Topological Modeling of Force Networks in Granular Materials

Advisor:

David B. Damiano

Author:

Reader:

Emily T. Winn

Andrew Hwang

College of the Holy Cross

Department of Mathematics and Computer Science

May 16, 2017

Abstract

Granular materials, or collections of solid macroscopic particles in contact with each other, play an important role in chemistry, pharmaceutical sciences, and agriculture. Despite their importance, little is understood about the force networks formed from inter-particle contacts in these materials because of the difficulty of modeling large networks. The challenge lies in finding a model that is specific enough to characterize the local topology of a cell while still being robust enough to capture information of the entire network. A recently developed method that may answer this dilemma is swatch and cloth, which uses adjacency graphs and statistics to model random cell networks. We discuss how the swatch and cloth model can be applied to force networks and examine the abilities of swatch and cloth to detect differences in networks extracted from simulations of granular materials with different sized particles. We also apply swatch and cloth to communities, or partitioned sub regions, of a force network to determine which behaviors are uniform throughout the network and which behaviors tend to cluster in specific regions.

1 Introduction

Granular materials refers to collections of "contacting solid, discrete, macroscopic particles" [3]. Examples of granular materials include sand, powder, and soil. While granular materials play significant roles in the chemical, pharmaceutical, and agricultural industries [3], little is known about how to predict *bulk properties*, such as a granular material's ability to retain heat or absorb water [1]. Through experimentation, materials scientists can classify the behavior of different materials, but there is still difficulty articulating why certain materials behave the way they do. If necessary and sufficient characteristics of bulk properties are identified, materials scientists will be able to create synthetic materials that better suit the needs of the various industries in which they are used.

Particles in contact with each other possess an interaction force between them; the collection of these forces is the *force network*. Figure 1b depicts force chains that form when
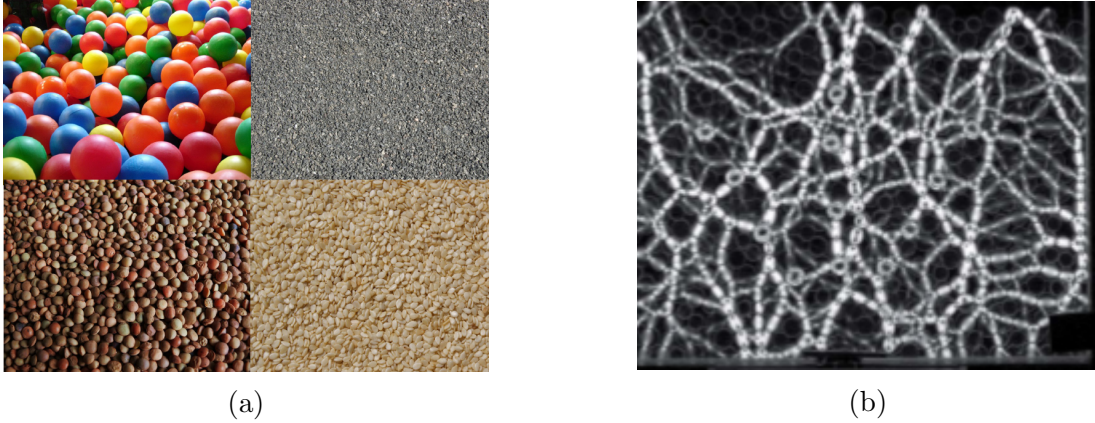
1

(a)   (b)

Figure 1: Images courtesy of Google and [1]

granular materials are compressed from outside forces. Subsets of the force network are sometimes refered to as *force chains*. A number of materials scientists believe that the key to predicting bulk properties lies in these force networks [1]. The ultimate goal for many materials scientists is to find a universal method to characterize such networks to enable the comparison of different materials; we believe statistical topology could provide such a method.

In topology, a *cell complex* is built by attaching n-dimensional cells; for example, a point is a 0-cell, a line segment is a 1-cell, a disk is a 2-cell, and so on. When the cells are simplices, complexes have simpler structure and are relatively easy to characterize. In dimension two, a simplicial complex consists of vertices, line segments, and triangles. Random cell complexes, where, for example, in dimension two the polygons have any number of sides, occur more frequently in nature and are more difficult to describe and analyze. These are the complexes that arise in the study of force networks. Figure 2 shows a simplicial complex and a random cell network. Until recently, there has been no method for characterizing random cell networks that was detailed enough to give information about the network but robust enough to apply cell networks in general. Attempts to apply statistical topology to force networks in granular materials has been met with mixed results. Kramar *et al* used persistent homology and persistence diagrams to understand changes in the force networks of simulated granular materials during compression. However, they found that while their
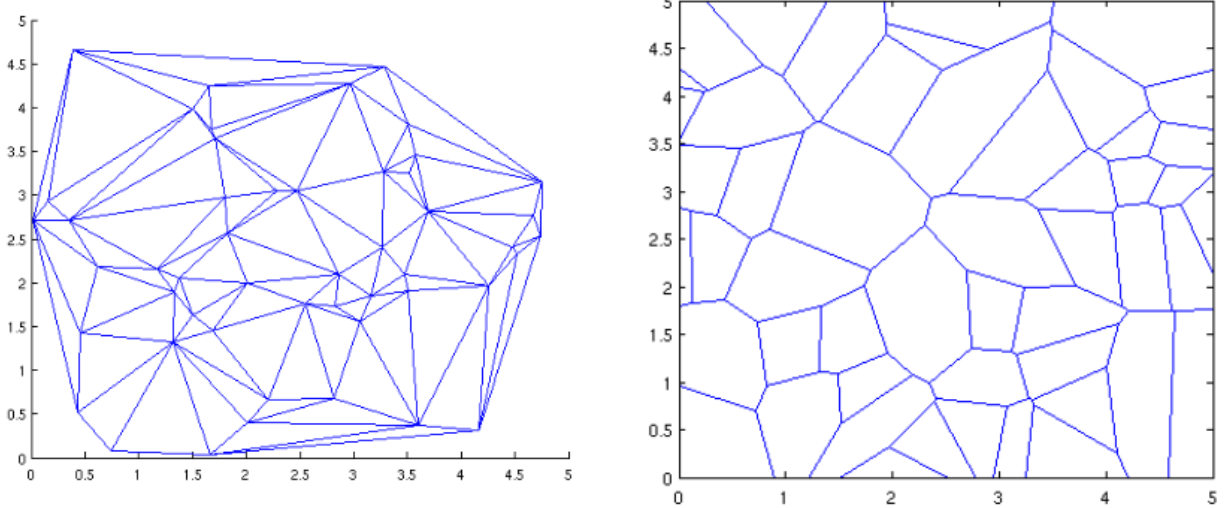
2

Figure 2: A simplicial versus a random cell complex.

analysis provided useful information about the stability of the simulated material, it was unable to "provide detailed information about the force network structure" [5]. Further, their probability density function revealed information about the large forces in the system but "not about the connectivity of the force network" [5]. To address the issue, we turn to the recently developed method of swatch and cloth.

1.1 Swatch and Cloth

The method of swatch and cloth was developed by Schweinhart, MacPherson, and Mason [9]. The method is designed for modeling complex systems represented by random cell complexes K by creating and analyzing the structure of the adjacency graph for K rather than K itself [8]. The *adjacency graph* $A(K)$ is a one-dimensional cell complex in which each vertex represents a cell from the original complex K (see Figure 3). In the figure, each vertex is color-coded based on the dimension of the cell it represents. Two vertices v_1 and v_2 in the adjacency graph only share an edge if the dimensions of the cells they represent are of dimension k and $k + 1$, $k \geq 0$, and if the cell of dimension k is a face of the cell of dimension $k + 1$. Note for the rest of this section, the terms *vertices* and *edges* will refer to

3

the vertices and edges of the adjacency graph of the cell complex. Schweinhart *et al* then use probability distributions of local configurations that are subsets of the adjacency graph (represented by swatches) to characterize the whole complex (represented by the cloth). To illustrate the defintions that follow, we use examples of swatches from the force network and its corresponding adjacency graph in Figure 7 unless otherwise stated.
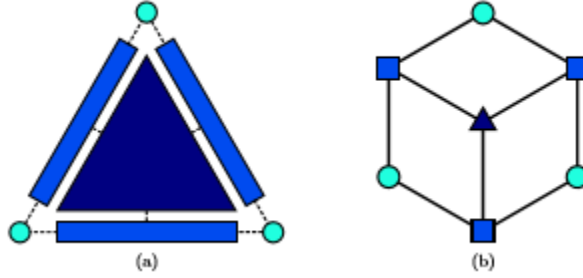

Figure 3: (a) A cell complex with three 0-cells, three 1-cells, and one 2-cell and (b) it's adjacency graph with vertices corresponding to the dimension of the cells they represent [8]. Two vertices of the adjacency graph share a vertex only if the cells they represent have a difference in dimension of one and if their cells interact in the composition of the original cell complex.

To define a swatch, we first define the distance between two vertices $v_1, v_2 \in A(K)$. Two vertices $v_1, v_2 \in A(K)$ are distance r apart if the shortest path between the vertices along the graph has r edges. For a vertex $v \in A(K)$ representing a 0-cell in K, a *swatch* of radius r rooted or centered at v is the subgraph of the adjacency graph $A(K)$ that contains all vertices of $A(K)$ within distance r of the root cell. Each swatch represents a local topological configuration of the original cell complex K [8]. To illustrate, a swatch with root cell v of radius 0 is the root cell by itself, a swatch at root cell v of radius 1 is the root cell and all vertices that are one edge away from v, a swatch at root cell v of radius 2 is the collection of vertices that are within two edges of the root cell, and so on. An *isomorphism* of swatches is a bijection between the vertex sets of the two swatches that maps the root cell of one to the root cell of the other and preserves the adjacency relationships of the swatches. Two swatches are *equivalent* if there exists an isomorphism between them.

A *subswatch* s_0 of a swatch s_1 is a swatch with the same root cell as s_1 but with a smaller
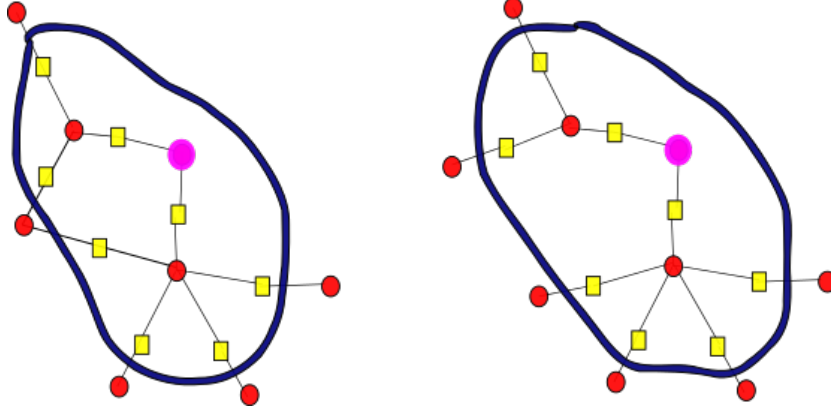
4

Figure 4: Two swatches of radius four with the root cell highlighted in pink, the other 0-cells in red, and the 1-cells in yellow. The largest common subswatch, circled in both swatches, has radius $r = 3$ and has $n = 11$ vertices, so therefore the distance between them is $1/11$.

radius [8]. We say that a swatch s_1 *descends* from an *ancestor* swatch s_0 if s_0 is a subswatch of s_1. The collection of distinct, non-isomorphic swatches in $A(K)$ can be given the structure of a rooted tree called the *swatch tree* of the complex. The depth or level of a swatch in the tree is its radius r. Each swatch of radius r is connected by an edge to the unique ancestor of radius $r - 1$ at the same root. Thus each swatch of radius r is connected to its descendant swatches of radius $r + 1$. Figure 5 shows the swatch tree up to radius 4 for the adjacency graph in Figure 7.

The *cloth* of a complex K is the swatch tree of $A(K)$ together with a probability distribution at each level r of the swatch tree. Recall that the level r of a swatch tree contains each unique swatch type of radius r that appears in $A(K)$. The frequency at which a swatch occurs in the network is the number of root cells in the adjacency graph $A(K)$ at which the swatch appears. The probability of the swatch is the frequency divided by the total number of root cells in the network. This defines a probability distribution on each level of the swatch tree. This distribution is the *cloth* at level r of the network. The cloth for the entire network is the collection probability distributions for every level r. Figure 6 shows a visual representation of two distinct cloths at level 4.

Note that the probabilities of each level of the cloth must add up to 1 and that the probabilities of the descendants of radius $r+1$ of a swatch s_0 of radius r sum to the probability
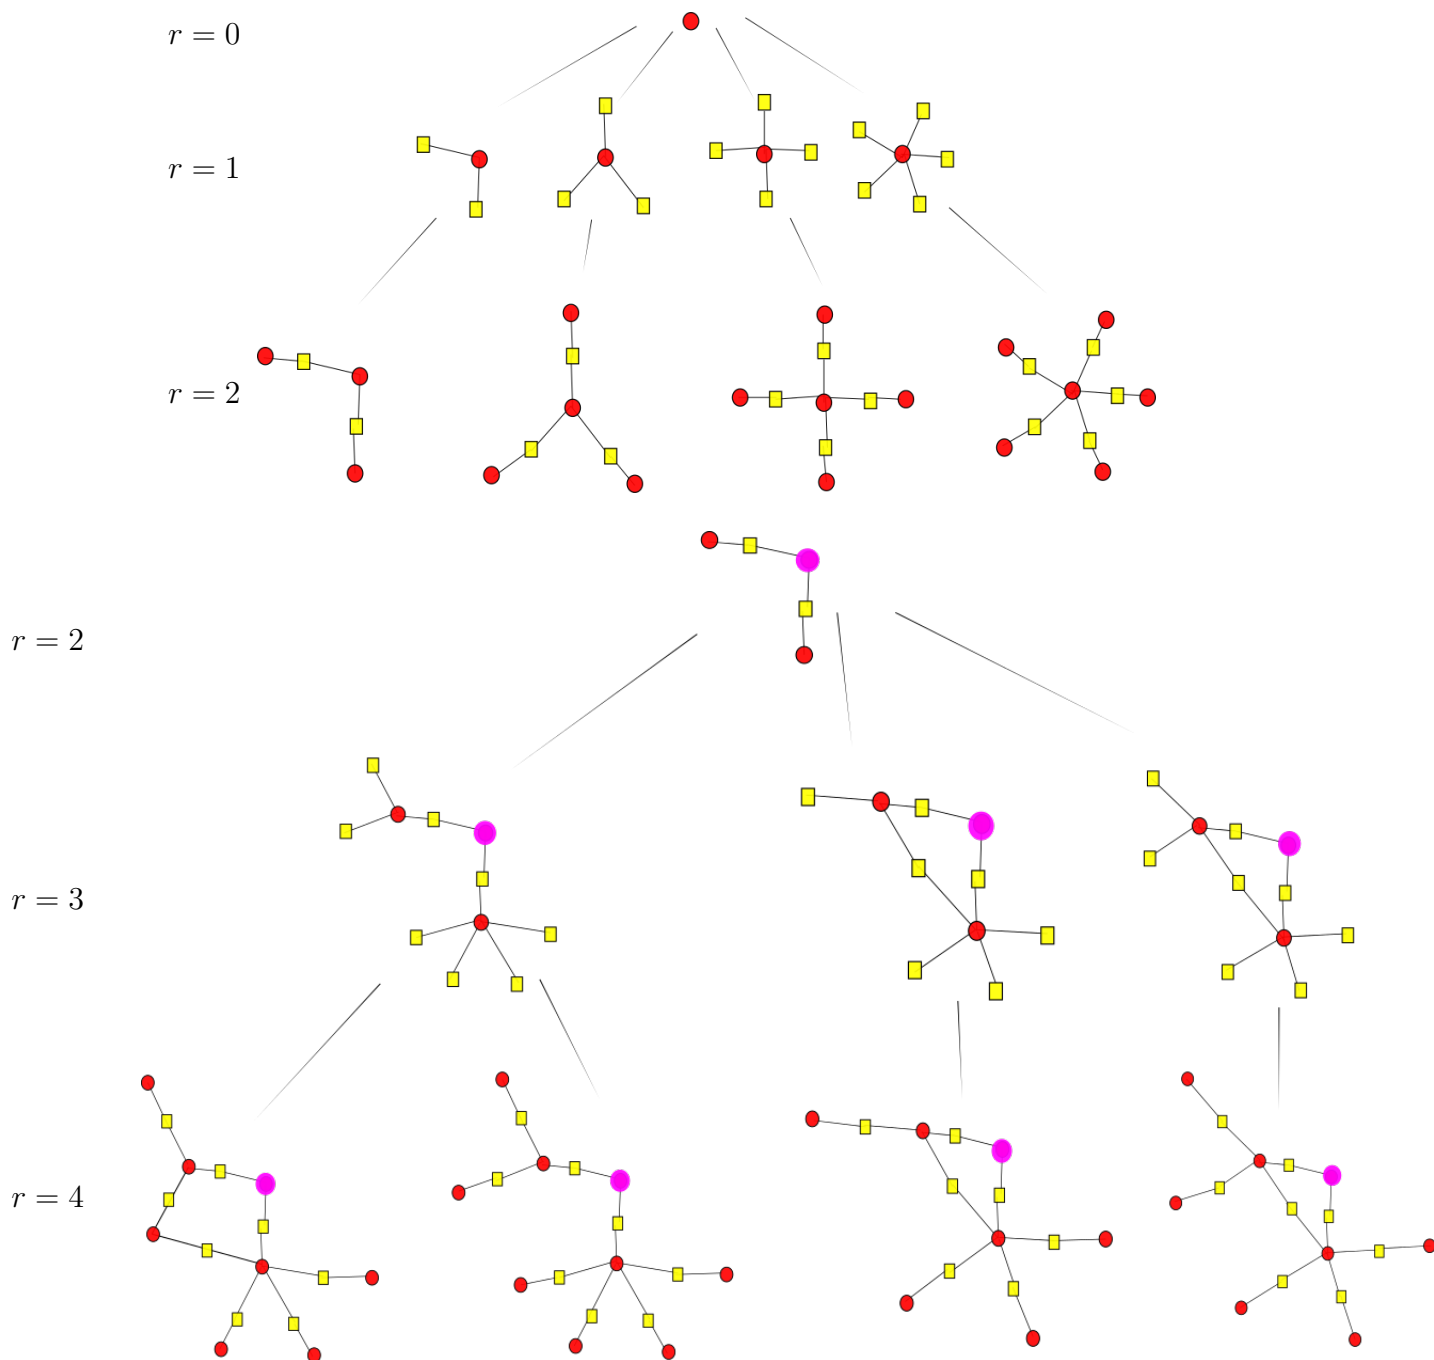
Figure 5: A swatch tree of up to radius 4; assigning probabilities to each swatch gives the cloth. Red vertices represent 0-cells, while yellow vertices represent 1-cells. In the second half of the tree, the root cells of the swatches are marked in pink. Note that radii 2-4 show the subtree as descends from the far left swatch at level 2. A similar branching happens for all swatches, and thus a swatch tree quickly becomes large.

of s_0 itself. To compare two cloths, we determine the distance d_r between cloths at each level r using the *earth mover's distance*. This can be interpreted as the minimum effort to transform one probability distribution into another [8]. We first define the distance between two swatches of radius r. The *distance* between two swatches of equal radii is the reciprocal of the number of vertices in their largest common subswatch [8]. If two swatches are equivalent, then the distance between them is 0. See Figure 4 to see an example of two swatches of radius four and the calculation of the swatch distance between them. The "effort" used to transform probability distributions depends on the finite set itself; for a cloth of level r, the effort is calculated using the swatch distance at radius r. The minimum effort for transforming a cloth at level r is the minimum sum of the costs of all of the operations needed to transform one probability distribution at level r to another [9]. The cost of each operation is "the probability mass transferred times the distance between two swatch types" [9].

To illustrate how to calculate the distance between two cloths at a given level, consider two level-four cloth distributions on the four swatches shown in Figure 5, displayed again in Figure 6 with the two cloth distributions.
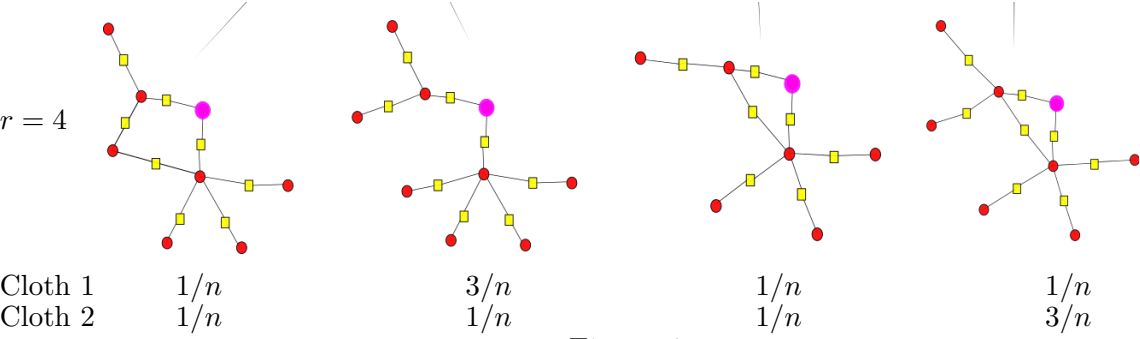

$r = 4$				
Cloth 1	$1/n$	$3/n$	$1/n$	$1/n$
Cloth 2	$1/n$	$1/n$	$1/n$	$3/n$

Figure 6

Assume that there are n total root cells in the network and that the swatches at the other $n-7$ root cells have the same frequency distribution in both cloths. For the probability distributions to be equal, the two probabilites assigned to the same swatch must be equal. In this example, the second and fourth swatches have different probabilities. To change Cloth 2 to Cloth 1, the probability of swatch four must change from $3/n$ to $1/n$ and simultaneously

7

the probability of swatch 2 must change from $1/n$ to $3/n$. This equates to "moving" two of the swatches from type four to type two, which transforms $2/n$ units of mass of the probability distribution. From the tree in Figure 5, we see that the largest common subswatch for these two takes place at level two and has five vertices. The swatch distance is thus $1/5$. It follows that the earth mover's distance is $(2/n)(1/5) = 2/(5n)$, as this modification of the distribution requires the minimum effort since it moves the minimum number of swatches over the minimum possible distance. Therefore, the distance d_4 between the two cloths at level 4 is $2/(5n)$.

The distance d between two cloths C_1 and C_2 is the limit of these earth mover's distances as the level r approaches infinity. That is,

$$d(C_1, C_2) = \lim_{r \to \infty} d_r(C_1, C_2).$$

The distance calculation permits the comparison of the local topological properties of a complex, global topological properties of a complex, and properties in between to fully detail how two complexes compare or do not compare to one another. We will apply swatch and cloth to the complexes that represent force networks in granular materials.

Given a force network, we will build a one-dimensional cell complex where a 0-cell is the center of a particle and a 1-cell represents a normal contact force between two particles. We then build the adjacency graph with two vertex types representing the 0- and 1-cells of the network. From this graph we perform our swatch and cloth analysis. Figure 7 shows a simulation of a compressed granular material, the force network, and the adjacency graph for the force network. As far as we know, swatch and cloth has not yet been used to analyze force networks in granular materials.

1.2 Communities

While the swatch and cloth method provides information about the local characteristics of a network, it does not reveal where such characteristics occur within the network. There is
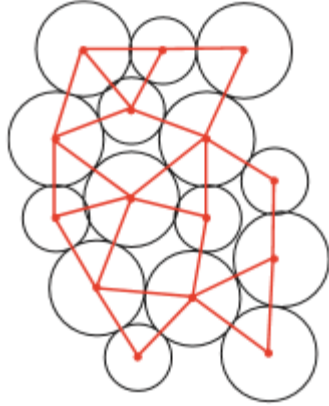
(a) An example of a position force network extracted from 15 bidisperse particles in contact. Image from [5]
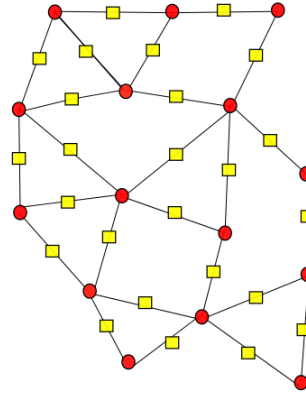
(b) Adjacency graph of network

Figure 7

no information as to whether a swatch type appears with higher frequency in a region of the network or is distributed throughout the network. As a result, when looking at larger networks comprised of more than a few vertices, it is difficult to determine where different network behaviors occur. To address this shortcoming, we use the concept of communities to understand the local behavior of a random cell network.

Communities are subsets of a network partitioned along areas of less or weaker force contacts. The partition relies on the value of the modularity Q of a network [1], which is determined by a number of variables. Define an *edge weight* between particles i and j as the magnitude of the contact force between them. Since the magnitudes of the contact forces were not considered in our simulations, we will say the edge weight between particles i and j is 1 if they are in contact and 0 if they are not in contact. These values can be represented in a matrix W where the ij^{th} entry is the edge weight between particles i and j. The *predicted edge weight* between i and j is the product of the number of contact forces involving i and the number of contact forces involving j divided by the total number of contact forces in the network. This calculation is known as the *Newmann-Girvan null model*, which is the most common model used for optimizing the modularity Q, although other models can be used [1]. The matrix of predicted edge weights is denoted by P, where the ij^{th} entry is the
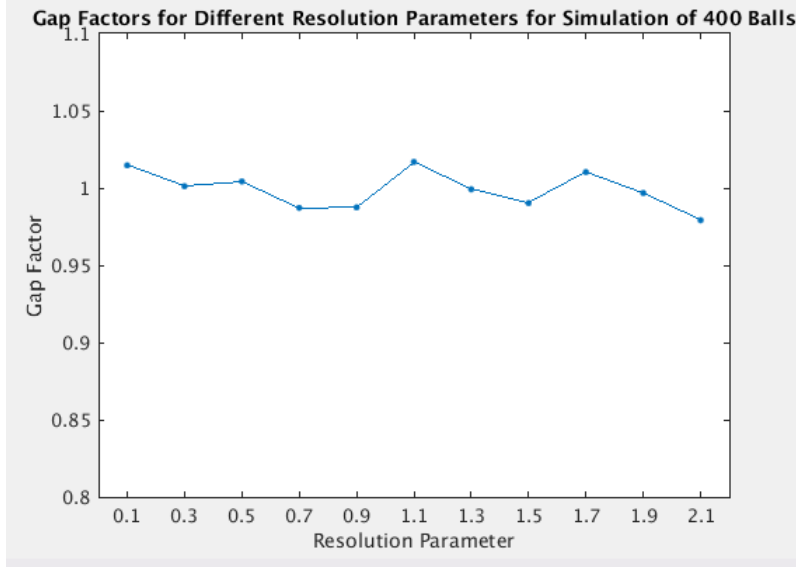
9

Figure 8: The gap factors for all resolution parameters tested for a bidipserse simulation of 400 balls. The maximum gap factor occured at $\gamma = 1.1$, so resolution parameter 1.1 was used to calculate communities.

predicted edge weight between particles i and j. The *resolution parameter* γ is an arbitrarily chosen value that determines the partition. The *clusters* of these partitions are the subsets of network determined by the resolution parameter. The cluster to which particle i belongs is denoted c_i. Note that $c_i = c_j$ if particles i and j are in the same cluster as determined by the partition. The modularity Q of a partition for resolution parameter γ is defined as

$$Q = \sum_{i,j} [W_{ij} - \gamma P_{ij}] \delta(c_i, c_j) \, ,$$

where i and j represent distinct particles, W is the symmetric matrix of edge weights, P is the symmetric matrix of expected edge weights, γ is the resolution parameter, c_i and c_j are the designated clusters of particles i and j, and δ is the Kronecker delta. The non-zero contributions to the calculation of modularity are from vertices that are in the same clusters in a particular partition. For each choice of γ, there is a partition c of the network that maximizes Q.

To determine which resolution parameter to use for the calculation of Q and thus de-

termine which network partitions gives the desired communities, a diagnostic called the gap factor was used. The *gap factor* g_c of a cluster c "measures the presence of gaps and the extent of branching in a community" [1]. It describes the relationship between the Euclidean distances between network vertices and the *hop distance*, or the path of the minimum number of network edges, between particles within the same cluster. This identifies communities with branching as opposed to communities that are compact or have linear chains, thus isolating force chains within the same community. The gap factor g_c for a cluster is calculated as

$$g_c = 1 - \frac{r_c s_c}{s_{max}} \, ,$$

where r_c is the Pearson correlation between the hop distance and the Euclidean distance between particle pairs in cluster c, s_c is the number of particles in cluster c, and s_{max} is the number of particles in the largest cluster of the partition [1]. The gap factor g for the entire network is given by

$$g = 1 - \frac{1}{n} \sum_c \frac{r_c s_c}{s_{max}} \, ,$$

where n denotes the number of clusters in the partition [1]. Selecting the partition that maximizes the gap factor produces the desired communities; Figure 8 shows the different gap factors for different gamma values tested for the simulation in Figure 9. We will further detail our methods for calculating communities in the Methods and Materials section.

The method of swatch and cloth can be applied to communities within a force network and communities between force networks. This allows us to determine whether behavior is localized to a particular area of the network or whether the behavior occurs consistently throughout the network. Figure 9 shows the stages a force network as it is extracted from the jammed simulation and partitioned into communities.

1.3 Purpose

The goal of our project is to apply swatch and cloth to both entire force networks and communities within force networks to determine which networks can be distinguished by swatch

(a) 400 Bidisperse particles jammed

(b) It's corresponding position network
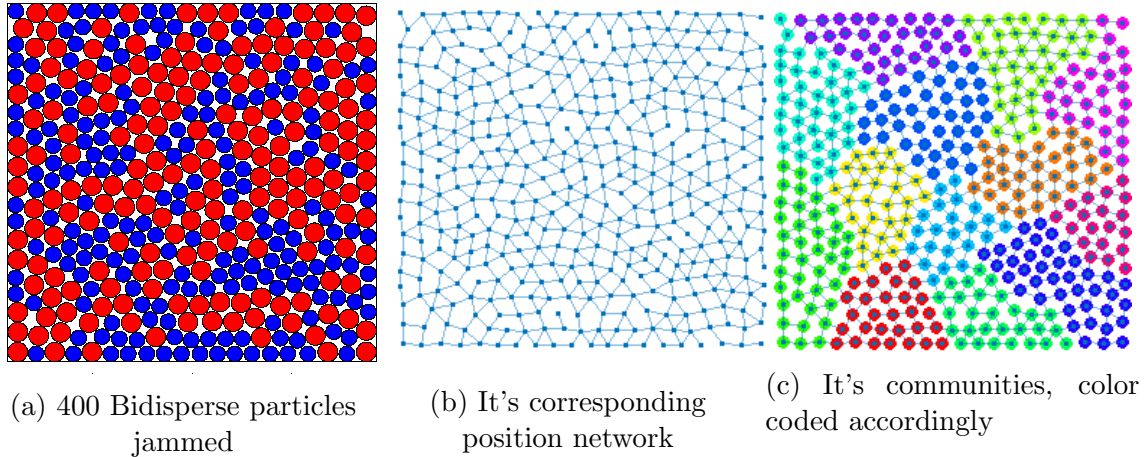
(c) It's communities, color coded accordingly

Figure 9: A bidisperse simulation with 400 balls of radius 4 (blue) and radius 5 (red) shown at jamming, the extracted position network from the jammed state, and the communities color coded for that network.

and cloth. The cell complexes for force networks were derived from computer simulations of particles in motion inside a walled container being compressed until no further compression is possible. We say that these particles are jammed. Note that due to limitations in software, we were only able to extract and analyze the position networks, or the networks in which all edges are evenly weighted, rather than calculating the true force magnitudes for all contact forces of an interaction network.

We examined monodisperse and bidisperse simulations of various sizes and applied swatch and cloth to examine whether the jammed networks of different types were significantly different. By monodisperse, we mean a simulation in which all particles have the same mass and radius, whereas by bidisperse, we mean there is a 50-50 distribution of two types of particles that have equal densities but different radii.

For communities, we ran bidisperse simulations with different numbers of balls, identified the communities, and then compared the networks of the communities to each other to observe whether communities were universally the same across networks or whether they were dependent on the simulation from which they came. We also compared the cloths of communities to cloths of whole networks to determine if certain characteristics of the whole networks were clustered or whether they were distributed throughout the network.

12

2 Methods and Materials

To collect and analyze data of force networks in granular materials, we used three major pieces of software: a simulation of granular materials based on the software written by Stepen Sahrun [7], a program written by Benjamin Schweinhart to calculate and compare cloths [8], and Generalized Louvain software to detect and identify communities [4]. Note that Schweinhart's software implements NAUTY [6] to analyze the adjacency graphs for cell complexes and to determine distinct swatch types at a given radius up to isomorphism. NAUTY (No AUTomorphism, Yes?) is a graph isomorphism testing progarm developed by McKay and Piperno [6].

The simulations of granular materials were run in MATLAB. The final software was based on Sahrun's program that simulates elastic collisions between balls. We modified his code to increase physical accuracy of the elastic, frictionless collisions and to implement moving compressing walls. In each simulation, a number of balls of predetermined sizes and densities were placed randomly in a rectangular box with randomly assigned velocities. Each collision between balls was assumed to be elastic, and each of the balls were considered to have no friction and to have perfect rigidity, meaning they did not change shape from collisions. At the start of each simulation, the boundaries of the rectangular box moved towards the center of the box at a constant velocity, thus decreasing the area in which the balls were moving. The collisions between the balls and the walls were considered completely elastic, and the walls were also frictionless. The simulation finished when it reached a point of *jamming*, or a state in which the balls were fully compressed and the walls could not move in any further. In the jamming state, almost all balls are in contact with at least one other ball and thus there are many contact forces appearing in the simulation. The information for each particle, including mass, radii, Euclidean coordinate positions, and velocity, as well as the final positions of the boundaries were recorded at the finish of each simulation. Figure 10 shows a simulation at the start and end of the program.

Because the simulation significantly slowed when the balls reached too high of a velocity and the time step between collisions approached zero, initial speeds were capped at 12.5
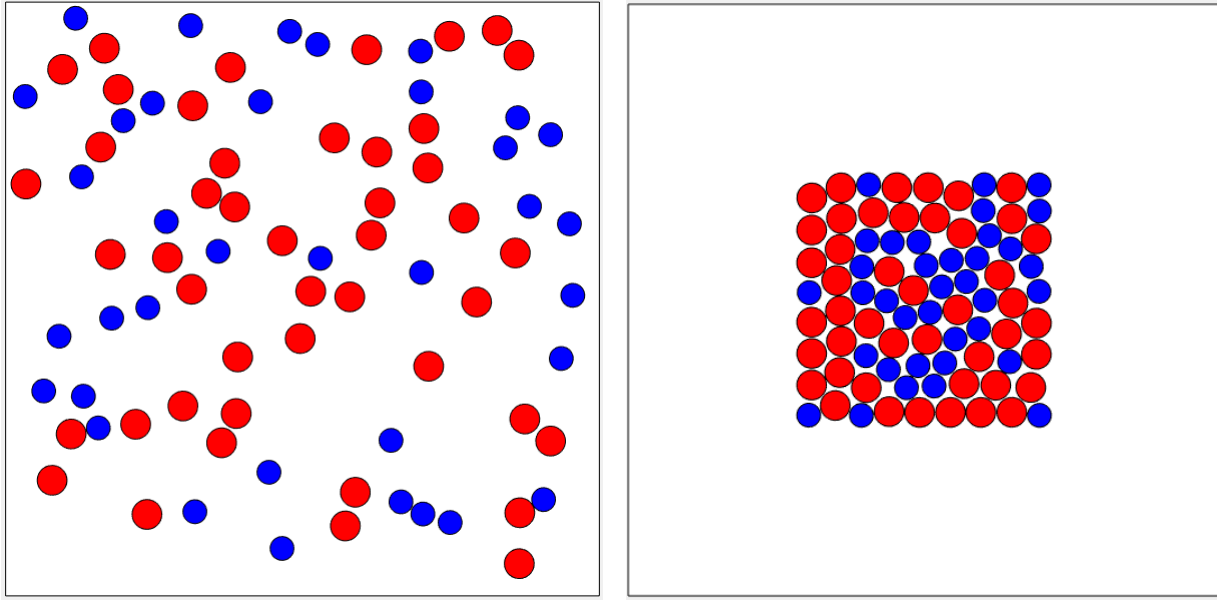
Figure 10: A bidisperse simulation of 100 balls at the start of the program and at the end of the program when it had reached a state of jamming.

units per second, and maximum speeds were restricted to 25 units per second. The slowing of the program was caused by the balls picking up speed with each collision, as collisions were elastic and thus no kinetic energy left the system. This caused the program to have to adjust for collisions in extremely small time increments, which prevented the simulation from reaching the jammed state. Finally, all particles in simulations were disks. Even though most macroscopic particles in granular materials are not perfectly round, most simulations of granular materials use disks as they are easiest to program and the shapes' effects on overall data are relatively small [3]. Specific parameters of the simulations ran for testing hypotheses are outlined below.

Since most balls were not in perfect contact at the end of a simulation, we introduced a small scalar value ϵ. Balls within 2ϵ of each other were considered to be in contact. The use of "growing" radii of balls is similar to that outlined in [2]. For the majority of our network calculations, an ϵ smaller than 20% of the radius of each ball was used. Schweinhart's swatch and cloth software calculated the cloths for the resulting networks at different radii starting at 2 and ending at 10, going only slightly further than Schweinhart *et al* did in

their analysis. They found that most calculations beyond radius 8 did not reveal novel or interesting information about a network or the relationship between two networks. [9]

Generalized Louvain software was used for community detection in MATLAB [4]. Since this program requires that networks be fully connected, that is, every vertex shares at least one edge with another vertex, the choice of ϵ's played an important role. For any given network, ϵ was chosen to be the minimum value at which all vertices shared at least one edge with another vertex. This provided the smallest possible single component network containing the jammed network. The genLouvain software calculates the partition for the Newmann-Girvan null model that maximizes the modularity of Q for the given γ. The resolution parameter γ that maximized the gap factor g and thus partitioned the communties correctly was calculated by the same methods outlined in Bassett et al [1]. For each network, values of γ from 0.1 to 2.1 were tested in increments of 0.2. The optimal partition of the network for the value of γ that optimized the value of g was the chosen partition for the network [1]. The networks for the resulting communities were also analyzed by Schweinhart's software and compared to one another.

To test whether swatch and cloth could detect differences between two types of networks, a reference set was selected from one type of simulation. This was compared to two test sets, one from the same type as the reference set and the other from the second simulation type. The reference set consisted of a random sample of 20 simulations of one type. The test sets consisted of 40 simulations of the same type as the reference set and 40 simulations of the second type. Two distance sets, one for each test set, were then calculated; each distance set contained the distances between all pairs of reference set simulations and test set simulations. A t-test that assumed different variances was applied to the two distance sets to detect a difference of means at the 5% level of significance. To account for possible outliers, for any two simulation types, the process was repeated 25 times with reference sets from the first type and 25 times with reference sets from the second type. Two types of networks were considered significantly different if the difference of means between distance sets were significant regardless of which simulation type was used for the reference set.

Simulations of size 100 balls were tested for whether swatch and cloth was able to detect differences in cell networks extracted from monodisperse simulations and bidiperse simulations varying in radii and number of balls in the simulation. The monodisperse simulations consisted of particles of radius 4. The bidisperse simulations had distributions of approximately 50% balls of a smaller radius and 50% balls of a larger radius. The bidisperse simulations were of granular materials with paired particle radii 2 and 7, 3 and 6, and 4 and 5. Figure 11 has images of the jammed simulations, the extracted network, and the adjacency graph of each type. The ϵ values used to determine the force networks were fixed for simulations of the same type. The following table shows these ϵ values for each type of simulation:

Sim. Type	2-7	3-6	4-5	4
ϵ	1.44	0.95	0.8	0.77

At these values of ϵ, the majority of networks were connected save for a few outlying particles in each network; increasing ϵ further would present the risk of losing network information. The distances between cloths were calculated both from the same and from different simulation types at every radii considered for each of the networks. The goal was to detect if there was a significant difference in networks extracted from particle simulations of different radii/size distributions. The results are listed in the next section.

We also ran large bidisperse simulations of radii 4 and 5 with 400, 600, and 1000 balls. Their cloths were compared to the simulations of 100 balls to determine whether the size of the simulation had any effect on the resulting cloth. Their ϵ values, which were calculated as above, were as follows:

Sim. Size	100	400	600	1000
ϵ	0.8	0.93	2.85	1.117

For these simulations, we also extracted their communities and applied swatch and cloth. The goal was to compare communities within the same network to each other and to large networks as well as communities across different networks. Most of the data and analysis
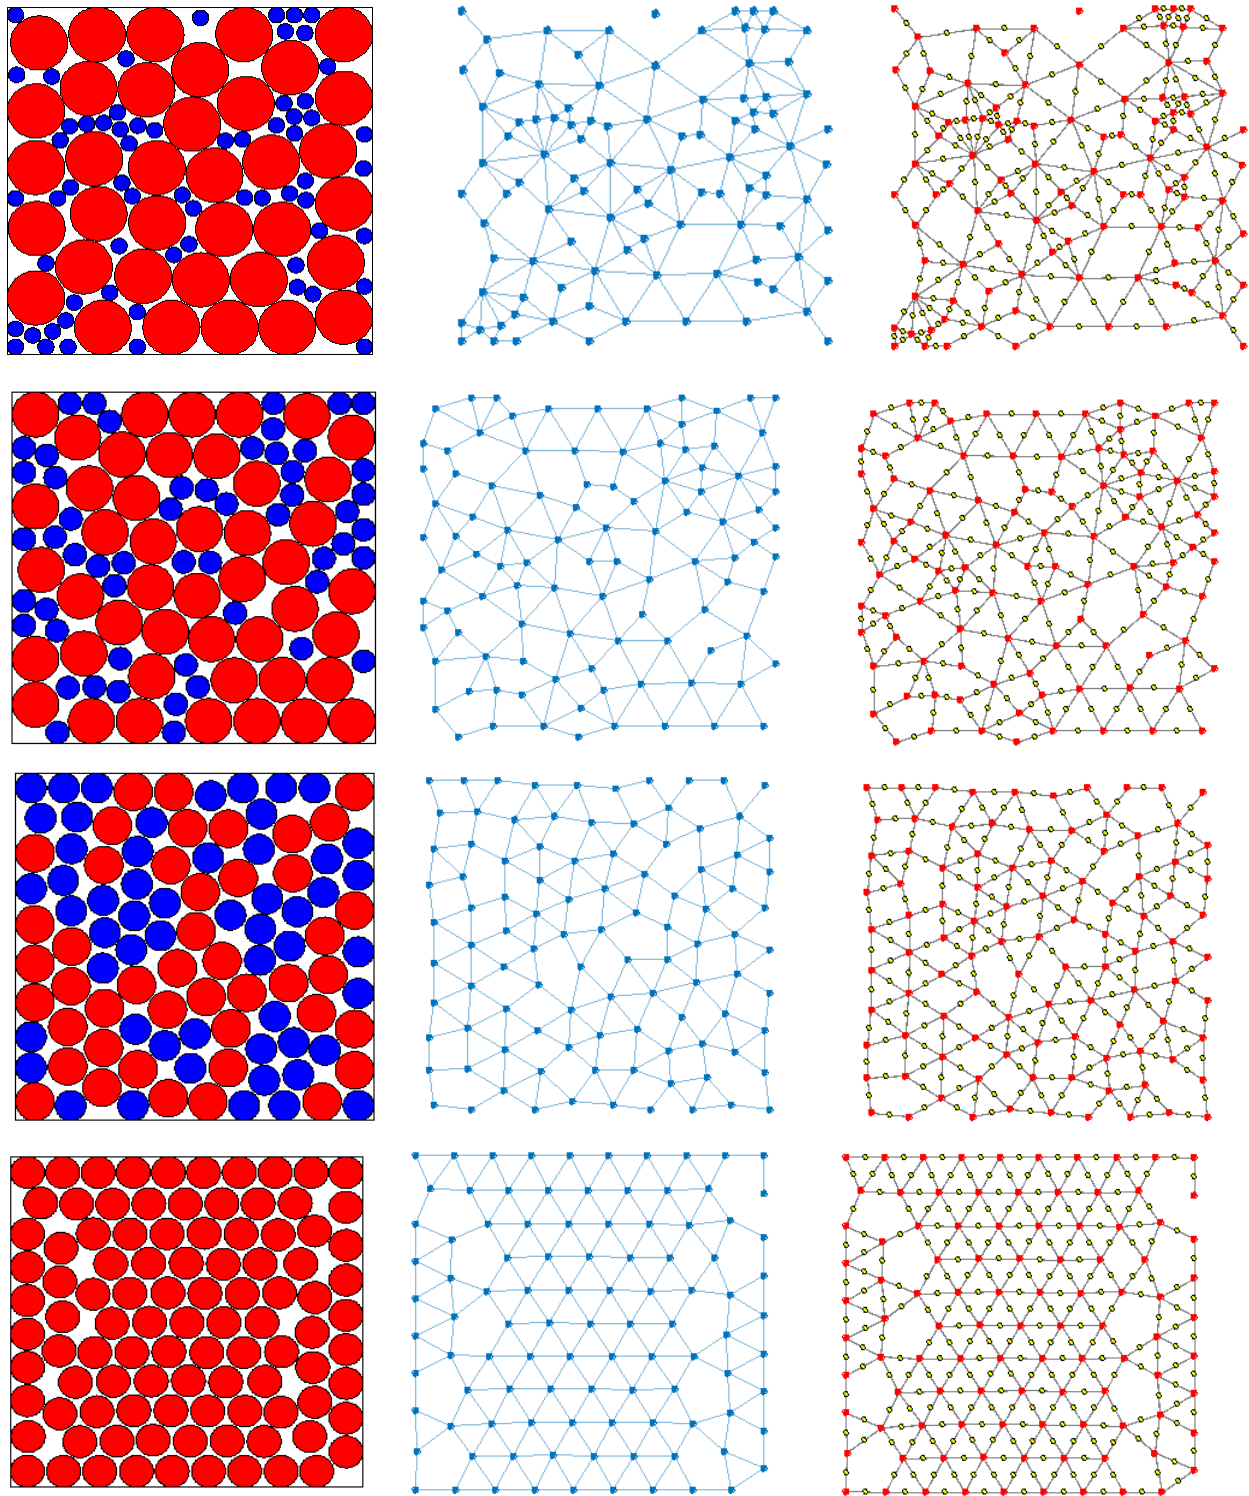
Figure 11: Examples of jammed states, corresponding networks, and adjacency graphs for a 2-7 Bidisperse distribution, a 3-6 Bidisperse distribution, a 4-5 bidisperse distribution, and a monodisperse distribution.

was from the simulations of 400 balls, as the 600 ball and 1000 ball simulations were computationally expensive. Thus, we were unable to produce the necessary amount of data for analyis in the time alloted for this project. However, we did analyze the available networks from larger simulations; their potential in further research is discussed in the conclusion of the paper. Additionally, the cloths for communities were only calculated up to radius 5 rather than radius 10 as some of the communities contained only 20 particles; the small size meant that any swatches larger than radius 5 would contain most of if not the entire community. Note that no community analysis was performed on the 100 ball simulations, as those networks were too small to produce meaningful communities for analysis. We also restricted our work to internal communties, or communities made up entirely of balls that were not adjacent to the compressing borders of the simulation, in order to avoid a "boundary effect."

3 Results and Discussion

3.1 Mono- and Bidisperse Simulations of 100 Balls with Different Particle Sizes

When calculating the cloths for each network, we noticed that by radius four, most of the swatches in the swatch trees only appeared once. Swatches that appeared more than once usually only appeared two or three times and rarely more. The uniqueness of swatches increased with radius; in most networks, at larger radii, there are perhaps one or two swatches that repeated if there are any at all. This was consistent across all simulation types and sizes. This is by no means surprising given the simulations; root cells had degrees varying from 0 to 7. However, considering the large number of swatches at levels greater than three, the following results become more meaningful.

In comparing the test sets for 100 ball simulations, we found that swatch and cloth was able to differentiate between some of the simulations of different types but not others. Figures 12 and 13 show the heat maps for distances between individual simulations of each type. In Figure 12, we see that the swatch and cloth method was able to distinguish between

bidisperse simulations with radii of 2 and 7 and all other simulation types at all cloth levels tested. In the statistical analysis, when reference sets were taken both from the 2-7 bidisperse simulations and the other simulations, the null hypothesis of no difference between mean distances between simulations of the same and different types was rejected every time at a negligible level of significance averaging 1×10^{-15}. This result is not surprising given the unique structure of the 2-7 networks visible in Figure 10. Swatch and cloth was also able to distinguish at all levels between the bidisperse simulations of radius 3 and 6 and bidisperse simulations of radius 4 and 5, with a p-value usually in the range of 1×10^{-8}; these results are visible in the first row of heat maps in Figure 13. Again, given the networks extracted from both simulation types, this result was not surprising.

A surprising result, however, involved the monodisperse distributions of radius four. There was inconclusive evidence at all levels that swatch and cloth distinguished between the monodisperse simulation networks and the bidisperse simulations network of radii 3 and 6 and of radii 4 and 5. While the statistical analysis with reference sets in either bidiperse simulations did demonstrate a significant difference of means of distances between the bidisperse and monodisperse networks, when the reference set was drawn from the monodisperse simulations, the two means were not significantly different. This may not come as a surprise for the 4-5 bidisperse simulation and the 4 monodisperse simulation, as they share a common radius and their networks appeared fairly similar, but the lack of difference in distances between the monodisperse simulations and the 3-6 bidisperse simulations was surprising. Not only did the simulations and the resulting force networks appear different but also swatch and cloth distinguished between 4-5 bidisperse simulations and 3-6 bidisperse simulations, so if 4-5 bidisperse was indistinguishable from monodisperse, one would think that swatch and cloth also would have distinguished between monodisperse and 3-6 bidisperse. Given that the difference was not statistically significant when the reference set came from monodisperse simulations, perhaps there was more variation in the networks extracted from monodisperse simulations than there were in the networks extracted from 4-5 and 3-6 bidisperse simulations. Thus, the mean distances could not be significantly different when the reference set

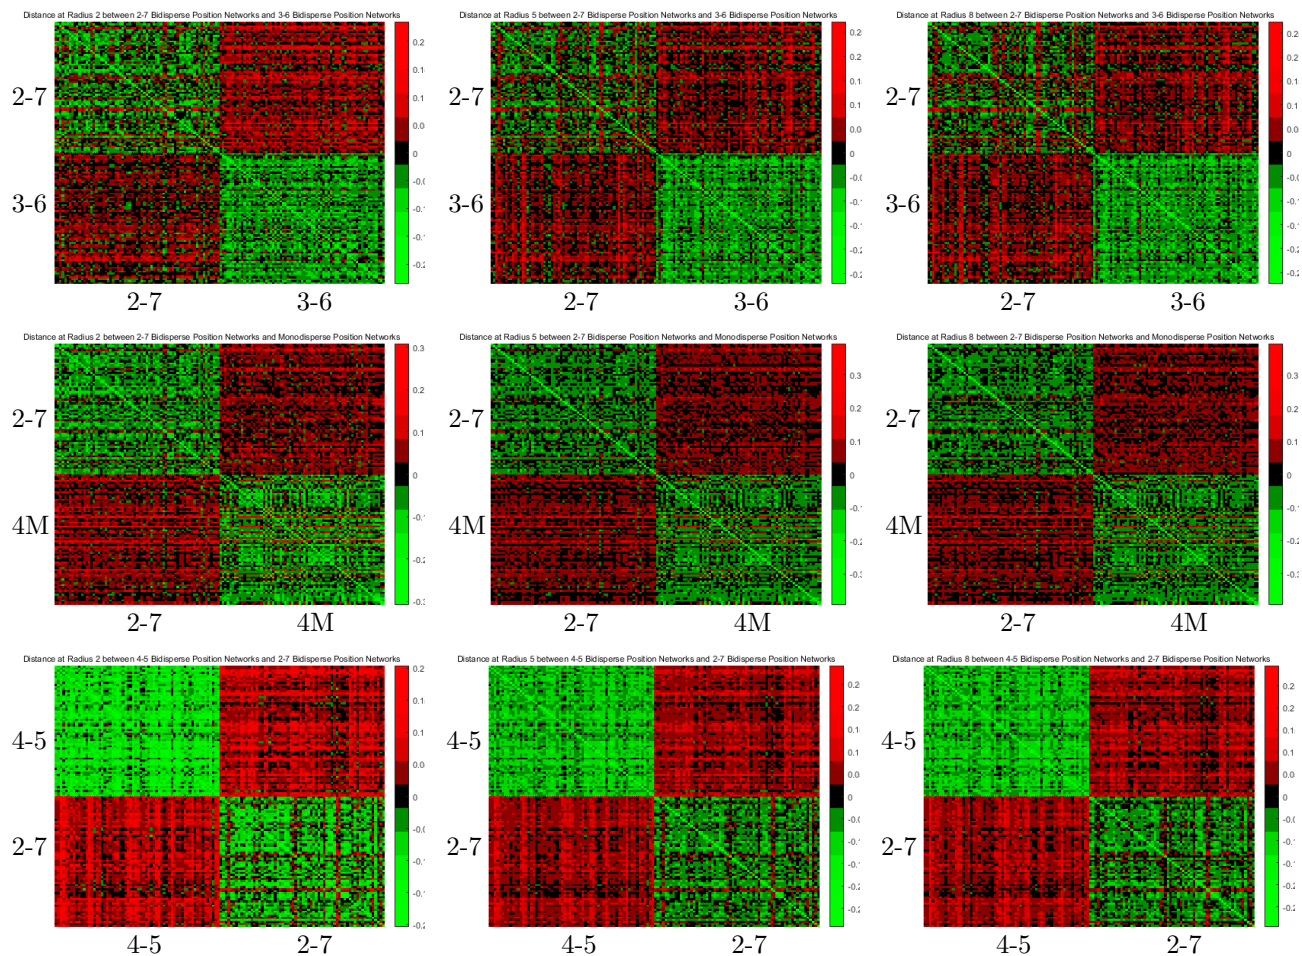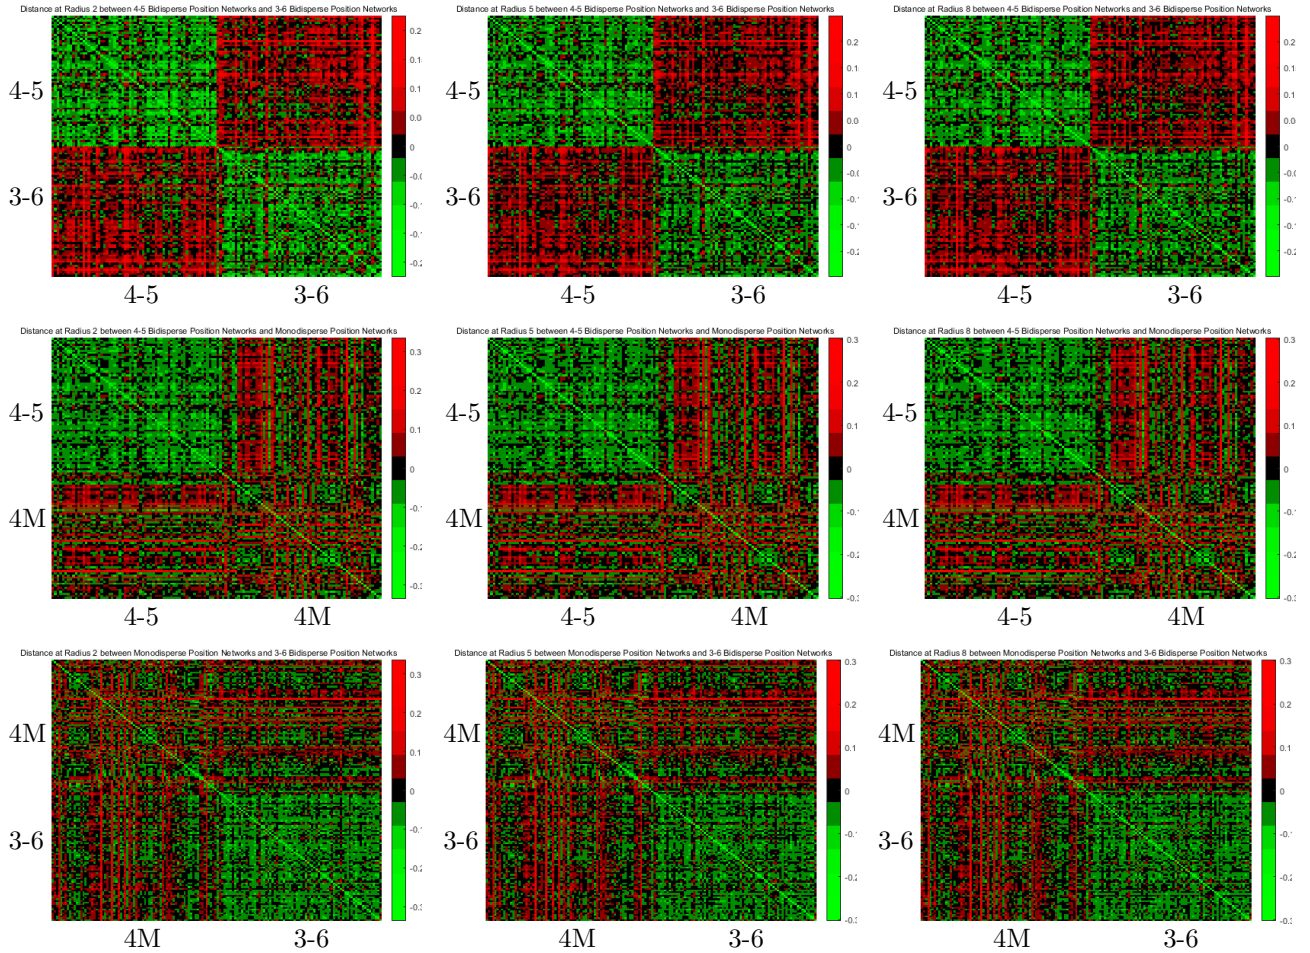consisted of monodisperse simulations.

Figure 12: Results of distances between 2-7 Simulations and 3-6 Bidisperse, Radius 4 Monodisperse, and 4-5 Bidisperse Simulations. Each image shows a heat map where each square represents a distance between two cloths at radius 2, 5, or 8 between simulations either of the same or different types. The green indicates that the distance is below the mean distance of the matrix, while the red indicates a distance above the mean. Swatch and cloth successfully discerned 2-7 bidisperse networks from all other simulation network types.

3.2   4-5 Bidisperse Simulations of 100, 400, 600, and 1000 Balls

We ran bidisperse simulations of radii 4 and 5 that had 100 balls, 400 balls, 600 balls, and 1000 balls to determine whether the size of the network affected the cloths. The resulting heat maps of the distance matrices are shown in Figure 14. For this statistical analysis,

Figure 13: Results of distances between 3-6 Bidisperse, Radius 4 Monodisperse, and 4-5 Bidisperse Simulations. Each image shows a heat map where each square represents a distance between two cloths at radius 2, 5, or 8 between simulations either of the same or different types. The green indicates that the distance is below the mean distance of the matrix, while the red indicates a distance above the mean. These figures show that while swatch and cloth was able to pick up a difference between the 3-6 and 4-5 bidisperse simulation networks, it was unable to distinguish between the 4 monodisperse and the 3-6 and 4-5 bidisperse simulation networks.

because there were only 27 jammed simulations of 400 balls, 8 jammed simulations of 600 balls, and 2 jammed simulations of 1000 balls, we took the reference set for comparisson only from the set of simulations of 100 balls. In testing the hypotheses, we found significant differences in the networks of simulations of size 100 balls and the networks of sizes 400 and 1000 balls, while the networks of simulations of size 100 balls and of size 600 balls had no significant difference. These results are surprising because while the simulations were sized differently, we hypothesized that there would be no distinguishable difference in the networks, as the simulations were of balls with the same radii and densities. The variance in the networks could be cited due to the increasing number of swatch types that follows from increasing the network size or the different epsilon values used across the simulation sizes.

3.3   Communities

In total, 113 internal communities were identified from the 27 bidisperse simulations of 400 balls, 49 internal communities were identified from the 8 bidisperse simulations of 600 balls, and 1 internal community was identified from the two simulations of 1000 balls. Communities were compared both to each other and to the whole networks from which they came using the same technique of reference sets and test sets as before. For the internal communities from networks of 400 particles, there was a significant difference in the distance between 400s communities and the distance between 400s communities and communities from networks of size 600 and 1000. On the other hand, communities identified in networks of size 600 particles were not significantly closer to themselves in distance than they were to communities from simulations of 400 and 1000 balls. In comparing communities to the whole networks from which they were identified, 400s communities were significantly different from the whole 400 networks, while 600s communities were not significantly different from the whole networks of 600 balls. This indicates that there may be isolated behavior happening within the 400 networks and that there may be more uniform behavior in the 600 networks, although there is insufficient data to conclude this definitively. There was also insufficient data to extensively analyze 1000s communities and whole networks, but the distance between the
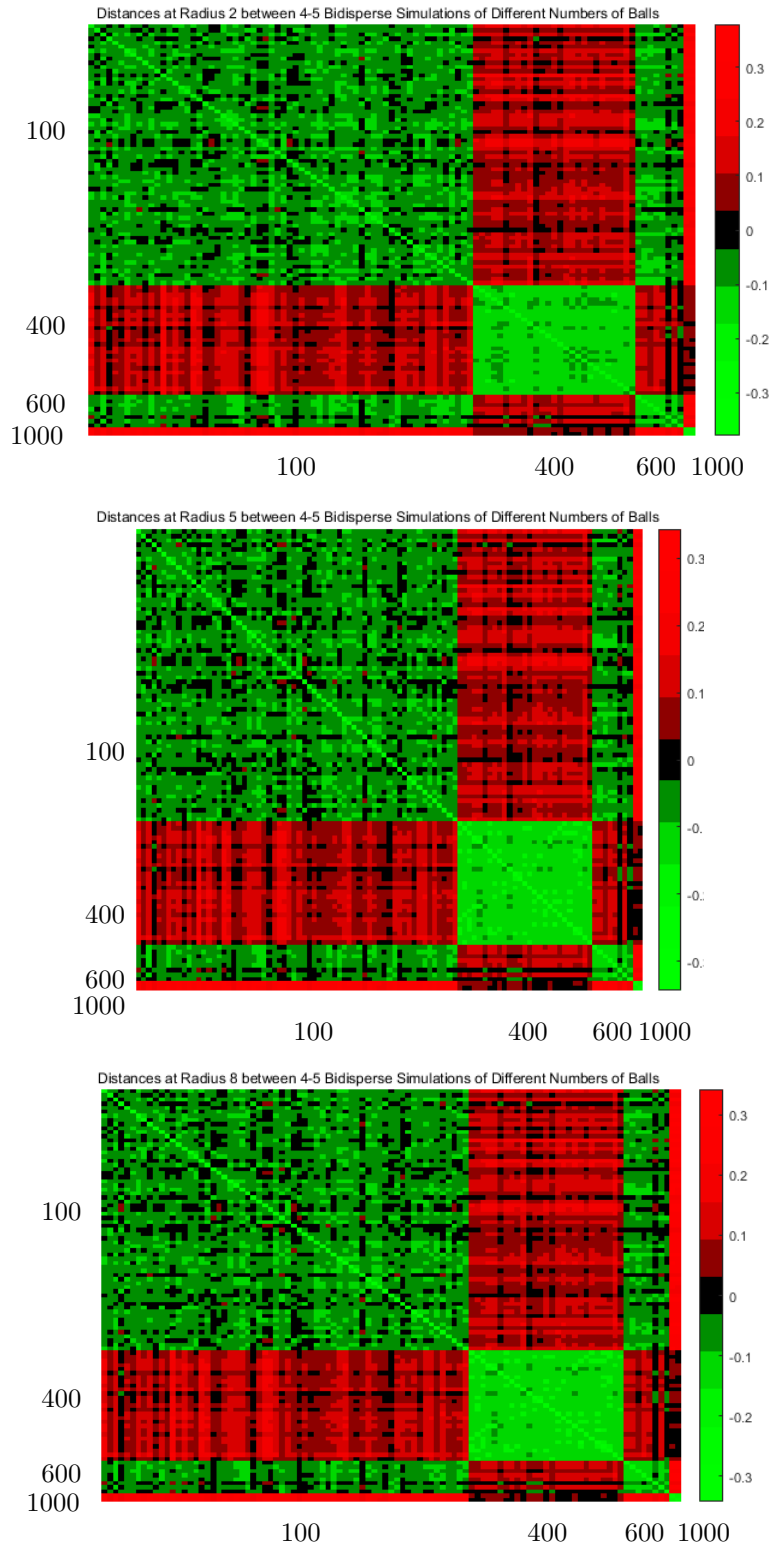
Figure 14: Heat maps for all bidisperse simulations of balls of radii 4 and 5, of sizes 100, 400, 600, and 1000. The heat maps support that the 400 and 1000 ball simulations were significantly distanced from the 100 simulations, while the 600 simulations were not.
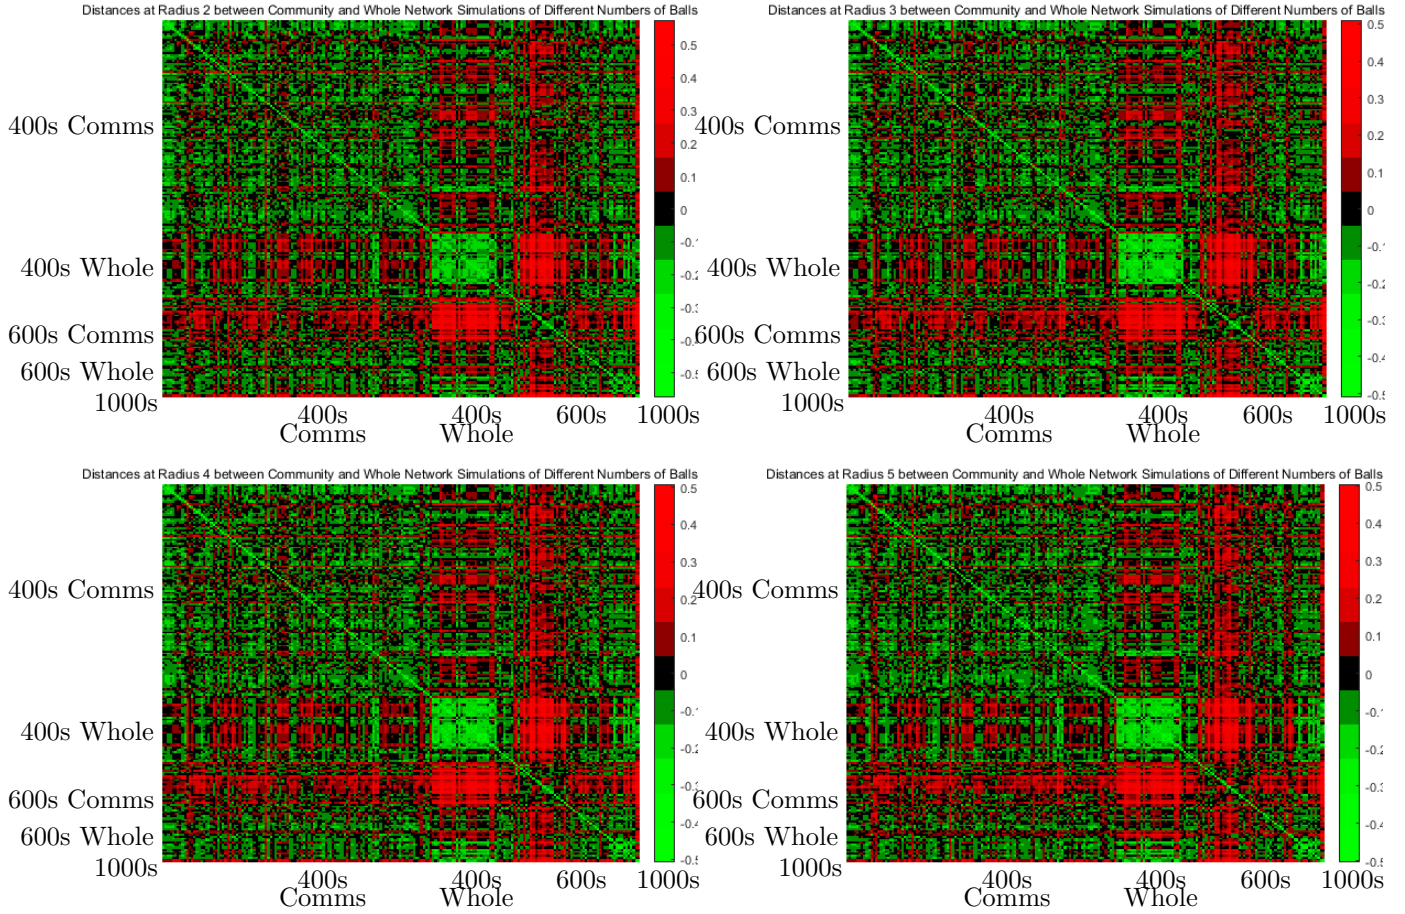
23

Figure 15

whole networks was almost eight times smaller than the distance between each network and the internal community. Figure 15 summarizes the distances between all communities and all their whole networks.

One interesting general result was that any networks compared were either close in distance across all level of the cloth or far in distance across all cloth levels. There were no networks that, when compared, were close or indistinguishable at the smaller levels of the cloth and then farther apart at larger levels. This poses the question of whether two such networks exist in granular materials in regards to monodisperse and bidisperse distributions.

4 Conclusion and Further Work

This work is a first step of applying swatch and cloth to force networks in granular materials to outline their bulk properties.

First, we believe that as larger networks hold more information regarding the behavior of a force network, especially since 100 macroscopic particles is a very small number in comparisson to the millions of Voronoi points Schweinhart *et al* observed [9]. We are interested to see how our observations compare to networks that have thousands and tens of thousands of macroscopic particles. In the same vein, observing the communties of these larger networks would be fascinating and comparing them to each other and to those from smaller networks. The communities calculated for most of the networks tended to be compact rather than contain true force chains; using a different value of ϵ or using a different null model may impact the types of communities that are partitioned and thus their distances from each other and the entire network.

Second, there is more to explore along the study of bidisperse simulations. In our simulations, we assumed a 50-50 distribution of small and large paticles in the bidisperse simulations. One could examine how the cloth of a network is affected when that distribution is changed, for example, to 20-80, or 40-60, or 75-25. In these simulations, one could look again at different particle sizes within a distribution or look at same particle sizes across distributions, using swatch and cloth to assess the differences in the networks arising from both cases. Further, in this data set, swatch and cloth was only able to distinguish bidisperse simulations of different sizes from each other. It would be interesting to consider simulations with three or more particle sizes. Not only could one look at the difference in networks along polydisperse distributions, but also one can examine whether swatch and cloth can distinguish polydisperse simulations from bidisperse or monodisperse simulations.

Finally, there are many questions surrounding characteristics other than sizing of the simulation particles that likely affect bulk properties. Another route of research would be to assess the ability of swatch and cloth to pick up on other possible characteristics of a granular material, such as inelastic collisions, particles of different shapes, various coefficients of static

and kinetic friction, and polydisperse simulations. Should swatch and cloth identify these characteristics, one could potentially predict swatch and cloth behavior for various materials and thus predict how to manipulate bulk properties. Another interesting approach would be to take actual force network data from granular materials and attempt to identify different cloths based on the bulk properties of the materials and characteristics of the material. While simulations give a glimpse as to how swatch and cloth assesses granular materials, they are merely simulations and are no substitutes for the granular materials themselves.

We hope that this research helps to solidify swatch and cloth as a legitimate method for mapping random cell networks in nature and that others will be inspired to apply it to other disciplines where such networks arise.

5 Acknowledgments

This thesis was completed at Department of Mathematics and Computer Science at the College of the Holy Cross.

References

[1] Danielle S Bassett, Eli T Owens, Mason A Porter, M Lisa Manning, and Karen E Daniels. Extraction of force-chain network architecture in granular materials using community detection. *Soft Matter*, 11(14):2731–2744, 2015.

[2] H. Edelsbrunner and J. Harer. *Computational Topology: An Introduction*. Applied Mathematics. American Mathematical Society, 2010.

[3] Stefan Hedman. Smooth and non-smooth approaches to simulation of granular matter. 2011.

[4] Inderjit S. Jutla, Lucas G. S. Jeub, and Peter J. Mucha. A generalized louvain method for community detection implemented in matlab. `http://netwiki.amath.unc.edu/GenLouvain`, 2011-2014.

[5] Mirolsav Kramar, Arnaud Goullet, Lou Kondic, and Konstantin Mischaikow. Persistence of force networks in compressed granular media. *Physical Review E*, 87(4):042207, 2013.

[6] Brendan D. McKay and Adolfo Piperno. Practical graph isomorphism, {II}. *Journal of Symbolic Computation*, 60(0):94 – 112, 2014.

[7] Stepen Sahrun. Elastic collision ball simulation in matlab. `https://zerocrossraptor.wordpress.com/my-shared-resources/matlab-fun-toolbox/`, 2012.

[8] Benjamin Schweinhart. *Statistical Topology of Embedded Graphs*. PhD thesis, Princeton University, 2015.

[9] Benjamin Schweinhart, Jeremy K. Mason, and Robert D. MacPherson. Topological similarity of random cell complexes and applications. *Physical Review E*, 93(6):062111, 2016.