

Quantifying force networks in particulate systems

Arnaud Goulet^a, Lou Kondic^b, Miroslav Kramár^a, Konstantin Mischaikow^a

^a*Department of Mathematics, Hill Center-Busch Campus, Rutgers University, 110 Frelinghusen Rd, Piscataway, NJ 08854-8019, USA*

^b*Department of Mathematical Sciences, New Jersey Institute of Technology, University Heights, Newark, NJ 07102*

Abstract

We present mathematical models based on persistent homology for analyzing force distributions in particulate systems. We define three distinct chain complexes: *digital*, *position*, and *interaction*, motivated by different capabilities of collecting experimental or numerical data, e.g. digital images, location of the particles, and normal forces between particles, respectively. We describe how algebraic topology, in particular, homology allows one to obtain algebraic representations of the geometry captured by these complexes. To each complex we define an associated force network from which persistent homology is computed. Using numerical data obtained from molecular dynamics simulations of a system of particles being slowly compressed we demonstrate how persistent homology can be used to compare the geometries of the force distributions in different granular systems. We also discuss the properties of force networks as a function of the underlying complexes, and hence, as a function of the type of experimental or numerical data provided.

1. Introduction

Particulate systems consisting of a large number of particles have attracted significant attention in the last decades. Despite significant research on these systems, their properties are still not well understood and some of them appear to be rather elusive. The fact that the forces do not propagate uniformly in the systems made of interacting particles has been established in a number of different systems, ranging from those where ‘particles’ are on atomic or molecular scales, to those with macroscopic particles, see, e.g., [1, 2, 3]. It is well accepted that the interparticle forces play a key role in determining the mechanical properties of static and dynamic systems; see e.g. [4] for an extensive review of the role of interaction networks in the context of amorphous solids. However there are no universal methods for describing and quantifying relevant aspects of the interparticle forces. For example, even the concept of a ‘force chain’ (loosely speaking, a connected set of particles interacting by a larger than average force), commonly used in granular community, is not precisely defined.

One approach is to consider particle interactions from statistical point of view. For example, the works by Radjai and collaborators, see, e.g. [5, 6], discussed the differences in the probability density functions of strong and weak force networks (distinguished by the forces being larger or smaller than the average one) arising in simulations; Behringer and collaborators explored these networks in

the systems built from photoelastic particles, see e.g. [3]. Only recently, attempts have been made to move beyond purely statistical description and consider in more detail the properties of these networks. Examples of recent studies include works by Tordesillas and collaborators where a detailed discussion of the forces between particles have been presented, see [7, 8] and the references therein. These studies have explained a number of features of the force networks, but have also uncovered significant complexity of the particle interactions, which is difficult to address by considering the features of the force network based on local-type of analysis. Very recently, initial attempts have been made to consider the interparticle forces based on network-type of analysis [9, 10], and of by exploring their topological properties [11].

With the goal of extending our ability to systematically explore network properties in greater detail, the present paper develops rigorous mathematical models capable of capturing geometric features of particle interactions. There are a variety of concerns that need to be addressed by these models. Three prominent issues are the ability to interrogate large data sets, the form of data that is available to be analyzed, and the ability to quantify a multitude of geometric structures that range over a variety of different force levels.

The approach that we present is based on algebraic topology and in particular persistent homology [12, 13]. This is a relatively new mathematical technique that provides a computationally efficient rigorous framework for multi scale analysis. The computational efficiency is essential since the goal is to apply these techniques to large data sets. Another key feature of persistent homology is that it reduces a scalar function to a persistence diagram, which is a collection of points in the plane where each

Email addresses: arnaud.goulet@gmail.com (Arnaud Goulet), kondic@njit.edu (Lou Kondic), miroslav@math.rutgers.edu (Miroslav Kramár), mischaik@math.rutgers.edu (Konstantin Mischaikow)

point encodes well defined geometric information about the function, but does not rely on a particular choice of threshold to do this. In the context of particulate systems, this means that a specific definition of ‘force-chains’ or similar objects is not required. Furthermore, there are a variety of metrics that can be imposed on the space of persistence diagrams such that in the context of particulate systems the application of persistent homology can be interpreted as a continuous non linear projection of the force networks to the space of persistence diagrams.

These properties of persistent homology suggest that it is a good tool to study both the static and the dynamical aspects of experimental and computational realizations of particulate systems. We have shown that persistent homology can be used to identify and quantify the difference in the force networks of dense granular materials (DGM) made of particles with different physical properties [14] and we are currently exploring its use in the analysis of dominant time scales and the structure of attractors in systems undergoing time dependent forcing. However, this analysis is predicated upon being able to take the experimental or numerical data, produce persistence diagrams, and understand the confidence with which one can interpret the results. This process is the focus of this paper.

As indicated above, persistent homology is based on algebraic topology and thus its computation is based on the construction of a finite complex. Therefore, the first issue that is addressed (see Section 2) is the construction of appropriate complexes from the experimental or numerical data. The answer depends upon the type of data that is provided, which in turn is dependent upon the method by which the data is obtained. With this in mind we propose three different complexes: *digital*, *position*, and *interaction*. One of the results that we present is that the information that can be extracted via the use of an interaction network is significantly more reliable than that of a digital or position network. The interaction network can be used in the setting of numerical simulations or particular types of experiments [3] where complete information about the forces between adjacent particles may be known. However, for many experiments only the total force experienced by a particle may be available [15]. This necessitates the use of a digital or position network, depending upon how the data is collected and physical properties of the individual particles.

The digital and position complex can, in a weak sense, be thought of as providing a geometric structure for the underlying particle network. The interaction complex is an abstract mathematical space. This complex serves as an underlying space for modeling all possible force interactions between the particles. To include the information about forces requires the construction of the associated force networks (see Section 4). Again, the construction is based on two constraints, the type of data available and the requirement that we can use the force networks to compute persistent homology.

We do not assume that the reader is familiar with alge-

braic topology. Section 3 provides a simple review of the basic ideas of homology theory and Section 5 describes persistent homology and persistence diagrams. In Section 6 we discuss the space of persistence diagrams, in particular the appropriate metrics on the space, and we provide a theorem that justifies the claim that the interaction force network is optimal (of course, this depends on having the appropriate data).

In Section 7 we conclude with a review of the developed concepts in the context of DGM data obtained by discrete element based simulations (DES). We choose to work with these numerical simulations since all the data is available with high precision and therefore we can process it through all three force network constructions (digital, position, and interaction). This allows for greater clarity in interpreting the geometric meaning of the persistence diagrams, greater ease in comparing the results of the different networks, and simplicity in testing for stability with respect to perturbations. The reader who is familiar with the language of persistent homology may wish to skip directly to this section, before examining the details of the constructions.

There are several points that we encourage the reader to keep on mind while reviewing Section 7. First, we provide a wide variety of figures of the force networks in this section to help develop intuition in interpreting the persistence diagrams. However, the information itself is contained completely within the persistence diagrams. This is important since current technology makes the study, either experimentally or numerically, of 3 dimensional particulate systems feasible. Visually extracting information from these higher dimensional systems is much more difficult (this is, in fact, one of the reasons that we choose to present our results with 2 dimensional systems of particles). Second, in our examples we make use of the magnitude of the normal force, but in the context of DGM there are other options such as the magnitude of tangential force and in principle any function that assigns a scalar value to every edge can be used. Finally, and even more generally, though this paper is focussed on DGM, the constructions are independent of the details of particle-particle interaction and could as well be applied to any other system consisting of interacting particles.

2. Particle Networks

As is indicated in the Introduction, the first step towards using algebraic topology to characterize the geometric structures associated with DGM is to create complexes through which the geometry of the particles can be expressed. We introduce three complexes motivated by the type of data commonly obtained from experiments or numerical simulations. Consider Figure 1(a) that shows a small portion of an image derived from physically shearing a collection of photoelastic disks (courtesy of R.P. Behringer, unpublished results).

Digital This figure arises from a digital image and thus the data can be viewed as a collection of a large number of pixels.

Positions Since this is a controlled experiment involving circular disks, to determine the configuration of the particles it is sufficient to know the the locations of their center points and their radii.

Interactions The particles are made of photoelastic material and thus the light intensities within the particles can be used to determine the normal forces between the particles.

We remark that modern imaging technologies make it easy to extract pixel or voxel data from two and even three dimensional granular systems. Furthermore, the natural a priori assumptions that need to be made concerning the geometry of the particles appear to be relatively modest; the imaging resolution should be fine enough to recognize individual particles. In contrast appropriate use of position data requires that the particles have a simple, easily characterized geometry. Clearly the interaction data is the most difficult to obtain, requiring clever use of special materials and/or technology by the experimentalist.

Our approach is to encode each of these data types into a different complex and one of the goals of this paper is to make clear the difference in form and quality of the geometric information that can be reliably extracted from each approach. One not particularly surprising conclusion is that the interaction data provides the best information and the digital data the worst, but it is worth quantifying these differences. With this in mind we begin with several formal definitions. Our focus is on physical systems, thus for the most part we restrict our discussion to two and three dimensional complexes (see [12, 16] for a more general discussion).

To define a simplicial complex CN we begin with a finite set of vertices $\text{CN}^{(0)} := \{v_i \mid i = 0, \dots, I\}$. An n -dimensional simplex in CN is a subset of $\text{CN}^{(0)}$ consisting of $n + 1$ vertices. The set of n -dimensional simplices in CN is denoted by $\text{CN}^{(n)}$. Given the set of vertices $\text{CN}^{(0)} := \{v_i \mid i = 0, \dots, I\}$ it is customary to denote the 0-dimensional simplices by $\langle v_i \rangle$, the 1-dimensional simplices by $\langle v_i, v_j \rangle$, and the 2-dimensional simplices by $\langle v_i, v_j, v_k \rangle$. A one and two dimensional simplices are referred to as *edges* and *triangles*. A *simplicial complex* CN is a collection of simplices that satisfies the following property: if $\sigma \in \text{CN}$ and $\sigma' \subset \sigma$, then $\sigma' \in \text{CN}$.

Definition 2.1. Given a collection of circular disks $\{p_i \mid i = 0, \dots, I\}$, location of their centers $\{x_i \in \mathbb{R}^2 \mid i = 0, \dots, I\}$, and their radii $\{r_i \mid i = 0, \dots, I\}$ the associated *position complex* CN_P is the simplicial complex consisting of vertices $\{v_i \mid i = 0, \dots, I\}$, where each vertex v_i is identified with particle p_i , and edges $\langle v_i, v_j \rangle$ if and only if $\|x_i - x_j\| \leq r_i + r_j$.

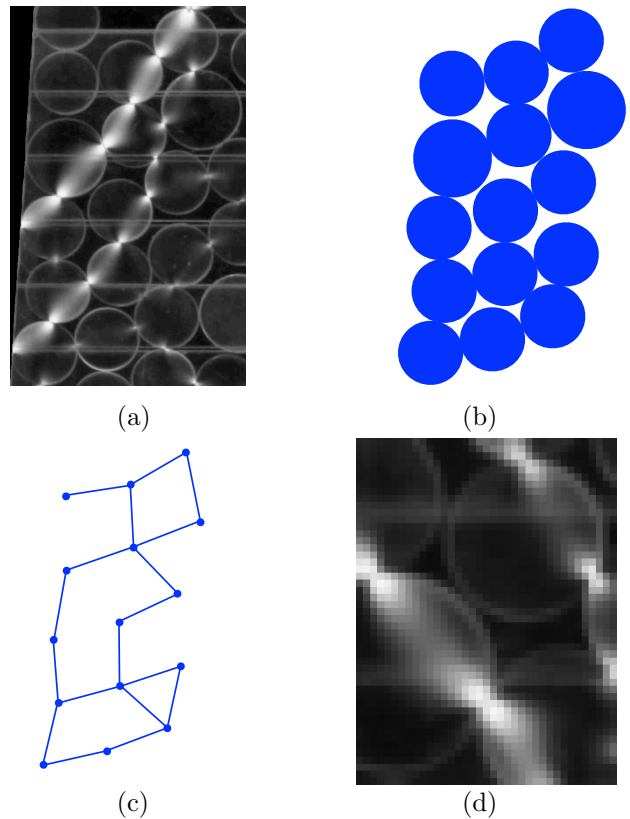


Figure 1: Different representations of the particle networks derived from experimental data. For simplicity we neglect the particles that intersect the edges of the picture. (a) Small portion of a digital image of an experimental system. (b) Idealized extraction of particles from image. (c) Position network. (d) Small portion of Digital network showing individual pixels. The interaction network has no immediate geometric meaning and hence is not shown.

Note that the position complex as defined above agrees with the unweighted (binary) network of [10] and the contact network of [8].

For the sake of clarity Definition 2.1 is presented in the context of the examples considered in this paper. More generally, one can consider spherical particles positioned in \mathbb{R}^d , $d = 3$ being the most relevant for physical applications. As presented the position complex is an abstract simplicial complex; that is, there is no specific geometric object associated with it. In the context of this work we can always geometrize the complex by declaring the vertices to be the points $\{x_i \in \mathbb{R}^2 \mid i = 0, \dots, I\}$ and the edges to be the line segments connecting the points. From now on we rarely distinguish between the abstract simplicial complex and its geometric realization.

Definition 2.2. Given a collection of particles $\{p_i \mid i = 0, \dots, I\}$ the *interaction complex* CN_I is the simplicial complex consisting of vertices $\{v_i \mid i = 0, \dots, I\}$ where each vertex v_i is identified with particle p_i , all edges $\langle v_i, v_j \rangle$, and all triangles $\langle v_i, v_j, v_k \rangle$.

There are two remarks to be made at this point. First, Definition 2.1 gives rise to simplicial complexes that are graphs. We introduce and make use of higher dimensional complexes at the end of this section. Second, while the position complex can be given a geometric interpretation, this is not the case for the interaction complex. The latter complex provides a suitable underlying space for studying force interactions between the particles. Justification of this statement is given in Section 5.

To interpret the pixel data we make use of cubical complexes. Observe that once the pixel data is obtained the actual size used to represent each pixel is no longer an issue. Thus, for the sake of simplicity of discussion and without loss of generality we assume that the pixel data is embedded in \mathbb{R}^2 with each pixel being represented by a square defined by the integer lattice. More precisely, a *2-dimensional cube (pixel)* is a square of the form $[n, n + 1] \times [k, k + 1]$, a *1-dimensional cube (edge)* is a unit interval of the form $[n, n] \times [k, k + 1]$ or $[n, n + 1] \times [k, k]$, where $n, k \in \mathbb{Z}$, and a *0-dimensional cube (vertex)* is a point with integer coordinates. As in the simplicial setting a *cubical complex* CN is a collection of cubes that satisfy the following property: if $\sigma \in \text{CN}$ and $\sigma' \subset \sigma$, then $\sigma' \in \text{CN}$.

Definition 2.3. Given a digital image of particles $\{p_i \mid i = 0, \dots, I\}$ the *digital complex* CN_D is the cubical complex consisting of squares $\{\sigma_j\}$ where each square σ_j represents a single pixel associated with some particle.

How pixels are associated with particles is intentionally left vague in Definition 2.3. The actual association depends on the particular characteristics of the imaging device, filtering, thresholding, etc used to obtain and process the data. Conceptually, the most straightforward approach is to discretize the domain of the image into squares, identify the squares with pixels, and declare the pixel to represent a particle if the associated square intersects the particle. A cubical particle network CN_D is shown in Figure 1(d). We provide more detail about the construction later.

The cubical structure can be applied to voxels, by representing each voxel as a unit cube of the integer lattice in \mathbb{R}^3 (see [16] for the general theory).

Having defined these complexes a reasonable first question is whether they correctly capture the topology of the particle configuration. We begin with the following positive result under the assumption that the particles cannot deform under the pressure induced by contacts with other particles.

Proposition 2.4. *Given a collection of circular hard disks $\{p_i \mid i = 0, \dots, I\}$, location of their centers $\{x_i \in \mathbb{R}^2 \mid i = 0, \dots, I\}$, and their radii $\{r_i \mid i = 0, \dots, I\}$ the associated position complex CN_P is homotopic to the union of the regions occupied by the particles, $\bigcup_{i=0}^I p_i$.*

The proof follows from retracting the set of particles onto the geometric realization of CN_P (see for example




Figure 2: With complete information the position complex CN_P has the same homotopy type as the configuration space of the particles $\bigcup_{i=0}^I p_i$. The proof involves collapsing the particles onto graph.




Figure 3: Failure of the digital complex to correctly capture the topology of the particle configuration. (a) The particle configuration consists of two components and contains no loops. The associate digital complex has one component and one loop. (b) The particle configuration contains one loop while the digital complex contains none. (c) The particle configuration contains one component and no loops. The associate digital complex has one component and one loop. Furthermore, doubling the resolution of the camera does not extinguish the existence of an unwanted loop.

Figure 2). We do not provide details of the proof because this result is of limited importance. In any experiment or numerical simulation the locations of the particles can only be given up to some specified precision. If we assume the particles to be hard, then two particles p_i and p_j are in contact if and only if $\|x_i - x_j\| = r_i + r_j$. Clearly, arbitrarily small errors in x_i and x_j can lead to an inequality which indicates that the particles are not in contact. The same argument applies to arbitrarily small errors in the measurements of the radii of the disks. Assuming that the particles are soft makes this result slightly more robust, but this is tempered by the fact that this stability depends on the existence of sufficiently large normal forces. We attempt to quantify these comments in Section 7.2.

To measure the topological fidelity of the digital complex requires the choice of a rule for determining if a pixel is included in the complex or not. For the sake of clarity we continue with the conceptually simple rule introduced earlier. As is indicated by Figure 3 the failure of the digital complex CN_D to correctly capture the topology of the particle configuration can be quite dramatic. Even the simplest setting of two particles with a high pixel resolutions does not guarantee a correct topological description. Figures 3(a) and (c) demonstrates that both the number




Figure 4: Two configuration spaces with position complexes CN_P . (a) Crystalline structure with many loops formed by three particles. (b) Noncrystalline structure with fewer loops than (a). (c) Flag complex $\text{CN}_P^\blacktriangle$ derived from (b). Observe that the only loops which remain are associated with defects.

of connected components and holes can be counted incorrectly. Figures 3(a) and (b) show that there is no particular direction to the error in the hole count. The fact that the number of components of the digital complex is never larger than that of the particle configuration arises from our assumption on how to identify pixels with particles. In particular our approach leads to an artificial expansion of the area covered by each particle. Thus, two separate particles can appear to be in contact, but two particles that are in contact can never appear to be separated.

To keep things in perspective we remind the reader that even though it is clear that the digital complex can fail to record the correct topology in a variety of ways it is the easiest means of collecting data and is applicable in situations in which the forces between the grains cannot be directly measured and without a priori assumptions about the geometry and rigidity of the grains.

As is indicated in Figure 4 if we restrict our definition of the position complexes to graphs, we will detect many holes that may not be of interest. In particular, Figure 4(a) suggests particles in a *crystalline* structure in which CN_P has 9 holes. This should be contrasted with Figure 4(b) in which there are only 4 holes, but three of them indicate that the particles are not as densely packed as possible. Since in a perfect densely packed crystalline structure made up of disks of the same size all holes would be made up of exactly 3 particles we refer to a hole involving four or more particles as *defect*.

To remedy this counting of ‘uninteresting’ loops or equivalently to only count defects we construct the flag complex $\text{CN}_P^\blacktriangle$ as follows. Set

$$\text{CN}_P^{\blacktriangle(0)} := \text{CN}_P^{(0)} \quad \text{and} \quad \text{CN}_P^{\blacktriangle(1)} := \text{CN}_P^{(1)}$$

and

$$\langle v_i, v_j, v_k \rangle \in \text{CN}_P^{\blacktriangle(2)}$$

if and only if

$$\{\langle v_i, v_j \rangle, \langle v_i, v_k \rangle, \langle v_j, v_k \rangle\} \subset \text{CN}_P^{(1)}.$$

Remark 2.5. If one is working with particles in \mathbb{R}^3 , then the same arguments can be used to justify the use of the

flag complex where $\langle v_i, v_j, v_k, v_l \rangle \in \text{CN}_P^{\blacktriangle(3)}$ if and only if

$$\{\langle v_i, v_j, v_k \rangle, \langle v_i, v_j, v_l \rangle, \langle v_i, v_k, v_l \rangle, \langle v_j, v_k, v_l \rangle\} \subset \text{CN}_P^{\blacktriangle(2)}.$$

It is worth noting that an analogous approach will not work in the context of digital complexes. A square can be missing from CN_D because of the phenomena indicated in Figure 3(a) or (c), but it can also be missing because it represents the loop formed by four distinct particles.

3. Homology

We pause in our development of the networks to review a few fundamental definitions from the classical theory of homology with a focus on the simple setting of the digital, position and interaction complexes introduced in Section 2. For a more general discussion the reader is referred to a standard text in algebraic topology or to [12, 16] for descriptions more closely associated with data analysis.

Recall that position complexes CN_P and interaction complexes CN_I are simplicial complexes. This leads to our use of simplicial homology. Since we are working with planar arrangements of particles it is sufficient to use \mathbb{Z}_2 coefficients, i.e. the set $\{0, 1\}$ with the standard binary addition and multiplication operations. Recall that $\text{CN}^{(n)}$ denotes the set of n -dimensional simplices in the simplicial complex CN . The n -chains of CN is defined to be the vector space

$$C_n(\text{CN}) := \left\{ \sum_{\sigma \in \text{CN}^{(n)}} m_\sigma \sigma \mid m_\sigma \in \mathbb{Z}_2 \right\}. \quad (1)$$

Observe that $C_n(\text{CN})$ is the vector space over \mathbb{Z}_2 with basis elements consisting of the n -dimensional simplices.

The associated *boundary maps* are linear maps (these are often represented as matrices using the simplicies as bases) $\partial_n: C_n(\text{CN}) \rightarrow C_{n-1}(\text{CN})$ ($C_{-1}(\text{CN}) := 0$) defined on the simplices as follows

$$\begin{aligned} \partial_0 \langle v_i \rangle &:= 0 \\ \partial_1 \langle v_i, v_j \rangle &:= \langle v_i \rangle + \langle v_j \rangle \\ \partial_2 \langle v_i, v_j, v_k \rangle &:= \langle v_i, v_j \rangle + \langle v_i, v_k \rangle + \langle v_j, v_k \rangle. \end{aligned}$$

A direct calculation making use of the linearity and the use of \mathbb{Z}_2 coefficients show that $\partial_{n-1} \circ \partial_n = 0$, e.g.

$$\begin{aligned} &\partial_1 \circ \partial_2 \langle v_i, v_j, v_k \rangle \\ &= \partial_1 (\langle v_i, v_j \rangle + \langle v_i, v_k \rangle + \langle v_j, v_k \rangle) \\ &= \partial_1 \langle v_i, v_j \rangle + \partial_1 \langle v_i, v_k \rangle + \partial_1 \langle v_j, v_k \rangle \\ &= \langle v_i \rangle + \langle v_j \rangle + \langle v_i \rangle + \langle v_k \rangle + \langle v_j \rangle + \langle v_k \rangle \\ &= 0. \end{aligned}$$

The boundary maps can be used to identify components and loops. To do this we focus on *cycles*, these are chains which are sent to the 0 vector under ∂_n . More formally,

$$Z_n(\text{CN}) := \ker \partial_n.$$




Figure 5: Three different 1-dimensional chains. Each of these chains corresponds to a loop and hence is a cycle. The [color used]red chain c_1 and the [color used]green chain c_2 correspond to the same hole in the particle network. The brown loop c_3 does not correspond to any hole and can be contracted to a point.

Observe that $\langle v_i \rangle \in Z_0(\text{CN})$ and $\langle v_i, v_j \rangle + \langle v_i, v_k \rangle + \langle v_j, v_k \rangle \in Z_1(\text{CN})$.

The power of homology is that we are able to move from geometric data to an algebraic format from which we can then extract geometric information. For example, the algebraic statement that $\langle v_i \rangle \in Z_0(\text{CN})$ can be interpreted as a statement that $\langle v_i \rangle$ identifies a component of CN. Similarly, $\langle v_i, v_j \rangle + \langle v_i, v_k \rangle + \langle v_j, v_k \rangle \in Z_1(\text{CN})$ can be identified with the path of edges $\langle v_i, v_j \rangle, \langle v_i, v_k \rangle, \langle v_j, v_k \rangle$ that makes up a loop. To emphasize the relationship between the algebra and geometry consider the simplicial complex indicated in Figure 5. There are three chains that form loops and hence cycles indicated in [color used]red, green and brown.

For obvious reasons it is important not to over count components or loops. In particular, if an edge $\langle v_i, v_j \rangle$ belongs to CN, then $\langle v_i \rangle$ and $\langle v_j \rangle$ belong to the same component and therefore we wish to identify them. This can be done algebraically by the relation $\partial_1 \langle v_i, v_j \rangle = \langle v_i \rangle + \langle v_j \rangle$. Similarly, if a 2-dimensional simplex $\langle v_i, v_j, v_k \rangle \in \text{CN}$, then the loop $\langle v_i, v_j \rangle, \langle v_i, v_k \rangle, \langle v_j, v_k \rangle$ does not enclose a hole and thus should not be counted. Again, this can be detected algebraically by the relation $\partial_2 \langle v_i, v_j, v_k \rangle = \langle v_i, v_j \rangle + \langle v_i, v_k \rangle + \langle v_j, v_k \rangle$. Observe that the relations in these examples are obtained via images of the boundary operator. This leads to the definition of the *boundaries* of CN,

$$B_n(\text{CN}) := \partial_{n+1}(C_{n+1}(\text{CN})).$$

Referring to the complex depicted in Figure 5 observe that there exists \bar{c} such that $\partial_2 \bar{c} = c_1 + c_2$ which implies that the cycles c_1 and c_2 represent the same hole in the complex. This motivates the following definition. The n -th *homology group* of the simplicial complex CN is defined by

$$H_n(\text{CN}) := \frac{Z_n(\text{CN})}{B_n(\text{CN})}$$

it is the vector space of equivalence classes of cycles identified by boundaries. To be more specific given a cycle $z \in Z_n(\text{CN})$ the associated homology class $[z] = [z]_{\text{CN}}$ is

the equivalence class of all cycles of the form $z + b$ where $b \in B_n(\text{CN})$.

The dimension of $H_n(\text{CN})$ is called the n -th *Betti number* $\beta_n(\text{CN})$. $\beta_0(\text{CN})$ counts the number of components and $\beta_1(\text{CN})$ counts the number of loops which encircle a void. Because the simplicial complex CN_P^Δ can be embedded in the plane, $\beta_2(\text{CN}_P^\Delta) = 0$. The homology of CN_I is more complicated, though from the definition one can determine that

$$\beta_n(\text{CN}_I) \cong \begin{cases} 1 & \text{if } n = 0 \\ 0 & \text{if } n = 1. \end{cases} \quad (2)$$

For the sake of simplicity we have presented homology in the context of a 2-dimensional simplicial complexes. However, homology can be extended to arbitrary dimension and defined in the context of very general topological spaces. One of the fundamental properties is that if two topological spaces are homotopic, then they have the same homology groups. A corollary of this is that under the hypothesis of Proposition 2.4 the Betti numbers of CN_P agree with the Betti numbers of the space defined by $\cup_{i=0}^I p_i$. However, as is discussed in Section 2 we cannot expect this to be the case for CN_I , and in practice this is rarely the case for CN_P or CN_D .

We do not present the details of computing homology with cubical complexes. Conceptually the ideas are the same, though the boundary operators are slightly different. The interested reader is referred to [16] for a complete presentation. Even more generally, simplicial and cubical complexes are examples of *chain complexes* and the individual simplices or cubes are examples of *cells*.

The reader may be somewhat underwhelmed by the fact that we have constructed a significant amount of algebra to essentially count components and loops, especially since there are extremely efficient graph theoretic algorithms for performing these operations. However, the algebra allows us to compare loops in different complexes. Let CN and CN' be two distinct chain complexes. A collection of linear maps $\phi_n: C_n(\text{CN}) \rightarrow C_n(\text{CN}')$ are *chain maps* if

$$\partial'_n \phi_n = \phi_{n-1} \partial_n$$

for all k where ∂_n and ∂'_n are the boundary maps for $C_n(\text{CN})$ and $C_n(\text{CN}')$. A fundamental result is that if ϕ_n is a chain map, then ϕ_n induces a linear map on homology $\phi_n: H_n(\text{CN}) \rightarrow H_n(\text{CN}')$ defined by

$$\phi_n([z]_{\text{CN}}) := [\phi_n(z)]_{\text{CN}'}$$

For the purposes of this paper it is sufficient to note that if $\text{CN} \subset \text{CN}'$, then the inclusion map induces, for each dimension, a chain map and hence a map on homology.

4. Force Networks

As is indicated in the Introduction it is well accepted that the geometry of force chains plays an important role

in determining the macroscopic properties of dense granular material. In this section we expand on the complexes constructed in Section 2 to include the forces between the particles into this mathematical framework. In the present work we will focus only on the normal force, that is, the component of the force projected on the line connecting the centers of interacting particles. We will view the magnitude of the normal force as a scalar field defined over the complex, i.e. a function $f: \text{CN} \rightarrow \mathbb{R}$. There are two constraints on the definition of f . The first arises from the use of persistent homology to capture the geometry of the force chains. The second is related to the type of information extracted from the experiment or simulation.

To understand the first constraint assume for the moment that we are given a complex CN and a scalar field $f: \text{CN} \rightarrow \mathbb{R}$. Based on the concept of a force chain we are interested in the geometry of a part of the complex on which the forces exceed a specified level. Thus we define a *force network* to be the super level set

$$\text{FN}(f, \theta) := \{\sigma \in \text{CN} \mid f(\sigma) > \theta\} \quad (3)$$

which corresponds to the part of the particle network experiencing force larger than θ . We are using homology to quantify the geometry of $\text{FN}(f, \theta)$ and hence we need $\text{FN}(f, \theta)$ to be a complex for every value of θ . Recall that the crucial property of a complex is: if $\sigma \in \text{CN}$ and $\sigma' \subset \sigma$, then $\sigma' \in \text{CN}$. Thus, in our construction of f we need to insure that this condition is satisfied. This leads to the following definition.

Definition 4.1. Given a complex CN , a function $f: \text{CN} \rightarrow \mathbb{R}$ is *monotone* if $f(\sigma') \geq f(\sigma)$ for every $\sigma', \sigma \in \text{CN}$ such that $\sigma' \subset \sigma$.

It is left to the reader to check that if f is monotone, then $\text{FN}(f, \theta)$ is a complex for every value of θ .

Definition 4.2. Given a complex CN and a monotone function $f: \text{CN} \rightarrow \mathbb{R}$, the associated *force network filtration* is the collection of all force network complexes

$$\{\text{FN}(f, \theta) \mid \theta \in \mathbb{R}\}.$$

The second constraint on the construction of $f: \text{CN} \rightarrow \mathbb{R}$ arises from the assumptions concerning the available information. The weaker assumption, which we associate with digital or position complexes, is that for each particle we can estimate the forces experienced by that particle. The stronger assumption, which leads to the use of an interaction complex, is that we can estimate the forces between any two particles. This leads to the following definitions.

Digital Force Networks. Let $\psi_i \in \mathbb{R}$ denote the magnitude of the force on particle p_i . Recall that a 2-dimensional cube $\sigma \in \text{CN}_D^{(2)}$ if it intersects at least one particle p_i . Define

$$f(\sigma) = \max \{\psi_i \mid \sigma \cap p_i \neq \emptyset\}.$$

The definition of f is extended to the edges and vertices as follows:

$$f(\sigma) = \max \left\{ f(\sigma') \mid \sigma \subset \sigma', \sigma' \in \text{CN}_D^{(2)} \right\}.$$

Position Force Networks. Let $\psi_i \in \mathbb{R}$ denote the magnitude of the force on particle p_i . For each $\langle v_i \rangle \in \text{CN}_P^{(0)}$ define

$$f(\langle v_i \rangle) := \psi_i.$$

Extend the definition of f inductively by

$$f(\langle v_i, v_j \rangle) = \min \{f(\langle v_i \rangle), f(\langle v_j \rangle)\}$$

and in the case of the flag complex CN_P^Δ ,

$$f(\langle v_i, v_j, v_k \rangle) = \min \{f(\langle v_i, v_j \rangle), f(\langle v_i, v_k \rangle), f(\langle v_j, v_k \rangle)\}.$$

Interaction Force Networks. Let $\psi_{i,j} \in \mathbb{R}$ denote the magnitude of the force experienced between particles p_i and p_j . For the interaction network CN_I the natural starting point for the definition of f is on the edges,

$$f(\langle v_i, v_j \rangle) := \psi_{i,j}.$$

The function is extended to the vertices by

$$f(\langle v_i \rangle) = \max \{f(\langle v_i, v_j \rangle) \mid \langle v_i, v_j \rangle \in \text{CN}_I\}$$

and to the 2-dimensional simplices by

$$f(\langle v_i, v_j, v_k \rangle) = \min \{f(\langle v_i, v_j \rangle), f(\langle v_i, v_k \rangle), f(\langle v_j, v_k \rangle)\}.$$

We use the following proposition to summarize the above discussion and constructions.

Proposition 4.3. *Given a complex CN_\bullet , $\bullet \in \{D, P, I\}$ and f defined as above, the associated super level set $\text{FN}_\bullet(f, \theta)$ is a complex for all values of $\theta \in \mathbb{R}$.*

Since we are assuming that there are only finite number of particles in our system, any force network filtration $\{\text{FN}_\bullet(f, \theta) \mid \theta \in \mathbb{R}\}$ contains only finitely many distinct complexes. We can use homology, in particular the Betti numbers, to characterize the geometry of each of the distinct force networks in the force network filtration.

To gain intuition into the force networks consider the interaction force network indicated in Figure 6. Figure 6(a) represents a collection of particles. The particles are represented by the vertices and the shown edges correspond to the non-zero forces between the particles. Figure 6(b) shows only the simplices of CN_I for which the value of the function $f: \text{CN}_I \rightarrow \mathbb{R}$ is positive. The value of f on the edges is determined by the forces between the particles. In Figure 6(a) the non-zero forces are color coded. In increasing value the force is denoted by blue, cyan, green, and red. The value of the function f is extended to the vertices in Figure 6(c) and to the 2-dimensional simplices in Figure 6(d). The figures (e)-(h) indicate the associated force network for non negative values of θ . If $\theta < 0$ then

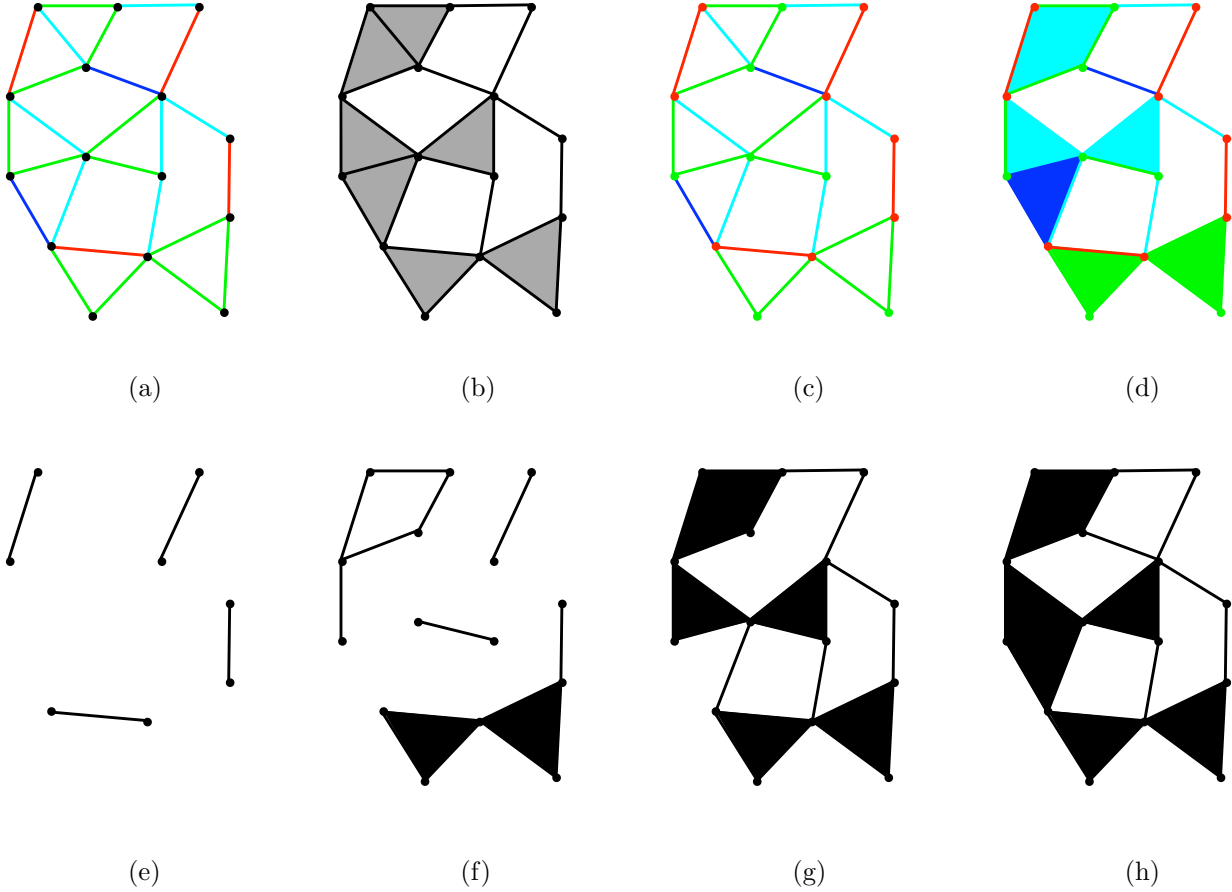


Figure 6: A representation of a simple interaction network FN_I^A . (a) Vertices represent the particles and the edges correspond to the non zero force between the particles. An increasing value of the force is denoted by blue, cyan, green, and red. (b) Collection of simplices on which the function $\text{CN}_I : f \rightarrow \mathbb{R}$ is positive. Extension of the function f to the vertices (c) and to the 2-dimensional simplices (d). (e)-(h) The complexes $\text{FN}(f, \theta_i)$ for positive θ equal to θ_4 (red), θ_3 (green), θ_2 cyan and θ_1 (blue), respectively.

$\text{FN}_I(f, \theta)$ is a full graph on the vertices corresponding to the particles and all the 2-dimensional simplices are filled in.

For example, referring to the force network filtration of Figure 6 we can extract the following data:

$$\begin{aligned} (\beta_0(\text{FN}_I(f, \text{red})), \beta_1(\text{FN}_I(f, \text{red}))) &= (4, 0) \\ (\beta_0(\text{FN}_I(f, \text{green})), \beta_1(\text{FN}_I(f, \text{green}))) &= (4, 1) \\ (\beta_0(\text{FN}_I(f, \text{cyan})), \beta_1(\text{FN}_I(f, \text{cyan}))) &= (1, 3) \\ (\beta_0(\text{FN}_I(f, \text{blue})), \beta_1(\text{FN}_I(f, \text{blue}))) &= (1, 4). \end{aligned}$$

It is worth noting that the H_0 homology information for $\text{FN}_I(f, \text{red})$ and $\text{FN}_I(f, \text{green})$ agree and yet the structure of the components has changed dramatically. Two distinct connected components become one and a new connected component is formed. To capture this information we make use of the fact that these complexes are nested by inclusion. This leads to the concept of persistent homology.

5. Persistent Homology

Given a force network filtration

$$\text{FN}(f, \theta) := \{\sigma \in \text{CN} \mid f(\sigma) \geq \theta\}.$$

generated by a finite number of particles, there is a finite number of values

$$0 = \theta_0 < \theta_1 < \dots < \theta_K = \max_{\sigma \in \text{CN}} f(\sigma)$$

such that $\theta_k = f(\sigma)$ for some $\sigma \in \text{CN}$. Though the Betti numbers characterize the topology of a given force network $\text{FN}(f, \theta_k)$, the vector space structure of homology plays an essential role in that it allows us to compare the topology of $\text{FN}(f, \theta_k)$ with any other force network $\text{FN}(f, \theta_j)$. Given $\theta_i < \theta_j$, $\text{FN}(f, \theta_j) \subset \text{FN}(f, \theta_i)$ and hence there is an inclusion map

$$\iota_{\theta_i, \theta_j} : \text{FN}(f, \theta_j) \rightarrow \text{FN}(f, \theta_i).$$

As is indicated at the end of Section 3, this defines maps

$$\iota_{\theta_i, \theta_j *} : H_*(\text{FN}(f, \theta_j)) \rightarrow H_*(\text{FN}(f, \theta_i))$$

on each homology group H_* . It is important to note that $\iota_{\theta_1, \theta_2 *}$ need not be an inclusion map on the level of the homology groups.

Persistent homology makes use of these inclusion maps to compare topological features within different force networks. The first observation, while trivial, is essential

for our discussion and follows directly from the fact that $\text{FN}(f, \theta) = \emptyset$ for all $\theta > \theta_K$.

Lemma 5.1. *If $\theta > \theta_K$, then $H_*(\text{FN}(f, \theta)) = 0$.*

Now consider a value θ_k such that $v \in H_n(\text{FN}(f, \theta_k))$ and let $v \neq 0$. If $n = 0$ or 1 , then v provides information about the existence of components or loops, respectively, in $\text{FN}(f, \theta_k)$. In light of Lemma 5.1, there exists a unique largest threshold $\theta_b(v) \geq \theta_k$ with the property that there exists $v_b \in H_n(\text{FN}(f, \theta_b))$ such that $\iota_{\theta_k, \theta_b}(v_b) = v$. The geometric feature associated with v is said to have been *born* at level $\theta_b(v)$.

It is also possible that for some $\theta < \theta_k$, $\iota_{\theta, \theta_k}(v) = 0$. In this case we define

$$\theta_d(v) := \max \{ \theta_j \mid \iota_{\theta_j, \theta_k}(v) = 0 \}$$

and we say that the geometric feature associated with v *dies* at level $\theta_d(v)$. Given our construction, not every geometric feature needs to die. In particular, for $n = 0, 1$,

$$H_n(\text{FN}(f, 0)) \cong H_n(\text{CN})$$

which, as the examples in this paper indicate, need not be trivial.

As is made clear shortly being able to identify birth and death levels to all geometric features is extremely useful. Thus we make use of the following convention

$$\text{if } \iota_{0, \theta_k}(v) \neq 0, \text{ then } \theta_d(v) = -1.$$

A remarkable fact [12, 13] is that given a finite filtration it is possible to choose a consistent set of bases for $H_n(\text{FN}(f, \theta_k))$, $k = -1, \dots, K$ such that each basis element has a well defined birth and death level (θ_b, θ_d) .

It is also worth noting that equation (2) implies that given any homology element $v \in H_n(\text{FN}_I(f, \theta_k))$, except one, $\theta_b(v) \leq \theta_K$ and $\theta_d(v) \geq 0$. Furthermore, if \bar{v} is a representative of the unique exception, then $\bar{v} \in H_0(\text{FN}_I(f, \theta_k))$ and $(\theta_b(\bar{v}), \theta_d(\bar{v})) = (\theta_K, -1)$

The collection of all pairs (θ_b, θ_d) associated with n -th homology group for the force networks are used to construct the n -th persistence diagram for the scalar field $f: \text{CN} \rightarrow [0, \infty)$. For the simple force network filtration presented in Figure 6 the associated persistence diagrams, shown in Figure 7, have a straightforward geometric interpretation. For the β_0 diagram a birth level $\theta_b(v)$ corresponds to the value of a local maximum that is associated with the birth of a connected component measured by the 0-homology class v . As θ decreases this component grows until it meets, at a point associated with a local minimum or saddle, another component. Assume this other component is measured by the homology class v' and that the value of the local minimum (or saddle) is $\underline{\theta}$. If $\theta_b(v) < \theta_b(v')$, then $\theta_d(v) = \underline{\theta}$. In this case, $\theta_b(v) - \theta_d(v)$ measures the difference in height between the local maximum and local minimum and hence this difference can be used as a measure of how robust a feature is.

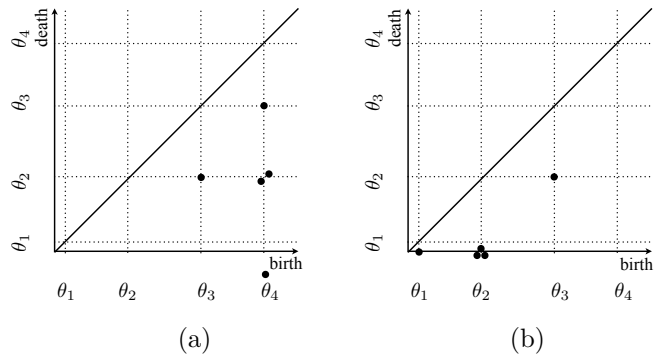


Figure 7: (a) β_0 and (b) β_1 persistence diagrams for the force network shown in Figure 6(c)

Let us be more precise and analyze the persistence diagrams for the interaction force network $\text{FN}_I(f, \theta)$ given by Figure 6. We start with the β_0 persistence diagram, see Figure 7(a). There are no points with the birth coordinate larger than θ_4 . This indicates the absence of components experiencing force larger than θ_4 . Four points with the birth coordinate θ_4 correspond to the four connected components that appear in $\text{FN}_I(f, \theta_4)$ (Figure 6(e)). The death coordinates of the points differ. This indicates that the components merge for different values of θ . The first merging appears for θ_3 (Figure 6(f)) and is represented by the dot (θ_4, θ_3) . Moreover a new component appeared at the level θ_3 and consequently merged with a preexisting component at θ_2 (Figure 6(g)) as indicated by the dot (θ_3, θ_2) . Also another two components that appear at θ_4 disappear at θ_3 hence there are two copies of the point (θ_4, θ_2) . Finally there is only one connected component for all $\theta < \theta_2$. This component appears for $\theta = \theta_4$ but does not disappear. In order to record it in the persistence diagram it is represented by $(\theta_4, -1)$.

Now we turn our attention to the β_1 persistence diagram. The first loop appears at θ_3 (Figure 6(f)) and is filled at θ_2 (Figure 6(g)) as shown by the point (θ_3, θ_2) in Figure 7(b). Another three loops appear at θ_2 and persist for all positive thresholds. Due to the definition of $\text{FN}_I(f, \theta)$ all the loops are filled in for $\theta = 0$. So the loops are represented by three copies of the point $(\theta_2, 0)$. The last loop appears at θ_1 (Figure 6(h)) and also persists for all positive thresholds hence the point $(\theta_1, 0)$ belongs to the β_1 persistence diagram.

Definition 5.2. Let $\Theta = \{\theta_k \mid k = -1, \dots, K\}$ and

$$\{\text{FN}(f, \theta_k) \mid \theta_k \in \Theta, \theta_k < \theta_j \Rightarrow \text{FN}(f, \theta_j) \subset \text{FN}(f, \theta_k)\}$$

be a force network filtration over a complex CN . The associated n -th persistence diagram $\text{PD}_n(f, \text{CN}, \Theta)$ is the multi set consisting of the following points:

1. one point for each n -th persistence point (θ_k, θ_j) ;
2. infinitely many copies of points (θ_k, θ_k) on the diagonal.

Condition (1) of Definition 5.2 arises because distinct geometric features can appear and disappear at the same thresholds and thus there may be multiple copies of the same persistence pair. The necessity of condition (2) is made clear shortly.

We conclude this section with an observation. Let \mathcal{PD}_n denote the set of all n -th persistence diagrams and \mathcal{PD} the set of all persistence diagrams. Given a chain complex CN , let $M(\text{CN}, [0, \infty))$ denote the set of monotone maps on CN . We can view persistence diagrams as a function

$$\text{PD}: M(\text{CN}, \mathbb{R}) \rightarrow \mathcal{PD} \quad (4)$$

or equivalently a collection of functions $\text{PD}_n: M(\text{CN}, \mathbb{R}) \rightarrow \mathcal{PD}_n$ defined by

$$\text{PD}_n(f) = \text{PD}_n(f, \text{CN}, \Theta)$$

where $\Theta = \{\theta_k \mid k = -1, \dots, K\}$ consists of the finite set of values obtained by f along with the convention that $\theta_{-1} = -1$ and $\theta_0 = 0$.

6. The space of Persistence Diagrams

The results concerning topological fidelity of the complexes CN_\bullet have, up to this point, been mostly negative. The introduction of persistence allows us to present positive results. In and of itself this suggests that force network filtrations and their associated persistent homology provide more appropriate metrics for understanding force networks than measurements performed at single thresholds. To obtain stability results we need to be able to compare persistence diagrams.

Definition 6.1. Let PD_n and PD'_n be two n -th persistence diagrams. The *bottleneck distance* between PD_n and PD'_n is defined to be

$$d_B(\text{PD}_n, \text{PD}'_n) = \inf_{\gamma: \text{PD}_n \rightarrow \text{PD}'_n} \sup_{p \in \text{PD}_n} \|p - \gamma(p)\|_\infty,$$

where $\|(a_0, b_0) - (a_1, b_1)\|_\infty := \max\{|a_0 - a_1|, |b_0 - b_1|\}$ and γ ranges over all bijections. Similarly, the *degree- q Wasserstein distance* is defined to be

$$d_{W^q}(\text{PD}_n, \text{PD}'_n) = \left[\inf_{\gamma: \text{PD}_n \rightarrow \text{PD}'_n} \sum_{p \in \text{PD}_n} \|p - \gamma(p)\|_\infty^q \right]^{1/q}.$$

As is indicated in [12], equipped with either the bottleneck or degree- q Wasserstein distance \mathcal{PD}_n is a metric space. From now on we always assume \mathcal{PD} is one of these metric spaces.

As is indicated in Figure 8 the ability to match points in persistence diagrams with points on the diagonal suggests that small perturbations lead to small distances between persistence diagrams. In fact, it is proven in [12] that given a complex CN and two monotone functions $f, g: \text{CN} \rightarrow \mathbb{R}$ the bottleneck distance satisfies

$$d_B(D(f), D(g)) \leq \sup_{x \in X} |f(x) - g(x)|. \quad (5)$$

A similar result holds for the degree- q Wasserstein distance [12, Section VIII.3]. A more formal statement is as follows.




Figure 8: (a) Two functions. Blue represents a noisy perturbation of red. (b) Associated persistence diagrams along with pairing of persistence points satisfying definition of bottleneck distance.

Theorem 6.2. Given a complex CN let $M(\text{CN}, \mathbb{R})$ denote the set of monotone functions on CN equipped with the sup norm $\|\cdot\|_\infty$. Then

$$\text{PD}: M(\text{CN}, [0, \infty)) \rightarrow \mathcal{PD}$$

defined by (4) is a Lipschitz continuous map.

Corollary 6.3. The map $\text{PD}: M(\text{CN}_I, [0, \infty)) \rightarrow \mathcal{PD}$ is Lipschitz continuous.

Corollary 6.3 implies that a small change in the forces, either through perturbation of the system or experimental error results in a small change in the associated persistence diagrams. This is the long promised stability result. The failure of CN_P^Δ and CN_D to be stable with respect to perturbations follows from the fact that small changes of particle positions can result in changes of the underlying complex and thus Theorem 6.2 is not applicable. Figure 9 demonstrates that it is possible, using CN_P^Δ , for an arbitrarily small change in the position of the particles to lead to an order one change in the bottleneck distance.

7. Numerical Experiments

The discussions of the previous section provide a mathematical framework for studying force networks associated with DGM. In this section we re-examine these concepts in the context of simulated data. We begin with a brief review of the numerical simulations and computational tools employed. We then analyze persistent homology on a variety of levels. First, we consider the stability of the persistence diagrams obtained from the digital, position, or interaction networks with respect to numerical and/or experimental error. Second, we discuss the question of relating the information in individual persistence diagrams to physical properties of the DGM and how this information is dependent on the particular choice of network. Finally, we demonstrate that DGM made up of particles characterized by different physical properties lead to different persistence diagrams.




Figure 9: An arbitrarily small change in positions can lead to an order one change in the bottleneck distance. 6 particles with the magnitude of the force field indicated. Thick edges have force value 2, thin edges have force value 1. (a) Because of configuration of particles we see two loops. One loop appears at $\theta = 2$, the second loop appears at $\theta = 1$. (b) A perturbation of the configuration in (a) but the forces on the particles do not change. Only one loop appears at $\theta = 1$. (c) β_1 persistence diagrams. The stars at $(1, -1)$ and $(2, -1)$ correspond to the persistence points for (a) and the dot at $(1, -1)$ is the single persistence point for b.

7.1. Simulations Used

We perform a series of molecular dynamics (MD) simulations similar to our previous works [17, 14]. For the present paper, we consider a set of about 2,000 circular particles contained in a square domain with rough walls composed of monodisperse particles. The system is slowly compressed allowing for a change of packing fraction, ρ , between 0.6 and 0.9. Initially the particles are given random velocities and are placed on a square lattice. The equations of motion are integrated using a fourth order predictor-corrector scheme. We implement the Cundall-Strack model for static friction which includes normal and tangential forces at the contact [18]. For frictionless system, the contact force reduces to a normal force with a spring and viscous damping term. In general, we use polydisperse particles where the particle sizes are chosen from a uniform distribution with width $r_p = (r_{max} - r_{min})/r_{ave}$, where r_{ave} is the mean particle radius. The coefficient of restitution measuring energy loss is given the value of $e_n = 0.5$, and the coefficient of static friction is either $\mu = 0.5$ for the frictional case or $\mu = 0.0$ for the frictionless one. See [17, 14] for more details.

For the most part we focus on a system of particles with $\rho \approx 0.86$. This is beyond the jamming transition for all the simulations considered here and the particles are packed close enough so that most of the particles belong to the same connected component of the position network CN_P . We use this data to address our first two goals: understanding the robustness of the persistent homology measurements as a function of the choice of complex; and understanding how the choice of complex effects the persistent homology. Our final goal is to demonstrate that persistent homology can capture and provide insight into the geometry of the associated force network. To do this we consider two particular cases where the force field

is known to be different [14]: a system of monodisperse frictionless particles ($r_p = 0.0$, $\mu = 0.0$) and a system of polydisperse frictional particles ($r_p = 0.4$, $\mu = 0.5$).

For the ρ 's of interest, we extract the magnitude $\psi_{i,j}$ of the normal force interaction between any two particles p_i and p_j . The values $\psi_{i,j}$ completely determine the interaction network FN_I . To construct FN_P and FN_D the positions of the particles need to be extracted as well. The value ψ_i assigned to the particle p_i is the total force experienced by this particle, i.e.

$$\psi_i := \sum_{\{j | \langle i,j \rangle \in \text{CN}_P^{(1)}\}} \psi_{i,j} = \sum_j \psi_{i,j}.$$

To simplify the comparisons between packing fractions, we normalize the function $f: \text{CN} \rightarrow \mathbb{R}$, defined in Section 4, by dividing it by the average force \hat{f} defined as follows: for the interaction force network

$$\hat{f}_I = \frac{1}{2M} \sum_{i,j=1}^N \psi_{i,j}, \quad (6)$$

and for the position and digital force networks

$$\hat{f}_P = \frac{1}{N} \sum_{i=1}^N \psi_i = \frac{1}{N} \sum_{i,j=1}^N \psi_{i,j}, \quad (7)$$

where M is the number of non zero force interactions $\psi_{i,j}$ and N is the number of particles. Note that the average number of contacts $Z = \frac{2M}{N}$ and

$$\hat{f}_P = Z \hat{f}_I. \quad (8)$$

We have produced open source software [19] that is used to encode this procedure and produce a force networks filtration $\{\text{FN}(f, \theta_k) \mid \theta_k \in \Theta\}$ (Θ is a function of the complex). The persistent homology of each filtration is computed using the open source software Perseus [20, 21]. We note that since the number of pixels representing a single particle is much larger than one, the size of the digital complex CN_D is considerably larger than the size of the position or interaction complexes, implying that the computational cost of analyzing CN_D is much larger than for CN_P or CN_I . To give a sense of the time needed to perform these types of computations we remark that using a 2.53 GHz processor to compute the persistence diagrams for the position force network, digital force network, and interaction force network required 25, 97, and 43 seconds, respectively.

7.2. Stability of Persistence Diagrams

Figure 10 shows the digital, position and interaction networks of a system of monodisperse frictionless particles at $\rho = 0.8696$. The associated persistence diagrams are depicted in Figure 11. In this section we use the following convention. If the feature persists until the zero

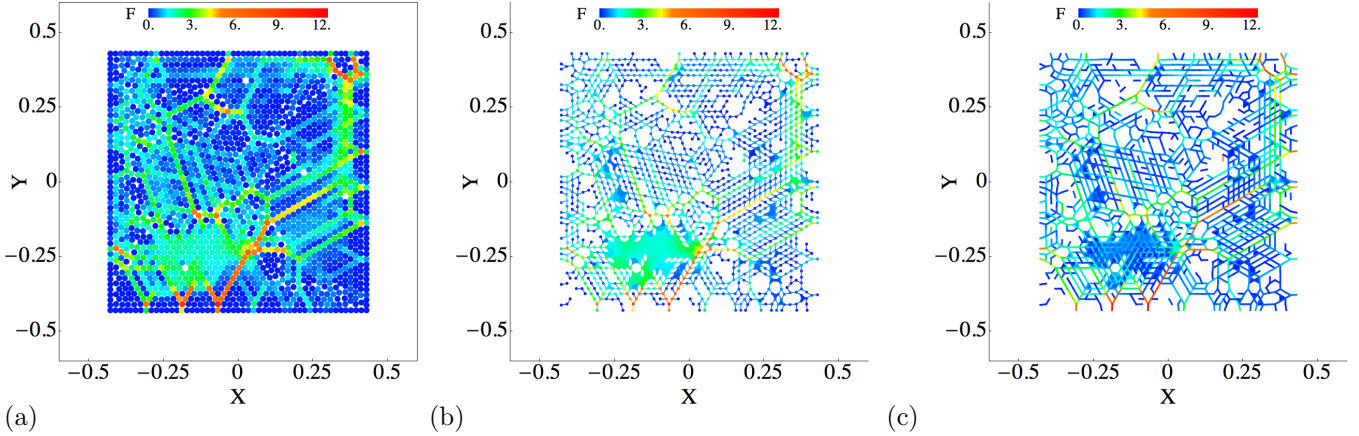


Figure 10: Different force networks for mono disperse frictionless system at the packing fraction $\rho = 0.86$. (a) digital force network (b) position force network (c) non zero simplices of the interaction force network.

threshold then we set the death coordinate to minus one. This only impacts the persistence diagram for the interaction network, allowing for simple visual identification of the defects. Since the information used to construct these various networks is different it is not surprising that the output differs. We investigate these differences below. For the moment we concentrate on the more fundamental question: how stable is the information contained in the persistence diagrams with respect to an error in input data?

The numerical simulations and the extraction of particle positions and normal interparticle forces are done using double precision floating point numerics. If we accept this as “truth” then the simplest form of error has to do with precision of measurement. With this in mind we introduce measurement error by truncating the position data of the particles to two or three significant digits (this impacts the position and digital networks) and by truncating the magnitude of the normal force acting between the particles to two or three significant digits (this impacts the interaction network as well). We then measure the distances between the original and perturbed persistence diagrams using the bottleneck and degree-1 Wasserstein metric. The results are indicated in Table 1. Some values in the table are in line with expectation. The relatively small values associated with the interaction network are predicted by Corollary 6.3. The fact that the values for d_{W^1} are significantly larger than d_B for each type of network is not surprising since the bottleneck distance measures the single largest change in the network while the degree-1 Wasserstein is sensitive to the many small local perturbations that may be occurring.

Other values in Table 1 are less easily explained. We have no theoretical results that explain the relative differences in perturbations of distances between the persistence diagrams associated with the position and digital networks. The digital network was constructed using resolutions of 1000×1000 and 2000×2000 pixels. Given the size of the particles in both cases each pixel represents a measurement to approximately two significant figures.

We hypothesize that this explains the relatively small (as compared with the position network) change in the persistence diagrams obtained by truncating to 3 significant digits. The sensitivity of the topological structure of cubical approximations of circular disks demonstrated in Figure 3 suggests that the larger distance value for the 2000×2000 digital complex with truncation at two significant digits should not come as a surprise.

The different distance values for different digital complexes raises another issue; how sensitive is the persistence diagram to the resolution of the digital network? We consider this issue using the system shown in Figure 10. Computing with the original numerical data at a resolution of 2000×2000 pixels the β_0 persistence diagrams for the digital networks are identical. Comparison of the β_1 persistence diagrams reveals that the number of loops is around 25% larger for the higher resolution. We have verified that this increase is caused by formation of extra holes at the places where the particles are close to each other; essentially the phenomenon indicated in Figure 3(a).

7.3. Force Networks as a Function of Complex Type

Figures 10 and 11 demonstrate that the digital, position and interaction force networks of a single system of particles can be quite different. The idea behind the construction of the digital and position networks is the same, the difference arises from the fact they are based on different complexes that provide different approximations of the geometry of the system of particles. Thus, to focus on the essential differences we restrict our discussion to a comparison of the position and interaction force networks. Figure 12 provides an enlarged view for three different subregions of the position and interaction force networks of Figures 10(b) and 10(c). The position force network is defined in terms of the vertices and thus the corresponding figures include the magnitude of the force on the vertices. The vertices are not highlighted in the interaction force network since the value of the force on the edges is used to define the values on the vertices.

metric	truncation	Network Type			
		interaction	position	digital (1000)	digital (2000)
d_B	3	0.00035	2.254	0.558	0.4724
	2	0.0064	3.37	1.6742	4.302
d_{W^1}	3	0.212	904.5	207.17	100.5
	2	2.19	1398.3	851.9	3197.3

Table 1: The distance between the persistence diagram of original network and the persistence diagram after truncation of numerical data to 2 or 3 significant digits.

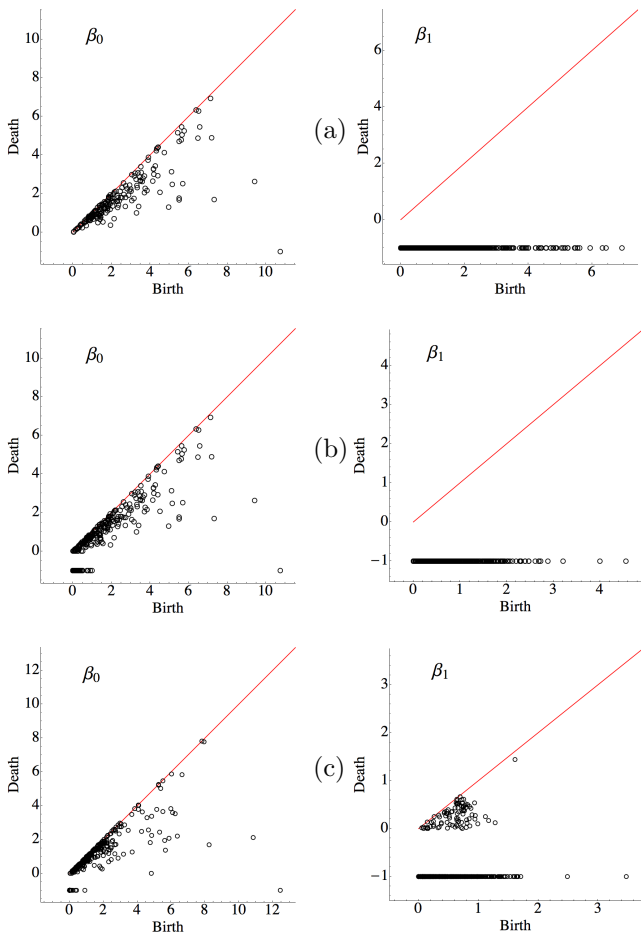


Figure 11: Persistence diagrams for mono disperse frictionless system at the packing fraction $\rho = 0.86$ shown in Figure 10. Persistence diagrams for (a) digital force network based on 1000×1000 pixels, (b) position force network, and (c) non zero simplices of the interaction network.

Figure 12(a) is typical of a region in which we see crystalline structure or equivalently a region over which there are no defects. Observe that in this crystalline region the normal forces for the position network are significantly larger than those of the interaction network. The explanation is that the force acting between the particles is small in these regions, but each particle has six contacts so the sum of the forces on each particle (which is what we are recording) is high. Note that the forces are rather uniform in the crystalline zone and $\psi_i \approx 6 \max_j \{\psi_{i,j}\}$. Let f_I and f_P denote the forces in the interaction and particle networks, respectively. Then

$$f_I(i) = \frac{\max_j \{\psi_{i,j}\}}{\hat{f}_I} = \frac{Z \max_j \{\psi_{i,j}\}}{\hat{f}_P} \approx \frac{Z \psi_i}{6 \hat{f}_P} = \frac{Z}{6} f_P(i).$$

Except for the perfect crystal the value of Z is less than 6. In our case $Z = 3$ and $f_I(i) \approx \frac{1}{2} f_P(i)$

We now consider a part of the domain where we find sets of particles interacting by large forces, resembling a 'force chain.' In this case, as can be seen along the red line/orange string in Figure 12(b), the position network tends to report a lower magnitude of force than the interaction network. To understand this recall that the packing fraction is $\rho = 0.86$, which is above the jamming transition. Thus we can assume that $Z \geq 3$. Observe that along the red/orange chain of particles in Figure 12(b) each particle typically has contact with 2 or 3 other particles. Therefore $\psi_i \leq 3 \max_j \{\psi_{i,j}\}$. By (8) and the inequalities stated in this paragraph we obtain

$$f_I(i) = \frac{\max_j \{\psi_{i,j}\}}{\hat{f}_I} = \frac{Z \max_j \{\psi_{i,j}\}}{\hat{f}_P} \geq \frac{\psi_i}{\hat{f}_P} = f_P(i).$$

An added effect is that a single continuous chain of strong force interactions in the interaction force network is reported to be a collection of shorter chains in the position force network (see Figure 12(c)). An immediate consequence is that we expect to see more points with relatively large birth values in the β_0 persistence diagram of the position force network than in the β_0 persistence diagram of the interaction network. This is confirmed by counting the number of points the β_0 persistence diagrams of Figures 11(b) and (c) (shown in Figure 11), with birth value greater than a given value, for example 3.

Figure 12(c) demonstrates another important difference between the position and interaction networks. In the interaction network a chain of particles experiencing

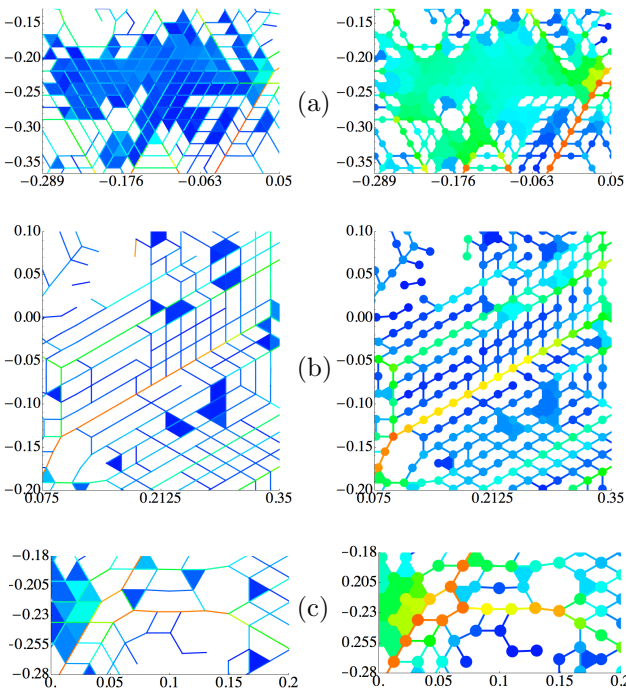


Figure 12: Enlarged views, using the same color scheme, of three different subregions of the interaction (left) and position (right) networks of Figures 10(b) and 10(c).

a strong force that splits will result in two chains consisting of weaker forces. Weaker forces need not result from splitting in the case of position networks, since the force on the particles is determined by the force between the neighboring particles. This relationship between splitting and forces implies that loops are formed at lower force levels in the interaction networks as compared to the position networks which, in turn, implies that there should be fewer points with relatively large birth values in the β_1 persistence diagram of the interaction network as compared with the β_1 persistence diagram of the position network. This is corroborated by Figures 11(a) and (b).

Another striking difference between the position and interaction networks is that in the β_1 persistence diagram of the position network the death value for all points is -1 , i.e. once a loop is formed it never dies. This is not the case for the β_1 persistence diagram of the interaction network. There is no reason a priori why loops in the position network cannot die, but a possible explanation is as follows. The death of loops is associated with the appearance of 2-simplices which is indicative of crystalline structure. Consider a single 2-simplex $\langle v_i, v_j, v_k \rangle$ and assume that

$$f(\langle v_i \rangle) > f(\langle v_j \rangle) > f(\langle v_k \rangle).$$

Given the definition of the position force network, the vertex $\langle v_i \rangle$ appears first, followed by the vertex $\langle v_j \rangle$ and the edge $\langle v_i, v_j \rangle$. Finally, the vertex $\langle v_k \rangle$, the edges $\langle v_i, v_k \rangle$ and

$\langle v_j, v_k \rangle$ and the 2-simplex $\langle v_i, v_j, v_k \rangle$ are all included at the same step. Thus there is no opportunity for a loop consisting of three edges to be generated. A similar argument can be made for the interaction network and hence loops that appear in the persistence diagrams must involve multiple edges. If we think of this sequence of edges as a force chain, then the previous argument suggests that for the position force network this chain is more likely to contain edges of lower magnitude than in the interaction network. At the same time, we have observed that in crystalline regions the force magnitudes at the particles are larger in the position network than the interaction network. These two observations suggest that in the position network it is difficult to construct a loop in a crystalline region that surrounds vertices with lower forces.

A final observation from the diagrams is that the β_0 persistence diagrams for the digital network have less points than those for the position or interaction network. This is due to the fact that, as is discussed in Section 2, our construction of the digital network artificially inflates area associated with each particle and hence it is possible for distinct components in the position or interaction network to form a single component in the digital network. This effect is particularly relevant in the context of rattlers, i.e. single particles that do not experience any force but are contained within regions of particles experiencing nontrivial normal forces. In contrast, the number of persistence generators in β_1 persistence diagrams is larger for digital networks due to the formation of the artificial holes, again as described in Section 2.

7.4. Comparison of different systems via persistence diagrams

The most direct means of applying persistence diagrams is to use them to distinguish and/or interpret the global force structures of systems of DGM composed of particles characterized by different physical properties. To do this, we use the interaction network and consider two systems, a monodisperse frictionless ($r_p = 0.0$, $\mu = 0.0$) and a polydisperse frictional ($r_p = 0.4$, $\mu = 0.5$) system.

We begin by assigning physical meaning to the location of persistence points in the persistence diagrams. Figure 13 shows a persistence diagram divided into five regions. With the exception of the region labelled defects, the location of the division lines is intended to be either system specific or conceptual. We explain these divisions in the context of the interaction force networks and persistence diagrams of the $r_p = 0.0$, $\mu = 0.0$ and $r_p = 0.4$, $\mu = 0.5$ systems at packing fraction $\rho = 0.86$ as indicated in Figures 10(c) and 14, respectively. A more complete analysis of DGM using these ideas is presented in [14].

There are at least two different interpretations of the points in the β_0 persistence diagram that lie close to the diagonal; the region we have labelled as roughness. The first is to treat this as noise, i.e. a byproduct of the imperfect measurements of the normal forces between parti-

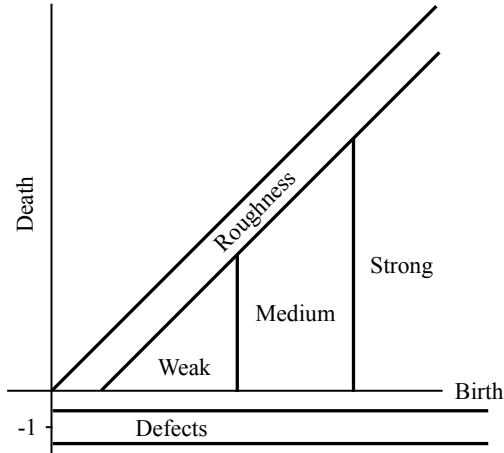


Figure 13: Persistence diagram divided into regions. Explanation of the regions is provided in the text.

cles. While this may be appropriate for many experimental settings, the data represented in Figures 10, 11 and 14 comes from numerical simulations. Thus, the errors are extremely small compared to the size of the normal forces. This leads to the second interpretation, which we adopt, that this region of the persistence diagram provides a measurement of how rough or bumpy the normal force landscape is, e.g. should we view the surface of the landscape as being made of glass or sandpaper? Alternatively, the points in the β_0 persistence diagram that lie outside the roughness region provide a means of measuring how non-uniform the normal force landscape is. Therefore by comparing 11(c) and 14 we conclude that $r_p = 0.4$, $\mu = 0.5$ system is rougher than the $r_p = 0.0$, $\mu = 0.0$.

To understand the region labelled as strong, observe that the image in Figure 14 of the forces for the $r_p = 0.4$, $\mu = 0.5$ system does not contain any red simplices, implying that there are no extremely strong force interactions. In contrast, such red simplices are present in the $r_p = 0.0$, $\mu = 0.0$ system displayed in Figure 10(c). This difference can be inferred from the β_0 persistence diagrams shown in Figures 11 and 14. As is indicated in Section 7.1 the function $f : \text{CN} \rightarrow \mathbb{R}$ is normalized. For FN_I the value 1 represents the average interaction force while for FN_P and FN_D it is the average total force experienced by particles. For the $r_p = 0.4$, $\mu = 0.5$ system there are no persistence points with the birth value larger than 3 and only a few points with birth value larger than 2.5. Thus, depending on the exact cut-off there are no or at most few points in the region marked strong for the $r_p = 0.4$, $\mu = 0.5$ system, in clear contrast to the $r_p = 0.0$, $\mu = 0.0$ system.

If we take the left division marker for the medium regime in Fig. 13 to be 1, then the persistence points in the medium and strong regions provide information about the geometry of what the DGM community typically refer to as a force chain. In the case of the $r_p = 0.4$, $\mu = 0.5$ system, we see a large number of β_0 persistence points that are born between 1 and 2.5 and die before 0.8. This suggests a landscape consisting of moderately high peaks separated by moderately high valleys. To continue the geographic metaphor, the $r_p = 0.4$, $\mu = 0.5$ force chain network takes place on a high plateau. In contrast, the $r_p = 0.0$, $\mu = 0.0$ system has fewer moderately high peaks, but they are separated by much deeper valleys since there are points with death values below 0.6. Therefore, we conjecture that the landscape for the $r_p = 0.0$, $\mu = 0.0$ system has fewer peaks (but some of them are strong) than that of the $r_p = 0.4$, $\mu = 0.5$ landscape, and these peaks are in general much more isolated and more likely to be separated by valleys of much weaker forces.

Finally, we consider the region labelled defects. In a β_0 persistence diagram each point in this region corresponds to a distinct connected component. In the context of the $r_p = 0.4$, $\mu = 0.5$ system these mostly correspond to rattlers. This conclusion is obtained by observing that aside from the single persistence point with a large birth force, that corresponds to the component containing most




Figure 14: Non-zero simplices of the interaction force network and the β_0 and β_1 persistence diagrams for the $r_p = 0.4$, $\mu = 0.5$ system at $\rho = 0.86$.

of the particles, the persistence points in the defects region have a birth value of 0, indicating that they are not experiencing any normal force. This is quite different from the $r_p = 0.0$, $\mu = 0.0$ system. In this case we have persistence points in the defects region with non-zero birth forces. This implies the existence of small clusters of particles (a single separated particle cannot have an interaction force) that are not interacting with the dominant particle cluster. Close inspection of the interaction force network in Figure 10(c) reveals these small components.

The defects region of the β_1 persistence diagrams provides additional information. As indicated in Section 7.3, β_1 persistence points lie in the defects region if and only if they correspond to loops that enclose non-crystalline regions. There are about twice as many persistence points in the defects region in the $r_p = 0.4$, $\mu = 0.5$ system as compared to the $r_p = 0.0$, $\mu = 0.0$ one. This suggests that the $r_p = 0.4$, $\mu = 0.5$ system is more likely to support defects than the $r_p = 0.0$, $\mu = 0.0$ one. At the same time there are 50% more points in β_1 persistence diagram that are not in the defects region for the polydisperse frictional system. These persistence points correspond to loops that are filled in by 2-dimensional simplices and thus must be contained within regions of crystallization. Thus this difference in the number of persistence points suggests that the polydisperse system contains a multitude of small crystalline regions as opposed to the $r_p = 0.0$, $\mu = 0.0$ system. This is corroborated by a careful examination of the force networks in Figures 10(c) and 14.

Up to this point we have focussed on the relationship between individual persistence diagrams and the physical characteristics of the DGM. While understanding this correspondence is essential to the claim that persistent homology provides a potentially important tool in the analysis of DGM, we believe that the most significant value of this technique will come through the analysis of large sets of persistence diagrams. We present a simple example of this idea here, leaving more detailed investigations for future work.

We extract a number (one hundred) of consecutive frames starting at $\rho = 0.87$ and $\rho = 0.79$ for the two considered systems, respectively (both systems are above jamming transition, see [14]). The increment of ρ between the consecutive frames is roughly 4×10^{-5} . We then compute the distance between every pair of persistence diagrams. These are presented in Fig. 15 in the form of a heat map. The first hundred points correspond to the sequence of samples from the $r_p = 0.0$, $\mu = 0.0$ system, and the second hundred points are from the $r_p = 0.4$, $\mu = 0.5$ system. The distinct blocks of blue and red are what we expect; the distances between persistence diagrams from the $r_p = 0.0$, $\mu = 0.0$ system are small (blue), the distances between persistence diagrams for the $r_p = 0.4$, $\mu = 0.5$ system are small (blue), but the distances between persistence diagrams for the $r_p = 0.0$, $\mu = 0.0$ and the $r_p = 0.4$, $\mu = 0.5$ system are large (red).

Another interesting observation comes from carefully




Figure 15: Distances between 100 consecutive frames of $r_p = 0.0$, $\mu = 0.0$ and $r_p = 0.4$, $\mu = 0.5$ systems. First hundred frames correspond to $r_p = 0.0$, $\mu = 0.0$ and the second hundred frames are from $r_p = 0.4$, $\mu = 0.5$. a) Bottleneck distance b) Wasserstein d_{W1} distance.

examining and comparing the blue regions. The lighter shades of blue indicate that the distances between the persistence diagrams corresponding to the $r_p = 0.0$, $\mu = 0.0$ system are larger than the distance for the polydisperse system. This indicates that the $r_p = 0.0$, $\mu = 0.0$ system evolves more and faster after the jamming transition than the polydisperse system.

8. Conclusion

Based on different methods for collecting data we have defined three different chain complexes and used them to construct force chain networks for particulate systems. Using the force chain networks we compute persistence diagrams and we discuss how one can use persistent homology to extract information about the geometric structure of the force distributions between the particles. We provide both theoretical and experimental arguments to show that the persistence diagrams obtained from interaction force networks are the most robust with respect to experimental or numerical errors. Using numerical data obtained from molecular dynamics simulations of a system of particles being slowly compressed we show that the persistence diagrams associated to the different force networks can have significant differences. This in turn implies that the geometry of the force distributions observed depends upon the methods by which the system is sampled. We provide some intuition concerning how in general the sampling method affects the geometry. We also demonstrate that using persistent homology of any of the three force networks allows one to draw meaningful distinctions between the behavior of the force distributions for systems made up of particles with different geometric and physical properties.

Acknowledgments

This work was partially supported by NSF-DMS-0835621, 0915019, 1125174, AFOSR and DARPA (A.G., M.K., and K.M) and NSF Grant No. DMS-0835611, and DTRA Grant No. 1-10-1-0021 (A.G. and L.K.).

References

- [1] C.-h. Liu, S. R. Nagel, D. A. Schecter, S. N. Coppersmith, S. Majumdar, O. Narayan, and T. A. Witten. Force fluctuations in bead packs. *Science*, 269:513, 1995.
- [2] M. E. Cates, J. P. Wittmer, J.-P. Bouchaud, and P. Claudin. Jamming, force chains and fragile matter. *Phys. Rev. Lett.*, 81:1841, 1998.
- [3] T. S. Majmudar and R. P. Behringer. Contact force measurements and stress-induced anisotropy in granular materials. *Nature*, 435:1079, 2005.
- [4] S. Alexander. Amorphous solids: Their structure, lattice dynamics and elasticity. *Phys. Rep.*, 296:65, 1998.
- [5] F. Radjai, M. Jean, J. J. Moreau, and S. Roux. Force distribution in dense two-dimensional granular systems. *Phys. Rev. Lett.*, 77:274, 1996.
- [6] F. Radjai, D. E. Wolf, M. Jean, and J.-J. Moreau. Bimodal character of stress transmission in granular packings. *Phys. Rev. Lett.*, 80:61, 1998.
- [7] A. Tordesillas, D. M. Walker, and Q. Lin. Force cycles and force chains. *Phys. Rev. E*, 81:011302, 2010.
- [8] D.M. Walker and A. Tordesillas. Taxonomy of granular rheology from grain property networks. *Phys. Rev. E*, 85:011304, 2012.
- [9] M. Herrera, S. McCarthy, S. Sotterbeck, E. Cephas, W. Losert, and M. Girvan. Path to fracture in granular flows: Dynamics of contact networks. *Phys. Rev. E*, 83:061303, 2011.
- [10] D.S. Bassett, E.T. Owens, K.E. Daniels, and M.A. Porter. Influence of network topology on sound propagation in granular materials. *Phys. Rev. E*, 86:041306, 2012.
- [11] R. Arévalo, L.A. Pugnaloni, I. Zuriguel, and D. Maza. Contact network topology in tapped granular media. *Phys. Rev. E*, 87:022203, 2013.
- [12] H. Edelsbrunner and J. L. Harer. *Computational topology*. AMS, Providence, RI, 2010.
- [13] G. Carlsson. Topology and data. *Bull. Amer. Math. Soc. (N.S.)*, 46:255, 2009.
- [14] M. Kramar, A. Goulet, L. Kondic, and K. Mischaikow. Persistence of force networks in compressed granular media. *Phys. Rev. E*, 87:042207, 2013.
- [15] R. R. Hartley and R. P. Behringer. Logarithmic rate dependence of force networks in sheared granular materials. *Nature*, 421:928, 2003.
- [16] T. Kaczynski, K. Mischaikow, and M. Mrozek. *Computational homology*, volume 157 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 2004.
- [17] L. Kondic, A. Goulet, C.S. O'Hern, M. Kramar, K. Mischaikow, and R.P. Behringer. Topology of force networks in compressed granular media. *Europhys. Lett.*, 97:54001, 2012.
- [18] P. A. Cundall and O. D. L. Strack. A discrete numerical model for granular assemblies. *Géotechnique*, 29:47, 1979.
- [19] M. Kramar. Chomp, 2013. <http://chomp.rutgers.edu/Topological%20Characterization%20of%20Dense%20Granular%20Media.html>.
- [20] K. Mischaikow and V. Nanda. Morse theory for filtrations and efficient computation of persistent homology. *Discrete & Computational Geometry*, to appear.
- [21] V. Nanda. Perseus, 2012. <http://www.math.rutgers.edu/~vidit/perseus.html>.