Molecular Dynamic Studies of the Fracture of Metals

Thesis Submitted in Partial Fulfillment
of the Requirements
of the Jay and Jeanie Schottenstein Honors Program

Yeshiva College
Yeshiva University
Fall, 2005

Emanuel A. Lazar

Mentors:
Professor Shlomo Ta'asan, Mathematical Sciences, Carnegie Mellon University
Professor Gabriel Cwilich, Physics, Yeshiva University

Table of Contents

Introduction

1.1 Metals Research

The importance of metals in our day-to-day lives can hardly be exaggerated. Since prehistoric man discovered that crafting bronze into tools could alleviate much of the physical strain hitherto involved in agriculture-related activities, metals have been a part of almost every realm of our history: agriculture, warfare, food-production, construction, transportation, industry, and modern electronics. Look around you for a moment. You probably cannot point to a single item, save yourself, that does not either contain metal nor was produced with it.

Considering the importance of metals in our day-to-day lives, it should not surprise us that people have been trying to understand the properties of metals and how to manipulate them for thousands of years. Understanding the laws that control the behavior of metals can assist us in developing new materials and in improving old ones.

Until the latter half of the twentieth century, the scientific study of metals was limited primarily to laboratories. Experiments were conducted with real samples – without melting real iron or bending bona fide copper, there was no effective way to study what might happen.

The development of computers over the last fifty years has made possible many new techniques for simulating metals and their behavior, including the molecular dynamic simulation method. Computers specialize in number-crunching and number-crunching lies at the heart of molecular dynamic simulation.

1.2 Molecular Dynamic Simulation

Many people, when they hear the term 'science experiment', immediately think of a laboratory—experiments are conducted in laboratories. Chemistry, biology and physics are all explored—experimentally—in special rooms and buildings dedicated to those sciences, with beakers and microscopes, incubators, and particle accelerators. Few people ever think about conducting science experiments inside computers. You simply cannot mix solutions inside IBMs or grow fungus inside a Macintosh. But while certain things cannot be done inside a computer, much else can be, and that is the focus of the research presented herein.

In the 1950's and 1960's, computers had advanced to a point where large numbers of calculations could be carried out by computing machines in a reasonable amount of time. This newfound power led to applications that were computationally intensive, such as numerical integration and simulations. Scientists realized the potential that computers held for studying materials that might be very difficult or impossible to study in laboratories. Complex proteins, for example, could be simulated at a much lower cost than they could be synthesized. The computer, then, became an ideal environment to study numerous topics of interest.

Molecular dynamic simulation is a method that employs computers to simulate large numbers of molecules, for 'numerical experiments' in fields including chemistry, theoretical and applied physics, and material science. Computers are given information necessary to simulate a sample. This information typically includes the positions and velocities of a large number of atoms and a number of rules (e.g. $F=ma$, $v_f=v_0+at$) by which to manipulate those atoms. Given the data for the atoms and the relevant rules, the computer is left to apply those rules in successive steps to the simulated material. It was using this technique that all of our work was done.

All of our 'experiments' were performed by computers. No real metals were stretched; no real samples were heated and cooled; no real beakers were broken. Instead, we instructed the computer to keep track of all of the atoms in a simulated piece of metal. Using the basic laws of classical physics,[1] computers computed relevant forces and velocities of thousands of atoms over the course of hundreds of thousands of time-steps. At every point in simulated time, a computer calculated the relevant forces, and recalculated the accelerations, velocities, and positions of every atom.

Within this simulation environment, we were able to perform numerous 'experiments' using simulated samples. We were able to 'heat' samples, 'cool' them, 'stretch' them, and observe every detail of interest to us. The simulated environment proved to be an ideal one; one that allowed us to change almost any environmental variable that we wished to study. 'Changing the temperature' required only tinkering a bit with parts of the preexisting program; 'identifying the individual grains' entailed coding a new subroutine; 'stretching the material'

1. The beginning of the twentieth century saw the theories of relativity and quantum mechanics replace classical physics as the accepted theories explaining the world around us. In theory, then, we should be considering modern physics in our research, and not classical physics, especially when working at the atomic level. The reasons that we, as well as do most people in the field, choose to work with classical physics are twofold: 1) The computational power necessary to calculate forces using modern physics is drastically restrictive; 2) The difference between using the two systems is insignificant for our experiments.

was accomplished by writing a more few lines of code. Pseudo-code for most of the important functions is presented in the paper itself.

1.3 Research Goals

Our research focuses on a very limited set of issues. In this paper we explore:

1) The relationship between the rate at which a molten metal is cooled and the average size of the grains in the resultant solid (Chapter 3);
2) The stress-strain behavior, including the ultimate tensile stress, the yield stress, and the shape of the stress-strain curve (Chapter 4);
3) The impact of average grain size on the ultimate tensile strength and Young's modulus (i.e. the modulus of elasticity) (Chapter 5).
4) The impact of grain orientation on the ultimate tensile strength and Young's modulus (Chapter 5).

1.4 Simulating the Materials

1.4.1 General

To carry out our experiments, we made use of an excellent package of code written by D. C. Rapaport of Bar-Ilan University and which is meant to supplement a book by the author entitled *The Art of Molecular Dynamics Simulation* [21]. Using his programs as a starting point allowed us to spend more time on our own experiments and less time on reinventing the wheel. All of our coding was done in the popular computer-programming language, "C". We used this language for a variety of reasons, primarily because the code from Rapaport was written in this language.

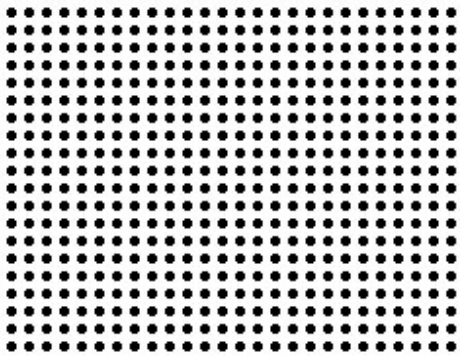

Here, we briefly explain what these simulations are about. Creating simulated metals begins with simulating individual atoms. For each atom that we wished to create, we had the computer generate a pair of numbers. The numbers in this pair represent the vertical and horizontal positions of the atom on an imagined Cartesian plane (i.e., the x,y-coordinates). All samples began with the 'atoms' positioned on either a

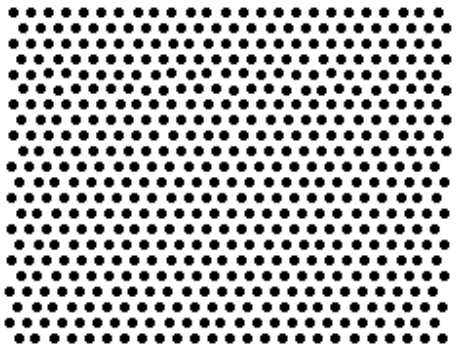

square lattice, illustrated above, or on a triangular lattice, illustrated to the left.

After the initial positions were set up, each atom was assigned a velocity. In most experiments, the magnitude of the velocity (a function of the desired temperature) was the same for every atom. The direction of the velocity, however, was different for each atom. For each atom, an angle whose measure lay between 0 and 2π radians was randomly chosen.[2] Using this angle, the computer computed the vertical and horizontal velocities for that particular atom. Once the initial configuration had been set up, the experiments could begin.

The atoms were then allowed to 'move around' in small time-steps. For every time step Δt, we changed the position of the atom using the equations:

$$x_{t+\Delta t} = x_t + v_x \Delta t \text{ and}$$
$$y_{t+\Delta t} = y_t + v_y \Delta t ,$$

where x_t and y_t are the x- and y-coordinates of the given atom at time t, and v_x and v_y are the x- and y-components of the velocity of the given atom at time t. That is, the final position after a given time step is the geometric sum of the previous position and the product of the velocity and the time step. We ran this calculation, of course, for every atom in the system at every step.

Aside from changes in position, we also needed to calculate changes in velocities and in accelerations of each atom. To do so, we calculated the net force on each atom after each step. Using Newton's equation of motion $F = ma$ (where F is the force on the particle of mass m and acceleration a) we were able to readjust the acceleration of each atom after calculating the net force on each atom. With this information, we were also able to adjust the velocities, using the equation:

$$v_{t+\Delta t} = v_t + a\Delta t$$

2. The generation of 'random' numbers by computers is a topic of much interest both to computer scientists and to theoretical mathematicians. The numbers generated appear random, and in some sense they certainly are, though, because of its deterministic nature, a computer cannot produce truly random numbers. The way in which we produced these 'random' numbers is similar to the way in which random numbers are generated for all such projects. The interested reader is referred to [16] for an overview of the use of random number in simulations.

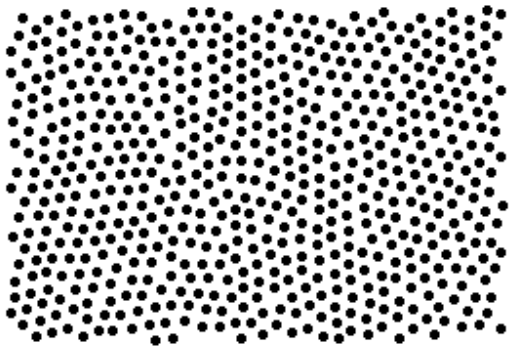

In this way, we simulated the way in which the many atoms would interact with one another and move around. Running these simulations with thousands of atoms for millions of time-steps allowed us to observe how the simulated materials might behave under specified conditions. Since the conditions were limited only by what could be programmed and run on computers, we were able to study simulated materials under almost every imaginable condition.

The diagram above shows a system of 20x30 atoms after 500 time-steps. At this point in time the atoms are "moving about" in random motion.

1.4.2 Periodic Boundaries

Properties of matter are noticeably different inside a material than they are at its surface. This causes us special problems when simulating a material. Due to limits of current computing power, regardless of how many atoms we could practically simulate, that number of atoms would be orders of magnitude smaller than the number of atoms in a real material, and the entire sample would behave as if it were entirely a surface sample. Because we wanted to study what occurs inside a metal when it fractures, simply simulating a material with typical boundaries would be unrealistic.

To solve this problem, we used ***periodic boundary conditions*** (PBC). Using PBC meant having the computer imagine that there were no free surfaces in the system. The left-most atoms in the system are imagined to be nearest neighbors of the right-most ones, and the upper-most ones of the bottom-most. Because of this, atoms inside the system act as if they were in an effectively infinite system.

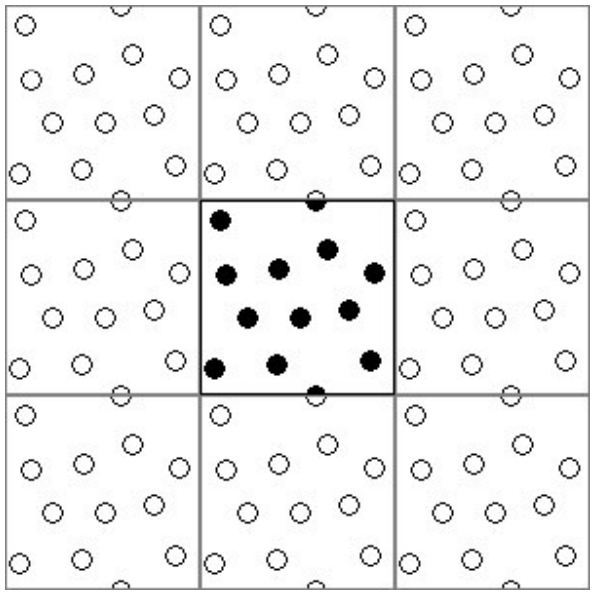

The picture to the right might help the reader visualize what was going on. The square in the middle contains the simulated

atoms that exist in the mind of the computer. The surrounding squares are merely imagined by the computer, in order to simulate an infinite system on all sides. The atoms on the left side of the main square interact with carbon copies of the atoms on the right side of the square, which are copied in an adjacent square to the left of the 'real' square.

In this way, all simulated atoms act as though they were positioned in the interior of a material, regardless of how close they may actually be to the surface. This technique helps us appropriately prepare the material we wished to study.

The nature of some of our experiments precluded using PBC at later stages of experimentation. Because we planned to fracture the metals by pulling them apart, we needed to have free space in the vertical plane. For this reason, we wrote special code that could turn on and off PBC as necessary. We generally used PBC when creating the sample, and then turned it off before attempting to study fracture.

1.4.3 The Potential Energy Function

The motion of each atom is largely determined by its neighboring atoms. Every atom exerts some force on every other atom, and the net force exerted on a given atom determines its acceleration and its eventually motion. Because all movement depends on correctly determining the forces that each atom exerts on its neighbors, properly modeling these forces is crucial to properly simulating a material.

The important question then is: how do we calculate the net force on a given atom? One of the first equations used in molecular dynamic simulation describes the Lennard-Jones Potential (LJP).[3] Owing to the nature of certain chemical bonds and to the manner in which electrons orbit the nuclei of the atoms and deform electron clouds of other atoms, various forces attract and repel the atoms to and from each other. When the atoms are relatively close to one another, the repellant force is very strong and prevails over the attractive force; when the atoms are far apart from each other, an attractive force pulls them closer together.

The following equation is typically used to describe this potential:

---

3. A note of possible interest to some readers: One of the first papers to discuss MDS in general, and the Lennard-Jones potential in specific, was written by the physicist Loup Verlet, at the time a professor in the Belfer Graduate School of Science of Yeshiva University. The most commonly used time-integration algorithm in MDS is called the Verlet Algorithm, named after this author, and introduced in that paper [27].

$$\phi_{LJ}(r) = 4\varepsilon\left[\left(\frac{\sigma}{r}\right)^{12} - \left(\frac{\sigma}{r}\right)^{6}\right]$$

The term *r* represents the distance between the two atoms. The terms ε and σ are material-dependant parameters.

Here is a graph of the potential as a function of distance between two atoms:

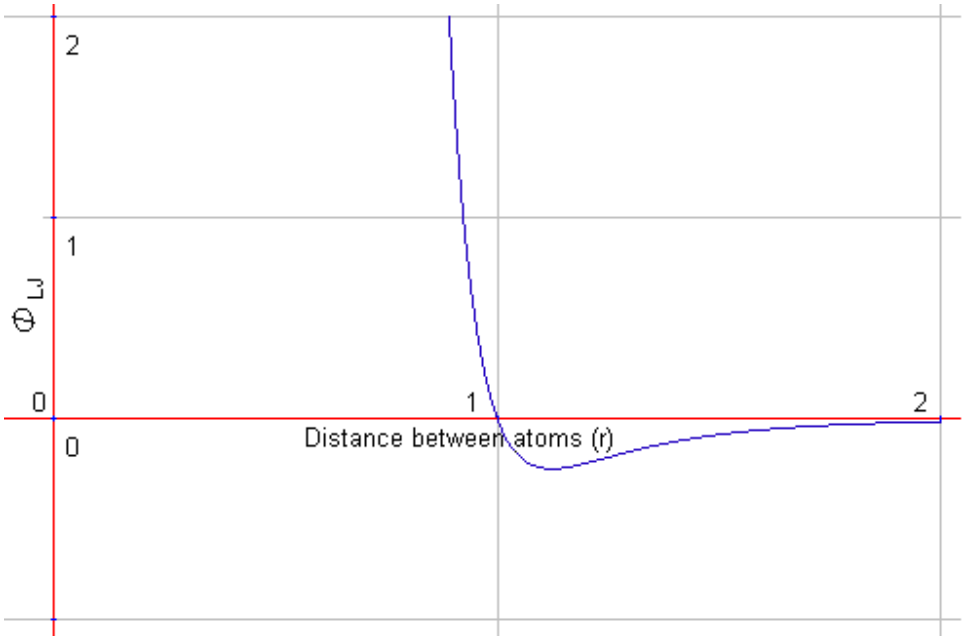

As you can see, when the atoms are close together, a positive potential is induced, pushing the atoms apart. When the atoms are further apart, an attractive force pulls them together. This attractive force, known as the Van der Waals attraction, is caused by the interactions between induced dipole moments.

It is important that we note here that although this potential is commonly used, it is not the only possible, or the only utilized, potential. While many other interatomic potentials have been developed, we limit consideration to this potential for the sake of simplicity.

1.4.4 Simulation Size

We used between 3000 and 40,000 atoms in most of our simulation cells. Depending on the nature of the individual experiment, various sample sizes proved to be computationally feasible and acceptable. Ideally, if our computing power had no limits, we would have run experiments with millions or billions of atoms. These sorts of simulations would more strongly resemble macroscopic

metals. However, our resources are limited and such experiments were not possible.

What then could we do, considering our limited resources? We have already introduced the idea of PBC that helps resolve some related problems. In addition, we sometimes ran tests on increasingly larger sets of atoms, paying attention to whether sample size affected the end results. If the results did not change with an increase in system size, we could safely ignore our limited system sizes. If results did change with increasing system sizes, we could sometimes conjecture what the data would look like in very large systems. Moreover, sometimes this helped us only identify very general trends rather than particular numbers. For example, we show that a decreasing cooling rate is correlated with a decrease in grain size. While the particulars of this correlation might depend somewhat on sample size, we have no reason to believe that this trend will reverse itself or otherwise significantly change when we consider larger sample sizes.

1.4.5 Accuracy

Awareness of the tension between computational efficiency and scientific accuracy informed our every step. Obviously, scientific accuracy is crucial if our research is to have any scientific meaning and value. On the other hand, pedantic attention to accuracy can seriously compromise the scope of our research. Because experiments can take hours, days, and sometimes weeks to run, it was critical to understand when accuracy was necessary and when it was gratuitous. For example, in certain tests, alterations could be made to the code in order to cut the run-time in half. If this would reduce the accuracy of the end result by 20%, for example, we would probably opt for the longer route to obtain more accurate results. If, however, such a cut would result in a reduction in accuracy of 1/10 of a percent, we would probably sacrifice the accuracy for the sake of the time saved. With the extra time, just think, we could run twice as many experiments!

Knowing which changes effected which reductions in accuracy was itself a complicated job. Oftentimes we would invest hours testing experiments with different parameters, so as to save tens of hours down the line. Sometimes it is crucial to invest heavily now if we are to gain in the long run.

One example of time-saving measures taken is the division of the system into cells to calculate the forces acting on each atom. This allowed us to reduce an $O(n^2)$ problem to an $O(n)$ problem.[4]

4. Computer scientists use a special notation to describe, in very general terms, how much time is necessary to implement a given algorithm. The letter n is used to denote the size of the input of a given function. For example, in sorting a list of words, n would denote the number of words

Consider the system, shown earlier in section 1.4.1, shown to the right. One method of calculating forces would be to calculate the force between every atom and every other atom in the system. We would end up making roughly n^2 calculation, if n is the number of atoms in the system, and so the algorithm would be described as $O(n^2)$. This might be acceptable if we were dealing with a small number of atoms, but it would quickly get out of hand when the number of atoms becomes large.

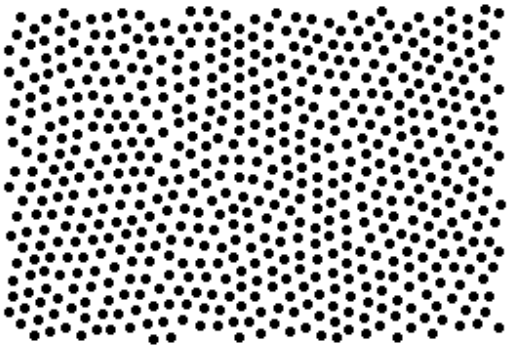

Moreover, it should be apparent from looking at the potential energy function shown in section 1.4.3, that the force existing between two atoms at large separations is very small compared with that existing between two atoms at small distances. The force that exists between an atom at the left side of the above sample and another atom at the right side of the sample is negligible. In order, then, to sensibly reduce the time necessary for calculations of force, we split the sample into numerous cells. We then have the computer calculate the forces that exist between atoms only if they are either in the same cell or in adjacent cells.

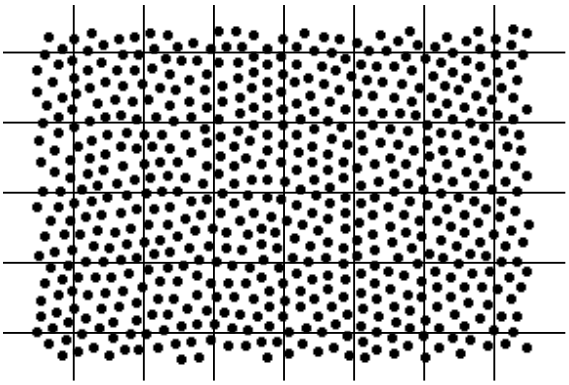

The picture to the left demonstrates how the sample looks after being divided into cells. Seeing that there are almost 50 cells, we can figure out how much time we are saving by using this cell-method. For each atom, we are now calculating

---

we'd like to sort. If we say that a given function runs in $O(n)$ time, we mean that the time necessary to execute that function grows linearly with the size of the input—sorting 200 numbers would take, theoretically, twice as long as would sorting 100 numbers. If we say that a given function runs in $O(n^2)$ time, we mean that the time necessary to execute the function grows as the square of the size of the input—sorting 200 numbers would take four times as long as would sorting 200 numbers. If we say that a given function runs in $O(\log n)$ time, we mean that the time necessary to execute that function grows as the log of the size of the input—sorting 200 numbers would take log200/log100, or about 1.15, times as long as would sorting 100 numbers.

Finding an appropriate algorithm to solve a certain problem is crucial for using computing resources efficiently. The study of algorithms is in itself an vast and fascinating topic, and we could easily written an book about the efficiency of various algorithms for MD experiments. Clearly, though, such an endeavor is well beyond the scope of our research.

the forces between it and atoms in 9 cells, instead of in all 48 cells. This turns out to be a saving of roughly 80% run-time. If the number of cells were even larger, as it would be in larger systems, that number could turn into 90%, 95%, and even higher. This method results in an $O(n)$ algorithm instead of the painfully-slow $O(n^2)$ necessary for calculating the forces between every pair of atoms. The time saved using this method surely is worth the slight loss in accuracy that could have been saved had we calculated the forces for *every* pair of atoms.

This size of the cells is large enough so that forces between all pairs of atoms separated by less than a distance *rCut*, would be calculated. In our experiments, we took *rCut* to be 2.5σ. The force between two atoms in our systems separated by such a distance turns out to be an attractive force of 0.00408, which, when compared to a maximum value of the function of 0.25000, is negligible in our final results.

1.4.6 Units of Measurement

Working with extremely small objects can be difficult when measuring quantities such as position and velocity is necessary. Because our particles are of atomic size, the number of meters and seconds that describe their positions and velocities are very small when expressed in standard SI units. For this reason we decided to use special dimensionless MD units.

When using these units, all measurements are described in units that are easier to deal with. Instead of saying that two atoms are 0.00000000034 meters apart, we can say that they are 1 MD length unit apart. Or instead of saying that the mass of an atom is $6.6904265 \times 10^{-23}$g, we can write instead that its mass is 1 MD unit of mass. Eventually we can convert our results back into standard units, but for the meantime, this notation will make our lives considerably easier. Throughout the thesis, all units should be understood to be MD units, where lengths are expressed in σ, energy in ε, and mass in m.

1.4.7 Hardware Specifications

Most of the computations were carried out on one of two standard off-the-shelf computers. At home in Pittsburgh, I used a PC equipped with an Intel Celeron 2.4 Ghz processor, and loaded with 256 MB of standard memory. While working in Israel, I used both an Intel Celeron 600MHz, loaded with 192MB or RAM, and later another laptop using a Mobile AMD Sempron 3000+, running at 1.8 Ghz, and loaded with 448 MB of memory.

Grain Identification

2.1 Background

Much of our interest lies in examining ensembles of atoms that are organized into groups called grains.[5] Properties of these grains directly affect many of a material's mechanical properties. In Chapter 5, we examine the strength of a metal, and its relation to the average size of the grains. Because grains are so important to our work, we devote an entire chapter to the problem of grain identification.

To study properties of the grains, it is crucial that we instruct the computer how to organize the many atoms into defined groups, or grains. After being able to 'identify' individual grains, we could then spend efforts studying those grains.

This task of 'identifying grains' is rather complicated, and often there exists no singular, 'correct' solution to the problem. Sometimes we have atoms about which it is not clear whether they should be counted as part of one grain or another. Other times, it is not obvious if we should count two areas towards one large grain or whether we should count them as two separate, albeit similar and adjacent, ones. Even when it is clear that two areas should be separated, it is often unclear where the separation between the two lies. Because of the inherent ambiguity in solving these problems, a degree of arbitrariness is introduced into the relevant algorithms.

In the following discussion, we define the problem of grain identification and explain how we went about solving it. We note in advance that there is often no singularly correct solution to this problem, and that every solution is plagued by quirky, unavoidable problems.

2.2 The Problem

The problem of grain identification is to specify criteria by which to divide a sample of atoms into a smaller number of groups called grains. A precise definition of grain might be given: A ***grain*** is a contiguous set of atoms sharing the same crystal symmetry and the same spatial orientation of the crystallographic axes.

---

5. The term grain is commonly used to refer to a phase of matter consisting of a regular arrangement of component units. The terms 'grain' and 'crystal' are used interchangeable throughout the literature to refer to the same entity. For the sake of uniformity, we will use only the term grain.

We need to decide what should constitute a grain, and how we are to determine whether a given atom should be considered part of a specified grain or not. We want a function that for any given atom can tell us whether or not it is in any grain and if it is, in which grain is it.

To solve this problem, we developed a number of techniques, the last of which we implemented in our numerical experiments. We leave for another time a detailed explanation of alternative methods that we did not use.

2.3 The Algorithm

In explaining our algorithm, we provide 5 examples and show pictures every step of the process.

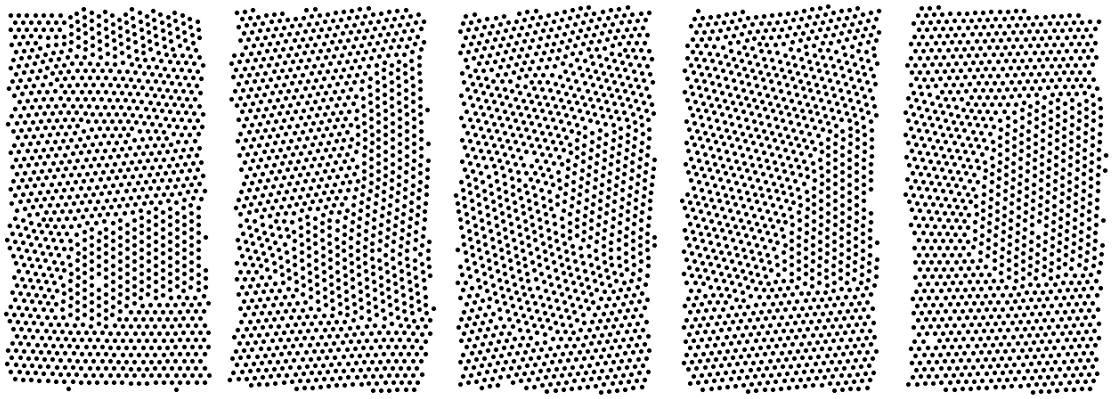

From merely glancing at the above pictures, you might notice that the samples have a number of different groups of atoms, each of which belongs to some grain. While it might not be entirely clear where one grain ends and where another one begins, it should be clear that the samples above can be divided into different grains.

The first task of our algorithm is to identify atoms that are most likely part of the interior of a grain. We can identify these atoms as those that have 6 neighbors that are approximately equidistant from itself. To the right is a picture of what such a situation might look like. The black atom in the middle is what we call 'well-centered' amongst the gray-colored ones.

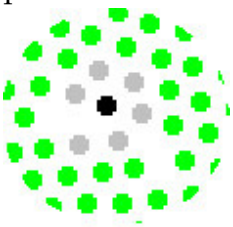

Atoms that do not have 6 approximately equidistant neighbors are either part of the exterior of a grain or are not part of any grain at all. To the left is a picture of what such a situation might look like. The black atom in the middle of the of the

gray-colored ones does not seem to be well-centered. Most likely, it is an atom belonging to a grain boundary.

```
Identify Internal Atoms ()
{
   for all atoms i
     if(neighbors(i)==6 AND stddev of distance from i
          to its neighbors is less than 0.05) then i
          is internal
     else i is external
}
```

If the standard deviation of the distance between the atom and its surrounding neighbors is too high, then chances are high that the atom is not well centered in the interior of a grain. The second image above, with the green surrounding atoms, is an example of such a phenomena.

Here are pictures of the 5 samples with the external atoms marked in black and the internal ones in gray:

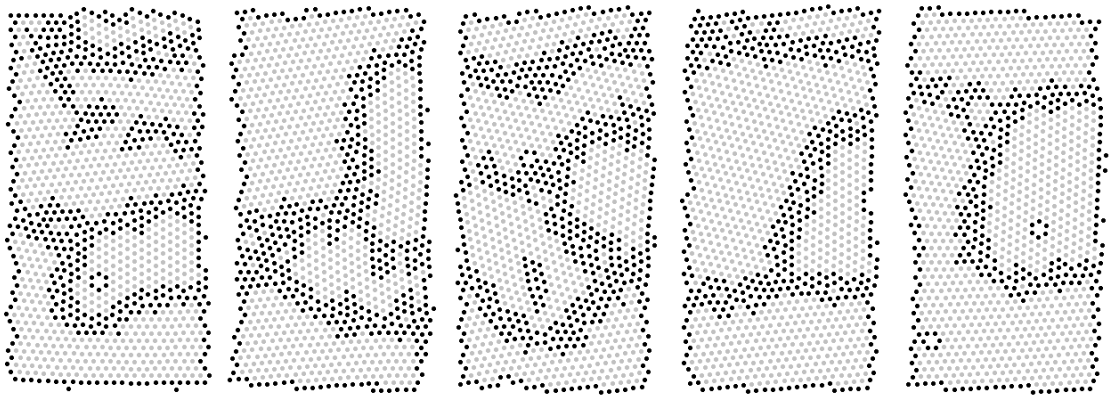

After marking off the external atoms, we also mark off all atoms that has at least two external atoms as neighbors. This is done to ensure that two separate grains, which happen to share two or three atoms, are not mistakenly joined as one. At this stage we eliminate grains below a certain threshold size. In our experiments, we eliminated any grains that had three or fewer atoms.

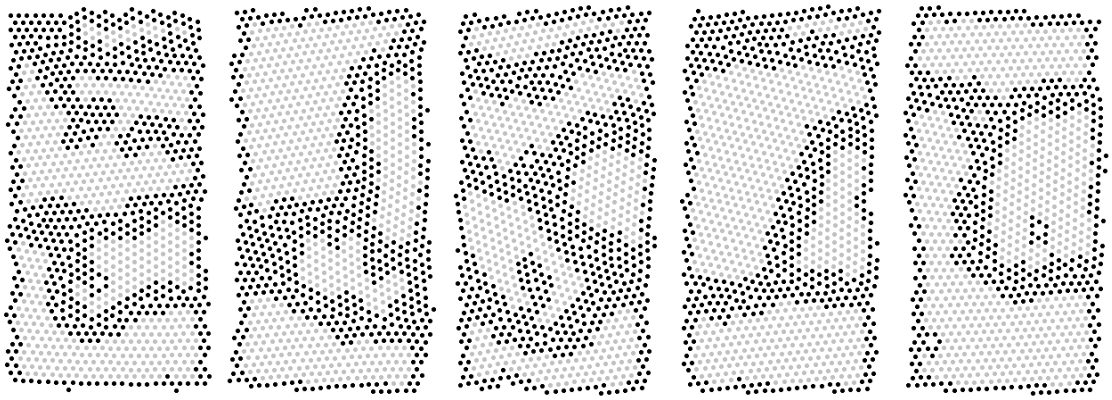

After the internal and external atoms are marked, the neighboring atoms are dealt with, and the mini-grains eliminated, we can begin coloring in the various grains. We start with any internal atom in the system and recursively color it and all its neighboring internal atoms the same color and call it a grain. When all internal atoms adjacent to that grain are colored, we start growing the next grain, beginning with an uncolored internal atom. We repeat this procedure until all internal atoms in the system are colored.

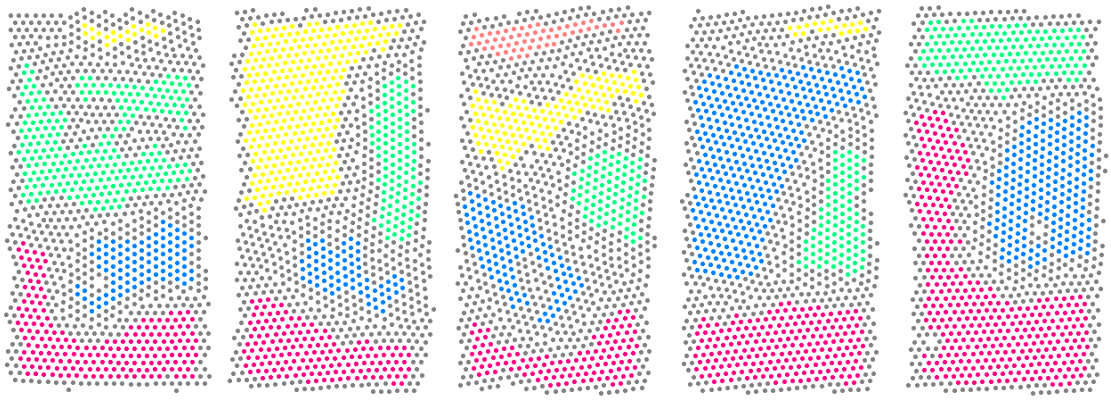

The final step in identifying the grains is expanding them, to let them subsume neighboring atoms that may have originally been marked off as grain-boundary atoms, but that could be considered part of the grain.

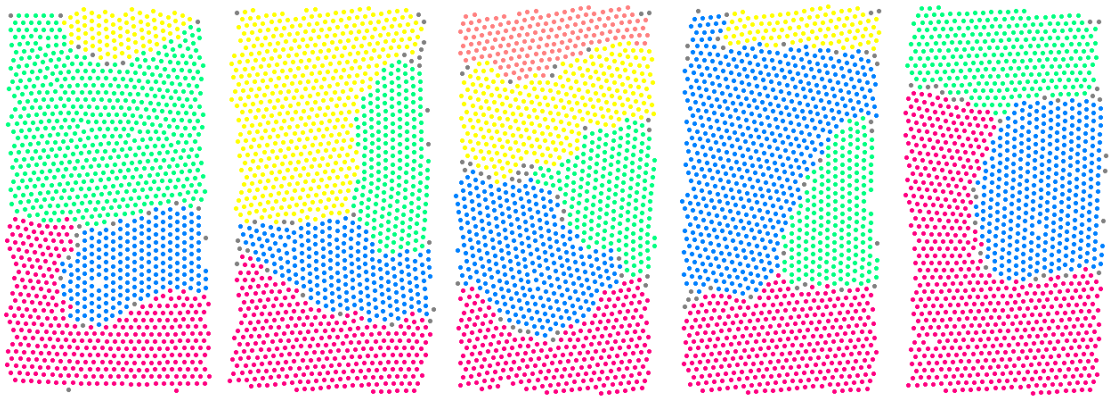

Our algorithm is thus able to read a set of coordinates and 'identify' patterns that we know as grains. It is important to note here, as we have done before, that the task of 'identifying grains' is an inherently complex one, and often there exists no single correct solution to the problem. Explaining these complications and describing their impact on solving this problem is beyond the scope of this thesis, and addressing it briefly would only do the topic injustice. I hope to address these issues at another time in more depth.

2.4 Other Methods

All of our experiments began with atoms moving around freely as a liquid and then cooling and settling into a crystalline grain structure. We then worked to identify the grains from what had formed 'naturally'. Another method used in MD simulations is the filling of a sample area with prefabricated grains. In this method, an average grain size is first chosen and then grain centers are distributed accordingly inside an area or volume. A method introduced by G.F. Voronoi in 1908 is used to divide the surrounding space in a manner that every point is associated with the grain center closest to it. The orientation of the grain is chosen beforehand and does not naturally develop. Two examples of research that use this method are [8] and [26].

An advantage of using such grains is the significant time saved in generating the grains, a process that takes almost no time when employing this method. However, this method also has a number of disadvantages. First, it is often difficult to obtain lattices with normal distributions of the grains [14]. Moreover, it is not clear that these artificially-created grains realistically resemble real grains of metals. It seems reasonable to believe that the shapes of the grains impact a number of their properties. Further work is necessary to determine whether grains generated in the two different methods produce different results in simulations.

Cooling Rates and Grain Sizes

3.1 Background

While heated in a gas or liquid state, the molecules of a metal move around quickly and do not settle into any structured (crystalline) form. As a metal cools, it solidifies, and grains begin to form.

It has been shown experimentally that the rate at which metals are cooled impacts the size of the grains that form. Quickly cooling metal produces small grains; slowly cooling them produces larger ones. The precise relationship between the cooling rate and the average grain sizes has not yet been fully documented. The published experimental results on the topic have been few, and there has been no published research detailing MDS reproduction of the experimental results.

In this section, using two-dimensional molecular dynamic simulations, we study how various cooling rates impact average grain sizes. Our simulations help us understand why the cooling rate affects the grain sizes of the metals. **We show that the average size of the grains is a function of the cooling rate.** More specifically, the function demonstrating the relation between the average grain size $\bar{d}$ and the cooling rate r is of the form:

$\bar{d} \propto r^x$, where x is about –0.18.

Our research confirms some prior, experimental data, which we cite later.

3.2 Terms and Concepts

The ***kinetic energy*** of each atom is its energy of motion. In these experiments we focused on the translational kinetic energy, a quantity which depends on two variables: the mass (m) and the magnitude of the velocity (v) of the atom. The following equation is commonly used to describe the kinetic energy of an atom,

$$KE = \frac{1}{2}mv^2$$

The ***temperature*** of a system is a measure of the average kinetic energy of the atoms in the system. To determine the temperature of the system, we sum the kinetic energy of each atom and take the average over all of the atoms in the system:

$$T \propto \frac{1}{n}\sum_{i=1}^{n}\frac{1}{2}m_i v_i^{\,2},$$

where n is the number of atoms in the system, m is the mass of each atom, and v is the velocity of each atom. Because m is 1 (in MD units) for all atoms, we can say:

$$T \propto \frac{1}{n}\sum_{i=1}^{n} v_i^{\,2}$$

Quenching is the rapid cooling of a metal from its initial temperature. This process, often accomplished by immersing a metal in oil or water, is used in processing metals in order to achieve certain properties, some of which we will study. Liquids can be transformed into solids by cooling them. The speed of this process is called the ***cooling rate*** and is measured in units of temperature per unit time. A typical unit for measurement is degrees per second.

Relaxing the system is a method in which we bring all the atoms in the system to a relaxed state—the net force acting on each atom is zero and the temperature is zero. In reality a system never fully relaxes and the atoms are always in motion, always vibrating around a fixed center.[6]

3.3 Survey of the Literature

Attempts to determine the relationship between cooling rates and grain sizes have already been made theoretically and experimentally. In [3], the authors propose two different dependences of the aluminum grain size on cooling rate. To the right is a graph [3] of these two relations: the upper line has $\bar{d} \propto r^{-0.75}$ and the lower line, which the authors ultimately prefer, has $\bar{d} \propto r^{-0.9}$.

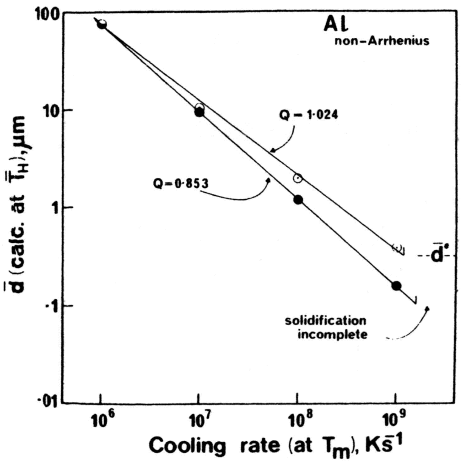

In [15] and [29], it was shown that the average grain sizes decreases with increasing cooling rate. In these papers, however, no quantitative analysis of this relationship is shown. In [25], experimental work is conducted in different iron-based systems. Among other things, the authors document

6. This is the so-called zero point motion.

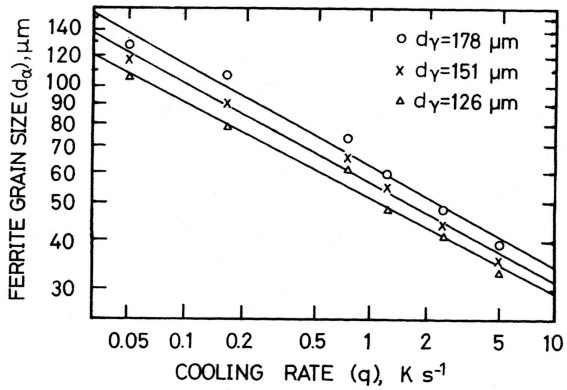

Effect of cooling rate on ferrite grain size for specimens with three different austenite grain sizes

the relationship between cooling rates and average grain sizes.

To the left is a graph [25] showing the relationship between cooling rates and average grain sizes of ferrite (body-centered cubic iron). Results from these experiments show that, $\bar{d} \propto r^{-0.26}$.

The authors [25] do not cite the earlier work presented in [3] and so have no need to explain the ostensibly large discrepancy between their work and the previous work [25].

To date, no research has been published that has reproduced the theoretical [3] or experimental [25] trends reported above. Perhaps this is the case because until now, most MD research constructs the sample artificially, and does not allow grains to develop 'naturally', as explained in Section 2.4. This thesis represents an important step in understanding the relationship between grain size and cooling rate and the influence of grain size and fracture properties in the context of MD simulations.

3.4 Simulating Quenching, Relaxation, and Grain Growth

3.4.1 Introduction

The simulations described in this part of the thesis are relatively simple: we prepared the samples as a liquid, quenched them, relaxed them, and then observed the size of the grains that formed.

For each trial, after choosing the desired dimensions of our sample, we had the computer place the atoms on a rectangular lattice and assign each atom a velocity. The magnitude of this velocity was the same for all atoms (a function of the desired initial temperature of the system); the direction of the velocity was randomized for each atom.

The atoms in the system were then allowed to move around following the laws of classical physics. We used the Leapfrog method [11b] to compute the velocities and coordinates of the atoms after every step. In this method, we adjusted the velocities for each atom, and then repositioned each atom according to its new velocities.

```
Leapfrog Step()
{
   for all atoms i
      v_ix(t + Δt) = v_ixt + a_ix*Δt
      v_iy(t + Δt) = v_iyt + a_iy*Δt
      x_i(t + Δt) = x_it + v_ixt*Δt
      y_i(t + Δt) = y_it + v_iyt*Δt
}
```

3.4.2 Cooling the System

After allowing the atoms to move about for a short time, we began quenching, or cooling, the material. We simulated quenching as follows: every 100 steps, we multiplied the velocity of every atom by some number between 0 and 1. We repeated this every 100 steps until the temperature of the system fell below 0.01 (MD units).

```
Quench(QuenchRate)
{
   for all atoms i
      v_ix = v_ix * QuenchRate
      v_iy = v_iy * QuenchRate
}
```

Below is a graph showing the temperature of a 60x40 atom system as a function of time. The different graphs show different cooling rates.

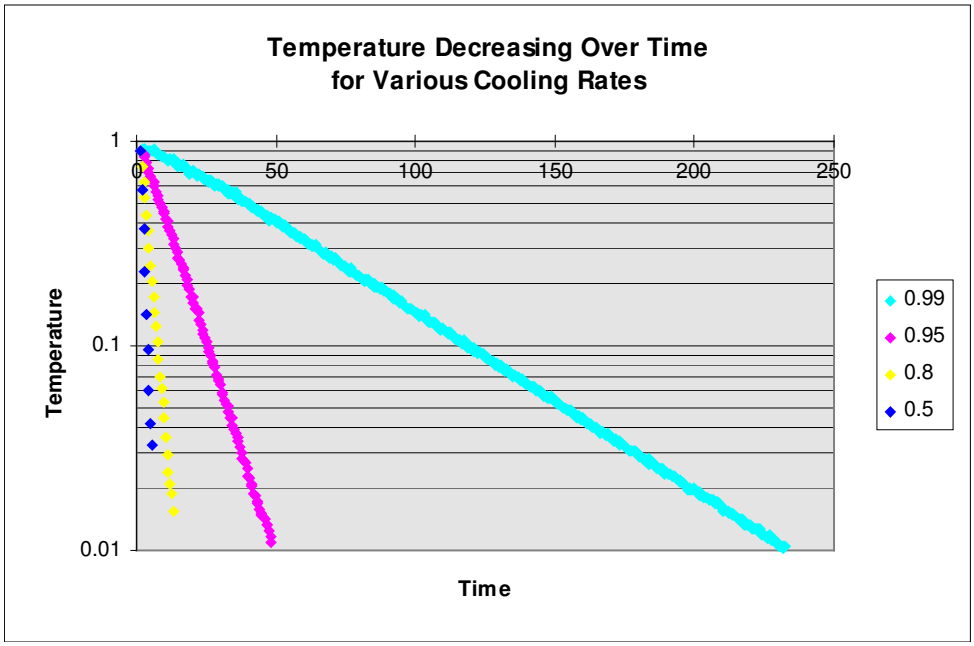

The graph above shows that the temperature of the system decreases exponentially with time. The numbers in the key are the value by which we multiplied the velocity of each atom every 100 time-steps. The higher the number, the longer it took for the system to cool down.

3.4.3 Relaxing the System

After cooling the system, we needed to then relax it, or reduce both the kinetic and potential energy of the system to near zero. This required two things: turning off the periodic boundary conditions (PBC) and changing the laws governing the motion of the atoms.

Turning off the PBCs was necessary to ensure that no external pressure was being exerted on the system. Turning off these boundary conditions was not overly complicated, but did require some tinkering with the pre-existing code. Because most of the code had been written for PBC simulations, and because we did not want to needlessly rewrite hundreds and thousands of lines of code, we changed a few parameters that would let the computer act as if there were no PBC, even though there really was. This was done by increasing the size over which the system was periodic. This effectively put a finite vacuum region around the crystal. In practice we multiplied the dimensions of the system by a factor of 3, so that if the system had been 100 x 100 in size, it now appeared to be 300 x 300 in size. An atom located at (50, 0), which before would effectively have been located right next to an atom situated at (-49, 0), was now located 200 away.

This change in dimensions effectively eliminated PBC effects while allowing us to keep the same code we had used until then.

Aside from tinkering with the boundary conditions, we also replaced the true MD method with a method that simply minimizes the energy of the system with respect to atomic coordinates (i.e., T=0). Maintaining MD methods would leave the atoms vibrating, as if at some finite temperature. If we let this happen, then the vibration rate of the atoms and the rate used in loading the sample to fracture (i.e., the stretching rate) would be similar. This is physically unrealistic because in real systems the vibration rate of the atoms is many orders of magnitude higher than the loading rate. Relaxing the system and abandoning MD methods lets us deal with the atoms as if they were quasi-static, corresponding to a much more realistic model of stretching.

While employing MD methods, we used the Leapfrog method, described above, to govern the motion of the atoms. While relaxing the system, we moved away from MD methods and used an artificial means of dynamics to govern the motion of the atoms:

```
Relax Step()
{
   for all atoms i
      x_{i(t + Δt)} = x_{it} + a_{ix}*Δt
      y_{i(t + Δt)} = y_{it} + a_{iy}*Δt
}
```

Thus, instead of using the equations $\dot{v} = F$ and $\dot{x} = v$, we used $-\dot{x} = F$. This is not molecular dynamics but simply relaxation, since this drains all of the kinetic energy from the system. Using these equations eventually relaxed the system, and since the PBC had already been turned off, we end up with a set of atomic configurations corresponding to zero net-stress. In our simulation experiments, we considered the system of atoms to be 'relaxed' when the average square of the net force on an atom was less than 10^{-12}.

Below is a graph showing the average square of the net force on an atom as a function of number of iterations during the relaxation calculation:

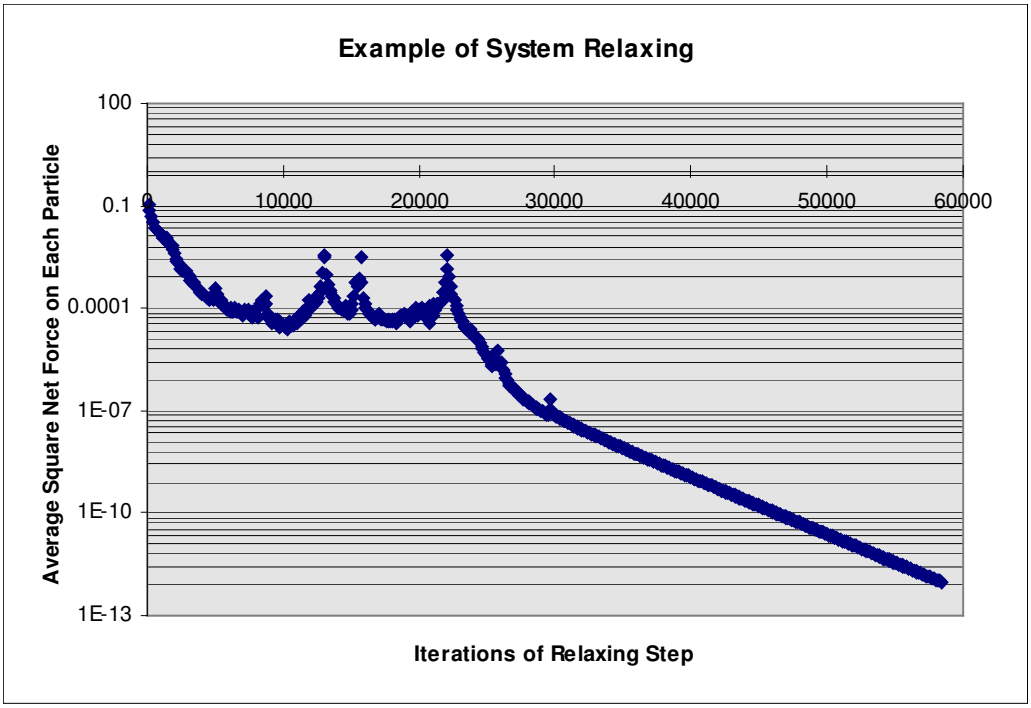

As this figure shows, initially, the average square force does not decrease monotonically with iteration number. The spikes in the graph indicate the times at which some grains dissolved into other grains or where two grains merged; these are events that take place early in the relaxation. Later, the number of grains tend to stay as they are and the atomic positions simply become more relaxed with the passage of more steps; this explains the relatively straight line at the tail of the function.

After the system had been cooled and relaxed, we identified the grains, as explained in Chapter 2, measured their sizes, and calculated averages.

3.4.4 Calculating Average Grain Size

In most experimental work, the size of a grain is typically determined by the grain's average diameter. Experimentally, this property is easier to measure than its surface area or volume. Even ignoring the questions of how to measure the diameter of non-circular entities, for us the easiest property to measure was the number of atoms in each grain. However, since we want to compare our computational results with experimental ones, we needed to talk in terms of diameter. Therefore, we took as our grain sizes the square root of the number of

atoms in each grain. The square root of the number of atoms in a grain is proportional to its radius and diameter.

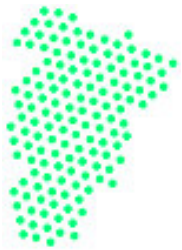

As an example, consider the image on the right. The displayed grain contains 144 atoms. We would say that the size of the grain is $\sqrt{144}$, or 12.

3.5 Simulation Experiments

3.5.1 Simulation Experiment 1

In this experiment, we began with a system of 50 x 50 atoms (2500 atoms in total) at a temperature of 1.0. Every 100 time steps, we multiplied the velocities of each atom by r, where $0.71 \leq r \leq 0.99$. We cooled the system until its temperature had fallen below 0.01, a point at which samples where considered to be low temperature solids. After cooling the systems, we relaxed them, as described in section 3.4.1.

To determine the cooling rate, we needed to measure the time (number of steps) taken for the system to cool. Because the samples all began at a temperature of 1.0, we could calculate the cooling rate as follows:

Cooling Rate = 1.0 / Time to Cool

After cooling and relaxing the system, we analyzed the grains, and calculated their average size. This was done as explained above in section 3.4.2. We ran 5 trials for each cooling rate. The results below are the average of the 5 trials. Further data from this experiment can be found in the Appendix.

Below is a graph showing the results of our simulation experiments.

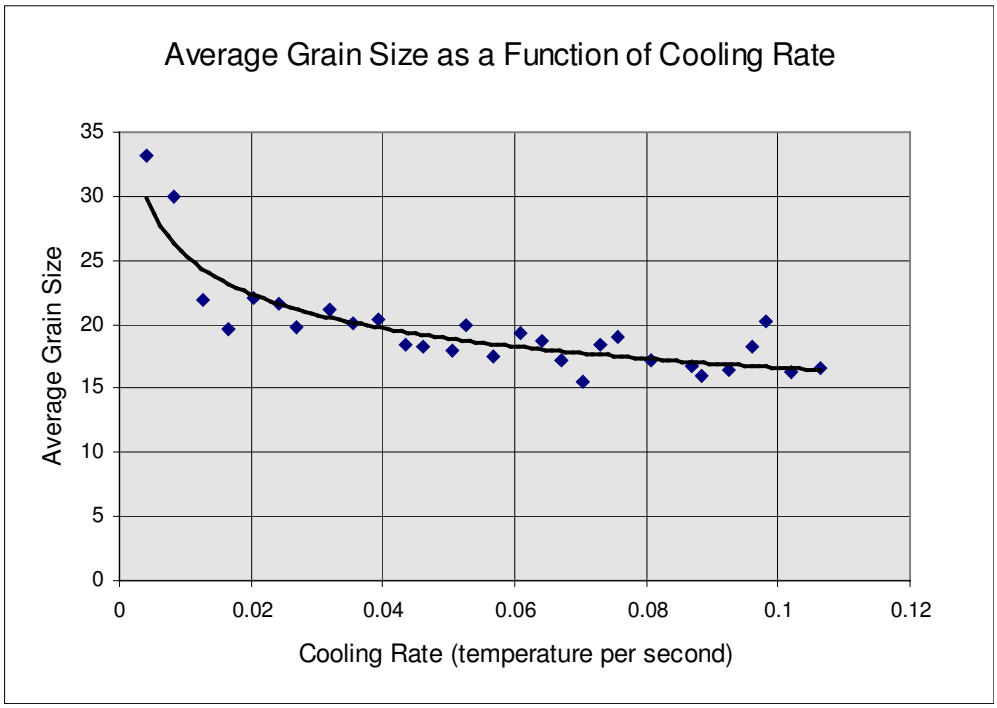

The solid line shown in this plot is the best-fit power-law, i.e.

$$f(r) = 10.872r^{-0.1846},$$

where r is the cooling rate. We compare this result with those provided in [3], [15], [25], and [29].

3.5.2 Simulation Experiment 2

While this simulation experiment produced considerably less data than the previous one, it is important in that it confirms the results of the previous simulation experiment by reproducing similar results with samples half the size.

In this experiment, we began with a system of 25 x 50 atoms (1250 atoms in total) at a temperature of 1.0. Every 100 time steps, we multiplied the velocities of each atom by r, where $0.70 \leq r \leq 0.95$. We cooled the system until its temperature had fallen below 0.01, a point at which samples could be considered solids. After cooling the systems, we relaxed them, as described in section 3.4.1.

We determined the cooling rates in the simulation experiments as described above. After cooling and relaxing the system, we analyzed the grains, and calculated their average size. This was done as explained above in section 3.4.2.

We ran 5 trials for each cooling rate. The results below are the average of the 5 trials. Further data from this simulation experiment can be found in the Appendix. Below is a graph showing results of our simulation experiments.

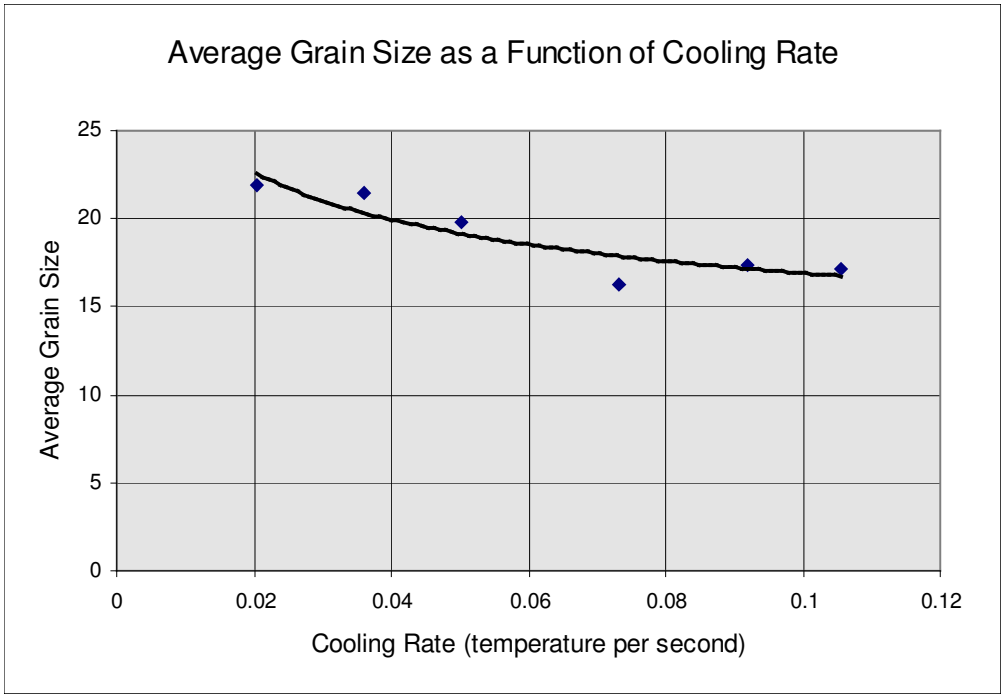

The best fit power law in this figure is of the form:

$$f(r) = 11.142r^{-0.1809},$$

where r is the cooling rate. In the analysis we compare this result with those provided in [3], [15], [25], and [29].

3.6 Observations and Analysis

The data from our two experiments confirms that the average size of the grains diminishes with increased cooling rates, and is consistent with results published in [3], [15], [25], and [29].

From the data collected in Simulation Experiment 1 and Simulation Experiment 2, it appears that the average grain size can be described as $r^{-0.180}$ and $r^{-0.185}$. These results are similar to the relationship $\bar{d} \propto r^{-0.26}$ reported in [25], and which was measured experimentally. The results are not consistent with the relationships $\bar{d} \propto r^{-0.75}$ or $\bar{d} \propto r^{-0.9}$ reported in [3], and which were predicted theoretically.

Further work must be conducted to explain the inconsistencies between the theory presented in [3] and the results published in [4-25] and that are reproduced here. Our work is the first to replicate the experiments conducted in [3], [15], [25], and [29] by means of MDS.

Stress-Strain Behavior

4.1 Background

In this section, we explain what stress-strain curves are and a number of the issues associated with drawing these curves in our simulated experiments. This will become important in the next chapter, where we look at the relationship between grain size and a number of important properties of metals.

When a tensile force is applied, all materials stretch; eventually, after sufficient force is applied, even the strongest of metals fail and break. Stress-strain curves are a means of visually representing the relationship between the amount of force applied to a metal and the way in which the metal changes shape in response. Various properties of these curves, for example, its domain and range, are of great interest to people studying metals.

4.2 Terms and Concepts

The term ***strain*** is used to measure the geometrical deformation of a sample. Strain, often represented by the letter ε, is computed as follows:

$$\varepsilon = \frac{L - L_0}{L_0}$$

where L_0 is the original length of the sample and L is the new length of the material after having been strained. For example, a meter-long sample of iron stretched out to 1.03 meters has been strained 3%.

The term ***stress*** is used to measure the relative force applied to a sample. Stress, often represented by the letter σ, is generally computed as follows:

$$\sigma = \frac{P}{A_0}$$

where P is the load, or force, applied to the sample, and A_0 is the initial cross-sectional area of the sample. In our simulations, where the samples were two-dimensional, the load P was divided by the length of the sample normal to the stress:

$$\sigma = \frac{P}{L_0}$$

The equation above, containing the term A_0 is used for computing what is known as the ***engineering stress-strain curve***. These calculations take into account only the initial cross-section area of the sample, and do not consider the deformation known as necking that occurs during stretching. Calculations taking this factor into consideration produce what are known as ***true stress-strain curves***. We leave until section 4.4 a detailed explanation of the difference between these two curves and a thorough explanation of necking.

Real, physical, experiments are usually performed by stretching (i.e., straining) a sample and measuring the resultant stress as a function of strain. In some experiments, the material is stretched at a fixed rate; in others, a constant force is applied and the resultant elongation of the sample is measured.

To the right is an image of a typical engineering stress-strain curve. Stress is shown as a function of strain. As the material is stretched, the stress varies.

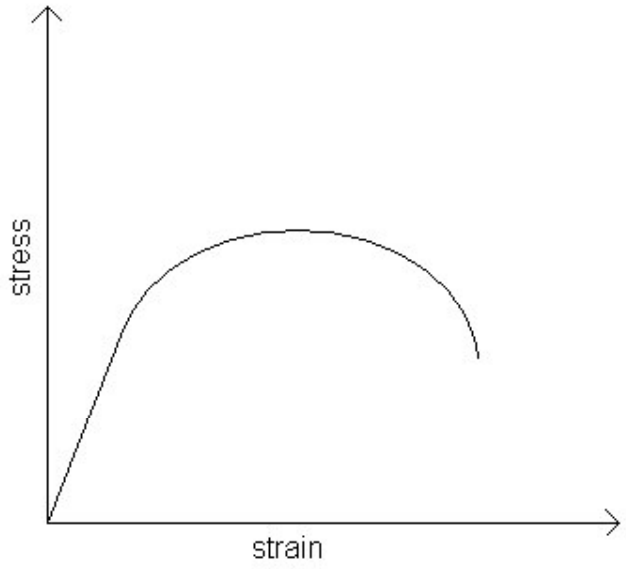

As you can see, the beginning of the curve shows an apparently linear relation between the stress and the strain. The slope of this initial segment of the curve is known as the ***Young's modulus*** (or modulus of elasticity) typically represented by the letter E:

$$E = \sigma / \varepsilon .$$

A steep slope indicates a stiff material. A material with a high value of E will require a great deal of stress to achieve a particular strain; a material with a small value of E deform significantly with very little stress. The value of E is material-specific and can be affected slightly by factors such as temperature.

As long as the stress-strain relationship can be expressed as a linear function, the deformation of the material is ***elastic***. This means that when the stress is removed the material will return to its original shape. The stretching of an elastic rubber band, is an excellent example of an elastic deformation; it returns to its original shape after the applied stress is removed.

Once the shape of the graph becomes nonlinear, the deformation is generally ***plastic***. This means that even after the stress is removed, the material will not return to its original shape. Plastics are good examples of materials that do not return to their original state after an applied stressed is removed.

At a certain point during stretching, deformation ceases to be elastic and begins to be plastic. On the graph, this is the point at which the curve ceases to be linear and begins to bend. Knowing this point is important because we often want to know how much stress an object can endure without sustaining permanent damage. A property of metals associated with this point is the ***tensile yield strength***, or just yield strength.

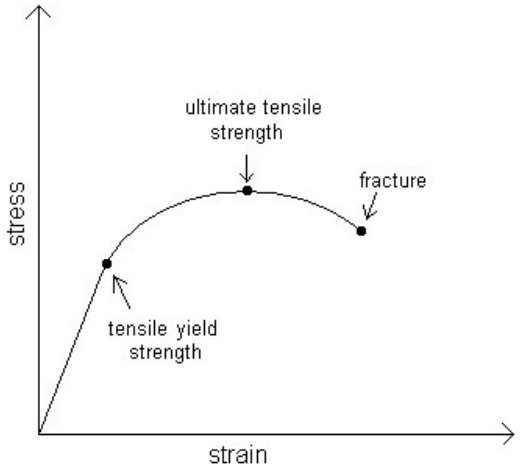

After a material is strained beyond where the stress exceeds the yield strength, it can still endure additional stress before breaking. The ***ultimate tensile strength*** is the maximum stress that a material can endure before fracturing. The ability to measure this property is important in planning the proper use of materials.

After hitting the ultimate stress, ductile metals begin to neck. In ***necking***, the material thins out in the middle; necking forestalls fracture. A brittle, or inflexible, material will fracture before it necks.

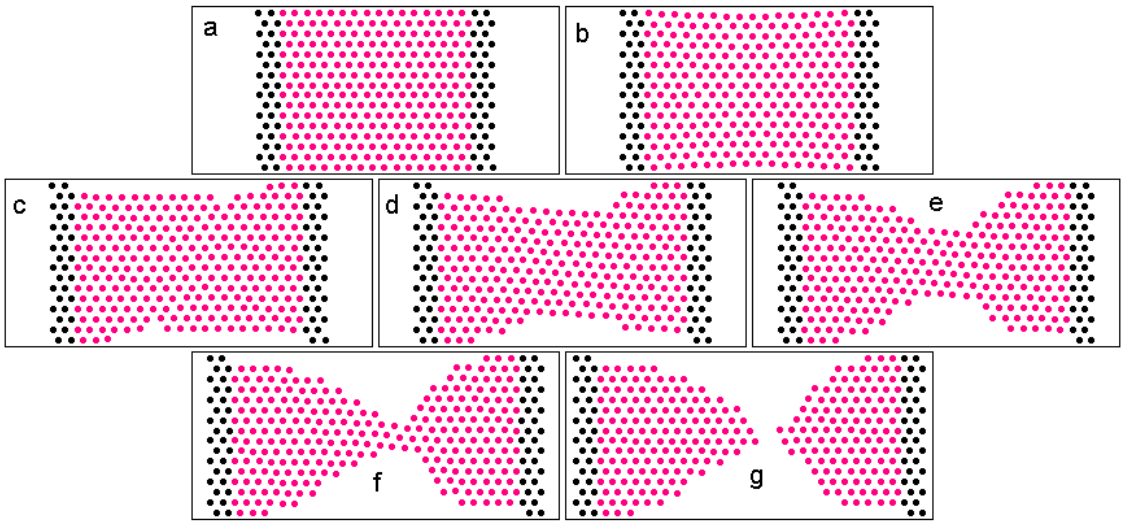

Above are seven pictures of a material undergoing strain: (a) No strain, (b) 10%, (c) 20%, (d) 30%, (e) 40%, (f) 50%, and (g) fracture at 58.5%.

Necking actually creates serious problems in computing the stress-strain curve. As explained in Section 4.2, an engineering stress σ is described by the formula:

$$\sigma = \frac{P}{A_0}.$$

This measurement of stress considers only the initial cross-sectional area A_0 (or cross-sectional length L_0) of the sample being stretched. However, when a material necks, its cross-sectional area shrinks, and is smaller than A_0. The true stress of the material is given by

$$\sigma_{true} = \frac{P}{A}$$

where A is the instantaneous cross-sectional area (or length) of the sample. In our experiments, we did not attempt to measure the instantaneous height of the sample and we thus settled for calculating the engineering stress rather than calculating the true stress.

4.3 Simulating Stress and Strain Behavior

4.3.1 Creating Sample Walls

In the previous chapter, we discussed preparing the samples, turning on and off the periodic boundary conditions (PBC), and relaxing the system. Once that is dealt with, we can soon begin experimenting with stretching the samples.

In order to stretch the samples, we must hold the samples by their edges, and to facilitate this, we needed to simulate edges. We did this by marking off the atoms that were in the right-most and left-most cells of the system and designated them as 'wall atoms'. When stretching the sample, the walls as whole entities were moved, but their internal atoms remained otherwise stationary. The free atoms in between the walls were free to move. The wall atoms here are colored black while the remaining atoms are colored pink.

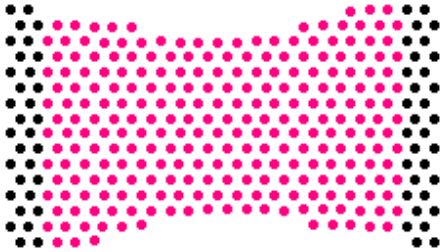

4.3.2 Stretching the Material

Stretching the material in our simulated experiments meant changing the x-coordinate of every atom of the system. To do so, we multiplied the x-coordinate of all interior atoms by some number $1+\alpha$, where α was typically 0.0025, which corresponds to a strain of 0.25%. Between every stretch, we allowed the sample to relax, as explained in Section 3.4.3. We continued stretching the material until it eventually fractured.

Since we wanted to increase the strain at a rate fixed relative to the initial configuration of the system, we could not scale the x-coordinates by the same value for every increment of the stretch. If we scaled the system by $1+\alpha$ for every strain, then after n stretches we will have scaled the system by $(1+\alpha)^n$ instead of an intended $1+n\alpha$. After 40 stretches, for example, instead of applying a strain of 10%, we would have applied a strain of 10.5%; after 100 stretches, instead of applying a strain of 25%, we would have applied a strain of 28%, a significant difference.

For this reason, we kept track of the number of stretch increments applied up to any given point. When we wanted to stretch the system again, we scaled the system as follows. If n gives the number of stretches including the current one, we would scale the system by a constant:

$$c = \frac{1+n\alpha}{1+(n-1)\alpha}$$

Making these calculations ensured that the system was being strained at a rate fixed relative to the original configuration of the system.

Wall atoms were always moved at the constant rate of α.

```
Stretch(α, n)
{
   const c = (1.+nα)/(1.+[n-1]α)

   for all internal atoms i
      xi = c * xi

   for all left-wall atoms i
      xi = xi - α

   for all right-wall atoms i
      xi = xi + α
}
```

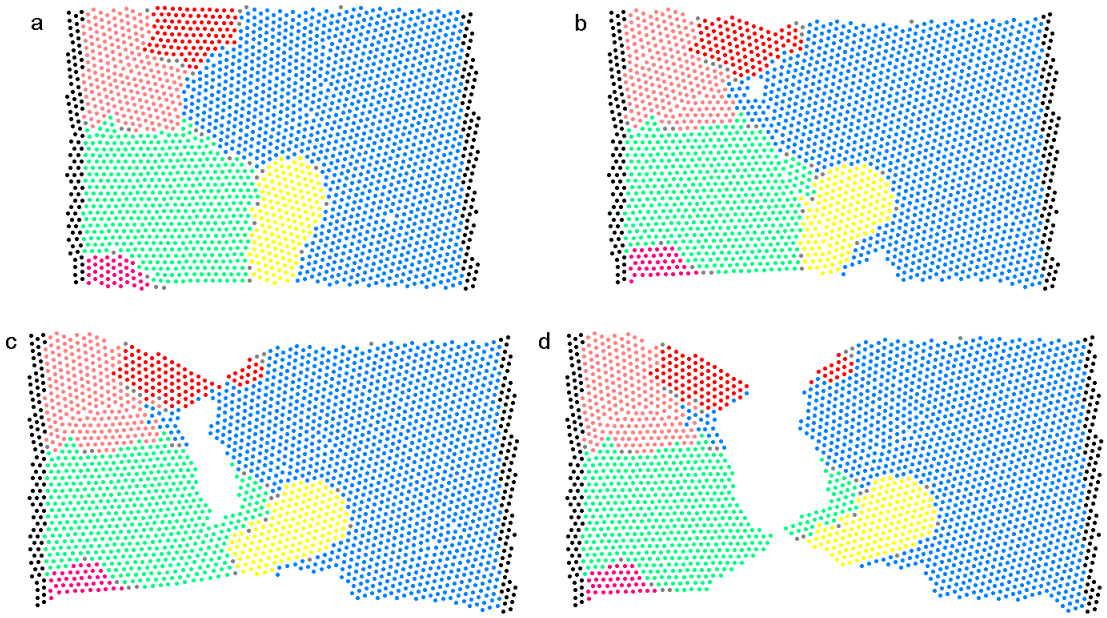

Above is 40x60 sample, with initial grains identified, shown at (a) No strain, (b) 10% strain, (c) 20%, and (d) fracture at 32.25%.

4.3.3 Measuring the Applied Stress

The most important test for measuring a metal's response to stress is the tensile test. In this test, one side of a sample, often shaped as a rod or wire, is held fixed

and the other end is pulled at a controlled rate. A machine connected to the sample measures the force associated with displacement. Alternatively, in some tests, the force applied is controlled and the displacement is the variable that is measured.[7]

Our experiments roughly simulated the tensile test on our simulated samples. Instead of holding one end fixed, we pulled the two ends in opposite directions. While in a physical laboratory, such a procure is difficult, while in our simulated experiments, this was straightforward. We controlled the strain on the system and measured the stress at every point, as is done in standard tests.

To measure the total stress we needed to calculate the net force on each of the two walls as we were stretching the sample. Dividing the total force on the walls by the length of the walls, we were able to compute the force per length, or stress σ, that had been applied in order to stretch the sample.

4.3.4 Determining Sample Fracture

Since we intend to study how much stress a material can endure before fracturing, we need to know when we should consider a material 'fractured'. Determining this was quite simple—when the sample fractured, the net force on the walls, or stress on the system, dropped to 0. Thus, when we calculated the stress to be 0, we could stop stretching the sample (i.e., it is already broken).

4.3.5 Simulation Procedure

Most of the simulation experiments followed a simple procedure. The first step was setting up the system, cooling and relaxing it, and identifying walls which we would hold fixed when stretching. Then, we iteratively stretched the system, measuring at every step the stress and strain, until the sample fractured. After every stretch we waited until the system relaxed before stretching it again.

```
SetupSystem()
CoolSystem()
while (sample not fractured)
{
   Stretch(α, n)
   RelaxSystem()
   Output Strain, Stress
}
```

7. Testing stress-strain properties of materials, as well as most experimental procedures used to test properties of metals, is commonly detailed by standards-setting agencies, in this case the American Society for Testing and Materials (ASTM). The test described here, the tensile test for metals, is detailed in [2].

4.4 Demonstrations

In this section we provide an example of our simulated stretching and explain many of the calculations that we prepared. We have already seen, in section 4.3.2, an example of a material stretching until fracture. Here we show the stress-strain graph associated with that sample's stretching:

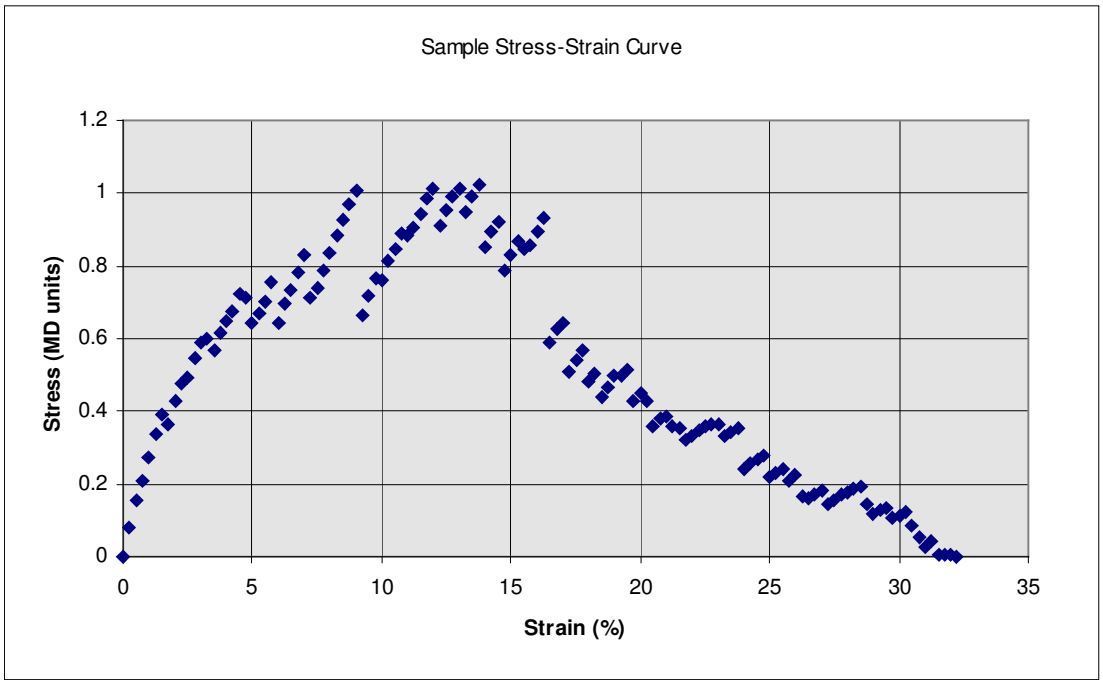

The graph begins relatively straight and then begins bending. As it moves higher, we notice a number of big drops in the stress of the system. Eventually, the stress peaks, stays there for a while, and then diminishes until it reaches 0. Why, you might wonder, does the stress drop suddenly a number of times throughout the procedure? Sharp stress drops occur at strains of 9.00% and 16.25%.

These drops have much to do with the stress that builds up in the system and the mechanism of plastic deformation of crystals, that is, the formation and propagation of defects known as dislocations. As the system is stretched, stress accumulates. Stress can be released when atoms reorganize themselves into a new arrangement in which the total potential energy of the system is lower. Necking is a highly-visible example of the manner in which metals deform as a result of stress. Necking is the results of plastic deformation that often precedes fracture. Ductile materials tend to deform more in this manner and thus can absorb more energy prior to failing than brittle metals.

Below are two pictures of the above sample, shown before and after the stretch which causes the drop in stress mentioned above:

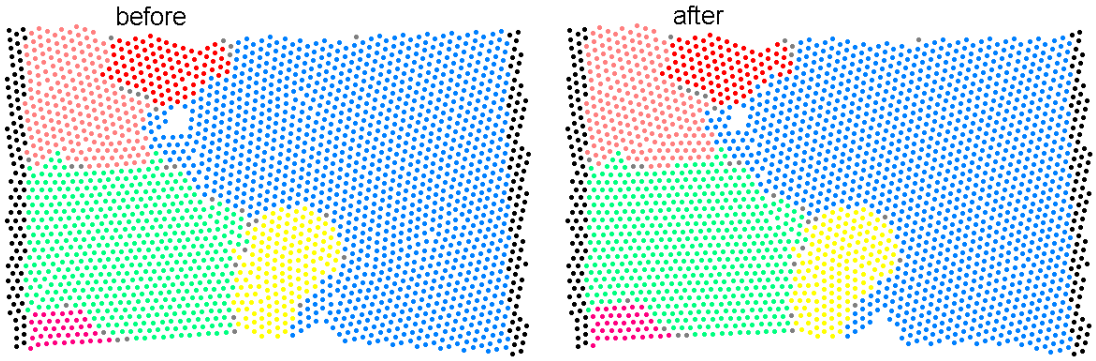

While it may be somewhat hard to see in the pictures above, the system here deforms significantly from the first picture to the second. Two changes might be noticeable when looking closely: 1) notice the tear at the bottom of the sample, just left of the center, increases significantly; 2) notice the bottom left part of the sample moves up slightly from the first picture to the second. This caving inward of the atoms is what, on a macroscopic level, appears as necking.

A much clearer example of the same phenomena can be seen in a much smaller sample of 8 x 20 atoms. Look at the following stress-strain graph:

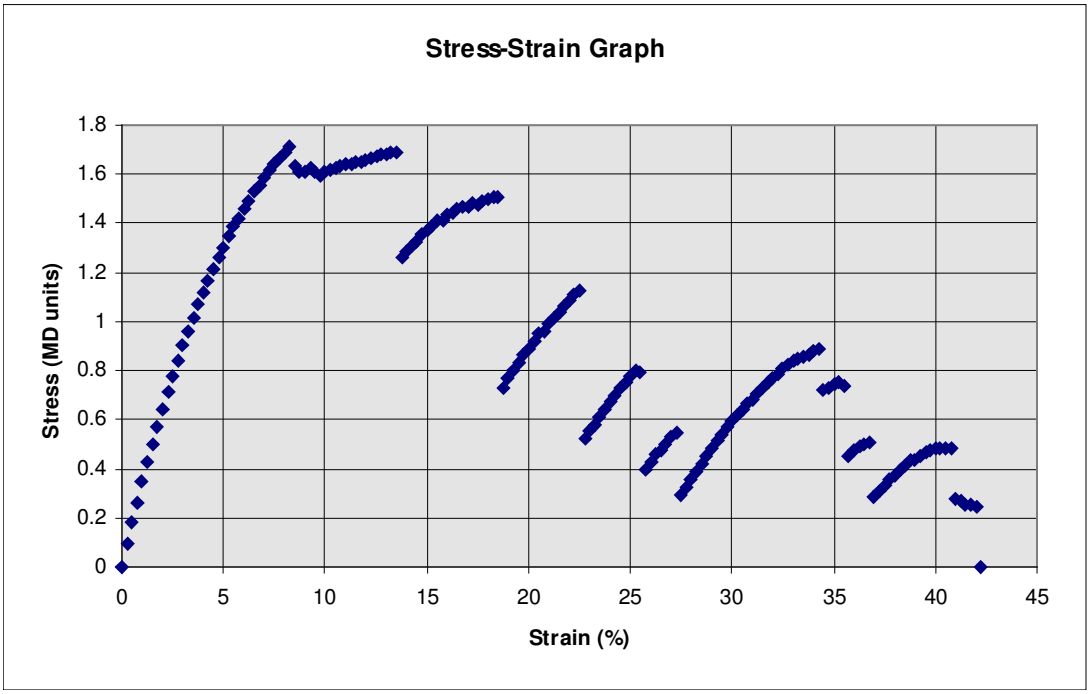

Here the large drops are much more visible than they were in the previous example. Consider the two pictures to the right. Notice that although the atoms are considerably further apart one from another in the second picture, the overall topography of the atoms stays the same. No significant changes have occurred; neither sliding nor dislocation formation or migration occurs to this very large strain 13.5%.

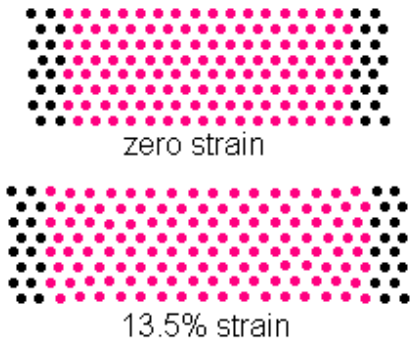

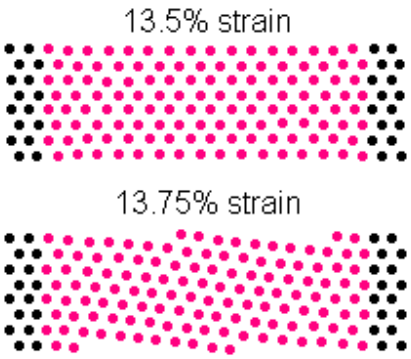

Consider, on the other hand, the two pictures on the left. After one more increment in the stretch, the accumulated stress cannot be accommodated through the minor movements of the atoms and an entire chunk of atoms move together—entire planes slip! Barely noticeable from the pictures is that the atoms on the whole are closer together after than they are before. This accounts for the large drop in stress at this point on the graph on the previous page.

Other drops in the system stress are explained in the same way. The three pictures to the right show the same sample at three various strain rates. The difference between the first two pictures is minimal and hardly noticeable. The two pictures mark the beginning and end of the second segment on the graph on the page, which extends from 13.75 to 18.75%. The third picture is noticeably different from the first two, and significant slip can be easily observed. The stress that had built up over the gradual straining finally was too much for the system to handle, and eventually slips occurred. When the grains are smaller, relative to the system, then stress typically does not have the same chance to build up, as the smaller grains can more easily change position and orientation than can the larger grains.

The drops also become smaller with an increasing system size. Note how much smaller the jumps are in the first example, a system of 2400 atoms, than they are in the second example, a system of only 160 atoms.

Grain Size, Orientation, Ultimate Tensile Strength, and Young's Modulus

5.1 Background

In Chapter 3 we examined how cooling rates affect the average size of grains when they form during solidification. Many important properties of metals depend on the size of their constituent grains, as well as on their orientation, another property we will discuss in this chapter.

When building structures, engineers must know how much stress their materials can endure before failing. This is known as the ultimate tensile strength, explained briefly in the preceding chapter. Another important property is the degree to which the material deforms under certain loads. One important constant associated with every material is the Young's modulus (also known as the modulus of elasticity). In this chapter we explore how grain size and orientation impact these two properties.

5.2 Terms and Concepts

As we explained briefly in Chapter 3, when a load is first applied to a material, the relationship between the stress on a material and its strain, or physical deformation, is linear. The slope of the line describing this relationship is known as ***Young's modulus***. After the load is removed, the material will return to its original shape.

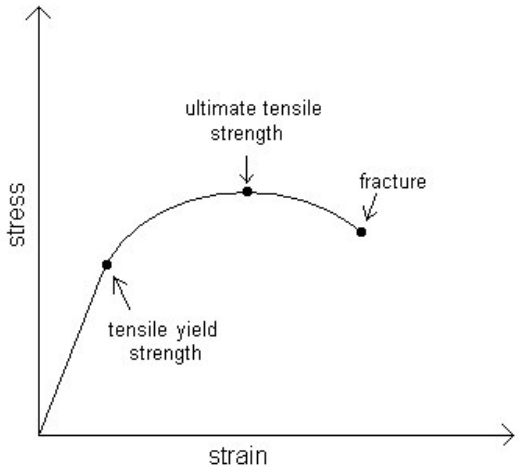

Above a certain stress, the material no longer behaves linearly and the material yields; this stress is called the ***tensile yield strength***. This property is taken to be the force at that point divided by the cross-sectional area of the material, with units such as N/m^2.

After reaching the yield tensile strength, a material can generally endure more stress before it fractures. The maximum stress it endures before fracturing is known as the ***ultimate tensile strength***. This property is taken to be the force at that point divided by the cross-sectional area of the material. It is thus measured in units of force per unit area. If the material was loaded by an increasing stress, it would fail at the ultimate tensile strength. If it were loaded by slowly

increasing the strain, it is possible to increase the strain beyond this point (with a corresponding drop in the macrostress). The end of the graph illustrates the point of apparent fracture (load controlled experiment).

The issue of ***grain sizes*** was already discussed in Chapter 3, and specifically in section 3.4.4. We will not repeat that information here.

Some properties of grains, as we will shortly see, depend on the orientation of the crystal lattice with respect to the applied load. Throughout this paper we will call this property ***grain orientation***. The two sample cuts to the right are identical in size and differ only in the orientation of the atoms (relative to the loading axis). We leave for another time the task of quantifying the orientation of a grain.

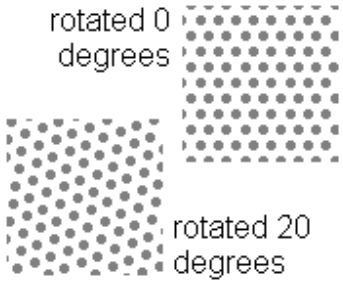

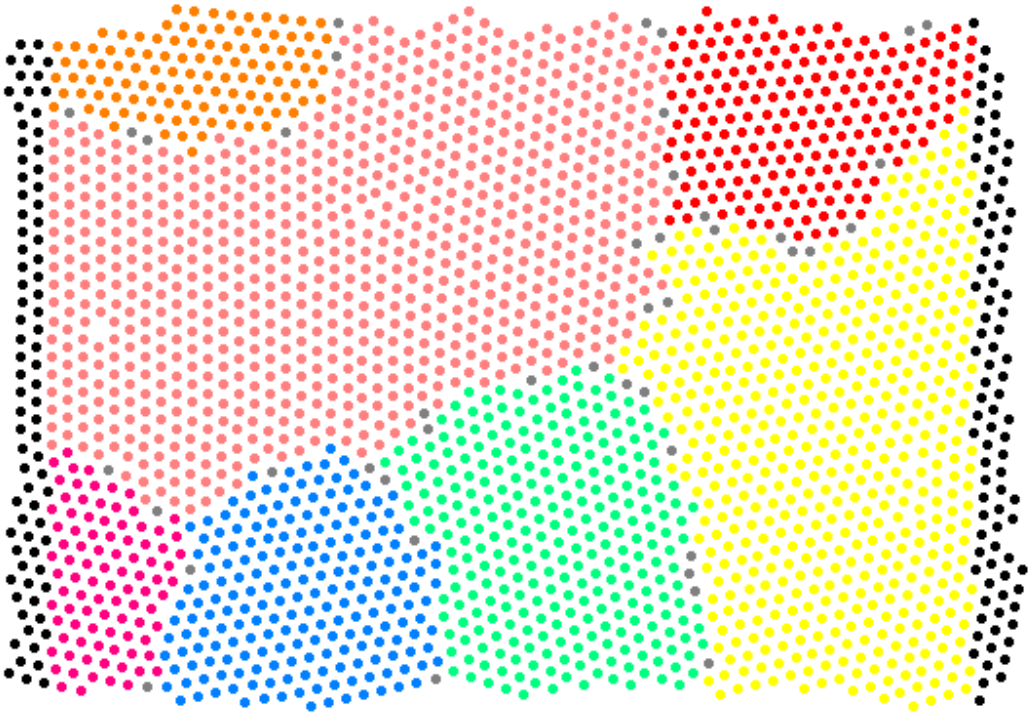

When materials solidify and grains form, their orientations may be random. Some grains are orientated one way, some another, all in the same material. The picture above shows a sample composed of seven grains, each with a different orientation. The way in which these grains are located and distributed affect many properties of the material. We limit ourselves in this section to looking at a very few simple cases.

5.3 Survey of the Literature

Since the early 1950's it has been known that a metal's strength is inversely proportional to the size of its grains. This relationship is known as the Hall-Petch relationship, named after the work the two researchers who first wrote on the topic in [11] and [20]. This relationship is often written:

$$\sigma_v = \sigma_{y0} + K_y d^{-1/2},$$

where σ_v is the yield strength of a material, d is the average grain size, and σ_{y0} and K_y are temperature-dependent constants. Further research has confirmed the validity of this relationship. See [12] an example of such work.

However, in the last fifteen or so years, much research has been published which alleges that this relationship fails, or is reversed, for very small grains, in the nanometer range. See [4], [18], and [23] for three examples of such experimental data. Computer simulations showed similar results; see [22].

The graph to the right is taken from [28]. The diagram suggests that grains larger than some critical value d_c follow the Hall-Petch relationship , while smaller grain exhibit an opposite relationship.

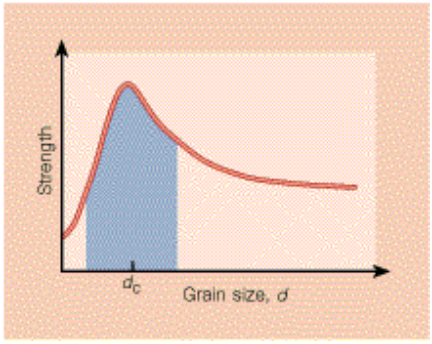

Subsequent research attempted to attribute this phenomenon to a number of factors, including the large fraction of atoms at grain boundaries [22], dislocation spacing in pileups [19], and an increased creep rate at the nanometer range of grains [4], [10].

However, many of these theories have not held up to close scrutiny and have been rejected for various reasons [13], [18], and [24]. It seems fair to say that, to date, no satisfactory explanation has been offered to explain this peculiar phenomenon.

On the relationship between average grain size and Young's modulus, I have not yet seen any published work. Although much data is available on the orientation dependence of Young's modulus and other elastic constants, not much is available about the yield strength.

I have not either seen much discussion of the relationship between the grain orientation and tensile strength and Young's modulus.

5.4 Simulating Experiments

5.4.1 Experiment 3 – Average Grain Size and Ultimate Tensile Strength

In this experiment, we look at the ultimate tensile strength of our materials as a function of average grain size. The samples for this experiment were prepared as explained in Chapter 3 and stretched as explained in Chapter 4. For this experiment, we used the samples already created in Experiment 2 (Section 3.5.2). These samples were 25 by 50 in size (1250 atoms in total). After each sample had been relaxed, it was stretched and a stress-strain curve for the stretching was recorded. Using this information, we were able to measure the maximum stress that the sample endured during the deformation.

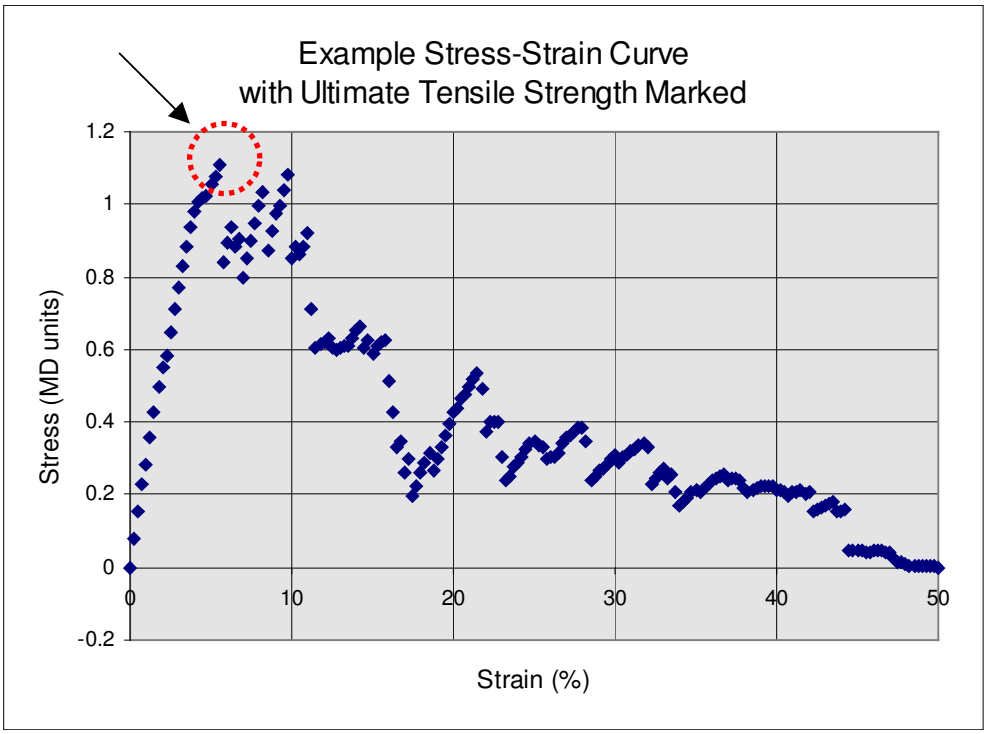

In this graph, the maximum stress is identified with the ultimate tensile strength. For each sample that we stretched, we monitored the average grain size (before applied stress) and the maximum stress. Results from our experiments are reported below.

These are stress-strain curves for samples with the largest and smallest grain sizes:

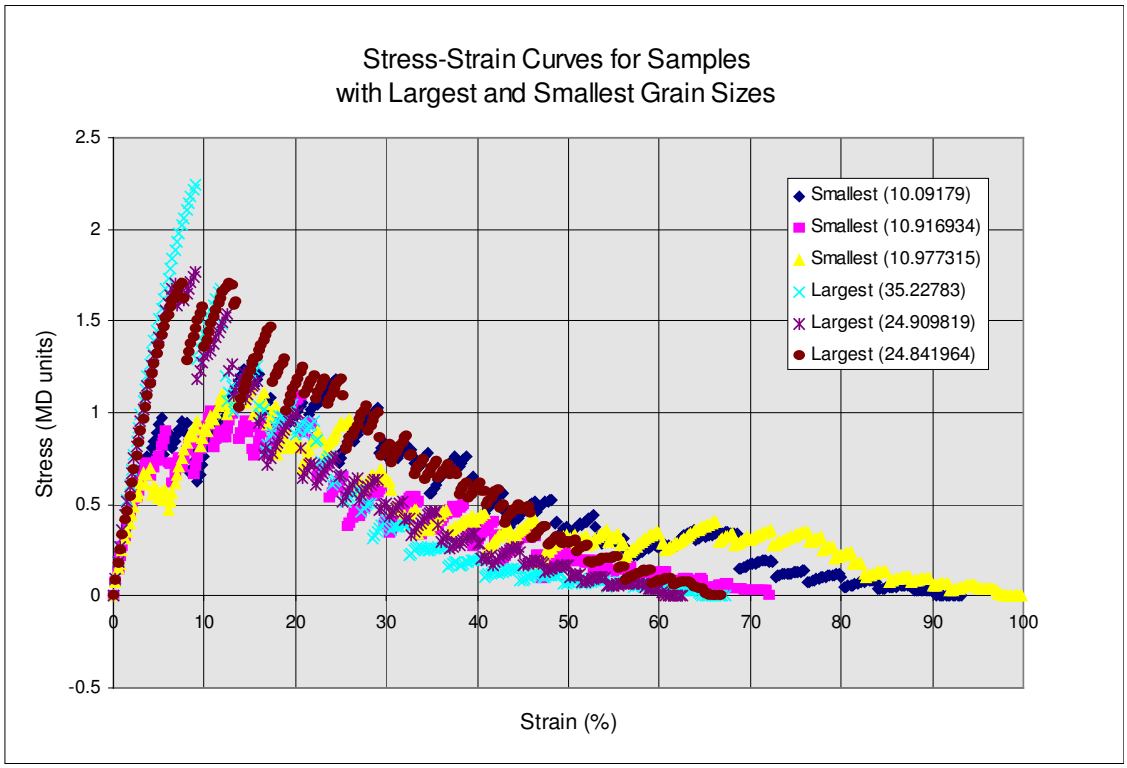

It is evident that the samples with larger grains reach higher maximum stresses than samples with smaller grains. The following graph demonstrates the maximum stress of various samples as a function of the average grain sizes of the sample:

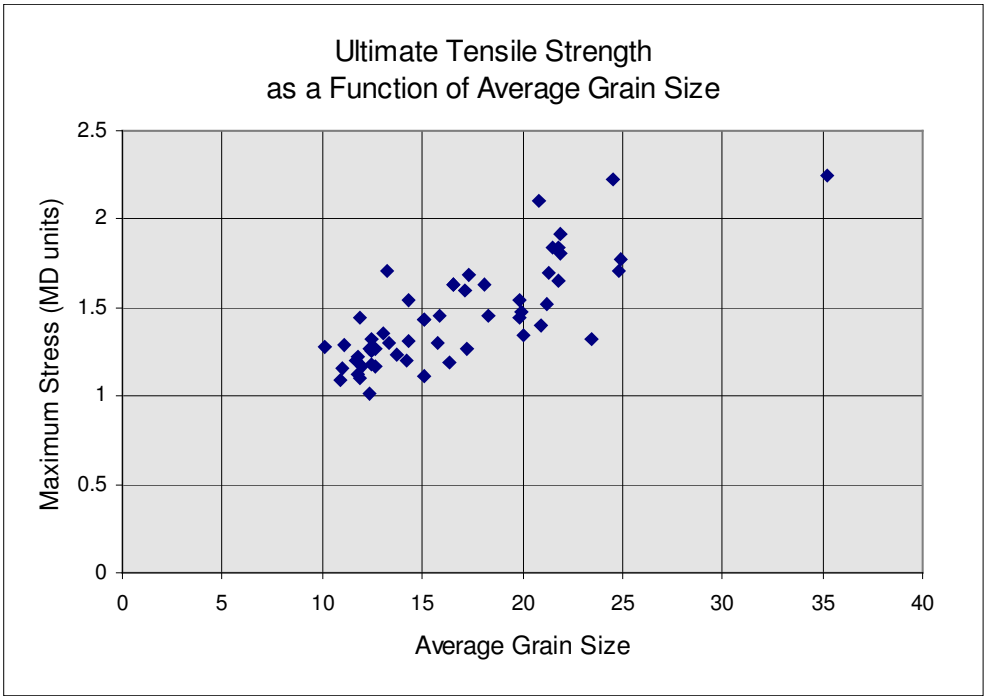

As is apparent from the graph, the bulk of the data was collected from samples with small average grain sizes. We could not fully control the average grain sizes of the samples—they are merely the results of cooling samples at various temperatures and taking the resulting samples.

While it is difficult to establish the reason for the relationship between grain size and ultimate tensile strength, the plot does show a clear trend: larger grain sizes mean stronger materials.

In the above graph you can see that only three samples exhibited maximum stresses greater than 2. Below are pictures of the initial configuration of these three samples:

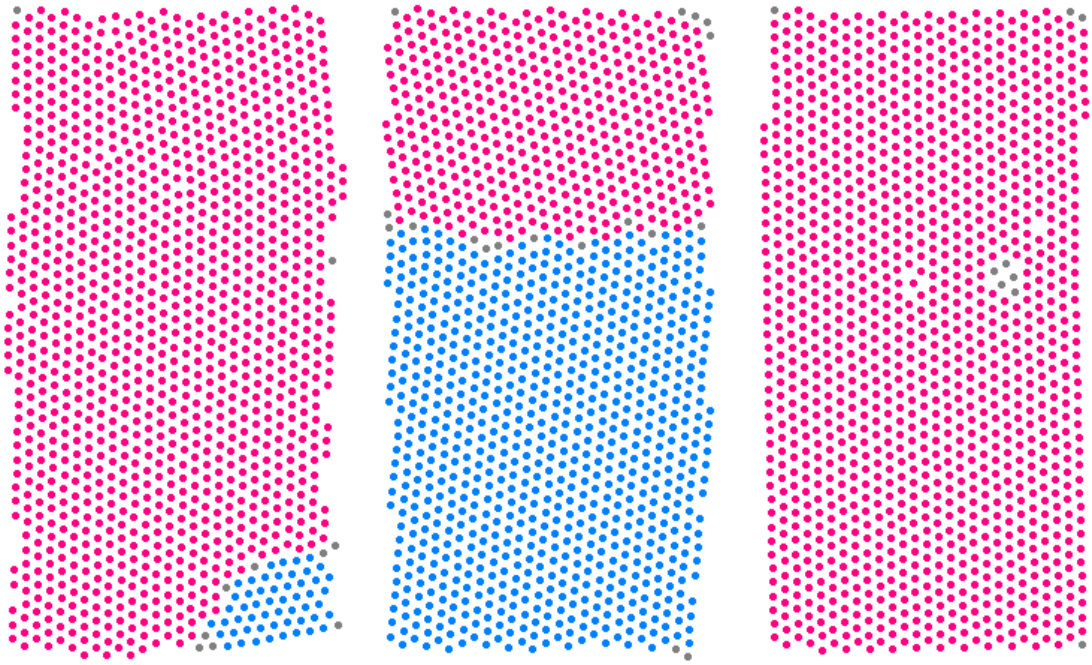

Here are pictures of the samples with the three lowest UTS:

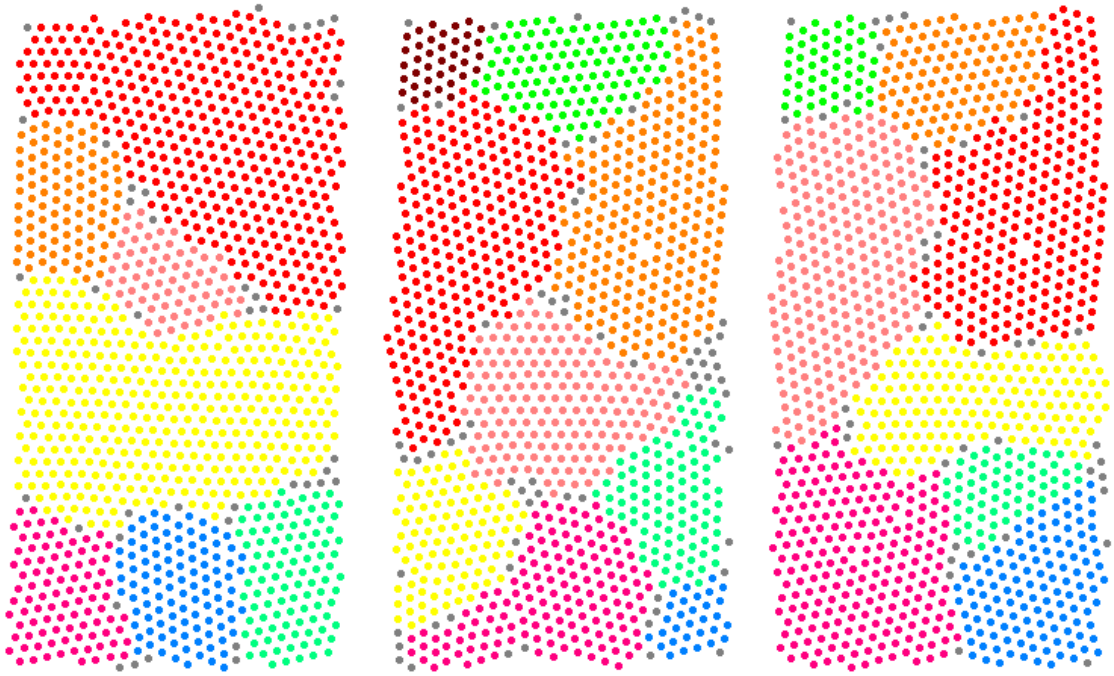

The evidence thus supports the thesis that ultimate strength increases with increasing average grain size. Further data relevant to this experiment and Experiment 4 can be found in the Appendix.

5.4.2 Experiment 4 – Average Grain Size and Young's Modulus

For this experiment we used the same data collected in Experiment 2 and used in Experiment 3. We calculated Young's modulus for each sample by measuring the slope of the graph over the first 0.5% strain. Over this range of strain, or negligible, dislocations or planar slips that resulted from the strain.

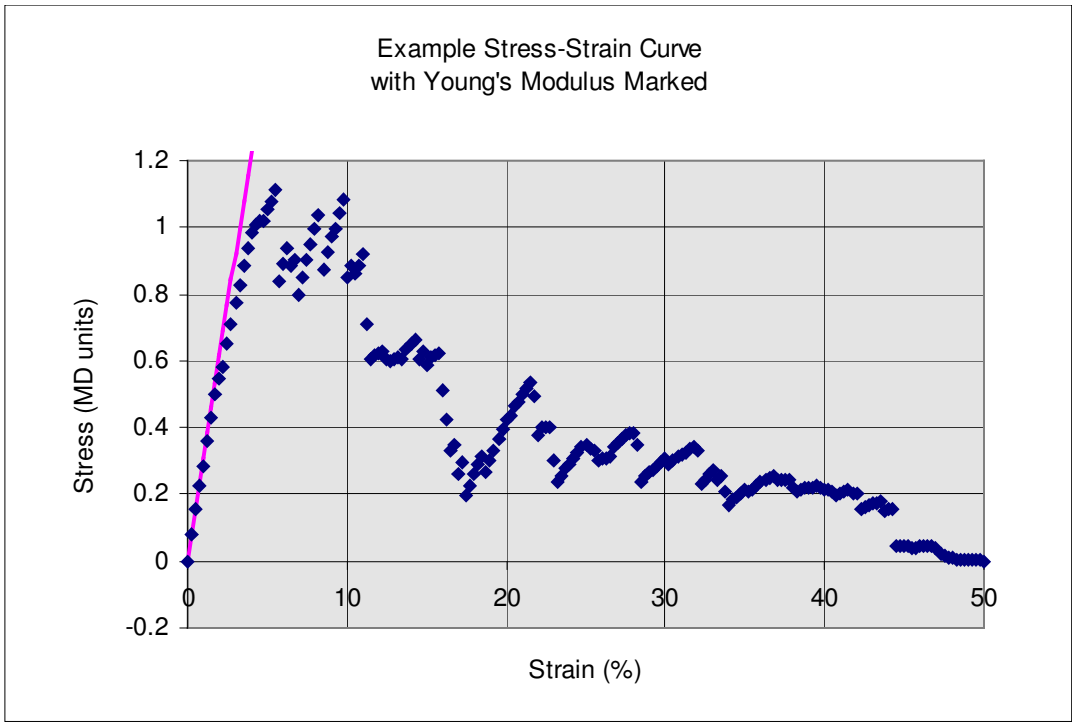

Below is a graph of Young's modulus for samples of various grain sizes:

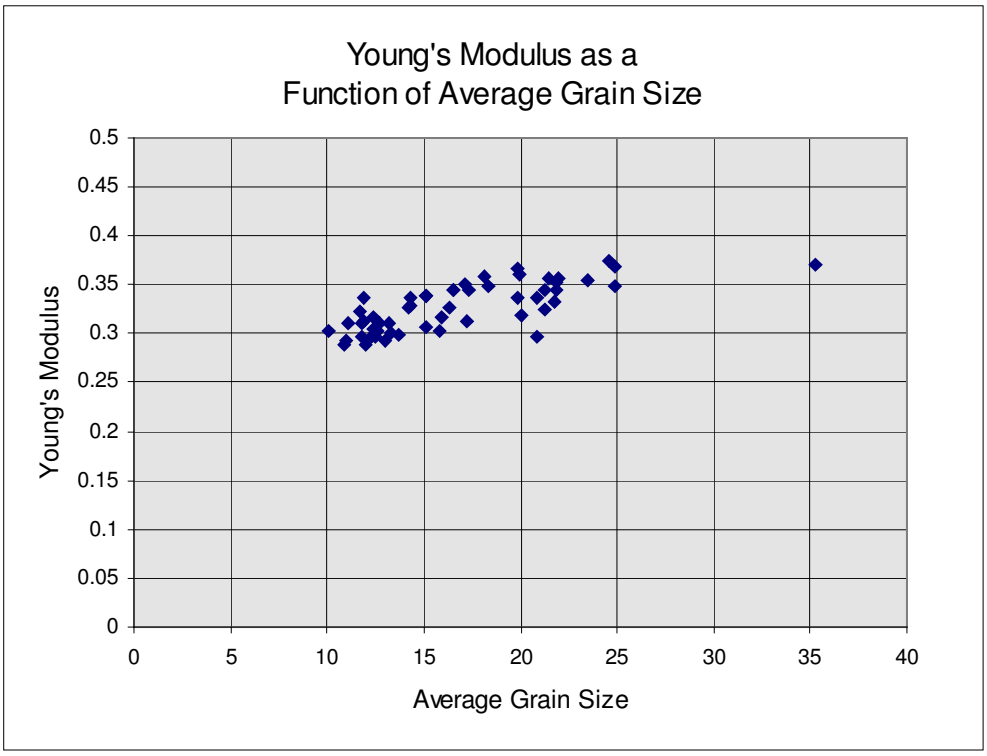

As we noted with regard to UTS, while a precise relationship between Young's modulus and grain size cannot be determined from his limited data set, it is relatively easy to see the general trend—larger grain sizes mean higher Young's modulus, although this dependence appears to be fairly weak.

Here are pictures of the initial configurations of the three samples with the highest moduli of elasticity:

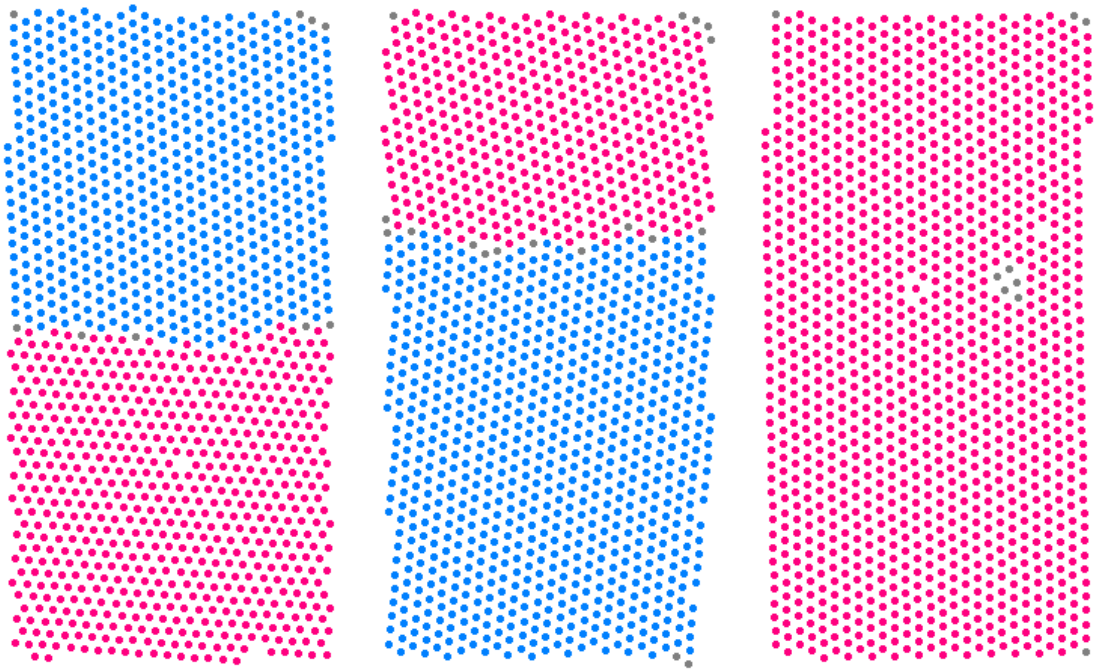

As you can readily see, two of the three samples are the same as those which exhibited maximal UTS. Below are the samples with the three lowest moduli of elasticity:

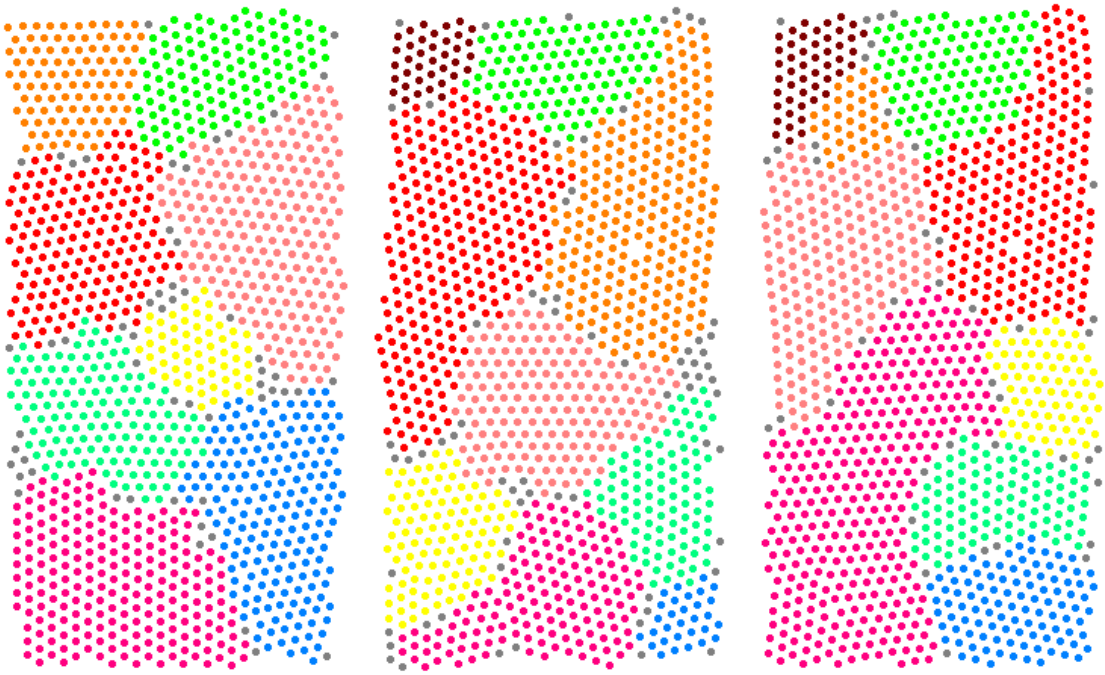

5.4.3 Experiment 5 – Grain Orientation and Ultimate Strength

One of the most important factors that influences the maximum stress a material can sustain is the orientation of its grains. In this section, we experiment with single-crystal samples fixed at various orientations, and examine the relationship between grain orientation and ultimate tensile stress.

Our first step in this experiment was preparing the samples. To do so, we first prepared a large circular mono-crystal sample as shown below:

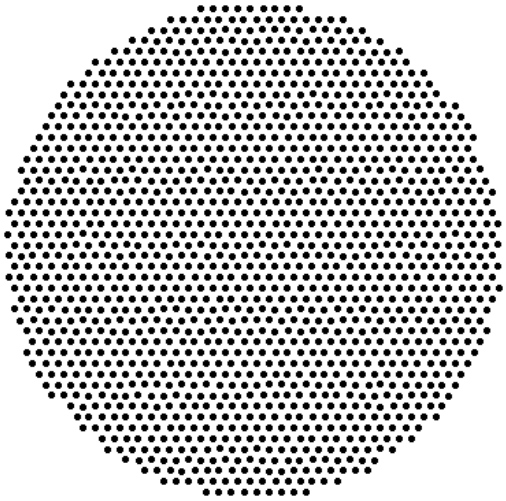

We then rotated (counterclockwise) the circular sample anywhere between 0 and 60^o in increments of 1^o. After rotating the larger sample, we cut smaller samples from the middle of the circle as shown:

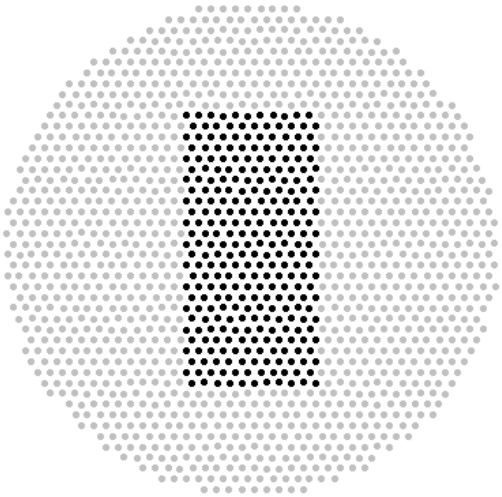

Here are a few examples of samples from circles rotated to various orientations:

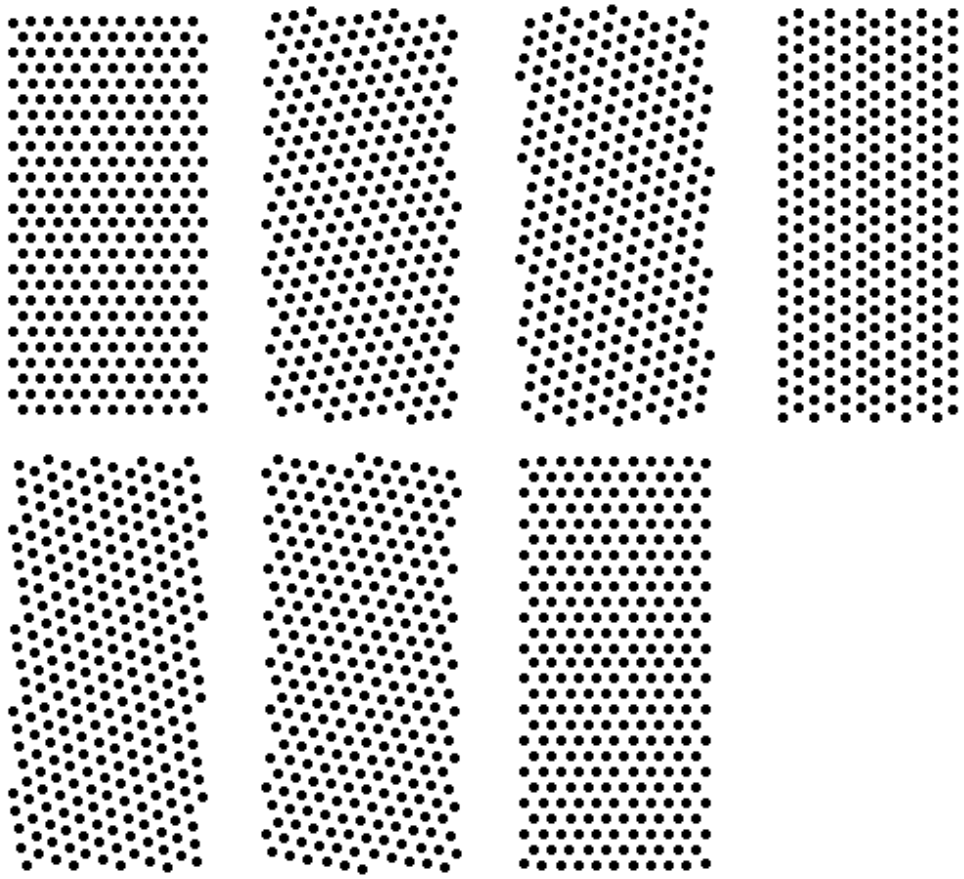

In order from right to left on the first row and then on the second row, we have mono-crystal samples at 0^o, 10^o, 20^o, 30^o, 40^o, 50^o, and 60^o. As is clear from the pictures and the symmetry of the lattices, the rotations are periodic, and 0^o and 60^o are identical.

We then strained the samples as in prior experiments and as described in Section 4.3.2. We measured various properties of the stretching procedure and recorded our results.

Below are stress-strain curves from samples at various orientations:

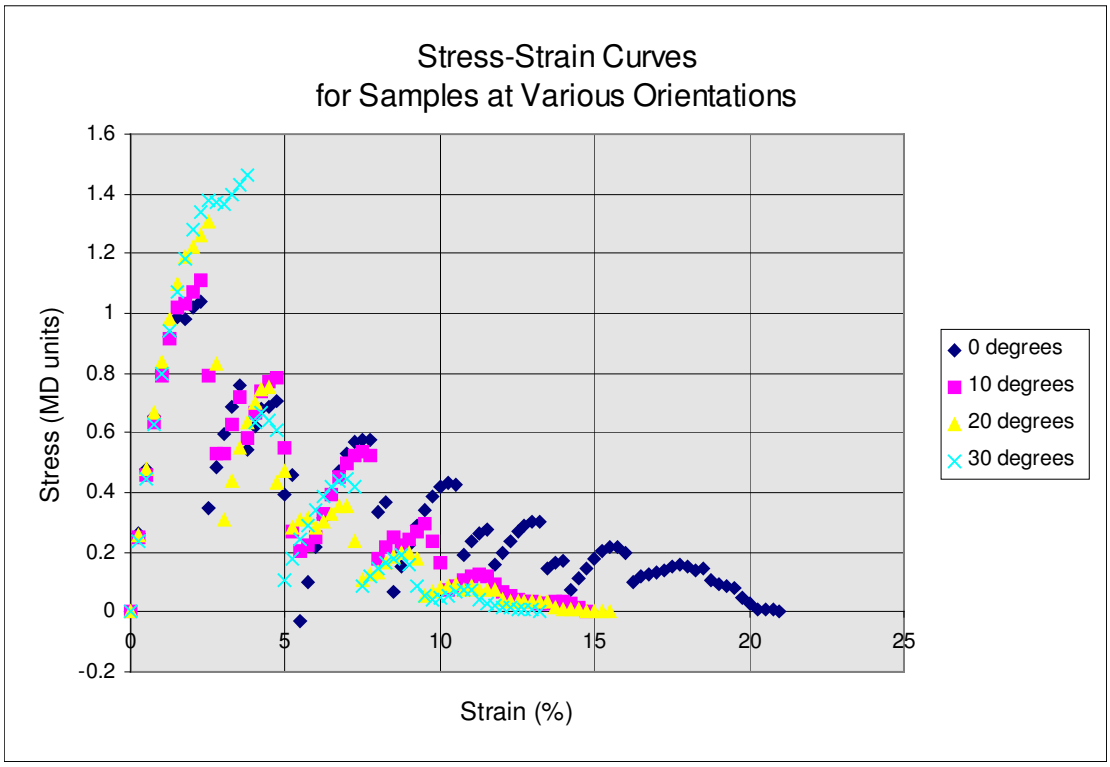

It is rather clear from the graph that the ultimate tensile strength increases with rotation angle from zero to thirty degrees and that the Young's modulus is nearly invariant with respect to rotation angle.

Here is a graph of the ultimate-stresses measured for each rotation of varying degrees:

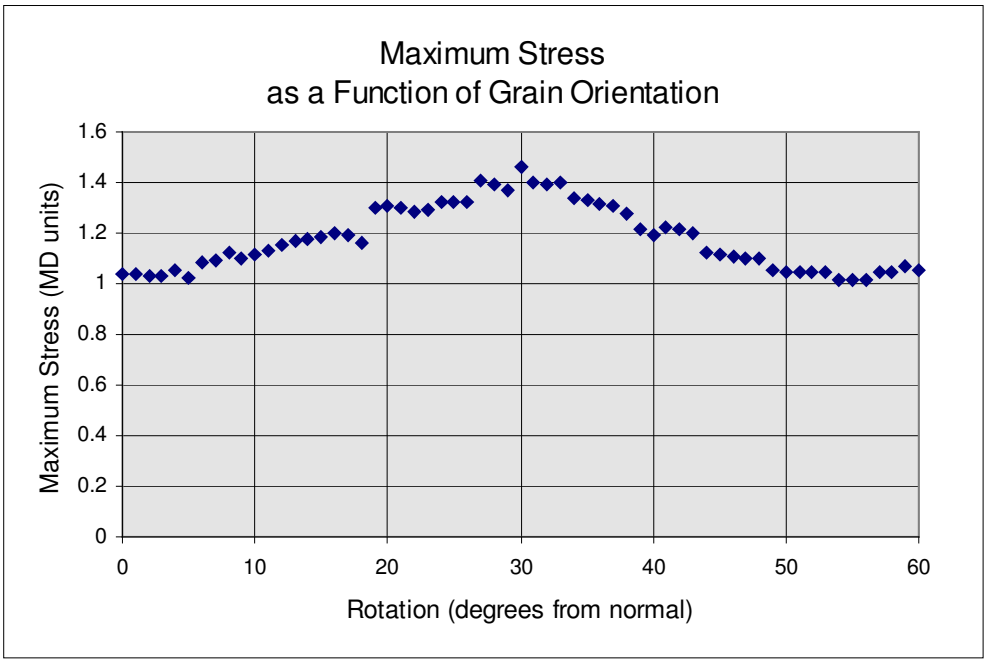

Below is a graph of Young's modulus for the various orientations. This value is calculated as it was in Experiment 4 (5.4.2).

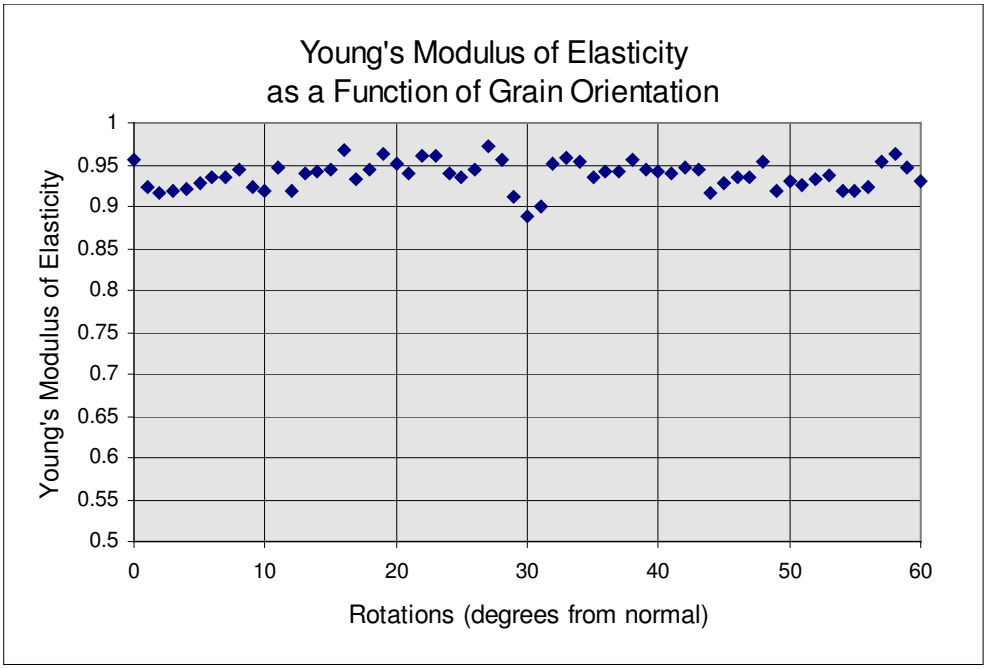

5.5 Observations and Analysis

5.5.1 Observations

Experiment 3 relates a decrease in tensile ultimate strength to a decrease in average grain size at the atomic range. Although the exact relationship between these two properties is unclear from our limited work, and although it seems clear that this relationship is probably not linear, that such a relationship exists seems beyond doubt.

5.5.2 Analysis

The reason for this relationship may be explained as follows. Stress applied to a grain can be 'absorbed' by the grain in one of two ways: (1) the atoms inside the grain can rearrange internally or (2) the grain as a whole can rotate or relocate. When a grain absorbs the stress by changing internally, the stress is distributed amongst the constituent atoms and each relocates slightly to absorb a small share of the stress. When a grain absorbs stress by rotating or relocating, its internal atoms remain stationary with respect to one another; the stress is entirely absorbed by changes in the interface between the grain and its neighbors.

When stress is applied to a sample, the stress will most likely be absorbed in the manner that requires the least rearranging of the atoms, minimal work. It seems to me that slightly moving every atom within a grain requires less work than rotating or relocating the entire grain. Grains will thus generally prefer internal rearrangement to rotating and relocating. However, internal rearrangements are not always possible and sometimes the grain as a whole must either reorient or relocate itself.

In both of the aforementioned stress-absorptions processes, a single large grain will tend to absorb more stress than will multiple smaller ones. If the stress is absorbed by rotating and relocating, it seems more difficult to rotate or relocate one large grain than it is to rotate or relocate multiple small ones. Moving a large rock can be substantially more difficult than moving numerous smaller rocks.

If the stress is absorbed by internal change, where stress is distributed amongst constituent atoms, then possessing more atoms amongst which to distribute the stress will allow for a larger total stress to be placed on the grain.

Our research does not address the relationship between average grain size and tensile yield strength, the subject of the Hall-Petch relationship. Our research does not either address recent debates concerning the relevance of the Hall-Petch relationship at this level and how to reconcile what we see at this level with what we see with larger grains.

Further research is necessary to explore a more exact relationship between average grain size and ultimate tensile strength; explore the relationship between average grain size and tensile yield stress; explain why the Hall-Petch relationship is observed in larger grains but not in smaller ones.

Experiment 4 indicates that an increase in average grain size corresponds with an increase in Young's modulus. The reason for this is probably similar to what has been mentioned above. When a single grain absorbs stress, the stress is distributed amongst the constituent atoms. The larger the grain, the more stress the entire grain can endure before needing to rotate, relocate, or fracture.

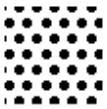

Experiment 5 indicates that when grains are oriented as in the picture on top (0 degrees), the strength of the material is minimal, while when the grain is oriented as in the picture on bottom (30 degrees), the strength of material is maximal. This has much to do with the way in which slips occur in the systems. In the sample pictured on bottom, the tensile stress is entirely perpendicular to the planes along which the system might slip. Enough stress is needed to rip along the entire vertical axis. In contrast, in the sample pictured on top, there exist slip planes not entirely perpendicular to the tensile stress. Slipping requires less stress than ripping through the entire piece.

In practice, when polycrystalline materials are used, there are probably so many grains at so many orientations, that these factors do not make a difference. This result, however, is important when single-crystal materials are produced and when the orientation can be controlled.

Young's modulus does not seem to depend in any way on the orientation of a grain. I don't know how to explain this.

Conclusions

6.1 Summary

In our work we have accomplished a number of important tasks:

- We have explained the challenge of grain identification and a number of inherent problems. We have also developed an algorithm to accomplish this task; as far as I know such an algorithm has not been yet developed. Such an algorithm is crucial for experimenting with 'naturally born' grains (as opposed to those created using Voronoi constructions).

- We have reproduced experimental data illustrating the influence of cooling rates on grain sizes. Such work has not yet been done using MDS. We have reproduced the general conclusions of previous research. It is shown that $\bar{d} \propto r^x$, where x is around –0.18, confirming experimental data.

- We have developed code to simulate stress-strain curves and to determine critical points along those curves, especially ultimate tensile strength.

- We have investigated the relationship between average grain size and ultimate tensile strength. We showed a correspondence between large grain sizes and higher UTS.

- We have investigated the relationship between average grain size and Young's modulus. We have shown that this value increases with grain size, perhaps linearly.

- We have investigated the relationship between grain orientation and ultimate tensile strength. We have shown that certain orientations are clearly associated with higher values of UTS and others with lower values.

- We have investigate the relationship between grain orientation and Young's modulus. We see no correlation between the two values.

6.2 Future Research

Future work interests include both improvements to the experiments we conducted as well as the exploration of new topics not touched upon in this research.

Improvements

- In this study, all experiments were conducted in two dimensions. We are interested in conducting similar research for material simulations in three dimensions.

- In this study, we used the Lennard-Jones potential to govern the interaction between all atoms. This indeed was the standard potential used for many initial MD simulations. However, over time the use of the Embedded Atom Method (EAM) has become very popular for modeling metals [5], [6]. This potential takes significantly more time to compute and thus we sufficed with the Lennard-Jones potential. In future work, we would like to explore how the change in potential affects the results.

- In this study, because of hardware limitations, we could only study samples of relatively small size. In future work, I would like to redesign the code such that it could be run on multiprocessor machines, thus allowing us to increase the possible sample sizes.

New Topics

- We did not investigate the relationship between average grain size and tensile yield stress. We have not either progressed in resolving recent controversy regarding the validity of the Hall-Petch relationship at various scales. Both of these topics I hope to investigate further at a later point in time.

- As explained earlier the stress-strain curves that we used here are 'engineering' curves and not 'true' curves. We could rewrite some of the code to determine the true stress and true strain.

- We have only studied tension so far. We would also like to apply the methods we have developed to studies of pressure and shear stresses as well.

Appendix

This appendix provides additional data for each of the five experiments. There we supply further details about how the experiment was run, including details deemed too technical to include in the paper, as well as further data collected but either not presented inside the paper, or not presented fully.

Experiment 1

Here we include data collected for each of the 5 trials of this experiment. First we provide a chart of the cooling rates from the various trials. Although we used the same scaling factor (between 0.71 and 0.99) for each trial, the samples varied slightly in the time needed for cooling. Here we show figures for the cooling rates, calculated by the function:

Cooling Rate = 1.0 / Time to Cool

	Cooling Rates (degrees per second)					
	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Average
71	0.111111	0.105263	0.105263	0.105263	0.105263	0.106433
72	0.105263	0.1	0.1	0.1	0.105263	0.102105
73	0.1	0.095238	0.095238	0.1	0.1	0.098095
74	0.1	0.095238	0.095238	0.095238	0.095238	0.09619
75	0.095238	0.090909	0.095238	0.090909	0.090909	0.092641
76	0.090909	0.086957	0.086957	0.086957	0.090909	0.088538
77	0.086957	0.086957	0.086957	0.086957	0.086957	0.086957
78	0.083333	0.076923	0.083333	0.08	0.08	0.080718
79	0.076923	0.076923	0.074074	0.08	0.071429	0.07587
80	0.071429	0.074074	0.076923	0.071429	0.071429	0.073057
81	0.071429	0.068966	0.074074	0.071429	0.066667	0.070513
82	0.068966	0.066667	0.068966	0.064516	0.066667	0.067156
83	0.0625	0.064516	0.064516	0.064516	0.064516	0.064113
84	0.058824	0.060606	0.0625	0.0625	0.060606	0.061007
85	0.057143	0.055556	0.057143	0.058824	0.055556	0.056844
86	0.04878	0.052632	0.054054	0.055556	0.052632	0.052731
87	0.051282	0.051282	0.04878	0.051282	0.05	0.050525
88	0.047619	0.044444	0.046512	0.047619	0.044444	0.046128
89	0.042553	0.043478	0.044444	0.042553	0.044444	0.043495
90	0.036364	0.039216	0.04	0.040816	0.040816	0.039442
91	0.037037	0.035088	0.033898	0.034483	0.036364	0.035374
92	0.031746	0.032258	0.03125	0.032258	0.031746	0.031852
93	0.025974	0.027027	0.028169	0.026316	0.027027	0.026903
94	0.025	0.024096	0.024096	0.02439	0.023256	0.024168
95	0.019229	0.021053	0.020619	0.020833	0.020619	0.02047
96	0.016948	0.016807	0.016393	0.016807	0.01626	0.016643
97	0.012738	0.012739	0.01227	0.012579	0.012579	0.012581
98	0.008	0.008511	0.008471	0.008547	0.008511	0.008408
99	0.003769	0.004311	0.004264	0.004313	0.004264	0.004184

We also show the average grain sizes for each of the various trials:

		Average Grain Size				
	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Average
71	14.62732	16.33784	15.17519	19.36442	17.18331	16.5376172
72	16.5892	19.41574	14.86254	14.8485	15.894383	16.3220712
73	20.98327	19.2189	15.63805	28.18103	17.26815	20.25788
74		27.46852	13.90845		13.506348	18.29443933
75	14.01696	19.40497	14.06386	18.78734	15.540289	16.362684
76	14.19266	17.81446		18.47232	13.528906	16.00208375
77	17.38449	19.41094	13.88832	18.83234	13.826116	16.6684402
78	13.44889	21.92666	14.11577	21.39168	15.228658	17.2223304
79		23.48725	15.36375	19.35759	18.004414	19.053251
80	18.49113	23.99001	13.84888	21.57495	14.361373	18.453267
81	13.7188	18.30396	12.98886	18.14652	14.502331	15.5320946
82	18.1833	18.43904	15.01207	19.81726	14.675684	17.225471
83	15.97456	23.90719	13.73248	19.61648	20.412798	18.7286992
84	23.14575	23.71337	14.23657	15.91583		19.2528775
85	16.90018	23.34457	14.08089	15.70073	17.107587	17.4267934
86	19.51613	28.04999	15.24552	19.25508	17.228789	19.859101
87	16.03843	20.03599	19.28927	21.30172	13.227421	17.9785678
88	16.2729	22.5976	15.63382	21.54201	15.462401	18.3017462
89	19.80453	27.2256	12.4898	17.55511	14.7896	18.3729266
90	24.70591	26.83557	17.24167	17.399	15.764038	20.3892368
91	22.22979	21.36036	16.99349	23.30695	16.575304	20.09318
92	21.4299	20.7988	15.59285	19.64623	28.177292	21.1290144
93	22.52358	22.58589	15.50561	17.16349	21.165846	19.7888832
94	20.22696	21.77952	19.99417	19.75795	26.429959	21.63771
95	22.75406	27.93534	18.28563	20.47299	20.841997	22.0580014
96	20.90477	22.14941	18.52274	20.28825	16.062011	19.5854376
97	19.46183	26.36667	20.57469	21.45219	21.571998	21.885474
98	49.94998	23.13038	20.13862	34.57912	22.373692	30.0343552
99	49.96999	21.87709	49.94998	21.7215	22.511757	33.2060622

Occasionally, experiments crashed for technical and so no information is available for them. Thos boxes are left blank and are ignored in calculated averages.

Experiment 2

Here we include data collected for each of the 5 trials of this experiment:

	Cooling Rates (degrees per second)					
	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Average
5	0.666667	0.666667	0.666667	0.666667	0.666667	0.666667
10	0.5	0.5	0.5	0.5	0.5	0.5
15	0.4	0.4	0.4	0.4	0.4	0.4
20	0.333333	0.333333	0.4	0.333333	0.4	0.36
25	0.333333	0.333333	0.333333	0.333333	0.333333	0.333333
30	0.285714	0.285714	0.285714	0.285714	0.285714	0.285714
35	0.25	0.25	0.25	0.25	0.25	0.25
40	0.222222	0.222222	0.25	0.222222	0.25	0.233333
45	0.2	0.2	0.222222	0.2	0.222222	0.208889
50	0.181818	0.181818	0.2	0.181818	0.2	0.189091
55	0.166667	0.153846	0.166667	0.166667	0.181818	0.167133
60	0.153846	0.133333	0.153846	0.142857	0.153846	0.147546
65	0.125	0.125	0.117647	0.133333	0.133333	0.126863
70	0.111111	0.111111	0.095238	0.105263	0.105263	0.105597
75	0.090909	0.090909	0.090909	0.090909	0.095238	0.091775
80	0.076923	0.074074	0.068966	0.076923	0.068966	0.07317
85	0.057143	0.04	0.055556	0.04878	0.05	0.050296
90	0.037037	0.033898	0.035088	0.039216	0.035088	0.036065
95	0.020833	0.02	0.020619	0.020833	0.019233	0.020304

		Average Grain Size				
	Trial 1	Trial 2	Trial 3	Trial 4	Trial 5	Average
5	19.81904	16.53859	15.10053	16.63517	18.03233	17.22513
10	19.81456	16.53859	15.10053	16.63517	18.03233	17.22424
15	14.17316	17.33506	10.91693	14.54381	18.04688	15.00317
20	12.42802	14.30136	14.75127	14.49592	15.70748	14.33681
25	11.99175	15.90233	13.69493	14.39483	15.47186	14.29114
30		13.04352	10.97732	17.21135	18.15114	14.84583
35	11.77582	20.86649	11.79752	13.81676	13.96635	14.44459
40	15.08452	21.81378	11.86655	16.42685	19.79672	16.99768
45	12.63181	17.1364	12.46613	16.67884	19.84826	15.75229
50	20.08202	21.80946	11.06875	12.56024	19.80192	17.06448
55	16.33568	21.46949	10.09179	17.12091	19.75488	16.95455
60	18.32713	21.27862	14.30954	15.47784	13.4164	16.56191
65	15.77844	20.89094	11.85538	16.64083	11.59714	15.35255
70	12.64852	21.86867	12.50029	19.42914	19.49623	17.18857
75	12.39592	21.92377	24.84196	11.61145	16.03948	17.36252
80	12.34012	17.19769	13.19823	19.91975	18.79201	16.28956
85	21.234	11.66739	24.90982	16.12034	24.80972	19.74825
90	18.08591	24.54392	19.92484	24.8644	19.65017	21.41385
95	23.47672	35.22783	19.81645	15.57882	15.51866	21.9237

Experiments 3 and 4

The results displayed in the paper for the two experiments originated in the same experiments. We do not provide here the stress-strain data from the more than 50 trials (3 trials for each of 19 different cooling rates), each of which takes hundreds of lines. The interested reader should contact the author for this information.

Here we provide some data from the experiments related to the Maximum Stress and to Young's Modulus, and their relation to Average Grain Size.

Grain Size	Maximum Stress	Youngs Modulus
10.09179	1.279136684	0.302586108
10.916934	1.08725245	0.287922322
10.977315	1.155425301	0.291884869
11.068752	1.289591084	0.310181876
11.667393	1.199705571	0.322570596
11.77582	1.218563899	0.296229103
11.797518	1.128524664	0.31003947
11.855381	1.43852408	0.337044442
11.866545	1.096370258	0.312741223
11.991749	1.168658741	0.288468929
12.340119	1.010707105	0.296625992
12.39592	1.26912417	0.305498931
12.428024	1.174752464	0.317325896
12.466131	1.257431699	0.300638901
12.500292	1.320024731	0.296740824
12.631808	1.166847129	0.302043557
12.648523	1.266671961	0.310250336
13.04352	1.349423782	0.292998207
13.198227	1.707891737	0.310310636
13.311108	1.296579467	0.300966507
13.694927	1.234213435	0.299241815
14.173164	1.200166315	0.326114475
14.301357	1.312488458	0.335735121
14.309539	1.539357859	0.328049276
15.08452	1.109619269	0.30704925
15.100533	1.427395097	0.338056935
15.100533	1.427395141	0.338056939
15.778443	1.304924552	0.302334651

Grain Size	Maximum Stress	Youngs Modulus
15.902334	1.45326366	0.317597169
16.335678	1.187032471	0.327536552
16.538591	1.625059516	0.344837237
16.538591	1.625059504	0.344837238
17.136402	1.595132661	0.351285722
17.197689	1.270315478	0.313348998
17.335056	1.684411945	0.344052552
18.085907	1.634413806	0.358124178
18.327127	1.450877738	0.348696165
19.816454	1.541189058	0.36722014
19.819038	1.441737652	0.336983555
19.924842	1.480735426	0.359845623
20.082022	1.340195426	0.317765245
20.866488	2.107587082	0.337078915
20.89094	1.397087977	0.296624335
21.234003	1.516019943	0.343657215
21.278621	1.693397402	0.325411404
21.469488	1.835242688	0.356340909
21.809463	1.835482477	0.333183147
21.813775	1.650311281	0.345417901
21.868672	1.806475753	0.351863754
21.923773	1.914970908	0.355577967
23.476715	1.323871181	0.354505318
24.543921	2.225725728	0.374415257
24.841964	1.707060918	0.347976044
24.909819	1.768204126	0.368358578
35.22783	2.246681714	0.370800974

Experiment 5

All relevant data is provided in the paper.

Bibliography

[0] Allen, Michael P., "Introduction to Molecular Dynamics Simulation," in Computational Soft Matter: from Synthetic Polymers to Proteins pp. 1-27, Editors: Attig N, Binder H, Grubmuller H, Kremer K (2004).

[1] Asaro, Robert J., "Micromechanics of Crystals and Polycrystals", Division of Engineering, Brown University, 1982.

[2] ASTM Standard E8-03, "Standard Test Methods for Tension Testing of Metallic Materials," Annual Book of ASTM Standards, Vol. 03.01, 2003.

[3] Boswell, P.G., and Chadwick, G.A., "The Grain Size of Splat-quenched Alloys," Scripta Metallurgica, Vol. 11, pp. 459-465, 1977.

[4] Chokshi, A.H., Rosen, A., Karch, J., and Gleiter, H., "On the Validity of the Hall-Petch Relationship in Nanocrystalline Materials", Scr. Metall. Vol. 23, no. 10, pp. 1679-1684, October 1989.

[5] Daw, M.S., and Baskes, M.I., "Semiempirical, Quantum Mechanical Calculation of Hydrogen Embrittlement in Metals", Phys. Rev. Lett. 50, 1285 (1983).

[6] Daw, M.S., and Baskes, M.I., "Embedded-atom method: Derivation and application to impurities, surfaces, and other defects in metals", Phys. Rev. B 29, 6443 (1984).

[7] Ercolessi, Furio, "A molecular dynamics primer", International School for Advanced Studies, Trieste, Italy, June 1997.

[8] Farkas, D., Van Swygenhoven, H., and Derlet, P.M., "Intergranular fracture in nanocrystalline metals", Phys. Rev. B 66, 2002.

[10] Gleiter, H., "Nanocrystalline Materials", Progress in Materials Science, Vol. 33, p. 223, 1989.

[11] Hall, E. O., "The Deformation and Ageing of Mild Steel: III Discussion of Results", Proc. Phys. Soc. Lond. B 64, pp. 747-753, September 1951.

[11b] Hockney, R. W. and Eastwood, J. W. "Computer simulations using particles." Adam Hilger, Bristol, 1988.

[12] Jang, J.S.C., Koch, C.C., "The Hall-Petch Relationship in Nanocrystalline Iron Produced by Ball Milling", Scripta Metallurgica et Materialia. Vol. 24, no. 8, pp. 1599-1604, August 1990.

[13] Koch, C.C., Narayan, J., "The inverse Hall-Petch effect—fact or artifact?", Symposium B: Structure and Mechanical Properties—Theory and Computer Simulations vs. Experiment as held at the 2000 MRS Fall Meeting; Boston, MA; USA; 28-30 Nov. 2000. 2001.

[14] Krivtsov, A.M., "Molecular Dynamics Simulation of Impact Fracture", Meccanica, Vol. 38, Number 1, January 2003.

[15] Leinonen, J.I., "Effect of cooling rate on ferrite grain size in a new heat treatment", Proceedings of the 5th ASM Heat Treatment and Surface Engineering Conference in Europe, 2000, pp. 333-337.

[16] L'Ecuyer, P., "Random numbers for simulation", Communications of the ACM, Vol. 33, Issue 10, pp. 85-97, 1990.

[17] MacKay, David J.C., "Strain and Stress", Cambridge University, http://www.inference.phy.cam.ac.uk/teaching/dynamics/tex/strain.pdf, November 2001.

[18] Nieh, T.G., Wadsworth, J., "Hall-Petch Relation in Nanocrystalline Solids", Scripta Metallurgica et Materialia. Vol. 25, no. 4, pp. 955-958, April 1991.

[19] Nieman, G.W., Weertman, J.R., and Siegel, R.W., Scripta Metallurgica et Materialia, Vol. 24, 145 (1990).

[20] Petch, N.J., "The cleavage strength of polycrystals", Journal of the Iron and Steel Institute, 174, pp. 25-28, 1953.

[21] Rapaport, D.C., *The Art of Molecular Dynamics Simulation*. Cambridge University Press, May 1997.

[22] Schiøtz, J., Di Tolla, F.D., and Jacobsen, K.W., "Softening of nanocrystalline metals at very small grain sizes", Nature 391, 561-563 (1998).

[23] Siegel, R.W. and Fougere, G.E., *Nanophase Materials: Synthesis - Properties – Applications*. (eds. Hadjipanayis, G. C. & Siegel, R. W.) pp. 233-261 (NATO-ASI Ser. E, Vol. 260, Kluwer, Dordrecht, 1994).

[24] Suryanarayana, C., Mukhopadhyay, D., Patankar, S.N., and Froes, F.H., "Grain Size Effects in Nanocrystalline Materials", Journal of Material Research, 7, pp. 2114-2117, 1992.

[25] Umemoto, M., Guo, Z.H., and Tamura, I., "Effect of cooling rate on grain size of ferrite in a carbon steel", Materials Science and Technology, April 1987, pp. 249-255.

[26] Van Swygenhoven and A. Caro, "Plastic behavior of nanophase Ni: a molecular dynamics computer simulation", Appl. Phys. Lett. 71, p. 1652, 1997.

[27] Verlet, L., "Computer Experiments on Classical Fluids. I. Thermodynamical Properties of Lennard-Jones Molecules", Physical Review, 159, pp. 98–103, July, 1967.

[27b] Verlet, L., "Computer experiments on classical fluids. II. Equilibrium correlation functions." Physical Review, 165, pp.201-214, 1968.

[28] Yip, S., "Nanocrystals: The strongest size", Nature 391, pp. 532-533, February 1998.

[29] Zhang, W.J., Evangelista, E., and Francesconi, L., "Effect of prior cooling rate on the grain size of fully-lamellar TiAl-base alloy developed by tempering/quenching", Scripta Materialia, Vol. 35, No. 1, pp. 41-45, 1996.