of the EBR-II irradiated 304 steel study	214
7	Transient grating spectroscopy and radiation materials science: moving forward	216
8	Appendix	220
8.1	Orientation distribution functions of the six cold-worked niobium samples	220
8.2	Automating MATLAB analysis of the thermal and acoustic components of the TGS traces	226
8.3	Additional results from the comparison of experimental $v_{SAW}(\theta)$ results with calculated $v_{SAW}(\theta)$ results for single crystal niobium	232
8.4	Pre-averaging thermal diffusivity and acoustic damping results for the 1500 lb and 2000 lb niobium samples	235
8.5	Additional notes on the material properties of B2-phase intermetallic NiAl	238
8.5.1	Elasticity of intermetallic NiAl	238
8.5.2	Hardness of intermetallic NiAl	239
8.5.3	Ductility of intermetallic NiAl	240
8.5.4	Stiffness	241
8.5.5	Anisotropy	242
8.5.6	Poly- versus single crystal	243
8.6	B2-NiAl TGS LAMMPS simulations: supplementary material	244
8.6.1	LAMMPS input structure	244
8.6.2	MATLAB code for building NiAl structures	244
8.6.3	Shell script for building NiAl test structures	245
8.7	MATLAB scripts for analyzing TGS data (SAW frequency, thermal diffusivity, and acoustic damping)	251
8.7.1	thermal_phase.m	253
8.7.2	find_start_phase.m	262
8.7.3	fit_spectra_peaks_interact.m	271
8.7.4	fit_spectra_peaks.m	273
8.7.5	make_fft_embed_time.m	275
8.7.6	param_extract_time.m	281
9	Acknowledgments	283

List of Figures

2.1	A basic schematic of ballistic-type radiation damage in a material. The impinging radiation knocks atoms from their place in the material’s crystal structure. These atoms may go on to knock other atoms out of place. The resulting disorder to the structure is the basic cause of radiation damage. Image adapted from [2].	38
2.2	Stopping power as a function of distance for 300 MeV protons in water. As the proton loses energy, it loses its remaining energy at a faster rate. This is an example of a Bragg curve. [10]	41
2.3	Stopping power divided by target material density for α -particles with varying energies (0.1 to 10 MeV). The graph shows how stopping power is a function of the incident radiation energy and the properties of the target material. [10]	41
2.4	Result of a SRIM simulation of 4.15 MeV α -particles into an aluminum 7075 target. The expected range of the particles is about 19 μm . A small number of particles diverge from the expected path, perhaps as a result of backscatter from the target nucleus. However, it’s clear that - probability-wise - this event is relatively rare, and the majority of the radiation’s energy would be deposited near the 19 μm depth.	42
2.5	Lifetime dpa and operating temperatures for different reactor types. Advanced reactor concepts - particularly molten salt reactors, sodium fast reactors, and fusion reactors - present major materials challenges, because radiation-facing components must be able to withstand much higher amounts of radiation damage. [13]	44
2.6	The initial damage caused by a PKA in the zircon lattice [25]	46
2.7	The PKA collides with other atoms, creating a damage cascade [25]	46
2.8	The damage begins to anneal out, as most atoms wind up back in a lattice spot [25]	46
2.9	Damage profiles in a nickel target as a function of depth for different radiation species. Heavy ions have higher associated damage but short range; neutrons have much lower associated damage but penetrate into the material bulk. [27]	49
2.10	The progression of radiation damage in time and space, from atomic-level interaction to component-scale response. Impinging radiation creates a damage cascade and thermal spike, characterized by high disorder in the localized lattice near the radiation impact. Most of this damage anneals out, but some is let behind. Defects combine to form larger defects, such as voids. These larger defects may be found uniformly throughout a component, significantly altering its behavior. Over time, radiation damage alters material properties, and frequently limits the lifetime of radiation-facing components in nuclear applications. Adapted from [28].	50
2.11	A dislocation is “a discontinuity at which the lattice shifts from the unsheared to the sheared state” [38]. The schematic of a crystal lattice shows an edge dislocation and a screw dislocation, which are the two basic dislocation configurations. The Burgers vector describes the magnitude and direction of the lattice strain induced by the dislocation. [39]	53

2.12	The Frank-Reed source is a pinned dislocation that forms a closed loop when force is applied. The loop becomes independent of the original dislocation source, which repeats the process. Adapted from [40]	54
2.13	Interstitial loops are often visible in a TEM as black “dots” that appear in geometric spacings. [43]	55
2.14	Interstitial loop growth observed <i>in situ</i> in austenitic stainless steel. [44]	55
2.15	Voids are visible in the microstructure of copper irradiated to 1.1 dpa at 300°C in the Oak Ridge Reactor. [48]	56
2.16	The dimensional effects of void formation are obvious in the sample on the right, which was irradiated to 1.5×10^{23} n/cm ² at 533°. [49]	56
2.17	The left image shows void formation due to neutrons (uniform throughout the bulk of the material) and the right image shows void formation as a result of ion irradiation (short range, material only affected near irradiation surface). Left image: Porollo, Konobeev, and Garner, 2000 (see also: [50]). Right image: [51].	56
2.18	Single crystal copper undergoes hardening due to neutron irradiation. <i>nvt</i> is an archaic unit for fluence. [56]	58
2.19	A “nano-tendrils bundle” that developed on the surface of tungsten exposed to a modulated helium plasma of 7.6×10^{25} and 1.6×10^{22} m ⁻² s ⁻¹ at 1020K. The tendrils were sliced with a FIB and the cross section imaged (right image). [61]	59
2.20	Atom probe tomography of cold-worked 316 stainless steel irradiated with 10 MeV Fe ⁵⁺ shows the effects of radiation-induced segregation near a grain boundary. Silicon and nickel are enriched, while chromium is depleted. [64]	59
2.21	HT9 samples irradiated to the same dpa with neutrons (in FFTF) and with Fe ²⁺ ions exhibit differences in defect population and typical size or density ratio of those defect populations, indicating that ion irradiation is not necessarily a perfect analog to neutron irradiation. [68]	61
3.1	TEM images of fluorapatite irradiated with 1 MeV krypton ions show the progression of amorphization as fluence increases. The TEM allows for imaging and analysis at the atomic level. [77]	66
3.2	Two pulses of laser (the “pump” beams) overlap on a sample (here, illustrated as a thin film on a substrate). The pulses generate an interference pattern on the sample, resulting in thermal expansion that is proportional to the intensity of the interference pattern. The non-uniform thermal expansion results in the propagation of acoustic waves throughout the sample. These acoustic waves interfere with each other, creating a standing acoustic wave (SAW) on the surface. The SAW can be probed with a second laser beam. This “probe” beam diffracts from the SAW. The signal collected from the diffracted probe beam can be used to analyze the properties of the SAW, and thus the properties of the sample. [114]	70

3.3	This schematic was use to explain the TGS setup in the early phases of the project. At (1) the probe and pump laser emit beams that enter the phase mask (2), which determines θ . All diffraction orders of the pump and probe beams are blocked except for the ± 1 . One order of the probe beam is used as the reference beam, while the other is referred to as the probe beam (3). (This is part of the heterodyne scheme, which is explained in greater detail in the text.) The beams enter lens 1, which sets them back on parallel paths. The reference beam passes through an adjustable neutral density filter to avoid detector saturation, while the probe beam passes through an adjustable glass slide that changes its path length (4). Lens 2 is used to recombine the beams (5), with the sample at the focal point. The SAW is created on the sample surface by the pulsed pump beam and the probe beam and reference beam diffract from the SAW. The diffracted light (the first orders of the probe and reference beams) carries the signal information (orange line), and optics are used to send it to the photodetector (6).	71
3.4	A long exposure photograph of a TGS experiment in progress in the MIT MNM laser lab. The green pump beams and red probe beam follow the same path through the setup optics to the sample. The probe beam diffracts from the pump-induced SAW and into the detector.	72
3.5	The probe beam spot overlaps the pump spot, which creates the transient grating on the sample surface. The probe spot is smaller than the pump spot to ensure that the probe is not partially capturing the sample surface outside of the excitation region. [119]	72
3.6	The positive and negative traces are amplified by heterodyning. [114]	74
3.7	The positive and negative heterodyne-amplified traces are collected simultaneously with the new setup, which cuts overall data collection time in half. [123]	75
3.8	A characteristic positive TGS signal.	76
3.9	Elastic, acoustic, and thermal properties can be extracted from the TGS signal.	77
4.1	Niobium, originally called Columbium, was discovered at the turn of the 19th century by Charles Hatchett, and is described here in an 1802 journal article. [126]	80
4.2	Modified image from [127] showing experimentally measured changes in Young's modulus for three electron-irradiated copper samples with varying levels of cold work.	81
4.3	Data from Figure 4.2 replotted (open circles and <i>x</i> s) so it can be compared with the theoretical prediction of $(\Delta E/E)^{-1/2}$ as a function of integrated flux (straight line). (See p.1417 of [127].)	81
4.4	A gradual decrease in modulus is observed with extended irradiation. This behavior is not immediately obvious if only studying material response at low dose. [127]	82
4.5	A typical niobium sample, prior to cold working. Samples are sectioned to be ≈ 1 mm thick from a 5mm diameter cylindrical single crystal. This sample is marked with permanent marker for orientation purposes (this can be easily removed with acetone or another solvent).	83
4.6	The Carver™ pellet press used to cold work the niobium samples. The lever is used to move the stage and apply force to the die.	84
4.7	A Carver™ 13mm ID pellet die was used for niobium cold working. The stage applies force to the die's pushing rod and core die.	84
4.8	A schematic showing the interior of the die. The Nb samples were centered in the inner faces of the interior steel pushing rod (top) and core die (bottom).	84

4.9	A circular niobium sample, photographed through the lens of an optical microscope, is mounted on a metallic block. The area polished with argon ions is visible in the upper half of the niobium - it is a dish-shaped area where the argon ions have begun to erode through the sample. It was very difficult to get a mirror polish surface without also inducing curvature, so the idea of adapting the cross-section polisher for small surface area polishing was abandoned.	85
4.10	Selected metallographs of the single crystal niobium polish after MasterPrep polishing, with bright field images denoted by BF and dark field images denoted by DF. The lefthand images of the control sample show a mostly smooth surface that nevertheless has uniform, if shallow, roughness (especially visible in the dark field view). The sample on the right has been cold-worked, and striations from the sectioning procedure are still visible. The dark field image shows a more noticeable degree of surface roughness (as well as some residue from the Masterprep®). The bottom edge of the metallograph covers an ≈ 3 mm distance across each sample. These images show that even with extensive polishing using the colloidal silica and the lapping fixture, the samples are not perfectly smooth, and it will be hard to obtain a clear TGS signal on many TGS setups, hence the diamond paper recommendation.	86
4.11	Constructive interference occurs when the Bragg condition is satisfied. By maximizing the intensity of a diffraction spot or pattern created with monochromatic X-rays, various information about a sample's crystal structure can be determined.	88
4.12	The arrangement of the X-ray source, the sample, and the X-ray detector in a typical XRD system used in this research. When the diffraction spot intensity is maximized, θ satisfies the Bragg condition, and characteristics of the crystal structure can be determined.	88
4.13	A schematic of a typical goniometer. The sample can be moved up (y), side-to-side (x), back-and-forth (z), rotated about z (ϕ), rocked along Ψ , and rotated relative to the detector (ω). Image courtesy of MIT CMSE.	89
4.14	The Bruker® system at MIT. The X-ray detector is on the left. The collimated X-ray beam is on the right. In the middle, the sample is mounted to a cradle that can be moved in six degrees of freedom. The detector and the cradle can be moved relative to each other and relative to the X-ray beam.	89
4.15	A frame from a coupled scan of the control sample. The diffraction spot is nearly centered ($\chi \approx -90^\circ$) but is somewhat stretched out, indicating some distortion or imperfection in the lattice.	90
4.16	A frame from a psi scan of the control sample, taken later in the orientation procedure. The diffraction spot has become more defined as alignment improves, but it still shows a small amount of distortion.	90
4.17	An integration of the intensity across diffraction spot in Figure 4.16.	91
4.18	A frame from a coupled scan of the 1000 lb sample. The diffraction spot is off-center ($\chi \approx -90^\circ$) and stretched, indicating some distortion or imperfection in the lattice and a significant amount of miscut.	91
4.19	A frame from a psi scan of the 1000 lb sample. The diffraction spot is more centered in this frame, but is faint and distorted.	91
4.20	A frame from one of the ϕ scans performed on the 1000 lb sample. The diffraction spot did not come into the center of the frame during these scans.	92
4.21	A frame from a coupled scan of the 1500 lb sample. The diffraction spot appears close to the center of the frame ($\chi \approx -90^\circ$) and is stretched more obviously than the previous two samples, indicating more distortion or imperfection in the lattice but not much miscut.	93

4.22	Integrating a frame from the coupled scan shows two distinct regions of intensity on either side of $\chi=90^\circ$.	93
4.23	A frame from a coupled scan of the 2000 lb sample. The diffraction spot appears close to the center of the frame ($\chi \approx -90^\circ$) and exhibits a level of distortion similar to that of the 1500 lb sample.	94
4.24	Integrating a frame from the coupled scan shows an asymmetric intensity peak: it is possible that this asymmetry is due to non-isotropic distortion of the sample. χ is about 4° off from -90° .	94
4.25	A frame from a coupled scan of the 2400 lb sample. The diffraction spot appears close to the center of the frame ($\chi \approx -90^\circ$) but note that the 2θ value is not optimized for the $\langle 011 \rangle$ orientation.	94
4.26	Integrating a frame from the coupled scan confirms that the sample is well-aligned to the reference orientation.	94
4.27	A frame from a coupled scan of the 3000 lb sample. The diffraction “spot” stretches along the arc defined by a single 2θ value, with varying intensity. Integration of intensity as a function of χ shows two main intensity peaks. One is about 30° off from $\chi = -90^\circ$, while the other is about 20° off.	95
4.28	Control ϕ scan	96
4.29	1000 lb ϕ scan	96
4.30	1500 lb ϕ scan	96
4.31	2000 lb ϕ scan	96
4.32	2500 lb ϕ scan	96
4.33	3000 lb ϕ scan	96
4.34	Pole figures for the control sample	99
4.35	Pole figures for the 1500 lb sample	99
4.36	Pole figures for the 2000 lb sample	99
4.37	Pole figures for the 1000 lb sample	100
4.38	Pole figures for the 2500 lb sample	100
4.39	Pole figures for the 3000 lb sample	100
4.40	HRXRD rocking curves for the control sample, 1500 lb sample, and 2000 lb sample. There is a significant difference between the control sample and the other two samples. Based on these curves, 1500 lb and 2000 lb likely have a very similar dislocation density.	101
4.41	The 1000 lb, 2000 lb, and 3000 lb samples are plotted separately, as they have a different orientation than the other samples. Note that the 1000 lb sample - the most intense peak here - is only about 20% the height of the control sample rocking curve.	102
4.42	Dislocation density in the cold-worked niobium samples, as estimated from Equations (4.5) and (4.4). A steady buildup of dislocations is observed until 1500 lbs of cold work, at which point a saturation seems to be reached. At 2500 lbs of cold work, dislocation density increases again.	103
4.43	The schematic shows the niobium sample mounted on the rotating stage used to make the measurements. For clarity, the sample depicted here is much larger relative to the stage than it is in practice.	105
4.44	All individual TGS measurements taken on the cold-worked Nb samples. Each measurement consisted of three traces, which are averaged here. This data is plotted prior to shifting data sets so that maxima align.	106

4.45	TGS data from the cold-worked niobium samples. This is the same data plotted in Figure 4.44, but the measurements at each position have been averaged and the data sets have been shifted relative to each other to align maxima. Error bars have also been added.	107
4.46	Calibrated grating spacing for each day of niobium measurements, showing that the grating spacing varied by as much as $0.04\mu\text{m}$ from day to day. The collection of calibration data for every day of measurement is therefore an important step to ensure that analysis of the TGS data reflects the physical properties of the samples as accurately as possible. The nominal grating spacing for all measurements was $5.5\mu\text{m}$	108
4.47	v_{SAW} results for the first group of cold-worked niobium samples, graphed alone for easier comparison. It is possible that the 2000 lb sample has undergone a small decrease in Young's modulus in the fast direction, but overall, the behavior of the samples with respect to SAW speed as a function of surface rotation is remarkably consistent despite the differences in cold work between them.	108
4.48	v_{SAW} results for the second group of cold-worked niobium samples. These samples show a more significant variation in behavior than do the group 1 samples.	109
4.49	The data for the cold-worked samples, separated by group and plotted as the ratio of Δv_{SAW} to v_{min} , where $\Delta v_{SAW} = v(\theta) - v_{min}$, and v_{min} is the minimum v_{SAW} observed for the control sample.	110
4.50	Rotational translations used in the numerical TGS predictions. For the purposes of this project, first two translations correspond to the offset of a sample's true alignment from an arbitrary alignment.	111
4.51	The black line shows experimental data for SAW speed versus degree of surface rotation for the control sample. The blue and red lines are the predicted SAW response for 0° and 15° of tilt. These results are given for rotations of 0 to 90° in 10° increments.	112
4.52	The black line shows experimental data of the SAW speed versus degree of surface rotation for the control sample. The colored lines are the predicted SAW speeds as a function of surface rotation for an $\langle 011 \rangle$ oriented single crystal, with a $0d$ rotational offset and varying amounts of tilt offset. The arrow shows the direction of increasing tilt.	114
4.53	The slowness surface for $\{110\}$ single crystal Nb shows that PSAW and SAW behavior will be observed.	114
4.54	At certain surface rotations, the SAW response for the 0° rotation condition for $\langle 011 \rangle$ single crystal Nb is small compared to the PSAW response. However, the algorithm "picks" the response by selecting the first response from the right, and so it only "sees" the small SAW response. This explains the dips observed in the calculated responses graphed in Figure 4.52. The experimental setup still captures the PSAW response, which is why the experimental data looks like a smooth sinusoid-esque curve.	114
4.55	The black line shows experimental data of the SAW speed versus degree of surface rotation for the 1000 lb sample. The blue and red lines are the predicted SAW response for 0° and 15° of tilt. These are given for rotations of 0 to 95° in 5° increments. Other tilt conditions have been omitted.	115
4.56	v_{SAW} was measured as a function of surface rotation of the control sample following irradiation with Si^{3+} ions to 0.01, 0.03, 0.1, and 1 dpa. Data sets have been shifted horizontally to align maxima.	119
4.57	v_{SAW} was measured as a function of surface rotation of the control sample following irradiation with Si^{3+} ions to 0.01, 0.03, 0.1, and 1 dpa. Data sets have been shifted horizontally to align maxima.	120

4.58	v_{SAW} was measured as a function of surface rotation of the control sample following irradiation with Si^{3+} ions to 0.01, 0.03, 0.1, and 1 dpa. Data sets have been shifted horizontally to align maxima.	121
4.59	v_{SAW} was measured as a function of surface rotation of the control sample following irradiation with Si^{3+} ions to 0.01, 0.03, 0.1, and 1 dpa over 180-200°. This graph shows the results of averaging all speed measurements taken on a given sample after each irradiation. Error bars are large in part due to the normal spread of v_{SAW} on the anisotropic single crystal. . . .	122
4.60	v_{SAW} in the fast direction (although some of the data points in the average are likely PSAW responses, as explained on page 113. All v_{SAW} measurements $\pm 5^\circ$ from the rotational position of the measured maximum v_{SAW} are averaged.	122
4.61	Thermal diffusivity results, control sample. Error bars are not shown for clarity. The associated error for each thermal diffusivity measurement was on the order of $10^{-7} m^2/s$	123
4.62	Acoustic damping results, control sample	124
4.63	Thermal diffusivity results for all samples after cold working (0 dpa) followed by irradiation to 0.01, 0.03, 0.1, 1, and 3 dpa. Points represent the average of all measurements. Error was an order of magnitude lower than the average and is not plotted for clarity, since the primary data points overlap in places.	125
4.64	The data in Figure 4.63 is replotted with dose on a log scale. The thermal diffusivity is observed to drop at low doses, rise again, and saturate near $2E-5 m^2/s$. The 0 dpa data are assigned a pseudo-dose of 0.001 dpa here to keep them on the graph. However, the response between 0 dpa and 0.01 dpa is unknown, so it may not be accurate to consider the 0.001 dpa response as a reasonable approximation for the known 0 dpa response.	125
4.65	Acoustic damping results for all samples after cold working (0 dpa) followed by irradiation to 0.01, 0.03, 0.1, 1, and 3 dpa. Points represent the average of all measurements. The procedure for calculating acoustic damping is still being refined, and these measurements should be considered as estimates only. Error is on the order of 10^{-9} , including for the (1 dpa, 1500 lb) point.	126
4.66	The data in Figure 4.65 is replotted with dose on a log scale. As before, the 0 dpa data are assigned a pseudo-dose of 0.001 dpa. The data plotted between 0.01 and 0.1 dpa indicate evidence of an exponential relationship between acoustic damping and dose.	126
4.67	SAW speed in the fast direction at 0 dpa for the control, 1500 lb, and 2000 lb single crystal niobium samples.	129
4.68	Oxygen impurities in single crystal niobium are mobile during neutron irradiation, and are trapped at radiation-induced defect clusters (dark spot). [147]	131
4.69	Dislocation density of niobium samples irradiated with 5 MeV protons rises sharply at low dose (0.01 dpa) as radiation damage builds up in the bulk. As dose increases, these defects concentrate into larger defects that are less dispersed through the bulk. These results were confirmed via TEM. Adapted from [148]	131
5.1	Thermal vacancy concentration in aluminum and niobium as a function of temperature. . . .	135
5.2	The NiAl phase diagram	136
5.3	The phase diagram above, annotated	137
5.4	The structure of B2-phase NiAl consists of a simple cubic lattice of Ni atoms interlocked with a simple cubic lattice of Al atoms. (image: http://minerva.mlib.cnr.it/)	138

5.5	The lattice parameter of B2-phase NiAl varies with composition. The lattice parameter reaches a maximum near the stoichiometric point. It decreases as the nickel content goes down relative to the stoichiometric point and the number of nickel vacancies goes up. It decreases at a more gradual slope as the nickel content increases relative to the stoichiometric point and the number of nickel anti-site defects on the aluminum sublattice goes up. Citations for the data plotted are as follows: “Cooper1963”, [164]; “Jacobi1971”, [165]; “Hughes1971,” [166], and “Kogachi1996”, [167].	139
5.6	Reported concentrations of structural vacancies present in NiAl as a function of composition from select literature. Note that the “Pike1997” and “Kogachi1996” report C_V^{Ni} values of ≈ 0 for at% Ni $\geq 50\%$, while the 2001 update to Kogachi1996 reports a C_V^{Ni} at 52 at% Ni that is $\approx C_V^{Ni}(50\% \text{ Ni}) \approx 9\%$. The updated Kogachi data is also significantly higher than the other data sets, with reported C_V^{Ni} of over 12% at lower values of nickel concentration. Interestingly, Kogachi also reports nonzero C_V^{Al} , although it never exceeds 2%. For less than 50% Ni, C_V^{Ni} decreases monotonically with increasing Ni concentration. Citations for the data plotted are as follows: “Kogachi1996”, [167]; “Kogachi1996 updated”, [172]; and “Pike1997,” [173].	140
5.7	The plot of magnetic susceptibility as a function of composition illustrates that NiAl has very different material properties on either side of the stoichiometric composition. Adapted from [174].	141
5.8	The plot of magnetic resistivity as a function of composition illustrates that NiAl has very different material properties on either side of the stoichiometric composition. Resistivity increases with increasing defect concentration: above the stoichiometric point, resistivity increases as the concentration of nickel vacancies increases, and below the stoichiometric point, resistivity increases as the concentration of nickel anti-site defects increases. [174]	141
5.9	Dislocations in an NiAl sample used in [159]	144
5.10	The NiAl phase diagram	145
5.11	Ni/Al powder mixture is transferred to the pellet press die using a funnel and a disposable scoop	146
5.12	A vacuum pump (connected to the tube) is used to remove air while samples are pressed.	146
5.13	The finished pellets are about 0.5” wide and 0.25” high. The pellets are much easier to handle than the loose powders, especially during the arc melting step.	146
5.14	The MIT MNM arc melter in use. The user controls the position of the arc and the power supplied to the electrode. Welders’ glass windows allow the user to see when the pellets have fully melted and formed a single button of molten material.	147
5.15	The five buttons made in the arc melter were heat treated in a high-temperature graphite furnace under an inert environment for 24h at 1200°C. The photograph shows the buttons arranged in the furnace prior to beginning the heat treatment. The purpose of the heat treatment was to grow large grains, so that the TGS measurements could be easily carried out within the boundaries of a single grain if the buttons were used for measurements. The furnace pictured here is in the Allinore lab at MIT.	147
5.16	NiAl sample #1 at 50x magnification in a polarized field. The etchant swab created shallow “scratches” that make it difficult to see the sample’s features.	148
5.17	NiAl sample #2 at 50x magnification in a polarized field. As with #1, the etching process obscured many of the features on this sample. This is exacerbated by the small voids throughout sample that seem to have formed during melting.	148

5.18	NiAl sample #3 at 50x magnification in a bright field. It is easier to see the large grains, some of which are several hundred μm in width and thus capable of fitting the entire TGS spot.	148
5.19	A second image of NiAl sample #3 at 50x magnification in a bright field, showing the large, clearly delineated grains.	148
5.20	NiAl sample #4 at 50x magnification in a bright field. As was the case with Sample #2, small voids developed during the arc melting process.	148
5.21	NiAl sample #5 at 50x magnification in a bright field. Voids appear to be distributed through this sample as well, although they appear smaller than the voids observed in #1 and #4. . . .	148
5.22	Degrees of freedom in the sample mount for the Bruker® D8 GADDS system, taken from the MIT CMSE SOP for the equipment.	150
5.23	A photograph of an NiAl sample in the XRD GADDS system. The wide-angle germanium detector is on the left and the X-ray source is on the right. The sample's cradle moves in $x, y, z, \phi, \psi,$ and ω . The entire cradle and the detector move independently to change the angle between the detector and the X-ray source (2θ) and the angle between the detector and the sample.	150
5.24	The schematic shows the mounted sample and the corrections required to make the $\langle 100 \rangle$ direction normal to the measurement surface. ω involves rotating the sample from side to side, whereas ψ indicates a back-and-forth tilt.	151
5.25	A not-to-scale schematic of the process of cutting the sample mount to correct for misalignment of the crystal plane relative to the TGS measurement surface. Sample planes (shown as light blue lines) have a misalignment of ψ . By making a cut at an angle of ψ from the back of the mount, we bring the crystal planes into their proper alignment when the back of the mount is level. A second cut is then made because the measurement surface needs to be flat. The new surface is repolished prior to measurement.	152
5.26	1A, bright field (41.18 Ni, 58.73 Al, 0.04 Fe, 0.05 Si)	157
5.27	1A, dark field (different region). The large voids are the most distinctive features here. . . .	157
5.28	1A, polarized (same region as dark field)	157
5.29	1B (41.98 Ni, 57.95 Al, 0.03 Fe, 0.04 Si) polarized. Grains are large and very distinct under polarized light for the samples with < 42 at% Ni.	157
5.30	#2, bright field (44.11 Ni, 55.76 Al, 0.08 Fe, 0.04 Si). Large population of mid-sized ($10\text{-}20\mu\text{m}$) voids.	157
5.31	#2, dark field (different region)	157
5.32	#2, polarized (same as dark field). Grain boundaries are more evident, but look distinctly different from those observed in Nos. 1A and 1B.	157
5.33	#3, bright field (44.51 Ni, 55.37 Al, 0.06 Fe, 0.05 Si)	158
5.34	#3, dark field (different region)	158
5.35	3, polarized (same region as dark field)	158
5.36	#6, bright field (44.58 Ni, 54.32 Al, 0.08 Fe, 0.03 Si)	158
5.37	#6, dark field (different region)	158
5.38	#6, polarized (different region)	158
5.39	#4, bright field (47.38 Ni, 52.53 Al, 0.05 Fe, 0.04 Si (at%))	158
5.40	#4, dark field (different region)	158
5.41	#4, polarized (different region)	158
5.42	#5, bright field (47.73 Ni, 52.16 Al, 0.07 Fe, 0.04 Si (at%))	159
5.43	#5, dark field (different region)	159

5.44	#5, polarized (different region)	159
5.45	#7, bright field (49.01 Ni, 50.90 Al, 0.06 Fe, 0.04 Si)	159
5.46	#7, dark field (different region)	159
5.47	#7, polarized (different region)	159
5.48	#8, bright field (49.97 Ni, 49.94 Al, 0.06 Fe, 0.04 Si)	159
5.49	#8, dark field (different region)	159
5.50	#8, polarized (different region)	159
5.51	#9, bright field (57.47 Ni, 42.43 Al, 0.07 Fe, 0.02 Si)	160
5.52	#9, dark field (different region)	160
5.53	#9, polarized (different region)	160
5.54	#10, bright field (59.63 Ni, 40.31 Al, 0.04 Fe, 0.02 Si (at%))	160
5.55	#10, polarized (different region)	160
5.56	SAW speed measurements on NiAl of varying compositions taken at a nominal $5.5\mu\text{m}$ grating spacing. These samples were not corrected for possible miscut or misalignment of the crystal planes. Data was shifted in degree so that the first maximum in each data set was matched.	164
5.57	An attempt to match the unknown, crystallographically corrected “blue” sample to a sample of known composition was made. The best match was Luvak 4. If Luvak4 and “Blue” were from the same sample, then this indicates that the angling step was probably unnecessary.	165
5.58	An attempt to match the unknown, crystallographically corrected “pink” sample to a sample of known composition was made. The best match was Luvak5. As with Figure 5.27, if we can assume that the unknown sample comes from the same single crystal rod as the identified sample whose response it most closely matches, then it would appear that the angling steps had little effect on the frequency (and therefore the speed) response.	165
5.59	SAW speed vs. at% NiAl for Batch I and Batch II samples. Results show a general decrease in speed on the Al-rich side of the compositional range as the stoichiometric point is approached (direction of decreasing vacancy concentration).	166
5.60	Linear fit to the v_{SAW} data for 44 to 50 at% Ni	167
5.61	Young’s modulus of NiAl versus vacancy concentration in NiAl, calculated from Equation (5.10).	169
5.62	Microhardness increases with vacancy concentration in intermetallic FeAl. [181]	170
5.63	Microhardness increases with vacancy concentration in intermetallic FeAl. [181]	170
5.64	Thermal diffusivity as a function of at% nickel in NiAl polycrystalline and single crystal samples.	171
5.65	Linear fit to the thermal diffusivity data for samples with composition 44-50 at% Ni.	173
5.66	Acoustic damping as a function of at% nickel in NiAl polycrystalline and single crystal samples.	174
5.67	The TGS simulation begins with a structure with a free surface.	176
5.68	Two regions on the free surface are specified for center-of-mass tracking.	176
5.69	A sinusoidal heat pulse is applied to the free surface.	176
5.70	Initial relaxation of nickel and aluminum test structures in LAMMPS was unsuccessful. The test structures were random mixes of Ni and Al, but did not replicate the actual structures of B2-phase NiAl. Instead, lattice parameter simply grew with the increasing nickel content, as nickel is the larger atom.	177
5.71	Results of convergence tests showed that the test structure should have at least 2 million atoms.	179

5.72	A stoichiometric NiAl structure constructed in MATLAB, and a Ni-poor NiAl structure with a significant concentration of constitutional vacancies constructed using the same code, are pictured. MATLAB is used to create input text files that rebuild these structures in LAMMPS.	179
5.73	Lattice parameter results for B2-NiAl obtained in LAMMPS compared against lattice parameters reported in the literature.	181
5.74	SAW speed measured from TGS simulations carried out on NiAl test structures of 2×10^6 lattice spots in LAMMPS.	182
5.75	SAW speed measured from TGS simulations carried out on NiAl test structures of 2×10^6 lattice spots in LAMMPS.	183
5.76	Measurements on the Batch I samples taken over the full surface rotation (with at least two spots per rotation position) during a previous data collection campaign are included with the data sets to provide a better picture of the single crystal behavior. Error values for these points were on the order of 1% of the averaged measurement value.	185
6.1	A radioactive hex block in a hot cell after removal from the EBR-II reflector. The block is 5.2 cm wide and about 20cm long. [189]	189
6.2	Radiation damage response is not uniform throughout each block due to neutron flux and temperature gradients that existed throughout the block. The temperature peaked near the middle, but the peak was shifted toward the side that was nearest to the reactor core. Gamma heating and dpa rate increased monotonically from the side furthest from the reactor core to the side closest. Schematic from [189]	189
6.3	The image, modified from [190], shows where the tested samples came from in the EBR core. The blocks (5 and 3) are marked with stars; the sections of those blocks relevant to this study are marked with yellow arrows. The blocks were stacked in a hexagonal duct with 1 mm thick 304 stainless steel walls. The duct was located in Row 8 of the core's reflector region [189].	190
6.4	The image, from [190], shows where Sample 3D1 was located in Coin 3D, as well as its dimensions. D signifies that the sample was immersion density measured.	191
6.5	Sample 3D1 was divided into five sections. The larger sections in the middle three (3DB1, 3D1C, and 3D1D) were each further divided into four subsections. One subsection from each of 3D1B, 3D1C, and 3D1D were used in this study.	191
6.6	The image, from [190], shows where Sample 3E1 was located in Coin 3E, as well as its dimensions. D signifies that the sample was immersion density measured.	191
6.7	Sample 3E1 was divided into five sections. Section 3E1B was divided into three sections. The middle section, 3E1B1, was divided into four more sections, two of which were thin slices intended for use in making TEM samples. One of these thin slices, 3E1B1A, is used here. [190].	191
6.8	The image, from [190], shows where Sample 5A1 was located in hex coin 5A. Sample 5A1 was divided into five sections, one of which, 5A1B, is relevant to this work.	192
6.9	Section 5A1B was divided into three more sections. The middle piece, 5A1B1, was divided into four more sections, two of which were thin slices intended for use in making TEM samples. One of these thin slices, 5A1B1B, is used in this work. Both this figure and Figure 6.8 are from [190].	192

6.10	5B1 of hex coin 5B was divided into four sections. The two middle sections, 5B1B and 5B1C, were each divided into four more sections. As in previous coins relevant to this work, two of these sections were thin slices. One thin slice from each of 5B1B and 5B1C were used in this work: 5B1B2B and 5B1C2B. [190].	192
6.11	5C1 of hex coin 5C was divided into three sections. The middle section, 5C1B, was divided into four more sections. As in previous coins relevant to this work, two of these sections were thin slices. One of the thin slices, 5C1B2B, was used in this work. [190].	192
6.12	Micrographs of the archival stainless steel hex block, which serves as the stock for the control sample used in the TGS project. The archival block indicates that the blocks were cold worked prior to placement in EBR-II. [189]	195
6.13	Flat-to-flat swelling for Blocks 2, 3, and 4; a schematic for visualizing the direction of the flat-to-flat measurements; and lengthwise swelling in Blocks 3 and 5. Block 3, in the center of the stack, exhibited the most dramatic swelling. Block 5 was at the top of the stack. Lengthwise measurements show that Block 3 exhibited significant expansion, whereas Block 5 did not. INL and WEC indicate the hot cell in which the measurements were made - Idaho National Laboratory or Westinghouse Electric Company. Negative swelling values in Block 2 and Block 4 indicate carbon densification has taken place. Flat-to-flat swelling data for Block 5 is presented in Figure 6.14. [189]	196
6.14	Post-irradiation flat-to-flat dimensions for Block 3 and Block 5. The solid line at 2.062" shows the nominal dimension of the block prior to irradiation. Block 3, which received the higher dose, clearly exhibits swelling behavior that clearly exceeds the nominal+tolerance flat-to-flat length. The amount of swelling from flat-to-flat varies with axial position. Block 5 exhibits negative swelling, which is associated with carbon densification. Note that left-to-right indicates bottom-to-top. [189]	197
6.15	Acoustic velocity mapping (lengthwise) for Block 3. Acoustic mapping for coin 3E from this angle and of Block 5 was not provided. Section 3D (the side of hex coin 3D) exhibits a variation of acoustic speed of about 20 m/s. The archival (unirradiated) acoustic velocity is 5735 m/s. [189]	198
6.16	An acoustic map of hex coin 3D shows that the lowest speeds are clustered slightly off-center - the same spot where the highest concentration of voids is expected to be and where radiation damage was highest. The coins from block 3 exhibit a large acoustic velocity gradient, with variations of up to 70 m/s. [189]	199
6.17	An acoustic maps of hex coins shows velocity increases relative to the nominal acoustic speed of the archival material (5735 m/s). The changes are less severe than those observed in block 3, with acoustic speed variations of about 30 m/s. [189]	199
6.18	An acoustic map of hex coin 3E, overlaid with density results. Note that the image is flipped relative to the image in Figure 6.16, with the flats on opposite sides. The low-velocity region occurs in the same place on each coin. Coin 3E was sectioned according to the dashed lines, and density measurements were made on each segment. The extent of the density change corresponded with the extent of the acoustic velocity reduction. [189]	200
6.19	TEM images of a sample from coin 3E show many voids dispersed throughout the material.Hex block 3 was in the center of the stack and received the highest radiation dose. [191]	200
6.20	Hex block 5 received a much lower dose than hex block 3. A few voids are visible in this coin 5D sample, but the difference in void density visible in this image and in the images of coin 3E is obvious. [191]	201

6.21	The table from [193] that quantitatively characterized voids, precipitates, and Frank-Reed loops, based on TEM of the EBR-II samples. Stars are added here to show which rows correspond to a sample that was immediately adjacent to the samples used in this study prior to the samples being cut from the hex blocks. It is assumed that the defect populations of the samples in this study are effectively equivalent to the defect populations of their corresponding samples in [193].	202
6.22	SAW speed measurements, averaged by sample and plotted by the dpa to which each sample was irradiated in EBR-II.	204
6.23	SAW speed measurements, with all data grouped by hex block. Hex block 5 (second point) is plotted at 2.625 dpa, which is the average of the dpa values reported for each sample measured here that originated in hex block 5. Hex block 3 (rightmost point) is plotted at 28 dpa, which was the reported dpa value for every sample measured here that originated in hex block 3.	205
6.24	Thermal diffusivity measurements, averaged by sample and plotted by the dpa to which each sample was irradiated in EBR-II.	205
6.25	Void density as measured in [193] versus TGS-measured SAW speed	206
6.26	Void density as measured in [193] versus TGS-measured thermal diffusivity	206
6.27	Void swelling as measured in [193] versus TGS-measured SAW speed	206
6.28	Void swelling as measured in [193] versus TGS-measured thermal diffusivity	206
6.29	Precipitate density as measured in [193] versus TGS-measured SAW speed	207
6.30	Precipitate density as measured in [193] versus TGS-measured thermal diffusivity	207
6.31	Frank loop density as measured in [193] versus TGS-measured SAW speed	207
6.32	Frank loop density as measured in [193] versus TGS-measured thermal diffusivity	207
6.33	Frank loop length density as measured in [193] versus TGS-measured SAW speed	207
6.34	Frank loop length density as measured in [193] versus TGS-measured thermal diffusivity	207
6.35	Neutron irradiated stainless steels undergo an increase in yield strength σ_y at doses that correspond to the doses incurred by the EBR-II samples, which corresponds to an increase in E . The irradiation temperatures to which the EBR-II samples were exposed was in the 400-420°C range, well below the range at which temperature would be expected to be associated with a decrease in σ_y . [197]	209
6.36	304 stainless steel stress-strain test results, prior to irradiation. [196]	209
6.37	304 stainless steel stress-strain test results, following irradiation to 30 dpa in the High Flux Reactor in Petten, the Netherlands. [196]	209
6.38	APT results for 304 steel irradiated to 5 dpa with protons shows evidence of radiation-induced segregation. A grain boundary in the APT sample is indicated for the Si scan with a black arrow. These results show slight depletion at the grain boundary for Fe, Cr, Mn, and Cu, and enrichment at the grain boundary for Si, P, B, and S. Enrichment is particularly dramatic for the Si, P, and B species. [203]	210
6.39	APT data of 304 stainless steel samples irradiated to 0.4 dpa (a) and 28 dpa (b) in EBR-II shows evidence of precipitation of Ni and Si from the Fe matrix. [209]	211
6.40	APT data of a 304 stainless sample irradiated to 0.4 dpa in EBR-II shows evidence of precipitation of P from the Fe matrix. [209]	212
6.41	APT data of 304 stainless steel samples irradiated to 28 dpa in EBR-II shows evidence of precipitation of P from the Fe matrix. [209]	212
6.42	Thermal diffusivity measured by TGS versus void density measured via TEM analysis in [193]	213
6.43	Average void diameter versus void density, both measured via TEM analysis in [193]	213

6.44	Dpa associated with samples versus their void density measured via TEM analysis in [193]	213
7.1	TGS enables more efficient radiation materials science research campaigns. Schematic by M. P. Short.	217
8.1	Orientation distribution function obtained from the pole figure data of the control Nb sample	220
8.2	Orientation distribution function obtained from the pole figure data of the 1000 lb Nb sample	221
8.3	Orientation distribution function obtained from the pole figure data of the 1500 lb Nb sample	222
8.4	Orientation distribution function obtained from the pole figure data of the 2000 lb Nb sample	223
8.5	Orientation distribution function obtained from the pole figure data of the 2500 lb Nb sample	224
8.6	Orientation distribution function obtained from the pole figure data of the 3000 lb Nb sample	225
8.7	The black line shows experimental data for SAW speed versus degree of surface rotation for the 1500 lb sample, which is expected to be well-aligned to the $\langle 110 \rangle$ plane. The blue and red lines are the predicted SAW response for 0° and 15° of tilt. These are given for rotations of 0 to 95° in 5° increments. Other tilt conditions have been omitted.	232
8.8	Experimental data and predicted SAW speed calculations as a function of surface angle for the 2000 lb sample	233
8.9	Experimental data and predicted SAW speed calculations as a function of surface angle for the 2500 lb sample	233
8.10	Experimental data and predicted SAW speed calculations as a function of surface angle for the 3000 lb sample	234
8.11	Thermal diffusivity results, 1500 lb sample	235
8.12	Acoustic damping results, 1500 lb sample	236
8.13	Thermal diffusivity results, 2000 lb sample	236
8.14	Acoustic damping results, 2000 lb sample	237
8.15	Elastic constants of NiAl as reported in [161].	239
8.16	Elastic constants of intermetallic NiAl of various concentrations as a function of temperature, from [161].	240
8.17	NiAl tends to deform by glide in the $\langle 100 \rangle$ family of directions in the $\{011\}$ and $\{001\}$ plane families. Above, the $[100]$ direction in the (011) and the (001) planes of a B2-structured system is illustrated.	241
8.18	On the left, Young's modulus is graphed as a function of temperature. The top three lines (45-50 at% Ni) correspond to NiAl compositions that we consider in our work. The modulus decreases with temperature. These values are close to the values computed from single crystal data. [175]	242
8.19	Example LAMMPS output code generated by the MATLAB script for a relatively small structure of about 2000 atoms showing the header script and the x-y-z specifications for the first seven atoms.	245

List of Tables

3.1	Advantages and disadvantages of various destructive radiation damage characterization methods, ✖ = no, ✔ = yes, ≈ = possible	67
3.2	TGS setup parameters used in this thesis	75
4.1	Crystallography peak list for pure niobium	90
4.2	Summary of XRD characterization of cold-worked niobium sample orientations	97
4.3	HRXRD rocking curves of the cold-worked single crystal Nb samples	102
4.4	Irradiation parameters for the cold-worked irradiated niobium study: integrated charge in 10^{-11}C to reach each dpa level (given in italics), and actual charge achieved during each irradiation step for each sample	118
5.1	Preparing NiAl specimens: some examples from six decades of literature	143
5.2	Composition plan for sample matrix, NiAl Batch I	145
5.3	Growing single crystal samples at LANL from the arc-melted and annealed NiAl buttons	149
5.4	Sample tilts as measured by XRD. Figure 5.24 illustrates ω and ψ relative to the mounted sample (oriented in ϕ). Each sample had a color label for easy identification.	151
5.5	Composition of NiAl Batch I	153
5.6	Planned compositions for NiAl Batch II	155
5.7	Composition of NiAl Batch II	155
6.1	EBR-II sample history and description, as provided by Westinghouse	194
6.2	Dose and temperature range for select hex coins, as reported in [191]	194
6.3	TGS-measured SAW speed and thermal diffusivity for the EBR-II samples	204
8.1	Elastic constants decrease with increasing vacancy concentration, as reported in [161]	239

Nomenclature

α	Alpha particle
$\bar{\nu}_e$	Antineutrino
β	β particle
β -	Indicates B2-phase intermetallic structure
β^+	Positron that has been emitted in β -decay
β^-	Electron that has been emitted in β -decay
β_Ω	Integral breadth of peak in HRXRD
χ	Tilt (front to back) of sample [XRD applications]
γ	Gamma ray
λ	Wavelength
\AA	Angstrom (10^{-10}m)
μ	Micro-; $\times 10^{-6}$
μ	Shear modulus
ν	Frequency
ν_e	Neutrino
ω	Degree of tilt (left to right) [XRD applications]
Φ	Radiation flux [bombardment applications]
ϕ	Degree of surface rotation [XRD applications]
ϕ	Radiation flux [bombardment applications]
σ	Standard deviation [error calculation applications]
σ	Tensional stress
σ_D	Displacement cross section
σ_y	Yield strength
θ	Bragg's angle [XRD applications]
θ	Surface rotation angle [TGS applications]
ε	Extensional strain
$^\circ$	degree
i	Indicates interstitial
v	Indicates vacancy
A	Atomic weight
b	Burgers vector
C_v	Vacancy concentration
C_{ijkl}	Elastic tensor entry
d	Average obstacle size [dislocation movement]
d	Lattice spacing [X-ray diffraction applications]

E	Energy
E	Young's modulus
E_d	Displacement energy
E_i	Initial or impinging energy of radiation
H_v	Enthalpy of vacancy formation
N	Number density of a material (number of atoms per unit volume)
N	Total number of measurements [error calculation applications]
N_d	Number of displacements per atom
R	Gas constant, $8.3144598 \text{ kg m}^2 \text{ s}^{-2} \text{ K}^{-1} \text{ mol}^{-1}$
R	Range of particles in a target [bombardment applications]
S	Stopping power
T_d	Damage energy; energy available for displacements
v_{SAW}	Speed of surface acoustic wave, typically reported in [m/s]
x	1D distance into target
Z	Atomic number
'	Feet
''	Inches
<x y z>	indicates a family of crystallographic directions
C	Coulomb
C	Elastic tensor
A	Ampère
APT	Atom probe tomography
ASD	Anti-site defect
at%	Atomic percentage
BF	Bright field [optical microscopy]
C	Degrees Celsius
c	Speed of light (2.99792458 m/s)
CANDU	Canada Deuterium Uranium reactor
CMSE	Center for Materials Science and Engineering at MIT
DF	Dark field [optical microscopy]
dpa	Displacements per atom; used as a pseudo-unit to describe radiation damage in material
dpa	Displacements per atom
e	Elementary charge ($1.602 \times 10^{-19} \text{ C}$)
EBR	Experimental Breeder Reactor
EDX	Energy-dispersive X-ray spectroscopy
eV	Electronvolt (work to accelerate an electron through a one volt potential difference)
f	Frequency
FIB	Focused ion beam
FWHM	Full width at half maximum
g	Grams
GADDS	General Area Detector Diffraction System
GFR	Gas-cooled fast reactor
GS	Grating spacing (typically used as a subscript, e.g. λ_{GS})
h	Hours
HOLOSLAM	Holographic scanning laser acoustic microscopy
HRXRD	High-resolution X-ray diffraction

HV	Hardness value
ICP	Inductively-coupled plasma
ID	Inner diameter
K	Degrees Kelvin
lb	Pound
LFR	Lead-cooled fast reactor
LSAW	Laser-induced surface acoustic wave
M	Mega-; $\times 10^6$
m	Meter
MD	Molecular dynamics
MIT	Massachusetts Institute of Technology
MIT MNM	MIT Mesoscale Nuclear Materials group
mol	Moles [1 mol = 6.022×10^{23}]
MSR	Molten salt reactor
PKA	Primary Knock-on Atom
ppm	Parts per million
rpm	Rotations per minute
RUS	Resonant ultrasound spectroscopy
s	Seconds
SAM	Scanning acoustic microscopy
SAW	Surface acoustic wave
SBS	Surface Brillouin scattering
SCWR	Supercritical water reactor
SEM	Scanning electron microscope
SFR	Sodium-cooled fast reactor
SLAM	Scanning laser acoustic microscopy
SRIM	Stopping Range of Ions in Matter computer program
T	Temperature
TEM	Transmission electron microscopy
TGS	Transient grating spectroscopy
V	Volt
VHTR	Very-high-temperature reactor
W	Weight
wt%	Weight percent
XRD	X-ray diffraction

Chapter 1

Prologue:

Breaking the bottleneck in nuclear materials research

The modern academic field of nuclear power engineering is organized around the principle that nuclear power is a reliable source of carbon-emissions free baseload energy, and that improvements in the safety, efficiency, and cost of existing nuclear technologies - or the implementation of new ones - increases the likelihood that nuclear power will be implemented more broadly across the world. The MIT Mesoscale Nuclear Materials (MIT MNM) laboratory exists at the intersection of nuclear engineering and materials science, and therefore, is concerned with solving problems in nuclear materials that otherwise hinder the above goal.

Nuclear power applications are characterized by harsh irradiation environments that present unique materials challenges. In a standard commercial power reactor - before even considering the impact of radiation - in-core materials must be able to withstand high temperatures and thermal stresses, large mechanical loads, corrosive effects of coolant and moderator fluids, and wear from mechanical vibration and fluid flow. Advanced reactor concepts and fusion technologies tend to utilize higher power outputs, temperature regimes, and more corrosive fluids, which exacerbate these challenges.

Now, to the list above concerns, consider the matter of radiation damage. Nuclear fission and fusion, the same phenomena that allow the reactor to produce power, are also the source of the most significant materials challenges. The core of the reactor is subject to high radiation fluxes. Typically, one thinks of the neutrons in the core of a commercial power reactor: the same neutral particles responsible for the fissioning of uranium can also knock atoms in structural materials out of place. Over time, the buildup of this damage changes the way the material performs, and can limit the component's useful lifetime. The radioactive processes of

a reactor core result in a host of other radiation species as well: structural components in the core must also contend with these alpha, beta, and gamma particles.

Nuclear materials research and development, therefore, is a subject of managing constraints. A material must be proven to retain its integrity - that is, its properties must not degrade beyond a certain acceptable bound - for the lifetime of the particular component being considered, for a projected history of temperature, stress, corrosive contact, and radiation. Furthermore, the material must do all of this without impeding the delicate neutronics balance that makes a reactor work: many otherwise suitable materials, for example, may have neutron absorption cross-sections that are too high to make the material usable in the reactor core.

It is not surprising, then, that materials often represent the major engineering hurdle in moving a new reactor concept closer to reality. The same materials that function adequately in a contemporary commercial reactor cannot be presumed to also work for a reactor that operates at a higher temperature, or which has a more corrosive coolant than water, or which has a different set of neutronics constraints. The difficulties are even more severe in the fusion field, where extremely high temperatures and radiation fluxes, coupled with the difficulties of maintaining a stable plasma, are an imposing roadblock between the academic setting and commercial deployment of fusion technology. Many suitable plasma-facing materials still sustain heavy damage that severely curtails their operating lifetime. The current projected cost of maintaining and replacing damaged components must be brought down if fusion reactors are to ever contribute to the world's energy production.

If nuclear materials research is the bridge between academic concept and commercial reality, it is frequently a long and expensive bridge to cross. In order to validate a new material for use in a specific reactor environment, one must test the material in representative conditions, or test the material in a sufficient number of conditions that the material's response to an arbitrary reactor environment can be accurately predicted.

Since this typically means testing the material in a representative radiation environment, one must contend with the problem of access to the necessary radiation source. Exposing samples in a research reactor core typically requires formal application process, extensive planning, adequate funding, and lots of time (application process, sample exposure, and radioactive cooling). Exposing samples using an accelerator is typically more accessible, but often still very expensive, with costs commonly in the hundreds-of-dollars-per-hour (not just for beamtime, but for setup and retrieval as well). These difficulties are compounded by the fact that it's usually not sufficient to run just one exposure experiment in order to gain a full understanding of the material's radiation response.

Of course, the exposure experiments themselves are simply the first step in the research process. Once the material has been exposed to radiation, it is then necessary to determine how the material has changed. A new set of time and cost obstacles must be surmounted.

One might choose to start with something straightforward and accessible, such as optical microscopy to check for any macro-level surface changes. However, such a technique is only useful for rough, qualitative observations. For samples exposed to lower radiation doses, there may be no obviously visible changes at all.

It is typically necessary to examine the sample via scanning electron microscopy (SEM). Such an instrument typically costs on the order of a hundred dollars an hour to use at a typical research institution. The SEM provides more detail of the sample surface, and the researcher can make more detailed observations about what has appeared to change on the sample.

To make visual observations of non-surface changes, it is necessary to use a focussed-ion beam (FIB) to prepare a sample for transmission electron microscopy (TEM). The TEM allows the researcher to view changes to the material structure on an atomic level. However, both the FIB and the TEM are even more specialized than the SEM, and typically require extensive training to use. They are also painstaking and expensive to operate. Furthermore, the SEM, FIB, and TEM are frequently in-demand instruments at a typical university, which tends to add even more time to a typical research campaign.

The TEM images provide more detailed images, but only over a tiny area of the sample. So, one must take the time to first prepare a sufficient number of FIB samples to be representative of the entire sample. Then, the TEM images themselves must be carefully analyzed - for example, in order to determine dislocation density, one must painstakingly count the visible dislocations in the image.

In order to characterize property changes, a battery of the typical measurement techniques can be introduced: tensile testing, resonant ultrasound spectroscopy, thermal flash methods, and so on. All of this requires time and money, or the selective exclusion of certain measurements.

Certain exposure experiments result in activated samples. Because the samples are now radioactively hot, significant precautions must be taken when analyzing them in order to ensure the safety of the researchers and to avoid possible contamination of workspaces. This means that the availability of the instruments mentioned above is even further curtailed: one must locate an instrument that is designated for handling hot samples and obtain the necessary radioactive sample training (or else hire a technician to perform the analysis), or else adapt and certify an existing instrument for hot use.

Either option incurs additional expense and time. At a typical university, this can be significant, when research budgets are constrained and there is competition for instrument time. In many cases, it might be necessary to go to another institution in order to find the necessary equipment. It is very much the exception for a researcher to have easy and affordable access to all the instruments necessary to perform a thorough analysis of materials properties on all the samples in their particular research campaign, especially if the samples are hot. Even at MIT - widely considered to be one of the best research institutions in the world, and one with a dedicated nuclear engineering department and access to accelerators and even a reactor - we frequently find ourselves leaving campus in order to access specialty instruments.

Advanced fission and fusion concepts will require many materials advances if they are to ever proceed to construction and operation. This means that nuclear materials research must be feasible and efficient at as many institutions as possible. Large-scale testing and validation of proposed materials cannot be the domain of the few labs who happen to have the funding and equipment access necessary to carry out exposure experiments and the subsequent analysis. The need for nuclear materials research far exceeds the available resources to carry it out. If we are to enable advanced nuclear concepts - for example, making a fusion

reactor a commercially viable enterprise - *we have to find a way to break this bottleneck.*

This is where transient grating spectroscopy comes in. Transient grating spectroscopy (TGS), in which lasers are used to induce and probe a surface acoustic wave on a sample surface, can provide information about thermal, elastic, and acoustic material properties quickly and non-destructively. It has been used to study thin films for decades. In 2014, the Mesoscale Nuclear Materials group at MIT began efforts to adapt TGS techniques to the study of solid metals and ceramics, particularly those of interest to the nuclear field.

Why bother at all? After all, there are plenty of ways to measure a material's elastic modulus or its thermal diffusivity. However, TGS appeared to be a promising candidate to fulfill our ideal measurement technique for studying radiation damage and its effects in nuclear materials.

First of all, it is non-destructive. We don't have to cut up the sample to study it. This makes delicate, time-consuming preparation techniques (e.g. TEM specimen preparation) unnecessary. Measurements can be made quickly - on the order of seconds to minutes - and used to provide information about elastic, thermal, and acoustic properties instantly. Because these properties can be correlated to radiation damage, TGS can be used to study radiation damage defect populations in a material *in addition to* the material effects of those defect populations. This makes analyzing radiation damage effects much faster.

Second, TGS is non-contact. Data collection does not require the use of transducers or traditional probes - the only requirement is optical contact between the sample surface and the lasers. Once the TGS facility is in place, the sample needs only be properly aligned with the focal point of the laser beams. There is no need for careful connection of a probe or modification of the sample to fit a specialized attachment. The sample only needs to be reflective enough where the measurement will be made that it is possible to obtain an adequate signal.

Third, it is affordable. The parts needed for a TGS setup can be procured for a relatively small amount of money in terms of research equipment. The parts are - with the exception of the phase mask - standard and easily obtainable from the usual research suppliers.

Fourth, it is versatile. The setup can be modified to fit the available space. Parts can be added to suit a particular research need. It is easy to modify.

The last two characteristics are particularly important because it means that TGS can be expanded to many labs. It does not require millions of dollars to build a TGS set up. Already, the group has begun setting up TGS experiments in other places, including Sandia National Laboratory and a university in Kazakhstan. The goal is to enable efficient nuclear materials research in as many places as possible. The TGS system is accessible and easy to replicate.

The first two characteristics - the non-destructive, non-contact nature of the technique - are crucial because it means that the TGS technique is adaptable for *in situ* use. Because it is non-contact, TGS can be used to monitor a material's properties while it is irradiated, as long as the surface is sufficiently reflective. There is no concern that a contact probe will block or interfere with the ion beam. Instead of irradiating a set of samples to different doses and performing destructive, time-consuming measurements on each, one can simply monitor a single sample continuously as it is exposed. This saves time and allows researchers

to access important information about material behavior that is lost when the material is studied *ex situ*. For example, *ex situ* studies will not include the effect of thermal vacancies, which are present in much higher concentrations at the high operating temperatures of a reactor than they are at room temperature. *In situ* studies allow insight into how the material behaves in a radiation environment without relying on assumptions.

Ultimately, by adapting these techniques for use with common nuclear materials, with an emphasis on using it to study radiation damage, we hope that transient grating spectroscopy will break open the bottleneck that has plagued the nuclear materials research community and which currently impedes advanced reactor concepts from making the leap from simulation to reality [1].

* * *

At this early stage of the overall TGS project, however, we are only just beginning to adapt the TGS system for *in situ* use on a beamline. The primary objective of this thesis project is to use TGS to study irradiated materials *ex situ*, as well as to characterize the material properties of samples with known defect populations. This will both validate that the TGS technique is appropriately sensitive to radiation-induced changes in material properties and demonstrate how TGS results can be used to analyze a sample's material properties as well as the nature of its likely defect populations.

This thesis begins with a general overview of radiation damage and its effects. The types and causes of radiation damage are considered, followed by how they impact material properties and - consequently - component performance. Existing methods of radiation damage detection, measurement, and analysis are considered. Their advantages and shortcomings are then discussed in detail in order to illustrate why TGS has so much potential as research tool for the nuclear materials community.

In the next chapter, the transient grating spectroscopy technique is explained. The use of TGS over the past decades in the broader materials field is summarized in order to contextualize its adaptation for nuclear materials. Particular attention is paid to the TGS setup built at the MIT Mesoscale Nuclear Materials group, as well as to the various modifications and adaptations that have been made to the setup since it was first constructed. The TGS signal will be explained, as are the methods used to analyze the signal. The benefits that TGS confers to a materials researcher, as compared to other methods, are also discussed. Information about the detailed operation of the setup and code for analyzing the signals is provided in an appendix to the thesis. This chapter concludes the "background" section of the thesis.

The research for this thesis was divided into three broad campaigns: the first involved the study of cold-worked niobium; the second, constitutional vacancies in intermetallic nickel aluminum; and third, the study of well-characterized samples from an in-core component of EBR-II. Each campaign is given its own chapter, instead of compiling all "methods" into one chapter and all "results" into another. Each chapter describes the goals of the particular study, the method of sample preparation, any sample characterization that was performed, the TGS experiments and results, and experimental analysis. Many of these experiments involved multiple iterations as techniques were refined and mistakes were made. So, also included in each

chapter are so-called “lessons learned,” which will hopefully save any readers interested in repeating similar experiments the need to make our mistakes a second time.

The “cold-worked niobium” chapter focuses on the impact of dislocations. The niobium samples are extensively characterized via XRD methods to measure dislocation populations. TGS is used to study the samples over 180° of surface rotation. The goal here is to understand how dislocations impact the TGS signal and the properties it measures. TGS results are also predicted mathematically and compared with experimental results. Select cold-worked niobium samples are irradiated to progressively higher doses, and TGS measurements repeated between doses. Dislocations are not only interesting in their own right - and a likely effect of irradiation - their presence also affects how the material behaves as it continues to be irradiated.

The “nickel aluminum” chapter uses the high concentration of constitutional vacancies in certain phases of intermetallic NiAl as a stand-in for the thermal vacancies that are present in most metals when they are held at high temperature. NiAl allows for the *ex situ* study of high vacancy concentrations at room temperature. NiAl TGS results are compared with molecular dynamics simulations of the TGS experiment on representative NiAl structures. The development of a code used to efficiently build these MD test structures is also explained, as this was a key factor in successfully carrying out the simulations. Whereas the niobium study is primarily focused on elastic effects of defects, the NiAl study is primarily focused on thermal effects of defects.

The “EBR-II” chapter is based on the TGS measurements of samples from the Experimental Breeder Reactor (II). These samples were in-core samples, and are still hot. They contain complex defect distributions due to their prolonged exposure to in-core operating conditions. They are also extremely well characterized by previous studies. This chapter, then, explores the utility of TGS for analyzing complex defect populations that exist in real nuclear materials as a result of exposure to a reactor core environment.

Finally, the conclusion summarizes the findings from each of these projects, and how each campaign represents progress toward the development of TGS as a standard technique for conducting nuclear materials research. The appendix contains relevant standard operating procedures used in this work, as well as any code that was developed specifically for this work - such as the codes used to generate NiAl structures for the molecular dynamics study and the codes used to analyze the TGS signals themselves.

Chapter 2

Introduction to radiation damage in metals and alloys

2.1 Introduction

This chapter is intended to provide an overview of what radiation damage is, how it is quantified, and current techniques for measuring radiation damage in the metals and alloys of interest to nuclear engineers. This will provide important context for the subsequent chapters, as well as elucidate some of the challenges with current detection and analysis techniques that were mentioned in the prologue to this thesis.

Radiation damage is often introduced using a simple ballistic model: picture a material, made up of atoms arranged in a regular pattern, and now picture a particle impinging upon it. If it has enough energy, the particle can knock one of the atoms out of its place. Figure 2.1 is an example of this heuristic. The radiation (green) is incident on a lattice of atoms (white circles). The radiation hits one of these atoms (now colored red). The radiation carries on through the material, transferring its energy to the material's atoms and maybe knocking some more of them out of place until it comes to a stop. The particles it hits might go on to do additional damage, either by knocking out other atoms or by coming to rest in an interstitial space in the lattice.

In a typical radiation environment, this process is multiplied many millions of times over: for example, the neutron flux in a typical commercial reactor is on the order of 10^{19} particles per cm^2 . The result is that, over time, radiation facing component materials undergo changes as they sustain radiation damage. The extent of the damage depends on the nature of the radiation, the operating environment parameters, and the properties of the material. Sometimes, these changes are significant enough that the component is no longer useful, and it must be replaced.

This section provides an overview of the types of radiation and where they come from. Section 2.4 provides a more detailed explanation of the mechanisms by which radiation damage occurs.





Figure 2.1: A basic schematic of ballistic-type radiation damage in a material. The impinging radiation knocks atoms from their place in the material’s crystal structure. These atoms may go on to knock other atoms out of place. The resulting disorder to the structure is the basic cause of radiation damage. Image adapted from [2].

2.2 The types of radiation

This subsection introduces the different types of radiation that cause radiation damage in materials.

2.2.1 Alpha radiation

Alpha particles (denoted α or He^{2+}) are comprised of two protons and two neutrons.¹ They are most commonly produced by the α -decay of heavy atoms. For example, uranium transmutes to thorium via the reaction



The α -particle produced by a typical transmutation process has a kinetic energy of ≈ 5 MeV.

High energy (≈ 16 MeV) α -particles can also result from fission reactions (ternary fission). This is relatively rare: less than 0.5% of fission event in a reactor will be ternary fission events. In the most typical instance of ternary fission, the heavy nucleus splits into two fission fragments of roughly equal size and a single α -particle. (This is in contrast to the vast majority of fissions, in which the split of the heavy nucleus yields two fission fragments of roughly equal size and, on average, two fast neutrons). In a reactor where thousands of trillions of fissions per second are occurring, though, this is a significant source of high-energy α -particles.

Typically, α -particles impinging on a component surface are stopped very quickly (this will be explained quantitatively in Section 2.4.2). This makes it very easy to shield α -particle radiation: they are readily stopped by human skin, for example. However, despite their short range, they are strongly ionizing. They present a major health risk if they are ingested because they cause major damage to chromosomes in the


¹Note that when an accelerator is used to produce He^{2+} ions, they are usually referred to as such. The term α -particle indicates that the particle is the result of some sort of nuclear reaction, be it α -decay or fission.

cells of biological tissue.²


In a reactor, α -particles can cause embrittlement and buildup of helium gas bubbles [3, 4, 5]. In a tokamak, they are known to cause surface blistering on plasma-facing components [6, 7, 8].

2.2.2 Beta radiation

Beta radiation (denoted β^+ or e^+ for positrons, and β^- or e^- for electrons) comes from radionuclides that undergo β -decay. In β^- decay, a neutron in an unstable nucleus of atomic number A is converted to a proton, an electron (e^-), and an electron antineutrino ($\bar{\nu}_e$). The electron (the β^- particle) and the antineutrino are emitted as radiation. The nucleus transmutes to an isotope of an element with atomic number $A + 1$. An example of β^- decay is



In β^+ decay, a neutron in an unstable nucleus of atomic number A is converted to a proton, a positron (e^+), and an electron neutrino (ν_e). The positron (the β^+ particle) and the neutrino are emitted as radiation. The nucleus transmutes to an isotope of an element with atomic number $A - 1$. An example of β^+ decay is x



β -decay is a common phenomenon in reactors because many fission products are unstable and transmute via β -decay processes to more stable elements. β -particles tend to create point defects in a material. Point defects may anneal out or they may combine into larger defects that impact material performance. (The neutrinos pass through the target material without interacting.)

2.2.3 Gamma radiation

Gamma radiation (denoted γ) is comprised of high-energy photons. In a reactor, γ -rays are primarily produced by γ -decay of certain radionuclides. γ -decay does not result in the transmutation of the radionuclide; rather, it simply leaves the nucleus in a less excited, more stable state [9].

γ -rays often have a long range and may pass through structural materials with relatively little direct damage. However, they may ionize atoms in the material, resulting in secondary β radiation [9].

2.2.4 Neutron radiation

In commercial reactors, neutrons present the primary radiation damage concern. Neutrons are chargeless and do not directly ionize the atoms of the target material. It is useful to consider the simple ballistic method illustrated in Figure 2.1 to understand the direct damage mechanism of neutrons.

²This illustrates an important semantic note for this thesis: in this project, the phrase “radiation damage” always refers to radiation-induced defects in the atomic lattice of metals or ceramics. It is not in reference to the radiation damage that occurs in biological tissue exposed to ionizing radiation.

However, if a neutron has a collision with an atom, that atom's nucleus may become ionized and behave as a directly ionizing particle when it recoils and interacts with other atoms. The neutron may also be absorbed by a target atom, causing the atom to become unstable and transmute, which then results in the creation of more ionizing radiation. (Of course, that target atom might also be one of the fuel atoms, in which case the transmutation is accomplished via fission, which produces more neutrons.)

The damage cascades associated with neutrons are discussed in further detail in Section 2.4.2.

2.2.5 Ion radiation

Linear accelerators are frequently used in nuclear materials research. Linear accelerators can produce parameter-controlled electron, proton, and ion beams, enabling the study of how materials behave when exposed to different radiation fluxes.

The right accelerator and ion source can be used to create a beam of nearly any kind of ion. In this subfield of nuclear materials research, this capability is important because it allows researchers to quickly induce a high level of radiation damage throughout a specific range in a target material. For example, in a later chapter of this thesis, niobium samples were irradiated with the goal of producing a detectable change in the materials properties. By using heavy ions, it was possible to complete the irradiations quickly - on a scale of minutes to hours, instead of the days to weeks that would be needed with lower-mass particles.

So, while one might not have to worry about a high flux of silicon ions in a reactor core, they can still be quite useful for experimental simulation of radiation damage, as use of a linear accelerator is nearly always going to be easier, faster, and cheaper than exposing samples in a research reactor. The damage associated with charged particles is discussed in more detail in Section 2.4.2.

2.3 Quantifying the interaction of radiation with materials

The range of radiation in a target material, and the amount and location of energy it deposits, depends on the characteristics of the radiation and of the target material. There are well-established methods to quantitatively predict the way this interaction is likely to unfold. A brief overview of the methods that will be relevant to this project is presented here.

2.3.1 Stopping power and range

Stopping power is used to describe how quickly radiation loses energy as it passes through a target material. In one dimension (x), stopping power S is described as a function of energy E as

$$S(E) = -\left(\frac{\partial E}{\partial x}\right)_{nuclear} - \left(\frac{\partial E}{\partial x}\right)_{electric} - \left(\frac{\partial E}{\partial x}\right)_{radiative} \quad (2.1)$$

Nuclear, electric, and radiative interactions all contribute to the radiation's loss of energy as it travels through the target material [9]. As an example, in the ballistic model, the radiation loses energy in a "collision" with an atom of the target material. This collision is actually a scattering reaction. A particle might undergo many such scattering reactions, losing some of its kinetic energy each time.

The inverse of the stopping power, integrated over energy, yields the expected range R of the radiation:

$$R = \int_0^{E_{max}} \frac{1}{S(E)} dE \quad (2.2)$$




Figure 2.2: Stopping power as a function of distance for 300 MeV protons in water. As the proton loses energy, it loses its remaining energy at a faster rate. This is an example of a Bragg curve. [10]

Figure 2.3: Stopping power divided by target material density for α -particles with varying energies (0.1 to 10 MeV). The graph shows how stopping power is a function of the incident radiation energy and the properties of the target material. [10]

Figure 4.11 is a plot of stopping power versus distance traveled for 300 MeV protons in water [10]. This is an example of a Bragg curve. Bragg curves show the rate of the radiation's energy loss as a function of distance traveled in the target material. The plot shows that the rate of energy loss increases the more energy the radiation loses. The Bragg curve also shows the particle's expected range: radiation and matter interactions are a stochastic process, but the curve shows that the expected range of 300 MeV protons in water is about 51 cm: some protons might travel a little further or a little short of this (this phenomenon is known as straggling), but the stopping distance of the protons would average out to this value. It is also evident that the proton loses energy at the fastest rate near the end of its expected range - after a certain amount of energy has been lost, the particle rapidly loses the rest of its energy and comes to a stop.

Figure 2.3 plots the rate of energy loss (divided by target material density) for 0.1 to 10 MeV α -particles in a variety of target materials. The graph shows that lighter target nuclei are more effective at stopping the

α -particles. It also shows, for each target material, there is a α -particle energy at which the target material is most effective at stopping the particles.

In planning radiation experiments, it's often necessary to determine the range and expected behavior of an ion of arbitrary energy in any possible target material. SRIM is a computer program that allows the user to model any number of any kind of ions of any energy and their predicted two-dimensional path through the target [11].




Figure 2.4: Result of a SRIM simulation of 4.15 MeV α -particles into an aluminum 7075 target. The expected range of the particles is about 19 μm . A small number of particles diverge from the expected path, perhaps as a result of backscatter from the target nucleus. However, it's clear that - probability-wise - this event is relatively rare, and the majority of the radiation's energy would be deposited near the 19 μm depth.

Figure 2.4 is an example from a SRIM study of 4.15 MeV α -particles impinging on an aluminum 7075 target. The figure is a plot of the trajectory (in x and y , with x being depth into target) taken by each α -particle simulated by SRIM. It shows that the expected range of these α -particles is about 19 μm into the aluminum alloy. Some particles stop slightly short of that; a small number appear to backscatter. Based on this simulation, one would expect the majority of the damage from these α -particles to be found 18-19 μm into the aluminum, since the majority of particle energy loss (e.g energy deposit into the target) occurs near the end of the particle's trajectory.

2.3.2 Quantifying damage in terms of atomic movement

Radiation damage is, fundamentally, caused by atoms being moved from their ideal lattice spots. This creates defects in the lattice that can become property-altering damage.

The **dpa** (an initialism of the phrase “**displacements per atom**”) is a pseudo-unit that is commonly used in radiation materials science to express the amount of radiation damage a material has sustained. The idea behind this unit is that it describes how many times, on average, a given atom in the target lattice will be knocked out of place by at least one lattice length by impinging radiation. So, if a component is expected to sustain “2 dpa per year” of radiation damage, that means that each atom would be expected to change its place on the component’s atomic lattice twice in a year. The dpa is often used a proxy for an amount of radiation, e.g. “The samples were irradiated to 1 dpa.” This is particularly common usage in nuclear materials, in which researchers are commonly interested in a specific amount of damage and the way it manifests in a given target.

To derive the formula for dpa, we begin with an expression for the rate of atomic displacements in a target material exposed to a radiation flux:

$$R = \int_0^{E_{max}} N \cdot \Phi(E_i) \cdot \sigma_D(E_i) dE_i \quad (2.3)$$

Here, N is the number density of the target, Φ is the radiation flux, and σ_D is the displacement cross section.³

A cross section has units of area: it is an expression of probability that a certain reaction will occur. So, $\sigma_D(E_i)$ quantifies the likelihood that radiation of energy E_i will cause one of the target atoms to be displaced. For most nuclear reactions, the probability of any given interaction type is dependent on the energy of the radiation; by integrating over the total energy spectrum of the radiation field, one can account for this fact.

In terms of units (using meters as the unit of length), this works out (quasi-conceptually) as:

$$\frac{\text{displacements}}{\text{m}^3\text{s}} = \frac{\text{atoms}}{\text{m}^3} \cdot \frac{\text{particles}}{\text{m}^2\text{s}} \cdot \text{m}^2$$

where we think of a displacement as being what happens when you multiply an “atom” by a “particle.” (Mathematically, of course, this is usually just left out; the unit of Φ is technically $\frac{1}{\text{m}^2\text{s}}.$) This calculation is repeated for every possible energy in the energy spectrum of the radiation, and these are all added together to calculate the total rate of displacements.

In order to get atom displacement rate, we simply divide R by N . This yields dpa per second. Multiplying by the total time the target material is exposed to the radiation field yields total dpa.

Figure 2.5 shows the typical dpa range and temperature regimes for radiation-facing components in various reactor types. It is clear from this figure that many advanced reactor concepts will face significant materials challenges as compared to commercial Gen-II reactors.

³The displacement cross section can be described as $\sigma_D E_i = \int \sigma(E_i, T) \nu(T) dT$. The displacement cross section for a radiative particle with energy E_i is the probability that that particle will transfer recoil energy T to the target nucleus multiplied by $\nu(T)$, the number of atoms that will be displaced due to a recoil nucleus with kinetic energy T , integrated over the possible values T can take. [12]

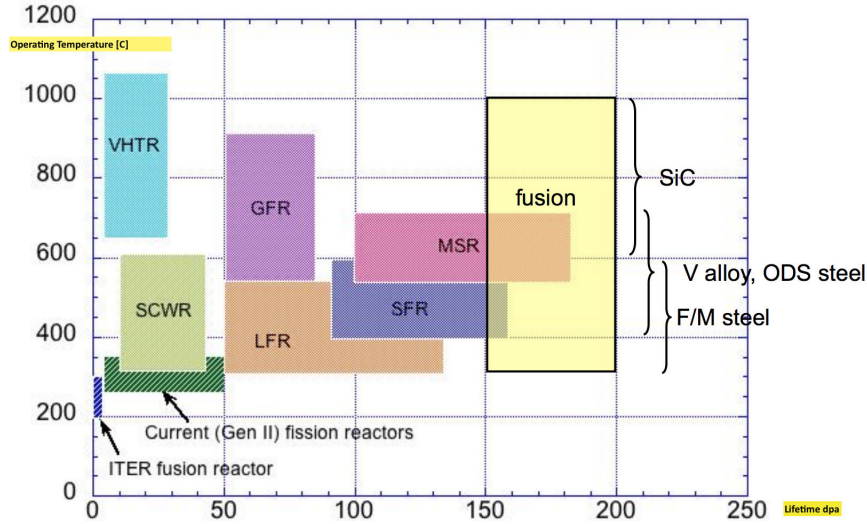


Figure 2.5: Lifetime dpa and operating temperatures for different reactor types. Advanced reactor concepts - particularly molten salt reactors, sodium fast reactors, and fusion reactors - present major materials challenges, because radiation-facing components must be able to withstand much higher amounts of radiation damage. [13]

Another common method for estimating the number of displacements is the **Kinchin-Pease model** [14]. The predicted number of displacements per atom (N_d) is calculated as a function of the energy available to cause displacements, T :

$$N_d(T) = \begin{cases} 0 & T_d < E_d \\ 1 & E_d \leq T < 2E_d \\ \frac{T}{2E_d} & 2E_d \leq T_d < E_c \\ \frac{E_c}{2E_d} & T \geq E_c \end{cases} \quad (2.4)$$

E_d is the displacement energy of the target atom (this is the minimum energy required to remove the atom from its lattice, and commonly has a value in the range of 20-25 eV [12]) and E_c is the cutoff energy for energy loss by electron stopping. The Kinchin-Pease model is most commonly used to estimate how many atomic displacements will occur in a single damage cascade. The Kinchin-Pease model makes many simplifications and assumptions to describe the interaction between impinging radiation and a target material, and many modifications to this model have been developed to better account for the complexities of the process. Some of these modifications include: utilization of more accurate energy transfer cross sections, treatment of energy loss due to electronic excitation, and consideration of phenomena specific to the target material (e.g. focussing). An overview and derivation of these modifications can be found in Ch. 2 of [12].⁴

⁴Equation (2.4) is from the version provided in the 2017 edition of [12], not the 2007 edition.

The limitations of using dpa as a true unit of radiation damage center around the fact that it doesn't elucidate the material property changes that take place in a material as a result of radiation damage. For example, a material irradiated to 1 dpa at one temperature might have completely different properties than the exact same material irradiated to 1 dpa at a different temperature.⁵

The parameters that define an irradiation experiment - the target material, the temperature, the ambient environment, the radiation type, the radiation energy, the radiation flux, the dose rate - all impact how a material will change when it is irradiated. The dpa does not tell us how much of the induced damage anneals out, or what kind of defects are expected to form. From an engineering perspective, it is more useful to consider the end state.

Knowing what changes in material properties have occurred during irradiation is more useful than the theoretical number of times a given atom in the material moved lattice spaces. However, it is still useful to have at least some standardized measure of radiation damage. After all, it is more illustrative to say that a material was irradiated to "2 dpa" than it is to simply state the integrated flux. It is important to remember, though, that the dpa is more useful as a measure of total radiation received than it is as an indicator of how the material may have changed.

G. Was makes this point clear in his textbook on radiation materials science:

... We cannot fully understand radiation effects by only calculating the number of displaced atoms. We cannot treat radiation effects as a black box. Rather, in order to understand the effect of the damage on the properties of the material, we must understand the fate of these defects after they are formed. ([12], p. 115)

In other words, predicting the number of displacements due to radiation is just the beginning of the story.

2.4 Mechanisms and effects of radiation damage

This section covers the mechanisms by which radiation damage occurs in greater detail, and identifies the common radiation-induced defects that affect nuclear materials.

⁵The following studies have found differences in the property response of identical materials exposed to the same dpa when the following parameters are varied [15].

- Temperature [16]
- Dose rate [17]
- Type of radiation [18]
- Beam rastering [19]
- Imposed stress [20]
- Processing and texture [21]
- Impurities [22]
- Coinjection of gas [23]

2.4.1 Radiation damage mechanisms

To begin, it's useful to return to the model presented in Figure 2.1. The likelihood that an impinging radiative particle will knock an atom from the crystal lattice of the target can be calculated using the models presented in the previous section. What happens next?

When an atom is knocked from the lattice, a Frenkel pair is created [24]. The Frenkel pair consists of two point defects: the vacancy (typically denoted with a v) and the interstitial (typically denoted with an i). The vacancy is the empty lattice spot left behind when the target atom is knocked out of place. The interstitial is what the target atom becomes - an atom wedged between the other atoms in the lattice, but not in a lattice spot itself. The vacancy and the interstitial cause distortions to the local lattice.

The damage path typically isn't limited to just the Frenkel pair, however. The interstitial atom had some amount of kinetic energy before it came to rest as an interstitial. It likely hit multiple other lattice atoms, possibly knocking more of them out of position as well. The target atom - which becomes a damage agent itself - is referred to as the **primary knock-on atom**, typically abbreviated as PKA. The chain process of atomic-level damage events set off by an impinging radiative particle is referred to as the **damage cascade**.




Figure 2.6: The initial damage caused by a PKA in the zircon lattice [25]

Figure 2.7: The PKA collides with other atoms, creating a damage cascade [25]

Figure 2.8: The damage begins to anneal out, as most atoms wind up back in a lattice spot [25]

Figures 2.6 - 2.7 are select screenshots of a molecular dynamics simulation of a damage cascade in zircon from [25]. The simulated damage was introduced by a zirconium PKA.

Figure 2.6 shows a mostly intact zircon lattice with a small defect in the lower left corner: a zirconium PKA (black dot) has impinged on the zircon lattice, knocking several atoms out of position (and creating a Frenkel pair for each damage site).

Figure 2.7 shows the disorder that arises in the immediate aftermath of the damage event: each atom that has been knocked out of its lattice site has some kinetic energy, and these atoms knock other atoms out. The damage and disorder multiplies exponentially. This stage is associated with a thermal spike due to the sharp increase of local atomic movement.

Figure 2.8 shows the quench state: this local temperature rise dissipates as heat is conducted away

through the lattice. Atomic movement begins to slow down, and some of the displaced atoms find new lattice spots. The lattice still looks quite disordered, but a significant amount of the damage has annealed out. This entire process happens on the order of picoseconds. Given more time, more annealing is expected to occur.

This illustrates something important: the majority of the radiation damage that occurs anneals out as atoms find their way back to lattice spots (typically, this is the more energetically favorable outcome). This process is enhanced by the local temperature spike associated with the initial damage cascade, and, in an operating reactor, by the high temperatures that are associated with atomic movement and annealing. This is why a sample irradiated to 1 dpa still typically looks like it did before irradiation, but with a few more flaws - each atom may have been displaced, on average, one time each, but most of them wound up back in a lattice site.

This phenomenon is typically referred to as **recombination** of a Frenkel pair: an interstitial winds up in an empty lattice spot, “annihilating” both itself and the vacancy. Other interstitials and vacancies might migrate to defect sinks like grain boundaries [12]. Point defects may also wind up combining into larger, more stable defects, which will be discussed in Section 2.4.3.

2.4.2 Damage cascades associated with different types of radiation

The behavior of the damage cascade described in the previous section is dependent not just on the energy of the impinging radiation, but on the type of radiation. Damage cascades created by electrons, protons, heavy ions, and neutrons with the same energy have different characteristics and material effects.

The 1 MeV electrons tend to create a single Frenkel pair defect. They lose energy in the target material almost entirely due to electronic interactions with the target atoms. 1 MeV protons tend to have multiple collisions before losing their kinetic energy. Each collision causes a small damage cascade. 1 MeV heavy ions have fewer collisions than protons, but are associated with a larger damage cascade. All of these particles lose energy via Coulombic interaction with the negatively charged electrons of the atoms in the target material [26]. The electrons of the target material may be excited to a higher electron shell, or they may be removed from their nucleus entirely, creating an ion (hence *ionization*).

The charged particles do not lose all of their energy in a single interaction. At most, the energy loss of a positively charged particle is equivalent to

$$\Delta E = \frac{4Em_{0,e}}{m} \quad (2.5)$$

with $m_{0,e}$ the rest mass of the target electron, m the mass of the charged particle, and E the kinetic energy of the particle. For an alpha particle, given that an electron has about 1/2000 the mass of a proton and a neutron has approximately the same rest mass as a proton, this means that the maximum energy an impinging alpha particle can transfer to an electron in the target material is $4Em_{0,e}/(8000m_{0,e}) = E/2000$. As a result, a single particle of ionizing radiation can affect thousands of atoms in the target before being fully stopped by

the target material.

Any charged particle will also lose energy via the radiation of electromagnetic energy (*bremsstrahlung*) as well as via Coulombic interactions [26]. The linear specific energy loss of a charged particle is given as

$$-\left(\frac{dE}{dx}\right)_r = \frac{NEZ(Z+1)e^4}{137m_0^2c^4} \left(4\ln\frac{2E}{m_0c^2} - \frac{4}{3}\right) \quad (2.6)$$

with Z the atomic number of the target material, m_0 the rest mass of the particle, and E its incident energy [12]. The energy loss due to radiation is low for heavy charged particles, but high for electrons, which has about 1/2000 the mass of a proton. *Bremsstrahlung* radiation has a continuous spectrum, so more energetic particles (like fast electrons) are capable of producing more energetic *bremsstrahlung* that can be capable of causing significant material or biological damage on its own.

Neutrons do not interact with the target material via the Coulombic force, and as a result tend to travel much further through a given absorber material than charged particles before interacting with a target atom [26]. When the neutron does interact with the target atom, it can transfer a large amount of kinetic energy to the target, now the PKA (or the recoil nucleus, as it is often called in the context of neutron-matter interactions). As discussed earlier, the PKA can be responsible for a significant amount of material damage and additional radiation, since they are effectively creating self-ion irradiation conditions in the target.

The neutron loses energy through each interaction with the target nuclei, and is eventually said to be thermalized when its kinetic energy is about 0.025 eV (roughly room temperature). In a typical thermal reactor, the majority of neutron-matter interactions will occur in the moderator (usually water).⁶ However, fission neutrons are born fast (about 2 MeV), and some of them will still interact with component materials. (This is even more of a concern in fast reactors and for advanced reactor concepts that utilize higher neutron fluxes.) After all, even a low-probability event can become significant when we expand our view from the travels of a single neutron to the 10^{13} neutrons per cm^2 per second that characterize the neutron flux in a typical reactor.

Figure 2.9 shows the predicted dpa per incident particle per cm^2 for different types of radiation impinging on a nickel target as a function of depth [27]. Heavy ions (tantalum, nickel) have the highest associated damages of the species considered here, but this damage is concentrated in the first few microns of the material. Lighter ions penetrate more deeply into the material. Neutrons, with their lack of charge, penetrate most deeply into the target, but the damage profile is effectively flat and doesn't display the characteristic Bragg peak behavior of the charged particles.




Figure 2.9: Damage profiles in a nickel target as a function of depth for different radiation species. Heavy ions have higher associated damage but short range; neutrons have much lower associated damage but penetrate into the material bulk. [27]

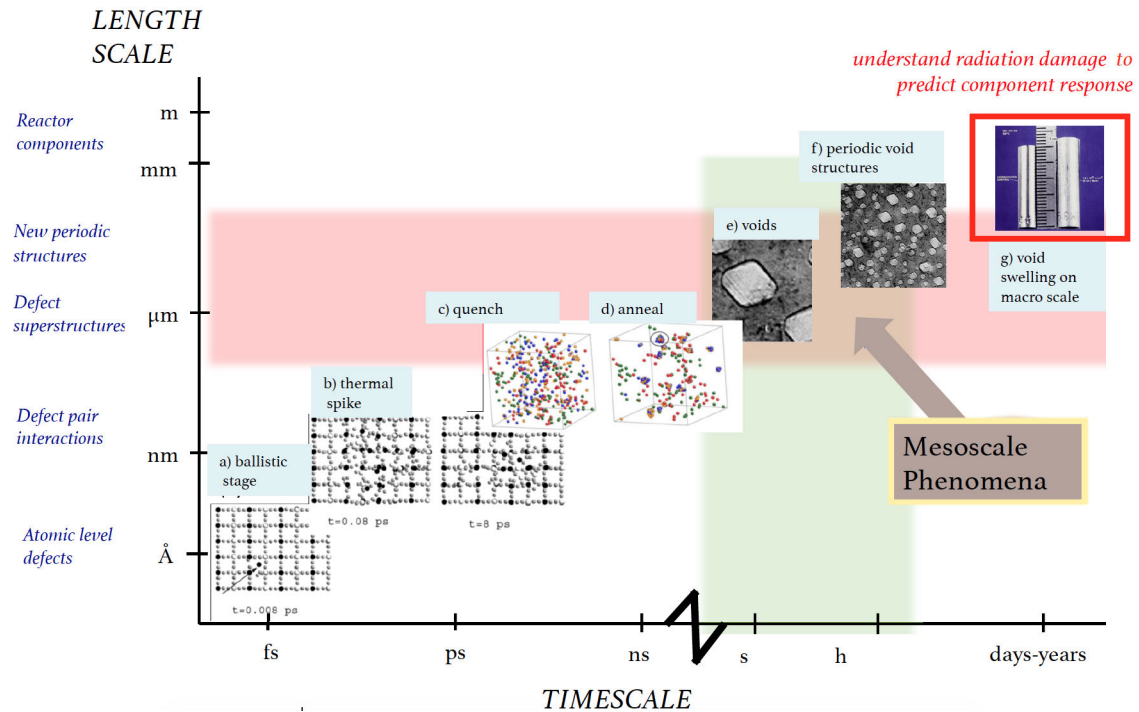


Figure 2.10: The progression of radiation damage in time and space, from atomic-level interaction to component-scale response. Impinging radiation creates a damage cascade and thermal spike, characterized by high disorder in the localized lattice near the radiation impact. Most of this damage anneals out, but some is let behind. Defects combine to form larger defects, such as voids. These larger defects may be found uniformly throughout a component, significantly altering its behavior. Over time, radiation damage alters material properties, and frequently limits the lifetime of radiation-facing components in nuclear applications. Adapted from [28].

2.4.3 Radiation-induced defects and resultant property changes

Most of the initial damage created in the cascade might anneal out, but some defects are left behind - and these can combine over time to create even larger defects with more significant impacts on material properties. This subsection is concerned with the various defect species that can arise in an irradiated material.

Figure 2.10 shows that these larger defects are the bridge from initial radiation impact to material changes that are significant on an engineering level. The damage that occurs on the picosecond scale (**b**, **c**) may mostly anneal out, but it still leaves behind small defects that combine over time to create larger mesoscale defects (**e**) that have significant material impacts. In some cases, these defects are found throughout the target (**f**), altering its structure and performance even more profoundly.

2.4.3.1 Point defects

The term “point defect” most commonly refers to the vacancies and interstitials created when atoms are knocked from their place in the target material lattice (see 2.4.1).⁷

Point defects are not stationary once formed, and move around the lattice via diffusion processes [12]. Vacancies “diffuse” when an atom moves from its lattice site to the vacancy site, filling the original vacancy and leaving a new one. For the purposes of modeling, the vacancy is treated like an interstitial (e.g. as a constant object), and diffusion is modeled from the perspective of the moving vacancy instead of from the perspective of the moving atom. Interstitial atoms diffuse by moving to new interstitial sites, although this is energy intensive and typically insignificant unless the interstitial is smaller than the lattice atoms [12].⁸

Most point defects anneal out, particularly due to recombination in the instants after formation. Others migrate through a material, and combine into larger defects discussed below or “disappear” at sinks. Grain boundaries, precipitates, voids, and dislocations are all examples of defect sinks in a typical material. Some sinks are classified as “biased” sinks: they cause a stress gradient that attracts one type of defect more than the others. Dislocations are biased sinks for interstitials because the stress gradient caused by a dislocation “attracts” interstitials to the dislocation core [12].

Mechanically speaking, an individual point defect only impacts its local lattice structure: a single interstitial, for example, imparts strain to the lattice immediately surrounding it, and has a negligible effect on the overall behavior of a component. However, if they exist in a material at sufficient concentrations, vacancies and interstitials can have measurable impacts on material properties even if they exist only as point defects (as opposed to clustering into the larger defects that are associated with more dramatic impacts on the lattice).

⁶The moderator is the substance used by thermal reactors to bring the neutrons born from fission to thermal energies where they are more likely to be absorbed by a fissionable atom. Water is the moderator in the majority of operating commercial reactors because neutron energy is efficiently transferred to the low-Z moderator atoms.

⁷It may also be used to refer to anti-site defects (ASD), in which an atom of an element that doesn't belong to the target material occupies a lattice spot. However, ASD are not typically a result of radiation damage.

⁸Mathematical treatment of diffusion is left out here, but Chapter 4 of [12] has clear descriptions of quantitative models of diffusion.

The lattice strain caused by point defects is not the only way they can change material behavior, however. For example, vacancies are known to scatter phonons and electrons [29], and this consequently decreases the thermal transport ability of the material [30]. They are also reported to affect the movement of dislocations [31]. In 1952, G. Dienes reported that vacancies decrease the elastic moduli of a metal, while interstitial defects increase it at a more substantial level, leading to a net increase of an elastic moduli [180]. Many other researchers have investigated the impact of vacancies and interstitials on elastic constants and described the elastic constants as a function of C_v and/or C_i [32, 33, 34, 35]. Not every paper is in agreement as to the extent or direction of these effects, but there is a general consensus that the contribution of point defects to material property changes is not automatically negligible.⁹

It is important to remember that, in addition to the Frenkel defects created by radiation, vacancies are often present in high concentrations due to elevated temperatures that are characteristic of many radiation environments. The number of vacancies in a material is described with an Arrhenius relationship:

$$n_v = n_0 \exp\left(\frac{-H_v}{RT}\right) \quad (2.7)$$

where n_0 is the number of atomic sites on the lattice, $-H_v$ is the enthalpy of vacancy formation, and T is temperature. It is reasonable, therefore, to assume that significant vacancy concentrations exist in component materials in operating reactors due to high temperature and a continuous radiation flux, and that this might have a corresponding impact on material properties. However, the same material, when removed from the source of heat and radiation flux, is likely to have a much lower C_v when examined *ex situ*. Therefore, even though the point defect concentrations *in situ* may be high enough to measurably change material properties (see page 52), this effect may not be observed if only *ex situ* measurements are carried out.

Vacancies and interstitials may cluster (or be formed in clusters). These clusters may be unstable and anneal out. If they are stable, they may diffuse as a cluster. They may also grow in size: for example, vacancy clusters may tend to join and form a void, while interstitial clusters tend to be more stable than vacancy clusters [12]. In this work, examples of the effect that point defects have on mechanical and thermal material properties will be discussed in 4.9, 5.1.2, and 6.4. Clusters will be treated in greater detail in Section 2.4.3.3.

2.4.3.2 Dislocations

Dislocations are larger-scale defects in the crystal's structure resulting from an "extra half-plane of atoms shoved into the lattice" [37]. The two basic types of dislocation are the **edge** and **screw** dislocations, and both of these are illustrated in Figure 2.11.

The edge dislocation is typically described as either an extra half-plane or a missing half-plane of atoms, such that the lattice is distorted near the mismatch. The Burgers vector is used to describe lattice distortion

⁹For example, it was predicted in [180] that a C_v of 1% leads to an elastic constant reduction of 1%, and a C_i of 1% can increase elastic constants by up to 10%. This estimation was refuted by F. Nabarro, who claimed that a C_v of 1% lead to an elastic constant reduction of 2.3% and C_i increased elastic constants by 3.8% [36]. Dienes refuted this refutation in the letter below Nabarro's in [36].

caused by the dislocation. Far from the defect, the distortion is negligible. A screw dislocation is typically visualized as the planes above the defect being rotated about an axis by one or more lattice sites - bringing to mind the spiraling threads of a screw. To be more specific, an edge dislocation is a dislocation that lies perpendicular to the Burgers vector; a screw dislocation is parallel to it [38].




Figure 2.11: A dislocation is “a discontinuity at which the lattice shifts from the unsheared to the sheared state” [38]. The schematic of a crystal lattice shows an edge dislocation and a screw dislocation, which are the two basic dislocation configurations. The Burgers vector describes the magnitude and direction of the lattice strain induced by the dislocation. [39]

Dislocations dramatically impact material properties, particularly as carriers of plasticity. The theoretical yield stress of a given real crystal is typically at least 1000 times greater than the crystal’s actual observed yield stress because of the presence of dislocations [38]. This is because material deformation fundamentally involves the breaking of bonds, the energy required for the dislocation to move is far lower than the energy required to break the unstrained atomic bonds of a perfect lattice. Dislocations move via glide or climb. Dislocation glide involves the movement of a single plane at a time in the slip direction; dislocation climb requires the diffusion of vacancies to the dislocation core [37]. The dislocation density of a material can be increased via cold work. At certain dislocation densities, the dislocations may begin to “pin” each other, leading to an effective decrease in plasticity and ductility and an increase in yield strength.

Of particular importance to nuclear materials is the dislocation loop, which is created by the clustering of interstitials or of vacancies as they migrate through a material under irradiation [41] (see Section 2.4.3.3 for more detail and literature examples regarding dislocation loops). Frank-Reed loops in particular act as dislocation loop generators in crystalline materials: they occur when a force acting upon a pinned dislocation causes it to bow out, eventually “touch” and form a loop, and restart the process over again. This process is best explained visually, and a schematic is shown in Figure 2.12. Frank-Reed loops are found in the EBR-II samples used in this project, and their effects are further explored in Section 6.3.




Figure 2.12: The Frank-Reed source is a pinned dislocation that forms a closed loop when force is applied. The loop becomes independent of the original dislocation source, which repeats the process. Adapted from [40]

2.4.3.3 Defect clusters

Clusters are configurations of adjacent multiple vacancy or interstitial point defects that act as a single unit. They typically form during the damage cascade process, during which a great deal of short-range diffusion and atom movement is taking place over a very short time in a highly localized area [12]. Clusters that are stable - typically, interstitial clusters - can migrate during irradiation to defect sinks, much like point defects. Vacancy clusters are usually unstable, but they may join together to form larger, more stable voids [12].

Vacancy clusters and interstitial clusters may also result in the formation of dislocation loops, which may grow to stable size if a sufficient number of additional defects migrate to the initial loop.^{10 11} The stable loops then interact with other dislocations in the crystal and can thus affect material behavior. Figure 2.13 is a TEM micrograph of interstitial loops in electron-irradiated ferrite [43]; Figure 2.14 shows the growth of interstitial loops in austenitic stainless steel, created by both electron beams and lasers and observed in real time in a laser-equipped high-voltage electron microscope [44].

2.4.3.4 Voids and swelling

Voids form from vacancy clusters that grow to stable mesoscale sizes. In many cases, a nascent vacancy cluster will shrink due to recombination of the vacancies with interstitials. In order for a void to grow, the defect generation and diffusion rates must be such that the cluster absorbs vacancies at a higher rate than it does interstitials [12]. Because void growth is so dependent on the point defect nucleation and diffusion processes, it is a highly temperature-dependent phenomenon [45].^{12 13}

¹⁰A mathematical treatment of the growth rate of dislocation loops in metals can be found in Ch. 7 of the 2017 edition of [12].

¹¹A dislocation loop has a dislocation line enclosed entirely within the crystal [42].

¹²A clear mathematical treatment of void nucleation rates can be found in Ch. 8 of [12].

¹³In [45], 316 steel exposed to a neutron flux in a fast reactor was examined *ex situ* for void swelling. As an example of this temperature dependence, voids were not observed in steel irradiated to a total dose of 5.2×10^{22} n/cm² at 380°C, but it was observed in steel irradiated to a lower total dose of 3.2×10^{22} n/cm² at a higher temperature of 560°C.




Figure 18. Micrograph showing $\langle 100 \rangle$ interstitial loops in α iron irradiated at 823 K with 1 MeV electrons in an HVM.

Figure 2.13: Interstitial loops are often visible in a TEM as black “dots” that appear in geometric spacings. [43]

Voids are the cause of radiation-induced swelling, which can be a macro-scale effect in which a component’s dimensions are significantly increased during irradiation, sometimes to the point of limiting component lifetime. Modeling and predicting the onset of void swelling is an issue of great importance in the nuclear materials field. Voids are also dislocation barriers, and as such affect the plasticity of a material. They are associated with a loss of strength and with an increase in embrittlement [46, 17, 47]. (See Section 2.4.3.6 for additional information on how defects affect mechanical properties.)

If a void cavity has a spherical shape, it is likely to be a bubble. A bubble is filled with gas (typically a byproduct of some nuclear transmutation process), and the pressure of the gas pushes outwards on the cavity walls. Was notes that there is not a clear consensus for the amount of gas concentration at which a void is more accurately defined as a bubble [12].

Figure 2.15 shows a TEM image of 50nm-diameter voids in copper that has been neutron-irradiated to 1.1 dpa [48]. Voids often have characteristic shapes that reflect the geometry of the material’s crystal structure.

The onset of void swelling is characterized by an initial hardening, followed by softening as swelling increases. Void swelling can also enhance crack growth and formation [52]. Most obviously, it can result in the aforementioned significant changes to the dimensions of a component. As a result of all of these factors, void swelling can limit component life in materials that are prone to it.




Figure 2 | Dislocation loop growth during different types of irradiation and thermal annealing. (a)–(d) Electron-beam irradiation at 726 K. (e)–(h) Laser irradiation at 806 K ($\lambda = 532$ nm). (i)–(l) Thermal annealing at 780 K after laser irradiation ($\lambda = 1064$ nm).

Figure 2.14: Interstitial loop growth observed *in situ* in austenitic stainless steel. [44]




Figure 2.15: Voids are visible in the microstructure of copper irradiated to 1.1 dpa at 300°C in the Oak Ridge Reactor. [48]




Figure 2.16: The dimensional effects of void formation are obvious in the sample on the right, which was irradiated to 1.5×10^{23} n/cm² at 533°. [49]




Figure 2.17: The left image shows void formation due to neutrons (uniform throughout the bulk of the material) and the right image shows void formation as a result of ion irradiation (short range, material only affected near irradiation surface). Left image: Porollo, Konobeev, and Garner, 2000 (see also: [50]). Right image: [51].

2.4.3.5 Irradiation-enhanced creep

Materials exposed to high temperatures and constant load may undergo a type of slow-developing, permanent deformation known as creep. Creep can occur at low loads - loads that result in a stress far below the yield stress of the material - which makes it of particular importance for high-temperature applications in which components will be exposed to mechanical loads (e.g. turbine blades or fuel cladding). High temperatures enhance the mobility of diffusions, which are the carriers of plasticity.

Irradiation tends to increase the rate of creep by increasing the number of point defects present, which enhance plasticity by increasing the number of dislocations (which carry plasticity) and the number of vacancy sites to which the dislocations can move. Importantly, this results in creep occurring at lower temperatures than it would occur in the absence of a radiation environment [12, 53, 54].

2.4.3.6 Radiation hardening

Radiation hardening occurs when the increased presence of impurities and defects due to irradiation increase the energy required to “unlock” a dislocation, thus allowing it to move and reducing plasticity. Once unpinned, the dislocation’s movement might continue to be slowed by the presence of defects in the matrix [12].

The increment in yield strength can be described as

$$\Delta\sigma_y \propto \alpha\mu b \sqrt{ND} \quad (2.8)$$

where μ is the shear modulus,¹⁴ b is the Burgers vector characterizing the moving dislocations, N is the number density of the obstacle and d is the average obstacle size [12, 55].

Radiation hardening can also lead to changes in the brittle behavior of the material. For example, because irradiation tends to cause an increase in σ_y , which leads to an increase in the ductile-to-brittle transition temperature [12]. (Additionally, radiation-induced segregation, discussed in Section 2.4.3.8, may also affect brittle behavior by depleting or enriching embrittling elements along grain boundaries and defect sinks.)

Figure 2.18 shows the increase in σ_y observed for neutron-irradiated copper single crystals that occurs due to radiation-induced defects blocking dislocation movement [56].

2.4.3.7 Blistering and fuzz on plasma-facing surfaces

Blistering is an issue in plasma applications. Plasma-facing components are exposed to high concentrations of helium ions as a byproduct of the fusion reaction. The combination of high temperatures and high concentrations of helium ions lead to impregnation of the component by the helium. The ions cluster together, forming bubbles near the surface. This creates “blisters” that rupture, leading to surface exfoliation [12].

¹⁴The shear modulus is the ratio of shear stress to shear strain in a material, or (Force/Area)/($\Delta L/L_0$), where L is length.




Figure 2.18: Single crystal copper undergoes hardening due to neutron irradiation. *nvt* is an archaic unit for fluence. [56]

“Fuzz” is a phenomenon known to occur in certain metals exposed to helium plasmas, including tungsten. Since many of the components used in fusion designs are made of tungsten, this damage mode is of particular interest to fusion engineers. The fuzz, so named because it resembles a fuzz-like coating when viewed in the SEM, occurs on the plasma-exposed surface and is comprised of small tendrils of the component material.¹⁵ A characteristic instance of fuzz is shown in Figure 2.19. Fuzz affects the thermal properties of the surface, and thus affects plasma stability. It can also break off as small particles of tungsten “dust,” which also cause plasma instabilities [61].

2.4.3.8 Phase instability and radiation-induced segregation

Radiation damage causes a net flux of vacancies and interstitials, (and, in an alloy, atoms of the alloying elements) within the bulk [12]. This occurs because of the presence of vacancies and interstitials, which are present in significant amounts due to radiation and high temperature. Oversized solute atoms in the lattice can diffuse via the vacancy by “swapping” spots. Undersized solute atoms in the lattice can diffuse as interstitials. In effect, this means that the vacancies diffuse in the opposite direction of the alloying elements (and interstitials diffuse in the same direction as the alloying elements). In both cases, the high temperature associated with the irradiation environment enhances the mobility of lattice atoms and of defects, and the differences in diffusion rate of each species can quickly lead to noticeable material effects.

Elements have different diffusion coefficients, and so the constituent atoms of an alloy will diffuse at

¹⁵There are several mechanisms that have been proposed to explain fuzz development, including: the rupture of helium bubbles [57], viscoelastic flow of near-surface atoms [58], and altered chemical potential near helium bubbles [59]. [60] postulates that fuzz is caused by a combination of previously proposed models, and is a result of helium bubble presence, surface defect diffusion, and surface roughening.




Figure 2.19: A “nano-tendrils bundle” that developed on the surface of tungsten exposed to a modulated helium plasma of 7.6×10^{25} and $1.6 \times 10^{22} \text{ m}^{-2} \text{ s}^{-1}$ at 1020K. The tendrils were sliced with a FIB and the cross section imaged (right image). [61]

different rates. This can lead to depletion of a particular alloying element near defect sinks such as grain boundaries: as vacancies move toward the sinks, the alloying elements move away. Since the elements move at different rates, the element that diffuses most quickly will become depleted near the defect sink [12].

This may lead to a process known as radiation-induced segregation, and it can have important effects on alloy performance. For example, chromium is added to many alloys to aid in oxidation resistance [62]. However, irradiation can lead to chromium depletion at grain boundaries via the mechanism described above. This leaves the alloy with weakened oxidation resistance at grain boundaries, which can lead to corrosion and cracking [12, 63]. Experimental evidence of radiation-induced chromium depletion at a grain boundary is shown in Figure 2.20 [64]. RIS effects are observed in iron-chrome alloys at low doses (0.01 dpa) [65].

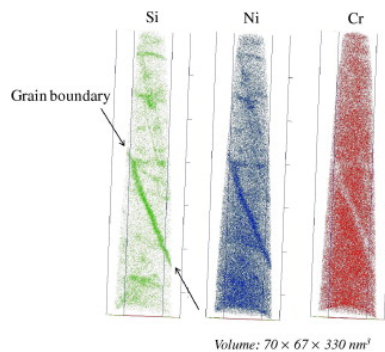


Figure 2.20: Atom probe tomography of cold-worked 316 stainless steel irradiated with 10 MeV Fe^{5+} shows the effects of radiation-induced segregation near a grain boundary. Silicon and nickel are enriched, while chromium is depleted. [64]

The movement of the alloying elements can also lead to the enrichment of one or more alloying elements (often near a grain boundary), which can lead to precipitation of one or more alloying elements if they become concentrated at a level above their solubility limit in the matrix element. The flux of alloying elements can also lead to the dissolution of existing precipitates.

Radiation-induced precipitation is another example of radiation-induced phase instability. In some cases, RIS enriches one element past its solubility in the alloy, and this causes the formation of precipitates [12]. As an example that is relevant to commercial reactors, neutron radiation is known to induce Ni_3Si , $\text{M}_6\text{Ni}_{16}\text{Si}_7$, FeTiP , and Cr_3P precipitation in austenitic stainless steel [66]. Radiation damage can also lead to lattice destabilization and loss of long-range order. In some case, this can result in the development of an amorphous phase [12, 67]. RIS tends to lead to the overall reduction in the concentration of solute elements in the bulk, a phenomenon that will be explored in greater detail in Chapter 6 as it is highly relevant to the observed TGS results of the EBR-II irradiated steel samples.

2.4.4 Neutron irradiation versus ion irradiation

Nuclear materials research tends to rely heavily on ion irradiation to test material response to radiation damage, even when the primary question of interest is how a material will respond to *neutron* irradiation specifically. Neutron exposure experiments are typically carried out in a research reactor core, and these are (relative to an ion accelerator) difficult to access and expensive to use. Obtaining permission and exposure time in the core is often a lengthy process in and of itself, and often requires a grant-like application process. It also takes a long time to carry out exposure experiments to reach representative lifetime dpa - often on the order of years. (This lengthy exposure time also increases the difficulty of successfully receiving permission to carry out the irradiation experiment in the first place.) Meanwhile, ion accelerators are usually more accessible, with a quicker on-ramp time to experiment, and a lower associated cost. The type of ion, its energy, and its current can usually be controlled easily, and high dpa can be achieved much faster.¹⁶ Furthermore, the neutron activation that complicates post-irradiation examination with samples irradiated in a reactor (as is the case with the EBR-II samples in this work) is eliminated as a concern when using an ion accelerator.

However, there are drawbacks to using ion accelerators to simulate neutron damage that arise from inherent differences in how each damage process proceeds. The radiation damage microstructure is evolving more rapidly, and this may lead to a material response that is not representative of how the material would behave in-core. Neutrons cause transmutation reactions that change how radiation damage proceeds due

¹⁶As an example of the flexibility provided by ion accelerators, in this work, Si^{3+} ions were used to irradiate niobium from 0.01 to 3 dpa. Silicon ions were chosen because they could yield high damage rates much faster - irradiations were complete on the order of minutes or tens of minutes, instead of the hours required by lighter ions. The current and power could be chosen to ensure that damage was uniform throughout the depth of material of interest. It is relatively straightforward to set the irradiation parameters to achieve a given dpa throughout a given sample depth in a given amount of time. (Meanwhile, the EBR-II samples are very useful for understanding how steel responds to neutron irradiation in a reactor core, but this involved 14 years of irradiation.) In [68], the necessity of accelerators for testing materials for advanced reactor concepts with very high expected lifetime dpa is described thusly: "200 dpa can be reached in days instead of decades."

to gas production [68]. Figure 2.21 shows some of the differences in defect populations observed in HT9 samples irradiated to the same dpa in the Fast Flux Test Facility (FFTF) and with Fe^{2+} ions in an accelerator. Not only are the volume fractions of each considered defect mode different depending on irradiation mode (graph (a)), but the characteristics of the defect populations themselves vary (graph (b)). However, the types of defects present are the same in both cases, and void swelling results in particular are comparable. Furthermore, the reported defect populations have the same order of magnitude for both the neutron irradiation and the ion irradiation cases. [68] discusses ways in which ion irradiation can be controlled to be a better simulacrum of neutron irradiation (e.g. by considering irradiation temperature effects or adding a gas implantation step to the ion irradiation experiment). The researcher must ask: “how close does the agreement need to be to successfully emulate the irradiated microstructure?” ([68], p.36)

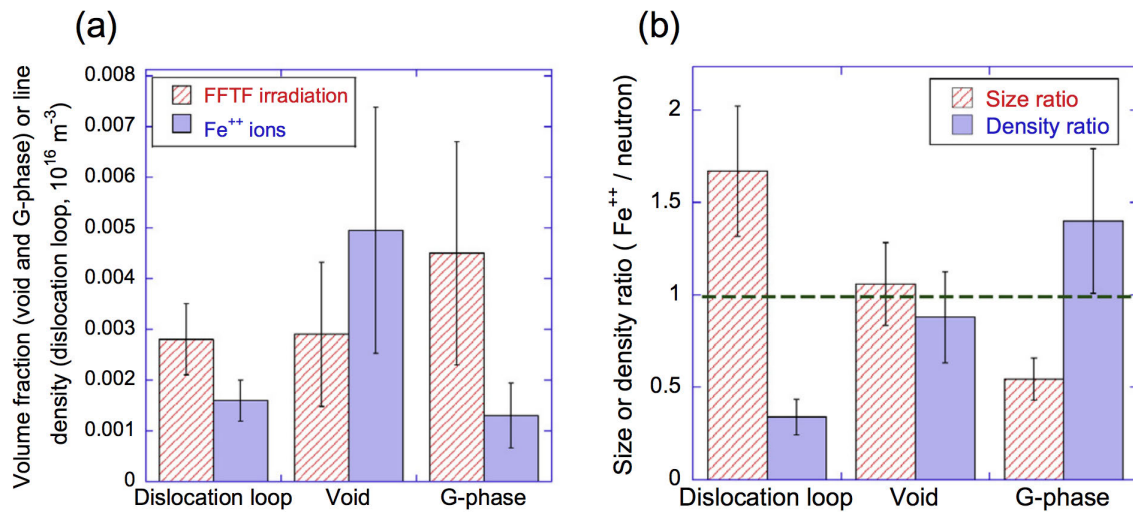


Figure 2.21: HT9 samples irradiated to the same dpa with neutrons (in FFTF) and with Fe^{2+} ions exhibit differences in defect population and typical size or density ratio of those defect populations, indicating that ion irradiation is not necessarily a perfect analog to neutron irradiation. [68]

Ion irradiation can also result in artifacts of the ion irradiation process that are not representative of radiation damage that would occur in a reactor, such as beam rastering artifacts and high ionization rates [69].¹⁷ In particular, ion irradiations are associated with a wider range of PKA recoil energy than neutron irradiations, and this affects the type of radiation damage caused by the PKA [69].¹⁸

These issues are brought up here because ion irradiation studies are so important in nuclear materials research, because *ex situ* studies of ion-irradiated materials with TGS are important to this work, and because

¹⁷Beam rastering is also known to suppress void swelling rates, which is important because void swelling is one of the important damage modes to consider when validating a material for use in a reactor [70, 71, 72].

¹⁸See also: [73, 74].

one of the eventual goals of the TGS project is to use it primarily as an *in situ* radiation damage technique in conjunction with an ion beam. As the TGS project proceeds, it is important to be mindful of these differences when connecting findings from ion-irradiated materials to the implications for the same materials in a neutron-radiation environment. Once TGS is fully capable of *in situ* real-time measurements on an ion beam, it will also be beneficial to the ongoing research regarding the differences between neutron and ion irradiation, and how to adapt ion irradiation to simulate neutron damage. This is because TGS methods will provide a great deal of insight into how radiation damage affects materials in real time, particularly at low doses, for any arbitrary combination of ion species, beam energy and current, and rastering vs. diffused spot conditions.

* * *

The notion of “radiation damage,” as can be seen from the very brief overview of its modes in this chapter, is almost too broad to be useful from an engineering standpoint. It encompasses defects that span the atomic to the meso-scale in terms of size, and which might form in a picosecond or over many months of sustained irradiation. The unit used to measure the extent of radiation damage, the dpa, tells us very little about the damage modes or their effects on a material - if anything, it is only useful as a shorthand for dose.

Nuclear materials engineers need to know how radiation damage has changed a material’s mechanical and thermal properties, as this will determine the material’s ability to perform properly inside a reactor or other nuclear application. An in-core component’s ability to withstand the mechanical loads in-core and transfer heat as intended is crucial to its successful performance and, frequently, the safe operation of the reactor. This information can’t be gleaned from the dpa (without extensive cross-referencing against literature-reported materials changes for the same material irradiated the same way to the same dpa). Transient grating spectroscopy, however, will be able to provide this information - changes in material properties due to changes in radiation-induced defect populations - and is therefore poised to become an important radiation damage diagnostic technique.

Chapter 3

Detecting and measuring radiation damage with existing methods and with transient grating spectroscopy

This chapter first considers the characteristics of an ideal radiation damage measurement system: what does the nuclear materials field need in order to make experimental radiation damage studies more efficient and more effective? With this context in mind, the next sections provide an overview of some of the most common experimental methods used to characterize and measure radiation damage. By first examining the uses, advantages, and limitations of existing measurement methods, the particular benefits of transient grating spectroscopy are more apparent.

The second half of the chapter is focused on describing the transient grating spectroscopy technique. The basic idea of the method and existing applications are explained. Next, the TGS experimental setup built by Cody Dennett at the MIT Mesoscale Nuclear Materials group is discussed, along with the procedures and codes used to analyze the TGS data. Modifications to the setup since its initial construction are also described.

3.1 Detecting and measuring radiation damage

The prologue to this thesis discussed the bottleneck problem in radiation materials science: validating a new material for use in a radiation environment requires extensive expensive and time-consuming exposure experiments followed by equally expensive and time-consuming analysis of the resultant damage. As discussed in Section 2.3.2, knowing dpa to which a component was irradiated is not enough on its own to elucidate how it responded to radiation: experimental characterization of the irradiated material remains a

vital part of nuclear materials research.

This section will provide an overview of existing radiation damage characterization methods, and explain why they are not sufficient on their own to break the nuclear materials research bottleneck. First, though, we should consider the “wishlist” for a method of radiation damage measurement that *would* be able to dramatically increase the efficiency of nuclear materials research.

- Ideally, it would be reasonably priced. If it is too expensive for all but a few labs to own and operate, it won't be able to gain traction as a standard method of measurement. (It should be noted that this project did not have a dedicated grant at any point during this thesis work, so this constraint very much applied to us as well.)
- It should be non-contact. If the equipment doesn't need to be in direct contact with the sample, we have fewer constraints on how it can be adapted. Techniques that require sample contact with a probe or transducer often have requirements on the geometry of the sample that limit their applicability. Furthermore, contact probes can block ions, precluding *in situ* measurements.
- It needs to have an adaptable geometry. This means both that it must not be too large and that it must be able to be modified, so that it can be adapted for *in situ* beamline use.
- Its probing mechanism must be radiation resistant, so that it can be adapted for *in situ* beamline use. The probing mechanism should also be usable in very low or very high temperatures.
- It needs to be non-destructive. The reason is twofold. First, the measurement method itself must not alter the sample in a way that affect subsequent measurements. Second, it cannot be eventually adapted for *in situ* component monitoring for operating nuclear technology if it damages the components.
- It needs to make measurements quickly. Part of breaking the bottleneck involves making measurements more efficient. However, continuous real-time measurements for *in situ* applications will not be possible if the equipment takes longer to make the measurement than it takes for the sample properties to measurably change.
- It needs to be able to probe near-surface phenomena. Charged particle irradiations tend to incur damage near the surface, and characterization methods that only probe the bulk of the material can miss these.
- It is capable of measuring inherent material properties - e.g. Young's modulus, or thermal diffusivity - instead of (or in addition to) only observing local features and characteristics.
- It should be able to measure multiple material properties of interest at once from the same measurement in order to increase the rapidity of a measurement campaign.

A goal of the TGS initiative at the MIT Mesoscale Nuclear Materials group is to develop a radiation damage characterization method that *will* meet all of these criteria, and which will be able to be readily adapted for the needs of any laboratory studying radiation damage in nuclear materials.

3.1.1 A brief overview of destructive methods of studying radiation damage

Destructive methods of measuring radiation damage are any method that involves altering the sample. Perhaps the most important example for the study of radiation damage is transmission electron microscopy, which involves the creation of a thin foil from the sample to be examined. (Commonly, this involves the use of a focused ion beam (FIB) to create a specialized thin foil from the irradiated sample that is then used for microscopy, although this has the additional complication of sometimes imparting damage to the sample that must be separated from the damage modes of interest [75, 76].) Electrons pass through the foil, and their interactions with the foil are used to create a very detailed image. Like a scanning electron microscope, the TEM utilizes the fact that electrons have a much smaller De Broglie wavelength than photons to achieve resolutions that are far higher than can be achieved with any optical equipment. Unlike an SEM, the TEM utilizes transmission instead of scattering, and so the resulting images can be used to study internal lattice phenomena such as dislocations and vacancies, instead of only revealing surface features.

Figure 3.1 shows TEM images of an ion-irradiated target that shows increasing amorphization as fluence is increased. TEM is valuable because it allows detailed imaging of materials at the level of individual atoms, allowing for a precise characterization of microstructural changes. However, the sample is so thin that it may be difficult to gain a full picture of radiation damage in the bulk without creating many samples (particularly for large components which may have axial and/or radial damage gradients). It may also make it difficult to study defects that are large relative to the sample thickness. It is not possible to solve this by simply creating thicker samples: the thicker samples require electron beam currents that may damage the samples, and the resolution limit in thicker samples is therefore only on the order of 1-3 nm.

Other examples of destructive radiation damage characterization methods include¹:

- Etching a sample surface with a corrosive agent to reveal etch pits to count dislocation density and other defects [78]
- EDX analysis of a sample surface to identify elements present (it tends to leave “burn” marks in the surface) [79]
- SEM analysis of microstructure - sometimes accidentally (charge can build up and destroy surface features) and other times as on purpose (SEM electrons used as the radiation source) [80]
- Hardness testing (leaves indents on the sample surface) [81]

¹Citations are examples of the method being used to study radiation damage, not examples of why the method is destructive




Figure 3.1: TEM images of fluorapatite irradiated with 1 MeV krypton ions show the progression of amorphization as fluence increases. The TEM allows for imaging and analysis at the atomic level. [77]

- Inductively-coupled plasma (ICP) mass emission spectra analysis of composition (sample must be dissolved first) [82]
- Tensile testing (permanently deforms sample) [16]
- Atom probe tomography (APT) (requires a section of the sample to be cut out) [83]

In many cases, it is possible to utilize a destructive characterization method that leaves the majority of the sample intact. It is typically not the “destructive” nature of these methods that is a limitation to efficient radiation damage measurements, but rather, the fact that they all involve direct contact with the samples. This makes it very difficult to adapt the method for *in situ* use, as it puts significant constraints on experimental setups.

Most of the listed methods - etching, hardness testing, ICP, and APT in particular - are explicitly *ex situ* methods that require direct alteration of the sample. The TEM has been adapted many times for *in situ* use to study samples as they are irradiated [84, 85, 86]. However, it remains time-consuming, painstaking and expensive [87]. One must have access to a TEM adapted for use with an ion beam or similar and access to a FIB. Furthermore, the TEM is also best suited for study of specific defects in the nano- and microstructure of a sample: it does not yield quantitative information about the sample’s material properties. Tensile testing has also been adapted for *in situ* use during irradiation, although this practice is confined to specific experimental campaigns and is not a standard method of material characterization [88].

Table 3.1 summarizes the various advantages and disadvantages of these destructive methods with regards to the criteria for an ideal measurement method.

Table 3.1: Advantages and disadvantages of various destructive radiation damage characterization methods, **✗** = no, **✓** = yes, **≈** = possible

Property	TEM	SEM	EDX	Hardness testing	ICP	Tensile testing	Etching	APT
Low cost	✗	✗	✗	✓	✗	✓	✓	✗
Non-contact	✗	✗	✗	✗	✗	✗	✗	✗
Adaptable geometry	✗	✗	✗	✓	✗	✓	✓	✗
Radiation-resistant probe	≈	✓	≈	✓	✗	✓	✗	≈
Temperature-resistant probe	≈	≈	≈	✓	✗	✓	✗	≈
Non-destructive	✗	✗	✗	✗	✗	✗	✗	✗
Real-time measurement	✗	✓	≈	✗	✗	≈	✓	✗
Detect near-surface phenomena	✓	✓	✓	✓	✗	✗	✗	✓
Measures inherent property	✗	✗	✓	✓	✓	✓	✗	✗
Measure multiple properties	✗	✗	✗	✗	✗	✗	✗	✗

3.1.2 Nondestructive methods

Nondestructive methods leave the samples unchanged by the measurement procedure, and are therefore better candidates for eventual development into a rapid-measurement, *in situ* measurement method that is capable of measuring radiation damage changes in real time.

Ultrasonic characterization methods encompass a broad class of techniques that measure the velocity of an ultrasonic acoustic wave as it travels through a material. Changes in wave velocity v_{us} can be correlated to changes in material density and elastic modulus, as well as to changes in void density and carbide precipitate population, according to the following equation [89]:

$$v_{us} = \sqrt{\frac{E(1-\nu)}{\rho(1+\nu)(1-2\nu)}} \quad (3.1)$$

where ν is Poisson's ratio, E is the elastic modulus, and ρ is density. Ultrasonic methods have been proposed as a possible method of continuous in-core monitoring of reactor components [90, 91]. They are already used industrially in the nuclear industry to monitor certain types of components - for example, CANDU fuel channel pressure tubes are inspected for flaws using ultrasonic systems [92], and ultrasonics can be used to test PWR pressure vessels for changes in embrittlement [93]. Various ultrasonic systems have been proposed for monitoring in-vessel structures of sodium fast reactors (SFR) due to the opacity of the sodium coolant

[94, 95, 96]. Similar systems have been proposed for use in lead-bismuth eutectic (LBE) cooled designs for the same reason [97, 98]. Nonlinear ultrasonics have been explored for real-time monitoring of radiation damage in pressure vessels [99, 100]. However, these techniques rely on empirical correlations developed based on changes in a single lumped parameter, and it has not been established that this lumped parameter is able to be sufficiently correlated to the mechanical response of the RPV.

A disadvantage of traditional ultrasonic methods include the difficulty in separating effects of E and ρ on the observed changes in v_{us} , which adds another layer of complexity to the matter of correlating speed changes to changes in different defect populations. Analysis requires explicit knowledge of the sample geometry, which adds another step to measurement and analysis. These methods typically require direct contact with the sample: while not destructive, this is also not ideal for many *in situ* measurement needs. Furthermore, and perhaps most importantly, conventional ultrasonic methods do not provide insight into thermal diffusivity, and their resolution ($\approx 10 \mu\text{m}$) is not sufficient to detect point defects and dislocations.²

The principles of acoustic waves have also been utilized to develop nondestructive microscopy techniques, including scanning acoustic microscopy (SAM), scanning laser acoustic microscopy (SLAM), and holographic SLAM (HOLOSLAM) [103, 104, 105]. However, all three techniques require direct coupling with the measurement sample, and can therefore not be classified as non-contact.

Surface Brillouin Scattering (SBS) is an example of a non-destructive, non-contact technique that has been used to measure radiation-induced changes in materials [106, 107]. To perform SBS, a laser is incident on a sample. The sample's acoustic modes, which are affected by any damage incurred by the material, determine how the incident light is inelastically scattered from the sample surface [108]. Changes in the frequency of inelastically scattered light can then be correlated to material changes. However, SBS measurements take hours to collect: it is not suited for rapid measurements. It is also best for probing materials at the atomic scale, whereas dislocations, voids, and many other radiation damage defect modes are characterized as mesoscale phenomena.

Resonant ultrasound spectroscopy (RUS) uses principles of internal friction to measure the elastic tensor of a sample by studying its mechanical resonances [109, 110]. However, it requires contact with the material, as well as regular, well-defined sample geometry, making it ill-suited for *in situ* use on any general component.

3.2 Transient grating spectroscopy: the basic idea

Transient grating spectroscopy (TGS) is a non-contact, non-destructive mechanical spectroscopy method that is capable of providing information about a material's acoustic, thermal, and elastic properties in a single measurement. TGS has been used for several in the materials science field to study thin films [111,

²Laser ultrasonic methods can be used for non-contact testing, but they still present the same issues with regards to defect resolution and inability to directly measure thermal properties [101, 102].

112, 113].³ ⁴ The MIT MNM group has postulated that TGS can be used to study nuclear materials, and that it can be adapted to fulfill the criteria identified on page 64. The high-level objective of this thesis is to prove that TGS is capable of measuring radiation-damage-induced changes in the materials commonly used in nuclear materials.⁵

TGS (as used in this work) uses pulsed laser beams (the “pump” beams) to generate a standing acoustic wave (SAW) on a sample surface, and a continuous laser beam (the “probe” beam) to study the SAW behavior. Figure 3.2 shows the two basic steps of TGS: excitation of the SAW and detection of the diffracted probe beam. More specifically, the SAW is created due to the interference of the pump beams with each other on the sample surface: this creates an interference pattern on the surface with periodic spatial variation in intensity. The intensity of the pattern correlates to the amount of heat energy imparted to the sample surface. This heat energy causes the sample to undergo thermal expansion, with more expansion occurring where the diffraction pattern is more intense. This in turn launches acoustic waves (material waves) in the sample, which interfere with each other and create the SAW. The pump beams are pulsed: this is important because the SAW decays away in between pulses of the pump beam. This allows the researcher to capture not just the frequency of the SAW, but the way in which the SAW dissipates. This will be explained in greater detail in the next pages.

This section will further elaborate upon the basic principles behind TGS. Most importantly, the TGS setup built by Cody Dennett of the MIT MNM group is presented and explained. Modifications that have been made to the setup since its initial build are detailed. Finally, the methods that are used in this thesis to analyze data obtained from the TGS setup are described.

3.3 The transient grating spectroscopy experimental setup in the MIT Mesocale Nuclear Materials group

The original setup layout. Figure 3.3 is a basic diagram that shows a simplified original layout of the first iteration of the TGS setup. Both laser beams (1) are first sent through the phase mask (2), which is also referred to as a diffraction grating. The diffraction grating is a specialized glass slide that can split the beams

³These citations are just a few examples; TGS of thin films is a subfield unto itself, largely due to the efforts of the Nelson group at MIT. Of particular usefulness to the MIT MNM group in beginning this project were the theses written by J. A. Johnson in 2011 [115] and J. K. Eliason in 2015 [116].

⁴Note also that TGS has been used in some metallurgical, non-thin-film applications, albeit with key differences from the way in which it is used here. For example, in [117], the authors used TGS techniques to study the orientation and elastic properties of micrograins in polycrystalline austenitic steels. The authors make multiple measurements of observed SAW velocities in a given propagation direction by measuring time-of-flight of the laser pulses and the known travel distance. They then determine the probability that a given velocity in the resulting distribution of observed SAW velocities has of occurring. This information can be used to obtain elastic information about the grains. However, this was done by inducing SAW at one point on the sample surface, and probing the sample at a different point, with the goal of detecting when the SAW arrived. In this work, the pump and probe spots overlap, and SAW speed is calculated from the known grating spacing and measured frequency of the SAW.

⁵The Hofmann group at Oxford university concurrently began studies of TGS for nuclear materials - namely ion-irradiated tungsten, and as such have become natural collaborators with MIT MNM. Their work was also assisted by the Nelson group at MIT. In particular, see [118], in which TGS techniques are used to measure the increase in elastic anisotropy of single crystal tungsten implanted with helium ions.




Figure 3.2: Two pulses of laser (the “pump” beams) light overlap on a sample (here, illustrated as a thin film on a substrate). The pulses generate an interference pattern on the sample, resulting in thermal expansion that is proportional to the intensity of the interference pattern. The non-uniform thermal expansion results in the propagation of acoustic waves throughout the sample. These acoustic waves interfere with each other, creating a standing acoustic wave (SAW) on the surface. The SAW can be probed with a second laser beam. This “probe” beam diffracts from the SAW. The signal collected from the diffracted probe beam can be used to analyze the properties of the SAW, and thus the properties of the sample. [114]

into the desired θ : the slide actually has multiple diffraction gratings, each one creating a different beam-splitting angle θ . Since θ controls the wavelength of the SAW (see Equation (3.2) below), we can move the phase mask to change the properties of the SAW and the depth of the sample it probes. All diffraction orders except the ± 1 are blocked. Not pictured in the diagram is the block for the direct reflection zeroth order light, which is in between the phase mask and lens 1 and whose alignment is important to avoid light contamination from this intense diffraction order. Lens 1 makes the beams parallel again. One order of the probe beam, the reference beam, is decreased in intensity via an adjustable neutral density filter that ensures that the detector does not oversaturate. The other order, still referred to as a probe beam, passes through a phase adjust (4), which is a small square glass slide whose angle can be changed to alter the path length of the probe beam. This is important for the heterodyning of the signal, which will be explained later. Lens 2 is used to recombine all beams on the sample surface, with the sample being on an adjustable holder that can be moved to align with the focal point of the re-crossed beams. The probe beam then refracts off the SAW created by the pulsed pump laser, and optics are used to direct this light into the detector (6).⁶ The detector signal was sent to the oscilloscope (not pictured).




Figure 3.3: This schematic was used to explain the TGS setup in the early phases of the project. At (1) the probe and pump laser emit beams that enter the phase mask (2), which determines θ . All diffraction orders of the pump and probe beams are blocked except for the ± 1 . One order of the probe beam is used as the reference beam, while the other is referred to as the probe beam (3). (This is part of the heterodyne scheme, which is explained in greater detail in the text.) The beams enter lens 1, which sets them back on parallel paths. The reference beam passes through an adjustable neutral density filter to avoid detector saturation, while the probe beam passes through an adjustable glass slide that changes its path length (4). Lens 2 is used to recombine the beams (5), with the sample at the focal point. The SAW is created on the sample surface by the pulsed pump beam and the probe beam and reference beam diffract from the SAW. The diffracted light (the first orders of the probe and reference beams) carries the signal information (orange line), and optics are used to send it to the photodetector (6).

Excitation region characteristics: wavelength and spot size. The wavelength Λ of the SAW is given as

⁶See also: [15]




Figure 3.4: A long exposure photograph of a TGS experiment in progress in the MIT MNM laser lab. The green pump beams and red probe beam follow the same path through the setup optics to the sample. The probe beam diffracts from the pump-induced SAW and into the detector.

$$\Lambda = \frac{\lambda_e}{2 \sin(\theta/2)} \quad (3.2)$$

where λ_e is the wavelength of the pump beam and θ is the angle set by the phase mask. Note that Λ is also a good indicator of the sample depth that TGS will probe, and so the phase mask can be used to increase or decrease this probing depth by changing θ .⁷ The overall size of the excitation spot is typically about $200 \mu\text{m}$ diameter. The probe spot is spatially aligned with the pump spot, but it is smaller (it was measured at about $140 \mu\text{m}$ diameter for the TGS experiments presented in this thesis) to ensure that it is capturing only SAW behavior, as shown in Figure 3.5.




Figure 3.5: The probe beam spot overlaps the pump spot, which creates the transient grating on the sample surface. The probe spot is smaller than the pump spot to ensure that the probe is not partially capturing the sample surface outside of the excitation region. [119]

The heterodyning scheme. The diffracted signal is amplified via heterodyning [120]. Heterodyning makes

⁷See [119] for additional notes on the size of the excitation region and its dependence on the setup optics.

use of the fact that if there is a phase difference ϕ between two oscillators, their combined intensity will be dependent on a cross term between them that can be used to amplify the signal via control of ϕ . The probe beam and the reference beam spot and their first order diffraction from the sample surface are spatially coincident. (Note that if there was no heterodyning, then we would only need one of the two beams in order to probe the sample.) Heterodyning allows for the amplification of the diffracted probe signal: the information encoded therein is not changed, but it is easier to detect and the signal-to-noise ratio is increased. The heterodyned signal intensity is given as:

$$I_S(t) = I_R + I_D(t) + 2\sqrt{I_R I_D(t)} \cos \phi \quad (3.3)$$

where $I_S(t)$ is the intensity of the signal, I_R is the intensity of the reference beam, $I_D(t)$ is the intensity of the diffracted signal, and ϕ is the phase difference between the reference and diffracted probe beam.⁸ I_R is constant throughout the measurement, although it can be adjusted with the neutral density filter to avoid detector saturation. ϕ is changed by altering the path length of the probe beam by changing the angle of the phase adjust.

The phase adjust is simply a glass slide (ours are squares about 1" by 1" in size and about 1/8" thick) that can be tilted back and forth. The probe beam always passes through the phase adjust. Glass has a different refractive index than air. The "neutral" position of the phase adjust is typically taken to be when it is perfectly perpendicular to the optical table: this represents the shortest distance that the probe beam can travel through the glass slide. When the phase adjust is tilted relative to the optical table, the probe beam travels a slightly longer distance through the glass. The overall distance the probe beam travels between lens 1 and lens 2 does not change: the phase adjust simply alters the portion of that distance that the probe travels through a medium other than air. So, when a user changes the phase adjust, they are changing the phase the probe beam has at the point in time that it reaches the sample surface. By changing the phase of the probe beam, we change the phase difference between the probe and reference beam ϕ , and thus change $I_S(t)$.

With no heterodyning, we have:

$$I_S(t) \propto |I_D(t)|^2 \quad (3.4)$$

With heterodyning, we have:

$$I_S(t) \propto |I_D(t) + E_{reference} e^{i\phi}|^2 \quad (3.5)$$

with $E_{reference}$ the magnitude of the electrical field of the reference beam. Of course, (3.5) indicates that the heterodyned signal could take on an entire range of values. As a matter of practical implementation, the phase adjust is slowly changed while observing the signal in the oscilloscope until the signal amplitude appears to be maximized. This can be verified more quantitatively by running a simultaneous Fast Fourier

⁸This formulation simplifies some of the complexity of the true signal. For an overview of this, see [119], and for a thorough physical treatment of the signal see [113].

Transform (FFT) on the signal in the scope and changing the phase adjust until the FFT peak is maximized. Figure 3.6 shows a visual comparison of a heterodyned trace versus a non-heterodyned trace.

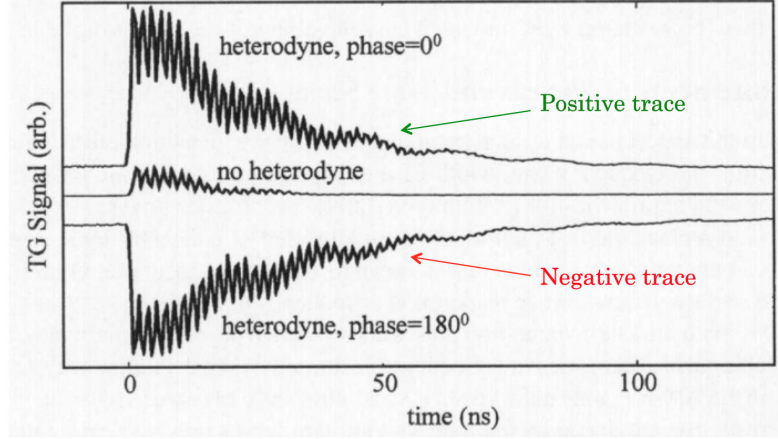


Figure 3.6: The positive and negative traces are amplified by heterodyning. [114]

Positive and negative traces. Figure 3.6 shows both a positive and a negative heterodyned trace. The latter is equivalent to the positive trace, but with the phase adjusted by 180° . In this work, both the positive and negative measurements are collected at each spot. This is because the signals can be subtracted from each other to further increase the amplitude of the signal while subtracting out some of the noise. The total signal amplitude becomes:

$$I_S^{tot}(t) = I_S^{POS}(t) - I_S^{NEG}(t) = 4\sqrt{I_R I_D}(t) \quad (3.6)$$

The dual-heterodyne configuration. Collection of the positive and negative traces allows for a better quality signal, but it also doubles the amount of time it takes to collect the TGS data at a given spot. In order to improve the efficiency of the setup, Cody Dennett built a modified version of the MIT MNM TGS facility that allowed for the simultaneous collection of the positive and negative traces. A schematic of the modified setup is shown in Figure 3.7. [121] describes the successful development of this dual-heterodyne TGS system. The dual-heterodyne system effectively cuts data collection time in half, which was incredibly important to collecting the data presented in this thesis in a reasonable amount of time. Note also that in this updated configuration, a chopper wheel was added that allows for the probe to be operated as a quasi-continuous wave.⁹ In earlier configurations, this quasi-continuous behavior was achieved via direct

⁹There is no reason for the probe to be “on” 100% of the time; it only needs to be continuous for the entire lifetime of the SAW wave that is generated by each pulse of the pump laser. The probe “on” cycle lasts $250\ \mu\text{s}$; the SAW observed in this work decay away completely on the order of $100\ \mu\text{s}$. The pump is pulsed every 1 ms to generate a new SAW.

current modulation, but this is not ideal for long-term operation of the diode. It also may have contributed significantly to signal noise, as it negatively affected the temperature stability of the diode

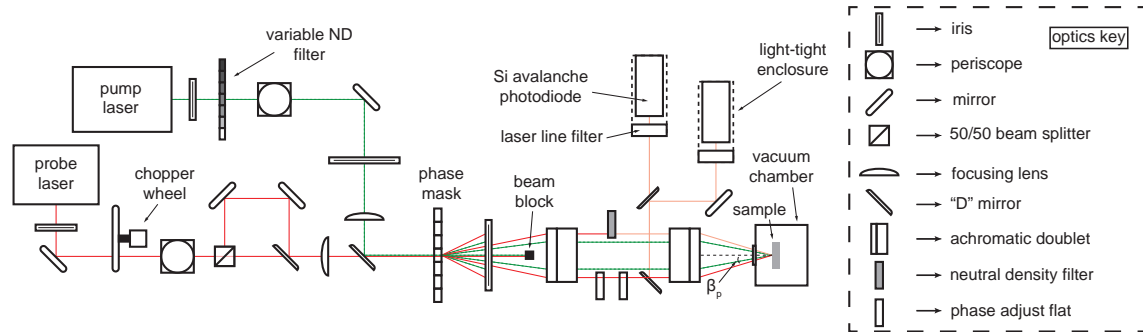


Figure 3.7: The positive and negative heterodyne-amplified traces are collected simultaneously with the new setup, which cuts overall data collection time in half. [123]

Configuration parameters for this work. Table 3.2 gives the setup parameters used to make the measurements presented in this thesis. It can be assumed that these values hold for all presented data unless otherwise noted. (Note that the nominal grating spacing is not the same as the calibrated grating spacing, which is close to the nominal value but which was found to vary slightly (on the order of $0.01\mu\text{m}$) from day to day.

Table 3.2: TGS setup parameters used in this thesis

Parameter	Value
Pump wavelength	532 nm
Pump repetition rate	1 kHz
Pump pulse length	300 ps
Pump pulse energy	$2.5 \mu\text{J}$
Probe power at sample surface	5 mW
Probe laser wavelength	785 nm
Chopper wheel frequency	1 kHz
Probe duty cycle	25%
Pump spot diameter	$140 \mu\text{m}$
Probe spot diameter	$120 \mu\text{m}$
Nominal grating spacing	$5.5 \mu\text{m}$

3.4 TGS signal analysis

A characteristic (positive) TGS trace is shown in Figure 3.8. The “spike” in the early stage of the trace corresponds to the thermal excitation region. The probe beam is diffracting from the oscillating SAW, and so we see these oscillations in the signal. The oscillations decay away as the heat energy imparted by the laser dissipates. Toward the end of the signal, we are only seeing noise. It should be remembered that the signal that appears on the oscilloscope during a measurement is the average of many traces (3 batches of 10,000 traces for all measurements in this thesis). The more traces, the more the noise will be minimized and the cleaner the signal will be. (This, of course, needs to be balanced with the need to make the measurement in a reasonable amount of time. The measurements in this thesis took just over five minutes each.) Figure 3.9 illustrates, at a general level, how material information can be extracted from this signal. The decay of the thermal spike can be correlated to the sample’s thermal diffusivity. The frequency of the signal, easily measured with a FFT, can be correlated to its elastic properties. The decay rate of the oscillation decay rate can be correlated to the acoustic damping of the signal. This is what makes the TGS technique so powerful: information about three independent material properties can be extracted from a single measurement. The calculation of these properties from the signal will be explained in brief here; note also that the full code used to analyze the signal can be found in the appendix of this thesis and on the MIT MNM Github repository.

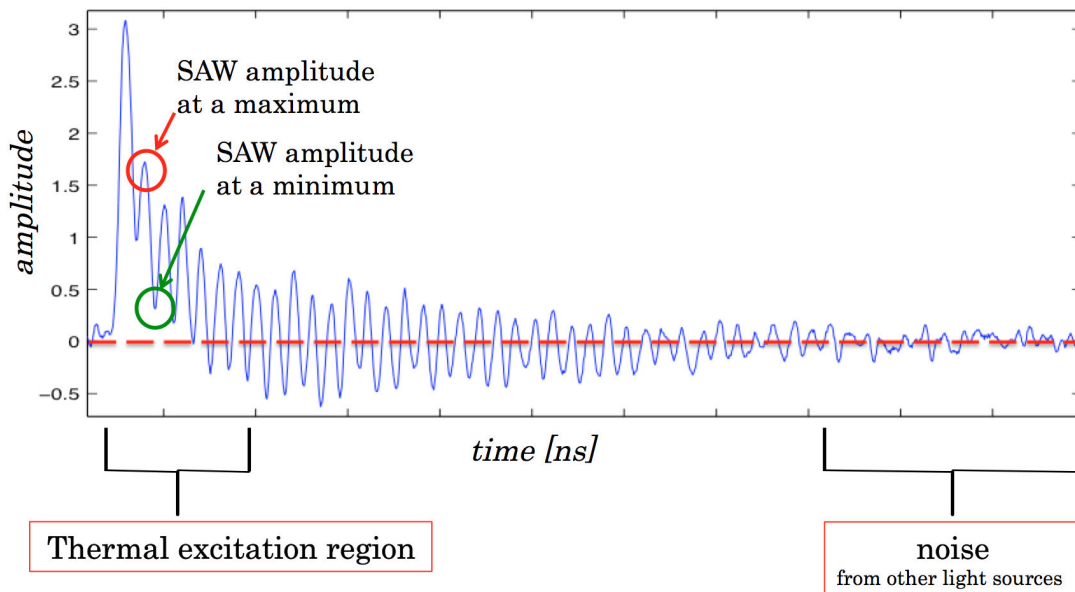


Figure 3.8: A characteristic positive TGS signal.

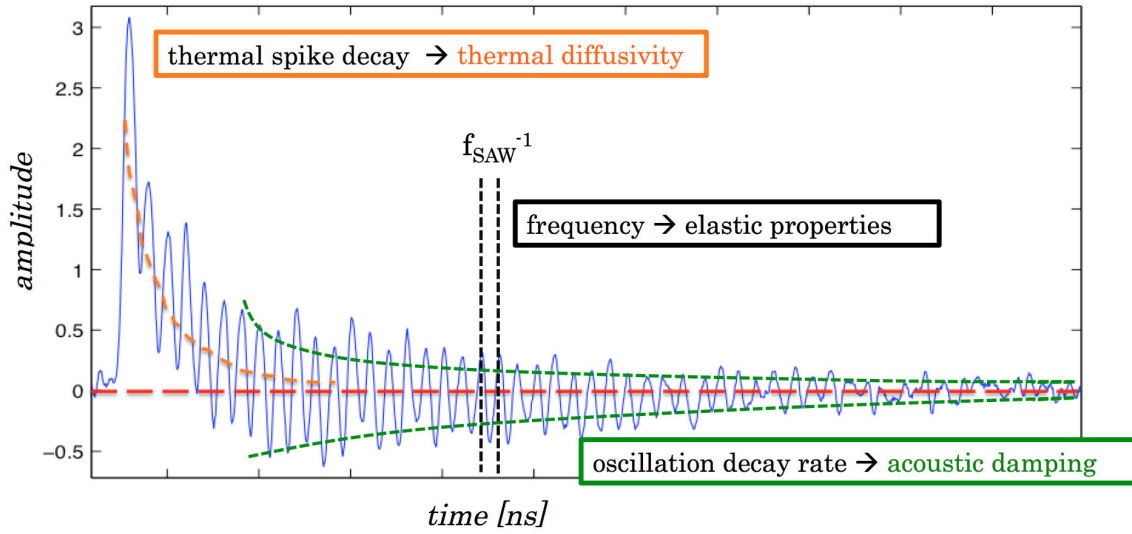


Figure 3.9: Elastic, acoustic, and thermal properties can be extracted from the TGS signal.

SAW speed and elastic properties from the TGS signal. As mentioned previously, the SAW frequency f_{SAW} can be extracted by performing a FFT on the TGS signal, as the oscillations of the diffracted beam correspond to the oscillations of the SAW.

The SAW speed, v_{SAW} , can be calculated from

$$v_{SAW} = f_{SAW}\Lambda \quad (3.7)$$

where Λ is the known wavelength of the diffraction grating.¹⁰ Since the SAW is a Rayleigh wave, v_{SAW} can be correlated to the sample's elastic properties by a polynomial fit to Rayleigh's wavespeed equation that was developed in [122]:

$$v_{SAW} = \sqrt{\frac{E}{2\rho(1+\nu)}}(0.874 + 0.196\nu - 0.043\nu^2 - 0.055\nu^3) \quad (3.8)$$

where E is the Young's modulus, ρ is the sample density, and ν is Poisson's ratio. If ρ and ν are known, E can be calculated. A higher E corresponds to a stiffer material (more resistant to elastic deformation under an applied stress); a lower E corresponds to a more compliant material. As will be shown in subsequent chapters, radiation damage can change E , and measuring v_{SAW} therefore gives us a way to probe radiation-

¹⁰One can use the nominal value from the known location of the spot on the phase mask, but more typically the value of Λ is calibrated on each measurement day by first taking a measurement spot on a single crystal tungsten (W) sample. The measured f_{SAW}^W and the known v_{SAW}^W , 2665.9 m/s, is used to calculate the true Λ .

induced changes in a sample's stiffness.

Thermal diffusivity from the SAW signal. The thermal diffusivity of the samples was calculated using the method developed by Cody Dennett in [123]. (For a complete treatment of this method, and the challenges involved in developing it, see that paper as well as Ch. 3 of [119].)

The equation used to calculate thermal diffusivity from the TGS signal was developed in [123], based on models presented in Kading 1995 transient. It is written as:

$$I_P(t) = A \left[\operatorname{erfc}(q \sqrt{\alpha t}) - \frac{\beta}{\sqrt{t}} \exp(-q^2 \alpha t) \right] + B \sin(2\pi f t + \theta) \exp(-t/\tau) + C \quad (3.9)$$

where A , B and C are constants, f is f_{SAW} , θ is the acoustic phase, and τ is the acoustic damping constant. These are all determined by the thermal fitting codes that can be found in Section 8.2. Based on a model of 1D thermal transport, we know what the form of the decay should look like. We use a nonlinear optimization process to extract the parameters in Equation (3.9) from the recorded data.

There are ways to either fix or pre-compute all the parameters in Equation (3.9) from physical characteristics of the signal that allow for a more robust optimization of their values. Blindly fitting the model will usually lead to a failure to properly fit the thermal diffusivity, but these pre-fitting procedures have improved this process significantly.

Acoustic damping from the SAW signal. There are two types of decay occurring in the signal: the thermal decay - observed most obviously as the thermal spike that decays over the timescale that a given trace is collected, but also observed as part of the oscillation decay (as the thermal energy imparted by the laser dissipates, the SAW decays as well). The oscillation decay is also affected by the acoustic damping that occurs due to energy attenuation of the acoustic wave by the sample.

Acoustic damping information is also extracted from Equation (3.9) as the parameter τ , which has units of time. However, it should be remembered that there is no definitive physical model for acoustic damping of a SAW. There are so many factors that go affect acoustic damping, such as surface preparation, grain boundary scattering, and lossy acoustic modes¹¹ that it is difficult to obtain a definitive physical model for the process. Equation (3.9) represents a phenomenological model of acoustic damping - acoustic damping appears to decay exponentially, and so it is fit to an exponential model - but it is far from fully generalizable to any TGS signal. The acoustic damping data presented in this thesis, therefore, should be considered as a "first pass" at collecting this information from TGS signals. A major part of the future work of adapting TGS for nuclear materials study will be improving the acoustic damping models that allow for accurate extraction of τ from the TGS signal.¹²

¹¹An example of a lossy acoustic mode is a PSAW, or pseudo-SAW, which looks like a SAW but which radiates energy into the bulk. It is not a true Rayleigh wave and it decays faster than a true SAW.

¹²[123] provides illustrated examples that show how the acoustic damping fit fails to properly fit certain TGS signals, and further elaborates on the difficulties that go into accurately fitting τ .

Chapter 4

Transient grating spectroscopy characterization of cold-worked and ion-irradiated single-crystal niobium

4.1 Introduction

TGS was used to characterize changes in single-crystal niobium samples that were first cold-worked, and then progressively irradiated to 3 dpa with silicon ions. The initial purpose of this phase of the project was to demonstrate the ability of TGS to detect changes in radiation damaged materials that occur due to the presence of dislocations. However, as will be explained in this chapter, TGS doesn't always reveal interesting material changes that occur solely as a result of the presence of dislocations. Rather, the real story is in how materials with different initial dislocation populations respond to irradiation.

This chapter will provide an overview of the motivation for cold-working and irradiating the samples. It will then describe the preparation and cold-working of the samples, as well as the extensive characterization of their orientations and dislocation densities that was carried out via numerical method and various X-ray diffraction (XRD) techniques. Next, the irradiation procedure, which was carried out after the XRD characterization was complete on the unirradiated cold-worked samples, will be discussed.

Next, the results of the TGS measurements will be discussed. Each sample was characterized via TGS through 180° of surface rotation in order to capture the full range of TGS response on the single crystals. In Section 4.6, TGS measurements of v_{SAW} as a function of surface rotation are examined closely for changes that arise as a result of increasing dislocation density.

In Section 4.8, the cold-worked samples are irradiated to progressively higher doses in a 5.3 MeV 100 nA Si³⁺ ion beam. During this phase of this study, the TGS characterization was repeated after each pro-

gressive irradiation. Samples were irradiated to 0.01 dpa, 0.03 dpa, 0.1 dpa, 1 dpa, and 3 dpa. $v_{SAW}(\theta)$ for the irradiated samples is considered first.¹ Next, thermal diffusivity and acoustic damping information is extracted from all data sets (including the 0 dpa data from Section 4.6).

Finally, the results will be discussed in context of what they convey about the elastic, thermal, and acoustic changes the samples undergo as a result of cold work and irradiation. These results provide valuable insight into the uses and current limitations of TGS for radiation damage characterization.

4.2 Original motivation for cold working and irradiating the samples

Dislocation formation and movement are a common effect of radiation damage, as detailed in Chapter 2. As a result, this phase of the project began as a "single-effect study" to understand how the presence of dislocations would affect the TGS signal. Single crystal niobium was chosen for this phase of the study because it was available and easily deformable compared to other spare single crystal specimens in lab (its Young's modulus is about 25% that of tungsten) [125]. However, it is also highly anisotropic, and so it was necessary to ensure that the full range of the TGS response was captured by measuring the samples over a range of surface angles.

The niobium samples would be cold-worked to different amounts to induce different levels of dislocation density, as described in the next section, and characterized via TGS. In theory, observed changes in the TGS signal could then be mapped to changes in the dislocation density, and used to predict how changing dislocation densities in a sample during irradiation would change the TGS signal during a real-time *in situ* measurement.

History of the mineral.	1. The mineral was sent with some iron ores to Sir Hans Sloane by Mr. Winthrop of Massachusetts, and there is therefore every reason to believe, that it came from some of the iron mines in that province.
External character.	2. It is heavy, and of a dark grey nearly black; it in some measure has the appearance of the Siberian chromate of iron.
Habitudes with acids.	3. The nitric, muriatic, and sulphuric acids act but very feebly upon the mineral; of these however, sulphuric acid produces the greatest effect, and dissolves some iron.

Figure 4.1: Niobium, originally called Columbium, was discovered at the turn of the 19th century by Charles Hatchett, and is described here in an 1802 journal article. [126]

Following the cold-work-only study, the samples would be irradiated in an ion beam to progressively larger dpa doses. The idea to carry out this phase of the niobium experiment came from a paper published by

¹For the 3 dpa samples, rotationally dependent data (v_{SAW}) was not collected in order to save time and because the thermal diffusivity data was of primary interest by the time they were irradiated. The v_{SAW} data from the 3 dpa samples is therefore not presented in Section 4.6.2, as it was not collected in a comparable way to the other v_{SAW} data sets for niobium (e.g. two points every 5° over at least 180°.)

Dieckamp and Sosin in 1956 [127]. In this paper, the Young's modulus of high-purity copper irradiated with 1 MeV electrons was measured. Three samples were used: sample **B** was well-annealed (low dislocation density), sample **A** was slightly cold-worked, and sample **C** was heavily cold-worked (high dislocation density).

Figure 4.2 shows the fractional change in the Young's modulus of each sample as a function of total flux (the electron beam was stopped periodically during the irradiations in order to perform the measurements of the Young's modulus). It is clear that the level of cold work in the samples affects how the sample responds to radiation (and radiation damage). Interestingly, the heavily cold-worked sample hews relatively closely to the behavior of the annealed sample. The lightly cold-worked sample, however, exhibits a significantly larger fractional change in Young's modulus (and a larger rate-of-response to the radiation at lower doses).

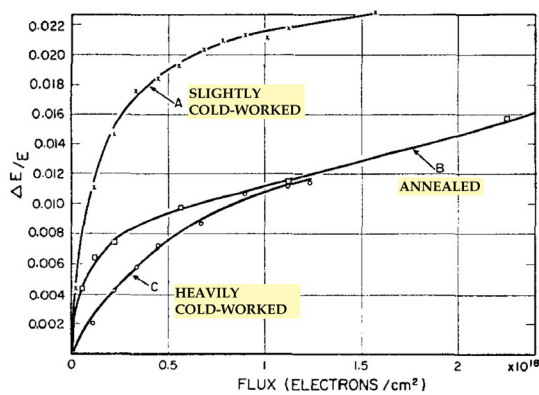


Figure 4.2: Modified image from [127] showing experimentally measured changes in Young's modulus for three electron-irradiated copper samples with varying levels of cold work.

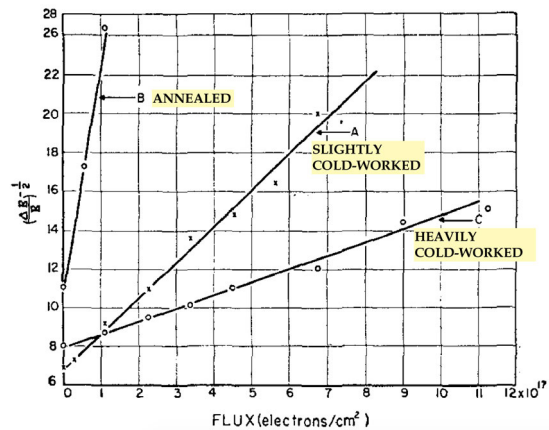


Figure 4.3: Data from Figure 4.2 replotted (open circles and xs) so it can be compared with the theoretical prediction of $(\Delta E/E)^{-1/2}$ as a function of integrated flux (straight line). (See p.1417 of [127].)

The fractional change in Young's modulus is related to dislocations as:

$$\frac{\Delta E}{E} = -\frac{1}{6}Nl^2 \quad (4.1)$$

with N the dislocation density and l the dislocation length (average).² The authors assume that the likelihood a defect will be pinned by a dislocation per length [cm] of dislocation is uniform, and can thus be described by a constant effective cross section k (where k describes the likelihood that a defect that will be pinned per cm length of dislocation). This assumption is used to derive the equation that fits the data plotted in Figure 4.3:

²This is identical to the equation given in [128], except with a defined constant.

$$\left(\frac{\Delta E}{E}\right)^{-1/2} = \left(\frac{N}{6}\right)^{-1/2} \left(k\phi + \frac{1}{l_0}\right) \quad (4.2)$$

where l_0 is the initial average dislocation length when flux $\phi=0$ Figure 4.3 shows the same data from Figure 4.2 replotted according to Equation (4.2). The authors assume a value for k by assuming that dislocations will be pinned by any defects that are present within three atoms of the dislocation, but note that their assumed value leads to unusually low values of N . Using measured values of N might lead to a better value of k for a specific material.

It is also shown in [127] that sustained irradiation of the copper samples leads to a steady decline in $\Delta E/E$, in contrast to the increases in Young's modulus measured at lower radiation doses. This behavior begins at a total flux of about 10^{14} electrons per cm^2 , as shown in Figure 4.4. The authors postulate that this is due to a bulk effect caused by distortion of the lattice by vacancies and interstitials, with dislocations having a negligible effect.




FIG. 3. Extended irradiation of 99.999+% copper followed by subsequent thermal annealing. Specimen was maintained at temperatures of 100°C or higher for 5-min intervals.

Figure 4.4: A gradual decrease in modulus is observed with extended irradiation. This behavior is not immediately obvious if only studying material response at low dose. [127]

The results of this paper indicate that we should expect different response behaviors from niobium samples that were subject to the same amount of radiation but different amounts of cold work. f_{SAW} (and v_{SAW}) is directly linked to the samples' elastic moduli. The authors of [127] used a driven resonant frequency method to measure Young's modulus, whereas this work will use TGS to measure v_{SAW} (an indicator of E) over 180° of surface rotation. TGS will also be used to collect thermal diffusivity data and acoustic damping data concurrently with the v_{SAW} data. Thermal diffusivity and acoustic damping are expected to vary as the radiation-induced defect population changes as well.

4.3 Preparation and characterization of the single-crystal niobium

This section first describes how the single-crystal niobium samples were prepared and cold-worked. Next, the extensive characterization of these samples via various XRD methods is discussed. Finally, the irradiation parameters used to induce real radiation damage in the cold-worked samples are described.

4.4 Cold-working the single-crystal niobium

The niobium samples to be cold-worked were cut from a ≈ 5 mm diameter rod of single crystal niobium of unidentified orientation. The rod was cut into 1 mm thick circular slices using a low-speed Buehler® diamond saw lubricated with IsoCut™ fluid. The rod cut easily, and so the sawblade was kept on low rotation to avoid deforming the samples. Samples were stored in individual sample boxes that could be labeled to avoid mixing up samples, and kept on double-sided tape post-polishing to avoid scratches.




Figure 4.5: A typical niobium sample, prior to cold working. Samples are sectioned to be ≈ 1 mm thick from a 5 mm diameter cylindrical single crystal. This sample is marked with permanent marker for orientation purposes (this can be easily removed with acetone or another solvent).

The initial cold working plan was to prescribe a "percent compression" for each sample, ranging from 2 to 10%, and then use calipers to measure the thickness of each sample before and after compression. With no good way to measure sample thickness while being pressed in the lab pellet press, however, it became

obvious that it made far more sense to simply prescribe an amount of force with which to press each sample. At this point, what mattered was producing samples with a range of dislocation densities - not samples with specific dislocation densities.

Calipers were used to measure the thickness of the cut niobium samples at various points around its circumference. Six samples were chosen that had the most uniform thickness, as there were variations in thickness in certain samples as a result of miscut during sectioning.




Figure 4.6: The Carver™ pellet press used to cold work the niobium samples. The lever is used to move the stage and apply force to the die.



Figure 4.7: A Carver™ 13mm ID pellet die was used for niobium cold working. The stage applies force to the die's pushing rod and core die .



Figure 4.8: A schematic showing the interior of the die. The Nb samples were centered in the inner faces of the interior steel pushing rod (top) and core die (bottom).

Samples were pressed using a 13mm inner diameter (ID) pellet press from Carver™. Samples were centered one at a time between the flat faces of the core die and the pushing rod. The samples were then pressed with different amounts of force. In total, six samples were used in the cold work matrix. One was kept unpressed as a control sample. The other five were pressed to 1000, 1500, 2000, 2500, and 3000 pounds. At 3000 pounds, the sample was visibly flattened, and it seemed likely that there would be a sufficient range of dislocation behavior in these six samples to execute the rest of the study.

Figures 4.6-4.8 depict the pellet press and the die used to apply the cold working compression force to the niobium samples.

The force levels applied by the pellet press became the de facto way to refer to the samples. For the remainder of this chapter, the “1000 lb” sample refers to the sample that was pressed to “1000 lbs” force according to the dial on the pellet press, and so on. The unpressed sample is referred to as the control sample.

4.4.1 Polishing the single crystal cold-worked niobium samples

It was important to polish the samples without inducing significant cold work, and due to the ductility of single crystal pure niobium, this step was not nearly as straightforward as it is with harder metals. Polishing was first attempted using a Jeol cross section polisher. This instrument uses argon ions to create a polished surface with extremely minimal cold work, but is best suited for examination of narrow cross sections and not a flat sample surface. Since it was not necessary to polish the entire sample surface, but just an area large enough for a TGS measurement to be made, parameters were varied to prevent argon ions from cutting all the way through the niobium sample. However, when exposed to argon ions long enough to get the sufficient mirror polish, the niobium would take on a hemispherical, dished appearance in the polished area. Because a flat surface was required for TGS, this strategy was abandoned. Figure 4.9 shows an example of a niobium sample with this dished appearance.




Figure 4.9: A circular niobium sample, photographed through the lens of an optical microscope, is mounted on a metallic block. The area polished with argon ions is visible in the upper half of the niobium - it is a dish-shaped area where the argon ions have begun to erode through the sample. It was very difficult to get a mirror polish surface without also inducing curvature, so the idea of adapting the cross-section polisher for small surface area polishing was abandoned.

This necessitated trying another method. Previous attempts to polish single crystal niobium with a spinning polisher powerhead had proved problematic, because at finer polishing grits it was difficult to avoid a slight shearing of the surface layers of niobium that led to a “wavy” surface appearance.

Instead of using the powerhead on the polisher, the niobium samples were polished using a lapping fixture and Buehler® MasterPrep™ alumina suspension. This ensured that minimal force was applied to the samples, because the sample rotated freely within the lapping fixture with only the weight of the lapping fixture sample holder pressing on it. Normally, samples would be polished using sandpapers before using MasterPrep™, but niobium is sufficiently soft, and the need to avoid additional cold work sufficiently impor-

tant, that these steps were foregone. The downside is that this procedure is very time consuming, because the MasterPrep™ particles are so fine and the weight on top of the sample so low. Furthermore, only one sample at a time could be fit onto the lapping fixture. However, this method was successful in avoiding the surface deformation experienced with the powerhead.

It was later found that the samples were still not as reflective as they needed to be for successful TGS measurements.³ Samples were repolished using diamond paper laid on top of a perfectly flat surface. A single gloved finger was used to gently move the sample in circles until it has an even polish. The samples were then rinsed in a sonicator using deionized water before polishing with the next-finest diamond paper grit. The diamond paper used in these experiments was purchased from South Bay Technologies™ in the grit sizes 30, 15, 9, 6, 3, 1, 0.5, and 0.1 μm. Figure 4.10 shows metallographs of polished samples.

For any future TGS experiments using single crystal niobium, it is recommended that an attempt is made to polish the samples using only the diamond papers. If this is sufficient, a significant amount of preparation time can be saved by avoiding the lapping fixture step.

Note that all polishing was carried out after the samples had been cold worked.




Figure 4.10: Selected metallographs of the single crystal niobium polish after MasterPrep polishing, with bright field images denoted by BF and dark field images denoted by DF. The lefthand images of the control sample show a mostly smooth surface that nevertheless has uniform, if shallow, roughness (especially visible in the dark field view). The sample on the right has been cold-worked, and striations from the sectioning procedure are still visible. The dark field image shows a more noticeable degree of surface roughness (as well as some residue from the Masterprep®). The bottom edge of the metallograph covers an ≈3 mm distance across each sample. These images show that even with extensive polishing using the colloidal silica and the lapping fixture, the samples are not perfectly smooth, and it will be hard to obtain a clear TGS signal on many TGS setups, hence the diamond paper recommendation.

³During the first set of TGS measurements - prior to irradiating the samples - the TGS setup was still in one of its earlier configurations. The chopper wheel had not been installed, and the probe laser was modulated differently. Most importantly, it was discovered that the setup alignment was not well optimized. Since making these modifications and realigning the setup, surface preparation has not needed to be so exacting, and it has been possible to obtain clear TGS signals even with moderate levels of imperfection in the surface polish of a sample.

4.5 X-ray diffraction (XRD) characterization of the cold-worked niobium samples

The single crystal niobium rod from which the samples were cut did not have any accompanying information to identify its orientation, and so it was necessary to determine this, since the TGS response of a single crystal is orientation-dependent. Samples were first characterized using a General Area Detector Diffraction System (GADDS) to determine their orientation and observe, qualitatively, how the dislocation density changed the diffraction spot. Samples were then measured using High-Resolution X-ray Diffraction (HRXRD) to obtain rocking curves. These curves are known to broaden with increasing dislocation density, and so can be used to estimate actual dislocation density in the samples. Finally, pole figures were constructed for each sample following extended scans in the GADDS system. The pole figures can be used to understand microstructural changes due to defect populations and sample orientation. They can also be used to obtain a second construction of the rocking curves, in order to obtain a measurement of dislocation density via a second method.

4.5.1 Basic principles of X-ray diffraction

X-ray diffraction (XRD) is a broad category of experimental techniques that utilize principles of wave interference to probe the crystal structure of materials. X-rays have a wavelength of 0.1 to 10 nm, and as such, their wavelength is on the order of the characteristic spacing distance of atom planes in materials. When X-rays are incident on a material, they are scattered elastically from the lattice planes, and undergo constructive and destructive interference as they interact. Bragg's law is used to describe the condition for constructive interference:

$$2d\sin\theta = n\lambda \quad (4.3)$$

where d is the lattice spacing, θ is the scattering angle, n is an integer, and λ is the wavelength of the incident wave [129]. Figure 4.11 shows the relationship between a crystal lattice, incident and scattered X-ray beams, and the Bragg angle. The XRD systems used in this research utilize the principles of Bragg's law to probe the crystal structure of a sample.

In XRD analysis, a collimated monochromatic X-ray beam is directed at a sample of interest, and the various angles between the beam, sample, and X-ray detector are adjusted until diffraction spots (single crystal material) or patterns (polycrystal material) are maximized. The relative positions of the X-ray source, the sample, and the X-ray detector are shown in Figure 4.12. The properties of the diffraction spot or pattern, combined with the known geometry of the setup and the known wavelength of the X-ray beam, can be used to study the sample's crystal structure.

For example, when examining a single crystal, there is a single diffraction spot to maximize in intensity by adjusting the positions of the detector and the sample. When this has been done, θ in Figure 4.12 is assumed to satisfy the Bragg condition. The lattice parameter of the sample can then be determined (and

a crystal orientation identified). In most modern XRD systems, the sample is on a remotely controlled goniometer cradle that has many degrees of freedom, shown schematically in Figure 4.13. This is useful for determining quantities like miscut (offset of lattice planes parallel to the surface from an expected orientation) by maximizing θ and then adjusting the goniometer in ϕ and χ to further maximize the spot intensity. Figure 5.23 shows the MIT GADDS system in use.




Figure 4.11: Constructive interference occurs when the Bragg condition is satisfied. By maximizing the intensity of a diffraction spot or pattern created with monochromatic X-rays, various information about a sample's crystal structure can be determined.




Figure 4.12: The arrangement of the X-ray source, the sample, and the X-ray detector in a typical XRD system used in this research. When the diffraction spot intensity is maximized, θ satisfies the Bragg condition, and characteristics of the crystal structure can be determined.

XRD was used to characterize the orientation of the six niobium samples, as well as to study the dislocation density in each sample. Dislocations result in a broadening of characteristic X-ray peaks, such as a rocking curve from a single crystal or the a broadening of the reflections obtained from a powdered sample [130, 131, 133, 134].

4.5.2 XRD characterization of the cold-worked niobium samples

Orientation was characterized on a Bruker® D8 General Area Detector Diffraction System (GADDS). Because the samples were already so thin, it was decided that miscut and misalignment would not be corrected since there was no good way to do this without grinding away the samples during the re-angling and re-polishing steps this would require. The initial characterization was instead used to orient the single crystal samples and determine the ϕ value at which the diffraction spots were brightest. This spot was then marked as the "up" orientation, and was used as the 0 degree point for the TGS measurements.

Because these were single crystals, the goal of the orientation was to maximize the brightness (and the centeredness) of the singular diffraction spot in order to determine the ideal orientation of the sample. However, as cold work increased, it was observed that faint diffraction rings were visible in addition to the diffraction spots. This phenomenon is associated with polycrystallinity. The diffraction rings indicate that the cold work had induced some sort of change in the single crystal lattice.

The generator of the XRD system was set to 40 kV and 40 mA. The working distance was 16 cm, and the sample was adjusted on the measurement stage so that the combined height of the sample and the stage




Figure 4.13: A schematic of a typical goniometer. The sample can be moved up (y), side-to-side (x), back-and-forth (z), rotated about z (ϕ), rocked along Ψ , and rotated relative to the detector (ω). Image courtesy of MIT CMSE.




Figure 4.14: The Bruker® system at MIT. The X-ray detector is on the left. The collimated X-ray beam is on the right. In the middle, the sample is mounted to a cradle that can be moved in six degrees of freedom. The detector and the cradle can be moved relative to each other and relative to the X-ray beam.

was 4 cm. The peak list for niobium, as given in the crystallography database accessed through HighScore Plus, is given in Table 4.1.

4.5.2.1 Control sample

The control sample was found to be relatively well-aligned. The conical cursor in the GADDS program is used to draw an arc, along which all points have the same 2θ value. Typically the conical cursor is used to investigate the properties of a Debye diffraction ring - but in a single crystal, only a diffraction spot is observed. The spot's intensity is maximized when the system is ideally aligned relative to the crystal planes.

in Figures 4.15 and 4.16, the conical cursor (the thin white arc in the frames) is used to study the diffraction spot. Figure 4.15 shows a frame from a coupled scan taken early in the alignment procedure. The diffraction spot is slightly distorted along the same 2θ ring, indicating some degree of imperfection in

Table 4.1: Crystallography peak list for pure niobium

Nº	h	k	l	d [Å]	2θ
1	0	1	1	2.33345	38.551
2	0	0	2	1.65000	55.660
3	1	1	2	1.34722	69.748
4	0	2	2	1.16673	82.634




Figure 4.15: A frame from a coupled scan of the control sample. The diffraction spot is nearly centered ($\chi \approx -90^\circ$) but is somewhat stretched out, indicating some distortion or imperfection in the lattice.




Figure 4.16: A frame from a psi scan of the control sample, taken later in the orientation procedure. The diffraction spot has become more defined as alignment improves, but it still shows a small amount of distortion.

the crystal lattice and some range of orientations of the niobium lattice. When the position marker of the conical cursor is centered, the **Chi** (χ) value is close to -90° : for ideal alignment, the center of the diffraction spot would be at the center of the frame and correspond to a χ value of -90° . The Angstrom value shown in this frame - 2.3293Å - is close to the value associated with the $\langle 011 \rangle$ plane given in Table 4.1, 2.3345Å.

Figure 4.16 shows a frame from a Ψ scan later in the alignment procedure. The diffraction spot is less spread out now that the alignment is more refined, but some distortion is still evident. The conical cursor indicates similar characteristics to those found in the coupled scan frame. In Figure 4.17, the intensity of the diffraction spot in Figure 4.16 is integrated over χ , showing a well-defined peak centered at about -92° . This indicates that a vector normal to the control sample surface is about 2° off in χ from the $\langle 011 \rangle$ aligned crystal planes.

4.5.2.2 1000 lb sample

The diffraction spot of the 1000 lb sample also exhibited lengthening along a 2θ arc. As with the scans of the control sample, the **2-Theta** and **Angstroms** values indicate that the data should be interpreted relative to the $\langle 011 \rangle$. An integration of intensity of the coupled scan shows a peak around -118° , indicating that the




Figure 4.17: An integration of the intensity across diffraction spot in Figure 4.16.

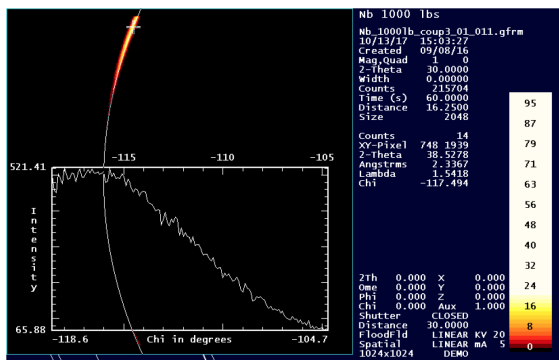


Figure 4.18: A frame from a coupled scan of the 1000 lb sample. The diffraction spot is off-center ($\chi \approx -90^\circ$) and stretched, indicating some distortion or imperfection in the lattice and a significant amount of miscut.

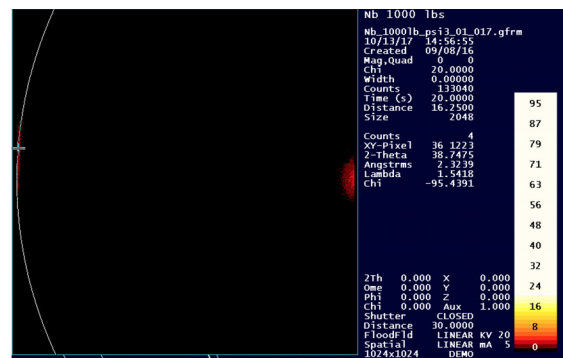


Figure 4.19: A frame from a psi scan of the 1000 lb sample. The diffraction spot is more centered in this frame, but is faint and distorted.

$\langle 110 \rangle$ plane is nearly 30° offset in χ from the surface normal. The ϕ scan in Figure 4.20 suggests a similar offset.

The scans of the 1000 lb sample indicate that the $\langle 110 \rangle$ planes are nearly parallel to the sample surface - there doesn't appear to be much miscut, as evidenced from the χ value shown in Figure 4.21. The diffraction spot is spread out along the optimized 2θ practically across the entire visible frame, indicating even more distortion of the lattice. This makes sense, as cold work has increased and a more significant distortion of the lattice is expected.

Integrating the intensity of the diffraction spot from the coupled scan shows that there are actually two distinct intensity peaks on either side of -90° . This could also be interpreted as a broad peak centered at -90° , and some effect of the cold work has resulted in distortion of the lattice that led to a warping of the previously symmetric peak.




Figure 4.20: A frame from one of the ϕ scans performed on the 1000 lb sample. The diffraction spot did not come into the center of the frame during these scans.

4.5.2.3 2000 lb sample

The 2000 lb sample exhibited diffraction behavior similar to that of the 1500 lb sample. The diffraction "spot" of the coupled scan, shown in Figure 4.23, lay entirely along one value of 2θ (within 0.1° of the theoretical 2θ value for the $\langle 011 \rangle$ orientation). The spot's intensity close to $\chi = -90^\circ$, indicating relatively little miscut, as shown in the integration in Figure 4.24.

The diffraction spot is spread along the visible 2θ arc, indicating some degree of lattice distortion. If this was not a single crystal, and instead was a niobium crystal composed of perfect grains that were identical except for orientation, with grains oriented in all directions, the "spot" would be a ring of equal intensity at all values of 2θ . It appears that increasing cold work increases distortions in the lattice that result in diffraction to an increasingly broad range of 2θ values.

4.5.2.4 2500 lb sample

Figure 4.25 shows the coupled scan for the 2500 lb sample. The spot is well-centered in the frame, as confirmed by the intensity integration in Figure 4.26. However, the **2-Theta** and **Angstroms** (lattice d-spacing) values displayed in the GADDS program no longer correspond to those that indicate a $\langle 011 \rangle$ orientation.

Table 4.1 confirms that the 2500 lb sample surface is oriented in the $\langle 112 \rangle$ direction. When the sample is well-aligned in the instrument (the diffraction spot intensity is centered at $\chi = -90^\circ$), as is the case in the frame from the coupled scan shown here, 2θ is 69.585° and the d-spacing is 1.3411 \AA . These agree well with the given values for $\langle 112 \rangle$ in Table 4.1, 1.34722 \AA and 69.748° .

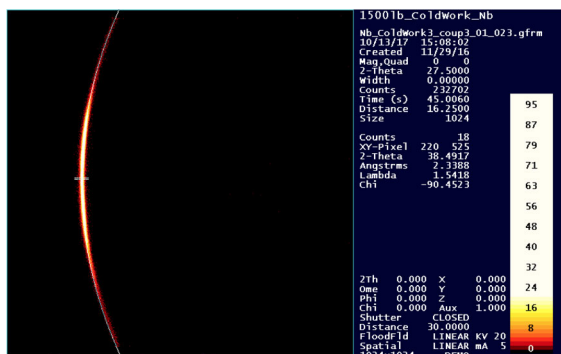


Figure 4.21: A frame from a coupled scan of the 1500 lb sample. The diffraction spot appears close to the center of the frame ($\chi \approx -90^\circ$) and is stretched more obviously than the previous two samples, indicating more distortion or imperfection in the lattice but not much miscut.

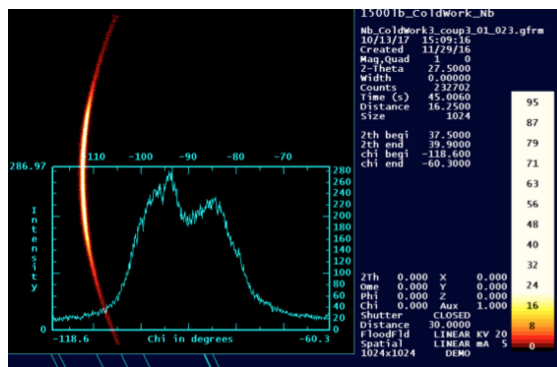


Figure 4.22: Integrating a frame from the coupled scan shows two distinct regions of intensity on either side of $\chi=90^\circ$.

4.5.2.5 3000 lb sample

The coupled scans from the 3000 lb sample show a single diffraction ring, instead of a more constrained spot, indicating a significant amount of lattice distortion has occurred. There isn't a clear maximum in intensity, in contrast to the coupled scans of the second-most cold-worked sample (see 2500 lb sample, Figure 4.25). Instead, there is an intensity peak at about $\chi = -120^\circ$ and -75° (although it should be noted that these peaks aren't sharply defined with respect to other χ values). The integration here was centered around a 2θ value of 38.5° , indicating that the primary "orientations" here can be interpreted relative to a $\langle 011 \rangle$ surface orientation: the lattice is primarily oriented about 30° from this surface in χ in the negative direction, and about 15° from this surface in χ in the positive direction.

It is possible that the 3000 lb sample has some sort of low-angle grain boundary, and isn't truly a single crystal anymore. However, attempts to confirm this via electron backscatter diffraction (EBSD) were unsuccessful due to the relatively rough surface preparation of the samples.⁴

⁴EBSD was attempted during the TGS measurement campaign, so the surface was not repolished to avoid inducing any new changes into the sample surface once TGS measurements began.




Figure 4.23: A frame from a coupled scan of the 2000 lb sample. The diffraction spot appears close to the center of the frame ($\chi \approx -90^\circ$) and exhibits a level of distortion similar to that of the 1500 lb sample.

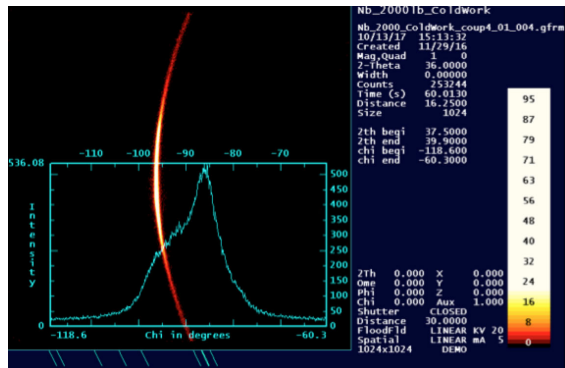


Figure 4.24: Integrating a frame from the coupled scan shows an asymmetric intensity peak: it is possible that this asymmetry is due to non-isotropic distortion of the sample. χ is about 4° off from -90° .




Figure 4.25: A frame from a coupled scan of the 2400 lb sample. The diffraction spot appears close to the center of the frame ($\chi \approx -90^\circ$) but note that the 2θ value is not optimized for the $\langle 011 \rangle$ orientation.

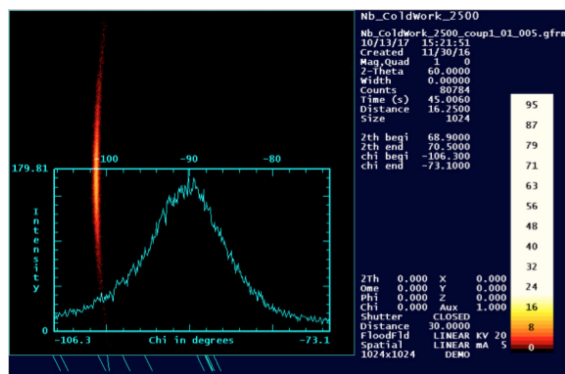


Figure 4.26: Integrating a frame from the coupled scan confirms that the sample is well-aligned to the reference orientation.

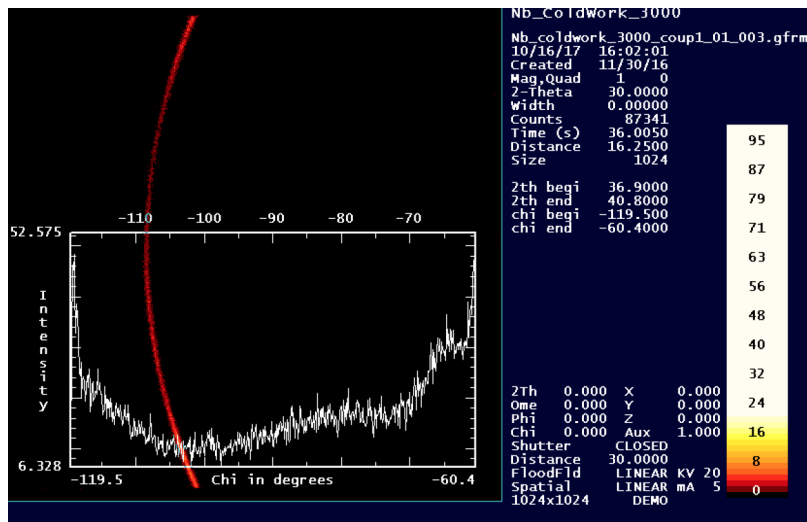


Figure 4.27: A frame from a coupled scan of the 3000 lb sample. The diffraction “spot” stretches along the arc defined by a single 2θ value, with varying intensity. Integration of intensity as a function of χ shows two main intensity peaks. One is about 30° off from $\chi = -90^\circ$, while the other is about 20° off.

4.5.2.7 Summary of XRD characterization of the niobium sample orientations

Table 4.2 summarizes the results of the XRD orientation characterization of the niobium samples. The control, 1500, and 2000 lb sample are well aligned to the $\langle 011 \rangle$ plane. The 1500, 2500, and 3000 lb samples appeared to be more closely aligned to the $\langle 112 \rangle$ plane.

Table 4.2: Summary of XRD characterization of cold-worked niobium sample orientations

Sample	Orientation
Control	2° off in χ from $\langle 011 \rangle$ plane ($\chi \approx 92^\circ$)
1000 lb	30° off in χ from $\langle 011 \rangle$ plane
1500 lb	Aligned to $\langle 011 \rangle$ plane, although there appear to be two regions of intensity (each about 1-2° off in χ from $\langle 011 \rangle$)
2000 lb	2° off in χ from $\langle 011 \rangle$ plane ($\chi \approx -88^\circ$)
2500 lb	Aligned to $\langle 112 \rangle$ plane
3000 lb	Two regions of intensity, one about 30° off from $\langle 011 \rangle$ ($\chi \approx -120^\circ$) and one about 15° off from $\langle 011 \rangle$ in the opposite direction ($\chi \approx -75^\circ$)

In hindsight, the ϕ -maximization could probably have been skipped. The orientation characterization is a time consuming step, and the TGS rotated data is often shifted in θ during the data analysis steps, as will be shown in subsequent sections. There is not an absolute need to predetermine which ϕ to use as the zero-point for TGS measurements if the sample will be rotated between measurements - data sets from different samples can simply be aligned after the measurements are taken. If working with anisotropic samples, however, and planning to only translate the sample in x and y between measurements, this step is necessary. Otherwise, the measurements between different samples will not be truly comparable. That being said, it was still very useful to get an idea of the sample orientation and to observe the changes in the diffraction spot that occurred with increasing cold work.

4.5.3 Pole figures of the cold-worked niobium

Pole figures were collected for each of the six niobium samples used in this project. This allowed for a more direct visual comparison of the orientations of each sample, as well as any textural changes incurred by the cold work. The pole figures were created using the Bruker® D8 GADDS system. The 2θ value is held within a specific range, typically with θ , the expected Bragg angle from literature or previous sample measurements, at the center of that range. The 2D area detector means that X-rays diffracting at a range of θ values will be collected. The detector is kept stationary and the sample is tilted along the Ψ axis (here, it was tilted to 25, 50, and 75°. At each tilt condition, the sample is rotated in ϕ in 6° increments for 360°. At every position, the area detector collects light scattering at near θ . These frames are then used to create the pole figures.

Figures 4.34-4.39 shows the results of the pole figure measurements for each niobium sample. The (200) pole figures further confirm that the control, 1500, and 2000 lb samples (two “spots”) effectively have a different orientation than the 1000, 2500, and 3000 lb samples (three “spots”). For clarity, the pole figures are grouped by orientation. On a qualitative level, it also appears clear that the increasing cold work leads to significant broadening of diffraction spots (for example, compare the 1000 lb and the 3000 lb samples).

Orientation distribution functions for each sample can be found in the Appendix, beginning on page 220.




Figure 4.34: Pole figures for the control sample




Figure 4.35: Pole figures for the 1500 lb sample




Figure 4.36: Pole figures for the 2000 lb sample




Figure 4.37: Pole figures for the 1000 lb sample




Figure 4.38: Pole figures for the 2500 lb sample




Figure 4.39: Pole figures for the 3000 lb sample

4.5.4 Characterizing dislocation density via HRXRD rocking curves

The primary goal of the HRXRD rocking curve measurements was to establish relative differences in dislocation densities between the samples in order to effectively compare observed changes in the TGS signals. It was necessary to characterize the dislocation density via a nondestructive method. Direct-dislocation-counting methods typically involve etching or TEM, and the samples needed to stay intact for TGS measurements. HRXRD rocking curve measurements allow a non-contact determination of dislocation density, because the broadening of the curves can be correlated to a dislocation density.

To collect a rocking curve from a sample, the Bragg peak is optimized in intensity by “rocking” the sample in ω about the location of the Bragg peak until its height is maximized. To characterize dislocations in the niobium samples, a Bruker™ D8 High-Resolution X-Ray Diffractometer (HRXRD) was used. The HRXRD was equipped with a Ge(022) incident beam monochromator. The HRXRD is ideal for obtaining high resolution Bragg peaks of a single crystal.




Figure 4.40: HRXRD rocking curves for the control sample, 1500 lb sample, and 2000 lb sample. There is a significant difference between the control sample and the other two samples. Based on these curves, 1500 lb and 2000 lb likely have a very similar dislocation density.

The rocking curves obtained for each of the six samples are shown in Figure 4.40 and 4.41. As described in the previous section, the rocking curves broaden when there are defects in the crystal lattice. The control sample, which underwent no additional cold work, has the sharpest, most intense rocking curve. The other samples have much broader curves with lower peak intensities. The breadth of the curve generally seems to increase with cold work, with the 3000 lb sample having an extremely low, broad rocking curve in comparison to the other five samples.

Table 4.3 gives the properties of the measured rocking curves for the six samples. The software used to analyze the rocking curves was PANalytical™ HighScore Plus. In the analysis, the halves of the peaks could be split (e.g. different full-width half-maximum (FWHM) for the left and right side of the rocking curve) but not the shape (i.e. both halves were analyzed using the same function fit).

The rocking curve shape can be correlated to dislocation density using a method outlined in [135] and based on the work of [136] and [137]. In this method, the density of dislocations is given as

$$N = \frac{\alpha_{\Omega}^2}{4.35b^2} \quad (4.4)$$

where b is the Burger’s vector of the dislocation and α_{Ω} is calculated as




Figure 4.41: The 1000 lb, 2000 lb, and 3000 lb samples are plotted separately, as they have a different orientation than the other samples. Note that the 1000 lb sample - the most intense peak here - is only about 20% the height of the control sample rocking curve.

Table 4.3: HRXRD rocking curves of the cold-worked single crystal Nb samples

Sample	ω [°]	Height [cts]	FWHM Left [°]	FWHM Right [°]	Shape Left	Shape Right	Integral Breadth [°]	Microstrain only [%]
Control	12.6846	13648.77	0.6833	0.9632	0.818	0.818	1.155860	4.39333
1000 lbs	13.2279	1232.93	2.2663	7.8434	0.792	0.792	5.026103	18.91530
1500 lbs	24.9251	573.05	2.6838	6.5047	0.352	0.352	4.563269	9.00986
2000 lbs	16.8881	271.81	3.5421	5.4028	0.652	0.652	5.767712	16.93718
2500 lbs	19.4578	303.75	9.2429	3.8707	0.694	0.694	7.480708	18.89529
3000 lbs	12.2043	79.68	15.5890	26.9302	-0.1	-0.1	20.363760	83.12192




Figure 4.42: Dislocation density in the cold-worked niobium samples, as estimated from Equations (4.5) and (4.4). A steady buildup of dislocations is observed until 1500 lbs of cold work, at which point a saturation seems to be reached. At 2500 lbs of cold work, dislocation density increases again.

$$\alpha_{\Omega} = \beta_{\Omega}[0.184446 + 0.812692(1 - 0.998497\eta)^{1/2} - 0.659603\eta + 0.44554\eta^2] \quad (4.5)$$

with η the shape factor and β_{Ω} the integral breadth of rocking curve.

Equations (4.4) and (4.5) were used to estimate the dislocation density in each single crystal niobium sample. The β_{Ω} term was taken to be equivalent to the calculated integral breadth and η was taken to be equivalent to the shape left factor as determined for each measured rocking curve in HighScore Plus. The Burger's vector was calculated in the typical manner for a BCC metal as $\sqrt{3}/2 \cdot a_0$, with a_0 the lattice parameter of niobium.

The results of calculating N using Equations (4.4)-(4.5) and the values in Table 4.3 are graphed in Figure 4.42. This shows a steady buildup of dislocations as cold work increases, reaches a saturation point at which dislocation density is not significantly increasing, and then begins to increase again. The dislocation densities reported here are reasonable based on the values reported in the literature for other cold worked metals [217, 139].

In order to verify the parameters given in Table 4.3, the pole figures were discretized and intensity points along a constant ω value were extracted. Each vector of intensity points along a constant ω became a column in a single matrix. The rows of the matrix were then plotted to form a series of rocking curves (ω versus intensity), with the rocking curve of highest intensity being the representative rocking curve for a given sample. These rocking curves were shown to have β_{Ω} values similar to those collected via HRXRD. The pole figures were collected on another instrument and at a different time than the HRXRD rocking curves were collected. Therefore, this provides a good validation that the rocking curves shown in Figures 4.40

and 4.41 are truly representative of the samples, and the dislocation density estimates in Figure 4.42 are reasonable.

4.5.5 Conclusions of XRD characterization campaign

The key observations of the XRD characterization campaign are as follows:

- The cold-worked niobium samples should be treated as two groups of samples: one group (including the control) is aligned to the $\langle 011 \rangle$ direction, and the other group (1000, 2500, and 3000 lb) is aligned to a different orientation - possibly $\langle 112 \rangle$. As such, TGS results should be compared within each subgroup. As they are single crystals, their orientation is expected to have a significant impact on TGS response.
- The impact of increasing cold work can be clearly seen by the texture change revealed in the pole figures.
- The dislocation density increases with increased force applied to the samples.⁵
- Within the $\langle 011 \rangle$ group, the 1500 lb and 2000 lb samples have similar concentrations of dislocations, with the 2000 lb sample having slightly more.

4.6 Transient grating spectroscopy of the cold-worked single-crystal niobium samples: $v_{SAW}(\theta)$

This section describes the experimental measurements of v_{SAW} as a function of surface angle on the niobium samples. First, aspects of the TGS measurement methods specific to this phase of the project are described. Second, results from the TGS measurements of the cold-worked samples are presented. Finally, the experimental results are compared with numerical calculations of the expected SAW speed as a function of surface angle and misalignment from a $\langle 110 \rangle$ crystal face.

4.6.1 Making the TGS measurements

The SAW speed as a function of crystal rotation was measured using the dual-heterodyne transient grating setup described in Chapter 3. The nominal grating spacing used for all niobium measurements was $5.5\mu\text{m}$. Each sample was measured using a rotating sample stage. Measurements were taken at increments between 1° and 10° over a range of 180° from the starting measurement point. Between 2 and 10 measurements were taken at each rotation position. Each measurement ("spot") consisted of three batches of 10,000 traces per

⁵In addition to XRD methods, an attempt was made at Los Alamos National Laboratories to use resonant ultrasound spectroscopy to characterize the dislocations as a third dislocation density characterization method, but the sample geometry wasn't sufficiently regular. Since the samples couldn't be altered, this initiative was abandoned.

batch. The sample was moved laterally and/or vertically between each measurement, so each averaged measurement point (e.g. the average v_{SAW} on the control sample at 10°) represents the average of measurements taken at different points on the sample surface while the sample was in the given rotational position.

The signal analysis technique described in Section 3.4 was used to identify the peak frequency. When there appeared to be a significant variation in SAW frequency measurements at a given position, more measurements were taken at that position in order to reduce error or confirm the variation reflected reality.

Figure 4.43 is a schematic of a niobium sample mounted on the rotating sample stage (the niobium sample is much larger relative to the stage in the schematic than it is in real life). The niobium sample is kept on a small glass slide with adhesive for easier handling (especially when mounting for XRD analysis). The glass slide and the sample are marked with the "up" direction as identified in the XRD. Crosshairs are drawn on the slide to make alignment with the sample stage earlier.

The rotating sample stage consists of a rotating circular stage in a stationary holder. This allows the sample to be rotated by well-defined increments. If the sample is properly centered on the stage, the spot location should remain the same throughout a 360° rotation unless the holder is moved laterally or vertically. The 0° point is aligned to the position indicator on the outer edge of the stage. The "up" direction is aligned to the 180° mark on the stage.

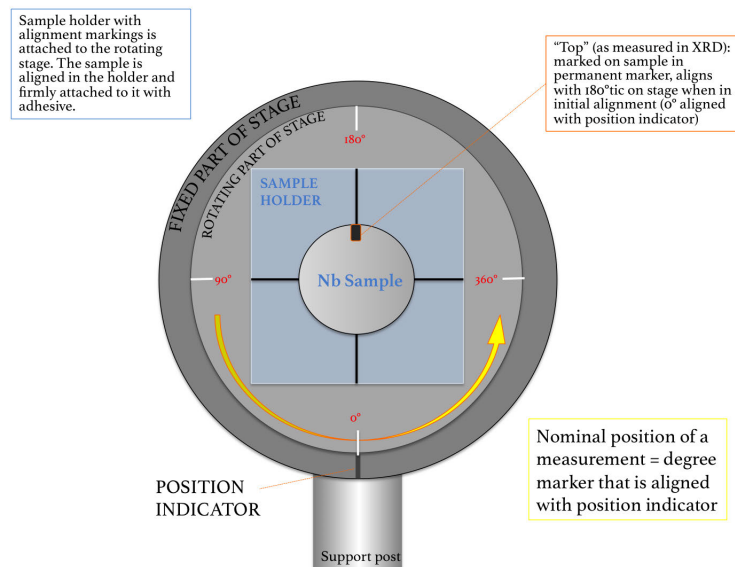


Figure 4.43: The schematic shows the niobium sample mounted on the rotating stage used to make the measurements. For clarity, the sample depicted here is much larger relative to the stage than it is in practice.

4.6.2 TGS results: $v_{SAW}(\theta)$ of cold-worked single crystal niobium

Figure 4.44 shows each individual measurement made on the six cold-worked niobium samples. Measurements were made over 180° for each sample. The plotted points represent the peak SAW frequency identified from the measured signal. The 0° orientation for each sample was identified in the XRD GADDS system as the orientation at which the sample's diffraction spot was maximized. Clearly, though, this didn't result in all six samples being oriented correctly relative to each other, since the "fast" directions occur at different rotational positions.




Figure 4.44: All individual TGS measurements taken on the cold-worked Nb samples. Each measurement consisted of three traces, which are averaged here. This data is plotted prior to shifting data sets so that maxima align.

To obtain Figure 4.45, all frequency measurements were first converted to speed using the relationship $c = \lambda v$, where v in this case is the peak SAW frequency isolated from the signal and λ is the grating spacing. Once all measurements had been converted to speed, points taken at a single rotational position on a given sample were averaged. The standard error was then calculated as σ / \sqrt{N} , where σ is the standard deviation of all the measurements taken at that position and N is the total number of measurements taken at the position. The dashed lines mark the position of (Average - Standard Error) and (Average + Standard Error) at each point. Each data set was then shifted so that the maximum of each set occurs at the same rotational position to better enable comparison of the data from each sample.

Note that it is important to convert the frequency measurements individually (instead of converting the average of frequency measurements taken at a specific rotational position) because the spacing varies slightly from day to day due to any subtle shifts that may occur with regard to how the spot falls on the phase mask.

The frequency ν that should be used is the frequency calculated from the tungsten calibration sample on the day the measurement was taken. To calculate the true grating spacing at the time of the measurement, a TGS signal from a standard single crystal tungsten calibration sample (surface orientation $\langle 001 \rangle$) is recorded. The peak frequency $\nu_{measured}$ is determined. The known speed of surface acoustic waves in $\langle 001 \rangle$ single crystal tungsten is 2665.9 m/s [140, 141, 142]. The grating spacing can then be calculated from $\lambda_{GS} = v_{SAW} / \nu_{measured}$. Figure 4.46 shows the calibrated grating spacing on each day of the niobium measurements. The measurements are clustered around the nominal grating spacing for these measurements ($5.5\mu\text{m}$) but are seen to vary by as much as $\pm 0.04\mu\text{m}$ from the nominal value - hence the importance of the calibration step at the beginning of a data collection session.




Figure 4.45: TGS data from the cold-worked niobium samples. This is the same data plotted in Figure 4.44, but the measurements at each position have been averaged and the data sets have been shifted relative to each other to align maxima. Error bars have also been added.

Figure 4.47 shows the $v_{SAW}(\theta)$ results for the control, 1500 lb, and 2000 lb samples on one graph. The response of each sample is consistent with the others. The 1500 lb sample shows the most irregularity in its response curve: there is a small local minima in SAW speed in the slow direction that is not observed in the control sample, and some irregularity in the fast direction. It is possible that the results show a small decrease in Young's modulus for the 2000 lb sample in the fast direction, but the error bars overlap with those of the 1500 lb sample.

Figure 4.48 shows the $v_{SAW}(\theta)$ results for the 1000 lb, 2500 lb, and 3000 lb samples on one graph. The response of these samples is markedly different: there is a significant sudden reduction in v_{SAW} just before the maximum SAW speed is attained that does not occur in the first group of samples. There is also a great




Figure 4.46: Calibrated grating spacing for each day of niobium measurements, showing that the grating spacing varied by as much as $0.04\mu\text{m}$ from day to day. The collection of calibration data for every day of measurement is therefore an important step to ensure that analysis of the TGS data reflects the physical properties of the samples as accurately as possible. The nominal grating spacing for all measurements was $5.5\mu\text{m}$.




Figure 4.47: v_{SAW} results for the first group of cold-worked niobium samples, graphed alone for easier comparison. It is possible that the 2000 lb sample has undergone a small decrease in Young's modulus in the fast direction, but overall, the behavior of the samples with respect to SAW speed as a function of surface rotation is remarkably consistent despite the differences in cold work between them.




Figure 4.48: v_{SAW} results for the second group of cold-worked niobium samples. These samples show a more significant variation in behavior than do the group 1 samples.

deal more variation within this group, with the most cold-worked sample showing significantly higher SAW speeds in the fast direction than the other two. Interestingly, the SAW speeds observed at the minima of the slow direction are very close.

Figure 4.49 depicts the same data, but with each sample on its own graph and with the y-axis the ratio of Δv_{SAW} to v_{min} , where $\Delta v_{SAW} = v(\theta) - v_{min}$, and v_{min} is the minimum v_{SAW} observed for the control sample.




Figure 4.49: The data for the cold-worked samples, separated by group and plotted as the ratio of Δv_{SAW} to v_{min} , where $\Delta v_{SAW} = v(\theta) - v_{min}$, and v_{min} is the minimum v_{SAW} observed for the control sample.

4.6.3 Comparison with predicted TGS response

The results presented in Section 4.6.2 can be compared to numerical calculations of the expected SAW speed in single crystal niobium as a function of surface rotation for varying degrees of misalignment from a given standard surface (such as $\langle 011 \rangle$). These calculations were carried out in MATLAB using the code reported in [143]. Figure 4.50 shows how this works in principle. One begins with a sample perfectly aligned to a given orientation - no offset or miscut. The sample can then be tilted relative to this plane and rotated about this tilt axis. Finally, the sample is rotated about the original axis - this is the “observed” surface, and this is the rotation graphed when v_{SAW} is plotted as a function of surface angle. The first two translations represent how the sample’s true alignment is offset from the ideal alignment.

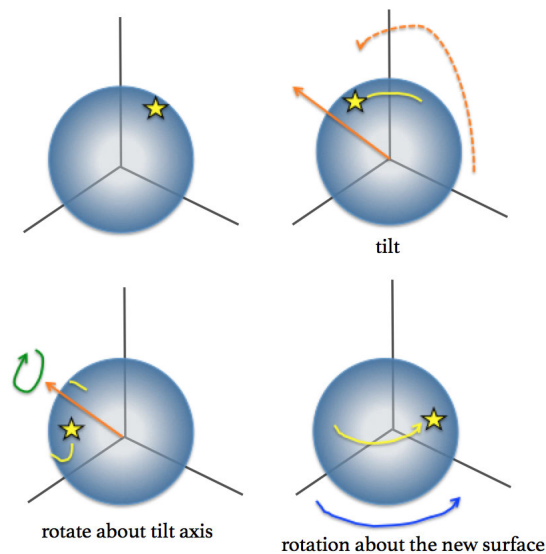


Figure 4.50: Rotational translations used in the numerical TGS predictions. For the purposes of this project, first two translations correspond to the offset of a sample’s true alignment from an arbitrary alignment.

Section 4.5.2.1 showed that the control sample was reasonably well aligned to the $\langle 011 \rangle$ direction (the vector associated with the $\langle 011 \rangle$ set of planes is normal to the sample surface). Scans suggested that the sample was offset from $\langle 011 \rangle$ in χ by about 2° .

Figure 4.51 shows the experimental TGS data from the control sample overlaid with select data from the calculated expected SAW speeds. The experimental data is plotted in black, and has been shifted in x to align with the calculated plots. In this case, the shift was determined by aligning the maximum of the experimental data with the maxima of the the calculated plots. The 0° (blue) and 15° (red) tilts for each of the selected rotations are plotted as a function of surface angle. Other tilt conditions are omitted for clarity. The graphs are tiled for easy visual comparison.




Figure 4.51: The black line shows experimental data for SAW speed versus degree of surface rotation for the control sample. The blue and red lines are the predicted SAW response for 0° and 15° of tilt. These results are given for rotations of 0 to 90° in 10° increments.

Rotations 0d, 5d, and 90d appear to be reasonable data fits. Figure 4.52 examines the experimental data against the calculated data with more tilt conditions, and it appears that a tilt of 0° most closely corresponds to the experimental data. This is in good agreement with the XRD observations for this sample, which indicated that the control sample surface was well-aligned to the <011> plane. The discontinuities on either side of 90° indicate that there is likely PSAW behavior at 90°. This is not inconsistent with the measured data, which shows a smooth curve, because it is likely that the experimental setup is simply capturing the PSAW behavior where the dips are observed.

The slowness surface for <011> single crystal Nb (Figure 4.53 indicates that we should see PSAW behavior at that orientation, and Figure 4.54 can be used to explain why the calculated curves for $v_{SAW}(\theta)$ show a dip in SAW speed on either side of $\theta = 90^\circ$ but the experimental curve does not. Figure 4.54 shows displacement as a function of slowness. The dominant displacement is actually a PSAW; the SAW displacement to the right is much smaller.⁶ (Which is which can be determined using Figure 4.53). The algorithm from [143] calculates speed from the first selected displacement curve from the right of the slowness axis in order to select the SAW behavior, so it does not “see” the PSAW response. So, the calculated curves in Figure 4.52 accurately reflect the SAW behavior.

The reason we don’t see the dip in the experimental data is that the setup isn’t sensitive enough to properly distinguish the PSAW from the SAW behavior. The SAW speed is extracted from the TGS data by taking the FFT of the signal. The PSAW and SAW frequencies are close to each other and most likely overlap in the FFT. Since the data analysis procedure calculates f_{SAW} by selecting the frequency at which the maximum intensity occurs in the FFT, we are probably just collecting PSAW data near 80° and 100°, and the SAW frequency is obscured by the PSAW peak.

Figure 4.55 shows predicted results for select rotations and 0° and 15° of tilt plotted with the experimental results for the 1000 lb sample. It is clear that the 0° rotation condition, which matched the control sample experimental data, does not match this data. A more significant rotation away from the <110> surface - particularly at 45° - yields predicted SAW speed responses that align more closely with the experimental data. This is consistent with expectations based on the XRD alignment results.

Comparisons of experimental rotation data with the predicted $v_{SAW}(\theta)$ response for the other four samples can be found in the Appendix beginning on page 232.⁷

⁶A PSAW, or pseudo-SAW, behaves much like the true SAW, only some of its energy is dissipated into the bulk (instead of only being lost through thermal diffusion and acoustic damping).

⁷These comparisons are given in the following Appendix figures.

- 1500 lb sample: Figure 8.7
- 2000 lb sample: Figure 8.8
- 2500 lb sample: Figure 8.9
- 3000 lb sample: Figure 8.10




Figure 4.52: The black line shows experimental data of the SAW speed versus degree of surface rotation for the control sample. The colored lines are the predicted SAW speeds as a function of surface rotation for an $\langle 011 \rangle$ oriented single crystal, with a $0d$ rotational offset and varying amounts of tilt offset. The arrow shows the direction of increasing tilt.




Figure 4.53: The slowness surface for $\{110\}$ single crystal Nb shows that PSAW and SAW behavior will be observed.




Figure 4.54: At certain surface rotations, the SAW response for the 0° rotation condition for $\langle 011 \rangle$ single crystal Nb is small compared to the PSAW response. However, the algorithm “picks” the response by selecting the first response from the right, and so it only “sees” the small SAW response. This explains the dips observed in the calculated responses graphed in Figure 4.52. The experimental setup still captures the PSAW response, which is why the experimental data looks like a smooth sinusoid-esque curve.

4.7 Analysis of $v_{SAW}(\theta)$ measurements for cold-worked niobium

Perhaps most obviously, the TGS results from this study confirmed what we already know - that v_{SAW} is highly sensitive to surface rotation and orientation in single crystals. The purpose of using single crystals in this study was to isolate dislocation effects, and not have to worry about effects from grain boundaries or accounting for grain size. Furthermore, at the beginning of this project, the early TGS setup performed

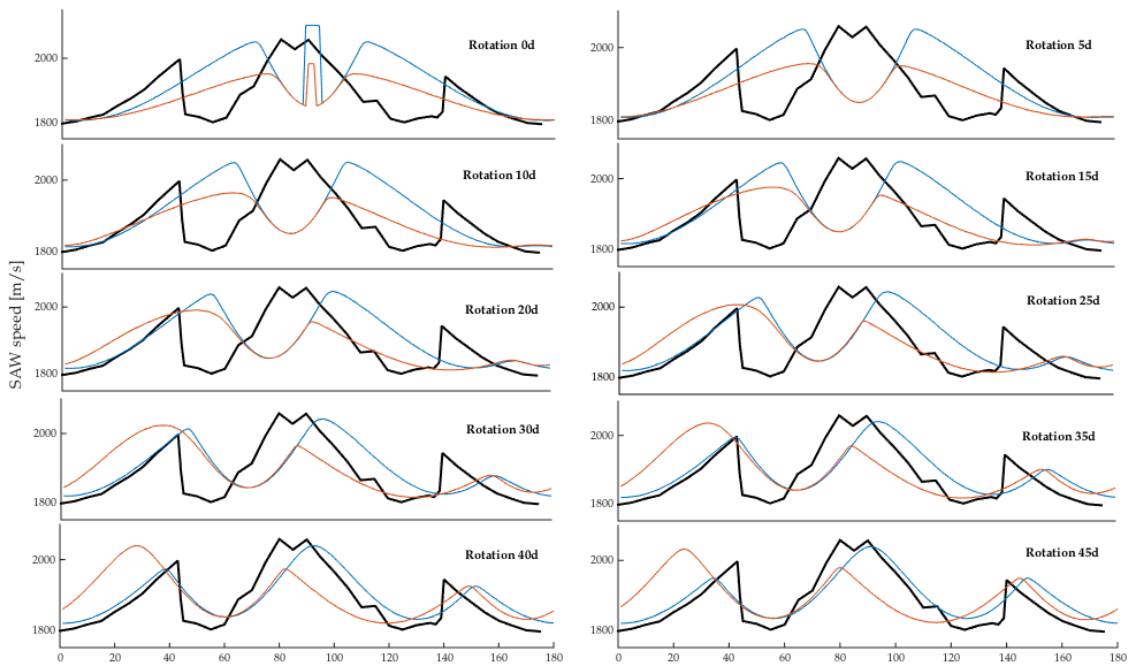


Figure 4.55: The black line shows experimental data of the SAW speed versus degree of surface rotation for the 1000 lb sample. The blue and red lines are the predicted SAW response for 0° and 15° of tilt. These are given for rotations of 0 to 95° in 5° increments. Other tilt conditions have been omitted.

poorly with all polycrystalline materials, and so projects were planned around single crystals.

However, in the years that have passed since then, the setup has been improved significantly, and is capable of producing clear TGS signals on polycrystalline materials. Polycrystalline materials have a major advantage over single crystal materials (unless they are effectively isotropic, like the tungsten calibration samples) in that they do not require careful measurements at regular known intervals of surface rotation in order to capture the full SAW speed response. If the probe spot captures TGS effects from a sufficient number of grains on an untextured surface, each measurement is the averaged behavior across many orientations. This results in a significant reduction of measurement time - for example, the single-crystal data presented in Section 4.6 took about 50 days to collect. In addition, it becomes less crucial to carefully analyze the orientation of each sample via XRD if they are polycrystalline. Finally, polycrystalline material behavior is more representative of real materials used in nuclear applications. So, now that the limitations of the setup do not predetermine the choice, the advantages of using single crystals should be weighed against the experimental expedience of polycrystals.

It was expected that dislocation density would have a measurable impact on E , which would be reflected in a change in v_{SAW} . This change was small for the $\langle 011 \rangle$ samples, which were of primary interest because this group included the control sample and because the SAW speed response followed a regular pattern. It is also difficult to separate any possible dislocation effects on SAW speed from the effects of slight misalignments relative to each other. In order to fully explore the potential utility of $v_{SAW}(\theta)$ as an indicator of dislocation density, this study should be repeated with polycrystalline samples (with grain size \ll spot size). This will remove the necessity of performing many measurements at each rotational position, it will remove the uncertainty that arises from the presence of variations in the orientation, and it will remove the requirement to carefully check the alignment of each sample via XRD methods. However, it does appear that - if any change was present - the 2000 lb sample becomes slightly *more* compliant in the fast direction (see Figure 4.47), which is at odds with the initial expectation that dislocation density would be proportional to stiffening.

The results in Figure 4.47 can also be examined in the context of the inverse relationship between dislocation density and dislocation segment length. These can be related to changes in Young's modulus as

$$\frac{\Delta E}{E} = -K\rho_D L^2 \quad (4.6)$$

, with K a constant, ρ_D the dislocation density, and L the dislocation segment length [127, 128]. For a given change in modulus, $\rho_D \propto L^{-2}$. For example, in the fast direction, the control sample and the 1500 lb sample exhibit, roughly, the same change in elastic modulus (the data points lie within each other's error bars). So, we have

$$\rho_D^{control} L_{control}^2 = \rho_D^{1500lb} L_{1500lb}^2 \quad (4.7)$$

From Figure 4.42, we know that $\rho_D^{control} \approx 10^{10} \text{ cm}^{-2}$, and $\rho_D^{1500lb} \approx 10^{12} \text{ cm}^{-2}$. From this, we can obtain a ratio of the dislocation segment lengths for the two samples: $L_{control}/L_{1500lb} = 10$. This could prove to be an interesting way of probing the characteristics of the dislocations in a material without resorting to destructive techniques like TEM. However, it does require previous knowledge of the dislocation density, which can be obtained non-destructively via XRD, but which does add time to the experimental analysis.

This phase of the study also showed that the mathematical model used to predict SAW speed as a function of miscut and surface rotation are in excellent agreement with the observed SAW speeds. The calculation results also showed how SAW speed can be used to gain insight into the orientation of a single crystal sample - for example, the $\langle 011 \rangle$ sample results lined up with the calculated results for SAW speed for 0° misalignment off the $\langle 011 \rangle$ face.

4.8 Transient grating spectroscopy of irradiated cold-worked niobium

Initially, the goal of this niobium study was to simply investigate how dislocations impacted the TGS signal, as radiation damage can cause dislocation buildup and migration in materials. Once rotation data was complete, though, and the cold-worked samples were no longer required to remain in their original states, it became of interest to irradiate the cold-worked samples and investigate a second phenomenon: instead of just using dislocations as a proxy for radiation damage, we could also investigate how radiation damage affected samples with varying dislocation densities.

Only three samples were selected for this phase of the experimental studies: the control sample, the 1500 lb sample, and the 2000 lb sample. These were chosen because the XRD analysis of the samples revealed that these three samples had the same orientation of $\langle 011 \rangle$ (at least, within a few degrees of tilt). The surface rotation (θ) measurements of these samples ($v_{SAW}(\theta)$) also revealed a more straightforward TGS response - e.g. a smooth sinusoid shape instead of the sharp fast-direction jump observed in the Group 2 samples. It seemed that it would be easier to discern changes due to radiation damage in the simpler behavioral responses of the former. Furthermore, the control sample was in this group.

Samples were irradiated to progressively larger dpa values on the CLASS (Cambridge Laboratory for Accelerator-based Surface Science) General Ionex 1.7 MV tandem ion accelerator at MIT.⁸ Between each successive irradiation, the rotation measurements ($v_{SAW}(\theta)$) would be repeated, and the results would be compared, enabling a consideration of how total irradiation and pre-existing dislocation populations affected the properties of the niobium samples.

5.3 MeV Si^{3+} ions were used to irradiate the samples. The beam current was 100 nA, and the beam passed through a circular aperture with a diameter of 8 mm. SRIM [11] was used to estimate that an average of 0.218227 displacements per Angstrom-ion through the first micrometer of niobium bombarded with 5.3 MeV Si^{3+} ions. This value was used to estimate the nominal integrated charge that each sample needed to receive from the beam in order to reach a given dpa level in the first micron from the surface.

⁸[cstar.MIT.edu/tandetron.php](http://cstar.mit.edu/tandetron.php)

Table 4.4 shows the parameters for each irradiation. During irradiations, the total dose received by each sample is tallied by keeping track of integrated charge. The nominal integrated charge is the charge needed to reach a given dpa, and is given in italics for each dpa level.⁹ During irradiation, an attempt was made to get the integrated charge as close to this value as possible for each sample. The TGS measurements were repeated after each irradiation. The nominal integrated charge given in Table 4.4, therefore, is the amount of charge required to reach the desired dpa based on the previous dpa level.

Table 4.4: Irradiation parameters for the cold-worked irradiated niobium study: integrated charge in 10^{-11}C to reach each dpa level (given in italics), and actual charge achieved during each irradiation step for each sample

total dpa	0.01	0.03	0.1	1	3
calculated charge	<i>614168</i>	<i>1228336</i>	<i>4299176</i>	<i>55275000</i>	<i>122833000</i>
control	615846	1227203	4299176	55316972	122883237
1500	614529	1228137	4299066	57020123	121658964
2000	615339	1228279	4299915	55295118	122855051

4.8.1 $v_{SAW}(\theta)$ results

The SAW speed response of the samples as a function of surface angle was measured after each irradiation up to 1 dpa. These results are shown in Figures 4.56-4.58. At least two separate measurements, taken at different x and y positions, were made at each rotational position. Data sets are shifted so that maxima and minima align. Figure 4.59 shows the data in Figures 4.56-4.58, but with all data for a given sample and irradiation level averaged. Figure 4.60 shows the average v_{SAW} in the fast direction for each sample and irradiation level, with v_{SAW} values $\pm 5^\circ$ to either side of the maximum v_{SAW} included in the average. This is the more important of the two graphs to consider, because the datasets are more meaningfully comparable. By looking only at the fast direction, we get a better idea of how v_{SAW} changed not only for changing dose for the same sample, but how it changed for different samples. In Figure 4.59, it isn't clear how the anisotropy of the response might be obscuring true differences between the samples. The averaging also leads to data sets at each dpa that are close in value but comprised of data points with lots of variation. The result of this is that the error bars overlap to the point that inter-sample differences are not discernible. Picking only one direction to look at leads to a more useful data set. In this case, the fast direction was chosen because the differences between samples appeared to be most pronounced there.

Note that v_{SAW} rises rapidly at low dose in Figure 4.60 for all three samples before leveling off. Note also that this behavior does not follow identical paths for each sample: the 1500 lb sample achieves a higher v_{SAW} at each dpa level than does the control or 2000 lb sample; the 2000 lb sample follows the control sample behavior closely (but just below it). This behavior is also observed in Figure 4.2 (electron-irradiated

⁹Charge is given in these units - of 10^{-11}C - because this corresponds to the "counts" readout on the accelerator equipment.

copper), with the control sample being the analog of the annealed sample, and the 2000 lb sample the analog of the heavily cold-worked sample. This will be discussed further in Section 4.9.1.




Figure 4.56: v_{SAW} was measured as a function of surface rotation of the control sample following irradiation with Si³⁺ ions to 0.01, 0.03, 0.1, and 1 dpa. Data sets have been shifted horizontally to align maxima.




Figure 4.57: v_{SAW} was measured as a function of surface rotation of the control sample following irradiation with Si^{3+} ions to 0.01, 0.03, 0.1, and 1 dpa. Data sets have been shifted horizontally to align maxima.




Figure 4.58: v_{SAW} was measured as a function of surface rotation of the control sample following irradiation with Si^{3+} ions to 0.01, 0.03, 0.1, and 1 dpa. Data sets have been shifted horizontally to align maxima.




Figure 4.59: v_{SAW} was measured as a function of surface rotation of the control sample following irradiation with Si^{3+} ions to 0.01, 0.03, 0.1, and 1 dpa over 180-200°. This graph shows the results of averaging all speed measurements taken on a given sample after each irradiation. Error bars are large in part due to the normal spread of v_{SAW} on the anisotropic single crystal.




Figure 4.60: v_{SAW} in the fast direction (although some of the data points in the average are likely PSAW responses, as explained on page 113. All v_{SAW} measurements $\pm 5^\circ$ from the rotational position of the measured maximum v_{SAW} are averaged.

4.8.2 Thermal diffusivity and acoustic damping results

The TGS data collected for the previous section was also analyzed for thermal diffusivity and acoustic damping. Because there were such a large number of data files to analyze - about 250 spots per sample, with each spot consisting of three positive and three negative trace files, a shell script was written to automate the MATLAB analysis. This script can be found in the Appendix on page 232.

Figures 4.61 and 4.62 show the results of this analysis for the control sample. (The corresponding results for the 1500 lb¹⁰ and 2000 lb samples¹¹ can be found in the Appendix beginning on page 235.) Note that the acoustic damping functions, at this time, are considered to be in a draft stage, and the results should be considered as an estimate, but not a definitive measurement. For each of the three irradiated samples, results for each irradiation level - 0 dpa (cold work data), 0.01 dpa, 0.03 dpa, 0.1 dpa, and 1 dpa - are plotted as a function of surface angle. Each data set was “aligned” to the others by applying the same degree shift used to align maxima of the $v_{SAW}(\theta)$ plots for the same sample and dpa. The data weren’t modulated by 180° after shifting, so any negative values are kept here.




Figure 4.61: Thermal diffusivity results, control sample. Error bars are not shown for clarity. The associated error for each thermal diffusivity measurement was on the order of 10^{-7} m²/s.

¹⁰Figures 8.11 and 8.12

¹¹Figures 8.13 and 8.14

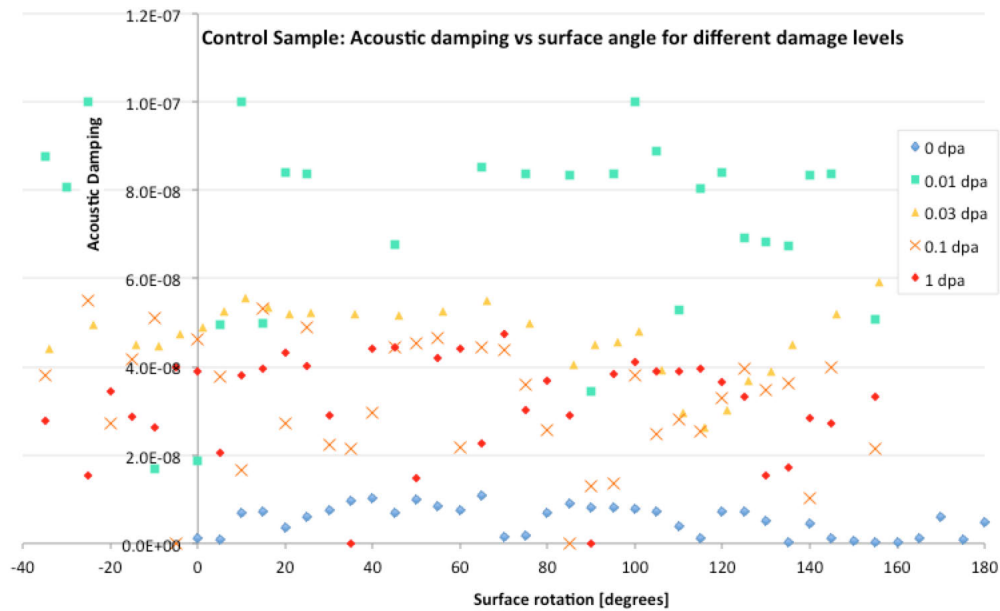


Figure 4.62: Acoustic damping results, control sample

The results in Figures 4.61-4.62 show that variation of these parameters with surface rotation is minimal, and - at this point - not likely to be meaningfully ascribed to a variation of thermal diffusivity with a specific rotation about the $\langle 011 \rangle$ axis. So, to save time, the samples were irradiated once again - this time to 3 dpa - and eight spots were measured at random positions and rotations on each sample. The averaged results of the thermal diffusivity and the acoustic damping measurements - including those for 3 dpa - are graphed in Figures 4.63 and 4.65.

Figure 4.64 replots the data in Figure 4.63 on a semilog plot, with the dpa replotted on a logarithmic axis. This better illustrates the changes in thermal diffusivity response observed at low doses. In order to keep the 0 dpa data on a log scale, 0 dpa is changed to 0.001 dpa for the purposes of the plot. However, it is not obvious at this point that the 0 dpa data is a good facsimile of what the 0.001 dpa behavior would be if it was measured.

Figure 4.66 replots the data in Figure 4.65 on a semilog plot, with the dpa replotted on a logarithmic axis. As in Figure 4.64, the 0 dpa data is assigned a pseudo-dose of 0.001 dpa to keep it on the chart, and as before, the same caveats apply - that we do not actually know what acoustic damping response would look like between 0 and 0.01 dpa.




Figure 4.63: Thermal diffusivity results for all samples after cold working (0 dpa) followed by irradiation to 0.01, 0.03, 0.1, 1, and 3 dpa. Points represent the average of all measurements. Error was an order of magnitude lower than the average and is not plotted for clarity, since the primary data points overlap in places.




Figure 4.64: The data in Figure 4.63 is replotted with dose on a log scale. The thermal diffusivity is observed to drop at low doses, rise again, and saturate near $2E-5 \text{ m}^2/\text{s}$. The 0 dpa data are assigned a pseudo-dose of 0.001 dpa here to keep them on the graph. However, the response between 0 dpa and 0.01 dpa is unknown, so it may not be accurate to consider the 0.001 dpa response as a reasonable approximation for the known 0 dpa response.




Figure 4.65: Acoustic damping results for all samples after cold working (0 dpa) followed by irradiation to 0.01, 0.03, 0.1, 1, and 3 dpa. Points represent the average of all measurements. The procedure for calculating acoustic damping is still being refined, and these measurements should be considered as estimates only. Error is on the order of 10^{-9} , including for the (1 dpa, 1500 lb) point.




Figure 4.66: The data in Figure 4.65 is replotted with dose on a log scale. As before, the 0 dpa data are assigned a pseudo-dose of 0.001 dpa. The data plotted between 0.01 and 0.1 dpa indicate evidence of an exponential relationship between acoustic damping and dose.

4.9 Analysis of the TGS results for irradiated and cold-worked niobium

4.9.1 $v_{SAW}(\theta)$ results

$v_{SAW}(\theta)$ results for the three samples used in the irradiated samples (Figures 4.56-4.58) show some speed response variation as the samples are irradiated. The control sample (Figure 4.56) behavior is very consistent, even at high doses. It is possible that the sample's compliance is increasing relative to its original state slightly at low doses (0.01 dpa) and decreasing relative to its original state as dose is increasing. Note that this behavior is only evident in the fast direction (here, at about 120° surface rotation).

The 1500 lb sample (Figure 4.57) shows what might be similar behavior in the fast direction - although it is difficult to discern, simply by looking at the response graphs, what response variations correlate to a physical change and which are the result of experimental uncertainty. The consistent responses of the control and 1500 lb samples underscore the need to repeat this experiment with a polycrystalline sample: it is easier to collect a representative range of sample responses, and the effects of anisotropy do not need to be separated from the effects of changes due to dislocations and/or radiation damage.

The 2000 lb sample (Figure 4.58) shows more variation in the response than the other samples do. Like the other two samples, the variation is concentrated in the fast direction.

In order to gain a more global view of this data, the SAW speeds were averaged for each data set (sample + irradiation condition) in Figure 4.59. However, the error bars are too large for this plot to be useful. This illustrates one of the complications with using single crystals in a study like this one: the normal variation between the speeds dominates the error. So, for Figure 4.60, only the speeds in the fast direction were averaged: this has the benefit of circumventing the error induced by the angular variation in speed and of focusing on the region of data that showed the most dramatic response to increasing radiation damage. The fast direction was considered to be the 10° centered on the top of the measured $v_{SAW}(\theta)$ curve.

Figure 4.60 shows a clearer pattern of response, without overlap between error bars. The samples appear to experience a sharp increase in v_{SAW}^{fast} at low dose - indicating a decrease in compliance - followed by a leveling-off in this behavior as dose increases past 0.03 dpa. It is possible that v_{SAW}^{fast} doesn't level off, but instead undergoes a slow decrease past 0.1 dpa. However, without additional data points between 0.1 and 1 dpa, or past 1 dpa, combined with the error bars at 1 dpa, it is difficult to determine with certainty that the long range behavior of v_{SAW}^{fast} would continue to show a decrease with increasing dose. [127] reports a slow decrease in modulus for a well-annealed sample (not graphed in Figure 4.2) with extended irradiation, however, and attributes this to a "bulk effect" dominated by the effects of point defects that distort the crystal lattice.

Since v_{SAW} is proportional to the square root of Young's modulus, this indicates that the samples exhibit a sharp decrease in compliance (increase in Young's modulus E) at low dose before a saturation point is reached, and this decrease levels off. This is similar to the low-dose behavior of the irradiated copper in [127], as shown in Figure 4.2.

With regards to the impact of cold work on the effect of increasing dose on Young's modulus, the re-

sponse shown in Figure 4.60 echoes that shown in Figure 4.2. The annealed sample in [127] corresponds to the control sample here - the sample with the lowest dislocation density. The slightly cold-worked sample in [127] corresponds to the 1500 lb sample: with a certain amount of cold work that is higher than the annealed state but less than a certain cutoff point, the sample responds particularly dramatically to low doses of irradiation, and has the most dramatic radiation-induced increase in E . The 2000 lb sample seems to correspond to the heavily cold-worked sample in Figure 4.2: its response to irradiation is similar to that of the control sample, but the increase in E at low dose occurs more gradually.

This interpretation would indicate that the 1500 lb sample dislocation density is below some dislocation density “cutoff,” below which the sample exhibits a more rapid increase in E at low dose. The 2000 lb sample, then, would be above this cutoff. However, experimental results have indicated that the 1500 lb and 2000 lb samples have relatively similar dislocation densities, with the 2000 lb sample having the greater of the two but not dramatically more so. In [127], the difference between the dislocation density of the slightly cold-worked sample and the dislocation density of the heavily cold-worked sample is more than two orders of magnitude. There may be another mechanism at work though - although the samples have similar dislocation densities, the 2000 lb sample received approximately 1/3 more force than the 1500 lb sample during the cold working process. So, the dislocations in the 2000 lb sample are likely to have a different distribution - namely, it seems likely that they have begun to pin each other to a greater extent than the dislocations in the 1500 lb sample have.

Consider L in Equation (4.6) as being flux dependent, such that $\Delta E/E \propto L(\phi)^2$. The 1500 lb and 2000 lb samples have different L to begin with, and therefore have different $L(\phi)$ values at each stage of the irradiation process. Therefore, we see different $\Delta E/E$ responses for each sample when ϕ is nonzero. Most noticeably, we observe a greater jump at low dose in $\Delta E/E$ for the 1500 lb sample than we do for the 2000 lb sample (which is less than the jump observed in the control sample). This makes sense because we expect the 1500 lb sample to have longer average L (segment length) than the 2000 lb sample, which has more pinning. It also has a much higher ρ_D than the control sample. The 2000 lb sample, while it has a much higher ρ_D than the control sample, has a smaller L due to pinning.

If one considers v_{SAW} in the fast direction at 0 dpa for the three samples, there appears to be a significant difference in the initial stiffness of the 1500 lb and 2000 lb sample, as shown in Figure 4.67. The error bars on the 1500 lb point are high, making it difficult to compare it to the control sample. While the data in this figure is not inconsistent with the trends observed in Figure 4.60, it does not show the same definitive trends. This illustrates why adding the irradiation step was necessary: by utilizing the dependence of L on ϕ , we can observe the $\Delta E/E$ trends more accurately by remeasuring v_{SAW} for each sample as dose is increased.




Figure 4.67: SAW speed in the fast direction at 0 dpa for the control, 1500 lb, and 2000 lb single crystal niobium samples.

4.9.2 Thermal diffusivity results

Figure 4.61 shows the raw thermal diffusivity (κ) data for the control sample. At a low dose of 0.01 dpa, α is decreased by $\approx 70\%$.¹² At a slightly higher dose of 0.03 dpa, α increases, but only to $\approx 60\%$ of its initial value. However, the thermal diffusivity data measured across 180° for 0.03, 0.1, and 1 dpa overlap. Thermal diffusivity response does not appear to vary with surface rotation, as expected.¹³ The same pattern is observed for the 1500 lb and the 2000 lb samples (Figures 8.11 and 8.13, respectively).

Figure 4.63 shows the averaged results of these measurements, plus the data from the TGS measurements carried out post-irradiation to 3 dpa. Unlike v_{SAW} , α doesn't have a strong surface angle dependence, and therefore the anisotropy complications of using single crystals to measure v_{SAW} (and E) are not as pressing. At the lowest dose (0.01 dpa), α decreases by more than 70%. At the next highest dose (0.03 dpa), α increases sharply, although its value is only about half what it was in the unirradiated state. It appears that α effects are then saturated: change is minimal as dose increases out to 3 dpa.

Figure 4.64 replots the data from Figure 4.63 on a log scale, which makes the response of α to dpa easier to discern. However, it would be useful to know more about the response of α for dpa between 0 and 0.01 dpa. As it is, the 0 dpa data have been assigned a dpa of 0.001 for the purposes of presenting all the data on a log plot, but it is not clear that the response of α at 0.001 dpa would actually be similar in magnitude to the 0 dpa data.

The level of cold work does not appear to have a significant impact on the thermal diffusivity results. All three samples have very similar thermal diffusivities at each irradiation level, and there is no clear pattern in

¹²This type of drastic decrease in thermal properties with irradiation is not unheard of - for example, it's been shown to occur in SiC thermal conductivity (which is directly proportional to thermal diffusivity) to an even greater extent at even lower dpa. See: [144]

¹³The variation in the 0 dpa sample might be attributable to experimental error; the 0 dpa measurements (the cold-work measurements) were taken a year previously, with a less refined version of the TGS setup.

which one sample's data point is always higher or lower than the rest.

At low doses, point defects dominate the radiation damage modes. Point defects are efficient scatterers of free electrons, which are the primary conductors of heat in metals, and so thermal diffusivity sharply decreases [145, 146]. As dose increases, point defects are lost at sinks or incorporated into larger mesoscale defects like voids and dislocations [12]. These defects are not as efficient at scattering electrons, and so α rises again. However, there is still some equilibrium concentration of point defects in the bulk, albeit much lower than the concentration of point defects at very low doses. This lower concentration of point defects, in conjunction with the mesoscale defects, still result in a significant reduction of thermal diffusivity in comparison to the unirradiated state.

Literature results are consistent with these findings. TEM results of neutron irradiated niobium from [147], shown in Figure 4.68, show how neutron irradiation ($E > 0.1$ MeV) affected single crystal niobium samples with varying starting concentrations of oxygen impurities. The oxygen impurities are mobile during irradiation. As irradiation dose increases, the number of radiation-induced defects increases, and these trap the oxygen impurities. Increased dose “cleans up” the microstructure of the bulk, with defects and impurities concentrating in cluster (the dark spots in the TEM images).

Similar results were found for pure polycrystalline niobium irradiated with 5 MeV protons to 0.01-4 dpa, which is comparable to the dpa range induced here with Si^{3+} [148]. Figure 4.69 shows the dislocation density as a function of dpa observed in [148]: dislocation density rises sharply at low dose (0.01 dpa) as radiation damage builds up in the bulk. As dose increases, these defects concentrate into larger defects that are less dispersed through the bulk.¹⁴ These results were confirmed via TEM.

The TEM results from [148] and [147] and the dislocation density measurements from [148] indicate that it is reasonable to assume we are seeing similar behavior in the cold-worked niobium samples.¹⁵ In this case, the initial spike in dislocation density at 0.01 dpa results in a sharp drop in thermal diffusivity. Dislocation density then decreases and remains steady, and this is reflected in the slight rise in α , followed by saturation, that is evident in Figure 4.64.

4.9.3 Acoustic damping results

It is anticipated that defects like dislocations will have a strong effect on acoustic damping, and acoustic damping has been used to investigate meso- to macro-scale damage like cracks and flaws [149] as well as micro- to meso-scale damage like creep and dislocations [150] in non-nuclear structural materials. However, at this stage of the project, the code and experimental setup are not able to extract acoustic damping information from the TGS signal data with the sufficient level of accuracy necessary for meaningful conclusions.

¹⁴In particular, it is proposed in [148] that interstitial point defects diffuse more quickly to sinks, leaving behind a greater concentration of vacancy point defects, which form clusters that collapse to dislocation loops - which, in turn, become defect sinks.

¹⁵It was not possible to perform TEM analysis on the samples for this project because they needed to remain intact between irradiations in order for the results from each step to be comparable, as there was only one of each sample. Any future studies building on this phase of this project should consider beginning with multiple samples for each level of cold work, in order to allow for one sample to be used for destructive analysis that can be used to analyze the TGS results.




Figure 4.68: Oxygen impurities in single crystal niobium are mobile during neutron irradiation, and are trapped at radiation-induced defect clusters (dark spot). [147]




Figure 4.69: Dislocation density of niobium samples irradiated with 5 MeV protons rises sharply at low dose (0.01 dpa) as radiation damage builds up in the bulk. As dose increases, these defects concentrate into larger defects that are less dispersed through the bulk. These results were confirmed via TEM. Adapted from [148]

Data presented in this thesis should be considered a rough estimate.

Figure 4.66 averages the data collected for each sample after each irradiation step. It is difficult to identify the sample behavior at low doses, so the data is replotted on a semilog plot, again assigning a pseudo-dose of 0.001 dpa to the 0 dpa data. The semilog plot shows evidence of exponential decrease of acoustic damping with dose between 0 and 0.1 dpa.

If the data presented in Figure 4.66 can be taken at face value, it is possible to interpret it in the context of the radiation damage evolution theory postulated above. At low doses, point defects dominate the radiation damage modes present. These point defects are efficient scatterers of phonons, and so at low doses there is a significant increase in acoustic damping as compared to the unirradiated condition [151, 152, 153]. As meso-scale damage begins to dominate the damage modes, phonon scattering becomes less efficient as compared to the point-defect-dominant condition, and a decrease in acoustic damping is observed. However, the acoustic damping is still greater than it was in the unirradiated condition.

In order to improve MIT MNM's current ability to probe acoustic damping response via TGS, it would be necessary to have a better physical model of acoustic damping. Many existing models describe acoustic damping empirically or semi-empirically, because there are so many factors that can influence acoustic damping. This makes it difficult to come up with a generalized model of the phenomenon. However, without such a model, it is difficult to write an analysis code for the TGS signal data that accounts for all of the experimental factors that could influence measured acoustic damping behavior but which are not due to properties of the sample being measured.

4.10 Conclusions of the irradiated, cold-worked niobium study

The study demonstrated how TGS is most useful for studying changes in irradiated samples. Characterizing the cold work alone, without irradiating the samples, yielded inconclusive results, probably due to the competing effects between dislocation density and dislocation segment length in cold-worked materials. When it comes to characterizing defect populations, the irradiation response is itself an important indicator, because the expected irradiation response will be modulated by the defect populations. For example, at 0 dpa, the group 1 fast direction SAW speeds didn't really provide enough information to elucidate anything particularly interesting or definitive about the samples. However, as the samples were irradiated, their elastic modulus changed at different rates as a function of dose due to differences in their initial microstructures. So, the initial *ex situ* TGS studies on various unirradiated materials don't always do a good job of underscoring how useful the TGS technique actually is for measuring radiation damage. Add in radiation, however, and a lot of interesting information is revealed.¹⁶

This leads to the other important conclusion from this study, which is that the level of cold work in a sample will affect its behavior at low dose, particularly with regards to changes in the elastic modulus. This means that low dose measurements of v_{SAW} can be used to elucidate the starting microstructure of a sample

¹⁶"As with most of life's problems, this one can be solved by a box of pure radiation." - Andy Weir, *The Martian*

if it isn't known. Of course, it's more likely that the starting microstructure *is* known; in that case, these results are important because it means that, when examining the low-dose response of E in a material, one should account for the effect of starting levels of cold work (and dislocation density, if also known) before attributing the rate of change of E to other defect populations.

This study also functioned to underline the importance of the thermal diffusivity measurement capabilities of the TGS technique, because these measurements provided insight into the evolution of defect populations as a function of dose. An initial decline in thermal diffusivity is associated with a high concentration of point defects in the niobium bulk at low dose. The thermal diffusivity rises and saturates as vacancies and interstitials conglomerate into larger defects and move toward sinks. TGS can be used to pinpoint when the point defects dominate the defect population, and when larger defects begin to dominate. The ability of TGS to quickly measure thermal diffusivity as it changes with defect populations, even over small changes in dose, is also important to the broader goals of nuclear materials research, because a decrease in thermal diffusivity due to radiation-induced defects is a major concern to nuclear engineers. If an in-core component (for example, fuel cladding) undergoes a major drop in thermal diffusivity due to radiation damage, this can be very detrimental to the safe operation of the core since the component can no longer transfer heat away effectively, and overheating (and possible component failure) may result. TGS can therefore be used to quickly validate that a given material's thermal properties will remain within specified margins for all expected radiation doses.

Chapter 5

Transient grating spectroscopy of intermetallic Nickel Aluminum

5.1 Introduction to B2-NiAl

This chapter is focused on TGS characterization of defect populations in intermetallic nickel aluminum samples of varying composition. Intermetallic NiAl has high concentrations of stable constitutional vacancies. The vacancies that exist in an irradiated material due to both radiation damage (point defects) and elevated temperature (thermal vacancies) tend to anneal out when removed from the radiation and/or heat source. This presents a complication for *ex situ* nuclear materials studies, as an irradiated sample *ex situ* may not be representative of an under-irradiation sample *in situ*. Intermetallic NiAl presents a way to study the impact of high vacancy concentrations on the TGS signal.

This chapter introduces the properties and crystal structure of B2-phase NiAl. It then describes the procedures used to create multiple groups of single crystal and polycrystal intermetallic NiAl of varying composition. Next, TGS results from these samples is presented and discussed. In the next major part of this chapter, the molecular dynamics code LAMMPS is used to simulate the TGS experiments, and these results are compared with experimental results. In order to carry out this phase of the project, it was necessary to first ensure that the MD structures reflected the true behavior of NiAl across its compositional range. This procedure is also detailed, as intermetallic NiAl is an important material in non-nuclear engineering applications, and so the MD study of it is of interest to other fields.

Finally, the results presented in this chapter are considered in the context of their implications for the future of TGS as a method of radiation damage measurement.

5.1.1 Motivation for using NiAl in this study

Intermetallic nickel aluminum is of particular interest to this project because of its high concentration of stable constitutional vacancies at certain compositions. This allows for the investigation of the effect of vacancies on the TGS signal *ex situ*, without any additional specialized equipment. For example, initial ideas for studying vacancies *ex situ* involved heating metal samples in a furnace to create thermal vacancies, and then immediately flowing liquid helium or nitrogen over the sample in an attempt to partially “freeze” the thermal vacancies in, as they tend to anneal out at room temperature. Figure 5.1 shows the concentration of thermal vacancies expected in aluminum and niobium as a function of temperature.




Figure 5.1: Thermal vacancy concentration in aluminum and niobium as a function of temperature.

Note, however, that Figure 5.1 shows that even near the melting point of a metallic sample, one only expects on the order of 0.01% thermal vacancies, meaning that relying on heat-induced vacancies is not likely to be a good proxy for studying radiation-induced vacancy concentrations. (This is especially true at this stage of the TGS research, since such small vacancy concentrations were not likely to result in a measurable change to the TGS signal - at least not one that could be discerned through typical measurement uncertainties.) Meanwhile, vacancy concentrations are known to reach up to 1% in irradiated samples [185]. So, materials with high concentrations of stable constitutional vacancies are - for now - the best proxy material we have for beginning to study vacancy-induced changes to the TGS signal *ex situ*.

It is also an interesting material in its own right, and its high strength at high temperature has made it an attractive material in the aerospace industry, among others. Intermetallic NiAl is an attractive material for certain engineering applications because of its high strength, low specific weight, thermal stability across a wide range of temperatures, strength and stiffness stability at high temperatures, high thermal conductivity, high melting point, and good corrosion and oxidation resistance at high temperatures [154, 155]. Proposed uses include oxidation resistant turbine blade coatings [162], coatings on Ni- and Co-based superalloys, metallization on III-IV semiconductors, high voltage electrodes in vacuum circuit breakers [156], and in some cases as a replacement for Ni-based superalloys [154].

Thus the studies of NiAl for this project have a twofold impact: first, they serve the ends of the TGS project, and second, the results are of interest to any engineering field considering intermetallic NiAl as a candidate material.¹

The NiAl experiments in this thesis are carried out on both single and polycrystal samples. Single crystal NiAl is itself of interest because it performs better than polycrystalline NiAl with regards to high-temperature strength and ductility when used as a turbine blade coating [159].




Figure 5.2: The NiAl phase diagram

Figure 5.2 shows the phase diagram that (anecdotally) seems to appear the most in literature concerning

¹Certain other intermetallics also have significant concentrations of constitutional vacancies, including CoAl and NiGa [157]. TiAl is reported to reach upwards of 10 at% constitutional vacancies for certain Ti-poor/Al-rich compositions [158]. NiAl was chosen because (1) it tends to have higher concentrations of constitutional vacancies than the other intermetallics that share this structural property, (2) is composed of two metals that are familiar to our research group due to their ubiquity in common metallurgical applications, (3) is composed of two metals that are easily obtainable and relatively affordable at high levels of purity, and (4) is a significant material to other engineering fields, as discussed above.




Figure 5.3: The phase diagram above, annotated

nickel-aluminum alloys. Of interest to us is NiAl, in the center of the phase diagram - sometimes referred to as β -phase NiAl, and more commonly referred to as B2-phase NiAl. “B2” is used to describe an intermetallic compound in which there is an equal number of atoms of each metal. Throughout this chapter, “NiAl” should be understood to mean B2-phase NiAl unless otherwise noted.

β -NiAl exists over a wide range of compositions, from under 45 at% Ni to just below 60 at% Ni at low temperatures. At higher temperatures, NiAl may exist as β -NiAl at at% Ni as low as 40% and as high as 65%. Interestingly, the phase diagram also shows that, in the low-temperature compositional range, the system will form β -NiAl once $T < T_{melt}$: there are no other phases present after solidification.

5.1.2 The structure and lattice parameter of B2-phase NiAl

In this section, the changes observed in the NiAl structure as its composition varies are described. A particular focus is placed on the NiAl lattice parameter, both because it illustrates the unusual behavior of this class of intermetallic material and because it is important to the molecular dynamics simulations of TGS experiments on NiAl carried out in Section 5.4. The second major focus of this section is the presence of constitutional vacancies in Ni-rich NiAl. Additional information about NiAl’s properties can be found in the Appendix beginning on page 244.

NiAl is often described as having a “CsCl”-type structure, analogous to that of cesium chloride salt

[160]. It consists of simple cubic lattice of nickel atoms interlocked with a simple cubic lattice of aluminum atoms (see Figure 5.4). Figure 5.2 shows that NiAl is stable over a wide range of temperatures, from about 45 to 60 at% Ni at lower temperatures and from about 43 to 67 at% Ni at higher temperatures [161, 162].




Figure 5.4: The structure of B2-phase NiAl consists of a simple cubic lattice of Ni atoms interlocked with a simple cubic lattice of Al atoms. (image: <http://minerva.mlib.cnr.it/>)

The lattice parameter of NiAl is typically reported as 2.88 Å. This value pertains to the lattice parameter reported for 50 at% Ni NiAl in [163]. However, this parameter varies for off-stoichiometric compositions of NiAl. Reported values from four different literature sources are collected and plotted in Figure 5.5. This variation is to be expected because the defects associated with the non-stoichiometric compositions of NiAl affect the lattice structure.²

The presence of constitutional vacancies in aluminum-rich compositions of B2-phase NiAl is generally accepted in the available literature [168, 169, 170, 171, 167]. Questions as to the actual concentration of vacancies as a function of composition remain, with theoretical models often at odds with experimental measurements [171]. The literature also shows variation in the reported concentration of vacancies in off-stoichiometric NiAl, as shown in Figure 5.6. Concentrations of vacancies are generally predicted to be fairly high at lower concentrations of nickel, with the lower reported values at about 5% and the higher reported values above 12%. Interestingly, one source, [167], reports nonzero concentrations of vacancies on the aluminum sublattice as well. Another interesting feature of this data set is the comparison between “Kogachi1996”, [167], and its 2001 update, referred to in the graph as “Kogachi1996 updated” [167]. The update reports high concentrations of vacancies on the nickel sublattice at 52% nickel, whereas the 1996 paper (as well as an unrelated 1997 paper) report vacancy concentrations that tend toward zero as the stoichiometric composition is reached.

The studies used to create Figure 5.5 and 5.6 made use of density measurements of NiAl samples to experimentally determine its vacancy concentration [167, 172, 169, 173, 161]. [169] and [161] state that the

²All concentration percentages of Ni or Al in NiAl are atomic percentages. If at% is not specified, it should be assumed that it is an atomic percentage. The clarification “at%” is dropped here in favor of “%” for readability.




Figure 5.5: The lattice parameter of B2-phase NiAl varies with composition. The lattice parameter reaches a maximum near the stoichiometric point. It decreases as the nickel content goes down relative to the stoichiometric point and the number of nickel vacancies goes up. It decreases at a more gradual slope as the nickel content increases relative to the stoichiometric point and the number of nickel anti-site defects on the aluminum sublattice goes up. Citations for the data plotted are as follows: “Cooper1963”, [164]; “Jacobi1971”, [165]; “Hughes1971,” [166], and “Kogachi1996”, [167].




Figure 5.6: Reported concentrations of structural vacancies present in NiAl as a function of composition from select literature. Note that the “Pike1997” and “Kogachi1996” report C_V^{Ni} values of ≈ 0 for at% Ni $\geq 50\%$, while the 2001 update to Kogachi1996 reports a C_V^{Ni} at 52 at% Ni that is $\approx C_V^{Ni}(50\% \text{ Ni}) \approx 9\%$. The updated Kogachi data is also significantly higher than the other data sets, with reported C_V^{Ni} of over 12% at lower values of nickel concentration. Interestingly, Kogachi also reports nonzero C_V^{Al} , although it never exceeds 2%. For less than 50% Ni, C_V^{Ni} decreases monotonically with increasing Ni concentration. Citations for the data plotted are as follows: “Kogachi1996”, [167]; “Kogachi1996 updated”, [172]; and “Pike1997,” [173].

density was used in conjunction with lattice parameter measurements in order to measure C_v . [169] notes that the density measurements lead to systematic errors in the determination of C_v . In particular, microvoids that are artifacts from melt solidification process used to make the NiAl samples decrease the density measurement. As a result, C_v was systematically overestimated. On the other side of the stoichiometric point, where at% Ni $> 50\%$, Ni antisite defects on the Al sublattice accommodate the excess of nickel, and constitutional vacancies on the Al sublattice are not observed.

Additional evidence for the concentration-dependent structural changes observed in B2-phase NiAl NiAl has a particular composition-dependent structure: there is a decreasing concentration of constitutional vacancies as at% Ni approaches 50%, followed by (in theory) an ideal CsCl structure at the stoichiometric point, followed by little or no constitutional vacancies but increasing concentration of Ni ASD on the Al sublattice. In Section 5.1.2, evidence for this compositional dependency could be seen in measurements of the NiAl lattice parameter, which were shown in Figure 5.5. Further evidence is provided by observations of other B2-NiAl material properties as a function of composition. Two electromagnetic properties are considered here as an example.

Magnetic susceptibility measures the intensity and direction (attraction or repulsion) of a material’s response to an external magnetic field. It is a dimensionless constant, typically denoted χ , and is calculated as the ratio of the material’s magnetization to the strength of the applied magnetic field. The magnetic




Figure 5.7: The plot of magnetic susceptibility as a function of composition illustrates that NiAl has very different material properties on either side of the stoichiometric composition. Adapted from [174].




Figure 5.8: The plot of magnetic resistivity as a function of composition illustrates that NiAl has very different material properties on either side of the stoichiometric composition. Resistivity increases with increasing defect concentration: above the stoichiometric point, resistivity increases as the concentration of nickel vacancies increases, and below the stoichiometric point, resistivity increases as the concentration of nickel anti-site defects increases. [174]

susceptibility of NiAl at room temperature varies with composition in [174], as shown in Figure 5.7, with sharply different behavior on either side of the stoichiometric point. In this case, as at% Ni increases toward the stoichiometric point, susceptibility gradually increases as the number of vacancies decreases. The magnetic susceptibility behavior changes sharply at the stoichiometric point: there is a sudden drop in susceptibility when NiAl is close to an ideal CsCl structure. It then increases more dramatically with at% Ni continues to increase. χ is affected by the electron density of states, which is in turn determined by the composition, structure, and other characteristics of a material. The results in Figure 5.7 - particularly the discontinuity at the stoichiometric point - indicate that NiAl's structure is markedly different for the Ni-poor and Ni-rich compositions in a way that can't be simply by having unequal concentrations of Ni and Al atoms (in which case one would expect a gradual, but continuous, trend in χ as a function of composition.)

Resistivity measurements from [174] are given in Figure 5.8, and these measurements also show an obvious change in behavior at the stoichiometric point. Resistivity decreases with increasing aluminum content for at% Al < 50% - e.g., resistivity decreases as the concentration of nickel anti-site defects decreases and the NiAl sublattice structures are closer to their defect-free versions. In the aluminum-rich zone, where

the off-stoichiometry composition is accommodated by nickel vacancies on the nickel sublattice, resistivity increases as at% Al increases, which corresponds to an increasing concentration of nickel vacancies. So, the data here can be interpreted as resistivity that increases monotonically with increasing defect concentration (be they ASD or vacancies).

5.2 The NiAl samples used in the TGS experiments

This section describes the fabrication of NiAl samples made for this study. There are two groups of samples that are used. The first group of samples that are used in the TGS results are referred to as "Batch I". The Batch I samples are single crystal samples. Batch I also includes a single crystal NiAl sample from a previous thesis work carried out in the MIT Materials Science and Engineering department [159].

"Batch II" made in the summer of 2017 to (a) increase the breadth of the sample matrix and (b) make use of lessons-learned while making Batch I. The two batches are also distinguished in this work due to slight differences in their respective fabrication procedures. Batch II samples were measured in their polycrystalline state.

5.2.1 NiAl sample fabrication: Batch I (Fall 2015)

5.2.1.1 Fabrication of B2-phase NiAl in the literature

Table 5.1 describes the procedures used to make NiAl in papers selected from five decades of research that used the B2-phase of the intermetallic for various research purposes. This information was used to plan the best procedure for making single crystal and polycrystalline NiAl samples of different compositions using available resources. In particular, the literature shows that it is important to use high purity starting materials and to melt them under an inert atmosphere to prevent contamination. It was also common to have lengthy annealing periods after NiAl was made.

5.2.1.2 Raw materials

Nickel and aluminum powders were ordered from Alfa Aesar. The nickel powder was sized at -120 mesh (90% of particles pass through a sieve with mesh openings 0.125 mm wide) and had a purity of 99.996% (metals basis). The aluminum powder was sized at -100 + 325 mesh (90% of particles pass through a sieve with mesh openings 0.149 mm wide and 90% of the same particles are retained by a sieve with mesh openings 0.044 mm wide) and had a purity of 99.97% (metals basis).

5.2.1.3 The Busso samples

One set of B2-phase NiAl samples were procured from Sam Allen (MIT Materials Science and Engineering). These samples were used in the 1990 doctoral thesis of E. Busso [159]. These samples were used to study

Table 5.1: Preparing NiAl specimens: some examples from six decades of literature

Year	Source	Make single crystals	Heat treat	Sample preparation
1963	[164]	Mix Ni and Al powders, compress in 0.25" die under 2-ton load, heat in Pt resistance furnace under vacuum.	Anneal for 36h at 1300°C, cool, and remove. Grind in mortar, repress, and anneal for 48h at 1300°C.	Mount in diakon and polish.
1966	[168]	Slow cool single crystal NiAl from a melt. ($\leq 0.06\text{wt}\% \text{ C}$)	Homogenize for 12h at 1300°C under argon.	Orient samples on the cutting jig using X-ray techniques and section with a diamond saw.
1969	[174]	Arc melt Ni and Al powders under gettered argon atmosphere.	<i>Not given</i>	<i>Not given</i>
1977	[161]	Induction melt Ni and Al under argon atmosphere.	Hold for 4-6wks at $T_{\text{solidus}} - 50\text{K}$ and cool slowly.	Use Laue backscatter to locate $\langle 110 \rangle$ direction and spark cut at minimum possible current.
1979	[175]	Vacuum induction melt Ni and Al powders and cast into rods.	Homogenize for 72h at 1150°C under Ar atmosphere. Use alumina/NiAl powder packing to prevent Al loss to oxidation. Cool at 1K/min to 600°C and hold for 24h. Furnace cool to room temperature.	Use centerless grinding to prepare rod-shaped specimens.
1989	[220]	<i>Not given</i>	Anneal after polishing at 800-1000°C to ensure well-ordered surfaces.	Cut samples to expose low-index faces. Mechanical polish is followed by Ne^+ ion sputtering to clean surfaces.
1994	[221]	Use mechanical alloying to create nanostructured B2-phase compounds (not single crystal). Alloy under Ar atmosphere.	N/A	N/A, use powder as-is post-milling for XRD analysis.
1996	[167]	Melt pure Ni and Al in arc furnace with non-consumable W electrode under Ar atmosphere.	Homogenize ingots at 1273K for 7d in sealed silica tube filled with Ar. Water quench.	Ingots were crushed in a mortar to 30-160 μm particle size, not sectioned.
1997	[177]	Use floating-zone technique and induction heating under Ar to create single crystal samples. Paper gives highly detailed description of necessary furnace setup.	N/A	(Prep. is for single crystal growth, not post-growth examination.) Melt powders by induction heating. Cast in Cu moulds. Clean by etching in 5:3:1:1 acetic/nitric/sulfuric/phosphoric acid

deformation in single crystal NiAl. A specimen from this set of samples was used in this study, in addition to the the NiAl samples fabricated from scratch.

The crystals were grown using the floating zone method, and were determined to have a composition of 50.6 ± 0.6 at% Ni and 48.4 ± 0.2 at% Al via an atomic absorption method. The samples were found to have a dislocation density of approximately $0.2 \mu\text{m}/\mu\text{m}^3$ via TEM analysis. The anisotropy of the NiAl matrix affects the dislocations, which form in a manner that minimizes their line energy.³ Figure 5.9 is an electron micrograph from [159] that shows the two types of dislocations the author observed in these samples.

The specimens described in the thesis are oriented to within $\pm 2^\circ$ of the $\langle 100 \rangle$ direction.



Fig. 4.6 Electron micrograph showing a helical dislocation and sharply kinked or zig-zag dislocations

Figure 5.9: Dislocations in an NiAl sample used in [159]

5.2.1.4 Preparation for arc melting

First, a sample matrix was planned as detailed in Table 5.2. 25g had been ordered of both Al and Ni powders, so five samples spanning the compositional range of B2-phase NiAl were planned. It was planned use 9.8g total of powder for each sample to ensure that there would be some powder left over for additional testing if needed.

The composition plan was based on at% of Ni and Al, per the convention of the typical NiAl phase diagram, so it was necessary to convert at% to wt%, and wt% to actual mass of each powder needed to create the samples. Equations (5.1) and (5.2) describe this calculation.

$$g_{\text{Ni}} = g_{\text{total}} \left(\frac{A_{\text{Ni}}/100 \times 58.6934 \frac{\text{g}}{\text{mol}}}{A_{\text{Ni}}/100 \times 58.6934 \frac{\text{g}}{\text{mol}} + A_{\text{Al}}/100 \times 26.9815 \frac{\text{g}}{\text{mol}}} \right) \quad (5.1)$$

³Busso explains in Section 4 of [159] how this results in square-shaped dislocations with edges parallel to $\langle 100 \rangle$ axes. Dislocations with sharp kinks were found to lie on a certain plane and have a certain Burger's vector associated with the screw resultant ((011) and [100] respectively). Busso explains this by showing that $\{100\}\langle 100 \rangle$ screw dislocations tended to be unstable in B2-NiAl.

$$g_{Al} = g_{total} \left(\frac{A_{Al}/100 \times 26.9815 \frac{g}{mol}}{A_{Ni}/100 \times 58.6934 \frac{g}{mol} + A_{Al}/100 \times 26.9815 \frac{g}{mol}} \right) \quad (5.2)$$

where g refers to grams (g_{total} refers to total mass of the powders used for one sample, here 9.8g) and A refers to nominal atomic percentage.




Figure 5.10: The NiAl phase diagram

Sample №	at% Ni	at% Al	wt% Ni	wt% Al	[g] Ni	[g] Al
1	46	54	64.9498	35.0502	6.3651	3.4349
2	49	51	67.6377	32.3623	6.6285	3.1715
3	52	48	70.2079	29.7921	6.8804	2.9196
4	55	45	72.6680	27.3320	7.1215	2.6785
5	58	42	75.0250	24.9750	7.3525	2.4475

Table 5.2: Composition plan for sample matrix, NiAl Batch I

Powders for each sample were weighed on a scale in a glove box. This was done partially to minimize the potential of contamination by other materials during initial stages of sample preparation. The primary reason for using a glove box was safety: this was the step that involved the most transferring of the finely milled powders between containers, and use of a glove box removed concerns of skin contact and inhalation.

The weighed powders were separated into five air-tight glass containers that were labeled 1-5. Basic mixing of the powders was carried out via gentle stirring of the powder in each container.

Next, the powders were pressed into pellets using a Carver® pellet press in a vacuum environment. Each pellet had a mass of 1-3g, so 3-5 pellets were created for each sample. This step was purely for convenience: it is much easier to transfer a few solid pellets into the crucible of an arc melter than it is to pour very fine powder into the same crucible without risking inhalation or powder loss. The pellet pressing process is depicted in Figures 5.11-5.13.

5.2.1.5 Arc melting and post-melt heat treatment

Arc melting was carried out in the lab in a Miller arc melter. Pellets were placed in a copper crucible designed for use with the arc melter, which was cleaned with alcohol prior to use. The melting chamber was purged of all air using a vacuum pump and refilled with argon. This purge-and-refill process was repeated four more times to ensure that the environment inside the melting chamber was fully inert.

A foot pedal is used to supply power to the electrode. Once power is being supplied, the electrode is slowly moved toward the sample until the arc appears. In general, the minimum amount of current required to create a successful arc was used.

The arc is moved over the raw material (Ni/Al pellets) until it is thoroughly melted, glowing hot, and constricted into a spheroid shape. At this point, the foot pedal is released, the arc melter power is turned off,



Figure 5.11: Ni/Al powder mixture is transferred to the pellet press die using a funnel and a disposable scoop




Figure 5.12: A vacuum pump (connected to the tube) is used to remove air while samples are pressed.



Figure 5.13: The finished pellets are about 0.5" wide and 0.25" high. The pellets are much easier to handle than the loose powders, especially during the arc melting step.

and the sample is allowed to cool for at least five minutes. The sample and crucible are then retrieved for inspection and cleaning, and the sample is flipped over. The melting process is then repeated.

This entire process is repeated three times, flipping the sample in between each melt, so that each sample is arc melted six times.⁴ The final product is a small metallic button of material, about 0.75" wide and 0.5" high. Figure 5.14 shows the Short lab arc melter in use: the user is controlling the electrode's position with his hands while supplying power with the foot pedal, and observing the sample's melt progress through the welding glass windows that surround the melting chamber.

The buttons were then heat treated in a high-temperature graphite furnace under an inert environment at 1200°C for 24 hours. Figure 5.15 shows the five sample buttons in the furnace prior to beginning the heat treatment. The first purpose of this heat treatment was to grow larger grains, so that TGS could be easily performed entirely within the boundaries of a single grain if necessary in case the single crystal fabrication plans were unsuccessful. The second purpose of this heat treatment was to give the five buttons a more identical heating history, since there would naturally have been variations in the temperatures and durations of the arc melting treatment each button received.

5.2.1.6 Metallographic examination of the as-arc-melted NiAl buttons

Prior to the 1200°C heat treatment, a small section was sliced from each button using a low-speed diamond saw and mounted in epoxy. The samples were then mechanically polished to 1200 grit sandpaper on a

⁴The three-times-arc-melt process was described in [178], and was repeated here since the arc melter used in that paper had similar properties to our own. This paper, which is specifically about arc melted NiAl, also pressed the Ni and Al powders into pellets prior to arc melting, versus transferring loose powder to the crucible.




Figure 5.14: The MIT MNM arc melter in use. The user controls the position of the arc and the power supplied to the electrode. Welders' glass windows allow the user to see when the pellets have fully melted and formed a single button of molten material.



Figure 5.15: The five buttons made in the arc melter were heat treated in a high-temperature graphite furnace under an inert environment for 24h at 1200°C. The photograph shows the buttons arranged in the furnace prior to beginning the heat treatment. The purpose of the heat treatment was to grow large grains, so that the TGS measurements could be easily carried out within the boundaries of a single grain if the buttons were used for measurements. The furnace pictured here is in the Allinore lab at MIT.

Buehler® Vector™ Power Head mechanical polisher, followed by polishing with 3 μ m and 1 μ m diamond suspension. The final polishing step was performed with a 0.01 μ m alumina suspension.

Samples were then swabbed with a 2% Nital solution to etch the surface and make microstructures more visible. Finally, the samples were imaged using a Zeiss microscope outfitted with an Axiocam that can take digital metallographs of the samples. These metallographs are shown in Figures 5.16-5.21.

The grains depicted in these metallographs are already quite large, and it appears that some of the samples would have grains that could be large enough to fit a TGS spot (120 μ m²). Annealing is useful, however, for growing the grains even larger and ensuring more consistency between samples by giving them more comparable heating histories.

5.2.1.7 Growth of single crystals at Los Alamos National Laboratory

Material from the buttons was used to grow NiAl single crystals using a modified Czochralski process. The process was carried out by Kenneth McClellan at Los Alamos National Laboratory. The surface of each button was ground to remove oxides before beginning the process. Argon gas (flow rate 1 SL/min(air)) was used as the atmosphere for all five samples. The seed used for all five samples was internally labeled as




Figure 5.16: NiAl sample #1 at 50x magnification in a polarized field. The etchant swab created shallow “scratches” that make it difficult to see the sample’s features.




Figure 5.17: NiAl sample #2 at 50x magnification in a polarized field. As with #1, the etching process obscured many of the features on this sample. This is exacerbated by the small voids throughout sample that seem to have formed during melting.




Figure 5.18: NiAl sample #3 at 50x magnification in a bright field. It is easier to see the large grains, some of which are several hundred μm in width and thus capable of fitting the entire TGS spot.




Figure 5.19: A second image of NiAl sample #3 at 50x magnification in a bright field, showing the large, clearly delineated grains.




Figure 5.20: NiAl sample #4 at 50x magnification in a bright field. As was the case with Sample #2, small voids developed during the arc melting process.




Figure 5.21: NiAl sample #5 at 50x magnification in a bright field. Voids appear to be distributed through this sample as well, although they appear smaller than the voids observed in #1 and #4.

F2-NiAl-04, and had a $\langle 100 \rangle$ orientation. The maximum pull rate was 15mm/h. Additional parameters and notes related to the single-crystal-growth process for each sample are given in Table 5.3.

5.2.1.8 Preparation of single crystal NiAl samples

Slices several millimeters thick were sectioned from each single crystal rod using a low-speed diamond saw. These slices were then mounted in epoxy resin using a Struers Prestopress-3 mounting system and Buehler Epomet GTM powder. Samples were polished on a Buehler MetaServ 250TM with a Vector PowerheadTM.

The first phase of polishing involved polishing samples with progressively finer grits of sandpaper. Each polishing step lasted between 1-3 minutes; samples were checked approximately once a minute, since the goal was to avoid polishing samples for longer than necessary to minimize cold work effects and the chance of damage to the sample surface during the polishing procedure. When all samples looked uniform, the

Table 5.3: Growing single crystal samples at LANL from the arc-melted and annealed NiAl buttons

Sample No	Hearth/Seed [rpm]	O ₂ [ppm]	Ken's notes (Copied verbatim from his notes on each run)
1	25/35	0.8	Substantial oxide release on initial melting. Extensive powder debris deposited throughout chamber during growth.
2	35/10	3E-11	Moderate oxide on melt, some at initial seed-on from seed surface. Substantial powder deposited in chamber but less than for #1.
3	35/15	1E-13	Problems with seed on, had to melt back seed. Less deposit in chamber than for #2.
4	35/10	3E-15	Melted charge upside down, had oxide on initial melt surface. Facets visible along length but surface has oxide along entire length.
5	35/10	1E-15	Minimal oxide on charge at initial melting. Facets visible along length, kept some of crystal for seed extension.

sample holder and the samples were rinsed to remove grit, and the next sandpaper in the sequence was swapped in. Sandpapers used were of 120, 240, 400, 800, and 1200 grit, in that order. All sandpapers were purchased from Buehler.

Next, diamond suspensions (both Buehler MetaDi™) were used to polish the samples. A 3 μ m suspension was used, followed by a 1 μ m suspension. Finally, a 0.01 μ m alumina suspension (Buehler Masterprep™) was used to perform a final mirror polish. Note that each step had a dedicated polishing cloth to avoid cross-contamination of the grits. As before, the sample holder and the samples were rinsed thoroughly between steps.

5.2.1.9 Mixing of samples and loss of identification (Batch I)

The single crystal rods and the arc melted buttons were kept in labeled small boxes within a larger box. The larger box was knocked open, leading to the samples being mixed up. Since they all look similar, it was not possible to match the sample to their nominal composition. There was one set of mounted samples that had been labeled previously that were not affected. These samples were oriented via X-ray diffraction as described in Section 5.2.1.10.

The single crystal rods were given arbitrary labels. Samples from each single crystal rod were sent out for elemental analysis so that the samples could be re-identified, as described in Section 5.2.1.11.

5.2.1.10 Orienting the samples using X-ray diffraction.

One set of mounted NiAl single crystal samples were examined at the MIT Center for Materials Science and Engineering X-Ray Shared Experimental Facility using a Bruker® D8™ X-ray diffraction (XRD) system equipped with a General Area Detector Diffraction System (GADDS).

Care was taken to cut perpendicular to the axis of the rod, so that the measurement surface would be as close to the $\langle 100 \rangle$ direction as possible (this was the intended direction of the single crystals, as described above). Ideally, one would be able to cut the samples in conjunction with Laue backscatter diffraction to ensure sample orientation, but as option was unavailable, the samples were cut and then checked in XRD.

The samples were moved, tilted, and rotated relative to the detector and the X-ray beam with the goal of maximizing the intensity diffraction spot from the single crystal. Figure 5.22 shows the degrees of freedom with which the sample is moved, excepting the 2θ angle that describes the relationship between the sample, the X-ray beam, and the detector. Figure 5.23 shows one of the NiAl samples being measured on the XRD GADDS system and the relationship between the detector, the sample, and the X-ray beam. The angle between the detector and the sample is ideally θ , one-half of 2θ . If the diffraction spot is maximized at an angle that deviates from θ , this indicates some degree of miscut in the sample.




Figure 5.22: Degrees of freedom in the sample mount for the Bruker® D8 GADDS system, taken from the MIT CMSE SOP for the equipment.




Figure 5.23: A photograph of an NiAl sample in the XRD GADDS system. The wide-angle germanium detector is on the left and the X-ray source is on the right. The sample's cradle moves in x, y, z, ϕ, ψ , and ω . The entire cradle and the detector move independently to change the angle between the detector and the X-ray source (2θ) and the angle between the detector and the sample.

The results of these measurements are shown in Table 5.4. Figure 5.24 shows the direction of the cor-

Sample [at% Ni]	Color label	ω	ψ
42	PINK	-1.6°	11.8°
45	YELLOW	-1.4°	27°
48	BLUE	0°	-5°
50.6	GRAY	-1.05°	15.6°
52	PURPLE	0.7°	3.8°
54	GREEN	1.2°	12.1°

Table 5.4: Sample tilts as measured by XRD. Figure 5.24 illustrates ω and ψ relative to the mounted sample (oriented in ϕ). Each sample had a color label for easy identification.




Figure 5.24: The schematic shows the mounted sample and the corrections required to make the $\langle 100 \rangle$ direction normal to the measurement surface. ω involves rotating the sample from side to side, whereas ψ indicates a back-and-forth tilt.

rections that need to be made. Note that the corrections are made about an imaginary red dashed line: this is because the ψ and ω corrections are made after the sample has been rotated to its optimal ϕ value. When the dashed line is level, the sample has been rotated to maximize the diffraction spot. (In practice, a line is drawn on a piece of tape on the mount surface once ϕ is optimized using a straightedge aligned with the X-ray source while the sample is still attached to the XRD cradle.)

It was decided that only the ψ corrections would be made, since they were significantly larger than the ω offsets, and the precision of the method used to correct the miscut was on the order of a degree anyways.

To make the ψ corrections, a slice was cut from the back of the mounted sample, as shown in Figure 5.25. The slice has an angle of ψ . If a positive ψ offset was measured, the sample needed to be tilted “back” (see Figure 5.24) and the “slice” had its maximum height at the top of the sample (again, “top” is defined by moving the sample to its optimum rotation in ϕ). Next, the top of the mount mount is cut so that both faces of the mount are again parallel. The sample face is now truly in the $\langle 100 \rangle$ direction (in theory). (A similar procedure can be followed to correct for ω offset, only now, the “slices” go from left to right instead of top to bottom. Again, left and right are defined relative to the optimum ϕ rotation of the sample.)

In practice, these cuts were made using simple trigonometry and calipers. The width of the sample mount was measured at the top and bottom (as defined by the ϕ rotation), as well as the diameter of the mount. To find the height of the “slice” x we have:

$$\text{diameter} \cdot \tan\psi = x \quad (5.3)$$




Figure 5.25: A not-to-scale schematic of the process of cutting the sample mount to correct for misalignment of the crystal plane relative to the TGS measurement surface. Sample planes (shown as light blue lines) have a misalignment of ψ . By making a cut at an angle of ψ from the back of the mount, we bring the crystal planes into their proper alignment when the back of the mount is level. A second cut is then made because the measurement surface needs to be flat. The new surface is repolished prior to measurement.

The calipers are then used to mark a distance of x on the top side of the mount, and a cut is made from the corner on the “bottom” side to the mark. (This is for a positive ψ correction; for negative ψ corrections, the procedure is the same, but x is marked on the “bottom” of the mount and not the top, since the corrective tilt is in the other direction.)⁵

A special swivel arm assembly was purchased for the low speed diamond saw that allowed mounts to be rotated and tilted. This allowed for the alignment of the cut line with the saw blade.

Once the angled cuts are made, a second cut is made through the top of the sample, as shown in the third image of Figure 5.25. This creates a flat surface for TGS measurements. The samples are then repolished with the usual steps.

During these steps, two of the samples were lost: yellow (Ni-45) and purple (Ni-52). This was because the samples that were mounted were thin, and they fell from the mount during cutting or were ground away during polishing. If it is known when preparing samples for sectioning and mounting that this procedure will be performed, it is a good idea to allow for 4-5mm of sample height if possible. Thin sections may not have enough “room” for resectioning and repolishing, and are more likely to fall from the mounts. During the resectioning of the top of the mount, care should be taken to minimize the chance that the sample will be ripped or jostled from its position in the mount - low saw speeds are suggested.

⁵If the mount is found to vary in height prior to making the corrective step, this has to be corrected for during this step because x assumes a uniform cylindrical mount. If the “top” height is greater than the “bottom” height, mark the excess height first (measuring from the back of the mount), and then mark x from that point. If the “bottom” height is greater than the “top” height, mark the excess height and draw a straight line to the other side of the mount. x can then be measured from that mark.)

5.2.1.11 Composition analysis of single crystal NiAl samples (Batch I)

Samples from each of the five single-crystal rods, as well as from one of the Busso samples, was sent to Luvak, Inc. for compositional analysis via direct current plasma emission spectroscopy. The results were given in wt% and converted to at% upon receipt. These results are given in Table 5.5.

To convert from wt% to at%, the following formula is used for each element:

$$\left(A_X \left[\frac{\text{g}}{\text{mol}}\right]\right)^{-1} \cdot W_X[\text{g}] \cdot 6.022140857 \cdot 10^{23} \left[\frac{\text{atoms}}{\text{mol}}\right] / N_{total}[\text{atoms}] \quad (5.4)$$

where A_X is the atomic weight, W_X is the weight of the sample, and N_{total} is the total number of atoms. W_X is calculated by assuming that there is 1 g total of the sample, so a sample that is 67.9 wt% Ni is given a mass of 0.679g for the purpose of this calculation. Note that the total number of atoms for each element in the sample must be calculated first, so that they can be summed to find N_{total} , before completing the conversion to at% for each element.

Table 5.5: Composition of NiAl Batch I

SAMPLE	Luvak 1 [wt%]	Luvak 2 [wt%]	Luvak 3 [wt%]	Luvak 4 [wt%]	Luvak 5 [wt%]	Luvak 6 [wt%]
Ni	67.9	68.5	70.7	70.8	75.4	73.7
Al	32	31.4	29.2	29.0	24.4	26.1
Co	0.038	0.036	0	0.038	0.039	0.039
Fe	0.062	0.054	0.024	0.061	0.071	0.068
Si	0	0	0.071	0.013	0.028	0.031
sum	100	99.99	99.924	99.899	99.91	99.938
	[at%]	[at%]	[at%]	[at%]	[at%]	[at%]
Ni	49.34	49.78	51.38	51.45	51.10	53.56
Al	50.58	49.64	46.16	45.84	45.53	41.26
Co	0.03	0.03	0	0.03	0.27	0.03
Fe	0.05	0.04	0.02	0.05	0.46	0.05
Si	0	0	0.11	0.02	0.20	0.05
sum	100.00	99.48	97.66	97.38	97.57	94.94

5.2.1.12 Making a second set of samples from Batch I (Batch1B)

A second set of samples were sectioned, mounted, and polished from the Luvak-identified single crystal rods (+ the Busso sample). Since the composition of these samples were fully identified, they provide a useful comparison against the angled samples (especially since two of those were lost). These samples were not XRD-characterized and re-angled, however, in order to save time (and investigate the difference between angled and uncorrected samples in the TGS results).

5.2.2 NiAl sample fabrication: Batch II (Summer 2017)

5.2.2.1 Motivation for making a second batch of NiAl samples

Batch I consisted of two sets of NiAl samples, made from single crystals made at LANL using a modified Czochralski process using material from the arc melted buttons made in lab. Each set also includes a sample made from the NiAl used in [159]. Table 5.2.1.4 gives the intended compositions for the non-Busso samples. However, it is clear from Table 5.5 that there is a wide discrepancy between the intended compositions and the actual compositions.

First, note that “Luvak 3” is the Busso sample here; it is reported to have a composition of 51.31 at% Ni and 46.1 at% Al (with trace amounts of Fe and Si). The Busso sample’s reported composition is 50.6 ± 0.6 at% Ni and 48.4 ± 0.2 at% Al. The Luvak measurements of the nickel is nearly in agreement with this, considering the upper bound of the expected nickel concentration (51.2 at%). The measured at% Al is significantly less than reported.

Next, consider that the intended range of compositions in the other samples was 42 to 54at% Ni, with aluminum being the balance. The intention was to have samples that were very far in the nickel-poor, high-constitutional-vacancy-concentration region, samples that were very far into the aluminum-poor, high-nickel-antisite-defect-concentration region, and a few near the stoichiometric point. Instead, the samples were found to have compositions much closer to the stoichiometric point, with a nickel range from 49.34 at% Ni to 53.56 at% Ni.

This narrow range showed that it was necessary to try another method to make NiAl samples in order to get far from the stoichiometric point. Simply measuring the correct amounts of nickel and aluminum was insufficient to achieve the desired results. At any rate, these compositional analysis results show the necessity of confirming compositions after sample fabrication, regardless of fabrication method.

The primary goal of fabricating Batch II was to expand the compositional range of the NiAl samples, since the compositional analysis of Batch I revealed that the sample compositions were clustered closer to the stoichiometric point than intended. These samples were heat treated to grow larger grains, but were not used to make single crystals.

5.2.2.2 Making the arc melted buttons for Batch II

The same Ni and Al powders described in Section 5.2.1.2 were reordered. The same arc melting procedure described in Section 5.2.1.4 was followed. However, a brand new copper hearth was ordered and used for arc melting the Batch II NiAl buttons. The Luvak analysis revealed that cobalt, iron, and silicon was present in the single-crystal Batch I samples made from the NiAl buttons arc-melted in lab (see Table 5.5.)⁶

Table 5.6 gives the amounts of Ni and Al powder measured for each sample in Batch II, and the at% that

⁶The sample labeled Luvak 3 was fabricated at MIT for a previous project (the “Busso” sample), and did not contain any cobalt. It did contain iron, and the highest amount of silicon of any of the six samples used for Batch I. All other Batch I samples were made from in-lab arc-melted NiAl buttons.

these amounts should correspond to. The conversion from powder weight to atomic percentage of Ni and Al were again calculated using equations (5.1) and (5.2). A range of at% was planned, from at% Ni = 40-60%.

Table 5.6: Planned compositions for NiAl Batch II

Sample №	Al [g]	Ni [g]	at% Ni	at% Al
1A	6.0174	4.0820	40.39	59.61
1B	5.9391	4.0609	40.20	59.80
2	6.1989	3.8867	42.30	57.70
3	7.2385	2.7955	45.66	54.34
4	6.4092	3.5945	45.05	54.95
5	6.4100	3.5900	45.08	54.92
6	6.2929	3.7007	43.87	56.13
7	6.6034	3.4033	47.14	52.86
8	6.6758	3.3227	48.01	51.99
9	7.3129	2.7341	55.15	44.85
10	7.4589	2.5779	57.08	42.92

The same arc melting procedure used to make Batch I was performed for each sample. The samples were annealed in an inert environment at 1200°C for 36h.

5.2.2.3 Compositional analysis of Batch II samples

Table 5.7: Composition of NiAl Batch II

Sample №	at%				<i>planned at%</i>		wt%				
	Ni	Al	Fe	Si	<i>Ni</i>	<i>Al</i>	Ni	Al	Fe	Si	Sum
1A	41.18	58.74	0.04	0.05	40.39	59.61	58.1	38.1	0.05	0.03	96.28
1B	41.98	57.95	0.03	0.04	40.20	59.80	58.3	37.0	0.04	0.03	95.37
2	44.11	55.76	0.08	0.04	42.30	57.70	60.4	35.1	0.11	0.03	95.64
3	44.51	55.37	0.06	0.05	45.66	54.34	61.2	35.0	0.08	0.04	96.32
4	47.38	52.53	0.05	0.04	45.05	54.95	62.4	31.8	0.07	0.02	94.29
5	47.73	52.16	0.07	0.04	45.08	54.92	62.5	31.4	0.09	0.02	94.01
6	45.58	54.32	0.08	0.03	43.87	56.13	60.6	33.2	0.10	0.02	93.92
7	49.01	50.90	0.06	0.04	47.14	52.86	64.3	30.7	0.07	0.02	95.10
8	49.97	49.94	0.06	0.04	48.01	51.99	65.3	30.0	0.07	0.02	95.40
9	57.47	42.43	0.07	0.02	55.15	44.85	71.3	24.2	0.09	0.01	95.60
10	59.63	40.31	0.04	0.02	57.08	42.92	73.7	22.9	0.05	0.01	96.66

Table 5.7 shows the compositional analysis results as performed by Luvak, Inc. via direct plasma emission spectroscopy. The range of Ni and Al at% was much wider: 41.2% to 59.6%, versus 49.3% to 53.6%

for Batch I. Samples 1A and 1B are likely slightly outside of the B2-phase NiAl region.

It is also now possible to directly compare the planned Batch II composition matrix with the actual composition matrix. (This was hard to do on a sample-to-sample basis for Batch I, due to the incident in which the samples were all dropped and jumbled prior to compositional analysis, so it was no longer possible to identify which sample corresponded to which arc melt, each of which had an intended composition.) It's evident that it's difficult to wind up with the precise desired concentration, although the real goal - a sample set that spanned the full range of NiAl compositions - was achieved. One of the primary drawbacks of Batch I was that the nickel-poor samples were intended to have a range beginning at 42 at%, but the lowest nickel content achieved was more than 49%. This was a significant issue because a primary goal of this study is to investigate the SAW signal's dependence on vacancy content, and constitutional vacancy concentrations are expected to be more significant the more nickel-poor the NiAl sample is.

A recommendation for making any future NiAl batches is, therefore, to plan a sample matrix with a wider-than-needed composition (e.g. planning NiAl samples with at% Ni starting at 40, even though that's outside of the NiAl range), and to order enough raw material to create many samples. This increases the chances of winding up with a good number of samples in the desired range. In Batch I, we made too few samples, started the sample matrix at too high an atomic percentage of nickel. In particular, it seems easier to make Ni-rich (ASD-rich) samples than Ni-poor (vacancy-rich) samples. When studying constitutional vacancies is important, like here, it's wise to plan more samples on the Ni-poor side in order to maximize chances of getting vacancy-rich samples.

5.2.2.4 Metallurgical examination of the post-heat-treatment Batch II samples

A low-speed diamond-blade saw was used to cut sections from the arc-melted and annealed buttons. These were mounted in Bakelite and mechanically polished following the same procedures that were used to prepare the Batch I samples for metallography.

All eleven samples contained trace amounts of iron and silicon. Iron content ranged from 0.03 to 0.08 at% and silicon content ranged from 0.02 to 0.05 at%.

Samples were not etched prior to metallography, so the grains are harder to distinguish as compared to the post-arc-melt metallographs of Batch I, which were etched. The samples used to take metallographs were used for TGS as well, so the etchant step was foregone to better preserve the finished surface.

Metallographs are presented in order of nickel content, not the order of their assigned numerical labels. (Labels were not re-ordered after compositional analysis in order to ensure consistency during experiments.) Compositions reported in the captions are given in at%.

Figures 5.26-5.28 show metallographs from Sample 1A, which is one of the samples that is expected to not be B2-phase NiAl due to its low nickel content. (Based on the NiAl phase diagram (Figure 5.2), it seems likely that samples 1A and 1B have the Al_3Ni_2 structure. Since they are made, though, they are included in the TGS analysis out of general interest: for example, is there a clear change in TGS results at the phase transition?




Figure 5.26: 1A, bright field (41.18 Ni, 58.73 Al, 0.04 Fe, 0.05 Si)




Figure 5.27: 1A, dark field (different region). The large voids are the most distinctive features here.




Figure 5.28: 1A, polarized (same region as dark field)




Figure 5.29: 1B (41.98 Ni, 57.95 Al, 0.03 Fe, 0.04 Si) polarized. Grains are large and very distinct under polarized light for the samples with < 42 at% Ni.




Figure 5.30: #2, bright field (44.11 Ni, 55.76 Al, 0.08 Fe, 0.04 Si). Large population of mid-sized (10-20μm) voids.




Figure 5.31: #2, dark field (different region)




Figure 5.32: #2, polarized (same as dark field). Grain boundaries are more evident, but look distinctly different from those observed in Nos. 1A and 1B.




Figure 5.33: #3, bright field (44.51 Ni, 55.37 Al, 0.06 Fe, 0.05 Si) **Figure 5.34:** #3, dark field (different region) **Figure 5.35:** 3, polarized (same region as dark field)




Figure 5.36: #6, bright field (44.58 Ni, 54.32 Al, 0.08 Fe, 0.03 Si) **Figure 5.37:** #6, dark field (different region) **Figure 5.38:** #6, polarized (different region)




Figure 5.39: #4, bright field (47.38 Ni, 52.53 Al, 0.05 Fe, 0.04 Si (at%)) **Figure 5.40:** #4, dark field (different region) **Figure 5.41:** #4, polarized (different region)




Figure 5.42: #5, bright field (47.73 Ni, 52.16 Al, 0.07 Fe, 0.04 Si (at%)) **Figure 5.43:** #5, dark field (different region) **Figure 5.44:** #5, polarized (different region)




Figure 5.45: #7, bright field (49.01 Ni, 50.90 Al, 0.06 Fe, 0.04 Si) **Figure 5.46:** #7, dark field (different region) **Figure 5.47:** #7, polarized (different region)




Figure 5.48: #8, bright field (49.97 Ni, 49.94 Al, 0.06 Fe, 0.04 Si) **Figure 5.49:** #8, dark field (different region) **Figure 5.50:** #8, polarized (different region)




Figure 5.51: #9, bright field (57.47 Ni, 42.43 Al, 0.07 Fe, 0.02 Si)




Figure 5.52: #9, dark field (different region)




Figure 5.53: #9, polarized (different region)




Figure 5.54: #10, bright field (59.63 Ni, 40.31 Al, 0.04 Fe, 0.02 Si (at%))




Figure 5.55: #10, polarized (different region)

Figure 5.26 reveals significant porosity in the sample, which tracks with general observations of the arc-melted buttons 1A and 1B: they feel less dense than samples 2-10, and are brittle and fragile - the porosity is similar to that of scoria, but the fragility is reminiscent of pumice. There appear to be two populations of bubbles - one large (20-50 μm diameter) and one much smaller (perhaps 0.1-1 μm diameter). Grain boundaries are easily visible in the bright field image, and the larger bubbles appear to occur on the borders or at the vertices of the grains, rather than in the grain bulk. Dark field imaging of an adjacent spot on the sample surface reveals that the samples are quite smooth between the large bubbles, indicating that TGS should be possible, provided the spot is centered on a pore-free spot. If we want a 200 μm diameter spot, it should be possible to find multiple suitable TGS measurement spots on this sample.

Figure 5.28 is a polarized view of the same sample area in the dark field metallograph. These polarization results are far more dramatic than typically observed for NiAl samples, likely a result of grains that exist at very different orientations relative to each other, further indicating that this sample isn't B2-phase NiAl. (Higher nickel-content samples reveal grains that are much thinner and elongated, and far less dramatically distinguishable under polarized light.)

Sample 1B was created from the same powder mix as Sample 1A, and has only a slightly different composition. Figure 5.29 shows that Sample 1B exhibits the same dramatic polarization behavior as Sample 1A, suggesting that they have the same crystal structure.

Figure 5.30 shows that Sample 2 has significant porosity as well, but less evident grain boundaries in the bright field view. The polarization behavior is distinctly different from that of Samples 1A and 1B: the grains are more angular and irregular in shape and the polarization variation across the surface is less intense. The porosity again indicates that care will need to be taken to find measurement spots that are sufficiently large and unmarred. At 44.11 at% Ni, Sample 2 is expected to be firmly in the compositional range of B2-phase NiAl.

Figure 5.33 is a bright field metallograph of Sample 3 taken near the sample slice's edge. There is an absence of the larger pores that characterize previous samples but the smaller pores seem to have organized in vein-like structures near the edge. The dark field shows surface irregularities that may require additional polishing before TGS measurements are carried out. It is difficult to distinguish individual grains in this sample.

Sample 6 is presented next, since it happened to have the next highest nickel content. Again, it is difficult to determine grains, although it is possible to see what appear to be several grains in relief in the polarized image (Figure 5.38).

Metallographs of Sample 4 (Figures 5.39-5.41) again indicate that care will need to be taken in finding appropriate measurement areas due to the dense distribution of pores - or that a large number of measurements will need to be taken in order to cancel out any unusual results from measurement spots that contain an unusually high number of voids. The bright field metallograph in Figure 5.39 shows distinct grains, which are easiest to see in the right-hand side of the image. The grains have a thin, extended appearance.

Sample 5 (Figures 5.42-5.44) has a smoother appearance than the other samples listed thus far in the

dark field metallograph, although the bright field image still shows porosity (if to a lesser extent than the other samples). Grains are more visible, particularly in the polarized metallograph. The grains continue to have an elongated appearance, consistent with the grains observed in the Batch I samples.

Sample 7 (Figures 5.45-5.47) came from a smaller button than the other samples, and the bright field metallograph indicates that the sample may not have been perfectly level, and thus parts of the sample did not receive the final polishing step. The dark field photograph shows interesting porosity, with lines of small voids arranged along right angles to each other (lower right-hand side of Figure 5.46). The polarized metallograph reveals distinct large grains. The large size of the grains observed here, and in the other samples, indicate that a larger number of TGS spots may be needed to achieve a representative view of the sample's material properties, as NiAl is anisotropic and it is possible (and in some cases, likely) that a TGS phase grating will be contained entirely, or nearly entirely, by a single grain. The porosity of Sample 7 is relatively mild, except for several very large voids near the sample edges.

Sample 8 (Figures 5.48-5.50) also has only mild porosity compared to the other samples. (It is interesting to wonder if this has to do with Samples 7 and 8 being close to the stoichiometric point, and thus perhaps more inclined to have a regular crystal structure that is less prone to defect development.) The grains are distinguishable in relief in the bright field metallograph and have the same narrow structure that seems to be typical of B2-phase NiAl samples that have been formed by arc-melting followed by a long anneal at 1200°C.

Sample 9 (Figures 5.51-5.53) has fewer large flaws and a mild porosity. Grains are visible in relief, but are difficult to distinguish even under polarized light, and continue to have the usual long narrow structure.

Sample 10 (Figures 5.54-5.55) has more porosity than Sample 9, making it hard to distinguish individual grains in the bright field micrograph. They are easier to distinguish in the polarized metallograph.

5.3 TGS of intermetallic NiAl

This section first examines results from the initial experiments performed on the Batch I single crystal samples, which helped shape the rest of the experimental campaign. Next, results of measuring v_{SAW} of the Batch I and II samples are examined. This is followed by the results and analysis of the thermal diffusivity measurements for the same samples.

5.3.1 Initial $v_{SAW}(\theta)$ measurements of Batch I and lessons learned

The initial set of TGS measurements on the Batch I samples was carried out in the summer of 2015. However, there were several issues with this data set. First, the planned compositions had a much wider compositional range than was actually achieved. This was realized later during the first Luvak analysis. Second, the samples were single crystal and as such the TGS measurements were orientation dependent. The samples were only moved in x and y , and were not rotated, so the measurements were not representative of the Batch I sample properties.

After the labeled Batch I samples were dropped, many (unsuccessful) attempts were made to identify them experimentally. The eventual ICP measurements carried out by Luvak revealed that the sample compositions were different than we believed anyways. Furthermore, during the re-angling of the Batch I samples to better align the crystallographic planes, two of the samples were ruined.

So, a second set of samples was prepared from the Batch I single crystal rods. These samples were able to be correlated directly to the Luvak measurements. Figure 5.56 shows $v_{SAW}(\theta)$ results for this set of NiAl samples. They were left uncorrected for possible miscut and misalignment. The data in Figure 5.56 has been shifted so that the first maximum in each data set is aligned. Interestingly, the samples do not appear to display the same periodicity. To take these measurements, each mounted specimen was placed on a rotating stage for measurement. The stage allowed for specific control of the degree of rotation between measurements. (See page 105 for more information.)

Two to four measurements were taken at different spots on the sample at each rotation position. The error associated with these measurements are not graphed here in order to make the data easier to see. The errors associated with the frequency measurements (e.g., even before accounting for slight differences in the calibrated grating spacing) were small: The average error associated with the data labeled Luvak 1 was 0.12%; for Luvak 2, 0.14%; for Luvak 3, 0.30%; for Luvak 4, 0.21%; for Luvak 5, 1.04%; and for Luvak 6, 0.67%. These errors are far too small to be a possible reason for some of the interesting differences observed here - e.g., the $\approx 40\%$ difference in wave speed between the minimum wavespeed of Luvak 2 and the minimum wavespeed of Luvak 5.

Figure 5.57 shows the SAW frequency vs. rotation results for one of the NiAl samples (“blue”) that was corrected via XRD alignment and reangling, but whose composition was unknown due to the mixup of the samples. It is compared with the SAW frequency vs. rotation results for the NiAl samples of known composition that were prepared from the same Batch I single crystal stock, but which were left uncorrected.




Figure 5.56: SAW speed measurements on NiAl of varying compositions taken at a nominal $5.5\mu\text{m}$ grating spacing. These samples were not corrected for possible miscut or misalignment of the crystal planes. Data was shifted in degree so that the first maximum in each data set was matched.

The original goal of this was to identify the composition of the corrected samples by matching them to responses from the known samples, but it isn't obvious how much of an impact the angling procedure had. It appeared that the "blue" sample matched best with Luvak 4 and possibly Luvak 5, but the period of the identified samples was slightly longer than that of the unknown one.

Figure 5.58 shows the SAW frequency vs. rotation results for another corrected but composition-unknown sample, compared with with SAW frequency vs. rotation results for six uncorrected but composition-known samples. Based on these results, it would appear that the "pink" sample is probably from the same rod as Luvak 5. If so, it doesn't appear that the angling procedure was necessary.

These measurements were undertaken because it seemed necessary to identify the dropped samples due to the time involved in carefully orienting them via XRD and re-angling them to the proper crystallographic direction. However, the measurements revealed that there probably wasn't that much difference between the uncorrected and the corrected samples. This meant that it was possible to use the uncorrected samples for the rest of the analysis in this chapter, all of which had known compositions.⁷

While this particular set of measurements yielded relatively little important material property information, it wasn't wasted labor. Most of these measurements were carried out fairly early in the experimental

⁷One thing that made the angle corrections challenging was the fact that the Laue backscatter system at the MIT CMSE was not operational during this work. The samples were aligned using the XRD GADDS system, and then angled *ex situ* based on those results. Ideally, one is able to use a system (such as an adapted Laue backscatter setup) that allows for careful alignment of the sample over a low-speed diamond saw within the X-ray system itself, ensuring an accurate re-alignment.




Figure 5.57: An attempt to match the unknown, crystallographically corrected “blue” sample to a sample of known composition was made. The best match was Luvak 4. If Luvak4 and “Blue” were from the same sample, then this indicates that the angling step was probably unnecessary.




Figure 5.58: An attempt to match the unknown, crystallographically corrected “pink” sample to a sample of known composition was made. The best match was Luvak5. As with Figure 5.57, if we can assume that the unknown sample comes from the same single crystal rod as the identified sample whose response it most closely matches, then it would appear that the angling steps had little effect on the frequency (and therefore the speed) response.

One polycrystalline sample (Sample 5, Batch II, Ni at% 47.73%) was measured at $4.8\mu\text{m}$ and $6.4\mu\text{m}$ grating spacings in order to check that the disparity between measurements was small, and that multiple grating spacings were not needed across the board in order to successfully compare the measurements relative to each other. The Batch I samples are not orientation-corrected.

The code in Appendix 8.7 was used to take the Fourier transform of each SAW trace and isolate the dominant frequency of the samples. This frequency was then multiplied by the calibrated grating spacing to obtain SAW speed. Error bars are only shown for the polycrystalline Batch II samples in Figure 5.59; for all other data points, the error bars were small enough to overlap with the markers.

5.3.2.2 Analysis

At first glance, Figure 5.59 shows what appears to be a linear decrease in SAW speed as nickel concentration increases. However, it is questionable (especially based on the optical microscopy analysis in Section 5.2.1.6) whether the two leftmost data points belong to the B2-phase. Furthermore, the rightmost data point indicates that the downward trend may not continue at high concentrations of nickel. If this study was to be expanded, it might be worthwhile to fabricate and measure more polycrystalline NiAl samples with compositions between 50 and 60 at% Ni.

However, the primary reason that NiAl was measured in the first place was its high constitutional vacancy concentrations at low nickel concentrations. A linear fit for the data with at% Ni 44 to 50 - data from samples that were B2-phase NiAl and which were expected to have high C_v - is shown in Figure 5.60. The equation for the linear fit to this data is:

$$v_{SAW} = -83C_{Ni} + 7300 \quad (5.5)$$

with C_{Ni} = at% Ni and v_{SAW} in m/s.




Figure 5.60: Linear fit to the v_{SAW} data for 44 to 50 at% Ni

An initial attempt can be made here to link changes in the elastic modulus of the NiAl to vacancy concentration. First, it is assumed that for at% Ni < 50%, the offset from stoichiometry is accommodated entirely by vacancies on the Ni sublattice. That is, $C_v = 50 - C_{Ni}$.⁸ (For the purposes of this analysis, concentrations are given as their percent, and not as a value from 0-1, in order to remain consistent with the x-axis of the TGS data.)

We assume that v_R , as described in Chapter 3, is equivalent to v_{SAW} . That is,

$$v_R = v_s(0.874 + 0.196v - 0.043v^2 - 0.055v^3) \quad (5.6)$$

$$v_s = \sqrt{\frac{E}{2\rho(1+v)}} \quad (5.7)$$

Equation (5.5) can be set equal to Equation (5.6), such that

$$-83(C_{Ni}) + 7300 = \sqrt{\frac{E}{2\rho(1+v)}}(0.874 + 0.196v - 0.043v^2 - 0.055v^3) \quad (5.8)$$

$C_{Ni} = 50 - C_v$ in this analysis, so this can be substituted into the equation above to get an expression in terms of C_v . v is taken to be 0.33 for NiAl and ρ is taken to be 5900 kg/m³, allowing for simplification of the expression.

$$-83(50 - C_v) + 7300 = \sqrt{\frac{E}{15694}}(0.9320) \quad (5.9)$$

Solving for E yields

$$E = 1.8068 \times 10^4(9.9225 \times 10^6 + 5.229 \times 10^5 C_v + 0.6889 C_v^2) \quad (5.10)$$

This equation is plotted in Figure 5.61. The Young's modulus values calculated from Equation (5.10) are in reasonable agreement with values reported in the literature, which range from 95 to 270 GPa [179]. This empirical relationship indicates that TGS is sensitive to changes in E that arise as a result of changes in vacancy concentration.

These results also suggest that vacancy concentration may have a sufficiently strong effect on E (and thus, f_{SAW} and v_{SAW}) that the contributions of thermal and radiation-induced vacancies that are present in a sample during irradiation (but not during *ex situ* analysis of a sample) cannot be ignored when TGS is used *in situ*. However, additional work remains to be done to determine if constitutional vacancies affect E differently than mobile thermal and radiation-induced defects.

⁸This is a somewhat facile approach, and according to the data in Figure 5.6, may actually underestimate the concentration of vacancies. However, $C_v = 50 - C_{Ni}$ could be substituted for another expression that expresses vacancy concentration as a function of nickel concentration (or at%) by fitting the literature data or experimentally determining the vacancy concentrations in these samples.




Figure 5.61: Young’s modulus of NiAl versus vacancy concentration in NiAl, calculated from Equation (5.10).

The results of Figure 5.61 are at odds with initial expectations based on literature. [180] and [36] indicated that vacancies should lead to an overall reduction of E in simple metals. However, their considerations were not extended to intermetallic materials. It is also acknowledged in [180] that E tends to increase, not decrease, when interstitials are present alongside vacancies in significant contributions, indicating that the expected decrease in E due to a concentration of vacancies in the bulk can be overridden by the presence of other defects.

Hardness is considered as a function of vacancy concentration in intermetallic FeAl in [181]. FeAl has the same CsCl structure type as NiAl, and also exists as FeAl beyond its stoichiometric composition, but it tends to form thermal vacancies instead of constitutional vacancies [182]. Figure 5.62 shows the microhardness of various compositions of FeAl versus vacancy concentration, and shows that it increases with vacancy concentration. Hardness and Young’s modulus are not always directly linked, as Young’s modulus is an intrinsic material property and hardness will depend on the microstructure of the material surface being measured, so this should not be considered as any sort of obvious “proof” that Figure 5.61 is truly representative of NiAl behavior. However, Young’s modulus and hardness are generally observed to increase with each other [183, 184]. So, Figure 5.62 lends credence to the idea that E could increase with vacancy concentration in a B2-phase intermetallic material.

Figure 5.63 shows Young’s modulus as calculated in [175] from single crystal data. These calculations support the observation here that Young’s modulus increases with decreasing nickel content (and thus increasing constitutional vacancy concentration). It is also noted in [175] that it is difficult to obtain intermetallic NiAl crystals with controlled compositions and perfectly random texture that allow for a thorough comparison of experimental measurements with calculated predictions for the elastic properties of the intermetallic.




Figure 5.62: Microhardness increases with vacancy concentration in intermetallic FeAl. [181]




Figure 5.63: Microhardness increases with vacancy concentration in intermetallic FeAl. [181]

5.3.3 TGS measurements of thermal diffusivity as a function of composition

5.3.3.1 Results




Figure 5.64: Thermal diffusivity as a function of at% nickel in NiAl polycrystalline and single crystal samples.

Figure 5.64 shows the results of the thermal diffusivity analysis of the same TGS signal traces analyzed in the previous section. Error bars are not shown because the associated thermal diffusivity error calculated by the code was on the order of $10^{-7} - 10^{-8} \text{ m}^2/\text{s}$.

5.3.3.2 Analysis

The two left-most polycrystalline samples (1A and 1B from Batch II) have significantly higher thermal diffusivities than the other samples examined here, but this is likely further indication that they are not B2-phase NiAl. This is a good example of why the multi-property measurement allowed by TGS is so useful:

here, the thermal diffusivity results elucidate something that the SAW speed did not. In Figure 5.59, these points seem like they fit the general downward trend in v_{SAW} observed in the Ni-poor samples as at% Ni increases; there's nothing obvious to indicate that their properties might be significantly different.

The polycrystalline samples shows a steady increase from 44 to just before 50 at% Ni (the direction of decreasing vacancy concentration). A linear fit to this data is shown in Figure 5.65. This linear fit can be used to create a rough empirical model for thermal diffusivity in B2-NiAl as a function of vacancy concentration by substituting $50 - C_v$ for C_{Ni} in the linear fit equation.

$$\alpha = 1.04 \times 10^{-5} - 2.8 \times 10^{-7} C_v \quad (5.11)$$

Another approach for estimating how sensitive TGS-measured thermal diffusivity might be to changes in vacancy concentration uses literature-reported vacancy concentrations for B2-NiAl instead of using the rough $50 - C_v = C_{Ni}$ estimation above. Figure 5.5 shows several such data sets. If a linear fit is performed on the data set labeled "Kogachi1996 updated", we obtain

$$C_v = -0.7566 C_{Ni} + 47.359 \quad (5.12)$$

This can then be used to obtain an expression for C_{Ni} in terms of C_v that can be substituted into the initial linear fit of Figure 5.65. To calculate sensitivity of thermal diffusivity to vacancy concentration, we pick an arbitrary value of C_{Ni} and calculate α (these will be C_{Ni}^A and α^A). Next, we consider the average error associated with the data points plotted in Figure 5.65 - about 0.7%. Now, we calculate $\alpha^B = 1.007\alpha^A$. C_{Ni}^B is the nickel concentration that corresponds to α^B , as calculated from Equation (5.11). Finally, we calculate C_v^A and C_v^B by plugging C_{Ni}^A and C_{Ni}^B , respectively, into Equation (5.12). Sensitivity is the difference between them, and it is calculated here to be $\approx 0.18\%$.

A sensitivity of 0.18% - e.g. a significant change in the TGS-measured thermal diffusivity will be measured once a 0.18% change in vacancy concentration has occurred. This corresponds to a very high vacancy concentration (as mentioned before, the maximum expected radiation-induced vacancy concentration a sample can sustain is about 1% [185]), indicating that - as it currently stands - this analysis would only be useful at extremely high-damage conditions. However, the sensitivity could likely be improved by taking more data points on each sample. The data points in Figure 5.65 are the average of ten spots per sample, so the uncertainty of 0.7% could be reduced by increasing the number of memberships. It would also be of interest to experimentally, definitively determine the vacancy concentration in these samples, so that there could be confidence in the function that describes TGS-measured thermal diffusivity as a function of vacancy concentration. Right now, this sensitivity calculation is predicated on an estimation of C_v based on literature-reported values of $C_v(C_{Ni})$.

Sample 8 (Batch II), which is the closest to the stoichiometric composition (based on the Luvak analysis), shows a jump in thermal diffusivity. This may be due to the fact that the stoichiometric composition of NiAl is expected to have the lowest defect concentration. On the nickel-rich side, the Batch II samples appear to




Figure 5.65: Linear fit to the thermal diffusivity data for samples with composition 44-50 at% Ni.

have generally lower thermal diffusivities, although it is hard to discern any trends since there are only two polycrystalline samples with at% Ni > 50%.

The single crystal samples are clustered near the stoichiometric point, and the thermal diffusivities of those on either side of this point are very high. This is not surprising, because grain boundaries in polycrystals act as defects that scatter electrons, decreasing thermal diffusivity.

5.3.4 TGS measurements of acoustic damping in NiAl

5.3.4.1 Results

Results of the acoustic damping analysis of the TGS data collected for Batch I and Batch II NiAl samples is shown in Figure 5.66.




Figure 5.66: Acoustic damping as a function of at% nickel in NiAl polycrystalline and single crystal samples.

5.3.4.2 Analysis

The two most nickel poor samples (1A and 1B) again appear to have behavior that is different from the NiAl samples with at% Ni greater than 44%. It might be possible to use acoustic damping to investigate metallic phase changes by looking for discontinuities of the kind that appear to exist between the data from Samples 1A and 1B and the sample at 44 at% Ni.

Acoustic damping appears to increase as the stoichiometric point is approached from the Ni-poor side. This is contrary to what is expected - that acoustic damping would decrease with decreasing defect populations. The single crystal samples in Batch I also show a very wide spread of acoustic damping values, making it difficult to draw any conclusions from this data set.

However, as has been noted in the previous chapter, there are many reasons that we do not have high confidence in the current method of calculating acoustic damping from the TGS data. Acoustic damping was a more useful parameter in the irradiated niobium data campaign because the same data collection procedure was being repeated on the same samples each time, only varying one parameter (radiation dose). So, it is possible to observe relative changes between acoustic damping measurements, but is more difficult to use it to draw conclusions based on one measurement set from a given sample.

5.4 Simulating TGS experiments on NiAl in LAMMPS

LAMMPS, or the Large-scale Atomic/Molecular Massively Parallel Simulator, is an open-source molecular dynamics (MD) program that was developed at Sandia National Laboratories [186]. Extensive online documentation of the program can be found at lammmps.sandia.gov. In an MD simulation, a test structure or group of atoms whose properties and compositions are specified by the user interact with each other for a set amount of time in accordance with Newton's equations of motion.

The interaction between atoms in the simulation is defined by the potential function. The development of the potential function for B2-NiAl used in this research is described in [171]. The authors used experimental data and *ab initio* results to create an embedded-atom potential optimized for describing B2-NiAl.⁹

The ultimate goal of this phase of the research was to simulate TGS experiments on B2-NiAl of varying compositions in LAMMPS and compare the SAW speeds in the TGS experiments to the SAW speeds measured experimentally. To simulate TGS in LAMMPS, a test structure is first constructed, and a free surface is specified (Figure 5.67). Next, two regions on this free surface are specified in the code (Figure 5.68). The center of mass (COM) of each region is calculated. Next, a sinusoidal heat pulse is applied to the free surface (Figure 5.69). This heat pulse mimics the heat energy imparted by the pulsed laser grating in the experiments and induces a simulated standing acoustic wave. Finally, the program tracks the movements of the COM of each of the specified region at each timestep (the length and increments of which are defined by the user). By determining $\Delta z = z_{COM_2} - z_{COM_1}$ at each timestep, and plotting Δz vs time, a simulated TGS signal can be generated. This signal can then be analyzed according to the usual procedures to extract frequency information.




Figure 5.67: The TGS simulation begins with a structure with a free surface.

Figure 5.68: Two regions on the free surface are specified for center-of-mass tracking.

Figure 5.69: A sinusoidal heat pulse is applied to the free surface.

⁹This potential can be downloaded at <https://www.ctcms.nist.gov/potentials/Al-Ni.html>

5.4.1 Building the NiAl test structures

5.4.1.1 Initial tests

In order to simulate TGS experiments in LAMMPS, it is necessary to create test NiAl structures that accurately simulate the B2-phase NiAl structure.¹⁰ Figure 5.5 shows how the lattice parameter of B2-phase NiAl changes with composition: it increases steadily from about 2.868Å to 2.8887Å from 43 to 50 at% Ni, and then decreases steadily at a slightly lower rate on the opposite side of the stoichiometric point.




Figure 5.70: Initial relaxation of nickel and aluminum test structures in LAMMPS was unsuccessful. The test structures were random mixes of Ni and Al, but did not replicate the actual structures of B2-phase NiAl. Instead, lattice parameter simply grew with the increasing nickel content, as nickel is the larger atom.

In initial LAMMPS tests, structures of 2000 atoms (1000 atoms per sublattice) were built directly in the LAMMPS input function by specifying a percentage of Ni, a percentage of Al, a total number of atoms, and the nickel aluminum potential and relaxing the structure to see if it replicated lattice parameters observed in the literature. (See page 247 for an example LAMMPS input script. In this case, LAMMPS randomly places the nickel and aluminum atoms into lattice sites, but not in the regular order exhibited in B2-phase NiAl. Examination of the relaxation results to determine the final lattice parameter of these relaxed structures showed

¹⁰The word “structure” was used in this research instead of “ensemble,” since the test lattices were deliberately built in a certain way prior to relaxation and TGS simulation, and so “structure” seemed more illustrative for this work.

that these structures did not accurately replicate the behavior of B2-phase NiAl of varying compositions. These results are shown in Figure 5.70. Although the simulated structures had lattice parameters that got bigger with increased nickel composition - as observed in the literature for NiAl with at% Ni < 50% - they were larger than any lattice parameter observed in the literature. The misfit was worse closer to the stoichiometric point, which was the composition for which the input in [171] was optimized. Clearly, this method did not result in simulations that accurately reflected NiAl behavior. Sections 5.4.1.2 and 5.4.1.3 describe the methods used to create test structures that were more representative of reality.

5.4.1.2 Convergence tests

First, a series of convergence tests were run to determine what size test structure to use in the TGS simulations. A series of NiAl test structures with perfect stoichiometric structure of varying size were relaxed in LAMMPS (meaning the system was allowed to come to equilibrium) and their half-lattice parameter in the relaxed state measured. These results are shown in Figure 5.71. For small structures, the relaxed half lattice parameter changes significantly as more atoms are added, and so test structures this small would not be considered to yield reliable results. For reference, the commonly accepted half-lattice parameter of the B2-NiAl stoichiometric composition is 1.44Å.

When the structures have on the order of 10^6 atoms, the rate of change of the half-lattice parameter versus structure size begins to slow down. As structures increase in size, their relaxed structure half-lattice parameter approaches 1.44Å. Based on these results, it was determined that the test structures should have at least 2×10^6 atoms (1 million atoms on each sublattice). A more ideal test structure size would be double that, but the choice of 2×10^6 atoms was made because it resulted in a half-lattice parameter that was still reasonably close to the converged value while considering the computational cost of building, relaxing, and simulating TGS on many large structures.

5.4.1.3 Making a tool to build NiAl test structures with appropriate concentrations of Ni, Al, vacancies, and anti-site defects

Instead, it's necessary to ensure that the test structures both have the necessary composition and obey the structure of B2-phase NiAl. In order to do this, MATLAB was used to build large test structures with interlocking simple cubic lattices of Ni and Al. Ni atoms are then randomly removed from the Ni sublattice to obtain the necessary Ni and Al atomic percentages and corresponding vacancy concentration for Ni-poor compositions, and Al atoms on the aluminum sublattice are randomly reclassified as Ni atoms to obtain the necessary Ni and Al atomic percentages and corresponding anti-site defect concentration. Figure 5.72 shows a MATLAB visualization of a stoichiometric NiAl structure built using this code, as well as a Ni-poor NiAl structure with a significant percentage of Ni vacancies. MATLAB is used to generate a text file that represents these structures that can be read by LAMMPS and used to build the test structure for the TGS simulations. This code, along with detailed comments, can be found in the Appendix on page 244.




Figure 5.71: Results of convergence tests showed that the test structure should have at least 2 million atoms.

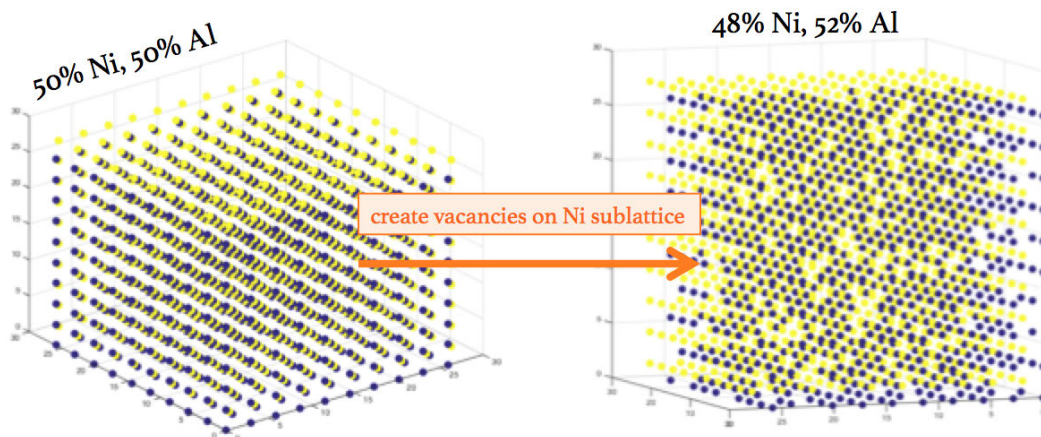


Figure 5.72: A stoichiometric NiAl structure constructed in MATLAB, and a Ni-poor NiAl structure with a significant concentration of constitutional vacancies constructed using the same code, are pictured. MATLAB is used to create input text files that rebuild these structures in LAMMPS.

5.4.2 LAMMPS lattice parameter tests

The test structures were run for at% Ni compositions in increments of 1% from 45 to 59 at%.¹¹ Near the stoichiometric point, the resolution was increased to fully capture the behavior. These results are shown in Figure 5.73. “Ideal behavior” corresponds to the structures for which at% Ni < 50% compositions are fully accommodated by constitutional vacancies on the Ni sublattice, and at% Ni > 50% compositions are fully accommodated by Ni antisite defects on the Al sublattice. f is a parameter that was used in the test structure generation code to model non-ideal behavior, e.g. Ni constitutional vacancies on the Ni sublattice even in the Al-rich compositions, or Al antisite defects on the Ni sublattice. Data shown in Figure 5.6 indicated that these non-ideal structures could be a possibility. f is defined as the ratio of ASD to vacancies on a test lattice. In the code used to generate the structures, f and overall composition are defined, and it is the number of Ni constitutional vacancies (in the Ni-poor regime) or the number of Ni ASD (in the Ni-rich regime) that deviate from their “ideal” behavior to preserve the composition. This was done primarily to investigate how more complex defect populations might affect the lattice parameter behavior, and see if they yielded a better match for the experimental lattice parameter than the ideal test structures did.

For at% Ni 44-40%, values of $f=0.01, 0.05, 0.1, 0.15, 0.2, 0.25,$ and 0.3 were considered. For at% Ni 51-50%, values of $f=10, 20, 100, 500, 1000, 5000, 10000,$ and 50000 were considered. Each test structure had a total of 2 million lattice sites (1 million per sublattice). Three cases were run for each composition. The first case was a repeatable case. This was accomplished by giving the random number generators used to assign spots to sublattice defects the same seed. The other two cases were not seeded. Near the stoichiometric point, the test matrix resolution was increased to capture any unusual behavior. A total of 373 test structures were built and relaxed for this step of the research.

The literature data points plotted in Figure 5.73 are presented as one data set because they were in good agreement with each other. (Figure 5.5 shows the individual data sets used to plot the “Literature” data set here.) The lattice parameters measured in the relaxed structures for each value of f are also plotted here. Note that each data point represents the averaged lattice parameter from the three test cases at that composition and f value.¹²

The simulated results show the same trend as the literature data, but they are offset from the literature data set by about 0.04\AA . This is not true for the stoichiometric composition, which has extremely good agreement with the literature data. This jump at the stoichiometric composition was the reason for the increased resolution of the test matrix near at% Ni = 50%: it was necessary to determine if this jump was a sudden discontinuity, or if it was a gradual climb from the 49% and 51% points. The resolution was increased as low as 0.001%, but the test structures for the ideal behavior at at% Ni = 49.999% and 50.001% were still continuous with the rest of the simulated data and not the stoichiometric point data. It is possible

¹¹Timesteps were 0.001s, with 10,000 total timesteps; boundary conditions were periodic; conditions were isothermal-isobaric (NPT ensemble). The full input file used here will be added to the MIT MNM GitHub repository and tagged with my name and “B2-NiAl” for the interested reader.

¹²The lattice parameter for a given f and composition was found to have very little variation between each of the three test cases, and so future tests may be able to save time by only testing one structure instead of three to obtain each data point.




Figure 5.73: Lattice parameter results for B2-NiAl obtained in LAMMPS compared against lattice parameters reported in the literature.

that result is due to the potential being optimized for the stoichiometric composition.

5.4.3 Results and analysis of TGS simulations on B2-NiAl in LAMMPS

The ideal NiAl structures were used for the TGS simulations, with only Ni constitutional vacancies for at% Ni < 50, and only Ni ASD on the Al sublattice for at% Ni > 50, as these structures yielded lattice parameters that were closest to the literature data. All structures yielded lattice parameters that followed, generally, the same trends as the experimental data, so there was no compelling reason at this stage to deviate from the ideal behavior. The results of the TGS simulations described on page 176 are shown in Figure 5.74.




Figure 5.74: SAW speed measured from TGS simulations carried out on NiAl test structures of 2×10^6 lattice spots in LAMMPS.

These results show an increase in SAW speed as vacancy concentration decreases, and a decrease after the stoichiometric point. Interestingly, the speed change in the Ni-rich regime appears more precipitous than it does in the Ni-poor regime, indicating that either the increased nickel content, or the nickel ASD on the Al sublattice - or a combination of both - have a stronger impact on the SAW speed than does the concentration of vacancies on the Ni sublattice.

The results in Figure 5.74 are overlaid with the experimentally measured SAW speeds in Figure 5.75. What is immediately obvious is that (1) the experimentally measured SAW speed for the polycrystalline NiAl decreases in the Ni-poor region with decreasing vacancy content, while the simulated SAW speeds increase and (2) the change in the experimentally measured SAW speed for the polycrystalline NiAl was measured over a much larger range of values than were the simulated SAW speeds. Ignoring the two left-most values, which are presumed to be from a different phase of NiAl (per the analysis in Section 5.3.3.2 and the NiAl Batch II metallographs), the experimentally measured SAW speeds for the polycrystal samples varied by about 600 m/s between 44 and 50 at% Ni. The simulated SAW speeds in the same compositional range vary by only about 100 m/s.




Figure 5.75: SAW speed measured from TGS simulations carried out on NiAl test structures of 2×10^6 lattice spots in LAMMPS.

These MD results also call into question the validity of the trends described in Section 5.3.2.2. However, as detailed in [175], it is challenging to compare calculations based on NiAl single crystals with experimental results based on NiAl polycrystals, due to the anisotropy of NiAl. The anisotropy associated with NiAl is evident in the great deal of variation in the speeds measured for the single crystal NiAl samples, even though they are clustered more tightly with respect to nickel concentration. These samples were cut carefully from the single crystal rods, and yet XRD analysis revealed miscut to be a common feature of the samples. Furthermore, these data points are the average of five spots. It would be useful to both correct these samples for miscut using a different method (as detailed previously, the previous procedure for measuring and correcting miscut was time consuming and frequently unsuccessful), or recut them to ensure that the alignment matched that of the simulated structures. Then, it would be advisable to take many more spots at a finer resolution of surface rotation changes in order to fully capture the SAW behavior.

In order to get a better picture of the single crystal sample behavior, results from older SAW speed measurements made on each of the single crystal sample were averaged and plotted with the rest of the data in Figure 5.76. These measurements were carried out over 180° and so give a more complete picture of the single crystal behavior. Some of them overlap almost perfectly with the averaged measurements from the more recent experimental campaign. Overall, by adding in these measurements, which capture the full range of the single crystal's anisotropic behavior, we find that the data points appear to lie closer to the trend line of the polycrystalline samples, which makes sense because polycrystals effectively yield averaged behavior, assuming the sample is free of texture and there are a sufficient number of grains that lie within the probe spot.

Another thing that would make this analysis more complete would be to have more polycrystalline samples with Ni-rich compositions. Since vacancies were the primary focus of this study, an emphasis was placed on creating as many samples in the Ni-poor compositional range as possible. However, it would be useful to see how the polycrystal experimental results compare with the LAMMPS results in the Ni-rich region, since overall we see more agreement between experimental and simulated results past the stoichiometric point. If the experimental data is consistent with the simulated results in the Ni-rich region, but not the vacancy-rich, Ni-poor region, that is an interesting result in and of itself. (It seems like that is the case already, but without more Batch II points past 50 at% Ni it is difficult to confirm for sure.) That result would possibly indicate that the potential used in the LAMMPS simulations ([171]) worked well to replicate physical B2-NiAl behavior for high concentrations of anti-site defects in the NiAl structure, but less so for high concentrations of vacancies.

The LAMMPS results are based on structures that are effectively single crystal and defect-free (save for the point defects imposed on each sublattice). It is possible that the grain boundaries in the polycrystal samples are impacting the signal result. It is also very possible that defects from the fabrication process are affecting the Batch II samples. The metallography results show that many of the samples had significant concentrations of large voids and other imperfections on the surface. While measurement spots were not taken *on* the voids, it is possible that voids below the surface are affecting the experimental data. Sectioning




Figure 5.76: Measurements on the Batch I samples taken over the full surface rotation (with at least two spots per rotation position) during a previous data collection campaign are included with the data sets to provide a better picture of the single crystal behavior. Error values for these points were on the order of 1% of the averaged measurement value.

the arc-melted buttons revealed that the voids tended to be distributed throughout the entirety of the sample, so this is possible.

5.5 Conclusions of the B2-NiAl study

This study demonstrated that vacancy concentrations must be accounted for when analyzing elastic properties with TGS (at least, when using our current setup configuration) if very high vacancy concentrations are expected, as we demonstrated that the TGS signal will be sensitive to vacancy concentration changes of 0.18% or greater. Since this is already on the order of the maximum vacancy concentration expected in very high dose applications, this may not be relevant for most TGS testing scenarios. Additional refinement of the setup and more data collection may lead to better sensitivities in the future, in which case the vacancy contribution to the TGS signal will need to be accounted in more typical radiation situations as well. It is also possible that, once *in situ* TGS is more established, we might find that TGS is affected by vacancy point defects even at low C_v (assuming we have a way to also track C_v in real time). At the moment, though, we can take it as a tentative finding that the contributions of C_v to E may be safely ignored for most radiation testing scenarios.

This study also demonstrated that it is possible to replicate the experimentally-observed lattice parameter behavior of B2-NiAl in LAMMPS by altering the concentration of vacancies on the test structures's Ni sublattice or the concentration of Ni anti-site defects on the test structure's Al sublattice. This provides further confirmation for the presence of constitutional vacancies in B2-NiAl at high concentrations. This work also provided the development of a MATLAB tool that can be used by other researchers interested in B2-NiAl to efficiently build many large test structures of B2-NiAl with any desired composition, defect population, and dimensions that can be easily imported into LAMMPS.

LAMMPS TGS simulations replicated the expected behavior with regards to v_{SAW} , but the experimental measurements of v_{SAW} on NiAl samples fabricated in LAMMPS did not. This underscored what has been reported in prior literature, which is that it is difficult to extrapolate single crystal NiAl behavior to polycrystalline NiAl behavior. However, after fabricating several groups of NiAl samples (both single and polycrystal) and measuring them in multiple TGS campaigns, we are well poised to modify the experiments in the future if desired to determine the cause of the discrepancy more conclusively, or possibly close it altogether. With this goal in mind, future work on this study would include:

- Experimental validation of vacancy concentration in fabricated samples should be carried out. (This will also contribute to the existing literature datasets for experimentally measured C_v in B2-NiAl.)
- Rebuild the test structures so that their compositions and C_v precisely matches the experimental samples, and rerun the LAMMPS TGS simulation.
- Find a way to fabricate the polycrystalline samples such that large voids (artifacts of the arc-melting and cooling process, not of radiation damage) are avoided .

- Take more than ten measurement spots.
- Instead of trying to grow the grains via annealing, aim to make polycrystalline NiAl with grains that are as small as possible, because of the known high anisotropy of the B2-NiAl crystal. This will allow a better averaging of the behavior with each spot.
- Consider repeating the study with a B2 intermetallic like FeAl that is associated with lower C_v , since this is more comparable to the C_v expected in most irradiated materials.

It would also be of interest to analyze the experimental data for acoustic damping results once the acoustic damping analysis techniques are finalized, since this would further elucidate the properties of the B2-NiAl and contribute new information to the literature on the same. While B2-NiAl might not have been the best proxy for studying C_v in irradiated materials after all, it is still an engineering material of great interest to fields outside of nuclear reactor design, and so novel characterizations are broadly relevant to materials science as a whole.

Chapter 6

TGS examination of stainless steel samples irradiated in EBR-II

In this chapter, we aim to move from model materials to a real engineering alloy irradiated in a reactor core, and see if TGS is still a useful method of radiation damage measurement even in this more complex, less “controlled” situation. This section of the thesis considers TGS research carried out on sample material from hex blocks that were irradiated in the core of Experimental Breeder Reactor-II (EBR-II) at Idaho National Laboratory, under the operation of Argonne National Laboratory. EBR-II was a sodium-cooled fast reactor that was in operation from 1965 to 1994. The core was designed to accommodate up to 65 subassemblies dedicated to experimentation at any one time, allowing for many important research studies to take place during its three decades of operation [187]. During that time, many material tests on stainless steel were carried out in the EBR-II core that proved to be major contributions to the field’s understanding of radiation-induced creep, swelling, and ductility changes [188].

This chapter uses samples obtained from AISI 304 steel hex blocks that were irradiated in the EBR-II in order to study the long-term effects of a fast reactor neutron flux on stainless steel. The hex blocks, which were initially 5 cm wide and ranging in length from 21.8-24.5 cm, were stacked in a duct in the EBR-II reflector [189]. Figure 6.1 is a photograph of one of the hex blocks in a hot cell after its removal from the reflector. The blocks are made of AISI 304 stainless steel of composition (by wt%): 12.26Cr, 8.81Ni, 1.57Mn, 0.43Si, 0.056C, 0.027P, 0.03S, with Fe as the balance.

6.1 EBR-II hex block section samples: original locations

The samples that were cut from the hex blocks have highly characterized initial locations in the hex blocks. This information is important because the radiation response of a given block is not uniform throughout.




Figure 6.1: A radioactive hex block in a hot cell after removal from the EBR-II reflector. The block is 5.2 cm wide and about 20cm long. [189]




Figure 6.2: Radiation damage response is not uniform throughout each block due to neutron flux and temperature gradients that existed throughout the block. The temperature peaked near the middle, but the peak was shifted toward the side that was nearest to the reactor core. Gamma heating and dpa rate increased monotonically from the side furthest from the reactor core to the side closest. Schematic from [189]

During their exposure, there were radial and axial gradients in neutron flux, temperature, and internal heating [189], as illustrated schematically in Figure 6.2. Therefore, it is useful to know where in the stack of hex blocks each sample was located when it was inside the EBR-II reflector.

Figure 6.3 shows where in the core the samples tested here were located. The hex blocks are labeled with numbers 1-5. Samples from Blocks 3 and 5 were used in this work. The blocks were cut into hexagonal coins, which are given a number identification. Coins 3D, 3E, 5A, 5B, and 5C were used in this work.

The test samples used for the TGS experiments have a six-digit alphanumeric identification number. The identification number provides the following information:

3D1C2A

- Hex block ID
- hex coin ID
- hex coin section ID (1 indicates the middle section)
- Subsection of the coin section
- Sample cut from the subsection
- Letter identifier for thin slice samples

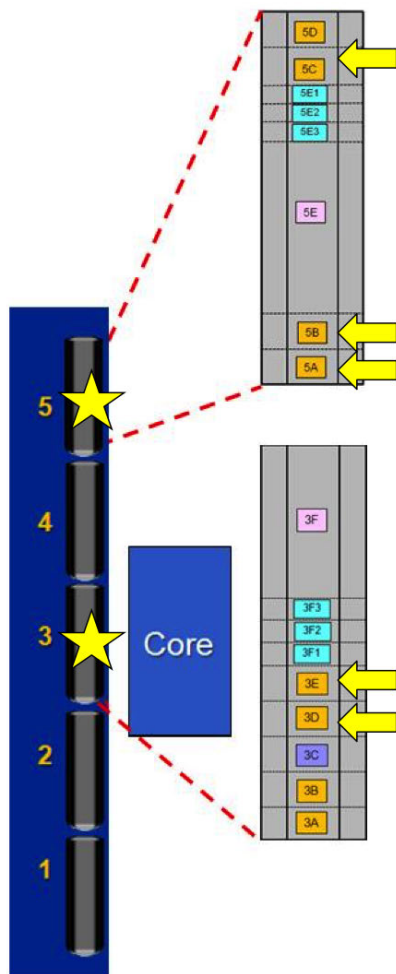


Figure 6.3: The image, modified from [190], shows where the tested samples came from in the EBR core. The blocks (5 and 3) are marked with stars; the sections of those blocks relevant to this study are marked with yellow arrows. The blocks were stacked in a hexagonal duct with 1 mm thick 304 stainless steel walls. The duct was located in Row 8 of the core's reflector region [189].




Figure 6.4: The image, from [190], shows where Sample 3D1 was located in Coin 3D, as well as its dimensions. **D** signifies that the sample was immersion density measured.

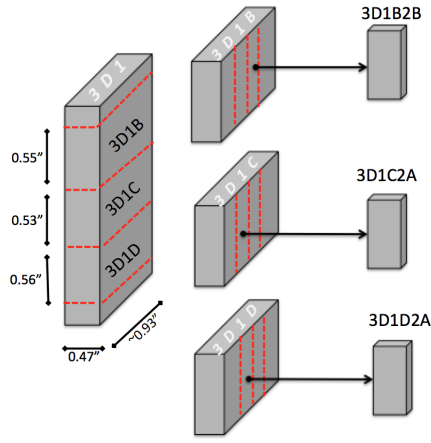


Figure 6.5: Sample 3D1 was divided into five sections. The larger sections in the middle three (3D1B, 3D1C, and 3D1D) were each further divided into four subsections. One subsection from each of 3D1B, 3D1C, and 3D1D were used in this study.




Figure 6.6: The image, from [190], shows where Sample 3E1 was located in Coin 3E, as well as its dimensions. **D** signifies that the sample was immersion density measured.




Figure 6.7: Sample 3E1 was divided into five sections. Section 3E1B was divided into three sections. The middle section, 3E1B1, was divided into four more sections, two of which were thin slices intended for use in making TEM samples. One of these thin slices, 3E1B1A, is used here. [190]




Figure 6.8: The image, from [190], shows where Sample 5A1 was located in hex coin 5A. Sample 5A1 was divided into five sections, one of which, 5A1B, is relevant to this work.

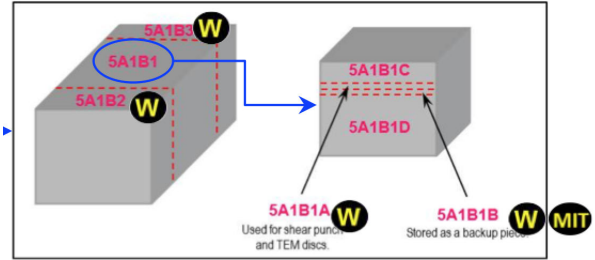


Figure 6.9: Section 5A1B was divided into three more sections. The middle piece, 5A1B1, was divided into four more sections, two of which were thin slices intended for use in making TEM samples. One of these thin slices, 5A1B1B, is used in this work. Both this figure and Figure 6.8 are from [190].




Figure 6.10: 5B1 of hex coin 5B was divided into four sections. The two middle sections, 5B1B and 5B1C, were each divided into four more sections. As in previous coins relevant to this work, two of these sections were thin slices. One thin slice from each of 5B1B and 5B1C were used in this work: 5B1B2B and 5B1C2B. [190].




Figure 6.11: 5C1 of hex coin 5C was divided into three sections. The middle section, 5C1B, was divided into four more sections. As in previous coins relevant to this work, two of these sections were thin slices. One of the thin slices, 5C1B2B, was used in this work. [190].

The control sample used in this study comes from an archival block that was identical to the pre-irradiation hex blocks but which was not exposed to the EBR-II core environment.

Section 3D1 comes from coin 3D, which was located near the middle of hex block 3, as shown in Figure 6.3. Figure 6.4 shows that Sample 3D1 is cut from the middle of coin 3D. 3D1 was subdivided into five sections. The middle three subsections were further divided into four sections as shown in Figure 6.5. One of the four sections from each of 3D1B, 3D1C, and 3D1D were used in this work. Section 3E1 came from coin 3E, which was located just above coin 3D in hex block 3, as shown in Figure 6.3. Figure 6.6 shows that sample 3E1 is cut from the middle of coin 3E. 3E1 was subdivided into five sections, one of which, 3E1B, was used in this work. 3E1B was divided into three sections. The middle section of 3E1B, 3E1B1, was divided into four sections, two of which were thin slices intended for use in making TEM samples. This is shown in Figure 6.7. One of these thin slices, 3E1B1A, was used in this work.

The Hex Block 5 samples used in this work came from three hex coins: coins 5A, 5B, and 5C. As shown in Figure 6.3, 5A and 5B were located at the bottom of hex block 5; 5C was near the top. Figures 6.8 and 6.9 show the relative original location of sample 5A1B1B in hex coin 5A. Two samples, 5B1B2B and 5B1C2B, are from hex coin 5B, a diagram of which is shown in Figure 6.10. One sample, 5C1B2B, is from hex coin 5C, a diagram of which is shown in Figure 6.11.

6.2 Radiation damage in the EBR-II samples

These samples have been previously studied and characterized, and their irradiation history is well known. Extensive research has been carried out on samples from the hex blocks, as there is significant interest in understanding how stainless steel responds to multi-year in-core exposure.

6.2.1 Radiation and temperature history

The hex blocks were stacked within a hexagonal duct made of 1 mm thick 304 stainless steel, which was placed in the reflector region of the EBR-II core. They were in the core for a total of thirteen years: for the first 4.5 years, the hex block stack was kept in in Row 8 of the reflector, and for 8.5 years the blocks were placed in Row 16, which was further back. Nearly all of the radiation damage the blocks sustained is expected to have been sustained while they were stacked in Row 8. As the hex blocks were stacked vertically in the core, samples from hex block 3 were exposed to the highest damage and gamma heating [189]. All of the blocks are still radioactive, as shown in Table 6.1, although they have cooled considerably since their removal from the EBR-II core. Table 6.2 gives the reported total dpa for select hex coins, as well as the operating temperatures to which they were exposed during irradiation, as reported in [191].

Table 6.1: EBR-II sample history and description, as provided by Westinghouse

Sample ID	Sample state upon receipt by MIT MNM	Width [in]	Length [in]	Thickness [in]	Volume [in ³]	Dose Rate* [mR/h]
3D1B2B	Used for Phase 2 shear punch specimens (plate has 4 punched out holes in it and there is also 1 punched disc)	0.4	0.5	0.014	0.003	7
3D1C2A	Used for Phase 2 shear punch specimens (plate has 3 holes punched out of it)	0.4	0.5	0.011	0.002	8
3D1D2A	Used for Phase 2 shear punch specimens (plate has 3 holes punched out of it)	0.4	0.5	0.010	0.002	6
3E1B1A	Intact	0.5	0.5	0.008	0.002	4
5A1B1B	Intact	0.5	0.5	0.017	0.004	15
5B1B2B	Used for Phase 2 shear punch specimens (plate has 3 holes punched out of it)	0.4	0.5	0.010	0.002	6
5B1C2B	Used for Phase 2 shear punch specimens (plate has 3 holes punched out of it)	0.4	0.5	0.010	0.002	6
5C1B2B	Used for Phase 2 shear punch specimens (plate has 3 holes punched out of it)	0.4	0.5	0.010	0.002	6

* This measurement was made in June 2014 at a distance of one foot from the sample.

Table 6.2: Dose and temperature range for select hex coins, as reported in [191]

Coin	dpa	T [°C]
3D	28	418-448
3E	28	420-450
5A	4	417-422
5B	3	416-420
5C	0.5	415




Figure 6.12: Micrographs of the archival stainless steel hex block, which serves as the stock for the control sample used in the TGS project. The archival block indicates that the blocks were cold worked prior to placement in EBR-II. [189]

6.2.2 The EBR-II study control sample

An archival block identical to Hex Blocks 1-5 but which was not placed in the reactor was used for comparison purposes. Figure 6.12 shows microscopy results of the archival hex block. The block displays evidence of cold work, with the microstructure indicating a level of $\approx 5\%$ cold work, and indicates that the blocks were not annealed prior to placement in the reactor [189]. A sample from this block is used as the control sample in the TGS measurements. It can be assumed that this block's microstructure is representative of the microstructure of hex blocks 3 and 5 prior to irradiation.

6.2.3 Hex block swelling and density changes due to radiation damage

When the blocks were removed from EBR-II, extensive characterization of the swelling they had undergone was performed. Select results are presented here to illustrate the extent of radiation damage the hex blocks sustained during the thirteen years in the core. Figure 6.13 shows selected profilometry, including the flat-to-flat swelling for Blocks 2-4 and lengthwise swelling for Blocks 3 and 5. Block 3, which was in the center of the stack, had the highest swelling -between 0.5% and 1.4%. Profilometry measurements were carried out using calipers and micrometers in hot cells at Idaho National Laboratory in [189].

Samples from Block 3 used in this work are 3D1B2B, 3D1C2A, 3D1D2A, and 3E1B1A. 3E1B1A is from hex coin 3E, which was very close to the center of Block 3. The first three come from hex coin 3D, which was just below 3E (nearer to Block 2 in the stack). In the plotted data in Figure 6.13, this corresponds

to an axial position of between 10 and 15 m, so we would expect between 0.5% and 1% swelling in the 3D samples in the flat-to-flat direction. The 3E sample is closer to the 15m axial position, and we would expect it to exhibit a swelling of at least 1% in the flat-to-flat direction.

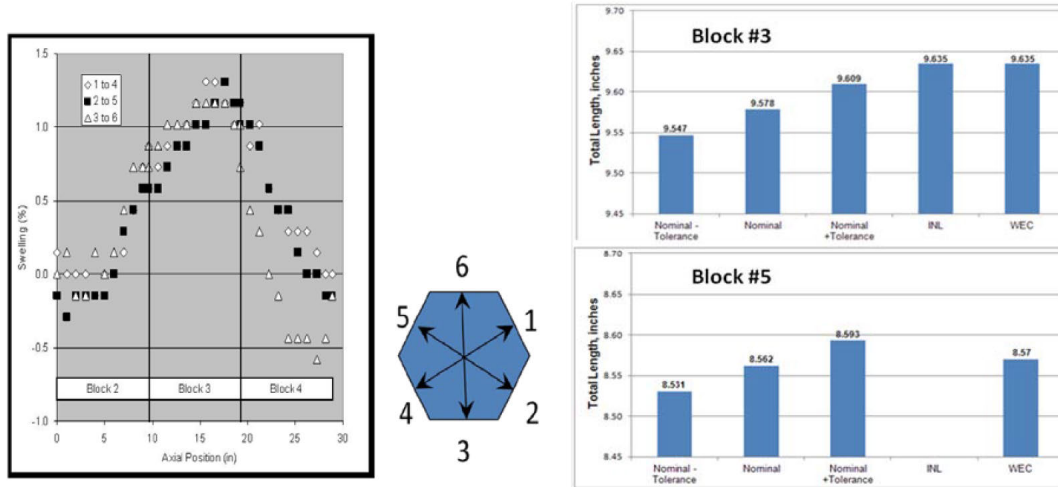


Figure 6.13: Flat-to-flat swelling for Blocks 2, 3, and 4; a schematic for visualizing the direction of the flat-to-flat measurements; and lengthwise swelling in Blocks 3 and 5. Block 3, in the center of the stack, exhibited the most dramatic swelling. Block 5 was at the top of the stack. Lengthwise measurements show that Block 3 exhibited significant expansion, whereas Block 5 did not. INL and WEC indicate the hot cell in which the measurements were made - Idaho National Laboratory or Westinghouse Electric Company. Negative swelling values in Block 2 and Block 4 indicate carbon densification has taken place. Flat-to-flat swelling data for Block 5 is presented in Figure 6.14. [189]

The authors of [189] find that swelling evidence in Block 5 is inconclusive, since the lengthwise measurements post-exposure are still within the block’s tolerance bounds.¹ Figure 6.14 shows the flat-to-flat swelling measurements taken along the axial length of Block 3 and Block 5. Block 3 data - as already shown in Figure 6.13 - shows an increase in the flat-to-flat distance from the bottom of the block to the top. Block 5 data, however, shows flat-to-flat distances that decrease as one moves from the bottom of Block 5 (closer to Block 3) to the top of Block 5 (further from Block 3). This indicates that carbon densification processes dominate over any swelling processes in Block 5 [189].

Profilometry results showed significant swelling in Block 3, which was in the center of the hex block stack and which received the highest damage dose. Block 5 exhibited “shrinkage” in some cases, indicating that carbon densification dominated over void production. Swelling in Block 3 was observed in both the flat-to-flat direction and in the axial direction. On a general level, this means that the samples from Block 3 should exhibit properties associated with void swelling and samples from Block 5 should exhibit properties

¹Later in this work, however, it is stated that voids were observed, indicating that swelling occurred to some degree, but was dominated by the densification process.




Figure 6.14: Post-irradiation flat-to-flat dimensions for Block 3 and Block 5. The solid line at 2.062” shows the nominal dimension of the block prior to irradiation. Block 3, which received the higher dose, clearly exhibits swelling behavior that clearly exceeds the nominal+tolerance flat-to-flat length. The amount of swelling from flat-to-flat varies with axial position. Block 5 exhibits negative swelling, which is associated with carbon densification. Note that left-to-right indicates bottom-to-top. [189]

associated with densification and with relatively less (or no) void swelling as compared to Block 3.

Other characterization tests involved the use of ultrasonic mapping. Ultrasonic mapping is a nondestructive technique that can be used to probe sample materials that has historically been used to locate cracks and other defects. By comparing ultrasonic maps to density changes, ultrasonic results can be validated, particularly as a method of studying more generalized microstructural changes such as radiation-induced swelling [192], making it a particularly interesting complement to a TGS study.

Results from the ultrasonic mapping of Hex Block 3 are shown in Figure 6.15. 3D and 3D are the hex coins of interest to this study; 3E was not mapped despite the fact that it received the highest dose in the block. This shows that the acoustic speed of coin 3D increases moving from Flat 5/6 to Flat 3/2. The variation of acoustic speeds measured in 3D is about 20 m/s. Several of the coins relevant to this work were also acoustically mapped. In Figure 6.16, the cross-sectional area of hex coin 3D is mapped.




Figure 6.15: Acoustic velocity mapping (lengthwise) for Block 3. Acoustic mapping for coin 3E from this angle and of Block 5 was not provided. Section 3D (the side of hex coin 3D) exhibits a variation of acoustic speed of about 20 m/s. The archival (unirradiated) acoustic velocity is 5735 m/s. [189]

The ultrasonic mapping results show that velocity is lowest where voids (and damage) were highest in these samples. Immersion density measurements, which were made on individual sections cut from the coins, were consistent with the acoustic mapping results: the greatest changes in density were observed in the same sections that had the greatest reductions in acoustic speed relative to the nominal archival acoustic speed of 5735 m/s. Coins from Block 3 had the most significant variation in acoustic speed, with a range of 70 m/s. The speeds reported here did not exceed 5720 m/s. Coins from Block 5 had a more uniform acoustic map, with a range of about 35 m/s. However, the lowest reported speed observed in Block 5 coins was about 5735 m/s. This makes sense, because Block 3 received the most damage, and experienced high rates of void swelling (which corresponds to lower acoustic speeds). Block 5 damage was dominated by densification, which corresponds to higher acoustic speeds.




Figure 6.16: An acoustic map of hex coin 3D shows that the lowest speeds are clustered slightly off-center - the same spot where the highest concentration of voids is expected to be and where radiation damage was highest. The coins from block 3 exhibit a large acoustic velocity gradient, with variations of up to 70 m/s. [189]




Figure 6.17: An acoustic maps of hex coins shows velocity increases relative to the nominal acoustic speed of the archival material (5735 m/s). The changes are less severe than those observed in block 3, with acoustic speed variations of about 30 m/s. [189]

6.2.4 TEM microscopy of the EBR-II samples and quantification of defect populations

TEM imaging revealed the presence of voids in coin 3E and the presence of dislocations and $M_{23}X_6$ precipitates in Coin 3D. This were carried out on Samples 3D1B2A, 3D1C2B, and 3D1D2B.

Carbides and voids were observed in Block 5 material [189]. The post-exposure microstructure is revealed to be completely different from that of the archival material: instead of dislocation loops, there are Frank-Reed loops, voids, and precipitates. TEM analysis was performed on Samples 3D1B2A, 3D1C2B, and 3D1D2B, which are the “sister” samples of Samples 3D1B2B, 3D1C2A, and 3D1D2A, which are used in this project [189].² These samples are thin slices that were immediately adjacent to each other, so it is reasonable to assume that the phenomena observed in 3D1(B2A, C2B, D2B) - e.g. precipitates - would also be present in 3D1(B2B, C2A, D2A). The carbides observed in Block 5 are consistent with the carbon densification process. The fact that voids are still observed indicates that carbon densification is the dominant effect - in terms of impact on swelling - but it is not the significant damage phenomenon present.

Figure 6.19 shows SEM images of coin 3E that reveal a significant density of voids. Figure 6.20 shows an TEM image of coin 5D. Hex block 5 received a significantly lower dose than hex block 3 due to their

²See Figures 6.4-6.11. The samples used in this project are from one of the two thin slices cut from the a larger section from the middle of the hex coins. The samples used in the TEM studies that are described here as “sister” samples are the other thin slice from the same piece of the same hex coin.




Figure 6.18: An acoustic map of hex coin 3E, overlaid with density results. Note that the image is flipped relative to the image in Figure 6.16, with the flats on opposite sides. The low-velocity region occurs in the same place on each coin. Coin 3E was sectioned according to the dashed lines, and density measurements were made on each segment. The extent of the density change corresponded with the extent of the acoustic velocity reduction. [189]





Figure 6.19: TEM images of a sample from coin 3E show many voids dispersed throughout the material. Hex block 3 was in the center of the stack and received the highest radiation dose. [191]



(f) Coin 5D at 25 mm (416°C).

Figure 6.20: Hex block 5 received a much lower dose than hex block 3. A few voids are visible in this coin 5D sample, but the difference in void density visible in this image and in the images of coin 3E is obvious. [191]

relative positions in the channel, and as a result, the difference in microstructure of samples from each hex block is stark. A few small voids are visible in the sample from coin 5D, but the number density of voids in coin 5D is obviously much lower than it is in coin 3E.

Figure 6.21 shows a table from [193] that quantitatively characterizes the defect population observed in various EBR-II samples via TEM analysis. Stars show which samples correspond to a sample (i.e., the used in this study).

Specimen Identity	dpa	Temp. (iâC)	Voids			Precipitates			Frank Loops		
			Diameter (nm)	Density ($10^{21} / m^3$)	Swelling (%)	Diameter (nm)	Number Density ($10^{21} / m^3$)	Volume (%)	Diameter (nm)	Number Density ($10^{21} / m^3$)	Length Density ($10^{14} / m^2$)
★ 3D1B2A	28	418	18.0	2.25	1.18	14	0.71	0.13	25.4	7.68	1.96
★ 3D1C2B	28	448	28.3	1.67	2.94	15	0.91	0.5	37.0	4.19	1.55
★ 3D1D2B	28	418	23.1	3.19	2.08	16	1.10	0.73	31.5	10.4	3.32
★ 3E1B1B	28	420	20.2	2.24	1.84	14	0.71	0.46	26.1	9.14	2.18
3E1C1A	28	448	22.5	2.61	2.76	15	0.91	0.44	28.5	8.83	2.32
3E1D1B	28	420	20.1	2.85	1.95	16	1.10	0.52	22.8	13	2.62
★ 5A1B1A	4	417	19.2	0.80	0.52	10	1.20	0.34	41.1	9.62	3.54
5A1C1A	4	422	18.1	0.80	0.47	12	0.64	0.31	44.5	10.6	4.13
★ 5B1B2A	3	416	19.4	0.46	0.31	10	1.20	0.32	45.1	7.24	3.32
★ 5B1C2A	3	420	21.1	0.59	0.47	12	0.64	0.39	44.3	8.34	3.65
★ 5C1B2A	0.6	416	18.3	0.42	0.29	13	0.63	0.4	53.0	2.78	1.52
5D1C1A	0.4	416	16.8	0.30	0.23	13	0.63	0.48	77.6	2.33	1.66

Figure 6.21: The table from [193] that quantitatively characterized voids, precipitates, and Frank-Reed loops, based on TEM of the EBR-II samples. Stars are added here to show which rows correspond to a sample that was immediately adjacent to the samples used in this study prior to the samples being cut from the hex blocks. It is assumed that the defect populations of the samples in this study are effectively equivalent to the defect populations of their corresponding samples in [193].

6.3 EBR-II TGS Experiments

The EBR-II samples were measured using the same conditions as the niobium and B2-NiAl samples: $5.5\mu\text{m}$ grating spacing, 10 spots per sample, and three batches of 10,000 traces per batch. The probe laser current was kept at 400.01 mA for each measurement, with a probe temperature of 18.71°C .³ The samples were radioactive, and so while no setup modifications were made, extra care was taken in handling them (e.g. ensuring that no one entered the TGS setup area while the measurements were taking place).

6.3.1 Results

TGS measurements were performed on each EBR-II irradiated sample provided to MIT MNM and on a control sample from the archival block. Ten spots were taken on each sample, with each spot measurement consisting of three batches of 10,000 traces each, such that the TGS data file for each measurement consisted of three positive signal traces and three corresponding negative signal traces. The negative trace is subtracted from the positive trace to create the trace used to extract the parameter measurements. The MATLAB codes provided in the appendix beginning on page 226 were used to isolate the dominant frequency of each signal (which is converted to SAW speed by multiplying it by the calibrated grating spacing for the day on which the measurement was taken) and the thermal diffusivity (via analysis of the thermal decay curve).

Figure 6.22 shows the results of the SAW speed measurements, plotted against dpa, for each individual sample. The plot is difficult to read due to the overlapping of error bars, especially for samples with identical reported total dpa. In Figure 6.23, the same data sets are used, but the results from all samples from a given hex block are grouped together. The dpa of the samples from hex block 5 are averaged, resulting in a value of 2.625 dpa. Every sample provided to MIT MNM for analysis from hex block 3 was irradiated to a reported 28 dpa. By presenting the data this way, the general trend is more obvious, with a SAW speed that appears to increase beginning at low dose and quickly saturate.⁴ Figure 6.24 shows the results of the thermal diffusivity measurements, plotted against dpa, for each individual sample. Table 6.3 provides the measurements plotted in Figures 6.22 and 6.24, listed by dpa and sample name.

The TGS results for v_{SAW} and thermal diffusivity were then compared against the quantitative defect populations reported in Figure 6.21. This was done by comparing which data point corresponded to which sample, and matching that with the defect population measurement for that sample (again, it is being assumed that for a given sample analyzed in [193], its measured defect populations will also be characteristic of the samples immediately adjacent to it). In Figures 6.25 and 6.26, TGS measurements of v_{SAW} and thermal diffusivity for the MIT MNM EBR-II samples are compared against the void density measurements made

³These are the rough values of the probe current and temperature for the other experiments, but it typically isn't necessary to be exacting. Here, because we were transitioning from model materials to complex materials, though, attempts were made to eliminate any identifiable source of possible noise or uncertainty.

⁴The error bars on the control sample data point are still large, but they do not overlap with the error bars of the other two samples. The surface of the control sample did not appear to be perfectly flat, however, leading to difficulty in obtaining a strong TGS signal and more noise in the signal traces than was typical. Further characterization of the EBR-II samples with TGS should either take more than ten control measurements, re-polish the sample surface, or obtain another sample from the archival block.

Table 6.3: TGS-measured SAW speed and thermal diffusivity for the EBR-II samples

Sample	dpa	v_{SAW} m/s	α m ² /s
CONTROL	0	2437.13	3.91E-06
5C1B2B	0.5	2530.99	5.10E-06
5B1B2B	3	2552.31	5.21E-06
5B1C2B	3	2559.30	5.26E-06
5A1B1B	4	2629.79	5.36E-06
3D1B2B	28	2497.95	5.14E-06
3D1C2A	28	2638.03	5.06E-06
3D1D2A	28	2501.64	5.30E-06
3E1B1A	28	2596.36	5.14E-06

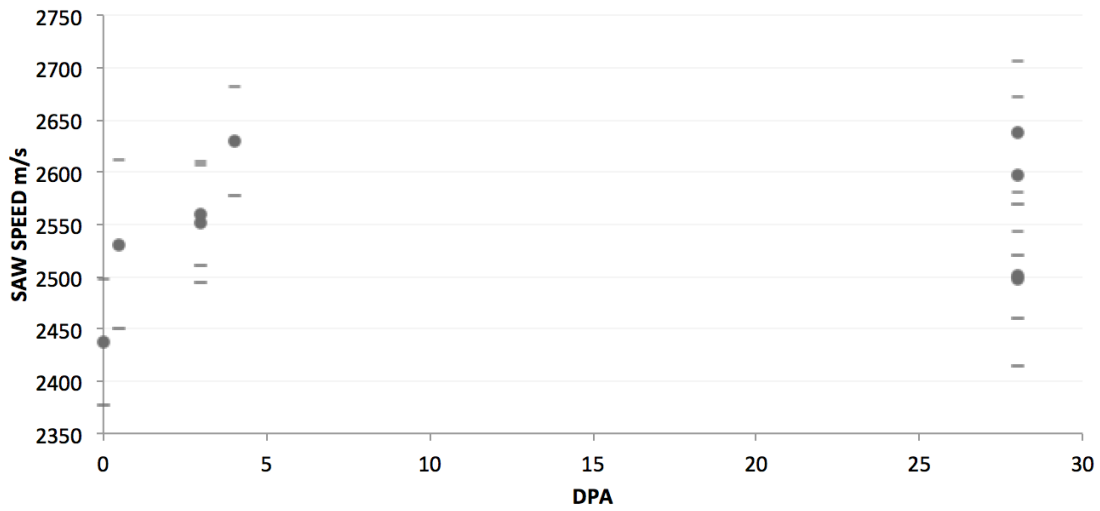


Figure 6.22: SAW speed measurements, averaged by sample and plotted by the dpa to which each sample was irradiated in EBR-II.




Figure 6.23: SAW speed measurements, with all data grouped by hex block. Hex block 5 (second point) is plotted at 2.625 dpa, which is the average of the dpa values reported for each sample measured here that originated in hex block 5. Hex block 3 (rightmost point) is plotted at 28 dpa, which was the reported dpa value for every sample measured here that originated in hex block 3.




Figure 6.24: Thermal diffusivity measurements, averaged by sample and plotted by the dpa to which each sample was irradiated in EBR-II.

on their sister samples in [193]. Figures 6.27 and 6.28 show the same for void swelling measurements; Figures 6.29 and 6.30 show the same for precipitate density measurements; Figures 6.31 and 6.32 show the same for Frank-Reed loop measurements, and Figures 6.33 and 6.34 show the same for Frank-Reed loop length density measurements.




Figure 6.25: Void density as measured in [193] versus TGS-measured SAW speed




Figure 6.26: Void density as measured in [193] versus TGS-measured thermal diffusivity




Figure 6.27: Void swelling as measured in [193] versus TGS-measured SAW speed




Figure 6.28: Void swelling as measured in [193] versus TGS-measured thermal diffusivity

6.4 Analysis of TGS measurements of EBR-II samples

Figure 6.23 shows what appears to be an increase in SAW speed that occurs at low doses, which corresponds to an overall stiffening of the material, as $v_{SAW} \propto \sqrt{E}$. Looking at this plot, with the Hex block 5 samples averaged together, we see an increase of about 120 m/s in SAW speed from 0 dpa to the averaged 2.625 dpa point. This corresponds to a 5% increase in SAW speed, which corresponds to a 25% increase in E .




Figure 6.29: Precipitate density as measured in [193] versus TGS-measured SAW speed




Figure 6.30: Precipitate density as measured in [193] versus TGS-measured thermal diffusivity




Figure 6.31: Frank loop density as measured in [193] versus TGS-measured SAW speed




Figure 6.32: Frank loop density as measured in [193] versus TGS-measured thermal diffusivity




Figure 6.33: Frank loop length density as measured in [193] versus TGS-measured SAW speed




Figure 6.34: Frank loop length density as measured in [193] versus TGS-measured thermal diffusivity

It would appear that this increase occurs at low doses, since the 0.5 dpa sample (5C1B2B) had a measured SAW speed of 2531 m/s, which corresponds to a 4% increase in SAW speed (16% increase in E). However, it is more instructive at this stage to average sample results by hex block and look for more general trends in SAW speed, as the uncertainties associated with individual sample measurements are very large. The average SAW speed of the samples from hex block 3 (the high dose hex block from the center of the stack)

This increase may be due to the observation that a material with dislocations and/or alloying elements - both of which are true of the control sample from the archival hex block - tends to undergo a stiffening at low irradiation doses. This is because the dislocations, initially pinned by alloying elements, may have their effective lengths shorted by the presence of defect clusters that are generated in the material during irradiation [194, 195]. In general, it's expected that void swelling leads to a decrease in the elastic modulus, but it's possible that the dislocation pinning effect is dominating at these doses.

Austenitic stainless steels have also been shown to exhibit an increase in yield strength at low- to mid-doses, as shown in Figure 6.35. Young's modulus E is equivalent to the slope of a material's stress-strain curve for $\sigma < \sigma_y$, with σ_y the stress at the point at which a material's elastic limit is reached. If $\sigma_y^{irrad} > \sigma_y^{control}$, and the strain at which σ_y occurs is also lower, then the slope of the elastic regime curve - i.e., E , must also be higher. Figure 6.35 shows that σ_y increases with irradiation at the doses and temperatures ($\approx 420^\circ$) that characterize the EBR-II samples measured here. Figures 6.36 and 6.37 show the experimentally measured stress-strain curves at different strain-rates for 304 stainless steel before and after irradiation over four years to 30 dpa in the High Flux Reactor in Petten [196]. The irradiated samples display a transition from elastic to plastic behavior at higher yield stresses and lower strain rates than is observed in the unirradiated samples. So, the results in Figure 6.23 are consistent with experimentally observed changes in 304 stainless steel. Note that with extended dose, σ_y is found to decrease. Unfortunately, the EBR-II samples do not exceed 28 dpa in dose, and so it is not possible to test whether E , as indicated by v_{SAW} , would have decreased with increasing dose in parallel with the reduction in σ_y indicated in Figure 6.35.

The thermal diffusivity results, with the averaged results for each sample shown in Figure 6.24, show an increase in thermal diffusivity that occurs at low dose (0.25 dpa) and appears to saturate. This may seem counterintuitive, because increased dose is associated with more radiation damage defects, and defects in turn are associated with less efficient conduction of heat by the material. However, this result can be interpreted in the context of what we know about how 304 stainless steel responds to irradiation. In steel alloys that have been engineered for radiation resistance, point defects tend to recombine and annihilate at high rates, but in steels susceptible to void swelling - like 304 - point defects recombine at slower rates and cluster into larger defect structures [198]. The result is microstructure clean up, similar to what was postulated in the analysis of the irradiated niobium TGS results in Chapter 4.

Radiation-induced segregation (RIS) of alloying elements may also have an effect on thermal diffusivity. In stainless steel, the alloying elements that are smaller than Fe (Si, Ni, P, B, S) tend to migrate preferentially toward sinks, and elements such as chromium tend to deplete [199, 200, 201, 202]. This is because point defect motion is associated with solute motion - in order for a vacancy to move, it displaces (or effectively




Fig. 6. Evolution of the yield stress change in solution annealed material. a) shows the variation of yield stress with dose and temperatures. b) shows the approximate variation of peak yield stress with irradiation temperature.

Figure 6.35: Neutron irradiated stainless steels undergo an increase in yield strength σ_y at doses that correspond to the doses incurred by the EBR-II samples, which corresponds to an increase in E . The irradiation temperatures to which the EBR-II samples were exposed was in the 400-420°C range, well below the range at which temperature would be expected to be associated with a decrease in σ_y . [197]

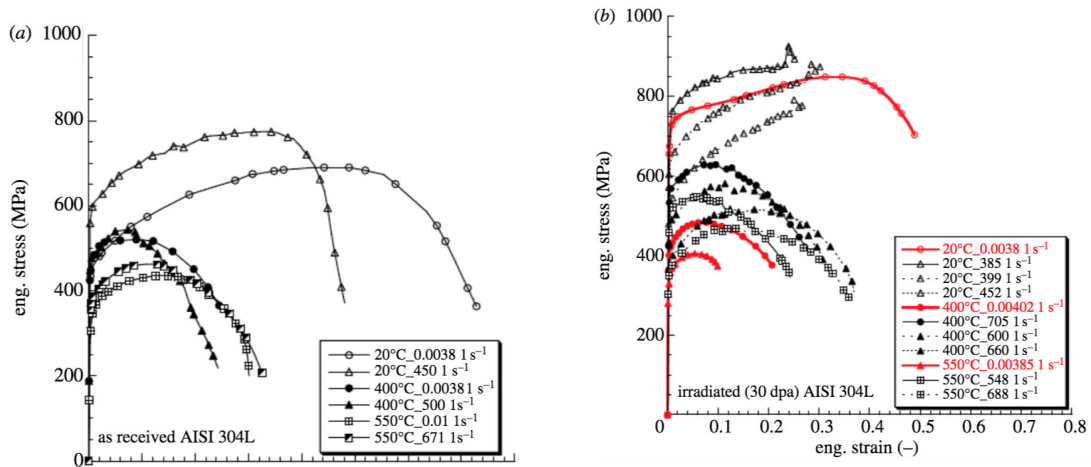


Figure 6.36: 304 stainless steel stress-strain test results, prior to irradiation. [196]

Figure 6.37: 304 stainless steel stress-strain test results, following irradiation to 30 dpa in the High Flux Reactor in Petten, the Netherlands. [196]

swaps locations) with an atom in the crystal lattice. The flux of vacancies toward sinks is balanced by a flux of atoms away from the sinks [12]. Interstitials move with the atom flux. In particular, solute atoms smaller than the matrix atoms tend to move toward grain boundaries and sinks. Figure 6.38 shows APT examination of stainless steel irradiated to 5 dpa with protons, and illustrates the strong enrichment of certain alloying elements at a grain boundary [203].

This is important because alloying elements are heterogeneities in the lattice which distort it and act as electron scatterers, which decreases thermal conductivity [204, 205]. Since thermal conductivity and thermal diffusivity are proportional, it follows that alloying elements decreases the thermal diffusivity.^{5 6} When stainless steels undergo RIS, the net effect is for a decrease in number of alloying element atoms present in the bulk of the grains, and this decrease would tend to increase the thermal diffusivity of the material.

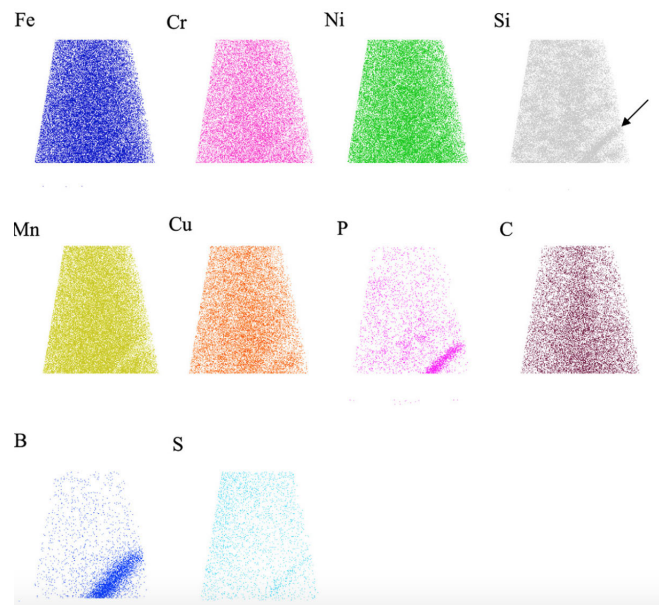


Figure 6.38: APT results for 304 steel irradiated to 5 dpa with protons shows evidence of radiation-induced segregation. A grain boundary in the APT sample is indicated for the Si scan with a black arrow. These results show slight depletion at the grain boundary for Fe, Cr, Mn, and Cu, and enrichment at the grain boundary for Si, P, B, and S. Enrichment is particularly dramatic for the Si, P, and B species. [203]

RIS can also lead to precipitation of alloying elements as they become enriched in certain areas past the solubility limit [198, 207, 208]. APT performed on another set of 304 stainless samples that were irradiated

⁵Thermal diffusivity can be calculated as $k(c_p\rho)^{-1}$, with k the thermal conductivity, c_p the specific heat, and ρ density.

⁶For example, the thermal conductivity in cast iron ferrite when alloying elements are present can be calculated as $k = k_0 - \ln \sum \chi_i$, with k_0 the thermal conductivity of pure iron and χ_i the solute concentration of the i^{th} solute species [206]. It's assumed here that a similar effect would hold for 304 stainless steels: alloying elements decrease thermal conductivity overall, and thus, decrease thermal diffusivity.

to 0.4 dpa and 28 dpa in EBR-II reveals precipitates rich in Ni, Si, and K, as shown in Figures 6.39-6.41 [209]. This precipitation was evident even at 0.4 dpa. Precipitation is further evidence of RIS, and an indicator that the concentration of alloying elements is decreased in the bulk of the grains. So, even though precipitates are themselves defects, the homogenization of the microstructure that results may outweigh their impact when it comes to the consideration of thermal diffusivity (this is especially true when precipitates are concentrated near grain boundaries).




Figure 6.39: APT data of 304 stainless steel samples irradiated to 0.4 dpa (a) and 28 dpa (b) in EBR-II shows evidence of precipitation of Ni and Si from the Fe matrix. [209]

Finally, the evolution of radiation-induced point defects into defect clusters, dislocation loops, and voids also acts to homogenize the microstructure of stainless steel. The transition from point defects to mesoscale defects occurs in the temperatures of interest here at relatively low dose: the transition is complete between 0.5 and 5 dpa [210]. Defects like voids, besides acting as sinks for point defects, also increase RIS in stainless steels [211, 212, 198]. All of this evidence serves to explain the thermal diffusivity response observed in Figure 6.24. At 0 dpa, the various inhomogeneities and defects in the archival material - imaged in Figure 6.12 - impede thermal diffusion. At the irradiation temperatures of interest, these inhomogeneities and defects begin to diffuse towards sinks, even at low doses. Point defects generated by the radiation begin to cluster into larger defects that create even more sinks, allowing for a more rapid “clean up” of the microstructure. So, while the evidence of radiation damage in the samples is apparent - we observe RIS at grain boundaries and defects, large increases in the number densities of voids and precipitates, and so forth - this does not automatically correspond to a reduction in thermal diffusivity.

The TGS results plotted in Figures 6.22 and 6.24 were then compared against various defect population measurements made on the sister samples of the ones provided to MIT MNM. The results of these defect populations can be found in Figure 6.21 [193]. First, the TGS results were compared against measurements of void density and void swelling. Figures 6.25 and 6.27 plot measured SAW speed versus void density and void swelling respectively, and do not show any obvious trend. Figure 6.26 plots thermal diffusivity

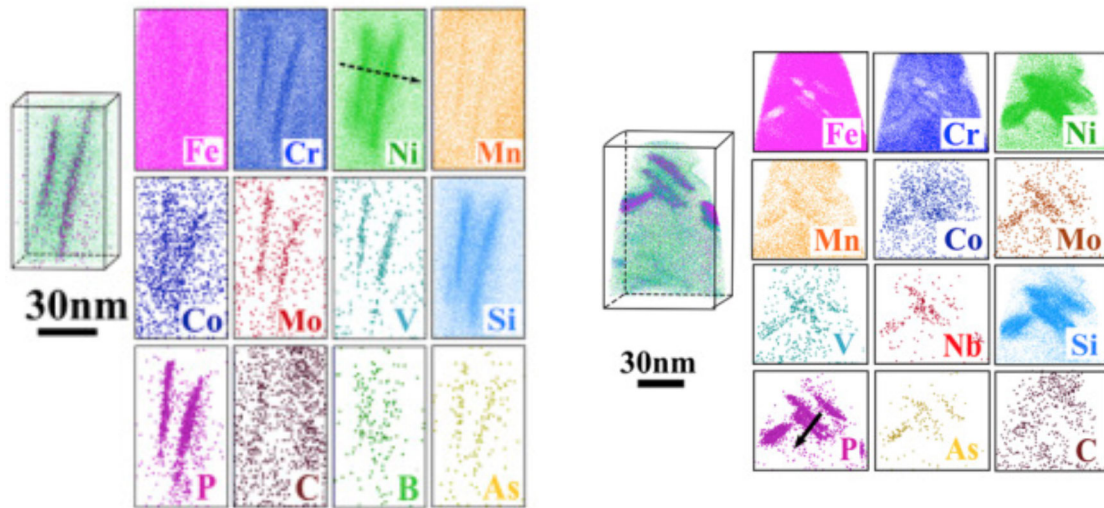


Figure 6.40: APT data of a 304 stainless steel sample irradiated to 0.4 dpa in EBR-II shows evidence of precipitation of P from the Fe matrix. [209]

Figure 6.41: APT data of 304 stainless steel samples irradiated to 28 dpa in EBR-II shows evidence of precipitation of P from the Fe matrix. [209]

results against void density, and appears to show two distinct trend lines. It is tempting to immediately assign significance to this, although it is not immediately obvious why we would observe a rapid increase in thermal diffusivity with void density at low void density ($\approx 0.5 \times 10^{21}$ per m^3), then a sharp decline when void density triples, followed by another steady increase.

To further investigate this behavior, the data is replotted in Figures 6.42-6.44 alongside average void diameter and dpa associated with each sample, all plotted against measured void density. These results indicate that the trends Figure 6.26 is likely a real result. In Figure 6.42, the two trend lines are identified as “A” and “B” for clarity. Trend line A - which shows the sharp increase of thermal diffusivity with void density for the first four data points - is associated with low dpa values. Next, we observe a sharp drop in thermal diffusivity with the fifth data point. Interestingly, this is associated with both a jump in dpa (to 28 dpa from 4 dpa) and with a sharp jump in average void diameter. In fact, Figure 6.43 shows that the average void diameter associated with this point is by far the largest in this data set. The other seven data points have data points clustered near 20 nm. Next, trend line B shows a steady increase in thermal diffusivity with void density following the drop when the void diameter and dpa both jump to their highest value.

One possible explanation for this behavior is that at low values of dpa, the microstructure “clean-up” behavior is observed. As void density increases, there are more sinks to attract point defects and alloying elements, and so we see an increase in thermal diffusivity. Then, we observe the drop in thermal diffusivity: the voids have an unusually large diameter at this step, and the dpa value has increased from 4 to 28. It is possible that the combination of void density and void size is dominating the thermal diffusivity here, and the




Figure 6.42: Thermal diffusivity measured by TGS versus void density measured via TEM analysis in [193]




Figure 6.43: Average void diameter versus void density, both measured via TEM analysis in [193]




Figure 6.44: Dpa associated with samples versus their void density measured via TEM analysis in [193]

voids are large and numerous enough to decrease thermal diffusivity more than the increased sink capacity helps to increase it. It is also possible that another damage mechanism - one that is not evident from this data set - has dominated the thermal diffusivity in the regime between 4 and 28 dpa. Without data points from that dose range, it is difficult to pinpoint the precise cause of this change. However, following the drop, trend line B shows once again an increase in thermal diffusivity with increasing void density, albeit at a less severe slope. All four data points on trend line B are associated with the 28 dpa dose level. However, note also that the other three data points on trend line B are associated with the same void diameters as the data points on trend line A. This points to a possibility that void diameter increase was the primary reason behind the thermal diffusivity drop between trend lines A and B.

Note also that in Figure 6.28, we observe an increase in thermal diffusivity with void swelling at low values of void swelling, possibly because the correlation between void swelling and void density is proportional at low doses (the first four points correspond to the samples with 0.6-4 dpa), and we are just seeing, effectively, trend line A from 6.42. However, at higher swelling rates, the trend is less obvious, although it generally shows that thermal diffusivity decreases with swelling after $\approx 1\%$ void swelling. Interestingly, considering the data in terms of void density appears to be more interesting - with regards to thermal diffusivity changes, anyways - than considering it in terms of void swelling, which is often the parameter of more direct interest in nuclear engineering.

These considerations of void swelling and void density also reveal some general observations about TGS. The first, which is a recurring story in this work, is that v_{SAW} is not always a good indicator of changing defect populations. The TGS data presented in Figure 6.25 and 6.26 came from the same set of TGS traces. However, what appeared to be random variation in the data when looking at the plot of v_{SAW} shows clear trends when thermal diffusivity is considered instead.

The same observation holds true with the next data sets, which consider precipitate density, Frank Reed loop density, and Frank Reed loop segment length. The v_{SAW} plots do not reveal anything especially interesting at first glance. However, Figure 6.30 appears to show a positive trend, with thermal diffusivity generally increasing with increased precipitate density. Since precipitate volume and diameter generally increase with

increasing precipitate density, it seems likely that this again due to an increased defect sink capacity working to homogenize the microstructure.

Thermal diffusivity is also shown to increase with increasing Frank Reed loop number density (Figure 6.32) and Frank Reed loop length density (Figure 6.34). This is further evidence for the idea that larger defects, by acting as sinks for small defects and alloying elements, increase the homogeneity of the bulk and increase the thermal diffusivity of the material. In Figure 6.34, the points are labeled with their associated dpa to show that this effect appears to be mostly independent of dose.

6.5 Conclusions of the EBR-II irradiated 304 steel study

The EBR-II campaign showed that it is possible to go from using TGS to measure model materials with carefully controlled microstructures and defect populations and compositions to measuring real, complex engineering alloys that have been irradiated for extended periods of time (14 years) in a real reactor.

It also showed that the thermal diffusivity measurements of these samples were consistent with the known microstructural changes of the samples. Because these samples have already been extensively characterized, we have the ability to compare TGS measurements with known microstructures and defect populations. In particular, the TGS results show that thermal diffusivity rises at low dose and saturates. This increase in thermal diffusivity is expected because the radiation acted to “clean up” the microstructure: we know that at 0 dpa, the alloy had a complex microstructure with many dislocation loops, and that as dpa increased, the alloy underwent a homogenizing process in the bulk as radiation-induced segregation began to occur and mesoscale defects like voids (which act as defect sinks) developed. These results are also important because it shows that thermal diffusivity appears to saturate in the samples and stop changing at low dpa. In terms of an engineering result, this means that, for 304 stainless steel, one might conclude that irradiation with fast neutrons will not lead to a detrimental decrease in thermal diffusivity. (If anything, the opposite has occurred - the material’s ability to transfer heat efficiently has improved, even though other damage effects like voids and RIS are still of concern to the nuclear engineer.)

These measurements also showed how TGS results could be correlated directly with specific types of defects (in this case, defect populations that were measured directly with TEM by prior researchers). This is very promising for the potential of TGS to determine defect species and measure actual defect populations in addition to measuring radiation-induced material property changes. Importantly, it was also demonstrated that one needs to examine the thermal diffusivity results in order for these dependencies to become apparent. In this case, looking at v_{SAW} alone does not yield any definitive trends. This is another case where, once the acoustic damping analysis techniques are finalized, it will be of interest to return to this data and reanalyze it to extract that information as well.

These samples spent fourteen years in a reactor core - not a practical amount of time for a researcher who wishes to validate a new material quickly. However, research has shown that preconditioning steels by irradiating them with neutrons to low doses, and then completing the rest of the exposure with ion irradiation,

is a valid (and faster) way to study neutron effects [223, 224]. This means that, once fully adapted for *in situ* use with an ion accelerator, TGS will be a valuable tool for studying neutron irradiation in alloys, even though it can't be used *in situ* in a reactor core. Interesting future work in this regard would be to use some of the EBR-II samples that have already been irradiated with neutrons and irradiating them in an ion accelerator, while using TGS to track radiation damage in the samples in real time. This would allow for the prediction of how 304 stainless steel might respond to more extreme dpa levels and lengthier exposures to neutron irradiation.

Chapter 7

Transient grating spectroscopy and radiation materials science: moving forward

The purpose of this thesis was to establish transient grating spectroscopy as a valid method for detecting radiation-damage-induced material property changes, particularly in metals and alloys. As was established in the prologue of this thesis, TGS has the power to “break the bottleneck” that currently frustrates efficient innovation in radiation materials science.

Right now, the validation of a material for use in a particular nuclear application involves designing and executing an extensive irradiation test matrix, followed by destructive analysis of each sample to build a picture of how the material will respond to a large range of doses, dose rates, radiation types, and operating temperatures. Both steps are very time consuming and very expensive.

TGS has the potential to be a fast, cost-effective means of characterizing radiation damage and its effects in a sample. The equipment required to build a TGS facility is relatively cheap compared to the cost of a full suite of destructive analysis equipment. The TGS measurements are quick: in this study, collection of one data point took about five minutes total, but that amount of time can easily be reduced for applications where less resolution is required. A single TGS measurement encodes data that can be used to measure a sample’s Young’s modulus, thermal diffusivity, and acoustic damping properties, all of which are affected differently by radiation damage. These three parameters provide a much fuller picture of how the material is being changed by radiation damage than one of them would alone.

Another important reason that TGS is so promising for improving the efficiency of radiation materials testing is the fact that it is non-contact and it has an adaptable geometry. This makes it ideal for modification to an *in situ* setup, in which the TGS facility is used concurrently with an ion accelerator that is irradiating a

sample, so that the TGS measurements are being performed in real time as the sample is irradiated. We have already shown in this thesis that the TGS setup presented here is sensitive to changes that occur in materials even at very low dose. This means that a single sample can be used to measure how a material will respond to radiation over a broad dose range (instead of the current practice of using many samples, each irradiated to an individual dose, and sacrificing dose resolution by decreasing the size of the irradiation matrix in order to save money and time). Figure 7.1 shows how a radiation materials science data campaign might proceed with and without TGS. Using traditional methods, one needs to plan on 1-2 weeks of preparation, irradiation, and data collection per sample in the matrix. Using TGS, the entire data set for the research campaign can be collected in one week or less.¹

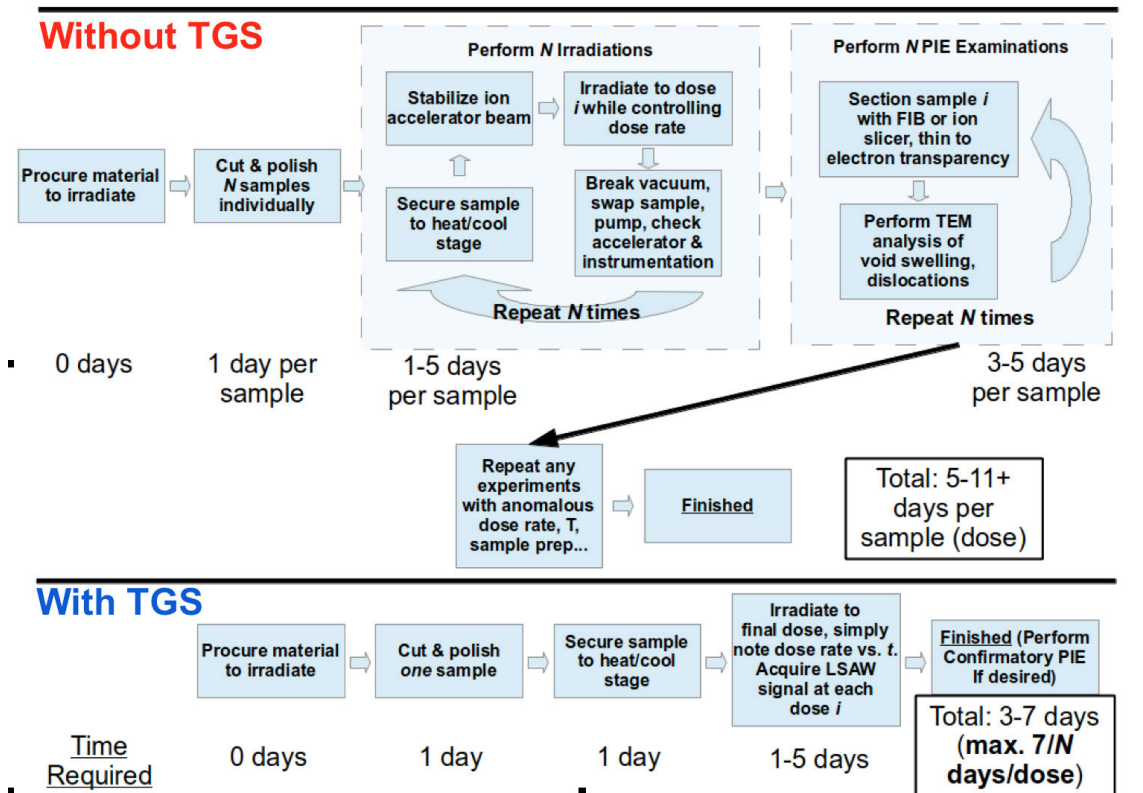


Figure 7.1: TGS enables more efficient radiation materials science research campaigns. Schematic by M. P. Short.

This thesis work has shown that TGS will be a valid method for characterizing the radiation damage response of both model materials and engineering alloys. (Summaries of the contributions from each of the three studies of this project can be found on page 132 (niobium), page 186 (B2-NiAl), and page 214

¹Of course, if one wanted to repeat the measurement at different temperatures or beam currents, that would extend the data collection time accordingly, but the overall research campaign would still be completed much faster than it would in the non-TGS scenario.

(EBR-II.) The immediate next step for this project will be using TGS to detect the onset of tungsten fuzz. Preliminary results (not presented in this thesis) have already been completed and used to inform the refinement of the experimental plan. Collection of this data is currently underway.

One highly recommended project for “future work” entails the refinement of the acoustic damping data collection. As has been explained previously, our current physical models for explaining acoustic damping in SAW are inadequate for our purposes, and our present code for extracting the acoustic damping parameter can only be considered as an estimation at best. A research contribution that both improved the physical model of acoustic damping and updated the analysis codes used to extract it from the TGS signal would vastly improve the already significant capabilities of TGS to characterize radiation damage. It is expected that acoustic damping will be highly sensitive to the presence of different types and densities of radiation defect. Right now, we can only use the elasticity and thermal data with confidence. Nailing down the acoustic damping part of the signal will fully flesh out the information we can extract from the TGS signal.

A second recommended future work project is to more fully use the capability of TGS to detect both SAW and PSAW. Because a PSAW radiates energy into the bulk, it behaves somewhat differently than a SAW, but it also may be able to add a fourth dimension to the information that is extractable from the TGS signal if the rate and magnitude of this bulk loss has components that are independent of the properties (like thermal diffusivity) that are already measurable from the TGS signal. This would involve improving the resolution of the signal and of the FFT, and making sure to define both the SAW peak and the PSAW peak. At the very least, even if the notion of using the PSAW bulk loss as a possible damage indicator is abandoned, this will further improve the SAW measurements by verifying that one is always using information from a SAW wave. This is particularly important for collecting elasticity information on certain materials (especially single crystals), since the frequency peaks of the SAW and PSAW often overlap.²

From an engineering perspective, it is worthwhile to consider how the setup might be further improved in the future for better efficiency. In particular, there is significant interest in automating the TGS focusing procedure, using the known geometry of the setup, so that the sample and lasers can be automatically aligned properly. The setup optics and the phase adjusts could then be adjusted automatically until the positive and negative signals were maximized. Currently, the sample is aligned and the signal is maximized manually, watching the oscilloscope to determine when the signal has been optimized. Automating this procedure would improve the consistency of the measurements from one spot to another (and one researcher to another). It would also allow for the eventual adaptation of the TGS system to a standardized “plug-and-play” piece of lab equipment, which would further expand its use. Another idea is to build on the dual-heterodyne modification that has already been completed on the MIT MNM TGS facility, and create a TGS system that is capable of taking dual-heterodyne measurements at more than one spot at a time. This is important for heterogeneous sample that are to be studied in real time, as one likely needs to consider more than one spot on the sample surface in order to gain a full picture of the sample’s characteristics. However, the current

²This phenomenon was evident in the $v_{SAW}(\theta)$ measurements carried out on the 0 dpa control niobium sample, as explained on page 113.

TGS system precludes the ability to measure more than one spot for a given dose and time during an *in situ* measurement.

Finally, the most important aspect of the future work for this project involves developing TGS for use *in situ* with an ion beam. This will allow for real-time measurements of radiation damage as it occurs, providing the ability to obtain unprecedented dose-and-time resolution in the measurement of radiation-damage-driven microstructure evolution. This phase of nuclear-materials-TGS development is already well underway, as Cody Dennett has recently completed construction of a TGS system at Sandia National Laboratories that can measure samples while they are irradiated with an ion beam. The first successful tests of the completed facility were recently carried out. This is extremely promising for the future outlook of TGS, as it demonstrates that TGS is fully capable of being adapted for *in situ* use.

Chapter 8

Appendix

8.1 Orientation distribution functions of the six cold-worked niobium samples

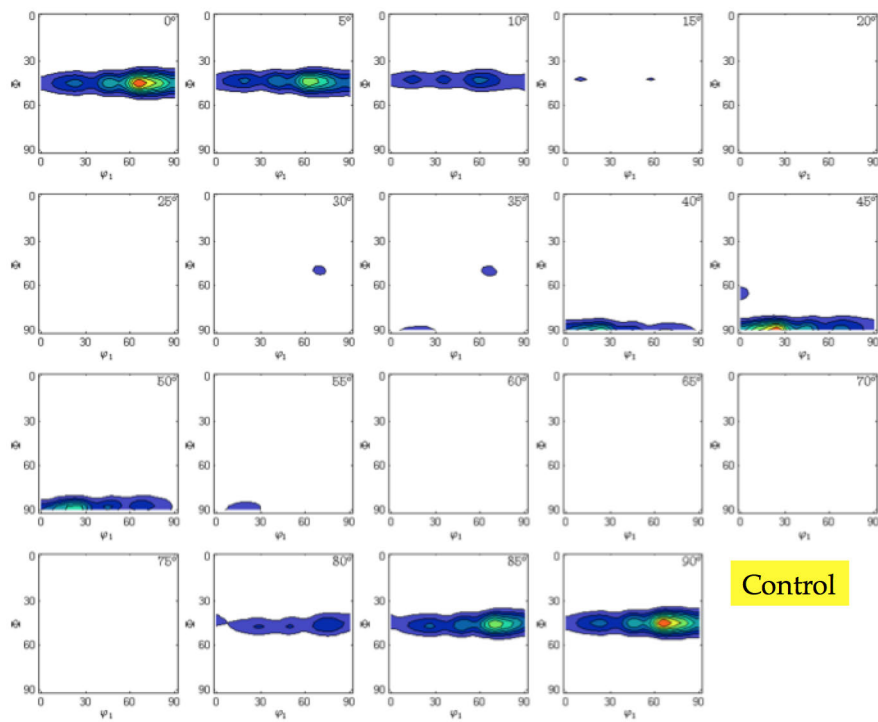


Figure 8.1: Orientation distribution function obtained from the pole figure data of the control Nb sample

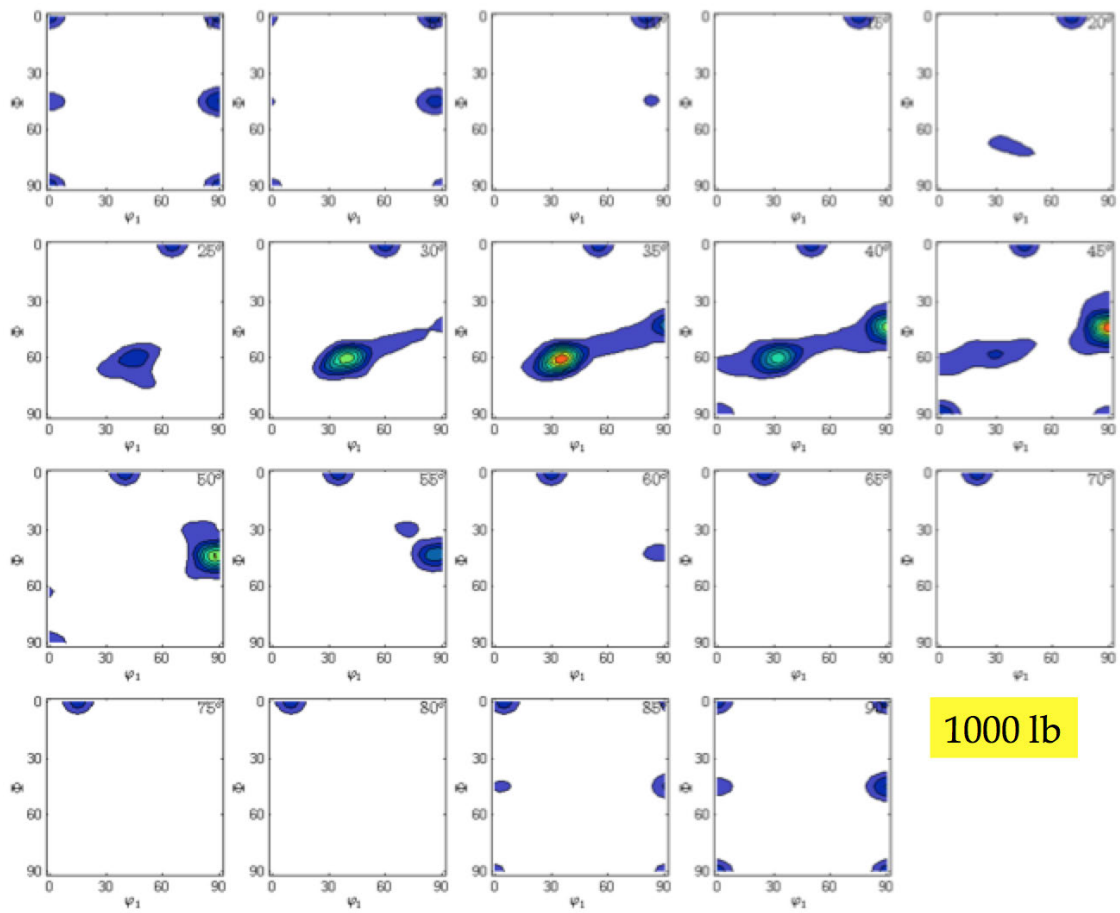


Figure 8.2: Orientation distribution function obtained from the pole figure data of the 1000 lb Nb sample

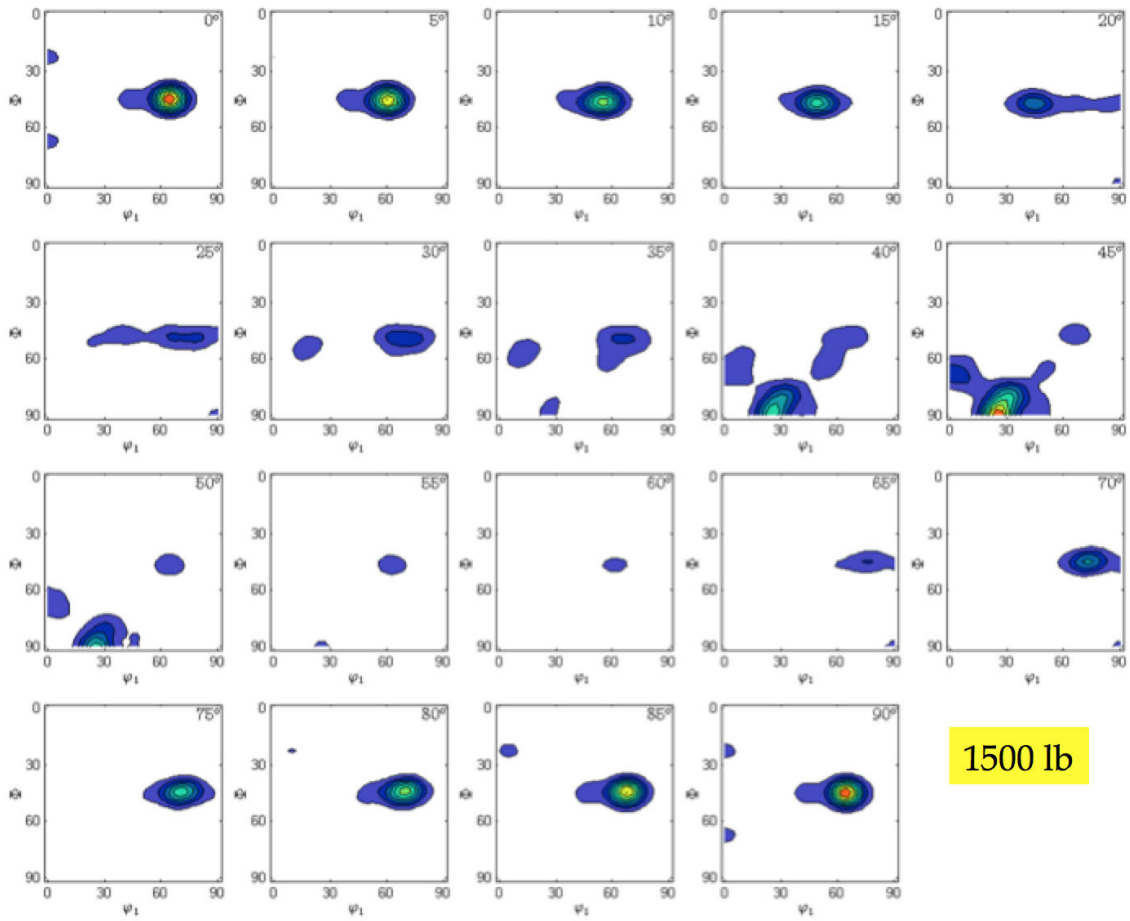


Figure 8.3: Orientation distribution function obtained from the pole figure data of the 1500 lb Nb sample

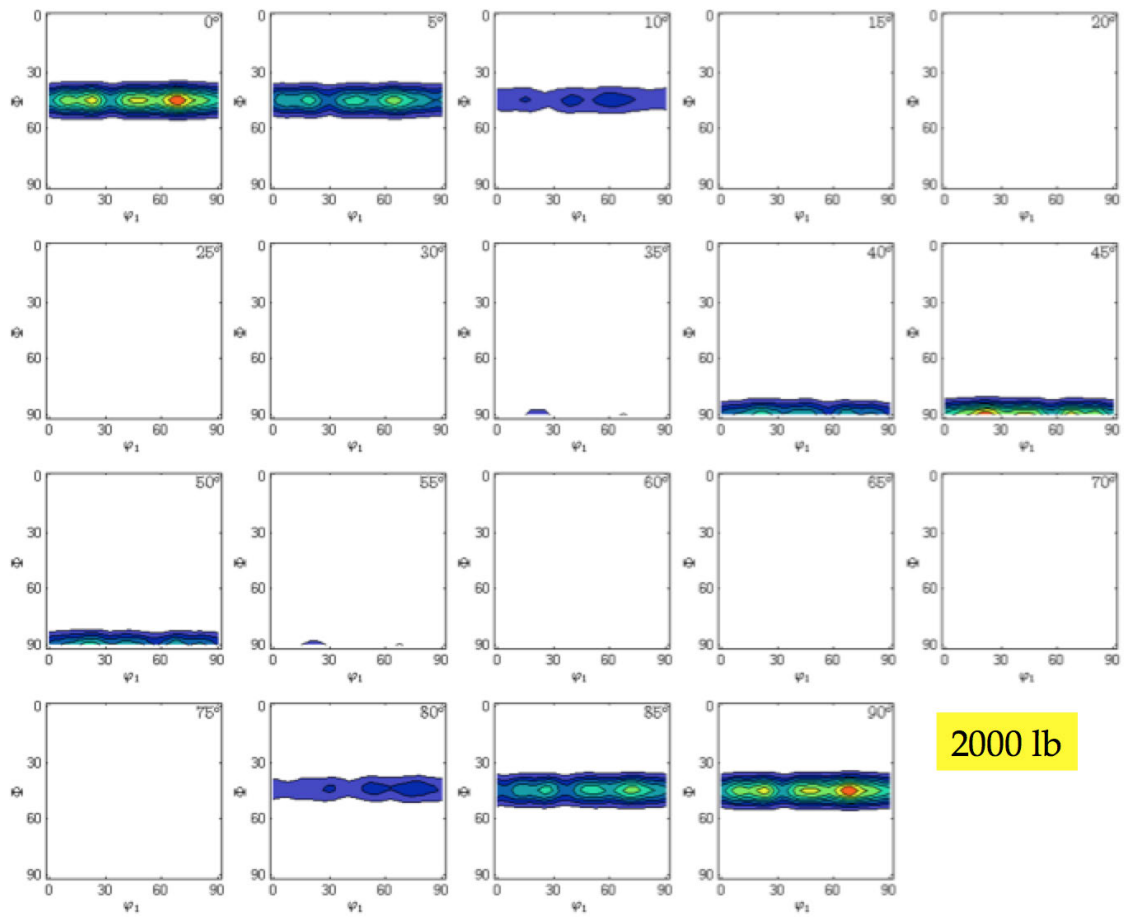


Figure 8.4: Orientation distribution function obtained from the pole figure data of the 2000 lb Nb sample

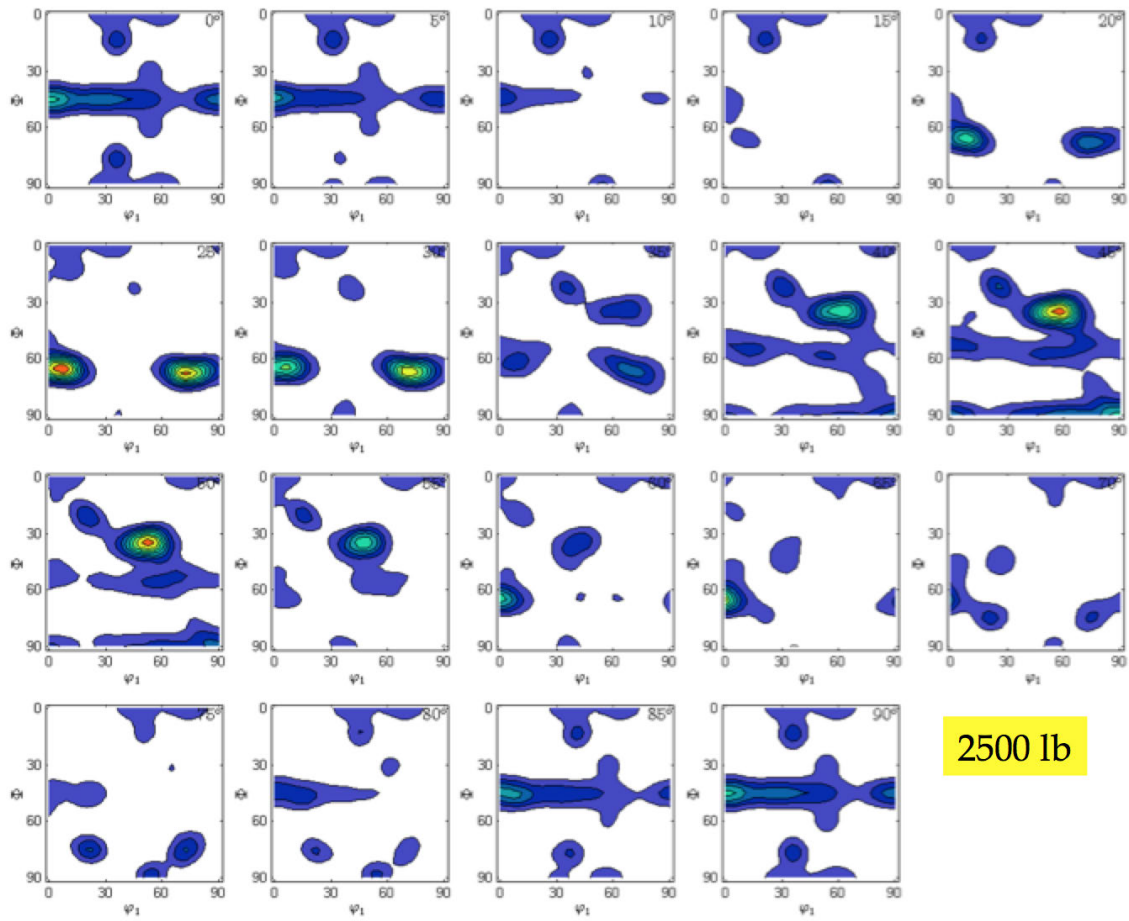


Figure 8.5: Orientation distribution function obtained from the pole figure data of the 2500 lb Nb sample

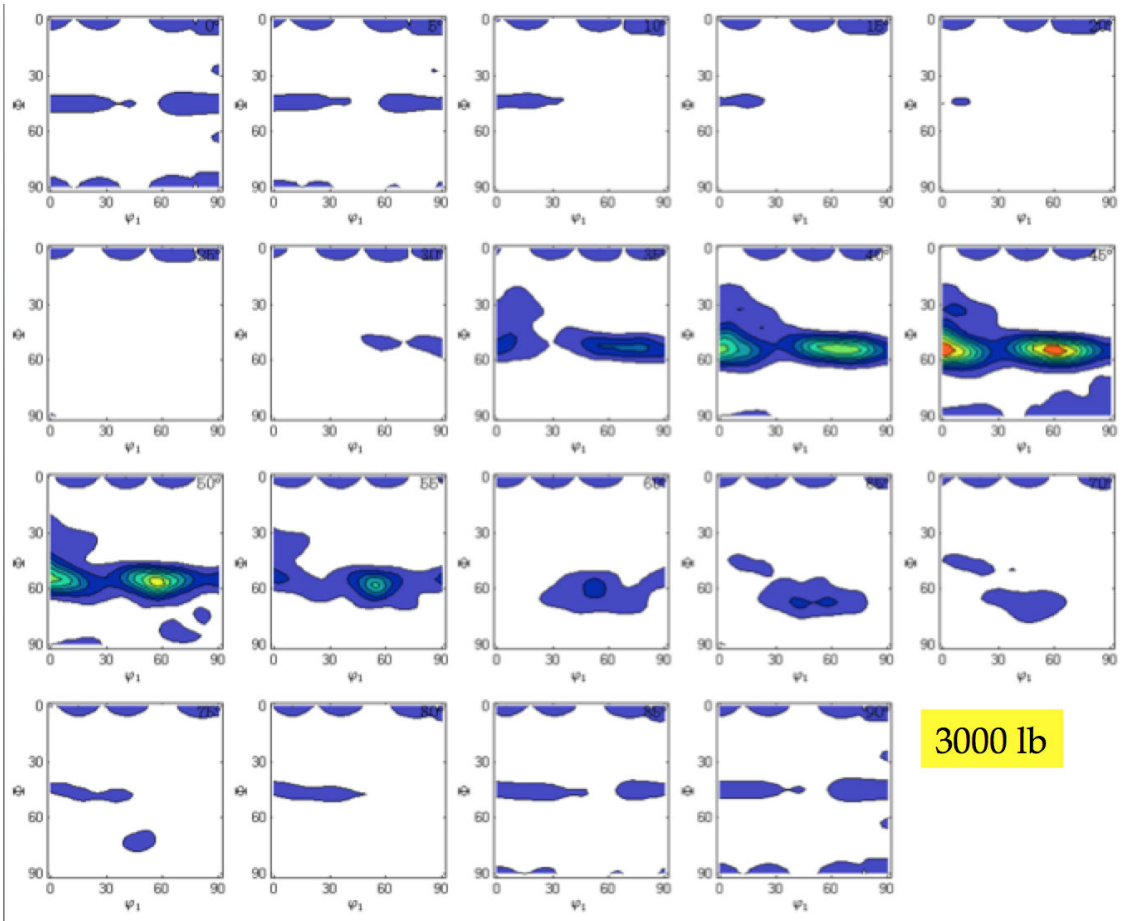


Figure 8.6: Orientation distribution function obtained from the pole figure data of the 3000 lb Nb sample

8.2 Automating MATLAB analysis of the thermal and acoustic components of the TGS traces

The shell script automates the MATLAB analysis used to extract the thermal diffusivity and acoustic damping parameters from the TGS traces. In order to use it, trace files should be stored in separate folders by date (unless the file naming convention is changed to not include the date in the file name). Each folder containing traces to be analyzed should include this shell script, in addition to the function files that MATLAB will need to call to complete the analysis. The grating spacing in the shell script should be updated to reflect the calibrated grating spacing for the particular day. The output files appear in .csv files with the basename of the trace files and the spot number. The .csv file will include the thermal diffusivity for each of the three traces and the acoustic damping for each of the three traces for a given surface rotation value. If multiple spots were measured for the same rotational position on the same day, each additional spot will have its own .csv file.

```

#!/bin/bash

## Set the real grating spacing in microns here

gs=5.5

## Iterate over all .txt files in the folder, assuming they are the
data files

## Loop over three variables: spot number, trace number, rotation
angle,
## noting that rotation angle can have a negative "neg" marker.

## Process filenames first. Loop over filenames, converting 'neg' to a
minus sign.

rename _neg _- *.txt

## Find the base part of the filenames in each folder which are common

set -- *.txt
files=$#
name=$1
while [ -n "$name" ]
do
    set -- $name*
    [ $# -eq $files ] && break
    name=${name%?}
done
echo "File name base is $name"
basename=$name

## First, iterate over all possible spots

for spot in {1..2}
do

## Next, loop over all possible angles, from -20 to 180. We'll skip
ones that
## don't show up. Assume every 5 degrees.

        for deg in `seq -20 5 180`
        do

## If a temporary script does not exist, create it.

## First check that this file exists.
## Construct what we think the filename should be.
## Naming template: Nb2000_Irrad_1dpa-2018-04-12-05.50um-spot02_175d-
POS-1.txt

```

```

FILE=$basename
FILE+=$spot
FILE+=_
FILE+=$deg
FILE+=d-POS-
FILE+=1
FILE+=.txt
if [ -f "$FILE" ]; then
    echo "Dataset $FILE exists."

    if [ -f "TempMATLAB.m" ]; then
        rm TempMATLAB.m
    else
        touch TempMATLAB.m
    fi

## Write out the first stuff for the MATLAB file

    echo "gs = $gs;" >> TempMATLAB.m
    echo "thermal = zeros(1,3);" >>
TempMATLAB.m
    echo "thermalerr = zeros(1,3);" >>
TempMATLAB.m
    echo "acoustic = zeros(1,3);" >>
TempMATLAB.m

## Finally, loop over all traces. Note that each has a POS and NEG, so
## we only care about matching the POS one.

    for trace in {1..3}
    do

## First check that this file exists.
## Construct what we think the filename should be.
## Naming template: Nb2000_Irrad_1dpa-2018-04-12-05.50um-spot02_175d-
POS-1.txt

        FILE=$basename
        FILE+=$spot
        FILE+=_
        FILE+=$deg
        FILE+=d-POS-
        FILE+=$trace
        FILE+=.txt
        if [ -f "$FILE" ]; then
            echo "File $FILE exists."

## Echo the lines to analyze this set of traces

```



```

                                echo " " >> TempMATLAB.m
                                echo "[a,aa,aaa] =
thermal_phase('$basename$spot'_"$deg"d-
POS-"$trace".txt', '$basename$spot'_"$deg"d-NEG-"$trace".txt', $gs, 2);"
>> TempMATLAB.m
                                echo " " >> TempMATLAB.m
                                echo "thermal(1,$trace)=a;"
>> TempMATLAB.m
                                echo
"thermalerr(1,$trace)=aa(1);" >> TempMATLAB.m
                                echo
"acoustic(1,$trace)=aaa;" >> TempMATLAB.m
                                fi

## Take the outputs of that MATLAB script, and append them to a CSV
file.

                                done

## Write out the array as an ASCII comma-delimited file

                                echo " " >> TempMATLAB.m
                                FileToWrite=$deg
                                FileToWrite+="d-spot0"
                                FileToWrite+=$spot
                                FileToWrite+=".csv"
                                echo
"dlmwrite('$FileToWrite',thermal);" >> TempMATLAB.m
                                echo
"dlmwrite('$FileToWrite',thermalerr,'-append');" >> TempMATLAB.m
                                echo
"dlmwrite('$FileToWrite',acoustic,'-append');" >> TempMATLAB.m
                                echo " " >> TempMATLAB.m
                                echo "quit;" >> TempMATLAB.m

## Run the MATLAB script

                                matlab -nodisplay -nosplash -
nodesktop -r "run('TempMATLAB.m')"

## End all the loops.

                                fi

## Finish up the MATLAB script

                                done
done

```

```

## Finally, fold all the .csv individual data files into one. Call a
separate script for this.
## Naming convention example: -20d-spot01.csv

## First, iterate over all possible spots

for spot in {1..2}
do

## Check to make sure this spot's file doesn't exist. If it doesn't,
create it.

    FILEOUT=$(basename
FILEOUT+=${spot}
FILEOUT+=.csv

    if [ -f $FILEOUT ]; then
        rm $FILEOUT
    fi
    touch $FILEOUT

## First, write out axis labels

    echo
"Degree,Thermal_1,Thermal_2,Thermal_3,Thermal_avg,Terr_1,Terr_2,Terr_3
,Terr_max,Acoustic_1,Acoustic_2,Acoustic_3,Acoustic_avg" > $FILEOUT

## Next, loop over all possible angles, from -20 to 180. We'll skip
ones that
## don't show up. Assume every 5 degrees.

    for deg in `seq -20 5 180`
do

## First check that this file exists.
## Construct what we think the filename should be.

        FILE=$deg
FILE+=d-spot0
FILE+=${spot}
FILE+=.csv
        if [ -f "$FILE" ]; then
            echo "Dataset $FILE exists."

## Start writing out each spot's file, in the following format:
## DEGREE   THERMAL_1   THERMAL_2   THERMAL_3   THERMAL_AVG
THERMAL_ERR_1   THERMAL_ERR_2   THERMAL_ERR_3   THERMAL_ERR_MAX

```

```
THERMAL_ERROR_TOTAL  ACOUSTIC_1  ACOUSTIC_2  ACOUSTIC_3  
ACOUSTIC_AVG
```

```
## Get the 1st line of the file as the thermal stuff, the 2nd line as  
the thermal errors, and the 3rd line as the acoustic stuff
```

```
thermal=`sed -n '1p' < $FILE`  
thermalerr=`sed -n '2p' < $FILE`  
acoustic=`sed -n '3p' < $FILE`
```

```
echo  
$deg,"$thermal", "$thermalerr", "$acoustic >> $FILEOUT
```

```
## End all the loops.
```

```
fi  
done  
done
```

8.3 Additional results from the comparison of experimental $v_{SAW}(\theta)$ results with calculated $v_{SAW}(\theta)$ results for single crystal niobium

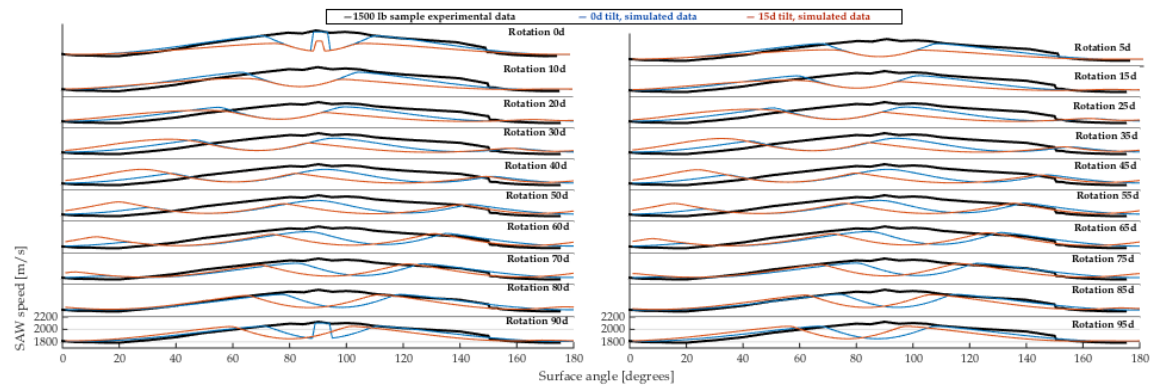


Figure 8.7: The black line shows experimental data for SAW speed versus degree of surface rotation for the 1500 lb sample, which is expected to be well-aligned to the $\langle 110 \rangle$ plane. The blue and red lines are the predicted SAW response for 0° and 15° of tilt. These are given for rotations of 0 to 95° in 5° increments. Other tilt conditions have been omitted.




Figure 8.8: Experimental data and predicted SAW speed calculations as a function of surface angle for the 2000 lb sample

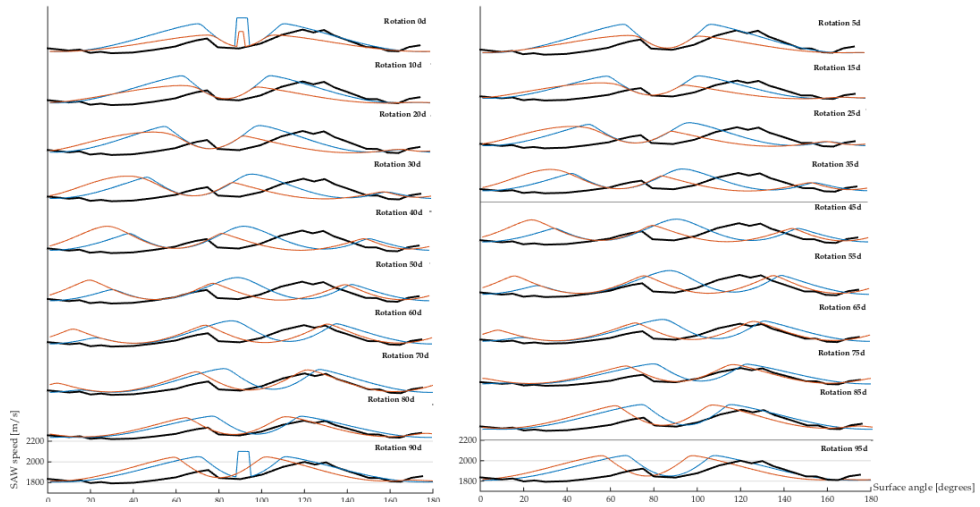


Figure 8.9: Experimental data and predicted SAW speed calculations as a function of surface angle for the 2500 lb sample

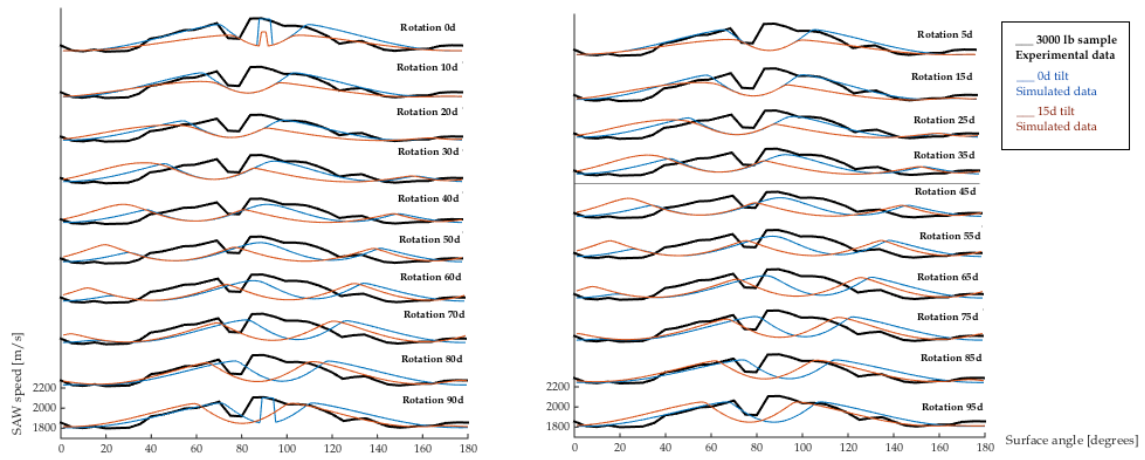


Figure 8.10: Experimental data and predicted SAW speed calculations as a function of surface angle for the 3000 lb sample

8.4 Pre-averaging thermal diffusivity and acoustic damping results for the 1500 lb and 2000 lb niobium samples




Figure 8.11: Thermal diffusivity results, 1500 lb sample




Figure 8.12: Acoustic damping results, 1500 lb sample




Figure 8.13: Thermal diffusivity results, 2000 lb sample




Figure 8.14: Acoustic damping results, 2000 lb sample

8.5 Additional notes on the material properties of B2-phase intermetallic NiAl

8.5.1 Elasticity of intermetallic NiAl

The general case of the stress-strain relationship in a linear elastic material can be expressed using Voigt notation as follows:

$$\sigma = \mathbf{C}\varepsilon \quad (8.1)$$

$$\begin{bmatrix} \sigma_{11} \\ \sigma_{22} \\ \sigma_{33} \\ \sigma_{23} \\ \sigma_{31} \\ \sigma_{12} \end{bmatrix} = \begin{bmatrix} C_{1111} & C_{1122} & C_{1133} & C_{1123} & C_{1131} & C_{1112} \\ C_{2211} & C_{2222} & C_{2233} & C_{2223} & C_{2231} & C_{2212} \\ C_{3311} & C_{3322} & C_{3333} & C_{3323} & C_{3331} & C_{3312} \\ C_{2311} & C_{2322} & C_{2333} & C_{2323} & C_{2331} & C_{2312} \\ C_{3111} & C_{3122} & C_{3133} & C_{3123} & C_{3131} & C_{3112} \\ C_{1211} & C_{1222} & C_{1233} & C_{1223} & C_{1231} & C_{1212} \end{bmatrix} \begin{bmatrix} \varepsilon_{11} \\ \varepsilon_{22} \\ \varepsilon_{33} \\ 2\varepsilon_{23} \\ 2\varepsilon_{31} \\ 2\varepsilon_{12} \end{bmatrix}$$

where C_{ijkl} are entries in the elasticity tensor.

\mathbf{C} is simplified to a second order tensor by asserting the equivalence of (23) to (32), (13) to (31), and (12) to (21):

$$C_{\alpha\beta} = \begin{bmatrix} C_{11} & C_{12} & C_{13} & C_{14} & C_{15} & C_{16} \\ C_{12} & C_{22} & C_{23} & C_{24} & C_{25} & C_{26} \\ C_{13} & C_{23} & C_{33} & C_{34} & C_{35} & C_{36} \\ C_{14} & C_{24} & C_{34} & C_{44} & C_{45} & C_{46} \\ C_{15} & C_{25} & C_{35} & C_{45} & C_{55} & C_{56} \\ C_{16} & C_{26} & C_{36} & C_{46} & C_{56} & C_{66} \end{bmatrix}$$

Because the structure of NiAl is cubically symmetric, there are only three independent elements in the elasticity tensor.

$$C_{\alpha\beta} = \begin{bmatrix} C_{11} & C_{12} & C_{12} & 0 & 0 & 0 \\ C_{12} & C_{11} & C_{12} & 0 & 0 & 0 \\ C_{12} & C_{12} & C_{11} & 0 & 0 & 0 \\ 0 & 0 & 0 & C_{44} & 0 & 0 \\ 0 & 0 & 0 & 0 & C_{44} & 0 \\ 0 & 0 & 0 & 0 & 0 & C_{44} \end{bmatrix}$$

It is the tensor above that is important to most papers that discuss the elasticity of NiAl. Most commonly, we see the elastic constants described with three terms: C , C_L , and C' . These definitions are from [161], but they are repeated throughout the literature.

$$C = C_{44} \quad (8.2)$$

$$C_L = \frac{1}{2}(C_{11} + C_{12} + 2C_{44}) \quad (8.3)$$

$$C' = \frac{1}{2}(C_{11} - C_{12}) \quad (8.4)$$

Figure 8.15 is a graph of C , C_L , and C' for various compositions of NiAl as reported in [161]. Rusovic et al. found that shear stiffness decreased with increasing nickel concentration (which corresponds to increasing concentration of nickel antisite defects on the aluminum sublattice) [161]. The elastic constants were found to decrease with increasing vacancy concentration (which corresponds to decreasing nickel concentration and increasing concentration of nickel vacancies on the nickel sublattice). These values are reported in Table 8.1.

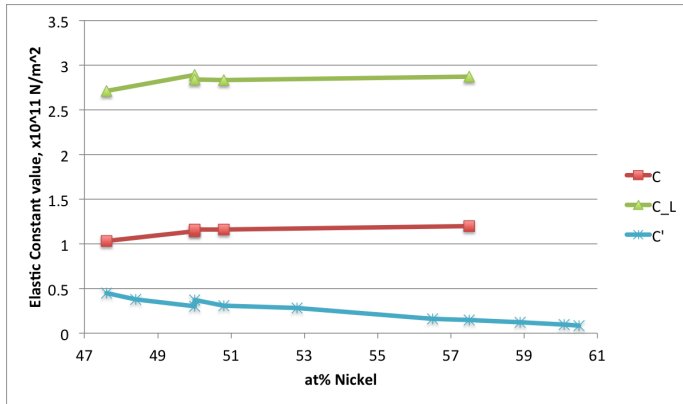


Figure 8.15: Elastic constants of NiAl as reported in [161].

Constant	% change in value per at% vacancy increase
C_{44}	-2.4
C_L	-1.3
C'	0
C_{11}	-0.5
C_{12}	-1.2

Table 8.1: Elastic constants decrease with increasing vacancy concentration, as reported in [161]

It is noted in [162] that NiAl elastic moduli will vary depending on how the samples were processed.

C , C_L , and C' are plotted in Figure 8.15 as a function of at% nickel. These values are compiled from reported values in [161], which includes values from [213] and [214]. [161] also reports that C_L and C decrease with temperature. C_1 decreases with temperature at low nickel concentrations (Ni-vacancy dominated), but increases at higher nickel concentrations (Al-ASD dominated). See Figure 8.16 for the plots from this paper.

8.5.2 Hardness of intermetallic NiAl

Hardness measures a sample's resistance to permanent deformation due to compressive force. A higher hardness value indicates a higher degree of resistance.

[173] reports that the hardness of a NiAl sample is relatively independent of quenching temperature. T_{quench} has the biggest impact on hardness near the stoichiometric composition (which makes sense, because

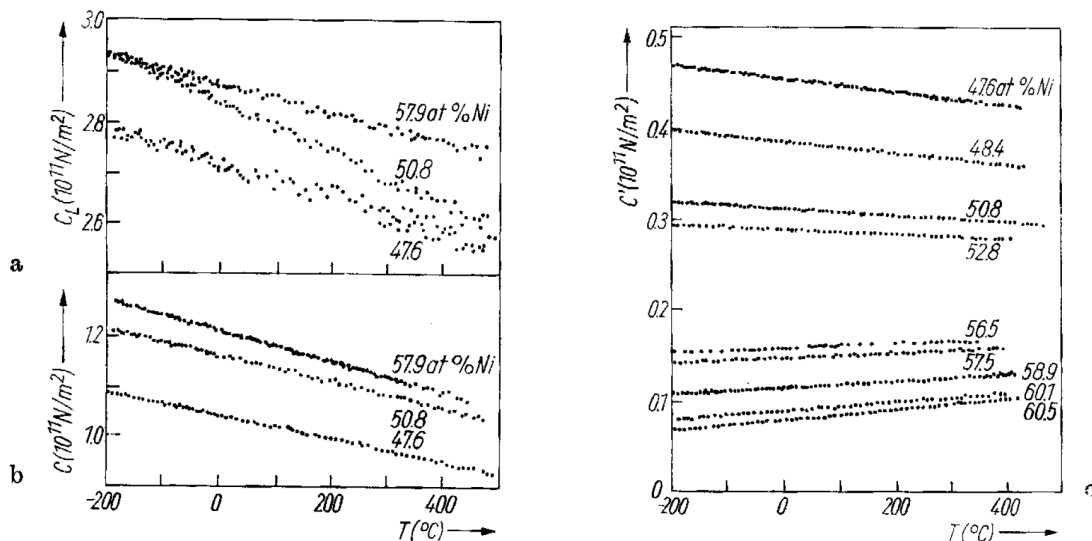


Figure 8.16: Elastic constants of intermetallic NiAl of various concentrations as a function of temperature, from [161].

this is where the material's structural behavior changes).

[173] also reports that vacancies increase NiAl hardness more strongly than antisites do. (In this case, we should experimentally observe that NiAl samples become harder with increasing aluminum content. If NiAl is observed to become harder with increasing nickel content, due to a hardening effect from nickel ASD on the aluminum sub lattice, then we expect that $d(\text{Hardness})/d(\text{Ni},+)$ for at% Ni > 50% is less than $d(\text{Hardness})/d(\text{Ni},-)$ for at% Ni < 50%. Polycrystalline samples of NiAl are hardest below $0.45T_{\text{melt}}$ (1676°C) and exhibit brittle behavior. Above that temperature, they exhibit more ductile behavior [168].

8.5.3 Ductility of intermetallic NiAl

Ductility measures a material's tendency to deform under tensile stress. Single crystal NiAl (which is what we are working with) is reported in [162] to have a limited tensile ductility. [168] contradicts this somewhat, stating that single crystal NiAl is ductile at low temperatures. [168] reports that polycrystalline NiAl is brittle, and tends to fail along the grain boundaries under tensile stress instead of plastically deforming.

[154] indicates that this brittleness is due to the B2-ordered structure having only three independent slip systems, whereas more ductile materials have at least five. When NiAl deforms, it tends to do so by glide along the $\langle 100 \rangle \{011\}$ and $\langle 100 \rangle \{001\}$ slip systems [162]. These are illustrated in Figure 8.17. [168] reports that single crystal NiAl is observed to exhibit compressible behavior along the $\langle 110 \rangle$ axes.

It is interesting to note that the literature is often contradictory when it comes to describing the ductility of NiAl. For example, the following assertions are made in different sources regarding the impact of

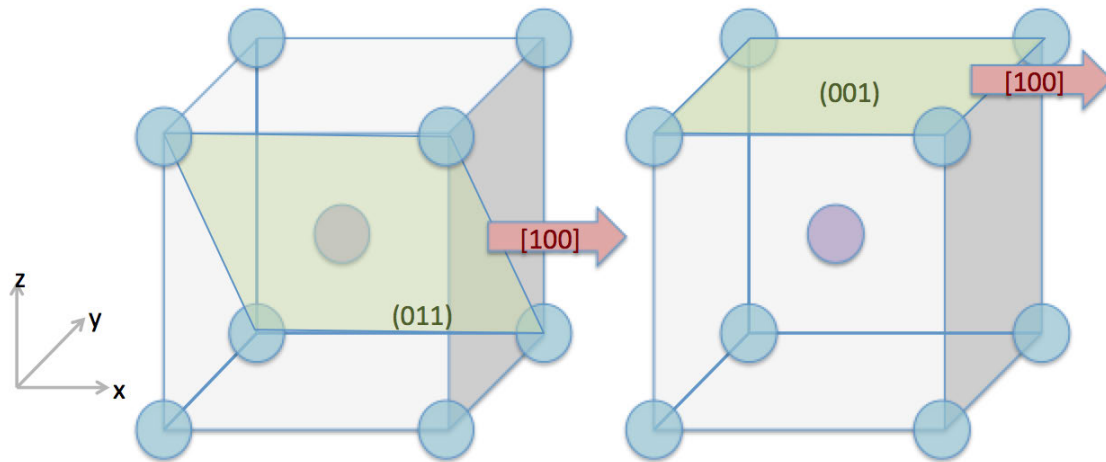


Figure 8.17: NiAl tends to deform by glide in the $\langle 100 \rangle$ family of directions in the $\{011\}$ and $\{001\}$ plane families. Above, the $[100]$ direction in the (011) and the (001) planes of a B2-structured system is illustrated.

temperature on NiAl ductility - some of which are in opposition to each other.

- NiAl exhibits high ductility at room temperature. [215]
- NiAl has low ductility at room temperature. [155]
- The ductility of NiAl increases with temperature. [176]
- There is good ductility at all NiAl compositions above 800°C. [155]
- Deformation is aided by diffusion at higher temperature. [168]

It's important to keep this in mind when using literature to explain or contradict experimental results. NiAl is particularly interesting because it exists over a wide range of composition and temperature, and both of these factors impact its material properties. Furthermore, the grain size of the sample and the method by which the sample was made also impacts material behavior.

8.5.4 Stiffness

Young's modulus (E) is a measure of material stiffness (resistance to deformation due to an applied force) and can be calculated as:

$$E = \sigma(\epsilon)/\epsilon \quad (8.5)$$

where σ is tensional stress and ε is extensional strain. Young's modulus of β 2-NiAl is found to decrease with increasing temperature for aluminum-rich compositions of NiAl, as shown in Figure 8.18 [175].¹

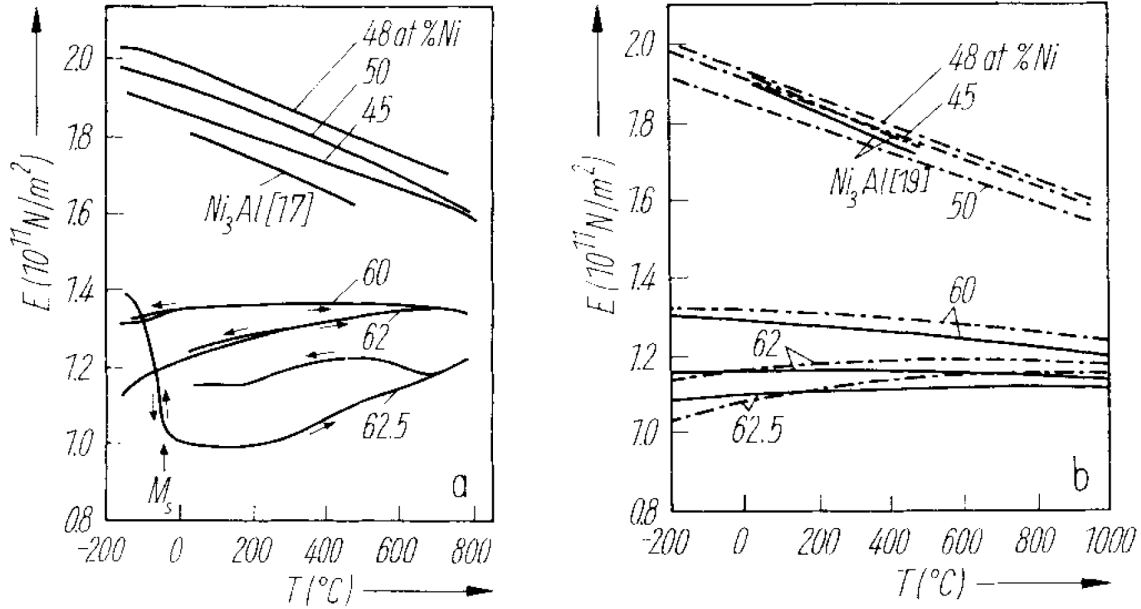


Figure 8.18: On the left, Young's modulus is graphed as a function of temperature. The top three lines (45-50 at% Ni) correspond to NiAl compositions that we consider in our work. The modulus decreases with temperature. These values are close to the values computed from single crystal data. [175]

8.5.5 Anisotropy

Elastically anisotropic materials have different elastic properties in different directions. Anisotropy matters for this project because it means that the TGS signal will be directionally dependent on single crystals, which means crystal orientation must be accounted for during measurements (as explained in the previous chapter).² NiAl is reported to have an anisotropy factor $E_{100}/E_{100} \approx 3.3$ by [162]. A value of ≈ 3.6 for NiAl at 300K is reported by [176].

¹The original caption for the right-hand graph in Figure 8.18 is "computed from single crystal data [5,19] for the same alloy compositions (- - - Kröner [12], - - - Hill [10])."

²This is why tungsten is such an attractive material for many TGS studies: its isotropy makes it easier to study, and also makes it a good material for calibration samples.

8.5.6 Poly- versus single crystal

Both polycrystalline and single crystal NiAl samples are characterized with TGS in this research. In general, polycrystalline NiAl is found to be less ductile and more hard than single crystal NiAl, although this is not necessarily true above $0.45T_M$ [168, 162] In [169], it was suggested that it can be easy to overestimate the number of vacancies in polycrystalline intermetallic samples, because there were often voids and microvoids in the samples that remained even after grinding the samples to fine powders. Single-crystal samples were found to have no large voids and few microvoids, making the vacancy measurements more accurate. Single crystal NiAl is orientation dependent, as noted in [168]. This is expected, since the material is anisotropic.

8.6 B2-NiAl TGS LAMMPS simulations: supplementary material

This section provides the following:

- An example LAMMPS input script
- The MATLAB code used to generate B2-NiAl test structures
- The shell script for running multiple instances of the structure generation code on the cluster

8.6.1 LAMMPS input structure

An example basic input file for LAMMPS is included on page 247. This input file is from early stages of the LAMMPS campaign for the NiAl project, when test structures were being built directly in LAMMPS in the “ATOM DEFINITION” section. Later, this input file would be modified to read in the MATLAB script provided later in this appendix. This allows for more effective structure building and more control over the NiAl structure (the method shown here simply arranges Ni and Al atoms at random in the concentrations specified).

8.6.2 MATLAB code for building NiAl structures

The following pages give the MATLAB code used to build the NiAl structures that were imported into LAMMPS. This code allows the specification of structure size, vacancy concentration, antisite defect concentration, and ratio of vacancies to antisite defects. This code is useful because it allows the specification of a structure that is difficult to create directly in LAMMPS, while still allowing a degree of randomization (e.g. in vacancy placement).

This code outputs a text file that is formatted to be readable by LAMMPS, as in Figure 8.19. The numbers 1 and 2 specify the atom type (Ni or Al). The text file needs to be compatible with the LAMMPS input, which will import the text file and read information beginning at specific line numbers.

In order to carry out the MD lattice parameter study in this work, it was necessary to build many large structures (on the order of millions of atoms per structure) of varying vacancy and antisite defect concentrations for each NiAl composition of interest before importing the relevant text files and relaxing the structures in LAMMPS. In order to do this in a time-effective manner, the lab cluster was used to run MATLAB and generate the structures. A shell script (copied below) was created to make multiple structures without having to modify the MATLAB script (page 250) for each individual structure.

Many thanks go to Dr. Penghui Cao for his extensive help in improving the efficiency and abilities of the structure-building script, the construction of the shell scripts, and tutelage in using LAMMPS. His expertise and generosity with his time was invaluable during the molecular dynamics phase of this research.

The annotated version of the MATLAB script used to generate the text files LAMMPS uses to build the test structures can be found on page 250.


```

Generate NiAl BCC in .lmp @ 08-Feb-2016 16:54:08

1961      atoms
2      atom types
0.000000 25.920000      xlo xhi
0.000000 25.920000      ylo yhi
0.000000 25.920000      zlo zhi

Masses

1 58.693400
2 26.981539

Atoms

1 1 0.000000 0.000000 0.000000
2 2 1.440000 1.440000 1.440000
3 1 0.000000 0.000000 9.000000
4 2 1.440000 1.440000 10.440000
5 1 0.000000 0.000000 18.000000
6 2 1.440000 1.440000 19.440000
7 1 0.000000 0.000000 27.000000

```

Figure 8.19: Example LAMMPS output code generated by the MATLAB script for a relatively small structure of about 2000 atoms showing the header script and the x-y-z specifications for the first seven atoms.

8.6.3 Shell script for building NiAl test structures

A useful shell script for building many test structures at once can be found on 250. This prevents having to manually change the parameters of interest in the structure building tool before running it for each individual structure, which was necessary here since the test matrix was so large (373 structures). Typically, this script was modified to create the test structure in smaller groups - e.g. all test structures of composition 45 to 46 at% Ni, all test structures of composition 47-48 at% Ni, and so on - and then each group was submitted to the cluster, so that the MATLAB code could be run on multiple nodes and proceed through building the full test matrix more quickly.

Input file for Ni-Al alloy equilibration

----- INITIALIZATION -----

units metal
dimension 3
boundary p p p
atom_style atomic
variable latparam equal 3.5

----- ATOM DEFINITION -----

read_data NiAl_50.data
group ni type 1
group al type 2

----- FORCE FIELDS -----

pair_style eam/alloy
pair_coeff * * NiAl02.eam.alloy Ni Al

----- SETTINGS -----

thermo 100
thermo_style custom step lx ly lz press pxx pyy pzz pe temp

----- Equilibration -----

reset_timestep 0
timestep 0.001

velocity all create 300.0 4928459 mom yes rot yes dist gaussian
fix 1 all npt temp 300 300 0.1 iso 0.0 0.0 2

dump 1 all xyz 100 relaxed_config.xyz

run 100000

clear all	
clc	
La = 1; %lattice length,	Lattice length. Usually set to La=2.88
x_ni = 0; y_ni = 0; z_ni = 0; %unit cell, Ni atom position	Define the position of the first Ni atom and the first Al atom. The unit cell is a cube with side length = La
x_al = 0.5*La; y_al = 0.5*La; z_al = 0.5*La; %unit cell, Al atom position	
Percent_Ni=45;%	Define the percentage of Ni.
f=0.1; %	Define f = N_ASD/N_vacancy.
	Note: f must be <1 or >1 depending on Ni %, or the numbers will come out negative
	Ni<50: vacancies dominate (f<1)
	Ni>50: ASD dominate (f>1)
NumLx = 4;	
NumLy=NumLx;	For lattice parameter studies, Lx=Ly=Lz. The total # of atoms per sublattice per side is NumLx+1.
NumLz=NumLx;	
total_atoms = 0;	Initialize this value (will be used as a counter in the loop)
entries=2*(NumLx+1)*(NumLy+1)*(NumLz+1);	Initialize M, the matrix that contains the info to build the structure for LAMMPS. Each atom has five entries in its row: ID #, type # (Ni or Al), x, y, z.
M=zeros(entries,5);	
count=0;	Initialize the count
for i = 0:NumLx	
for j = 0:NumLy	
for k = 0:NumLz	Want to assign a row to M for each atom, so iterate over each direction. 0 to NumLx = total # of atoms per side = NumLx+1, since indexing from 0.
%atom type 1	
M(total_atoms+1,1) = total_atoms+1;	
M(total_atoms+1,2) = 1;	The first atom will have ID # 1 (first entry in its row). This value is increased by one with each pass through the loop. The nickel atoms all have Type ID = 1. The initializing x/y/z position and the lattice parameter * the index number are used to assign a unique position to each atom in the structure.
M(total_atoms+1,3) = x_ni + La*i;	
M(total_atoms+1,4) = y_ni + La*j;	
M(total_atoms+1,5) = z_ni + La*k;	
%atom type 2	
M(total_atoms+2,1) = total_atoms+2;	
M(total_atoms+2,2) = 2;	
M(total_atoms+2,3) = x_al + La*i;	
M(total_atoms+2,4) = y_al + La*j;	The same is repeated for the aluminum atoms. They are assigned a type of 2, and they occupy the even numbered rows in the initial M matrix. Each pass through this loop assigns one Al atom for one Ni atom.
M(total_atoms+2,5) = z_al + La*k;	
total_atoms = total_atoms + 2;	Update the count (assigned a Ni and an Al, so advance by 2)
end	
end	
count=count+1;	Update the loop count (helpful to see loop progress if desired)
end	
[Ni_vacancies]=Calc_VplusASD(NumLx,NumLy,NumLz,Percent_Ni,f);	Calculate the # of vacancies, based on %Ni and f
Ni_ASD=f*Ni_vacancies;	Calculate the # of ASD, based on # of vacancies and f
if Ni_ASD<1 && Ni_ASD>0	
Ni_ASD=1;	Can't have fractions of ASD or vacancies. In the ideal cases, ASD or vacancy # = 0. We round up to one when between 0 and 1 ASD or vacancy is calculated. If the # of ASD or V is > 1, can just round to nearest integer value.
else	
Ni_ASD=round(Ni_ASD);	
end	
Ni_ASD=round(Ni_ASD);	
if Ni_vacancies <1 && Ni_vacancies>0	
Ni_vacancies=1;	
else	
Ni_vacancies=round(Ni_vacancies);	
end	
Remove_Ni_atoms=Ni_vacancies;	Variable reassignment for consistency with earlier versions
total=2*(NumLx+1)*(NumLy+1)*(NumLz+1);	Total # of lattice sites in the structure

number_of_Ni=total/2-Ni_vacancies+Ni_ASD; number_of_Al=total/2-Ni_ASD; Conc_Ni=number_of_Ni/(number_of_Ni+number_of_Al)	This is included as a check to make sure the function returned the correct # of vacancies + ASD to get the desired Ni%.
delete_vector=zeros(Remove_Ni_atoms,1); asd_vector=zeros(Ni_ASD,1);	Initialize the vectors that will be used to remove Ni atoms, or swap Al atoms to Ni ASD.
for i=1:Remove_Ni_atoms a=randi(total); while mod(a,2)==0 a=randi(total); end delete_vector(i,1)=a; end	a will be used to index to a row in M and remove that atom. To ensure that the selected row corresponds to a nickel atom, check that it is an odd value. If it is even, mod(a,2)==0, and the value is reselected. If the value is odd, it is assigned to an entry in the delete_vector for vacancy creation. It is possible that multiple entries in the delete vector have the same value.
check=length(unique(delete_vector));	
b=length(delete_vector)-check;	b will return the number of entries that must be reassigned.
b=zeros(b,1);	initilaize a vector for "b"
c=length(b);	
d=length(delete_vector);	These will be used in the next loop
delete_vector=[delete_vector; b];	Concatenate [b] to the delete vector
for i=1:c	Only want to iterate through the length of b, since those are the only values that we need to reassign
while length(unique(delete_vector))<Remove_Ni_atoms;	Want a number of unique, odd values that = the number of vacancies that must be created
a=randi(total); count=i;	Follow the same procedure to assign new values of a
d=d+count; while mod(a,2)==0 a=randi(total); end	Index into the end of the delete_vector, where we added the zeros. The number of zeros = the number of non-unique entries from the first loop
delete_vector(d,1)=a; end	assign the new value to the delete_vector (in the concatenated "b" vector of zeros)
end	This repeats assigning values to the concatenated end of the delete vector until the overall delete_vector has the correct number of unique values
delete_vector=unique(delete_vector);	Can now just simplify the delete_vector and delete duplicates
delete_vector=sort(delete_vector);	Sort from low to high
n=size(delete_vector);	n = the number of vacancies that need to be created
count=i; b=delete_vector(count,1);	Index into rows M (row # = a value in the delete_vector) ans make the whole row 0s.
if b == 0 b = 1; end	Prevent tryign to index into row 0 (dne)
M(b,:)=0; end	
for i=1:Ni_ASD a=randi(total); while mod(a,2) ~= 0 a=randi(total); end asd_vector(i,1)=a; end	This is very similar to the procedure above for vacancy creation. However, Ni ASD are only created on the Al sublattice, so now we want to make sure we are only getting values corresponding to even rows in M. Note that we do this here, BEFORE deleting the zeroed rows, so that we can index into M easily (easier with the intact M matrix, can just go to the even rows)
check=length(unique(asd_vector)); b=length(asd_vector)-check; b=zeros(b,1); c=length(b); d=length(asd_vector); asd_vector=[asd_vector; b];	Setting up the asd_vector creation commands in the same way as we did for the delete_vector
for i=1:c	Finish creating the asd_vector. As with the delete_vector, we need to

<pre> while length(unique(asd_vector))<Ni_ASD; a=randi(total); count=i; d=d+count; while mod(a,2)~=0 a=randi(total); end asd_vector(d,1)=a; end end </pre>	<p>have N unique entries where N = the number of anti-site defects that are to be created. These entries should only correspond to rows in M for Al atoms (so even rows).</p>
<pre> asd_vector=unique(asd_vector); asd_vector=sort(asd_vector); zerocheck_asd=length(asd_vector)-Ni_ASD; n=size(asd_vector); for i=1:n count=i; if b < 1 %prevent indexing into zeroth row b=1; end b=asd_vector(count,1); M(b,2)=1; end </pre>	<p>Clean up the asd_vector, can do another check that it's the right number of unique entries</p> <p>Now, as before when we indexed to the determined Ni rows and zeroed them out, want to index to the determined Al row and change the second entry (the type identifier) to 1 (change 2=Al to 1=Ni to create the ASD).</p>
<pre> [indexRow,indexColumn]=find(M(:,2)==0); M(indexRow,:)=[]; original=(NumLx+1)*(NumLy+1)*(NumLz+1)*2; correct=original-Remove_Ni_atoms; test=length(M); zerocheck_vacancies=correct-test; %GENERATE THE 3D PLOT (commented for large files) %scatter3(M(:,3),M(:,4),M(:,5),100,M(:,2),'filled') </pre>	<p>All rows have a 1 or a 2 in their second column, unless they were changed to vacancies (whole row zeroed). [] deletes the row.</p> <p>Total number of atoms How many atoms should be left in the structure Number of atoms left = length of M Check that this value is 0</p> <p>Generate the scatter plot that can be rotated to examine the structure For large structures, this will take too much computational time, so it's best to comment it out</p>
<pre> xlo = -0.5*La*(NumLx+1); xhi = 0.5*La*(NumLx+1); </pre>	<p>Define the boundaries of the simulation box</p>
<pre> ylo = -0.5*La*(NumLy+1); yhi = 0.5*La*(NumLy+1); </pre>	<p>Centered at 0, -/+ direction = 1/2 * number of cells * La</p>
<pre> zlo = -0.5*La*(NumLz+1); zhi = 0.5*La*(NumLz+1); % write the lammps file </pre>	<p>Imports into LAMMPS, which uses it to build the structure</p>
<pre> fidout = fopen('Ni_#_f_#', 'w'); new_comment = ['Generate NiAL BCC in .Imp @ ',datestr(now)]; fprintf(fidout,'%s\n\n',new_comment); fprintf(fidout,'%d \t %s\n',length(M),'atoms'); fprintf(fidout,'%d \t %s\n',Conc_Ni,'percentage Ni'); fprintf(fidout,'%d \t %s\n',2,'atom types'); </pre>	<p>('XXXXX') choose file name</p> <p>Describe and date the file</p> <p>Record number of atoms Record information about the %Ni Record number of atom types</p>
<pre> fprintf(fidout,'%f %f \t%s\n',xlo,xhi,'xlo xhi'); fprintf(fidout,'%f %f \t%s\n',ylo,yhi,'ylo yhi'); fprintf(fidout,'%f %f \t%s\n',zlo,zhi,'zlo zhi'); fprintf(fidout,'%s\n','Masses'); </pre>	<p>Record dimensions of simulation box</p> <p>Record masses of the atom types</p>
<pre> fprintf(fidout,'%d %f \n',1,58.6934); %ni mass 58.6934u fprintf(fidout,'%d %f \n',2,26.981539); %al mass 26.981539 fprintf(fidout,'%s\n','Atoms'); % the data </pre>	<p>Print the M matrix (ID, type, x, y, z for each atom in structure)</p>
<pre> for i = 1:length(M) fprintf(fidout,'%d \t %d \t %f \t %f \t %f \n', i,M(i,2),M(i,3),M(i,4),M(i,5)); end fclose(fidout); </pre>	

Shell script for running multiple iterations of the structure-building code on the lab cluster

```
1  #!/bin/bash
2
3  NiSteps=$(awk 'BEGIN{ for (i=59;i <59.1;i +=0.2) print i }')
4  ASDsteps=$(awk 'BEGIN{ for (j=100;j <101;j +=80) print j }')
5
6  for Ni in $NiSteps
7  do
8
9      for ASD in $ASDsteps
10     do
11
12         sed "s/<NI-PERCENT>/$Ni/" Generate_VandASD_structures > temp1.txt
13         k=$(bc <<< "scale=0; $ASD/0.01")
14         ## echo $k
15         j=$(bc <<< "scale=0; $Ni/0.1")
16         ## echo $j
17         sed "s/<ASD-FRACTION>/$ASD/" temp1.txt > "Generate_NiAl_Ni_"$j"_ASD_"$k"_Sept.m"
18
19         rm temp1.txt
20
21     done
22 done
23
24 NiSteps=$(awk 'BEGIN{ for (i=59;i =59.1;i +=0.2) print i }')
25 ASDsteps=$(awk 'BEGIN{ for (j=100;j <101;j +=80) print j }')
26
27 for Ni in $NiSteps
28 do
29
30     for ASD in $ASDsteps
31     do
32
33         k=$(bc <<< "scale=0; $ASD/0.01")
34         j=$(bc <<< "scale=0; $Ni/0.1")
35         matlab -nodisplay -nosplash -nodesktop -r "run('/home/seferry/NiAl_MD/
36             Summer2016_VandASD/Generate_NiAl_Ni_"$j"_ASD_"$k"_Sept.m');"
37         ## matlab -nodisplay -nosplash -nodesktop -r "quit"
38     done
39 done
```

8.7 MATLAB scripts for analyzing TGS data (SAW frequency, thermal diffusivity, and acoustic damping)

The following suite of MATLAB scripts is used to analyze TGS data output using our experimental setup. These scripts are used to obtain SAW frequency (which is then used to calculate SAW speed), thermal diffusivity, and acoustic damping. (At the time of writing, the acoustic damping script is not considered to be accurate, but is included here as it is part of the group's overall work toward a complete computational package for fully and accurately analyzing TGS data.) Credit for the bulk of these scripts goes to Cody Dennett. It is also important to acknowledge Jeffrey Eliason of the Nelson Group at MIT, who provided Cody and the group with the initial set of scripts for pinpointing SAW frequency from a TGS setup.

In this work, I typically called these functions using a script that is modified for the particular set of files I'm working with. First, to obtain the frequency associated with a particular trace, I usually use a script like this to call the `param_extract_time` function.

Example: Get SAW frequency

```
1 gs=4.0951;
2 temp=zeros(11,1);
3
4
5 temp(1,1)=param_extract_time(1,'EBR_Reference-2018-04-06-04.10um-spot01-NEG-3.txt','
   EBR_Reference-2018-04-06-04.10um-spot01-POS-3.txt',gs,0);
6 temp(2,1)=param_extract_time(1,'EBR_Reference-2018-04-06-04.10um-spot02-NEG-3.txt','
   EBR_Reference-2018-04-06-04.10um-spot02-POS-3.txt',gs,0);
7 temp(3,1)=param_extract_time(1,'EBR_Reference-2018-04-06-04.10um-spot03-NEG-3.txt','
   EBR_Reference-2018-04-06-04.10um-spot03-POS-3.txt',gs,0);
8 temp(4,1)=param_extract_time(1,'EBR_Reference-2018-04-06-04.10um-spot04-NEG-3.txt','
   EBR_Reference-2018-04-06-04.10um-spot04-POS-3.txt',gs,0);
9 temp(5,1)=param_extract_time(1,'EBR_Reference-2018-04-06-04.10um-spot05-NEG-3.txt','
   EBR_Reference-2018-04-06-04.10um-spot05-POS-3.txt',gs,0);
10 temp(6,1)=param_extract_time(1,'EBR_Reference-2018-04-06-04.10um-spot06-NEG-3.txt','
   EBR_Reference-2018-04-06-04.10um-spot06-POS-3.txt',gs,0);
11 temp(7,1)=param_extract_time(1,'EBR_Reference-2018-04-06-04.10um-spot07-NEG-3.txt','
   EBR_Reference-2018-04-06-04.10um-spot07-POS-3.txt',gs,0);
12 temp(8,1)=param_extract_time(1,'EBR_Reference-2018-04-06-04.10um-spot08-NEG-3.txt','
   EBR_Reference-2018-04-06-04.10um-spot08-POS-3.txt',gs,0);
13 temp(9,1)=param_extract_time(1,'EBR_Reference-2018-04-06-04.10um-spot09-NEG-3.txt','
   EBR_Reference-2018-04-06-04.10um-spot09-POS-3.txt',gs,0);
14 temp(10,1)=param_extract_time(1,'EBR_Reference-2018-04-06-04.10um-spot10-NEG-3.txt','
   EBR_Reference-2018-04-06-04.10um-spot10-POS-3.txt',gs,0);
```

```
15 temp(11,1)=param_extract_time(1,'EBR_Reference-2018-04-06-04.10um-spot11-NEG-3.txt','  
EBR_Reference-2018-04-06-04.10um-spot11-POS-3.txt',gs,0);
```

In this example, there were 11 spots to analyze. In the frequency scripts, it is necessary to first click to the left and the right of the peak that appears in the plot of the fast Fourier transform of the data - the script will find the maximum point in the range defined by that action. All of the data from a given sample are kept in one folder, along with this “call” script and its related functions. The user must run the call script and define the bounds on each plot (one for each point). The calculated frequency is then stored in a vector (called “temp” here).

(Note that in order to obtain SAW speed from the SAW frequency, it is necessary to multiply by the calibrated grating spacing corresponding to that day of measurements, in accordance with $v_{SAW} = \lambda_{GS} v_{SAW}$.)

In the data used for this thesis, there were always three traces obtained at a given spot. In the event of multiple traces, it is necessary to either run the script again, with updated file names (our system is set up to automatically append -1, -2, or -3 to each file name, depending on which trace it corresponds to in that spot), or modify the script to run all three traces and store them in an appropriate vector. I found it easier to do each trace separately, since it is easy to make a mistake when clicking to define the search bounds on the plot, and then the entire script exits.

It isn't necessary to click on the bounds of a plot in order to extract the thermal and acoustic data, so it is easier to create a call script for that data runs through the calculations for all three traces if desired. That call script uses the following format to call the necessary functions, where **a** is the thermal diffusivity, **b** is the thermal error, and **c** is the acoustic damping parameter:

[a, b, c] =thermal_phase('filename-NEG-1.txt', 'filename-POS-1.txt', gs, 2)

a, **b**, and **c** can then be stored in the appropriate vectors. Care should be taken when assigning values to a vector position to ensure that the user is confident in which vector entry goes with which measurement spot on a sample. The **thermal_phase** script and its associated scripts can be found in the subsequent pages.

8.7.1 thermal_phase.m

```
1 function [diffusivity ,diffusivity_err ,tau] = thermal_phase( pos_file , neg_file , grat ,
    start_phase , end_time)
2 % Function to determine thermal diffusivity from phase grating TGS data
3 % Data is saved in two files , positive (with one heterodyne phase) and negative (with
    another), must provide both files
4 % pos_file: positive phase TGS data file
5 % neg_file: negative phase TGS data file
6 % grat: calibrated grating spacing in um
7 % start_phase: provide integer between 1 and 4 to pick the null-point start from which
    fit will begin
8 % end_time: a shortened fit end time if you do not want to fit the whole profile.
9 %           If argument not given, will be set to default for 200ns data collection
    window
10
11 %%%Write this to include a sine variation in the fit by default, but to
12 %%%start the fits from a fixed null point, not time, relative to
13 %%%the initial SAW maximum.
14
15 %Settings for various plotting and output options to be set by boolean arguments
16 find_max=0;
17 plotty=0;
18 plot_trace=0;
19 plot_final=0;
20 print_final_fit=0;
21 two_detectors=1;
22 q=2*pi/(grat*10^(-6));
23 timestep=5e-11; %Set by scope used for data collection
24 no_pre_calc=0;
25
26 derivative=0;
27
28 %How far from guessed values for diffusivity and beta do you vary in the
29 %end, d for diffusivity b for beta. Diffusivity is most important
30 percent_range_d=0.45;
31 percent_range_b=2.1;
32
33 if nargin < 5
34     %     end_time=10e-7; %for 50ns base on scope
35     end_time=2e-7; %for 20ns base on scope
36 end
37
38 %Difference in file write format based on newer or older acquisition. hdr_len should be
    16 for the Ge dataset
```

```

39  if two_detectors
40      hdr_len=16;
41  else
42      hdr_len=15;
43  end
44
45  %Generate filtered power spectrum using make_fft_embed_time and find the peak frequency
      from that profile using fit_spectra_peaks
46  [fft,~]=make_fft_embed_time(1, pos_file , neg_file , 'therm' ,5, grat ,0);
47  [peak_freq ,~,~]=fit_spectra_peaks(fft ,0);
48
49  %read in data files for this procedure
50  pos=dlmread(pos_file ,'',hdr_len ,0);
51  neg=dlmread(neg_file ,'',hdr_len ,0);
52
53  %sometimes written data is off by one time step at the end, chop that off if they do not
      match
54  if length(pos(:,1))>length(neg(:,1))
55      pos=pos(1:length(neg(:,1)),:);
56  elseif length(neg(:,1))>length(pos(:,1))
57      neg=neg(1:length(pos(:,1)),:);
58  end
59
60  %normalize each set of data to the zero level before the pump impulse
61  pos(:,2)=pos(:,2)-mean(pos(1:50,2));
62  neg(:,2)=neg(:,2)-mean(neg(1:50,2));
63
64  time_index=186; %From peak in amp grating data
65  time_naught=neg(time_index ,1);
66
67  end_index=floor(end_time / tstep);
68
69  %re-normalize data to end signal decayed state if grating decays entirely during
      collection window, if not, do not re-normalize
70  if grat <8
71      base_index=floor(length(pos(:,2))/5);
72      long_base=mean((pos(end-base_index :end ,2)-neg(end-base_index :end ,2)));
73  else
74      long_base=0;
75  end
76
77  if plot_trace
78      figure()
79      plot(neg(:,1)*10^9,(pos(:,2)-neg(:,2)-long_base)*10^3,'-', 'Color',[0 0 0.75], '
          LineWidth',1.25)

```

```

80     hold on
81     plot([neg(1,1) neg(end,1)]*10^9,[0 0], 'k--', 'LineWidth',1.5)
82     xlim([0 end_time*10^9])
83     set(gca,...
84         'FontUnits','points',...
85         'FontWeight','normal',...
86         'FontSize',16,...
87         'FontName','Helvetica',...
88         'LineWidth',1.25)
89     ylabel({'Amplitude [mV]'},...
90         'FontUnits','points',...
91         'FontSize',20,...
92         'FontName','Helvetica')
93     xlabel({'Time [ns]'},...
94         'FontUnits','points',...
95         'FontSize',20,...
96         'FontName','Helvetica')
97 end
98
99 fixed_short=[pos(time_index:end_index,1)-time_naught pos(time_index:end_index,2)-neg(
    time_index:end_index,2)-long_base];
100
101 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
102 % Plot block for fixed null point figure (Figure 4)
103 % null_pts=4;
104 % null_pt_vec_x=zeros(1,null_pts);
105 % null_pt_vec_y=zeros(1,null_pts);
106 % for kk=1:length(null_pt_vec_x)
107 %     temp=find_start_phase(fixed_short(:,1),fixed_short(:,2),kk,grat);
108 %     temp_ind=floor(temp/tstep)+1;
109 %     null_pt_vec_x(kk)=fixed_short(temp_ind,1);
110 %     null_pt_vec_y(kk)=fixed_short(temp_ind,2);
111 % end
112 %
113 % null_pt_vec_x=null_pt_vec_x*10^9; %scale for units
114 % null_pt_vec_y=null_pt_vec_y*10^3;
115 %
116 % display(null_pt_vec_x)
117 %
118 % figure()
119 % plot(fixed_short(:,1)*10^9,fixed_short(:,2)*10^3,'-','Color',[0 0 0.75],'LineWidth
    ',1.25)
120 % hold on
121 % plot(null_pt_vec_x,null_pt_vec_y,'rd','MarkerSize',8,'MarkerFaceColor','r')
122 % hold on

```

```

123 % xlim([0 5])
124 % ylim([0 65])
125 % set(gca,...
126 %     'FontUnits','points',...
127 %     'FontWeight','normal',...
128 %     'FontSize',16,...
129 %     'FontName','Helvetica',...
130 %     'LineWidth',1.25)
131 % ylabel({'Amplitude [mV]'},...
132 %     'FontUnits','points',...
133 %     'FontSize',20,...
134 %     'FontName','Helvetica')
135 % xlabel({'Time [ns]'},...
136 %     'FontUnits','points',...
137 %     'FontSize',20,...
138 %     'FontName','Helvetica')
139 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
140
141 if derivative
142     der_len=length(fixed_short(:,1))-1;
143     fixed_derivative=zeros(1,der_len);
144     for jj=1:der_len
145         fixed_derivative(jj)=(fixed_short(jj+1,2)-fixed_short(jj,2))/tstep;
146     end
147     figure()
148     plot(fixed_short(1:der_len,1),fixed_derivative,'k-')
149     title('This is the derivative of fixed short')
150 end
151
152 %if you don't want to automatically find the peak t_0 of the profile, setting find_max to
153 %true above will allow
154 %you to select a region on the plot within with to search. Useful if there are initial
155 %transients.
156
157 if find_max || plotty
158     figure()
159     plot(fixed_short(:,1),fixed_short(:,2),'k-')
160     xlim([0 1.5*10^-7])
161     title('this is fixed short');
162     if find_max
163         hold on
164         [x_cord,~]=ginput(2);
165         neg_x_cord=x_cord(1);
166         pos_x_cord=x_cord(2);
167         pos_x_ind=floor(pos_x_cord/tstep);
168         neg_x_ind=floor(neg_x_cord/tstep);

```

```

166         [~,max_index]=max(fixed_short(neg_x_ind:pos_x_ind,2));
167         time_max=fixed_short(max_index+neg_x_ind,1);
168         close(gcf)
169     end
170 end
171
172 %Otherwise, find t_0 from the profile directly
173 if ~find_max
174     time_offset_index=20; %was 20 before
175     [~,max_index]=max(fixed_short(time_offset_index:end,2));
176     max_index=max_index+time_offset_index-1;
177     time_max=fixed_short(max_index,1);
178
179 start_time_phase=find_start_phase(fixed_short(:,1),fixed_short(:,2),start_phase,grat);
180 start_index_master=floor(start_time_phase/tstep)+2;
181 start_time_master=fixed_short(start_index_master,1);
182
183 %%%%%%%%%%%
184 % %Block to write the start time for each measurement based on the fixed null
185 % cd('start_times')
186 % dlmwrite(streat('start_time_null_2_',num2str(grat),'.txt'),start_time_master,'-append')
187 % cd('..')
188 %%%%%%%%%%%
189
190 %Fitting parameters for initial naive fit
191 LB=[0 0];
192 UB=[1 10^-4];
193 ST=[.05 5*10^-6];
194
195 OPS=fitoptions('Method','NonLinearLeastSquares','Lower',LB,'Upper',UB,'Start',ST);
196 TYPE=fittype('A.*erfc(q*sqrt(k*(x+time_max)))','options',OPS,'problem',{ 'q','time_max'},'
    coefficients',{ 'A','k'});
197 [f0,gof]=fit(fixed_short(:,1),fixed_short(:,2),TYPE,'problem',{q,time_max});
198
199 diffusivity=f0.k;
200 error=confint(f0,0.95);
201 diffusivity_err=[diffusivity-error(1,2) error(2,2)-diffusivity];
202
203 if plotty
204     figure()
205     plot(fixed_short(:,1),fixed_short(:,2),fixed_short(:,1),f0(fixed_short(:,1)))
206     hold on
207     title('First naive fit')
208 end
209

```

```

210 %We'll call the parameter beta the ratio of the amplitudes of the
211 %displacement versus temperature grating. beta should be a small number.
212
213 for jj=1:10
214
215     beta=q*sqrt(diffusivity/pi)*(q^2*diffusivity+1/(2*time_max))^(−1);
216
217     start_time=start_time_master;
218     start_index=start_index_master;
219
220     %Conduct initial parameter estimation without using an sin(x) contribution to the fit
221
222     LB1=[0 0];
223     UB1=[1 10^−4];
224     ST1=[.05 10^−5];
225
226     OPS1=fitoptions('Method','NonLinearLeastSquares','Lower',LB1,'Upper',UB1,'Start',ST1)
227     ;
228     TYPE1=fittype('A.*(erfc(q*sqrt(k*(x+start_time)))-beta*exp(−q^2*k*(x+start_time))./
229         sqrt((x+start_time)))','options',OPS1,'problem',{ 'q','beta','start_time'},'
230         coefficients',{ 'A','k'});
231     [f1,gof]=fit(fixed_short(start_index:end,1),fixed_short(start_index:end,2),TYPE1,'
232         problem',{q,beta,start_time});
233
234     diffusivity=f1.k;
235     error=confint(f1,0.95);
236     diffusivity_err=[diffusivity−error(1,2) error(2,2)−diffusivity];
237
238     if plotty
239         figure()
240         plot(fixed_short(:,1),fixed_short(:,2),fixed_short_old(:,1),f1(fixed_short_old
241             (:,1)))
242         hold on
243         title(strcat('Fit number ',num2str(jj+1),' − fixed beta'))
244     end
245 end
246
247 %If you've elected not to pre-compute, provide hard initial guesses for diffusivity and
248 %beta based on
249 %the bulk value of Ge diffusivity. Set ranges for final fit.
250 if no_pre_calc
251     diffusivity=0.3636*10^−4;
252     beta=2e−5;
253     low_bound=[1e−5 0];

```

```

249     up_bound=[1e-3 1e-4];
250 else
251     low_bound=[ diffusivity*(1-percent_range_d) beta*(1-percent_range_b) ];
252     up_bound=[ diffusivity*(1+percent_range_d) beta*(1+percent_range_b) ];
253     if percent_range_d > 1
254         low_bound(1)=0;
255     end
256     if percent_range_b > 1
257         low_bound(2)=0;
258     end
259 end
260
261
262 start_time2=start_time_master;
263 start_index2=start_index_master;
264
265 %After testing, initial guesses for acoustic damping are more appropriate for different
    grating spacings
266 LB2=[0 low_bound(1) low_bound(2) 0 -2*pi 0 -5e-3];
267 if grat < 4
268     UB2=[1 up_bound(1) up_bound(2) 10 2*pi 2e-7 5e-3];
269     ST2=[.05 diffusivity beta 0.05 0 15e-8 0];
270 else
271     UB2=[1 up_bound(1) up_bound(2) 10 2*pi 1e-7 5e-3];
272     ST2=[.05 diffusivity beta 0.05 0 1e-8 0];
273 end
274
275 OPS2=fitoptions('Method','NonLinearLeastSquares','Lower',LB2,'Upper',UB2,'Start',ST2);
276 TYPE2=fittype('A.*(erfc(q*sqrt(k*(x+start_time)))-beta*exp(-q^2*k*(x+start_time))./sqrt((
    x+start_time)))+B.*sin(2*pi*(peak_freq)*(x+start_time)+p)*exp(-(x+start_time)/t)+D',
    'options',OPS2,'problem',{ 'q','start_time','peak_freq'},'coefficients',{'A','k','beta'
    ,'B','p','t','D'});
277 [f2,gof]=fit(fixed_short(start_index2:end,1),fixed_short(start_index2:end,2),TYPE2,'
    problem',{q,start_time2,peak_freq});
278
279 if print_final_fit
280     display(f2)
281 end
282
283 fit_err=fixed_short(:,2)-f2(fixed_short(:,1));
284 sum_squared_err=sum(fit_err.^2);
285 deg_free=length(fixed_short(:,1))-3;
286 rms_err=sqrt(sum_squared_err/deg_free);
287
288 diffusivity=f2.k;

```

```

289 tau=f2.t;
290 error=confint(f2,0.95);
291 diffusivity_err=[diffusivity -error(1,2) error(2,2)-diffusivity];
292
293 %first checks that diffusivity has not pegged to fit bounds, second checks that beta
294 %has not pegged. If either do, it is a bad fit
295 if isnan(diffusivity_err(1))
296     display(strcat('Bad fit for: ',pos_file,'~re: alpha'))
297 elseif isnan(error(1,3))
298     display(strcat('Bad fit for: ',pos_file,'~re: beta'))
299 end
300
301 if plot_final
302     %Plotting factor for generation of traces in Figure 6
303     amp_factor=1;
304     %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
305     %Block to reconstruct the best-fit model without the sinusoidal
306     %contribution, for comparison
307     f_remove_sine=cfit(TYPE2,f2.A,f2.k,f2.beta,f2.B,f2.p,0,f2.D,q,start_time2,peak_freq);
308     %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
309     figure()
310     % plot((fixed_short(:,1))*10^9,(fixed_short(:,2))*10^3,'k-', 'LineWidth',1.5)
311     plot((neg(:,1)-time_naught)*10^9,(pos(:,2)-neg(:,2)-long_base)*10^3/amp_factor,'k-', '
        LineWidth',1.35)
312     hold on
313     %plot vertical line at start time
314     % plot([fixed_short(start_index_master,1) fixed_short(start_index_master,1)]*10^9,
        ylim,'b--')
315     % hold on
316     plot(fixed_short(:,1)*10^9,(f2(fixed_short(:,1)))*10^3/amp_factor,'r--', 'LineWidth'
        ,1.45)
317     hold on
318     plot(fixed_short(:,1)*10^9,(f_remove_sine(fixed_short(:,1)))*10^3/amp_factor,'-', '
        Color',[0 0 0.75], 'LineWidth',1.45)
319     hold on
320     xlim([-5 end_time*10^9])
321     set(gcf,'Position',[0 0 400 300])
322     hold on
323     set(gca,...
324         'FontUnits','points',...
325         'FontWeight','normal',...
326         'FontSize',16,...
327         'FontName','Helvetica',...
328         'LineWidth',1.25)
329     % 'YTickLabel','')

```



```
330 ylabel({'Amplitude [a.u.]'},...
331         'FontUnits','points',...
332         'FontSize',20,...
333         'FontName','Helvetica')
334 xlabel({'Time [ns]'},...
335         'FontUnits','points',...
336         'FontSize',20,...
337         'FontName','Helvetica')
338 end
339
340 end
```

8.7.2 find_start_phase.m

```
1 function [diffusivity ,diffusivity_err ,tau] = thermal_phase(pos_file ,neg_file , grat ,
    start_phase ,end_time)
2 % Function to determine thermal diffusivity from phase grating TGS data
3 % Data is saved in two files , positive (with one heterodyne phase) and negative (with
    another), must provide both files
4 % pos_file: positive phase TGS data file
5 % neg_file: negative phase TGS data file
6 % grat: calibrated grating spacing in um
7 % start_phase: provide integer between 1 and 4 to pick the null-point start from which
    fit will begin
8 % end_time: a shortened fit end time if you do not want to fit the whole profile.
9 % If argument not given, will be set to default for 200ns data collection
    window
10
11 %%%Write this to include a sine variation in the fit by default, but to
12 %%%start the fits from a fixed null point, not time, relative to
13 %%%the initial SAW maximum.
14
15 %Settings for various plotting and output options to be set by boolean arguments
16 find_max=0;
17 plotty=0;
18 plot_trace=0;
19 plot_final=0;
20 print_final_fit=0;
21 two_detectors=1;
22 q=2*pi/(grat*10^(-6));
23 timestep=5e-11; %Set by scope used for data collection
24 no_pre_calc=0;
25
26 derivative=0;
27
28 %How far from guessed values for diffusivity and beta do you vary in the
29 %end, d for diffusivity b for beta. Diffusivity is most important
30 percent_range_d=0.45;
31 percent_range_b=2.1;
32
33 if nargin <5
34     % end_time=10e-7; %for 50ns base on scope
35     end_time=2e-7; %for 20ns base on scope
36 end
37
38 %Difference in file write format based on newer or older acquisition. hdr_len should be
    16 for the Ge dataset
```

```

39  if two_detectors
40      hdr_len=16;
41  else
42      hdr_len=15;
43  end
44
45  %Generate filtered power spectrum using make_fft_embed_time and find the peak frequency
      from that profile using fit_spectra_peaks
46  [fft,~]=make_fft_embed_time(1, pos_file , neg_file , 'therm' ,5, grat ,0);
47  [peak_freq ,~,~]=fit_spectra_peaks(fft ,0);
48
49  %read in data files for this procedure
50  pos=dlmread(pos_file ,'',hdr_len ,0);
51  neg=dlmread(neg_file ,'',hdr_len ,0);
52
53  %sometimes written data is off by one time step at the end, chop that off if they do not
      match
54  if length(pos(:,1))>length(neg(:,1))
55      pos=pos(1:length(neg(:,1)),:);
56  elseif length(neg(:,1))>length(pos(:,1))
57      neg=neg(1:length(pos(:,1)),:);
58  end
59
60  %normalize each set of data to the zero level before the pump impulse
61  pos(:,2)=pos(:,2)-mean(pos(1:50,2));
62  neg(:,2)=neg(:,2)-mean(neg(1:50,2));
63
64  time_index=186; %From peak in amp grating data
65  time_naught=neg(time_index ,1);
66
67  end_index=floor(end_time / tstep);
68
69  %re-normalize data to end signal decayed state if grating decays entirely during
      collection window, if not, do not re-normalize
70  if grat <8
71      base_index=floor(length(pos(:,2))/5);
72      long_base=mean((pos(end-base_index :end ,2)-neg(end-base_index :end ,2)));
73  else
74      long_base=0;
75  end
76
77  if plot_trace
78      figure()
79      plot(neg(:,1)*10^9,(pos(:,2)-neg(:,2)-long_base)*10^3,'-', 'Color',[0 0 0.75], '
          LineWidth',1.25)

```

```

80     hold on
81     plot([neg(1,1) neg(end,1)]*10^9,[0 0], 'k--', 'LineWidth',1.5)
82     xlim([0 end_time*10^9])
83     set(gca,...
84         'FontUnits','points',...
85         'FontWeight','normal',...
86         'FontSize',16,...
87         'FontName','Helvetica',...
88         'LineWidth',1.25)
89     ylabel({'Amplitude [mV]'},...
90         'FontUnits','points',...
91         'FontSize',20,...
92         'FontName','Helvetica')
93     xlabel({'Time [ns]'},...
94         'FontUnits','points',...
95         'FontSize',20,...
96         'FontName','Helvetica')
97 end
98
99 fixed_short=[pos(time_index:end_index,1)-time_naught pos(time_index:end_index,2)-neg(
    time_index:end_index,2)-long_base];
100
101 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
102 % Plot block for fixed null point figure (Figure 4)
103 % null_pts=4;
104 % null_pt_vec_x=zeros(1,null_pts);
105 % null_pt_vec_y=zeros(1,null_pts);
106 % for kk=1:length(null_pt_vec_x)
107 %     temp=find_start_phase(fixed_short(:,1),fixed_short(:,2),kk,grat);
108 %     temp_ind=floor(temp/tstep)+1;
109 %     null_pt_vec_x(kk)=fixed_short(temp_ind,1);
110 %     null_pt_vec_y(kk)=fixed_short(temp_ind,2);
111 % end
112 %
113 % null_pt_vec_x=null_pt_vec_x*10^9; %scale for units
114 % null_pt_vec_y=null_pt_vec_y*10^3;
115 %
116 % display(null_pt_vec_x)
117 %
118 % figure()
119 % plot(fixed_short(:,1)*10^9,fixed_short(:,2)*10^3,'-','Color',[0 0 0.75],'LineWidth
    ',1.25)
120 % hold on
121 % plot(null_pt_vec_x,null_pt_vec_y,'rd','MarkerSize',8,'MarkerFaceColor','r')
122 % hold on

```

```

123 % xlim([0 5])
124 % ylim([0 65])
125 % set(gca,...
126 %     'FontUnits','points',...
127 %     'FontWeight','normal',...
128 %     'FontSize',16,...
129 %     'FontName','Helvetica',...
130 %     'LineWidth',1.25)
131 % ylabel({'Amplitude [mV]'},...
132 %     'FontUnits','points',...
133 %     'FontSize',20,...
134 %     'FontName','Helvetica')
135 % xlabel({'Time [ns]'},...
136 %     'FontUnits','points',...
137 %     'FontSize',20,...
138 %     'FontName','Helvetica')
139 %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
140
141 if derivative
142     der_len=length(fixed_short(:,1))-1;
143     fixed_derivative=zeros(1,der_len);
144     for jj=1:der_len
145         fixed_derivative(jj)=(fixed_short(jj+1,2)-fixed_short(jj,2))/tstep;
146     end
147     figure()
148     plot(fixed_short(1:der_len,1),fixed_derivative,'k-')
149     title('This is the derivative of fixed short')
150 end
151
152 %if you don't want to automatically find the peak t_0 of the profile, setting find_max to
153 %true above will allow
154 %you to select a region on the plot within with to search. Useful if there are initial
155 %transients.
156
157 if find_max || plotty
158     figure()
159     plot(fixed_short(:,1),fixed_short(:,2),'k-')
160     xlim([0 1.5*10^-7])
161     title('this is fixed short');
162     if find_max
163         hold on
164         [x_cord,~]=ginput(2);
165         neg_x_cord=x_cord(1);
166         pos_x_cord=x_cord(2);
167         pos_x_ind=floor(pos_x_cord/tstep);
168         neg_x_ind=floor(neg_x_cord/tstep);

```

```

166         [~,max_index]=max(fixed_short(neg_x_ind:pos_x_ind,2));
167         time_max=fixed_short(max_index+neg_x_ind,1);
168         close(gcf)
169     end
170 end
171
172 %Otherwise, find t_0 from the profile directly
173 if ~find_max
174     time_offset_index=20; %was 20 before
175     [~,max_index]=max(fixed_short(time_offset_index:end,2));
176     max_index=max_index+time_offset_index-1;
177     time_max=fixed_short(max_index,1);
178
179 start_time_phase=find_start_phase(fixed_short(:,1),fixed_short(:,2),start_phase,grat);
180 start_index_master=floor(start_time_phase/tstep)+2;
181 start_time_master=fixed_short(start_index_master,1);
182
183 %%%%%%%%%%%
184 % %Block to write the start time for each measurement based on the fixed null
185 % cd('start_times')
186 % dlmwrite(streat('start_time_null_2_',num2str(grat),'.txt'),start_time_master,'-append')
187 % cd('..')
188 %%%%%%%%%%%
189
190 %Fitting parameters for initial naive fit
191 LB=[0 0];
192 UB=[1 10^-4];
193 ST=[.05 5*10^-6];
194
195 OPS=fitoptions('Method','NonLinearLeastSquares','Lower',LB,'Upper',UB,'Start',ST);
196 TYPE=fittype('A.*erfc(q*sqrt(k*(x+time_max)))','options',OPS,'problem',{ 'q','time_max'},'
    coefficients',{ 'A','k'});
197 [f0,gof]=fit(fixed_short(:,1),fixed_short(:,2),TYPE,'problem',{q,time_max});
198
199 diffusivity=f0.k;
200 error=confint(f0,0.95);
201 diffusivity_err=[diffusivity-error(1,2) error(2,2)-diffusivity];
202
203 if plotty
204     figure()
205     plot(fixed_short(:,1),fixed_short(:,2),fixed_short(:,1),f0(fixed_short(:,1)))
206     hold on
207     title('First naive fit')
208 end
209

```

```

210 %We'll call the parameter beta the ratio of the amplitudes of the
211 %displacement versus temperature grating. beta should be a small number.
212
213 for jj=1:10
214
215     beta=q*sqrt(diffusivity/pi)*(q^2*diffusivity+1/(2*time_max))^(−1);
216
217     start_time=start_time_master;
218     start_index=start_index_master;
219
220     %Conduct initial parameter estimation without using an sin(x) contribution to the fit
221
222     LB1=[0 0];
223     UB1=[1 10^−4];
224     ST1=[.05 10^−5];
225
226     OPS1=fitoptions('Method','NonLinearLeastSquares','Lower',LB1,'Upper',UB1,'Start',ST1)
227     ;
228     TYPE1=fittype('A.*(erfc(q*sqrt(k*(x+start_time)))-beta*exp(−q^2*k*(x+start_time))./
229         sqrt((x+start_time)))','options',OPS1,'problem',{ 'q','beta','start_time'},'
230         coefficients',{ 'A','k'});
231     [f1,gof]=fit(fixed_short(start_index:end,1),fixed_short(start_index:end,2),TYPE1,'
232         problem',{q,beta,start_time});
233
234     diffusivity=f1.k;
235     error=confint(f1,0.95);
236     diffusivity_err=[diffusivity−error(1,2) error(2,2)−diffusivity];
237
238     if plotty
239         figure()
240         plot(fixed_short(:,1),fixed_short(:,2),fixed_short_old(:,1),f1(fixed_short_old
241             (:,1)))
242         hold on
243         title(strcat('Fit number ',num2str(jj+1),' − fixed beta'))
244     end
245 end
246
247 %If you've elected not to pre-compute, provide hard initial guesses for diffusivity and
248 %beta based on
249 %the bulk value of Ge diffusivity. Set ranges for final fit.
250 if no_pre_calc
251     diffusivity=0.3636*10^−4;
252     beta=2e−5;
253     low_bound=[1e−5 0];

```

```

249     up_bound=[1e-3 1e-4];
250 else
251     low_bound=[ diffusivity*(1-percent_range_d) beta*(1-percent_range_b) ];
252     up_bound=[ diffusivity*(1+percent_range_d) beta*(1+percent_range_b) ];
253     if percent_range_d > 1
254         low_bound(1)=0;
255     end
256     if percent_range_b > 1
257         low_bound(2)=0;
258     end
259 end
260
261
262 start_time2=start_time_master;
263 start_index2=start_index_master;
264
265 %After testing, initial guesses for acoustic damping are more appropriate for different
    grating spacings
266 LB2=[0 low_bound(1) low_bound(2) 0 -2*pi 0 -5e-3];
267 if grat < 4
268     UB2=[1 up_bound(1) up_bound(2) 10 2*pi 2e-7 5e-3];
269     ST2=[.05 diffusivity beta 0.05 0 15e-8 0];
270 else
271     UB2=[1 up_bound(1) up_bound(2) 10 2*pi 1e-7 5e-3];
272     ST2=[.05 diffusivity beta 0.05 0 1e-8 0];
273 end
274
275 OPS2=fitoptions('Method','NonLinearLeastSquares','Lower',LB2,'Upper',UB2,'Start',ST2);
276 TYPE2=fittype('A.*(erfc(q*sqrt(k*(x+start_time)))-beta*exp(-q^2*k*(x+start_time))./sqrt((
    x+start_time)))+B.*sin(2*pi*(peak_freq)*(x+start_time)+p)*exp(-(x+start_time)/t)+D',
    'options',OPS2,'problem',{ 'q','start_time','peak_freq'},'coefficients',{'A','k','beta'
    ,'B','p','t','D'});
277 [f2,gof]=fit(fixed_short(start_index2:end,1),fixed_short(start_index2:end,2),TYPE2,'
    problem',{q,start_time2,peak_freq});
278
279 if print_final_fit
280     display(f2)
281 end
282
283 fit_err=fixed_short(:,2)-f2(fixed_short(:,1));
284 sum_squared_err=sum(fit_err.^2);
285 deg_free=length(fixed_short(:,1))-3;
286 rms_err=sqrt(sum_squared_err/deg_free);
287
288 diffusivity=f2.k;

```



```

289 tau=f2.t;
290 error=confint(f2,0.95);
291 diffusivity_err=[diffusivity -error(1,2) error(2,2)-diffusivity];
292
293 %first checks that diffusivity has not pegged to fit bounds, second checks that beta
294 %has not pegged. If either do, it is a bad fit
295 if isnan(diffusivity_err(1))
296     display(strcat('Bad fit for: ',pos_file,'~re: alpha'))
297 elseif isnan(error(1,3))
298     display(strcat('Bad fit for: ',pos_file,'~re: beta'))
299 end
300
301 if plot_final
302     %Plotting factor for generation of traces in Figure 6
303     amp_factor=1;
304     %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
305     %Block to reconstruct the best-fit model without the sinusoidal
306     %contribution, for comparison
307     f_remove_sine=cfit(TYPE2,f2.A,f2.k,f2.beta,f2.B,f2.p,0,f2.D,q,start_time2,peak_freq);
308     %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
309     figure()
310     % plot((fixed_short(:,1))*10^9,(fixed_short(:,2))*10^3,'k-', 'LineWidth',1.5)
311     plot((neg(:,1)-time_naught)*10^9,(pos(:,2)-neg(:,2)-long_base)*10^3/amp_factor,'k-', '
        LineWidth',1.35)
312     hold on
313     %plot vertical line at start time
314     % plot([fixed_short(start_index_master,1) fixed_short(start_index_master,1)]*10^9,
        ylim,'b--')
315     % hold on
316     plot(fixed_short(:,1)*10^9,(f2(fixed_short(:,1)))*10^3/amp_factor,'r--', 'LineWidth'
        ,1.45)
317     hold on
318     plot(fixed_short(:,1)*10^9,(f_remove_sine(fixed_short(:,1)))*10^3/amp_factor,'-', '
        Color',[0 0 0.75], 'LineWidth',1.45)
319     hold on
320     xlim([-5 end_time*10^9])
321     set(gcf,'Position',[0 0 400 300])
322     hold on
323     set(gca,...
324         'FontUnits','points',...
325         'FontWeight','normal',...
326         'FontSize',16,...
327         'FontName','Helvetica',...
328         'LineWidth',1.25)
329     % 'YTickLabel','')

```

```
330 ylabel({'Amplitude [a.u.]'},...
331         'FontUnits','points',...
332         'FontSize',20,...
333         'FontName','Helvetica')
334 xlabel({'Time [ns]'},...
335         'FontUnits','points',...
336         'FontSize',20,...
337         'FontName','Helvetica')
338 end
339
340 end
```

8.7.3 fit_spectra_peaks_interact.m

```
1 function [peak, errs, ft]=fit_spectra_peaks_interact(fft, plotty)
2 %Feed function a power spectral density and it will return the position of
3 %the peak and the error in the gaussian fitting
4
5 fft=fft(1:floor(0.2*end),:);
6
7 % fft=fft(2500:end,:);%added 7/31/16 for irr_cu_16 processing
8 %
9 % [mx_val,ind]=max(fft(:,2));
10 %
11 % min_pct_fit=0.50;
12 %
13 % pos_ind=ind;
14 % go_p=1;
15 % while go_p
16 %     pos_ind=pos_ind+1;
17 %     if fft(pos_ind,2)<=min_pct_fit*mx_val
18 %         go_p=0;
19 %     end
20 % end
21 %
22 % neg_ind=ind;
23 % go_n=1;
24 % while go_n
25 %     neg_ind=neg_ind-1;
26 %     if fft(neg_ind,2)<=min_pct_fit*mx_val
27 %         go_n=0;
28 %     end
29 % end
30
31 figure()
32 plot(fft(:,1),fft(:,2),'k-');
33 hold on
34 [x_cord,y_cord]=ginput(2);
35
36 neg_x_cord=x_cord(1);
37 pos_x_cord=x_cord(2);
38
39 [~,neg_ind]=min(abs(fft(:,1)-neg_x_cord));
40 [~,pos_ind]=min(abs(fft(:,1)-pos_x_cord));
41
42 pk_trace=fft(neg_ind:pos_ind,:);
43
```

```

44 % This is the gaussian fitting for the peaks, which doesn't do so hot but
45 % is good enough
46 ft=fit(pk_trace(:,1),pk_trace(:,2),'gauss1');
47 peak=ft.b1;
48 error_mat=confint(ft,0.95);
49 errs=[peak-error_mat(1,2); error_mat(2,2)-peak];
50
51 if plotty
52     figure()
53     plot(pk_trace(:,1),ft(pk_trace(:,1)),pk_trace(:,1),pk_trace(:,2));
54 end
55
56 %Below is a record of some previous peak fitting procedures that I tried at
57 %some point.
58
59 % Try a center of mass weighting method to find the peak position
60 %% peak=(sum(pk_trace(:,1).*pk_trace(:,2)))/(sum(pk_trace(:,2)));
61 %% errs=[];
62 %% ft=[];
63
64 close all
65 end

```

8.7.4 fit_spectra_peaks.m

```
1 function [peak, errs, ft]=fit_spectra_peaks(fft, plotty)
2 %Feed function a power spectral density and it will return the position of
3 %the peak and the error in the gaussian fitting
4
5 %In case there is still any DC bit remaining in the spectrum, do not search there for the
6 %peak
7
8 fft=fft(3200:end,:); %how far past zero we are looking in the power spectrum, bigger
9 %spike in the DC level at zero
10
11 [mx_val, ind]=max(fft(:,2));
12
13 %How far down the peak to fit the gaussian to
14 min_pct_fit=0.75;
15
16 pos_ind=ind;
17 go_p=1;
18 while go_p
19     pos_ind=pos_ind+1;
20     if fft(pos_ind,2)<=min_pct_fit*mx_val
21         go_p=0;
22     end
23 end
24
25 neg_ind=ind;
26 go_n=1;
27 while go_n
28     neg_ind=neg_ind-1;
29     if fft(neg_ind,2)<=min_pct_fit*mx_val
30         go_n=0;
31     end
32 end
33
34 pk_trace=fft(neg_ind:pos_ind,:);
35
36 % This is the gaussian fitting for the peaks, which is good enough to give the peak
37 % location
38
39 ft=fit(pk_trace(:,1),pk_trace(:,2),'gauss1');
40 peak=ft.b1;
41 error_mat=confint(ft,0.95);
42 errs=[peak-error_mat(1,2); error_mat(2,2)-peak];
43
44 if plotty
45     figure()
```

```
41     plot(pk_trace(:,1), ft(pk_trace(:,1)), pk_trace(:,1), pk_trace(:,2));  
42 end  
43  
44 end
```

8.7.5 make_fft_embed_time.m

```
1 function [fft,diffusivity] = make_fft_embed_time(num_in, pos_file , neg_file ,mod, strt , grat ,
    peak)
2 % For this use only, generate filtered fft for input data
3 % num_in: legacy command, set to unity for all processing
4 % pos_file: positive phase data pos_file
5 % neg_file: negative phase data neg_file
6 % mod: provide model with which to subtract background for filtering trace before fft ,
    read code to see options
7 %     only relevant choice for the purpose here is the 'therm' option, will chose this
    by default if none slected
8 % strt: no longer used. set to 5 or 6 for processing , default to 6 if not provided
9 % grat: grating, either calibrated or estimated, in um
10 % peak:
11
12 %Boolean options for plotting and processing
13 derivative=0;
14 plotty=0;
15 plotfft=0;
16 saveout=0;
17
18 if nargin==3
19     mod='therm';
20     strt=5;
21 end
22
23 if nargin==4
24     strt=6;
25 end
26
27 hdr_len=16;
28
29 pos=dlmread(pos_file ,'',hdr_len,0);
30 neg=dlmread(neg_file ,'',hdr_len,0);
31
32 pos(:,2)=pos(:,2)-mean(pos(1:50,2));
33 neg(:,2)=neg(:,2)-mean(neg(1:50,2));
34
35 [~,time_index]=max(neg(1:1000,2));
36 time_naught=neg(time_index,1);
37
38 %If recorded traces differ in length, fix them to shorter of the two
39 if length(pos(:,1))>length(neg(:,1))
40     pos=pos(1:length(neg(:,1)),:);
```

```

41 elseif length(pos(:,1))<length(neg(:,1))
42     neg=neg(1:length(pos(:,1)),:);
43 end
44
45 fixed_short=[pos(:,1) pos(:,2)-neg(:,2)];
46
47 [~,fix_index]=max(fixed_short(:,2));
48 fixed_short=fixed_short(fix_index:end,:);
49
50 if saveout
51     dlmwrite('dat_temp.txt',fixed_short);
52 end
53
54 if plotty
55     figure()
56     plot(fixed_short(:,1),fixed_short(:,2),'r')
57     title('this is fixed_short');
58 end
59
60 %normalize to initial level before pump pulse
61 fixed_short(:,2)=fixed_short(:,2)-mean(fixed_short(end-50:end,2));
62
63 %Model option controls for trace filtering
64 custom=0;
65 therm=0;
66 thermsaw=0;
67 thermsaw_mod=0;
68 thermsaw2=0;
69
70 if strcmp(mod,'expl') %Single Exponential
71     mod_str=mod;
72     fixed_short(:,1)=fixed_short(:,1)-fixed_short(1,1); %slide peak back to 0 to fit more
        easily
73 elseif strcmp(mod,'exp2') %Sum of Exponentials
74     mod_str=mod;
75     fixed_short(:,1)=fixed_short(:,1)-fixed_short(1,1); %slide peak back to 0 to fit more
        easily
76 elseif strcmp(mod,'power2') % Two term power fit
77     mod_str=mod;
78 elseif strcmp(mod,'qexp') %Forced two term decaying exponential fit
79     custom=1;
80     fixed_short(:,1)=fixed_short(:,1)-fixed_short(1,1); %slide peak back to 0 to fit more
        easily
81 elseif strcmp(mod,'therm') %fit thermal decay to find diffusivity
82     therm=1;

```



```

83     fixed_short(:,1)=fixed_short(:,1)-fixed_short(1,1); %slide peak back to 0 to fit more
        easily
84     elseif strcmp(mod, 'thermsaw')
85         thermsaw=1;
86     %     thermsaw_mod=1;
87     fixed_short(:,1)=fixed_short(:,1)-fixed_short(1,1); %slide peak back to 0 to fit more
        easily
88     elseif strcmp(mod, 'thermsaw2')
89         if length(peak)<=1
90             thermsaw=1;
91         else
92             thermsaw2=1;
93         end
94     fixed_short(:,1)=fixed_short(:,1)-fixed_short(1,1); %slide peak back to 0 to fit more
        easily
95     else
96         mod_str='power1'; %single term power fit (goes something like 1/x)
97     end
98
99     %Fit data to one of a couple of models determined by the mod argument in
100    %the function
101
102    if custom
103        fixed_short(:,1)=fixed_short(:,1)-time_naught;
104        q=2*pi/(grat*10^(-6));
105        time_offset=fixed_short(strt,1);
106
107        LB=[0 0];
108        UB=[1 10^-4];
109        ST=[.05 10^-5];
110
111        OPS=fitoptions('Method','NonLinearLeastSquares','Lower',LB,'Upper',UB,'Start',ST);
112        TYPE=fitype('A.*erfc(q*sqrt(k*(x+time0)))','options',OPS,'problem',{ 'q', 'time0' },'
            coefficients',{ 'A', 'k' });
113
114        [f0 gof]=fit(fixed_short(strt:end,1),fixed_short(strt:end,2),TYPE,'problem',{q,
            time_offset});
115        diffusivity=f0.k;
116    elseif therm
117        q=2*pi/(grat*10^(-6));
118        LB=[0 0 0];
119        UB=[10 5*10^-4 0.1];
120        ST=[0.05 5*10^-5 0];
121        OPS=fitoptions('Method','NonLinearLeastSquares','Lower',LB,'Upper',UB,'Start',ST);
122        TYPE=fitype('A.*erfc(q*sqrt(k*(x+2.5e-9)))+c;','options',OPS,'problem','q',

```

```

123         coefficients','A','k','c'));
124     [f0_gof]=fit(fixed_short(strt:end,1),fixed_short(strt:end,2),TYPE,'problem',q);
125     diffusivity=f0.k;
126     elseif thermsaw
127         q=2*pi/(grat*10^(-6));
128         LB=[0 0 0 0 0 0];
129         UB=[10 10^-3 10 2*pi 10^-6 1];
130         ST=[0.005 5*10^-5 0.001 pi 10^-5 0];
131         OPS=fitoptions('Method','NonLinearLeastSquares','Lower',LB,'Upper',UB,'Start',ST);
132         TYPE=fitype('A.*erfc(q*sqrt(k*(x+2.5e-9)))+B.*sin(2*pi*f*(x+2.5e-9)+p)*exp(-(x+2.5e
            -9)/t)+C;', 'options',OPS,'problem',{ 'q','f'}, 'coefficients',{ 'A','k','B','p','t',
            'C'});
133
134     [f0_gof]=fit(fixed_short(strt:end,1),fixed_short(strt:end,2),TYPE,'problem',{q,peak})
            ;
135
136     diffusivity=f0.k;
137
138     elseif thermsaw_mod
139         q=2*pi/(grat*10^(-6));
140         LB=[0 0 0 0 0 0 0];
141         UB=[10 10^-3 1 10 2*pi 10^-6 1];
142         ST=[0.005 5*10^-5 0.05 0.01 0 10^-5 0];
143         OPS=fitoptions('Method','NonLinearLeastSquares','Lower',LB,'Upper',UB,'Start',ST);
144         TYPE=fitype('A.*(erfc(q*sqrt(k*(x+2.5e-9)))-D.*exp(q^2*k*(x+2.5e-9)))+B.*sin(2*pi*f*
            x+p)*exp(-x/t)+C;', 'options',OPS,'problem',{ 'q','f'}, 'coefficients',{ 'A','k','D',
            'B','p','t','C'});
145
146     [f0_gof]=fit(fixed_short(strt:end,1),fixed_short(strt:end,2),TYPE,'problem',{q,peak})
            ;
147     diffusivity=f0.k;
148
149     elseif thermsaw2
150         q=2*pi/(grat*10^(-6));
151         LB=[0 0 0 0 0 0 0 0];
152         UB=[10 10^-3 10 2*pi 10^-6 10 2*pi 10^-6 1];
153         ST=[0.05 0.5*10^-4 0.01 0 10^-5 0.01 0 10^-5 0];
154         OPS=fitoptions('Method','NonLinearLeastSquares','Lower',LB,'Upper',UB,'Start',ST);
155         TYPE=fitype('A.*erfc(q*sqrt(k*(x+2.5e-9)))+B.*sin(2*pi*f1*(x+2.5e-9)+p1)*exp(-(x+2.5
            e-9)/t1)+C.*sin(2*pi*f2*(x+2.5e-9)+p2)*exp(-(x+2.5e-9)/t2)+D;', 'options',OPS,'
            problem',{ 'q','f1','f2'}, 'coefficients',{ 'A','k','B','p1','t1','C','p2','t2','D'
            });
156
157     [f0_gof]=fit(fixed_short(strt:end,1),fixed_short(strt:end,2),TYPE,'problem',{q,peak}

```

```

        (1),peak(2) });
158
159
160     diffusivity=f0.k;
161
162     else
163         f0=fit(fixed_short(:,1),fixed_short(:,2),mod_str);
164         diffusivity=0;
165     end
166
167     if plotty
168         %plot fit to see how well it matches
169         figure()
170         plot(fixed_short(strt:end,1),f0(fixed_short(strt:end,1)),'b',fixed_short(strt:end,1),
171             fixed_short(strt:end,2),'r');
172     end
173     %this is the thermal decay filtering of the signal recorded vs time. This will clean up
174     %the DC end of the power spectrum
175     flat=[fixed_short(strt:end,1) fixed_short(strt:end,2)-f0(fixed_short(strt:end,1))];
176
177     if plotty
178         figure()
179         plot(flat(:,1),flat(:,2),'b-')
180     end
181     % Time step info necessary for differentiation and flat padding
182     timestep=flat(end,1)-flat(end-1,1);
183
184     % If option selected, take transform of derivative of recorded signal
185     % to filter out DC even more than just the background subtraction
186     if derivative
187         d_flat=diff(flat(:,2))/timestep;
188         flat=[flat(1:length(d_flat),1) d_flat];
189     end
190
191     % Find the stuff we need to take the spectral profile
192     num=length(flat(:,1));
193     fs=num/(flat(end,1)-flat(1,1));
194     p=18; %magnitude of zero padding to increase resolution in power spectrum
195     pdsiz=2^p-num-2; %more padding = smoother transform
196
197     %Only pad on the positive end
198     pad_val=mean(flat(end-50:end,2));
199     pad=zeros(pdsiz,1);

```

```

200 pad(1:end)=pad_val;
201 tpad=flat(end,1):tstep:flat(end,1)+(pdsiz-1)*tstep;
202
203 flat_pad=[flat(:,1) flat(:,2);tpad' pad];
204
205 nfft=length(flat_pad(:,2));
206
207 %Find the Power Spectral density
208
209 %Use a hamming window and a Welchs method. Hamming does the best of the
210 %ones I've tried and Welch does slightly better than the normal
211 %periodogram.
212 [psd freq]=periodogram(flat_pad(:,2),rectwin(nfft),nfft,fs); %periodogram method
213
214 %Don't save out DC spike in FFT
215 amp=sqrt(psd(5:end));
216 fft=[freq(5:end) amp];
217
218 if saveout
219     dlmwrite('dat_spec.txt',out);
220 end
221
222 if plotfft
223     figure()
224     hold on
225     plot(freq(5:end),amp,'r');
226     xlim([0 1.7e9]);
227 end
228
229 end

```

8.7.6 param_extract_time.m

```
1 function [freq ,therm , errs]=param_extract_time(num, pos_file , neg_file , grat , two_peak)
2
3 peak_fit_plot=0;
4
5 if nargin <5
6     two_peak=0;
7 end
8
9 [fft_1 ,~]=make_fft_embed_time(num, pos_file , neg_file , 'therm' ,6, grat ,0);
10
11 % [peak1 , errs1 , ft_1]=fit_spectra_peaks(fft_1 , peak_fit_plot);
12 [peak1 , errs1 , ft_1]=fit_spectra_peaks_interact(fft_1 , peak_fit_plot);
13
14
15 %Include case of meaningfully double-peaked spectra. This is the case where
16 %both a SAW and a PSAW are detectable
17 if two_peak
18     minus_peak_1=[fft_1 (: ,1) fft_1 (: ,2)-ft_1(fft_1 (: ,1))];
19     figure()
20     plot(fft_1 (: ,1) , fft_1 (: ,2) , 'b-');
21     hold on
22     plot(minus_peak_1 (: ,1) , minus_peak_1 (: ,2) , 'r-');
23     [peak2 , errs2 , ft_2]=fit_spectra_peaks(minus_peak_1 , peak_fit_plot);
24     [peak2 , errs2 , ft_2]=fit_spectra_peaks_interact(minus_peak_1 , peak_fit_plot);
25
26     freq=[peak1 peak2];
27     if freq(1)>freq(2)
28         errs (: , : ,1)=errs1;
29         errs (: , : ,2)=errs2;
30     else
31         freq=[freq(2) freq(1)];
32         errs (: , : ,1)=errs2;
33         errs (: , : ,2)=errs1;
34     end
35
36 else
37     freq=peak1;
38     errs=errs1;
39 end
40
41 if two_peak
42     [~, therm_2]=make_fft_embed_time(num, pos_file , neg_file , 'thermsaw2' ,6, grat , [peak1 peak2
43         ]);
```

```
43 else
44     [~, therm_2]=make_fft_embed_time(num, pos_file , neg_file , 'thermsaw' ,6, grat , peak1);
45 end
46
47 therm=therm_2;
48
49 end
```

Chapter 9

Acknowledgments

I have to begin by thanking my family for their support, love, and companionship, especially my partner Toomas, my parents, and my sister Taylor. And of course, the same must be said to Gram, Papa, Mimi, Bompa, and my many aunts, uncles and cousins. We are so lucky to have such a big, close family. Thank you especially to my uncles Bill and Keith, who have let me hide out with them countless times in D.C. over the past eleven years whenever I needed a respite from MIT.

Thanks also to my best friends, who are the center of my days, especially Ashley, Rosie, Cortney, Addy, and Marguerite. Thank you to Lucile Dezerald, who I met when we were both nuclear engineering interns in Paris back when the idea of getting a doctorate seemed to be just a hopeful wish for both of us. We used to daydream about her visiting me in Boston. In the years since we convinced each other to apply to grad school, she's been my roommate in Cambridge, had a postdoc at MIT, and become a professor back in France. Thank you to my other roommates from the Tosci's apartment - Alex, Brendan, Jim, and Brian - who always made returning home from lab a joy. Thank you to my "ski team:" Toom, Mangle, Katie, Mike, and Danielle. Thank you to Tim Lucas for many things, including the skunk beer (which has been passed on to Cody) and a perfect week aboard Slick (in the company of a particularly nosy barracuda) when I really needed it.

Thank you to Apollo, for reminding me that one should always make time for frisbee in the park, and that being exceptionally fluffy covers all manner of shoe-chewing related sins. Thank you to Maisie and Sidney, who are quite literally dreams-come-true, if one considers the years that Rosie and I spent wishing we had shih tzus to distract us from our problem sets.

Thank you to Riverside Boat Club - especially Lauren, Kalmia, Alexandra, Sam, Sarah, and Ilana - for flat-water mornings in the basin and sprint races and pick drills and goose-dodging and beers on the porch and lazy steady-state loops and all the other things that make rowing, but particularly rowing on the Charles, so special.

Thank you to Marguerite Siboni, Jeremy Conrad, and Grant Jordan, who have come through for me many times in the years since undergrad when I needed support or advice, and to the rest of Burton Third, who made my first years at MIT so much fun.

Thank you to the staff and patrons of the Thirsty Ear, especially Mike Grenier.

Thank you to the staff, faculty, and students of my department, MIT Nuclear Science and Engineering, which has been my academic home for more than a decade. Thank you to especially to Heather Barry: listing all the ways in which she has supported me since 2007 would probably take an entire chapter of this thesis. Thank you to Mike Driscoll, who gave me my first UROP, and whose encouragement convinced me that nuclear engineering was the field for me.

Thank you to Sarah Rankin, MIT Title IX, Blanche Staton, Dennis Whyte, Richard Lester, Ian Waitz, Doug Jonart, Sebastien Tesseyre, Michael Short, and Marina Dang for helping me with some difficult decisions, and an especially big thank you to Miriam Dreyer for helping me sort through the aftermath.

Thank you, thank you, thank you to my fellow members of the Collier Memorial design team and the Collier Strong 2014 Boston Marathon team. I'm so grateful to have been given the opportunity to be a part of those efforts. Thank you to the MIT PD for your friendship, especially Chief DiFava, Sergeant Turco, and of course Sean.

Now, for the thank yous specific to this work:

To the MIT Mesoscale Nuclear Materials group - especially Reid, Cody, Andrew, Weiyue, Miao Miao, Penghui, Rachel, Charlie, Max, and Sam - thank you for being excellent officemates, collaborators, and friends. Thank you also to Jennifer Dumel for spending a summer with us working on the TGS project.

To Marina Dang, thank you for your exceptional advice when it came to improving my presentations and papers.

To Charlie Settens of the X-Ray Diffraction Shared Experimental Facility at MIT, thank you for spending so much time - years, really - helping me run seemingly endless analyses on a seemingly endless supply of samples. Thank you for keeping the task of optimizing one diffraction spot after another from becoming painfully boring, and for your patience when I asked a million questions. Thank you for only making fun of me a little for my need to write everything down.

To Penghui Cao, thank you for your patience and guidance when it came to setting up and running molecular dynamics simulations, and for your help in improving and troubleshooting my MATLAB codes.

To Rachel Batista, thank you for helping our group run smoothly and always taking the time to answer my questions.

To Kevin Woller, thank you for always making the time to get the irradiations of my samples done so quickly, and for your idea to use TGS to study tungsten fuzz.

Thank you to Leigh Ann Kesler, Zach Hartwig, and Brandon Sorbom for their friendship and support throughout my time in graduate school. Also, thank you for letting me take some of your tungsten.

To Ken McClellan at Los Alamos National Laboratory, thank you for growing the single crystal NiAl samples used in this work.

To the Allinore group at MIT, especially Bradley Nakanishi and Andrew Caldwell, thank you for letting me anneal samples in your furnaces, and thank you for taking the time out of your own days to help me.

To Emilio Baglietto, thank you for being so supportive of both me and the people I care about. Thank you for letting us make fun of your cargo pants, which are full of secrets. I am sorry I am so bad at thermal fluids engineering. I am not sorry for making fun of your cargo pants.

To Ben Forget, thank you for always being there for me - literally since I showed up here - even when I was being *really* whiny. Thank you for chairing my defense.

To Kord Smith, thank you for making me good (better than I was, at any rate) at MATLAB. To Will Boyd and Sam Shaner, thank you for making sure I did not fail Kord's class back in 2012.

To Keith Nelson of the MIT chemistry department, this project owes a great deal of debt to your research on TGS. Thank you to you and your group members past and present, especially Jeff Eliason, for your continued advice, guidance, and collaboration.

To Michael Demkowicz, thank you for agreeing to be one of my committee members. Your input over the past four years has been invaluable, and you provided many of the ideas that wound up driving the final shape of this thesis.

To Ju Li, thank you for also agreeing to be one of my committee members. I've been very lucky to have you both as a professor in class and as a mentor on this project, and I am grateful that you have been such an important part of my education in nuclear materials science.

To Mike Short, thank you for taking me on, both in 2008 as your UROP and again - some years later - as your graduate student. I'm glad we both got to be part of each other's doctoral research. Thank you for being such a good boss, but most importantly, thank you for being one of my most trusted friends. Thank you for the embryo tea. Thank you for only occasionally accidentally setting things on fire. Thank you for the many food adventures, both here in Boston and around the world. Thank you for always being so open to changing your mind. It is a rare quality in a person, and one which I believe is the ultimate mark of a true scientist. Thank you for letting me do the type of science I dreamed I'd get to do one day when I was a little kid obsessed with physics.

Finally, to Cody Dennett, my project-mate and lab-mate and partner in grumpy cynicism, who built our TGS facility and wrote the codes that make all those little waves in the scope so useful. It isn't an overstatement to say that your efforts are responsible for many of the advances that the MIT MNM TGS project has made. When all is said and done, my most important contribution to the nuclear materials field might have been taking you out for a beer when you were still an undergrad at Cornell and convincing you that you absolutely *needed* to come to MIT and join Mike's lab. See, I was right! Cheers.

Bibliography

- [1] Cody A. Dennett, Penghui Cao, Sara E. Ferry, Alejandro Vega-Flick, Alexei A. Maznev, Keith A. Nelson, Arthur G. Every, and Michael P. Short. Bridging the gap to mesoscale radiation materials science with transient grating spectroscopy. *Physical Review B*, 94(21):214106, 2016.
- [2] Materials reliability and degradation management issues in nuclear power plants.
- [3] D. Kramer, H. R. Brager, C. G. Rhodes, and A. G. Pard. Helium embrittlement in type 304 stainless steel. *Journal of Nuclear Materials*, 25(2):121–131, 1968.
- [4] K. Matsumoto, T. Kataoka, M. Terasawa, M. Shimada, S. Nakahigashi, H. Sakairi, and E. Yagi. Embrittlement of austenitic stainless steel irradiated with α -particles. *Journal of Nuclear Materials*, 67(1-2):97–104, 1977.
- [5] S. I. Golubov, R. E. Stoller, S. J Zinkle, and A. M. Ovcharenko. Kinetics of coarsening of helium bubbles during implantation and post-implantation annealing. *Journal of Nuclear Materials*, 361(2-3):149–159, 2007.
- [6] W. Bauer, K. L. Wilson, C. L. Bisson, L. G. Haggmark, and R. J. Goldston. Alpha transport and blistering in tokamaks. *Nuclear Fusion*, 19(1):93, 1979.
- [7] W. Bauer, K. L. Wilson, C. L. Bisson, L. G. Haggmark, and R. J. Goldston. Helium induced blistering during simultaneous sputtering. *Journal of Nuclear Materials*, 76:396–401, 1978.
- [8] Shon B. Gilliam, Shaun M. Gidcumb, Nalin R. Parikh, D. G. Forsythe, Bijoy K. Patnaik, John D. Hunn, Lance Lewis Snead, and G. P. Lamaze. Retention and surface blistering of helium irradiated tungsten as a first wall material. *Journal of Nuclear Materials*, 347(3):289–297, 2005.
- [9] James E. Turner. Atoms, radiation, and radiation protection. *Edition, New York*, 1995.
- [10] Stefaan Tavernier. *Experimental techniques in nuclear and particle physics*. Springer Science & Business Media, 2010.

- [11] James F. Ziegler, Matthias D. Ziegler, and Jochen P. Biersack. SRIM—the stopping and range of ions in matter (2010). *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 268(11-12):1818–1823, 2010.
- [12] Gary S. Was. *Fundamentals of radiation materials science: metals and alloys*. Springer, 2016.
- [13] Steven J. Zinkle. Advanced materials for future nuclear plants. In *fission energy workshop: Opportunities for fundamental research and breakthrough in fission, Global Climate & Energy Project, MIT, Cambridge, MA*, 2007.
- [14] GH Kinchin and RS Pease. The displacement of atoms in solids by radiation. *Reports on progress in physics*, 18(1):1, 1955.
- [15] Michael P. Short, Cody A. Dennett, Sara E. Ferry, Yang Yang, Vikash K. Mishra, Jeffrey K. Eliason, Alejandro Vega-Flick, Alexei A. Maznev, and Keith A. Nelson. Applications of transient grating spectroscopy to radiation materials science. *JOM*, 67(8):1840–1848, 2015.
- [16] P. M. Rice and S. J. Zinkle. Temperature dependence of the radiation damage microstructure in V–4Cr–4Ti neutron irradiated to low dose. *Journal of Nuclear Materials*, 258:1414–1419, 1998.
- [17] Nikolai I. Budylnkin, Elena G. Mironova, V. M. Chernov, V. A. Krasnoselov, S. I. Porollo, and Francis A. Garner. Neutron-induced swelling and embrittlement of pure iron and pure nickel irradiated in the BN-350 and BOR-60 fast reactors. *Journal of Nuclear Materials*, 375(3):359–364, 2008.
- [18] G.S. Was and R. S. Averback. *Comprehensive Nuclear Materials*. Elsevier, Oxford, 2012.
- [19] N. Ghoniem and G. L. Kulcinski. Swelling of metals during pulsed irradiation. *Journal of Nuclear Materials*, 69:816–820, 1978.
- [20] I. A. Portnykh, A. V. Kozlov, V. L. Panchenko, V. M. Chernov, and Francis A. Garner. The mechanism of stress influence on swelling of 20% cold-worked 16Cr15Ni2MoTiMnSi steel. *Journal of Nuclear Materials*, 367:925–929, 2007.
- [21] Torben Leffers, Bachu Narain Singh, S. N. Buckley, and S. A. Manthorpe. Void-swelling in cold-worked copper during HVEM irradiation. *Journal of Nuclear Materials*, 118(1):60–67, 1983.
- [22] Louis Kenneth Mansur. Void swelling in metals and alloys under irradiation: an assessment of the theory. *Nuclear Technology*, 40(1):5–34, 1978.
- [23] L. K. Mansur, E. H. Lee, P. J. Maziasz, and A. P. Rowcliffe. Control of helium effects in irradiated materials based on theory and experiment. *Journal of Nuclear Materials*, 141:633–646, 1986.

- [24] J. Frenkel. Thermal movement in solid and liquid bodies. *Z. Physik*, 53:652–669, 1926.
- [25] Kostya O. Trachenko, Martin T. Dove, and Ekhard K. H. Salje. Atomistic modelling of radiation damage in zircon. *Journal of Physics: Condensed Matter*, 13(9):1947, 2001.
- [26] Glenn F. Knoll. *Radiation detection and measurement*. John Wiley & Sons, 2010.
- [27] G. L. Kulcinski, J. L. Brimhall, and H. E. Kissinger. Production of voids in pure metals by high energy heavy ion bombardment. Technical report, Battelle-Northwest, Richland, Wash. Pacific Northwest Lab., 1971.
- [28] Michael P Short and Sidney Yip. Materials aging at the mesoscale: Kinetics of thermal, stress, radiation activations. *Current Opinion in Solid State and Materials Science*, 19(4):245–252, 2015.
- [29] C. A. Ratsifaritana and P. G. Klemens. Scattering of phonons by vacancies. *International journal of thermophysics*, 8(6):737–750, 1987.
- [30] John M Ziman. *Electrons and phonons: the theory of transport phenomena in solids*. Oxford university press, 1960.
- [31] H. Ogi, A. Tsujimoto, M. Hirao, and H. Ledbetter. Stress-dependent recovery of point defects in deformed aluminum: An acoustic-damping study. *Acta materialia*, 47(14):3745–3751, 1999.
- [32] Hassel M. Ledbetter and Richard Palmer Reed. Elastic properties of metals and alloys, i. iron, nickel, and iron-nickel alloys. *Journal of Physical and Chemical Reference Data*, 2(3):531–618, 1973.
- [33] J. Holder, A. V. Granato, and L. E. Rehn. Effects of self-interstitials and close pairs on the elastic constants of copper. *Physical Review B*, 10(2):363, 1974.
- [34] R. Siems. Mechanical interactions of point defects. *physica status solidi (b)*, 30(2):645–658, 1968.
- [35] Seung-Hoon Jhi, Steven G. Louie, Marvin L. Cohen, and Jisoon Ihm. Vacancy hardening and softening in transition metal carbides and nitrides. *Physical Review Letters*, 86(15):3348, 2001.
- [36] F. R. N. Nabarro. Effect of radiation on elastic constants. *Physical Review*, 87(4):665, 1952.
- [37] 22.14 Nuclear Materials. Lecture 4 defects and deformation in crystals.
- [38] Robert E. Reed-Hill, Reza Abbaschian, and Reza Abbaschian. *Physical metallurgy principles*. 1973.
- [39] Cees W. Passchier and Rudolph A. J. Trouw. *Microtectonics*, volume 1. Springer Science & Business Media, 2005.

- [40] Craig R. Barrett, William D. Nix, and Alan S. Tetelman. *The principles of engineering materials*. Prentice hall, 1973.
- [41] G. S. Was and P. L. Andresen. Radiation damage to structural alloys in nuclear power plants: mechanisms and remediation. In *Structural Alloys for Power Plants*, pages 355–420. Elsevier, 2014.
- [42] F. Kroupa. Dislocation dipoles and dislocation loops. *Le Journal de Physique Colloques*, 27(C3):C3–154, 1966.
- [43] B. L. Eyre. Transmission electron microscope studies of point defect clusters in fcc and bcc metals. *Journal of Physics F: Metal Physics*, 3(2):422, 1973.
- [44] Zhanbing Yang, Norihito Sakaguchi, Seiichi Watanabe, and Masayoshi Kawai. Dislocation loop formation and growth under in situ laser and/or electron irradiation. *Scientific reports*, 1:190, 2011.
- [45] C. Cawthorne and E. J. Fulton. Voids in irradiated stainless steel. *Nature*, 216(5115):575, 1967.
- [46] V. S. Neustroev and F. A. Garner. Severe embrittlement of neutron irradiated austenitic steels arising from high void swelling. *Journal of Nuclear Materials*, 386:157–160, 2009.
- [47] S. I. Porollo, A. N. Vorobjev, Yu V. Konobeev, A. M. Dvoriashin, V. M. Krigan, N. I. Budylnkin, E. G. Mironova, and F. A. Garner. Swelling and void-induced embrittlement of austenitic stainless steel irradiated to 73–82 dpa at 335–365 c. *Journal of Nuclear Materials*, 258:1613–1617, 1998.
- [48] So J. Zinkle and K. Farrell. Void swelling and defect cluster formation in reactor-irradiated copper. *Journal of Nuclear Materials*, 168(3):262–267, 1989.
- [49] J. L. Straalsund, R. W. Powell, and B. A. Chin. An overview of neutron irradiation effects in LMFBR materials. *Journal of Nuclear Materials*, 108, 1982.
- [50] A. M. Dvoriashin, S. I. Porollo, Yu V. Konobeev, and F. A. Garner. Influence of cold work to increase swelling of pure iron irradiated in the BR-10 reactor to 6 and 25 dpa at 400°irc c. *Journal of Nuclear Materials*, 283:157–160, 2000.
- [51] Lin Shao, C.-C. Wei, J. Gigax, A. Aitkaliyeva, D. Chen, B. H. Sencer, and F. A. Garner. Effect of defect imbalance on void swelling distributions produced in pure iron irradiated with 3.5 MeV self-ions. *Journal of Nuclear Materials*, 453(1-3):176–181, 2014.
- [52] Edward A. Berndt and Igor Sevostianov. Micromechanical analysis of the effect of void swelling on elastic and electric properties of irradiated steel. *Journal of Nuclear Materials*, 446(1-3):148–154, 2014.

- [53] George Julian Dienes and A. C. Damask. Radiation enhanced diffusion in solids. *Journal of Applied Physics*, 29(12):1713–1721, 1958.
- [54] Louis Kenneth Mansur. Irradiation creep by climb-enabled glide of dislocations resulting from preferred absorption of point defects. *Philosophical Magazine A*, 39(4):497–506, 1979.
- [55] A. L. Bement. Fundamental materials problems in nuclear reactors. Technical report, Battelle-Northwest, Richland, Wash. Pacific Northwest Lab., 1970.
- [56] T. H. Blewitt, R. R. Coltman, R. E. Jamison, and J. K. Redman. Radiation hardening of copper single crystals. *Journal of Nuclear Materials*, 2(4):277–298, 1960.
- [57] S. Takamura and Y. Uesugi. Experimental identification for physical mechanism of fiber-form nanostructure growth on metal surfaces with helium plasma irradiation. *Applied Surface Science*, 356:888–897, 2015.
- [58] A. Lasa, S. K. Tähtinen, and K. Nordlund. Loop punching and bubble rupture causing surface roughening: a model for W fuzz growth. *EPL (Europhysics Letters)*, 105(2):25002, 2014.
- [59] Yu V. Martynenko and M. Yu Nagel. Model of fuzz formation on a tungsten surface. *Plasma physics reports*, 38(12):996–999, 2012.
- [60] Kevin B. Woller, Dennis G. Whyte, and Graham M. Wright. Isolated nano-tendrils on tungsten surfaces exposed to radiofrequency helium plasma. *Nuclear Materials and Energy*, 12:1282–1287, 2017.
- [61] Kevin Benjamin Woller. *Characterization of the dynamic formation of nano-tendrils surface morphology on tungsten while exposed to helium plasma*. PhD thesis, Massachusetts Institute of Technology, 2017.
- [62] Denny A. Jones. *Principles and prevention of corrosion*. Macmillan, 1992.
- [63] T. R. Allen, J. T. Busby, G. S. Was, and E. A. Kenik. On the mechanism of radiation-induced segregation in austenitic Fe–Cr–Ni alloys. *Journal of Nuclear Materials*, 255(1):44–58, 1998.
- [64] A. Etienne, B. Radiguet, N. J. Cunningham, G. R. Odette, and P. Pareige. Atomic scale investigation of radiation-induced segregation in austenitic stainless steels. *Journal of Nuclear Materials*, 406(2):244–250, 2010.
- [65] L Thuinet, M Nastar, E Martinez, GF Bouobda Moladje, A Legris, and F Soisson. Multiscale modeling of radiation induced segregation in iron based alloys. *Computational Materials Science*, 149:324–335, 2018.

- [66] P. J. Maziasz and C. J. McHargue. Microstructural evolution in annealed austenitic steels during neutron irradiation. *International materials reviews*, 32(1):190–219, 1987.
- [67] Kostya Trachenko. Understanding resistance to amorphization by radiation damage. *Journal of Physics: Condensed Matter*, 16(49):R1491, 2004.
- [68] G. S. Was, Z. Jiao, E. Getto, K. Sun, A. M. Monterrosa, S. A. Maloy, O. Anderoglu, B. H. Sencer, and M. Hackett. Emulation of reactor irradiation damage using ion beams. *Scripta Materialia*, 88:33–36, 2014.
- [69] S. J. Zinkle and L. L. Snead. Opportunities and limitations for ion beams in radiation effects studies: Bridging critical gaps between charged particle and neutron irradiations. *Scripta Materialia*, 143:154–160, 2018.
- [70] E. Getto, Z. Jiao, A. M. Monterrosa, K. Sun, and G. S. Was. Effect of irradiation mode on the microstructure of self-ion irradiated ferritic-martensitic alloys. *Journal of Nuclear Materials*, 465:116–126, 2015.
- [71] Jonathan G. Gigax, Eda Aydogan, Tianyi Chen, Di Chen, Lin Shao, Y. Wu, W. Y. Lo, Y. Yang, and F. A. Garner. The influence of ion beam rastering on the swelling of self-ion irradiated pure iron at 450 C. *Journal of Nuclear Materials*, 465:343–348, 2015.
- [72] E. H. Lee, N. H. Packan, and L. K. Mansur. Effects of pulsed dual-ion irradiation on phase transformations and microstructure in Ti-modified austenitic alloy. *Journal of Nuclear Materials*, 117:123–133, 1983.
- [73] C. Abromeit. Aspects of simulation of neutron damage by ion irradiation. *Journal of Nuclear Materials*, 216:78–96, 1994.
- [74] M. B. Lewis, N. H. Packan, G. F. Wells, and R. A. Buhl. Improved techniques for heavy-ion simulation of neutron radiation damage. *Nuclear Instruments and Methods*, 167(2):233–247, 1979.
- [75] Jinshan Yu, Junliang Liu, Jinxu Zhang, and Jiansheng Wu. Tem investigation of fib induced damages in preparation of metal material tem specimens by fib. *Materials Letters*, 60(2):206–209, 2006.
- [76] Joachim Mayer, Lucille A. Giannuzzi, Takeo Kamino, and Joseph Michael. TEM sample preparation and FIB-induced damage. *MRS bulletin*, 32(5):400–407, 2007.
- [77] <https://sites01.lsu.edu/faculty/luf/projects/>, listed as (Lu et al. RCS adv. 2013).
- [78] D. A. Young. Etching of radiation damage in lithium fluoride. *Nature*, 182(4632):375, 1958.

- [79] Y. Zhang, G. R. Lumpkin, H. Li, M. G. Blackford, M. Colella, M. L. Carter, and E. R. Vance. Recrystallisation of amorphous natural brannerite through annealing: The effect of radiation damage on the chemical durability of brannerite. *Journal of Nuclear Materials*, 350(3):293–300, 2006.
- [80] L. M. Howe, J. F. McGurn, and R. W. Gilbert. Direct observation of radiation damage produced in copper, gold and aluminum during ion bombardments at low temperatures in the electron microscope. *Acta Metallurgica*, 14(7):801–820, 1966.
- [81] E. H. Lee, J. D. Hunn, G. R. Rao, R. L. Klueh, and L. K. Mansur. Triple ion beam studies of radiation damage in 9Cr–2WVTa ferritic/martensitic steel for a high power spallation neutron source. *Journal of Nuclear Materials*, 271:385–390, 1999.
- [82] J. Ignacio GarciaáAlonso et al. Determination of fission products and actinides in spent nuclear fuels by isotope dilution ion chromatography inductively coupled plasma mass spectrometry. *Journal of Analytical Atomic Spectrometry*, 10(5):381–393, 1995.
- [83] P. Pareige, M. K. Miller, R. E. Stoller, D. T. Hoelzer, E. Cadel, and B. Radiguet. Stability of nanometer-sized oxide clusters in mechanically-alloyed steel under ion-induced displacement cascade damage conditions. *Journal of Nuclear Materials*, 360(2):136–142, 2007.
- [84] Khalid Hattar, Daniel Charles Bufford, and Daniel L. Buller. Concurrent in situ ion irradiation transmission electron microscope. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 338:56–65, 2014.
- [85] Daniel Bufford, Sarah H Pratt, Timothy J Boyle, and Khalid Hattar. In situ tem ion irradiation and implantation effects on au nanoparticle morphologies. *Chemical Communications*, 50(57):7593–7596, 2014.
- [86] Meimei Li, M. A. Kirk, P. M. Baldo, Donghua Xu, and B. D. Wirth. Study of defect evolution by TEM with in situ ion irradiation and coordinated modeling. *Philosophical Magazine*, 92(16):2048–2078, 2012.
- [87] David B. Williams and C. Barry Carter. The transmission electron microscope. In *Transmission electron microscopy*, pages 3–17. Springer, 1996.
- [88] M. Stoica, J. Das, J. Bednarcik, H. Franz, N. Mattern, W. H. Wang, and J. Eckert. Strain distribution in Zr 64.13 Cu 15.75 Ni 10.12 Al 10 bulk metallic glass investigated by in situ tensile tests under synchrotron radiation. *Journal of Applied Physics*, 104(1):013522, 2008.

- [89] J. Etoh, M. Sagisaka, T. Matsunaga, Y. Isobe, F. A. Garner, P. D. Freyer, Y. Huang, J. M. K. Wiezorek, and T. Okita. Development of a nondestructive inspection method for irradiation-induced microstructural evolution of thick 304 stainless steel blocks. *Journal of Nuclear Materials*, 440(1-3):500–507, 2013.
- [90] David A. Parks and Bernhard R. Tittmann. Ultrasonic NDE in a reactor core. In *Sensors, 2011 IEEE*, pages 618–622. IEEE, 2011.
- [91] Jeffrey W. Griffin, Timothy J. Peters, Gerald J. Posakony, Hual-Te Chien, Leonard J. Bond, Kayte M. Denslow, Shuh-Haw Sheen, and Paul Raptis. Under-sodium viewing: a review of ultrasonic imaging technology for liquid metal fast reactors. Technical report, Pacific Northwest National Laboratory (PNNL), Richland, WA (US), 2009.
- [92] Michael Trelinski. Inspection of CANDU reactor pressure tubes using ultrasonics. In *17th World Conference on Nondestructive Testing*, pages 25–28, 2008.
- [93] Robert E. Green Jr. Non-contact ultrasonic techniques. *Ultrasonics*, 42(1-9):9–16, 2004.
- [94] K. Swaminathan, A. Rajendran, and G. Elumalai. The development and deployment of an ultrasonic under-sodium viewing system in the fast breeder test reactor. *IEEE Transactions on nuclear science*, 37(5):1571–1577, 1990.
- [95] Y.-S. Joo and J.-H. Lee. Beam steering technique of ultrasonic waveguide sensor for under-sodium inspection of sodium fast reactor. In *The 13th international conference on nuclear engineering abstracts*, 2005.
- [96] A. Rajendran, C. Asokane, G. Elumalai, and K. Swaminathan. Development of an ultrasonic under-sodium scanner for the fast breeder test reactor. In *Trends in NDE science and technology: proceedings of the fourteenth world conference on NDT. V. 2*, 1996.
- [97] Rymantas Kazys, Algirdas Voleisis, Reimondas Sliteris, Liudas Mazeika, Rudi Van Nieuwenhove, Peter Kupschus, and Hamid Ait Abderrahim. High temperature ultrasonic transducers for imaging and measurements in a liquid Pb/Bi eutectic alloy. *IEEE Transactions on ultrasonics, ferroelectrics, and frequency control*, 52(4):525–537, 2005.
- [98] Elena Jasiūnienė. Ultrasonic imaging techniques for non-destructive testing of nuclear reactors, cooled by liquid metals. *Ultragarsas" Ultrasound"*, 62(3):39–43, 2007.
- [99] K. H. Matlack, J. J. Wall, J.-Y. Kim, Jianmin Qu, L. J. Jacobs, and H.-W. Viehrig. Evaluation of radiation damage using nonlinear ultrasound. *Journal of Applied Physics*, 111(5):054911, 2012.

- [100] K. H. Matlack, J.-Y. Kim, J. J. Wall, J. Qu, L. J. Jacobs, and M. A. Sokolov. Sensitivity of ultrasonic nonlinearity to irradiated, annealed, and re-irradiated microstructure changes in RPV steels. *Journal of Nuclear Materials*, 448(1-3):26–32, 2014.
- [101] Hideo Cho, Shingo Ogawa, and Mikiyo Takemoto. Non-contact laser ultrasonics for detecting subsurface lateral defects. *NDT & E International*, 29(5):301–306, 1996.
- [102] R. J. Dewhurst and Q. Shan. Through-transmission ultrasonic imaging of sub-surface defects using non-contact laser techniques. *Optics and lasers in engineering*, 16(2-3):163–178, 1992.
- [103] B. T. Khuri-Yakub. Scanning acoustic microscopy. *Ultrasonics*, 31(5):361–372, 1993.
- [104] L. W. Kessler and D. E. Yuhas. Principles and analytical capabilities of the scanning laser acoustic microscope (SLAM). *Scanning Electron Microscopy*, 1:555, 1978.
- [105] A. C. Wey and L. W. Kessler. Review of progress in quantitative nondestructive evaluation, ed. by D. O. Thompson, D. E. Chimenti, 1990.
- [106] J. Zuk, H. Kieft, and M. J. Clouter. Investigation of ion-implanted GaP layers by Brillouin scattering. *Journal of applied physics*, 73(10):4951–4954, 1993.
- [107] I. Motochi, B. A. Mathe, S. R. Naidoo, and T. E. Derry. Surface Brillouin scattering in ion-implanted chemical vapor deposited diamond. *Materials Today: Proceedings*, 3:S145–S152, 2016.
- [108] M. G. Beghi, C. E. Bottani, P. M. Ossi, T. A. Lafford, and B. K. Tanner. Combined surface Brillouin scattering and x-ray reflectivity characterization of thin metallic films. *Journal of applied physics*, 81(2):672–678, 1997.
- [109] Robert G. Leisure and F. A. Willis. Resonant ultrasound spectroscopy. *Journal of Physics: Condensed Matter*.
- [110] Albert Migliori, T. W. Darling, J. P. Baiardo, and F. Freibert. 5. resonant ultrasound spectroscopy (RUS). In *Experimental Methods in the Physical Sciences*, volume 39, pages 189–220. Elsevier, 2001.
- [111] John A Rogers, Lisa Dhar, and Keith A Nelson. Noncontact determination of transverse isotropic elastic moduli in polyimide thin films using a laser based ultrasonic method. *Applied Physics Letters*, 65(3):312–314, 1994.
- [112] John A Rogers, Martin Fuchs, Matthew J Banet, John B Hanselman, Randy Logan, and Keith A Nelson. Optical system for rapid materials characterization with the transient grating technique: Application to nondestructive evaluation of thin films used in microelectronics. *Applied physics letters*, 71(2):225–227, 1997.

- [113] Jeremy A Johnson, Alexei A Maznev, Mayank T Bulsara, Eugene A Fitzgerald, TC Harman, S Calawa, CJ Vineis, G Turner, and Keith A Nelson. Phase-controlled, heterodyne laser-induced transient grating measurements of thermal transport properties in opaque material. *Journal of Applied Physics*, 111(2):023503, 2012.
- [114] John A. Rogers, Alex A. Maznev, Matthew J. Banet, and Keith A. Nelson. Optical generation and characterization of acoustic waves in thin films: Fundamentals and applications. *Annual Review of Materials Science*, 30(1):117–157, 2000.
- [115] Jeremy Andrew Johnson. *Optical characterization of complex mechanical and thermal transport properties*. PhD thesis, Massachusetts Institute of Technology, 2011.
- [116] Jeffrey Kristian Eliason. *Optical transient grating measurements of micro/nanoscale thermal transport and mechanical properties*. PhD thesis, Massachusetts Institute of Technology, 2015.
- [117] D. Gasteau, N. Chigarev, L. Ducouso-Ganjehi, V. E. Gusev, F. Jenson, P. Calmon, and V. Tournat. Single crystal elastic constants evaluated with surface acoustic waves generated and detected by lasers within polycrystalline steel samples. *Journal of Applied Physics*, 119(4):043103, 2016.
- [118] RA Duncan, F Hofmann, A Vega-Flick, JK Eliason, AA Maznev, AG Every, and KA Nelson. Increase in elastic anisotropy of single crystal tungsten upon he-ion implantation measured with laser-generated surface acoustic waves. *Applied Physics Letters*, 109(15):151906, 2016.
- [119] Cody Andrew Dennett. *In-situ investigation of the oxidation kinetics of Fe-12Cr-2Si using time-resolved transient grating spectroscopy*. PhD thesis, Massachusetts Institute of Technology, 2017.
- [120] Reginald A Fessenden. Apparatus for signaling by electromagnetic waves., August 12 1902. US Patent 706,747.
- [121] Cody A. Dennett and Michael P. Short. Time-resolved, dual heterodyne phase collection transient grating spectroscopy. *Applied Physics Letters*, 110(21):211106, 2017.
- [122] Peter G. Malischewsky. Comparison of approximated solutions for the phase velocity of Rayleigh waves (comment on “Characterization of surface damage via surface acoustic waves”). *Nanotechnology*, 16(6):995, 2005.
- [123] C. A. Dennett and M. P. Short. Thermal diffusivity determination using heterodyne phase insensitive transient grating spectroscopy. *J. Appl. Phys.*, 123(21), 2018. In Press.
- [124] O. W. Käding, H. Skurk, A. A. Maznev, and E. Matthias. Transient thermal gratings at surfaces for thermal characterization of bulk materials and thin films. *Applied Physics A*, 61(3):253–261, 1995.

- [125] Izabela Nowak and Maria Ziolek. Niobium compounds: preparation, characterization, and application in heterogeneous catalysis. *Chemical Reviews*, 99(12):3603–3624, 1999.
- [126] Outline of the properties and habitudes of the metallic substance, lately discovered by Charles Hatchett, Esq. and by him denominated Columbium. *Journal of Natural Philosophy, Chemistry, and the Arts*, I, January:32–34, 1802.
- [127] H. Dieckamp and A. Sosin. Effect of electron irradiation on Young's modulus. *Journal of Applied Physics*, 27(12):1416–1418, 1956.
- [128] Hassel M. Ledbetter and Richard Palmer Reed. Elastic properties of metals and alloys, i. iron, nickel, and iron-nickel alloys. *Journal of Physical and Chemical Reference Data*, 2(3):531–618, 1973.
- [129] William Henry Bragg, B. A. W. L. Bragg, et al. The reflection of X-rays by crystals. *Proc. R. Soc. Lond. A*, 88(605):428–438, 1913.
- [130] B. E. Warren and B. L. Averbach. The effect of cold-work distortion on X-ray patterns. *Journal of applied physics*.
- [131] M. Wilkens. The determination of density and distribution of dislocations in deformed single crystals from broadened X-ray diffraction profiles. *Physica status solidi (a)*, 2(2):359–370, 1970.
- [132] M. Wilkens. Broadening of X-ray diffraction lines of crystals containing dislocation distributions. *Crystal Research and Technology*, 11(11):1159–1169, 1976.
- [133] T. Ungar, H. Mughrabi, D. Rönnpagel, and M. Wilkens. X-ray line-broadening study of the dislocation cell structure in deformed [001]-orientated copper single crystals. *Acta Metallurgica*, 32(3):333–342, 1984.
- [134] T. Metzger, R. Höpler, E. Born, O. Ambacher, M. Stutzmann, R. Stömmmer, M. Schuster, H. Göbel, S. Christiansen, M. Albrecht, et al. Defect structure of epitaxial GaN films determined by transmission electron microscopy and triple-axis X-ray diffractometry. *Philosophical magazine A*, 77(4):1013–1025, 1998.
- [135] T. Metzger, R. Höpler, E. Born, O. Ambacher, M. Stutzmann, R. Stömmmer, M. Schuster, H. Göbel, S. Christiansen, M. Albrecht, et al. Defect structure of epitaxial GaN films determined by transmission electron microscopy and triple-axis X-ray diffractometry. *Philosophical magazine A*, 77(4):1013–1025, 1998.
- [136] P. Gay, P. B. Hirsch, and A. Kelly. The estimation of dislocation densities in metals from X-ray data. *Acta metallurgica*, 1(3):315–319, 1953.

- [137] C. G. Dunn and E. F. Koch. Dislocation density calculation. *Acta Metall*, 5:584, 1957.
- [138] G. K. Williamson and W. H. Hall. X-ray line broadening from filed aluminium and wolfram. *Acta metallurgica*, 1(1):22–31, 1953.
- [139] M. J. Hordon and B. L. Averbach. X-ray measurements of dislocation density in deformed copper and aluminum single crystals. *Acta Metallurgica*, 9(3):237–246, 1961.
- [140] Frank H Featherston and JR Neighbours. Elastic constants of tantalum, tungsten, and molybdenum. *Physical Review*, 130(4):1324, 1963.
- [141] Robert Lowrie and AM Gonas. Single-crystal elastic properties of tungsten from 24 to 1800 c. *Journal of applied Physics*, 38(11):4505–4509, 1967.
- [142] D Io Bolef and J De Klerk. Elastic constants of single-crystal mo and w between 77 and 500 k. *Journal of applied Physics*, 33(7):2311–2314, 1962.
- [143] Xinpeng Du and Ji-Cheng Zhao. Facile measurement of single-crystal elastic constants from polycrystalline samples. *npj Computational Materials*, 3(1):17, 2017.
- [144] M Rohde. Reduction of the thermal conductivity of sic by radiation damage. *Journal of Nuclear Materials*, 182:87–92, 1991.
- [145] A. F. Mayadas and M. Shatzkes. Electrical-resistivity model for polycrystalline films: the case of arbitrary reflection at external surfaces. *Physical review B*, 1(4):1382, 1970.
- [146] Nils G. Weimann, Lester F. Eastman, Dharanipal Doppalapudi, Hock M. Ng, and Theodore D Moustakas. Scattering of electrons at threading dislocations in GaN. *Journal of Applied Physics*, 83(7):3656–3659, 1998.
- [147] B. A. Loomis and S. B. Gerber. Effect of oxygen impurity on defect agglomeration and hardening of neutron-irradiated niobium. *Acta Metallurgica*, 21(2):165–172, 1973.
- [148] Argha Dutta, N. Gayathri, S. Neogy, and P. Mukherjee. Microstructural characterisation of proton irradiated niobium using X-ray diffraction technique. *Philosophical Magazine*, pages 1–22, 2018.
- [149] Anne Jüngert. Damage detection in wind turbine blades using two different acoustic techniques. *The NDT Database & Journal (NDT)*, 2008.
- [150] Toshihiro Ohtani, Hirotsugu Ogi, and Masahiko Hirao. Acoustic damping characterization and microstructure evolution in nickel-based superalloy during creep. *International Journal of Solids and structures*, 42(9-10):2911–2928, 2005.

- [151] C. T. Walker and R. O. Pohl. Phonon scattering by point defects. *Physical Review*, 131(4):1433, 1963.
- [152] G. H. Zhu, H. Lee, Y. C. Lan, X. W. Wang, G. Joshi, D. Z. Wang, J. Yang, D. Vashaee, H. Guilbert, A. Pillitteri, et al. Increased phonon scattering by nanograins and point defects in nanostructured silicon with a low concentration of germanium. *Physical review letters*, 102(19):196803, 2009.
- [153] A. A. Maradudin. Theoretical and experimental aspects of the effects of point defects and disorder on the vibrations of crystals. In *Solid state physics*, volume 18, pages 273–420. Elsevier, 1966.
- [154] K. Bochenek and M. Basista. Advances in processing of NiAl intermetallic alloys and composites for high temperature aerospace applications. *Progress in Aerospace Sciences*, 79:136–146, 2015.
- [155] C.-S. Oh and C.-S. Han. The effects of Ni content on mechanical and oxidation resistance properties of B2-NiAl intermetallic compounds. *Journal of the Korean Society for Heat Treatment*, 26(1):1–6, 2013.
- [156] D. B. Miracle. Overview no. 104, The physical and mechanical properties of NiAl. *Acta Metallurgica et Materialia*, 41(3):649–684, 1993.
- [157] A. H. Cottrell. Constitutional vacancies in NiAl. *Intermetallics*, 3(5):341–345, 1995.
- [158] H. Wang, R. C. Reed, J. C. Gebelin, and N. Warnken. On the modelling of the point defects in the ordered B2 phase of the Ti–Al system: Combining CALPHAD with first-principles calculations. *Calphad*, 39:21–26, 2012.
- [159] E. P. Busso. *Cyclic deformation of nickel aluminide single crystals and high temperature coatings*. PhD thesis, PhD. thesis, Department of Mechanical Engineering, MIT, Cambridge, MA, May, 1990.
- [160] K. H. Johnson and J. W. D. Connolly. The electronic structures of cesium chloride type intermetallic compounds i. preliminary energy bands of β' AuZn and β' NiAl. *International Journal of Quantum Chemistry*, 4(S3B):813–825, 1969.
- [161] N. Rusović and H. Warlimont. The elastic behaviour of β 2-NiAl alloys. *Physica status solidi (a)*, 44(2):609–619, 1977.
- [162] M. L. Weaver. Investigation of strain aging in the ordered intermetallic compound beta-NiAl. Ph. D. thesis final contractor report. Technical report, Florida Univ., Gainesville, FL (United States), 1995.
- [163] A. J. Bradley and A. Taylor. An X-ray analysis of the nickel-aluminium system. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, pages 56–72, 1937.

- [164] M. J. Cooper. An investigation of the ordering of the phases CoAl and NiAl. *Philosophical Magazine*, 8(89):805–810, 1963.
- [165] H. Jacobi and H. J. Engell. Defect structure in non-stoichiometric β -(Ni Cu) Al. *Acta Metallurgica*, 19(7):701–711, 1971.
- [166] T. Hughes, E. P. Lautenschlager, J. B. Cohen, and J. O. Brittain. Erratum: X-ray diffraction investigation of β' -NiAl Alloys. *Journal of Applied Physics*, 43(6):2922, 1972.
- [167] M. Kogachi, T. Tanahashi, Y. Shirai, and M. Yamaguchi. Determination of vacancy concentration and defect structure in the B2 type NiAl β -phase alloys. *Scripta materialia*, 34(2):243–248, 1996.
- [168] A. Ball and R. E. Smallman. The deformation properties and electron microscopy studies of the intermetallic compound NiAl. *Acta metallurgica*, 14(10):1349–1355, 1966.
- [169] X. Ren and K. Otsuka. A unified model for point-defect formation in B2 intermetallic compounds. *Philosophical Magazine A*, 80(2):467–491, 2000.
- [170] S. M. Kim. Vacancy properties in β -NiGa and β -NiAl. *Acta metallurgica et materialia*, 40(10):2793–2798, 1992.
- [171] Y. Mishin, M. J. Mehl, and D. A. Papaconstantopoulos. Embedded-atom potential for B2- NiAl. *Physical Review B*, 65(22):224114, 2002.
- [172] T. Kogachi, M. and Haraguchi. Point defects in B2-type intermetallic compounds. *Materials Science and Engineering: A*, 312(1):189–195, 2001.
- [173] L. M. Pike, Y. A. Chang, and C. T. Liu. Point defect concentrations and hardening in binary B2 intermetallics. *Acta materialia*, 45(9):3709–3719, 1997.
- [174] S. R. Butler, J. E. Hanlon, and R. J. Wasilewski. Electric and magnetic properties of B2 structure compounds: NiAl, CoAl. *Journal of Physics and Chemistry of Solids*, 30(8):1929–1934, 1969.
- [175] N. Rusović and H. Warlimont. Young's modulus of β 2-NiAl alloys. *Physica status solidi (a)*, 53(1):283–288, 1979.
- [176] L. Liu, X. Wu, R. Wang, and Q. Li, W. and Liu. First principle study on the temperature dependent elastic constants, anisotropy, generalized stacking fault energy and dislocation core of NiAl and FeAl. *Computational Materials Science*, 103:116–125, 2015.
- [177] U. Essmann, R. Henes, U. Holzwarth, F. Klopfer, and E. Büchler. Containerless growth and annealing behaviour of NiAl single crystals. *Physica status solidi (a)*, 160(2):487–497, 1997.

- [178] A. Gungor, R. Ulu, and H. Demirtas. Structure and properties of arc melted dilute NiAl intermetallic compounds. *Engineering Science & Technology, an International Journal*, 16(3), 2013.
- [179] Ram Darolia. NiAl alloys for high-temperature structural applications. *JoM*, 43(3):44–49, 1991.
- [180] G. J. Dienes. Effect of radiation on elastic constants. *Physical Review*, 87(4):666, 1952.
- [181] Y. A. Chang, L. M. Pike, C. T. Liu, A. R. Bilbrey, and D. S. Stone. Correlation of the hardness and vacancy concentration in FeAl. *Intermetallics*, 1(2):107–115, 1993.
- [182] A. H. Cottrell. Vacancies in FeAl and NiAl. *Intermetallics*, 5(6):467–469, 1997.
- [183] Hsiu-Ying Chung, Michelle B. Weinberger, Jenn-Ming Yang, Sarah H. Tolbert, and Richard B. Kaner. Correlation between hardness and elastic moduli of the ultraincompressible transition metal diborides RuB₂, OsB₂, and ReB₂. *Applied Physics Letters*, 92(26):261904, 2008.
- [184] Yang-Tse Cheng and Che-Min Cheng. Relationships between hardness, elastic modulus, and the work of indentation. *Applied physics letters*, 73(5):614–616, 1998.
- [185] RS Averback and KL Merkle. Radiation-annealing effects in energetic displacement cascades. *Physical Review B*, 16(9):3860, 1977.
- [186] Steve Plimpton, Paul Crozier, and Aidan Thompson. Lammmps-large-scale atomic/molecular massively parallel simulator. *Sandia National Laboratories*, 18:43–43, 2007.
- [187] Reactors designed by Argonne National Laboratory, EBR-II, howpublished=<https://www.ne.anl.gov/About/reactors/frt.shtml>, note=Accessed May 2018.
- [188] L. C. Walters. Thirty years of fuels and materials information from EBR-II. *Journal of nuclear Materials*, 270(1-2):39–48, 1999.
- [189] F. A. Garner, T. Okita, Y. Isobe, M. Sagisaki, J. Etoh, T. Matsunaga, P. D. Freyer, Y. Huang, J. M. K. Wiezorek, and D. L. Porter. Measurement of void swelling in thick non-uniformly irradiated 304 stainless steel blocks using nondestructive ultrasonic techniques. Technical report, Idaho National Laboratory (INL), 2001.
- [190] Hex block specimens: combined sketches of inventory.
- [191] Ellen M. Rabenberg, Brian J. Jaques, Bulent H. Sencer, Frank A. Garner, Paula D. Freyer, Taira Okita, and Darryl P. Butt. Mechanical behavior of AISI 304SS determined by miniature test methods after neutron irradiation to 28 dpa. *Journal of Nuclear Materials*, 448(1-3):315–324, 2014.

- [192] J. Etoh, M. Sagisaka, T. Matsunaga, Y. Isobe, F. A. Garner, P. D. Freyer, Y. Huang, J. M. K. Wiezorek, and T. Okita. Development of a nondestructive inspection method for irradiation-induced microstructural evolution of thick 304 stainless steel blocks. *Journal of Nuclear Materials*, 440(1-3):500–507, 2013.
- [193] Y. Huang, J. M. K. Wiezorek, F. A. Garner, P. D. Freyer, T. Okita, M. Sagisaka, Y. Isobe, and T. R. Allen. Microstructural characterization and density change of 304 stainless steel reflector blocks after long-term irradiation in EBR-II. *Journal of Nuclear Materials*, 465:516–530, 2015.
- [194] C. A. Dennett, K. P. So, A. Kushima, D. L. Buller, K. Hattar, and M. P. Short. Detecting self-ion irradiation-induced void swelling in pure copper using transient grating spectroscopy. *Acta Materialia*, 145:496–503, 2018.
- [195] J. Friedel. XLVI. anomaly in the rigidity modulus of copper alloys for small concentrations. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 44(351):444–448, 1953.
- [196] Carlo Albertini, Ezio Cadoni, and George Solomos. Advances in the hopkinson bar testing of irradiated/non-irradiated nuclear materials and large specimens. *Phil. Trans. R. Soc. A*, 372(2015):20130197, 2014.
- [197] G. E. Lucas. The evolution of mechanical property change in irradiated austenitic stainless steels. *Journal of Nuclear Materials*, 206(2-3):287–305, 1993.
- [198] P. J. Maziasz. Overview of microstructural evolution in neutron-irradiated austenitic stainless steels. *Journal of Nuclear Materials*, 205:118–145, 1993.
- [199] S. M. Bruemmer and E. P. Simonen. Radiation hardening and radiation-induced chromium depletion effects on intergranular stress corrosion cracking in austenitic stainless steels. *Corrosion*, 50(12):940–946, 1994.
- [200] E. A. Kenik. Radiation-induced segregation in irradiated type 304 stainless steels. *Journal of Nuclear Materials*, 187(3):239–246, 1992.
- [201] E. A. Kenik, T. Inazumi, and G. E. C. Bell. Radiation-induced grain boundary segregation and sensitization of a neutron-irradiated austenitic stainless steel. *Journal of Nuclear Materials*, 183(3):145–153, 1991.
- [202] Heishichiro Takahashi and Naoyuki Hashimoto. Radiation-induced segregation and grain boundary migration in Fe–Cr–Ni model alloy under irradiation. *Materials Transactions, JIM*, 34(11):1027–1030, 1993.

- [203] Z. Jiao and G. S. Was. Novel features of radiation-induced segregation and radiation-induced precipitation in austenitic stainless steels. *Acta Materialia*, 59(3):1220–1238, 2011.
- [204] Zhiting Tian, Jivtesh Garg, Keivan Esfarjani, Takuma Shiga, Junichiro Shiomi, and Gang Chen. Phonon conduction in PbSe, PbTe, and PbTe $1-x$ Se x from first-principles calculations. *Physical Review B*, 85(18):184303, 2012.
- [205] M. J. Peet, H. S. Hasan, and H. K. D. H. Bhadeshia. Prediction of thermal conductivity of steel. *International Journal of Heat and Mass Transfer*, 54(11-12):2602–2608, 2011.
- [206] Erich Nechtelberger. The properties of cast iron up to 500*c. *Technicopy Ltd, xvi+ 176, A 4, illustrated*(30), 1980.
- [207] P. J. Maziasz. Formation and stability of radiation-induced phases in neutron-irradiated austenitic and ferritic steels. *Journal of Nuclear Materials*, 169:95–115, 1989.
- [208] Alan J. Ardell and Pascal Bellon. Radiation-induced solute segregation in metallic alloys. *Current Opinion in Solid State and Materials Science*, 20(3):115–139, 2016.
- [209] Y. Dong, B. H. Sencer, F. A. Garner, and E. A. Marquis. Microchemical and microstructural evolution of AISI 304 stainless steel irradiated in EBR-II at PWR-relevant dpa rates. *Journal of Nuclear Materials*, 467:692–702, 2015.
- [210] S. J. Zinkle, P. J. Maziasz, and R. E. Stoller. Dose dependence of the microstructural evolution in neutron-irradiated austenitic stainless steel. *Journal of Nuclear Materials*, 206(2-3):266–286, 1993.
- [211] C. m Shepherd and S. M. Murphy. The association of compositional fluctuations with clusters of cavities in irradiated alloys. *Journal of Nuclear Materials*, 172(2):143–150, 1990.
- [212] C. M. Shepherd. Radiation-induced segregation in ion-irradiated stainless steels: The origin of compositional fluctuations in irradiated alloys. *Journal of Nuclear Materials*, 175(3):170–176, 1990.
- [213] R. J. Wasilewski. Elastic constants and youngs modulus of NiAl. *AIME MET SOC TRANS*, 236(4):455–457, 1966.
- [214] E. S. Fisher. private communication, Argonne National Laboratory. See Rusovic, 1977.
- [215] G. Vaerst, W. Löser, D. Elefant, R. Kucharkowski, W. Michel, and L. Schultz. Stoichiometry control and mechanical properties of NiAl single crystals. *Intermetallics*, 7(3):237–239, 1999.
- [216] H. Ding, B. Medasani, W. Chen, K. A. Persson, M/ Haranczyk, and M. Asta. PyDii: a Python framework for computing equilibrium intrinsic point defect concentrations and extrinsic solute site preferences in intermetallic compounds. *Computer Physics Communications*, 193:118–123, 2015.

- [217] G. K. Williamson and R. E. Smallman. Iii. dislocation densities in some annealed and cold-worked metals from measurements on the X-ray Debye-Scherrer spectrum. *Philosophical Magazine*, 1(1):34–46, 1956.
- [218] P. A. Korzhavyi, A. V. Ruban, A. Y. Lozovoi, Y. K. Vekilov, I. A. Abrikosov, and B. Johansson. Constitutional and thermal point defects in B2 NiAl. *Physical Review B*, 61(9):6003, 2000.
- [219] R. Krachler and H. Ipser. Triple-defect complexes in the B2 intermetallic compound NiAl. *Physical Review B*, 70(5):054113, 2004.
- [220] M. Mostoller, R. M. Nicklow, D. M. Zehner, S.-C. Lui, J. M. Mundenar, and E. W. Plummer. Bulk and surface vibrational modes in NiAl. *Physical Review B*, 40(5):2856, 1989.
- [221] Z. G. Liu, J. T. Guo, L. L. He, and Z. Q. Hu. Formation of B2 intermetallic NiAl and FeAl by mechanical alloying. *Nanostructured materials*, 4(7):787–794, 1994.
- [222] Tamás Ungár, Jenő Gubicza, Gábor Ribárik, and A. Borbély. Crystallite size distribution and dislocation structure determined by diffraction profile analysis: principles and practical application to cubic and hexagonal crystals. *Journal of Applied Crystallography*, 34(3):298–310, 2001.
- [223] F. A. Garner, J. J. Laidler, B. Mastel, and L. E. Thomas. Influence of high voltage electron microscope irradiation on neutron-produced microstructures in stainless steels. In *Properties of Reactor Structural Alloys After Neutron or Particle Irradiation*. ASTM International, 1975.
- [224] Y. V. Konobeev, V. A. Pechenkin, and F. A. Garner. On the role of internal stresses in void denuded zone formation under HVEM-irradiation of metals and alloys preconditioned by reactor exposure. In *Effects of Radiation on Materials: 19th International Symposium*. ASTM International, 2000.