

CONFIGURATION SPACES OF DISKS IN AN INFINITE STRIP

HANNAH ALPERT, MATTHEW KAHLE, AND ROBERT MACPHERSON

ABSTRACT. We study the topology of the configuration spaces $\mathcal{C}(n, w)$ of n hard disks of unit diameter in an infinite strip of width w . We describe ranges of parameter or “regimes”, where homology $H_j[\mathcal{C}(n, w)]$ behaves in qualitatively different ways.

We show that if $w \geq j + 2$, then the homology $H_j[\mathcal{C}(n, w)]$ is isomorphic to the homology of the configuration space of points in the plane, $H_j[\mathcal{C}(n, \mathbb{R}^2)]$. The Betti numbers of $\mathcal{C}(n, \mathbb{R}^2)$ were computed by Arnold [2], and so as a corollary of the isomorphism, $\beta_j[\mathcal{C}(n, w)]$ is a polynomial in n of degree $2j$.

On the other hand, we show that if $2 \leq w \leq j + 1$, then $\beta_j[\mathcal{C}(n, w)]$ grows exponentially with n . Most of our work is in carefully estimating $\beta_j[\mathcal{C}(n, w)]$ in this regime.

We also illustrate, for every n , the homological “phase portrait” in the (w, j) -plane—the parameter values where homology $H_j[\mathcal{C}(n, w)]$ is trivial, nontrivial, and isomorphic with $H_j[\mathcal{C}(n, \mathbb{R}^2)]$. Motivated by the notion of phase transitions for hard-spheres systems, we discuss these as the “homological solid, liquid, and gas” regimes.

1. INTRODUCTION

We study the topology of configuration spaces $\mathcal{C}(n, w)$ of n non-overlapping disks of unit diameter in an infinite strip of width w .

In other words, for non-negative integers n, w we define

$$\begin{aligned} \mathcal{C}(n, w) = \{ (x_1, y_1, x_2, y_2, \dots, x_n, y_n) \in \mathbb{R}^{2n} : \\ (x_i - x_j)^2 + (y_i - y_j)^2 \geq 1 \text{ for every } i \neq j, \text{ and} \\ 1/2 \leq y_i \leq w - 1/2 \text{ for every } i. \} \end{aligned}$$

The coordinates (x_i, y_i) are the centers of disk i , the inequalities $(x_i - x_j)^2 + (y_i - y_j)^2 \geq 1$ ensure that the disks have disjoint interiors, and the inequalities $1/2 \leq y_i \leq w - 1/2$ ensure that the disks of unit diameter (or radius $1/2$) stay in the closed strip $0 \leq y \leq w$. These spaces generalize the well-studied configuration space of points in the plane, which we denote $\mathcal{C}(n, \mathbb{R}^2)$.

Our main result describes the asymptotics for the Betti numbers $\beta_j[\mathcal{C}(n, w)]$, for fixed j and w , as $n \rightarrow \infty$. Our results do not depend on the choice of coefficient

2010 *Mathematics Subject Classification.* Primary: 55R80, Secondary: 82B26.

Key words and phrases. configuration spaces, phase transitions.

M.K. thanks IAS for hosting him as a member in 2010–11, and for several visits since then. He gratefully acknowledges the support of a Sloan Research Fellowship, NSF #DMS-1352386, and NSF #CCF-1839358. H.A. is supported by the National Science Foundation under Award No. DMS 1802914.

All three authors thank ICERM for hosting them during the special thematic semester “Topology in Motion” in Autumn 2016.




FIGURE 1. Theorem 1.1 describes the rate of growth of $\beta_j[\mathcal{C}(n, w)]$, for fixed j and w , as $n \rightarrow \infty$. The results are up to a constant factor, e.g. $\beta_8[\mathcal{C}(n, 3)] \asymp 5^n n^{12}$.

field for the homology. We use the notation $f \asymp g$ to indicate that there exist positive constants c_1, c_2 such that

$$c_1 g(n) \leq f(n) \leq c_2 g(n)$$

for all sufficiently large n . In the following, the implied constants depend on j and w but not on n .

Theorem 1.1 (Asymptotic rate of growth of the Betti numbers as $n \rightarrow \infty$).

- (1) If $w \geq 2$ and $0 \leq j \leq w - 2$ then the inclusion map $i : \mathcal{C}(n, w) \rightarrow \mathcal{C}(n, \mathbb{R}^2)$ induces an isomorphism on homology

$$i_* : H_j[\mathcal{C}(n, w)] \rightarrow H_j[\mathcal{C}(n, \mathbb{R}^2)].$$

So if $n \rightarrow \infty$ then the asymptotic rate of growth is given by

$$\beta_j[\mathcal{C}(n, w)] \asymp n^{2j}.$$

- (2) If $w \geq 2$ and $j \geq w - 1$ then write $j = q(w - 1) + r$ with $q \geq 1$ and $0 \leq r < w - 1$. Then we have that

$$\beta_j[\mathcal{C}(n, w)] \asymp (q + 1)^n n^{qw + 2r}.$$

If $w = 1$ and $j = 0$, then $\beta_0 = n!$.

- (3) If either $w = 0$, or $w = 1$ and $j \geq 1$, then $\beta_j = 0$.

Configuration spaces of disks arise naturally as the phase space of a 2-dimensional “hard-spheres” system, so are of interest in physics as well. See, for example, the

discussion of hard disks in a box by Diaconis in [14], and the review of the physics literature by Carlsson et al. in [11].

The topology of configuration spaces of particles with thickness has been studied earlier, for example in [1], [3], [13], and [21], but so far, not much seems to be known. Some of this past work is also inspired in part by applications to engineering, for example motion planning for robots.

Inspired by the statement of Theorem 1.1, we suggest the following definitions for “homological solid, liquid, and gas” regimes in the (w, j) plane.

- We define the “homological solid” phase to be wherever homology is trivial. The motivation for this definition is that one expects that in a crystal phase, things are fairly rigid and that the configuration space is simple.
- We define the “homological gas” phase to be where homology agrees with the configuration space of points in the plane. In other words, through the lens of this homology group, the particles are indistinguishable from points, corresponding to the assumption of atoms acting as point particles in an ideal gas. Arnold [2] showed that the Poincaré polynomial of $\mathcal{C}(n, \mathbb{R}^2)$ is given by

$$\beta_0 + \beta_1 t + \cdots + \beta_{n-1} t^{n-1} = (1+t)(1+2t) \cdots (1+(n-1)t).$$

It follows that the Betti numbers are given by the unsigned Stirling numbers of the first kind.

$$\beta_j [\mathcal{C}(n, \mathbb{R}^2)] = \left[\begin{matrix} n \\ n-j \end{matrix} \right].$$

For a self-contained overview of the homology and cohomology of $\mathcal{C}(n, \mathbb{R}^2)$, see Sinha [26].

One can use a standard recursive formula for Stirling numbers to write $\left[\begin{matrix} n \\ n-j \end{matrix} \right]$ as a polynomial in n of degree $2j$. See, for example, Section 1.3 of Stanley’s book [27]. Formulas for the first few Betti numbers are given by:

$$\begin{aligned} \beta_0 [\mathcal{C}(n, \mathbb{R}^2)] &= 1 \\ \beta_1 [\mathcal{C}(n, \mathbb{R}^2)] &= \frac{n(n-1)}{2} \\ \beta_2 [\mathcal{C}(n, \mathbb{R}^2)] &= \frac{(3n-1)n(n-1)(n-2)}{24} \\ \beta_3 [\mathcal{C}(n, \mathbb{R}^2)] &= \frac{n^2(n-1)^2(n-2)(n-3)}{48} \end{aligned}$$

- Finally, we define the “homological liquid” phase to be everything else. This is the most interesting regime topologically, and we are somewhat surprised to find that there is a lot of homology. Another physical metaphor for the homological liquid regime, suggested to us by Jeremy Mason, is a turbulent fluid in a pipe.

Most of our work in this paper is in estimating the Betti numbers in the homological liquid regime. For lower bounds, we use the duality between the homology of $\mathcal{C}(n, w)$ and its homology with closed support. For upper bounds, we first prove that $\mathcal{C}(n, w)$ is homotopy equivalent to a cell complex $\text{cell}(n, w)$, and then apply

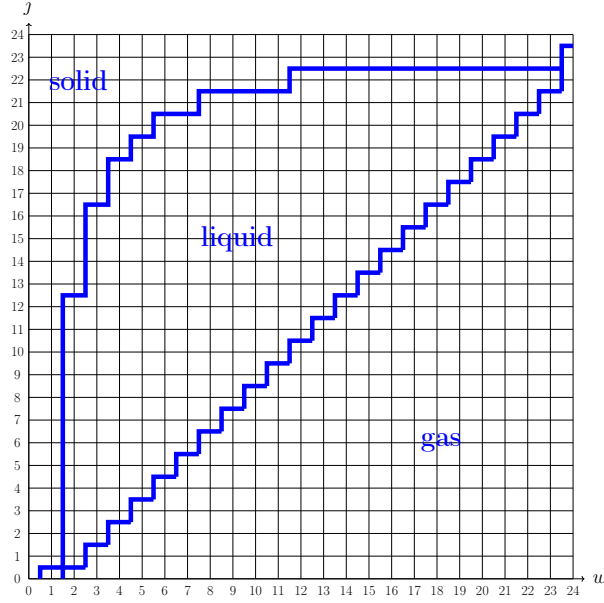


FIGURE 2. Theorem 1.2 describes the shapes of the homological solid, liquid, and gas regimes for every n . We illustrate here the case $n = 24$.

discrete Morse theory.

Some advantages of the definitions of homological solid, liquid, and gas include their simplicity, their generality, and being well defined for every finite n and not merely asymptotically. All three regimes are already visible when $n = 3$ and $j = 1$. The following describes the shapes of the regimes for every n . We emphasize that the boundary between solid and liquid regimes is more interesting for finite n than it appears to be in Theorem 1.1.

Theorem 1.2 (The phase portrait for every n).

- (1) (*Gas regime.*) If $w \geq 2$ and $0 \leq j \leq w - 2$, then the inclusion map $i : \mathcal{C}(n, w) \rightarrow \mathcal{C}(n, \mathbb{R}^2)$ induces an isomorphism on homology

$$i_* : H_j[\mathcal{C}(n, w)] \rightarrow H_j[\mathcal{C}(n, \mathbb{R}^2)].$$

Moreover, if $w \geq n$ then $\mathcal{C}(n, w)$ is homotopy equivalent to $\mathcal{C}(n, \mathbb{R}^2)$.

- (2) (*Liquid regime.*) If $1 \leq w \leq n - 1$ and $w - 1 \leq j \leq n - \lceil n/w \rceil$ then $H_j(\mathcal{C}(n, w)) \neq 0$, but the inclusion map $i : \mathcal{C}(n, w) \rightarrow \mathcal{C}(n, \mathbb{R}^2)$ does not induce an isomorphism on homology

$$i_* : H_j[\mathcal{C}(n, w)] \rightarrow H_j[\mathcal{C}(n, \mathbb{R}^2)].$$

- (3) (*Solid regime.*) If either $w = 0$, or $w \geq 1$ and $j \geq n - \lceil n/w \rceil + 1$, then

$$H_j[\mathcal{C}(n, w)] = 0.$$

Although the statement is for every n and not only asymptotically as $n \rightarrow \infty$, Theorem 1.2 is actually easier than Theorem 1.1 and is essentially a corollary of intermediate results.

The rest of the paper is organized as follows.

In Section 2, we give definitions and notation used throughout the rest of the paper. In particular, we describe a polyhedral cell complex $\text{cell}(n, w)$, a subcomplex of the Salvetti complex, which is homotopy equivalent to $\mathcal{C}(n, w)$.

In Section 3, we prove the homotopy equivalence of $\mathcal{C}(n, w)$ and $\text{cell}(n, w)$. Parts (1) and (3) of Theorem 1.1 follow immediately from the homotopy equivalence.

In Section 4, we prove lower bounds on the Betti numbers in the liquid regime, giving one direction of part (2) of Theorem 1.1.

In Section 5, we prove Theorem 1.2.

In Section 6, we describe a discrete gradient vector field on $\text{cell}(n, w)$. This allows us to collapse $\text{cell}(n, w)$ to a regular CW complex with far fewer cells.

In Section 7, we use the results from Section 6 to prove upper bounds, giving the other direction of part (2) of Theorem 1.1.

In Section 8, we close with comments and open problems.

Finally, in an appendix by Ulrich Bauer and Kyle Parsons, we include calculation of the Betti numbers for $n \leq 8$.

2. DEFINITIONS AND NOTATION

We first describe a ranked poset which we denote $\text{poset}(n)$, which is the face poset of a regular CW complex $\text{cell}(n)$ called the *Salvetti complex*. Then, afterward, we define $\text{cell}(n, w)$ as a subcomplex of $\text{cell}(n)$. The Salvetti complex and related constructions have appeared implicitly or explicitly many times—see Section 3 of [10] for a brief review of the literature. The complex was apparently first described explicitly by Salvetti in [24].

Definition 2.1. The poset which we denote $\text{poset}(n)$ has as its underlying set $\mathcal{A}(n)$, defined as follows. We call the elements of $\mathcal{A}(n)$ “symbols”. A symbol is a permutation in one-line notation $(\sigma_1 \ \sigma_2 \ \dots \ \sigma_n)$, where between each consecutive pair of elements $\sigma_i \ \sigma_{i+1}$, there can either be a bar or not.

The bars separate the permutation into pieces that we call *blocks*. The partial order on $\text{poset}(n)$ is characterized as follows: the covers in the Hasse diagram of a symbol α are the symbols obtained from α by the operation of removing a bar and merging the two adjacent blocks by a shuffle—the shuffle must preserve the relative order within each block.

The Hasse diagram of $\text{poset}(3)$ is illustrated in Figure 3. For example, $(1 \mid 3 \mid 2)$, $(3 \ 1 \mid 2)$, and $(3 \ 2 \ 1)$ are all symbols in $\mathcal{A}(3)$. Moreover, they form a chain in the poset.

There are $n - 1$ positions between consecutive pairs of elements, so there are exactly $n! 2^{n-1}$ symbols in $\mathcal{A}(n)$.

It is useful to consider “block notation” for a symbol. If we write

$$\alpha = (c_1 \mid c_2 \mid \dots \mid c_m),$$

it means that each c_i is a block of the permutation, separated from the rest of the permutation by bars. Forgetting the order of permutation elements within a block,




FIGURE 3. The Hasse diagram of $\text{poset}(3)$. This is the face poset of the Salvetti complex for the configuration space of 3 points in the plane.

we may also regard a block as a subset of $[n] := \{1, 2, \dots, n\}$. So we may write without ambiguity such statements as “ σ_k and σ_ℓ are in the same block”.

It is well known that $\text{poset}(n)$ is the face poset of a regular CW complex $\text{cell}(n)$ —see for example [10], usually called the Salvetti complex. In the more general context of complexifications of real hyperplane arrangements, it was shown in [24] that $\text{cell}(n)$ is homotopy equivalent to the configuration space of points in the plane $\mathcal{C}(n, \mathbb{R}^2)$.

The cell complex $\text{cell}(n)$ has $n! \binom{n-1}{i-1} = n! \binom{n-1}{n-i}$ i -dimensional faces, indexed by permutations with $n-i-1$ bars. If a cell is indexed by a symbol $\alpha = (c_1 | c_2 | \dots | c_m)$ with m blocks, then the cell has dimension $j = n - m$.

We will be mostly concerned with certain subcomplexes of $\text{cell}(n)$, described as follows.

Definition 2.2. For every $n, w \geq 1$, we define $\text{poset}(n, w)$ to be the sub-poset of $\text{poset}(n)$ where every block has width at most w , that is, at most w elements. We note that $\text{poset}(n, w)$ is an order ideal in $\text{poset}(n)$. Then since $\text{poset}(n)$ is the face poset of $\text{cell}(n)$, we have that $\text{poset}(n, w)$ is the face poset of a subcomplex which we denote $\text{cell}(n, w)$.


The cell complex $\text{cell}(3, 2)$ is illustrated in Figure 4.

In the remainder of the section, we define some spaces closely related to $\mathcal{C}(n, w)$. We use these spaces in Section 3 for the proof that $\mathcal{C}(n, w)$ and $\text{cell}(n, w)$ are homotopy equivalent, and in Section 4 to find lower bounds on the Betti numbers of these spaces.

Let $\mathcal{C}_d(n, w)$ denote the closed configuration space of n disks of diameter d in a strip of width 1, so that $\mathcal{C}(n, w) = \mathcal{C}_1(n, w)$. Rescaling gives homeomorphisms between $\mathcal{C}_d(n, w)$, $\mathcal{C}(n, w/d)$, and $\mathcal{C}_{d/w}(n, 1)$, so we can identify $\mathcal{C}(n, w)$ with $\mathcal{C}_{1/w}(n, 1)$. Let $\mathcal{C}_0(n, 1)$ denote the union of all $\mathcal{C}_d(n, 1)$, i.e., the configuration space of n points in the strip of width 1.

Definition 2.3. Given a symbol $\alpha \in \text{poset}(n, w)$, we define an open set U_α in $\mathcal{C}_0(n, 1)$ as follows. Write α in block notation $\alpha = (c_1 | c_2 | \dots | c_m)$, and then define the open set U_α to be the set of points $(x_1, y_1, x_2, y_2, \dots, x_n, y_n) \in \mathbb{R}^{2n}$ such that the following conditions are met.

- We have $0 < y_k < 1$ for all y_k .

FIGURE 4. The cell complex $\text{cell}(3, 2)$.

- Whenever σ_k and σ_ℓ are in the same block and $k < \ell$, we have $y_{\sigma_k} > y_{\sigma_\ell}$.
- Whenever σ_k and σ_ℓ are in different blocks and $k < \ell$, we have $x_{\sigma_k} < x_{\sigma_\ell}$.
- If σ_k and σ_ℓ are in the same block, and $\sigma_{k'}$ and $\sigma_{\ell'}$ are in different blocks, then

$$|x_{\sigma_k} - x_{\sigma_\ell}| < |x_{\sigma_{k'}} - x_{\sigma_{\ell'}}|.$$

The indices are not assumed to be distinct—in particular it may be that $k = k'$. Intuitively, elements in the same block must cluster by x -coordinate.

Given a symbol α , let $w(\alpha)$ denote the largest number of elements in any block of α . For any natural number w , let $U(n, w)$ denote the union in $\mathcal{C}_0(n, 1)$ of all U_α for which $w(\alpha) \leq w$. So $U(n, w)$ excludes exactly those configurations in $\mathcal{C}_0(n, 1)$ that have more than w points on the same vertical line. For the sake of completeness and clarity later, we rewrite this definition as follows.

Definition 2.4. We define $U(n, w)$ to the set of points $(x_1, y_1, x_2, y_2, \dots, x_n, y_n) \in \mathbb{R}^{2n}$ such that the following conditions are met.

- $(x_k, y_k) \neq (x_\ell, y_\ell)$ whenever $1 \leq k < \ell \leq n$,
- $0 < y_k < 1$ for every $1 \leq k \leq n$, and
- no $w + 1$ of the points have the same x -coordinate.

3. HOMOTOPY EQUIVALENCE

The main goal of this section is to prove the following theorem.

Theorem 3.1. *For every $n, w \geq 1$, we have a homotopy equivalence $\mathcal{C}(n, w) \simeq \text{cell}(n, w)$. Moreover, these homotopy equivalences for w and $w + 1$ commute up to homotopy with the inclusions $\text{cell}(n, w) \hookrightarrow \text{cell}(n, w + 1)$ and $\mathcal{C}(n, w) \hookrightarrow \mathcal{C}(n, w + 1)$.*

In Section 3.1 we give an overview of the parts of the proof, and in Section 3.2 we prove the technical lemmas needed to finish the proof. Then, in Section 3.3 we list a few of the immediate consequences of the homotopy equivalence.

3.1. Proof overview of Theorem 3.1. Our strategy is to use the nerve theorem. We briefly review some of the terminology and ideas of nerve theory. We say that an open cover $\mathcal{U} = (U_k)_{k \in \mathcal{K}}$ of a topological space X is *good* if every intersection of elements of \mathcal{U} is either empty or contractible. The *nerve* $N(\mathcal{U})$ of \mathcal{U} is the simplicial complex built by taking a vertex for each open set U_k and a simplex for every collection of open sets with nonempty intersection. For our purposes here, the indexing set \mathcal{K} will always be finite, so the nerve is always finite-dimensional.

The nerve theorem says that the nerve $N(\mathcal{U})$ is homotopy equivalent to the original space X . In particular, let $\{\phi_k\}_{k \in \mathcal{K}}$ be a partition of unity subordinate to \mathcal{U} , and let $\{v_k\}_{k \in \mathcal{K}}$ be the vertices of the nerve. Then the map $r: X \rightarrow N(\mathcal{U})$ defined by

$$r(x) = \sum_{k \in \mathcal{K}} \phi_k(x) v_k$$

is a homotopy equivalence.

To prove that the homotopy equivalences in Theorem 3.1 commute with the inclusions into wider strips, we use the following “extended” nerve theorem, which appears in Section 4 of Bauer, Edelsbrunner, Jablonski, and Mrozek’s [4]. (See also Proposition 4.2 in Ferry, Mischaikow, and Nanda’s [15].)

Theorem 3.2 (Extended nerve theorem [4]). *Suppose that $\mathcal{U} = (U_k)_{k \in \mathcal{K}}$ is a good open cover of a topological space X , and that $\mathcal{V} = (V_\ell)_{\ell \in \mathcal{L}}$ is a good open cover of a topological space Y . Suppose that $f: X \rightarrow Y$ is continuous, and that $g: \mathcal{K} \rightarrow \mathcal{L}$ is such that $f(U_k) \subseteq V_{g(k)}$ for every $k \in \mathcal{K}$. Let $\bar{g}: N(\mathcal{U}) \rightarrow N(\mathcal{V})$ be the linear simplicial map induced by g . Then the following diagram commutes, up to homotopy.*

$$\begin{array}{ccc} X & \xrightarrow{f} & Y \\ \downarrow r & & \downarrow r \\ N(\mathcal{U}) & \xrightarrow{\bar{g}} & N(\mathcal{V}) \end{array}$$

The other two pieces we need for the proof of Theorem 3.1 are the following two theorems, which we prove in Section 3.2.

Theorem 3.3. *$U(n, w)$ deformation retracts to $\mathcal{C}_{1/w}(n, 1)$.*

Theorem 3.4. *The cover $\{U_\alpha\}_{w(\alpha) \leq w}$ of $U(n, w)$ is good, and its nerve $N(n, w)$ is the barycentric subdivision of the regular cell complex $\text{cell}(n, w)$.*

Assuming these two theorems, we can finish the proof of Theorem 3.1.

Proof of Theorem 3.1. The homotopy equivalence between $\mathcal{C}(n, w)$ and $\text{cell}(n, w)$ is a composition of homotopy equivalences

$$\mathcal{C}(n, w) \rightarrow \mathcal{C}_{1/w}(n, 1) \rightarrow U(n, w) \rightarrow N(n, w) \rightarrow \text{cell}(n, w).$$

Specifically, let $\phi: \mathcal{C}(n, w) \rightarrow \mathcal{C}_{1/w}(n, w)$ be the rescaling map, which is a homeomorphism. The inclusion map $i: \mathcal{C}_{1/w}(n, w) \hookrightarrow U(n, w)$ is a homotopy equivalence because its homotopy inverse is the deformation retraction from Theorem 3.3. The map $r: U(n, w) \rightarrow N(n, w)$ is the homotopy equivalence given by the nerve theorem. And, applying Theorem 3.4 we let $b: N(n, w) \rightarrow \text{cell}(n, w)$ be the homeomorphism that undoes the barycentric subdivision.

We check that each of these maps commutes up to homotopy with the inclusions resulting from mapping each space to the corresponding space with w replaced by $w + 1$. The diagram

$$\begin{array}{ccc} \mathcal{C}(n, w) & \xrightarrow{i} & \mathcal{C}(n, w + 1) \\ \downarrow \phi & & \downarrow \phi \\ \mathcal{C}_{1/w}(n, 1) & \xrightarrow{i} & \mathcal{C}_{1/(w+1)}(n, 1) \end{array}$$

commutes because inclusion commutes with rescaling. The diagram

$$\begin{array}{ccc} \mathcal{C}_{1/w}(n, 1) & \xrightarrow{i} & \mathcal{C}_{1/(w+1)}(n, 1) \\ \downarrow i & & \downarrow i \\ U(n, w) & \xrightarrow{i} & U(n, w + 1) \end{array}$$

commutes because all of the maps are inclusions. The diagram

$$\begin{array}{ccc} U(n, w) & \xrightarrow{i} & U(n, w + 1) \\ \downarrow r & & \downarrow r \\ N(n, w) & \xrightarrow{i} & N(n, w + 1) \end{array}$$

commutes by the extended nerve theorem, Theorem 3.2. And, the diagram

$$\begin{array}{ccc} N(n, w) & \xrightarrow{i} & N(n, w + 1) \\ \downarrow b & & \downarrow b \\ \text{cell}(n, w) & \xrightarrow{i} & \text{cell}(n, w + 1) \end{array}$$

commutes because barycentric subdivision is functorial. \square

3.2. Technical lemmas for Theorem 3.1. First we prove Theorem 3.3, which says that $U(n, w)$ deformation retracts to $\mathcal{C}_{1/w}(n, 1)$. Let $\tau: \mathcal{C}_0(n, 1) \rightarrow (0, \infty)$ denote the function defined as follows; the paper [3] calls this the *tautological function*. For any configuration p , we set $\tau(p)$ to be the maximum value d such that $p \in \mathcal{C}_d(n, 1)$. Intuitively, we take the points in the configuration p , and consider disks of growing radius with those points as centers. When the disks first become tangent to each other or to the lines $y = 0$ or $y = 1$, the diameter of those disks is $\tau(p)$.

Because $\mathcal{C}_{1/w}(n, 1)$ is the subset of $U(n, w)$ where $\tau \geq 1/w$, our strategy for the deformation retraction is to flow along a vector field that increases τ . In the following lemma, we construct such a vector field for each U_α , and then we combine the vector fields for various α together.

Lemma 3.5. *For each set U_α , there is a continuous vector field $\{v_\alpha(p)\}_{p \in U_\alpha}$, with the property that if $\tau(p) < 1/w(\alpha)$, then τ is increasing at rate at least $2\sqrt{1/w(\alpha) - \tau(p)}$ in the direction $v_\alpha(p)$, that is, $D\tau_p(v_\alpha(p)) \geq 2\sqrt{1/w(\alpha) - \tau(p)}$.*

(Although τ is not a smooth function, it turns out that its directional derivative in every direction is well-defined, so the expression $D\tau_p(v_\alpha(p))$ makes sense.) Before proving this technical lemma, we show how it gives the deformation retraction needed for Theorem 3.3.

Proof that Lemma 3.5 implies Theorem 3.3. Let w be a natural number. To define the deformation retraction from $U(n, w)$ to $\mathcal{C}_{1/w}(n, 1)$, we combine the vector fields $v_\alpha(p)$ from Lemma 3.5 as follows to produce a vector field $v(p)$ on $U(n, w) \setminus \mathcal{C}_{1/w}(n, 1)$. Let ϕ_α be a partition of unity subordinate to the covering $\{U_\alpha\}_{w(\alpha) \leq w}$ of $U(n, w)$. That is, for any $p \in U(n, w)$, each $\phi_\alpha(p)$ is between 0 and 1, with $\phi_\alpha(p) = 0$ if $p \notin U_\alpha$, and $\sum_\alpha \phi_\alpha(p) = 1$. For each $p \in U(n, w) \setminus \mathcal{C}_{1/w}(n, 1)$, we define

$$v(p) = \sum_\alpha \phi_\alpha(p) v_\alpha(p).$$

We claim that for each p , if $\tau(p) < 1/w$ then we have $D\tau_p(v(p)) \geq 2\sqrt{1/w - \tau(p)}$; that is, $\tau(p)$ increases at rate at least $2\sqrt{1/w - \tau(p)}$ in the direction $v(p)$. We already have that each vector field $v_\alpha(p)$ increases $\tau(p)$ at this rate:

$$D\tau_p(v_\alpha(p)) \geq 2\sqrt{1/w(\alpha) - \tau(p)} \geq 2\sqrt{1/w - \tau(p)}.$$

For the convex combination $v_\alpha(p)$, we look at the configuration p of disks of diameter $\tau(p)$, and find all the places where two disks are tangent or where a disk is tangent to the boundary of the strip. For each of these tangencies, there is a linear functional on the tangent space at p , measuring how the distance between those two disks or between that disk and the boundary changes as p varies. The set of vectors at p that increase τ is the intersection of half-spaces, one for each of the tangencies. And, the set of vectors at p that increase τ at rate at least $2\sqrt{1/w - \tau(p)}$ is the intersection of half-spaces, one for each of the tangencies, cut out by planes that are parallel to the corresponding planes for finding the set of vectors at p that increase τ . Thus, this latter set is convex. The vector $v(p) = \sum_\alpha \phi_\alpha(p) v_\alpha(p)$ is a convex combination of vectors v_α that lie in the convex set, so $v(p)$ is also in the set.

The deformation retraction is given by flowing along $v(p)$ until we reach the set $\mathcal{C}_{1/w}(n, 1)$. To see how fast it finishes, let $p(t)$ be one of the trajectories, and let $\delta(t) = 1/w - \tau(p(t))$. Then we have that $\delta(t)$ is decreasing at a rate

$$|\delta'(t)| \geq 2\sqrt{\delta(t)},$$

so we have

$$\frac{\delta'(t)}{2\sqrt{\delta(t)}} \leq -1,$$

or

$$\frac{d}{dt} \sqrt{\delta(t)} \leq -1.$$

Thus the quantity $\sqrt{\delta(t)}$ decreases from at most $\sqrt{1/w}$ to 0 at rate at least 1, and so the deformation retraction finishes in time at most $\sqrt{1/w} \leq 1$. \square

The bulk of Theorem 3.3 is in constructing the vector field $v_\alpha(p)$ to prove Lemma 3.5.

Proof of Lemma 3.5. Given any α , we construct $v_\alpha(p)$ as

$$v_\alpha(p) = \lambda(p)(y(p) - p) + (x(p) - p),$$

where $x: U_\alpha \rightarrow U_\alpha$ and $y: U_\alpha \rightarrow U_\alpha$ specify configurations, and $\lambda: U_\alpha \rightarrow \mathbb{R}_{\geq 0}$ specifies a nonnegative scaling. The configuration $y(p)$ differs from p by moving the points vertically, and the configuration $x(p)$ differs from p by moving the points horizontally such that points in the same block of α keep their relationships but the horizontal space between blocks may increase. We first construct $y(p)$, then

construct $\lambda(p)$, and then construct $x(p)$ so that the resulting vector has the desired properties.

To analyze whether the vector $v_\alpha(p)$ increases $\tau(p)$ quickly enough, we observe that $\tau(p)$ is the minimum of all the distances between pairs of points in p and twice all the distances between the points and the boundary of the strip. We refer to these distances as the *measurements*, and refer to each measurement less than $1/w(\alpha)$ as *short*. It suffices to construct $v_\alpha(p)$ such that for each short measurement $m(p)$, moving along the vector $v_\alpha(p)$ increases $m(p)$ at rate at least $2\sqrt{1/w(\alpha) - m(p)}$.

We start by constructing $y(p)$ in such a way that for each short measurement within any block, moving from p to $y(p)$ increases that short measurement. Let $y(p) \in U_\alpha$ be a configuration in which the points have the same x -coordinates as in p , but the y -coordinates are evenly spaced within each block, in the following way. If a block has k elements, then the y -coordinates of those points in $y(p)$ are $1 - \frac{1}{2}k, 1 - \frac{3}{2}k, \dots, \frac{3}{2}k, \frac{1}{2}k$, so that the intervals of size $\frac{1}{k}$ around these values exactly tile the interval from 0 to 1. In the new configuration $y(p)$, each distance between points in the same block is at least $1/w(\alpha)$, and twice the distance from each point to the boundary of the strip is at least $1/w(\alpha)$.

We still need to check that the vector at p given by $y(p) - p$ infinitesimally increases each short measurement within any block. Consider any short measurement in p given by distance from a point to the boundary of the strip. In $y(p)$, this measurement is no longer short, so the vector $y(p) - p$ must move the relevant point away from the boundary of the strip, increasing that measurement.

Next, consider any short measurement in p given by distance between two points in the same block. Let a and b be the points in p , and let a' and b' be the corresponding points in $y(p)$. To determine whether the vector $y(p) - p$ increases this measurement, we need to look at the triangle formed by vectors $b - a$ and $b' - a'$. The measurement increases if and only if the angle at $b - a$ is obtuse; that is, if we have the inequality of inner products $\langle b' - a', b - a \rangle > \langle b - a, b - a \rangle$. Because our two points are in the same block, the x -coordinates of $b - a$ and $b' - a'$ are the same, while the y -coordinate of $b' - a'$ has the same sign as that of $b - a$ but has greater magnitude because our measurement is no longer short in $y(p)$. Thus, we do have the desired inequality.

Because the vector $y(p) - p$ infinitesimally increases each short measurement within any block, we can choose the scaling $\lambda(p)$ such that if $m(p)$ is a short measurement within any block, then $\lambda(p)(y(p) - p)$ increases $m(p)$ at a rate of at least $2\sqrt{1/w(\alpha) - m(p)}$. We choose the minimum possible scaling $\lambda(p)$ with this property. In particular, if p has no short measurements, or if every short measurement in p is a distance between points in different blocks, then $\lambda(p) = 0$.

Next we choose the configuration $x(p)$, which differs from p by sliding some pairs of consecutive blocks away from each other. The vector $x(p) - p$ does not change any measurements within any blocks. Thus, our task is to choose $x(p)$ such that the resulting vector $v_\alpha(p) = \lambda(p)(y(p) - p) + (x(p) - p)$ increases each short measurement between blocks by the desired amount; the measurements within blocks are already taken care of. We choose $x(p)$ such that the left-most block of p does not move. From this assumption, the configuration $x(p)$ is determined by specifying the amount of horizontal space between each pair of consecutive blocks.

Consider two consecutive blocks. If none of the distances between one point in the first block and another point in the second block are short measurements, then

we leave the distance between those two blocks the same. Otherwise, consider such a short measurement. Let a and b be the points in p that give the measurement, with a in the left block and b in the right block. Suppose our vector $v_\alpha(p)$ moves our points toward a' and b' ; that is, for some small $\varepsilon > 0$ such that $p + \varepsilon \cdot v_\alpha(p) \in U_\alpha$, let a' and b' be the corresponding points in the configuration $p + \varepsilon \cdot v_\alpha(p)$. (Note that whether $p + \varepsilon \cdot v_\alpha(p)$ is in U_α depends on $\lambda(p)$ and $y(p)$ but not on $x(p)$.) Then, as above, the measurement increases if and only we have $\langle b' - a', b - a \rangle > \langle b - a, b - a \rangle$, or equivalently if $\langle (b' - a') - (b - a), b - a \rangle > 0$. In fact, the rate that τ increases in the direction $v_\alpha(p)$ is $1/\varepsilon$ times the length of the projection of $(b' - a') - (b - a)$ onto the direction $b - a$. We claim that we can choose $x(p)$ such that this projection is sufficiently long.

Let $a = (a_1, a_2)$, $b = (b_1, b_2)$, $a' = (a'_1, a'_2)$, and $b' = (b'_1, b'_2)$. Then a_2 , b_2 , a'_2 , and b'_2 are determined by our choice of $\lambda(p)$, $y(p)$, and ε . We have that $b_1 - a_1$ is positive (because a is to the left of b), and that $(b'_1 - a'_1) - (b_1 - a_1)$ is equal to ε times the amount of additional space in $x(p)$ between the two blocks, compared to the space in p . By choosing this amount of additional space to be large, we may cause the quantity $((b'_1 - a'_1) - (b_1 - a_1))(b_1 - a_1)$ to be arbitrarily large, while keeping the quantity $((b'_2 - a'_2) - (b_2 - a_2))(b_2 - a_2)$ the same, thus making the inner product $\langle (b' - a') - (b - a), b - a \rangle$ as positive as we want. Thus, there is some choice of how much more space $x(p)$ should have than p between the two blocks, in order to have the property that for each short measurement $m(p)$ between the two blocks, the vector $v_\alpha(p)$ increases this measurement at rate at least $2\sqrt{1/w(\alpha) - m(p)}$; we choose the least possible such amount of additional space.

Repeating this computation for each pair of consecutive blocks gives $x(p)$ and thus completes the construction of $v_\alpha(p)$. The selection of $\lambda(p)$ and $y(p)$ guarantees that $v_\alpha(p)$ increases each short measurement $m(p)$ within a block at rate at least $2\sqrt{1/w(\alpha) - m(p)}$, and the selection of $x(p)$ guarantees that $v_\alpha(p)$ increases each short measurement $m(p)$ between two blocks at rate at least $2\sqrt{1/w(\alpha) - m(p)}$. Thus, $v_\alpha(p)$ increases the function $\tau(p)$, equal to the minimum of all these measurements, at rate at least $2\sqrt{1/w(\alpha) - \tau(p)}$. \square

This completes the proof of Theorem 3.3, the deformation retraction from $U(n, w)$ to $\mathcal{C}_{1/w}(n, 1)$. Next we prove Theorem 3.4, which implies using the nerve theorem that $U(n, w)$ is homotopy equivalent to $\text{cell}(n, w)$. The main thing to check is that the intersection

$$U_{\alpha_1} \cap U_{\alpha_2} \cap \cdots \cap U_{\alpha_k}$$

is nonempty if and only if the symbols

$$\{\alpha_1, \alpha_2, \dots, \alpha_k\}$$

form a chain in $\text{poset}(n, w)$; thus, a simplex in the nerve $N(n, w)$ corresponds to a chain of incident cells in $\text{cell}(n, w)$.

Proof of Theorem 3.4. Every U_α is a convex open subset of \mathbb{R}^{2n} , so every U_α is contractible and since the intersection of convex sets is convex, every nonempty intersection is contractible. Thus, the sets U_α form a good cover of $U(n, w)$.

Given a point $p \in U(n, w)$, we first describe an algorithm for finding \mathcal{A}_p , the set of symbols $\alpha \in \mathcal{C}(n, w)$ such that $p \in U_\alpha$. Along the way, we will see that \mathcal{A}_p is a chain.

We first define the poset of ordered partitions $\text{part}(n)$. An element of $\text{part}(n)$ is an ordered sequence (S_1, S_2, \dots) of non-empty subsets of $[n]$ such that the subsets S_j are pairwise disjoint, and their union is all of $[n]$.

The partial order on $\text{part}(n)$ is characterized as follows: the covers of an ordered partition π are the ordered partitions obtained from π by the operation of replacing two adjacent subsets by their union at the same place in the order. We remark that $\text{part}(n)$ is somewhat similar to $\text{poset}(n)$, but in $\text{part}(n)$ we forget the order of the elements within a block.

Now let a point

$$p = (x_1, y_1, x_2, y_2, \dots, x_n, y_n) \in U(n, w)$$

be given. We find the set of U_α containing p in the following way.

Step 1 produces a chain $\pi_1, \pi_2, \dots, \pi_m$ in the poset $\text{part}(n)$. This step uses the x coordinates of p but not the y coordinates. We say that x_k and x_ℓ are *consecutive x values of p* if $x_k < x_\ell$ and there is no k' for which $x_k < x_{k'} < x_\ell$.

For $i = 1, 2, \dots, n$ and $\rho \geq 0$, define $K_i(\rho) := [x_i - \rho/2, x_i + \rho/2]$, i.e., the closed interval of length ρ centered on x_i . Let $K(\rho) := \bigcup K_i(\rho)$ be the union of all the intervals $K_i(\rho)$. Every path component of $K(\rho)$ is a union of finitely many closed intervals and is connected, hence is a closed interval itself. We cluster the integers $[n]$ according to which path connected component of K they lie in. In other words, for $i \in [n]$ we say that $i \in S_k(\rho)$ if x_i is in the k th connected component of $K(\rho)$, counting left to right.

When $\rho = 0$, k and ℓ lie in the same cluster if and only if $x_k = x_\ell$. When ρ is sufficiently large, there is only one cluster and $S_1(\rho) = [n]$. In general, the ordered partition $\pi(\rho)$ changes only at certain values of ρ , namely the differences of consecutive x values. So as ρ increases, we get a finite sequence of distinct ordered partitions π_1, π_2, \dots . This sequence is the desired chain in $\text{part}(n)$.

Step 2 produces a chain $\tilde{\pi}_1, \tilde{\pi}_2, \dots, \tilde{\pi}_{m'}$ in $\text{poset}(n, w)$. Here $m' \leq m$ and for every $1 \leq i \leq m'$, the symbol $\tilde{\pi}_i \in \text{poset}(n, w)$ is a “lift” of the ordered partition $\pi_i \in \text{part}(n)$. This step uses the y coordinates of p but not the x coordinates.

Given a partition $\pi_i = (S_1, S_2, \dots)$ produced in step one, order the elements within each subset S_k to produce $\tilde{\pi}_i \in \text{poset}(n, w)$, in such a way that if σ_ℓ and $\sigma_{\ell'}$ are elements of S_k with σ_ℓ before $\sigma_{\ell'}$ in the ordering (that is, $\ell < \ell'$), then $y_{\sigma_\ell} > y_{\sigma_{\ell'}} + 1$. If for some S_k this can not be done, then discard π_i from the chain and exclude it from further consideration. We can also discard, then, π_k for any $k > i$; any such ordered partition is just made by merging elements of π_i , so it is still impossible to order within a part by y coordinate.

It is immediate that the chain constructed from steps one and two is the desired \mathcal{A}_p , from the definition of U_α .

Finally, we check that an intersection

$$U_{\alpha_1} \cap U_{\alpha_2} \cap \dots \cap U_{\alpha_k}$$

is nonempty if and only if the symbols

$$\{\alpha_1, \alpha_2, \dots, \alpha_k\}$$

form a chain in $\text{poset}(n, w)$.

First of all, if the intersection is nonempty then we apply the algorithm to any point p in the intersection. The set of U_α containing p gives a chain that includes all of $\alpha_1, \dots, \alpha_k$, and any sub-poset of a chain is a chain, so $\alpha_1, \dots, \alpha_k$ must form a chain.

Now suppose we have a chain $\alpha_1 < \alpha_2 < \dots < \alpha_m$ in $\text{poset}(n, w)$. We produce a point p such that $p \in U_{\alpha_i}$ for every i . First we choose the y -coordinates. None of the blocks in any symbol have width more than w , so there is enough room vertically in $U(n, w)$ to ensure that if σ_k and σ_ℓ are in the same block of α_m with $k < \ell$, then we can choose $y_{\sigma_k} > y_{\sigma_\ell} + 1$.

To choose the x -coordinates, we start by assuming without loss of generality that the chain is maximal, so getting to each symbol α_i from the previous symbol α_{i-1} corresponds to merging two consecutive blocks. We start with α_1 and add restrictions on the x -coordinates one step at a time, so that on the i th step we will have fixed the differences between x -coordinates within each block of α_i , but we think of the separate blocks sliding freely from side to side. After all the steps, we will have specified the configuration up to horizontal translation.

More precisely, at step 1 we require that if $k < \ell$ —that is, if σ_k appears before σ_ℓ in α_1 —then $x_{\sigma_k} < x_{\sigma_\ell}$, with no other restrictions. Any such configuration is in U_{α_1} . At step 2, two consecutive elements in α_1 together form a block of width 2, and we introduce the restriction that their x -coordinates have difference $9 = 3^2$. Then, continuing in the same way, at step i two consecutive blocks c_{k_i} and $c_{k_{i+1}}$ in α_{i-1} merge to give α_i . We introduce the restriction that the difference in x -coordinates between the first element of c_{k_i} and the last element of $c_{k_{i+1}}$ —where “first” and “last” are still taken in terms of the first symbol α_1 —should be 3^i .

Any configuration that satisfies the restrictions up through step i and also leaves horizontal gaps larger than 3^i between the blocks of α_i is in U_{α_i} . Note that step i sets the gap between blocks c_{k_i} and $c_{k_{i+1}}$ of α_{i-1} to be more than 3^{i-1} , which is what we need in order for the final configuration to be in $U_{\alpha_{i-1}}$. This is because, if we use the word “width” here to mean the range of x -coordinates, the widths of c_{k_i} and $c_{k_{i+1}}$ have been set to be distinct powers of 3 less than 3^i , or to be 0 if the block has only one element. Thus, the gap has size at most $3^i - (3^{i-1} + 3^{i-2}) > 3^{i-1}$.

In the final step, step n means merging two blocks to get α_n which has only one block, and at step n we set the difference $x_{\sigma_n} - x_{\sigma_1}$ to be 3^n (here, the numbering $\sigma_1, \dots, \sigma_n$ is still taken in terms of the first symbol α_1). At this stage we have specified the configuration p up to translation, and it is in $U_{\alpha_1} \cap \dots \cap U_{\alpha_n}$.

The barycentric subdivision of $\text{cell}(n, w)$ has one vertex for every cell in $\text{cell}(n, w)$, and one simplex for every chain of incident cells in $\text{cell}(n, w)$. The nerve $N(n, w)$ has one vertex for each U_α , and thus for each cell in $\text{cell}(n, w)$. And, we have just shown that every set of U_α with nonempty intersection—corresponding to a simplex in $N(n, w)$ —corresponds to a chain of incident cells in $\text{cell}(n, w)$, and vice versa. Thus, $N(n, w)$ is equal to the barycentric subdivision of $\text{cell}(n, w)$. \square

3.3. Consequences of the homotopy equivalence. One immediate consequence of the homotopy equivalence is Part (1) of Theorem 1.1, i.e., given a sufficiently wide strip we have an isomorphism on homology with the configuration space of points in the plane.

Proof of Part (1) of Theorem 1.1. We show that if $w \geq j + 2$, then we have an isomorphism between $H_j[\mathcal{C}(n, w)]$ and $H_j[\mathcal{C}(n, \mathbb{R}^2)]$. Moreover, this isomorphism is induced by the inclusion map $i : \mathcal{C}(n, w) \rightarrow \mathcal{C}(n, \mathbb{R}^2)$.

Every cell in $\text{cell}(n)$ but not in $\text{cell}(n, w)$ is indexed by a symbol with at least one block of width at least $w + 1$. Hence every such cell has dimension at least w . Therefore, the subcomplex $\text{cell}(n, w)$ of $\text{cell}(n)$ has the same $(w - 1)$ -skeleton as $\text{cell}(n)$, so the inclusion $i : \text{cell}(n, w) \hookrightarrow \text{cell}(n)$ induces an isomorphism of homology in degrees $\leq w - 2$.

By Theorem 3.1, the homotopy equivalences

$$\text{cell}(n, w) \rightarrow U(n, w) \rightarrow \mathcal{C}(n, w)$$

and

$$\text{cell}(n) = \text{cell}(n, n) \rightarrow U(n, n) = \mathcal{C}(n, \mathbb{R}^2)$$

commute up to homotopy with the inclusions $\text{cell}(n, w) \hookrightarrow \text{cell}(n, n)$, $U(n, w) \hookrightarrow U(n, n)$, and $\mathcal{C}(n, w) \hookrightarrow \mathcal{C}(n, \mathbb{R}^2)$, and so the inclusion $i : \mathcal{C}(n, w) \hookrightarrow \mathcal{C}(n, \mathbb{R}^2)$ also induces an isomorphism on homology in degrees $\leq w - 2$. \square

Another consequence of the homotopy equivalence is that $\text{cell}(n, 2)$ is an Eilenberg–MacLane space. This cell complex is a cube complex because for each dimension j , each j -cell is labeled by a symbol that has j blocks of size 2 and all other blocks of size 1. We can take this cell to be a j -dimensional cube.

Theorem 3.6. *The cubical complex $\text{cell}(n, 2)$ admits a locally-CAT(0) metric. As a corollary, $\mathcal{C}(n, 2)$ is aspherical, i.e., has a contractible universal cover. So $\pi_j(\mathcal{C}(n, 2)) = 0$ for $j \geq 2$.*

Proof. This follows immediately from Gromov’s criterion for a cube complex to admit locally-CAT(0) metric [18]. The only thing to check is that the link of every vertex in $\text{cell}(n, 2)$ is a “flag” simplicial complex. A simplicial complex is said to be flag if it is the clique complex of its underlying graph—i.e., if it is maximal with respect to its 1-skeleton. A precise statement and complete proof of Gromov’s criterion can be found in Appendix I.6 of Davis’s book [12].

Checking that the link of a vertex in $\text{cell}(n, 2)$ is flag is straightforward. Let v be a vertex in $\text{cell}(n, 2)$, and consider the link of v , $L = \text{lk}(v)$. The vertex v corresponds to a symbol in $\text{poset}(n, 2)$ where every block has width 1.

The vertices of L correspond to elements $\sigma \geq v$ in $\text{poset}(n, 2)$ where every block in σ has width 1 except one block of width 2. Similarly, edges in L correspond to symbols $\tau \geq v$ where every block in τ has width 1 except two blocks of width 2. These two blocks of width 2 are disjoint pairs in $[n]$.

Suppose v_1, v_2, \dots, v_k span a k -clique in $\text{lk}(v)$. Then every pair of vertices corresponds to a disjoint pair of elements in $[n]$, and then concatenating these k disjoint pairs (and respecting the order within each pair) gives a symbol $C \in \mathcal{A}(n, 2)$ with k blocks of width 2 and all the remaining blocks of width 1. This symbol C indexes a k -dimensional cube in $\text{cell}(n, 2)$, which corresponds to a $(k - 1)$ -dimensional simplex in $\text{lk}(v)$.

For example, let $v = (4 \mid 6 \mid 5 \mid 1 \mid 3 \mid 2)$. The vertices $(46 \mid 5 \mid 1 \mid 3 \mid 2)$, $(4 \mid 6 \mid 15 \mid 3 \mid 2)$, and $(4 \mid 6 \mid 5 \mid 1 \mid 23)$ span a clique in $\text{link}(v)$. Since the symbol $(46 \mid 15 \mid 23)$ corresponds to a 3-dimensional cube in $\text{cell}(6, 2)$, the clique in $\text{link}(v)$ is filled in by a 2-dimensional face. \square

On the other hand, $\mathcal{C}(n, w)$ is not an Eilenberg–MacLane space if $3 \leq w \leq n - 1$. Indeed, the 2-skeleton of $\text{cell}(n, w)$ is the same as the 2-skeleton of $\text{cell}(n)$ when w is in this intermediate range. So if $\mathcal{C}(n, w)$ were a $K(\pi, 1)$, its homology would have to agree with the configuration space of points but we will see in Section 5 that it does not.

4. ASYMPTOTIC LOWER BOUNDS

In this section, we exhibit a large number of linearly independent cycles to prove lower bounds on Betti numbers. The following is well known.

Lemma 4.1. *Suppose that M is an open d -dimensional manifold, with submanifolds Z_1, Z_2, \dots, Z_k and $Z_1^*, Z_2^*, \dots, Z_k^*$ satisfying the following.*

- (1) *Every Z_i is a compact orientable j -dimensional submanifold without boundary,*
- (2) *every Z_i^* is a closed orientable $(d - j)$ -dimensional submanifold without boundary,*
- (3) *whenever $a \neq b$ we have that $Z_a \cap Z_b^* = \emptyset$, and*
- (4) *Z_a intersects Z_a^* transversely in a point for every a .*

Then for any choice of coefficient field for the homology, we have $\dim H_j(M) \geq k$.

Proof. Choose orientations of each Z_i and let $[Z_i]$ in $H_j(M)$ be the fundamental class of Z_i . Choose orientations of each Z_i^* and let $[Z_i^*]$ in $H_{d-i}^{BM}(M)$ be the fundamental class of Z_i^* . (Here H_*^{BM} denotes homology with closed supports, or Borel–Moore homology.)

Choose an orientation of M so that the intersection pairing

$$p : H_i(M) \times H_{d-i}^{BM}(M) \rightarrow \mathbb{R}$$

is defined. By the stated properties of the manifolds Z_i and Z_i^* , this intersection pairing satisfies

- $p([Z_a], [Z_b^*]) = 0$ for $a \neq b$,
- $p([Z_a], [Z_a^*]) = \pm 1$.

Therefore, the homology classes $[Z_1], [Z_2], \dots, [Z_k]$ are linearly independent in $H_i(M)$, so the dimension of $H_i(M)$ is at least k . □

In what follows, rather than working with the space $\mathcal{C}(n, w)$ directly, it is most convenient to apply Lemma 4.1 to the space which we denote by $wU(n, w) - (0, \frac{w}{2})$, consisting of configurations of points in the strip $\mathbb{R} \times [-\frac{w}{2}, \frac{w}{2}]$ in which no $w+1$ points have the same x -coordinate. Because $wU(n, w) - (0, \frac{w}{2})$ is an open subset of \mathbb{R}^{2n} , it is an open manifold and thus is appropriate for Lemma 4.1. It is homeomorphic to $U(n, w)$, which we have shown in Section 3 is homotopy equivalent to $\mathcal{C}(n, w)$. Lemma 4.1 gives a lower bound on the Betti numbers of $wU(n, w) - (0, \frac{w}{2})$, giving the same lower bound on the Betti numbers of $\mathcal{C}(n, w)$.

Definition 4.2. Let $j = q(w - 1) + r$ with $0 \leq r < w - 1$. A *special symbol* $\alpha \in \mathcal{A}(n, w)$ is a symbol $(c_1 \mid c_2 \mid \dots \mid c_m)$ such that

- (1) α has q blocks of width w , r blocks of width 2, and all other blocks of width 1,
- (2) in every block, the largest element appears first, and




FIGURE 5. Upper picture: a point on the closed 34-dimensional submanifold Z_α^* . Middle picture: a point on the compact 12-dimensional submanifold Z_α . Bottom picture: the single point of transverse intersection $Z_\alpha^* \cap Z_\alpha$. In all three pictures, $\alpha \in \mathcal{A}(23, 5)$ is the special symbol

$$\alpha = (19 \mid 13 \ 4 \mid 5 \mid 23 \ 6 \ 11 \ 1 \ 7 \mid 17 \mid 14 \mid 10 \ 3 \mid 9 \ 8 \mid 21 \ 18 \ 15 \ 12 \ 2 \mid 22 \ 16 \mid 20).$$

- (3) if c_i and c_{i+1} are consecutive blocks of width strictly less than w , then the first element of block c_i is greater than the first element of block c_{i+1} .

Which symbols are special depends on n, j, w , but for the sake of simplicity we suppress these in the notation since these parameters are always implicit.

Definition 4.3. For every special symbol α , we define a closed submanifold Z_α^* in $wU(n, w) - (0, \frac{w}{2})$ as follows.

- (1) If σ_k and σ_ℓ are in the same block and $k < \ell$, then $x_{\sigma_k} = x_{\sigma_\ell}$ and $y_{\sigma_k} > y_{\sigma_\ell}$.
- (2) If σ_k and σ_ℓ are in different blocks and $k < \ell$, and either σ_k or σ_ℓ is in a block of width w , then $x_{\sigma_k} < x_{\sigma_\ell}$.

See, for example, the upper picture in Figure 5. It is clear that every Z_α^* is closed in $wU(n, w) - (0, \frac{w}{2})$, by Definition 2.4. Indeed, the strict inequalities $y_{\sigma_k} > y_{\sigma_\ell}$ and $x_{\sigma_k} < x_{\sigma_\ell}$ in Definition 4.3 could be replaced by weak inequalities $y_{\sigma_k} \geq y_{\sigma_\ell}$ and $x_{\sigma_k} \leq x_{\sigma_\ell}$, since at most w of the x -coordinates can be equal.

Now, for every special symbol α we describe a cycle with the desired intersection properties.

Theorem 4.4. *Given a special symbol α , there exists a cycle Z_α represented by an j -dimensional torus embedded in $wU(n, w) - (0, \frac{w}{2})$, such that whenever $\alpha' \neq \alpha$ we have that $Z_\alpha \cap Z_{\alpha'}^* = \emptyset$, and such that for every α , Z_α intersects Z_α^* transversely in a point.*

Proof. The cycle Z_α actually lies in the configuration space of disks of diameter 1, which we can denote by $\mathcal{C}(n, w) - (0, \frac{w}{2})$ and which is a subspace of $wU(n, w) - (0, \frac{w}{2})$. The general idea of the construction is illustrated in the middle picture in Figure 5. For every block c_i , we construct a $(w(c_i) - 1)$ -dimensional torus in the configuration space of only the disks appearing in that block. Then we put the configurations for the different blocks horizontally next to each other in the strip in sequential order, making a j -parameter family of configurations that is an embedded j -dimensional torus.

Looking more closely, for each individual block c_i , the corresponding $(w(c_i) - 1)$ -dimensional torus is constructed roughly as follows. We can spin the first two disks around each other inside a disk of diameter 2. Then we can spin the third disk around the first two disks, all inside a disk of diameter 3. Iterating this process, the final result is a disk of diameter $w(c_i)$, with the final disk of the block circling around its inner edge, and with the remaining disks moving around inside a disk of diameter $w(c_i) - 1$ tangent to the final disk. In this way, $w(c_i)$ disks can move around inside a disk of diameter $w(c_i)$ to make a $(w(c_i) - 1)$ -dimensional torus in their configuration space.

More precisely, we construct the cycle Z_α as follows. We parameterize the torus as

$$(S^1)^j = \{(\theta_1, \theta_2, \dots, \theta_j) \mid \theta_i \in [0, 2\pi), i = 1, 2, \dots, j\}.$$

Given a symbol α and angles $(\theta_1, \theta_2, \dots, \theta_j)$, we need to specify a configuration in $\mathcal{C}(n, w) - (0, \frac{w}{2})$.

We compute the coordinates $x_1, y_1, \dots, x_n, y_n$ as follows. Let $\alpha = (c_1 \mid c_2 \mid \dots \mid c_m)$ be a special symbol. Let $w(c_i)$ denote the width of c_i . By the definition of special symbol, the blocks in α are all of width 1, 2, or w .

Let

$$\begin{aligned} X_1 &= w(c_1)/2 \\ X_2 &= w(c_1) + w(c_2)/2 \\ &\dots \\ X_i &= w(c_1) + w(c_2) + \dots + w(c_{i-1}) + w(c_i)/2 \end{aligned}$$

This variable tells us how far to horizontally shift the center of the torus for the next block. So, the imagined disk of diameter $w(c_i)$ that contains the disks of the i th block has center at $(X_i, 0)$.

Let

$$D_i = w(c_1) + w(c_2) + \dots + w(c_i) - i.$$

This is a counter which tells us which angle we are on. That is, the first $w(c_1) - 1$ angles correspond to the first block, the next $w(c_2) - 1$ angles correspond to the second block, and so on, so that angle D_i is the last angle corresponding to the i th block.

- (1) If $w(c_i) = 1$, that is, c_i is the block with a single permutation element $c_i = |\sigma_m|$, then set

$$(x_{\sigma_m}, y_{\sigma_m}) = (X_i, 0).$$

- (2) If $w(c_i) = 2$, that is, $c_i = |\sigma_m \sigma_{m+1}|$, then set

$$(x_{\sigma_m}, y_{\sigma_m}) = (X_i - \frac{1}{2} \cos \theta_{D_i}, -\frac{1}{2} \sin \theta_{D_i}),$$

$$(x_{\sigma_{m+1}}, y_{\sigma_{m+1}}) = (X_i + \frac{1}{2} \cos \theta_{D_i}, \frac{1}{2} \sin \theta_{D_i}),$$

so that θ_{D_i} is the direction of the vector from disk σ_m to disk σ_{m+1} .

- (3) If $w(c_i) = w$, i.e., $c_i = |\sigma_m \sigma_{m+1} \dots \sigma_{m+w-1}|$, then

(a) Initialize $(u_0, v_0) = (X_i, 0)$.

(b) For $k = 1, 2, \dots, w$, we let (u_k, v_k) be the center of the imagined disk containing the first $w - k$ disks of the block, and $\theta_{D_{i+1-k}}$ will be the direction of the vector from (u_k, v_k) to the k th-to-last disk of the block. That is, for $k = 1$ we let

$$(u_1, v_1) = (u_0, v_0) - \frac{1}{2}(\cos \theta_{D_i}, \sin \theta_{D_i}),$$

and

$$(x_{\sigma_{m+w-1}}, y_{\sigma_{m+w-1}}) = (u_0, v_0) + \frac{w-1}{2}(\cos \theta_{D_i}, \sin \theta_{D_i}),$$

and for general k we let

$$(u_k, v_k) = (u_{k-1}, v_{k-1}) - \frac{1}{2}(\cos \theta_{D_{i+1-k}}, \sin \theta_{D_{i+1-k}}),$$

and

$$(x_{\sigma_{m+w-k}}, y_{\sigma_{m+w-k}}) = (u_{k-1}, v_{k-1}) + \frac{w-k}{2}(\cos \theta_{D_{i+1-k}}, \sin \theta_{D_{i+1-k}}).$$

This completes the construction of the cycle Z_α .

Now we must check that whenever $\alpha' \neq \alpha$ we have that $Z_\alpha \cap Z_{\alpha'}^* = \emptyset$, and that Z_α intersects Z_α^* transversely in a point for every α .

Suppose that

$$p = (x_1, y_1, \dots, x_n, y_n) \in Z_\alpha \cap Z_{\alpha'}^*.$$

Define an equivalence relation on $[n]$ by setting $k \sim \ell$ if $x_k = x_\ell$ in p . By the definition of cycle Z_α , if $k \sim \ell$ then k and ℓ are in the same block of α . By the definition of $Z_{\alpha'}^*$, if k and ℓ are in the same block of α' , then $k \sim \ell$. So then if $p \in Z_\alpha \cap Z_{\alpha'}^*$, if k and ℓ are in the same block of α' , then they are in the same block of α .

By assumption, both α and α' are special symbols in $\mathcal{A}(n, w)$, so they both have q blocks of width j , r blocks of width 2 and the remaining blocks of width one. So it must be that the converse is also true, that if k and ℓ are in the same block of α , then they are in the same block of α' .

Moreover, the partition of $[n]$ given by the equivalence relation \sim must be the same as the partition into blocks given by α and α' . So the elements within every block are vertically aligned. In the special symbol α' , the first element of a block is greatest in the underlying permutation, and in $Z_{\alpha'}^*$ it corresponds to the element at the top of the column (i.e., has the largest y -coordinate).

Consider any point on Z_α in which the elements in each block are vertically aligned, with the greatest element of each block (and hence first element, because α is special) on top. We claim that there is only one such point, and that the order of the elements in each block is the same as the order of their corresponding disks from top to bottom.

The first element in the block is on top by assumption. Then since the first two elements are vertically aligned and in an imagined disk of diameter 2, the next element of the block must lie immediately below the first element. Continuing by induction, if the first k elements of the block are vertically aligned and in an imagined disk of diameter k , then the k th element of the block must be immediately below the $(k-1)$ st element. Thus, the configuration is completely determined by the assumption that the first element of the block is on top, proving that if Z_α and $Z_{\alpha'}^*$ intersect, then $\alpha = \alpha'$ and the intersection is a single point.

The only thing left to verify is that in this case the intersection is transverse. Since Z_α and Z_α^* intersect at a single point and are of complementary dimension in the ambient manifold, the claim of transversality is equivalent to checking that the tangent space to $U(n, w)$ is the direct sum of the tangent spaces to Z_α and Z_α^* . Roughly, the tangent space to Z_α^* corresponds to the set of ways to assign a vector to each disk such that for each block, the horizontal components are equal; in the tangent space to Z_α , the vector corresponding to each disk is horizontal and for each block there is a single linear relation on the vectors. We omit the details. \square

Finally, we are ready to prove lower bounds on the Betti numbers of $\mathcal{C}(n, w)$.

Proof of lower bounds for part (2) of Theorem 1.1. We have just verified the conditions of Lemma 4.1, which then implies that if $n \geq qw + 2r$ (i.e., for sufficiently large n) then

$$\beta_j \geq \binom{n}{\underbrace{w, \dots, w}_q \text{ times}, \underbrace{2, \dots, 2}_r \text{ times}, n - qw - 2r} q! ((w-1)!)^q (q+1)^{n-qw-r}.$$

This counts the number of special symbols in $\mathcal{A}(n, w)$. The multinomial coefficient counts the number of ways to partition n into q subsets of size w , r subsets of size 2, and $n - qw - 2r$ subsets of size 1. There are $q!$ ways to order the subsets of size w , and $((w-1)!)^q$ ways to order the terms in each subset, considering the restriction that the largest element must come first within each part. Finally, we place the blocks of width 2 and 1 between the blocks of width w , and there are $(q+1)^{n-qw-r}$ ways to do this.

If j and w are fixed and $n \rightarrow \infty$, then we write the simpler asymptotic expression

$$\beta_j [\mathcal{C}(n, w)] = \Omega((q+1)^n n^{qw+2r}).$$

\square

Here $f = \Omega(g)$ means that there exists a positive constant c such that $f(n) \geq cg(n)$ for all sufficiently large n .

5. THE PHASE PORTRAIT FOR EVERY n

In this section we prove Theorem 1.2. Everything follows quickly from the homotopy equivalence $\mathcal{C}(n, w) \simeq \text{cell}(n, w)$ in Section 3, and the non-triviality of the

cycles constructed in Section 4.

Proof of Theorem 1.2.

- (1) Gas: This is identical to the proof of (1) of Theorem 1.1, in Section 3.3. The proof of isomorphism on homology holds for every n .
- (2) Liquid: If $1 \leq w \leq n - 1$ and $0 \leq j \leq n - \lceil n/w \rceil$ we see first that $H_j[\mathcal{C}(n, w)] \neq 0$. Indeed, the cycles constructed in Section 4 are already enough. One can partition $[n]$ into at most $\lceil n/w \rceil$ blocks of width at most w . By ordering elements within a block, and reordering blocks if necessary, then we have a special symbol α with at most $\lceil n/w \rceil$ blocks. This indexes a cycle Z_α of dimension at least $n - \lceil n/w \rceil$.

We next see that if $j \geq w - 1$ then the inclusion map $i : \mathcal{C}(n, w) \rightarrow \mathcal{C}(n, \mathbb{R}^2)$ does not induce an isomorphism on homology

$$i_* : H_j[\mathcal{C}(n, w)] \rightarrow H_j[\mathcal{C}(n, \mathbb{R}^2)].$$

We observe that the kernel of i_* is nontrivial. Consider two different torus cycles Z_α and $Z_{\alpha'}$, indexed by two different special symbols $\alpha, \alpha' \in \mathcal{A}(n, w)$ where α' is obtained from α by transposing two blocks (keeping the order of the elements within a block). Since $n \geq w + 1$ and $j \geq w - 1$, this is always possible. Indeed, the condition that $j \geq w - 1$ ensures that α and α' have at least one block of width w , and the condition that $n \geq w + 1$ ensures that there is at least one other block.

We have shown that Z_α and $Z_{\alpha'}$ are not homologous in $\mathcal{C}(n, w)$, but $i_*(Z_\alpha)$ and $i_*(Z_{\alpha'})$ are homologous in $\mathcal{C}(n, \mathbb{R}^2)$, so we conclude that $Z_\alpha - Z_{\alpha'}$ is in the kernel of i_* .

- (3) Solid: Finally, we check that if $w \geq 1$ and $j \geq n - \lceil \frac{n}{w} \rceil + 1$, then

$$H_j[\mathcal{C}(n, w)] = 0.$$

We know from Section 3 that $\mathcal{C}(n, w) \sim \text{cell}(n, w)$. The largest dimension of a cell in $\text{cell}(n, w)$ is $n - \lceil \frac{n}{w} \rceil$, since the minimum number of blocks is $\lceil \frac{n}{w} \rceil$. So if $j \geq n - \lceil \frac{n}{w} \rceil + 1$, then there are no j -dimensional cells, in which case there is no nontrivial j -dimensional homology.

□

6. DISCRETE MORSE THEORY

In this section, we describe a discrete gradient vector field on $\text{cell}(n, w)$. Then in the next section, we prove an upper bound on the number of critical cells, thus giving an upper bound on the Betti numbers. This upper bound completes the proof of Theorem 1.1, the asymptotic rate of growth of Betti numbers.

A *discrete vector field* V on a regular CW complex X is a collection of pairs of faces $[\alpha, \beta]$ where α is a face of β and $\dim \alpha = \dim \beta - 1$, and such that every face can be in at most one pair. The discrete vector field V is said to be *gradient* if there are no closed V -walks. A V -walk is a collection of pairs of faces $[\alpha_1, \beta_1], [\alpha_2, \beta_2], \dots, [\alpha_r, \beta_r]$ where $[\alpha_i, \beta_i] \in V$ for every i and α_{i+1} is a codimension 1 face of β_i other than α_i , and the V -walk is *closed* if $\alpha_r = \alpha_1$.

We call a face *critical* if it is not in any pair. The fundamental theorem of discrete Morse theory [16] is that X is homotopy equivalent to a CW complex X' , where X' has exactly one cell for every critical face in V . Any discrete gradient vector field gives an upper bound on the Betti numbers of the cell complex: each Betti number is at most the number of critical cells in the corresponding dimension. So, we give an asymptotic upper bound on the number of critical cells to get an asymptotic upper bound on the Betti numbers.

We begin by describing which cells will be critical with respect to the discrete gradient vector field that we will construct. In the symbol of a cell in $\text{cell}(n, w)$, we say that a block is *top-heavy* if the largest element of that block is the first element. We designate some pairs of blocks as leader/follower pairs, as follows. Roughly, we look at pairs of consecutive blocks, starting with the first two blocks and moving to the right. We may designate a pair of blocks to be a leader (the first block) and a follower (the next block), in which case we look next at the two blocks after these, so that no follower gets immediately also labeled a leader, and the leader/follower pairs are disjoint.

More precisely, we say that a block is a *leader* if it is not a follower and its first element is larger than all the other elements of that block and also all the elements of the next block; we say that a block is a *follower* if the previous block is a leader. These definitions allow us to describe the critical cells of our discrete gradient vector field. We say that a cell of $\text{cell}(n, w)$ is *k-crit* if the following is true for the first k blocks: every block that is not top-heavy is a follower, and every leader/follower pair has greater than w elements, combined. Our goal is to verify that this definition of *k-crit* agrees with which cells are critical with respect to the discrete gradient vector field we will construct.

Theorem 6.1. *There is a discrete gradient vector field V on $\text{cell}(n, w)$, such that the critical cells are exactly those that are k -crit for all k .*

In order to define the discrete vector field V , we describe how to find the matching cell for each non-critical cell of $\text{cell}(n, w)$. We define a function v that sends each cell to its matching cell; that is, if $[\alpha, \beta]$ is a pair in V , then we will have $v(\alpha) = \beta$ and $v(\beta) = \alpha$, and for any critical cell α , we will have $v(\alpha) = \alpha$. The definition of v is as follows. Given a cell α , if α is k -crit for all k , then we set $v(\alpha) = \alpha$. Otherwise, we find k such that α is $(k - 1)$ -crit but not k -crit. There are two possibilities:

- (1) The $(k - 1)$ st block is a leader, the k th block is a follower, and their combined number of elements is at most w ; or
- (2) The k th block is not a follower and is not top-heavy.

We refer to the first case as the “match-up at $k - 1$ ” case, and we refer to the second case as the “match-down at k ” case. In the first case, we obtain $v(\alpha)$ by swapping the $(k - 1)$ st block with the k th block and removing the bar between them. In the second case, we obtain $v(\alpha)$ by adding a bar just before the largest element of the k th block, to separate it into two blocks, and then swapping those two blocks. In order to be able to use v to define V , we need to check that v actually matches the cells in pairs.

Lemma 6.2. *The function v is an involution; that is, we have $v(v(\alpha)) = \alpha$ for every cell α of $\text{cell}(n, w)$.*

Proof. Suppose that α is a cell in the match-up at $k - 1$ case. We want to show that $v(\alpha)$ is in the match-down at $k - 1$ case. We know that $v(\alpha)$ is $(k - 2)$ -crit because

α and $v(\alpha)$ agree in the first $k - 2$ blocks. Suppose for the sake of contradiction that block $k - 1$ of $v(\alpha)$ is a follower. Then block $k - 1$ of α is also a follower, because in both cases the previous block is the same and the current block has the same largest element. But we know that block $k - 1$ of α is a leader and thus not a follower, giving a contradiction. So block $k - 1$ of $v(\alpha)$ is not a follower. It is clear from the construction that block $k - 1$ of $v(\alpha)$ is not top-heavy, so $v(\alpha)$ is in the match-down at $k - 1$ case, and it is also clear from the construction that applying v to $v(\alpha)$ gives α again.

Now suppose that α is a cell in the match-down at k case. We want to show that $v(\alpha)$ is in the match-up at k case. We know that $v(\alpha)$ is $(k - 1)$ -crit because α and $v(\alpha)$ agree in the first $k - 1$ blocks. To show that $v(\alpha)$ is k -crit, we need to check that block k of $v(\alpha)$ is top-heavy and is not a follower. It is clear from the construction that block k of $v(\alpha)$ is top-heavy. Suppose for the sake of contradiction that block k of $v(\alpha)$ is a follower. Then block k of α is also a follower, because in both cases the previous block is the same and the current block has the same largest element. But we know that block k of α is not a follower, because α is in the match-down at k case. Thus block k of $v(\alpha)$ cannot be a follower, and so $v(\alpha)$ is k -crit. Knowing that block k of $v(\alpha)$ is not a follower, it is clear from the construction that this block is a leader and that its combined number of elements with the next block is at most w , so $v(\alpha)$ is in the match-up at k case. Then it is also clear from the construction that applying v to $v(\alpha)$ gives α again.

Thus if α is in any of the three cases—critical, match-up, or match-down—we have $v(v(\alpha)) = \alpha$. \square

Having shown that every orbit of v has either one or two elements, we can define V to be the set of two-element orbits; that is, if $v(\alpha) = \beta$ and $v(\beta) = \alpha$, with $\beta \neq \alpha$, then the definition of v implies that we may swap the labels if necessary so that α is a codimension 1 face of β , and we let $[\alpha, \beta]$ be one of the pairs in V . To finish the proof of Theorem 6.1, we need to show that V is gradient.

Lemma 6.3. *The discrete vector field V is gradient; that is, it does not admit any closed V -walks.*

Proof. Suppose for the sake of contradiction that $[\alpha_1, \beta_1], [\alpha_2, \beta_2], \dots, [\alpha_r, \beta_r]$ is a closed V -walk. We define a function

$$\text{key}: \text{poset}(n, w) \rightarrow \bigoplus_{i=1}^{\infty} \mathbb{Z}$$

and show that if we compare the various $\text{key}(\alpha_i)$, they are in strictly decreasing lexicographical order. This gives a contradiction with the assumption that the V -walk is closed with $\alpha_r = \alpha_1$.

The key function is defined as follows. Given the symbol α of a cell in $\text{cell}(n, w)$, we consider each block, and we set entry $2k - 1$ of $\text{key}(\alpha)$ to be the first element of the k th block, unless that block is a follower, in which case we set that entry to be zero; in either case, we set entry $2k$ of $\text{key}(\alpha)$ to be the number of elements of the k th block. Past twice the number of blocks, all entries of $\text{key}(\alpha)$ are zero. The lexicographical order on $\bigoplus_{i=1}^{\infty} \mathbb{Z}$ is defined as follows: to compare two elements, we find the first entry where they differ, and we order the elements by their values in \mathbb{Z} at that entry.

We claim that for any i , we have $\text{key}(\alpha_{i+1}) < \text{key}(\alpha_i)$. Let k be the block where α_i merges to make β_i ; that is, α_i is match-up at k and β_i is match-down at k . Some block k' of β_i is split to form α_{i+1} , and there are three cases: it is the same block $k' = k$, it is an earlier block $k' < k$, or it is a later block $k' > k$.

Suppose $k' = k$. We know that block k of α_i is the longest subblock of block k of β_i that begins with the largest element of that block, so comparing entries $2k - 1$ and $2k$ of $\text{key}(\alpha_i)$ and $\text{key}(\alpha_{i+1})$, we find $\text{key}(\alpha_{i+1}) < \text{key}(\alpha_i)$ in this case.

Suppose $k' < k$. Because β_i is $(k - 1)$ -crit, the block k' that is split is either top-heavy or a follower, and block k' of α_{i+1} is a subblock of block k' of β_i . In the top-heavy case, comparing at entries $2k' - 1$ and $2k'$ gives $\text{key}(\alpha_{i+1}) < \text{key}(\beta_i)$, and because β_i and α_i agree past block k' , this implies $\text{key}(\alpha_{i+1}) < \text{key}(\alpha_i)$. In the follower case, block k' of α_{i+1} remains a follower, so comparing at entry $2k' - 1$ gives $\text{key}(\alpha_{i+1}) < \text{key}(\beta_i)$ and so $\text{key}(\alpha_{i+1}) < \text{key}(\alpha_i)$.

Suppose $k' > k$. Then block k of α_{i+1} is the same as block k of β_i , which has a smaller first element than block k of α_i (which is a leader and not a follower). Thus, comparing at entry $2k - 1$ gives $\text{key}(\alpha_{i+1}) < \text{key}(\alpha_i)$.

Thus, in all three cases the sequence $\text{key}(\alpha_i)$ is strictly decreasing and so cannot be a cycle, contradicting the existence of a closed V -walk, and so V is gradient. \square

Together, Lemma 6.2 and Lemma 6.3 imply Theorem 6.1.

Proof of Theorem 6.1. Lemma 6.2 shows that the discrete vector field V specified by the function v is well-defined: each cell can be in at most one pair in V . The construction of v automatically implies that the critical cells of V are those that are k -crit for all k , because those are the only cells that are fixed points of v . Lemma 6.3 shows that the discrete vector field V is gradient. \square

7. ASYMPTOTIC UPPER BOUNDS

In order to finish the proof of Theorem 1.1, we need to prove an asymptotic upper bound on the number of critical cells of each dimension. To do this, we group the critical cells of each dimension j into finitely many groupings and prove that each grouping satisfies the asymptotic bound. The groupings are called *skylines*. Roughly, the skyline retains the information about which blocks form leader/follower pairs and about the sequence of sizes of blocks, but forgets the numbers (corresponding to labels of disks) and all the blocks of size 1 that are neither leaders nor followers. Given the symbol of a critical cell in $\text{cell}(n, w)$, we refer to each leader/follower pair as a 2-block *barrier*. We find the *skyline* of that cell by the following process: we delete all the blocks that have just one element and are neither leaders nor followers (along with a bar adjacent to each), we replace the first element of each leader block by 1, and we replace all the other numbers in the symbol by 0.

The resulting skyline is a kind of symbol in which all of the numbers are 0 or 1. If the original cell was j -dimensional, then j is the number of zeros and ones in the skyline minus the number of blocks in the skyline, much as in the original cell. Any block with only one element is part of a barrier, so there are only finitely many different skylines for each j , independent of n . For each skyline S , we let $b(S)$ (“barriers”) denote the number of barriers, equal to the number of ones in S , and we let $z(S)$ (“zeros”) denote the number of zeros in S . In preparation for proving

Theorem 1.1, the following lemma implies an upper bound on the number of critical cells with a given skyline.

Lemma 7.1. *For every skyline S , there is an injective function code_S from the set of critical cells with skyline S into the set $[n]^{z(S)} \times [b(S) + 1]^n$.*

Proof. The function code_S is defined as follows. Given a critical cell α with skyline S , we can map α to an element of $[n]^{z(S)}$ by recording the original number in α corresponding to each zero in S , in the order these numbers appear in α . For the second coordinate, we divide the symbol of α into $b(S) + 1$ intervals: all the blocks up through the first barrier, all the blocks after the first barrier and up through the second barrier, and so on, with the last interval being all the blocks after the last barrier. Then we can map α to an element of $[b(S) + 1]^n$ by recording, for each number in α , which of the $b(S) + 1$ intervals it appears in.

To show that the function code_S is injective, we show how to recover α from $\text{code}_S(\alpha)$. The $[n]^{z(S)}$ coordinate specifies the original number for each 0 in S , so what remains is to find the original number for each 1 in S and to figure out where to insert the remaining numbers as one-element blocks. We can recover the original number for each 1 in S by finding which of the $b(S) + 1$ intervals ends with that barrier, selecting all the numbers in that interval, and taking the greatest of those numbers—the preceding blocks in the interval are top-heavy with initial elements in increasing order, and the 1 corresponds to the initial element of a leader block. Then, for all the numbers that do not correspond to zeros or ones in S , we find which of the $b(S) + 1$ intervals each number belongs to, and insert it as a one-element block into that interval in such a way that the initial elements of all the blocks in that interval (excluding the follower block at the end) are in increasing order. Because we can use this process to recover α from $\text{code}_S(\alpha)$, the function code_S is injective. \square

Putting these bounds together for all finitely many skylines, we can finish the proof of Theorem 1.1.

Proof of Theorem 1.1. The statements about the gas regime and the solid regime have already been addressed, and in Section 4 we have shown that if $j = q(w - 1) + r$ with $q \geq 1$ and $0 \leq r < w - 1$, then we have

$$\beta_j[\mathcal{C}(n, w)] = \Omega((q + 1)^n n^{qw + 2r}).$$

Thus, what remains is to prove that in this case we also have

$$\beta_j[\mathcal{C}(n, w)] = O((q + 1)^n n^{qw + 2r}).$$

Lemma 7.1 implies that for any skyline S , the number of critical cells with that skyline is at most $(b(S) + 1)^n n^{z(S)}$. Because the Betti number β_j is bounded by the number of critical cells of dimension j , and because there are finitely many skylines for each j , it then suffices to prove that for any skyline S corresponding to j -dimensional cells, we have

$$(b(S) + 1)^n n^{z(S)} = O((q + 1)^n n^{qw + 2r}).$$

Thinking of each block of size k as contributing a value of $k - 1$ to j , we observe that each 2-block barrier in S contributes a combined value of at least $w - 1$ to j . Thus we have $b(S) \leq q$. In the case where $b(S) < q$, we certainly have

$(b(S) + 1)^n n^{z(S)} = O((q + 1)^n n^{qw+2r})$, because the factor that is exponential in n overwhelms the factor that is polynomial in n .

Thus, it suffices to prove that if $b(S) = q$, then $z(S) \leq qw + 2r$. The number of zeros in S is j plus the number of blocks in S without a 1. Because $j = q(w - 1) + r$, this means that it suffices to show that the number of blocks in S without a 1 is at most $q + r$. Each barrier contains exactly one block without a 1, so there are q such blocks. The other blocks without a 1 are not part of barriers, so they have size at least 2. Each of these contributes at least 1 to j , and the barriers together contribute at least $q(w - 1)$ to j , so there are at most r of these non-barrier blocks in S . Thus, together the number of blocks in S without a 1 is at most $q + r$, so we have $z(S) \leq qw + 2r$, and thus

$$\beta_j[\mathcal{C}(n, w)] \leq \#(\text{crit cells of dim } j) = O((q + 1)^n n^{qw+2r}),$$

completing the proof of Theorem 1.1. \square

8. COMMENTS

- (1) In principle, one could compute homology of $\mathcal{C}(n, w)$ exactly. For example, $\mathcal{C}(n, w)$ is homotopy equivalent to $U(n, w)$, which in turn is homeomorphic to the complement of a certain real subspace arrangement. We define $\mathcal{A}_{n, w}$ to be the collection of $\binom{n}{2}$ subspaces of codimension 2

$$\{(x_1, y_1, \dots, x_n, y_n) \in \mathbb{R}^{2n} \mid x_k = x_\ell \text{ and } y_k = y_\ell\}$$

together with the $\binom{n}{w+1}$ subspaces of codimension $w + 1$

$$\{(x_1, y_1, \dots, x_n, y_n) \mid x_{i_1} = x_{i_2} = \dots = x_{i_{w+1}}\}.$$

Since $\mathcal{C}(n, w)$ is homotopy equivalent to the complement of this subspace arrangement, the homology is determined by the intersection lattice of the arrangement [17]. One might apply essentially combinatorial formulas to derive a formula for $\beta_j[\mathcal{C}(n, w)]$. Such an exact formula might be nice to have, even if in a complicated or recursive form.

- (2) The definitions of homological solid, liquid, and gas make sense even for 0th homology, especially for bounded regions. Determining the threshold for the solid-liquid phase transition for 0th homology passing from trivial to nontrivial is equivalent to the well-studied “sphere packing” problem.

There is another transition for 0th homology, the homological liquid-gas phase transition, where the configuration space becomes connected. This seems to be much less well studied, but the threshold for connectivity is a natural and important question for a number of reasons. For example, in his survey article [14], Diaconis writes about it in the context of ergodicity of Markov chains, a requirement for being able to effectively sample a configuration space by making small random movements of disks. See also [19] for discussion of the connectivity threshold.

- (3) We show in Section 4 that certain toruses generate a positive fraction of the homology, but on the other hand we also know that even if one considers all of the toruses that one can make in similar ways, they do not seem to generate all of the homology. Consider the example $n = 3$, $w = 2$, $j = 1$,

illustrated in Figure 4. We know that $\beta_1 = 7$, but only 6 cycles are accounted for by rotating a pair of disks around each other, and having the third disk on either one side or the other. The “outside circle” in the figure is visibly not in the span of the six smaller cycles.

- (4) Discrete Morse theory has been studied on the Salvetti complex before. For a more geometric approach to discrete gradients on $\text{cell}(n)$, see [25], [23], and [22]. We do not know whether the techniques from these papers can improve the upper bounds on $\beta_j[\mathcal{C}(n, w)]$, or perhaps even produce perfect discrete Morse functions or minimal CW complexes for $\mathcal{C}(n, w)$.
- (5) A related family of spaces is the “no k -equal space” $M_{n,k}^{\mathbb{R}}$ studied by Björner et al. in [7, 8, 9]. In particular, there is a natural map $\mathcal{C}(n, w) \rightarrow M_{n,w+1}^{\mathbb{R}}$ by projecting onto the x -coordinates. We do not know much about the induced map on homology in general. We point here out a coincidence we notice in the data that we do not currently have a good explanation for.

Comparing Table 1 in our appendix with the first table in the appendix of [9], it seems possible that $H_1(\mathcal{C}(n, 2))$ is isomorphic to $H_1(M_{n+1,3}^{\mathbb{R}})$ —at least the Betti numbers are equal for $n \leq 8$.

We emphasize that $\mathcal{C}(n, 2)$ is a configuration space of n points, and $M_{n+1,3}^{\mathbb{R}}$ is a configuration space of $n + 1$ points, so we do not even have an obvious candidate of map to induce such an isomorphism. Supposing that there were such a map, we might wonder if it also induces an isomorphism on π_1 but apparently not, as follows.

We showed that $\mathcal{C}(n, 2)$ is a $K(\pi, 1)$ in Section 3.3. The question of whether $M_{n,3}^{\mathbb{R}}$ is a $K(\pi, 1)$ was asked by Björner [6] and answered affirmatively by Khovanov [20]. Khovanov describes this as a real analogue of the fact that $M_{n,2}^{\mathbb{C}}$ (the configuration space of points in the plane) is a $K(\pi, 1)$. Since both spaces are $K(\pi, 1)$ ’s, if they had isomorphic fundamental groups then they would be homotopy equivalent. But the Betti number tables rule out the higher homology groups $j \geq 2$ being isomorphic.

APPENDIX BY ULRICH BAUER AND KYLE PARSONS

We computed the Betti numbers $\beta_j[\text{cell}(n, w)]$ for $n \leq 8$, for homology with $\mathbb{Z}/2$ coefficients, using the software PHAT [5]. The results of the computations appear in Table 1. For a point of reference, we note that $\text{cell}(8)$ has over 5 million cells.

REFERENCES

- [1] Hannah Alpert. Restricting cohomology classes to disk and segment configuration spaces. *Topology Appl.*, 230:51–76, 2017.
- [2] Vladimir I. Arnold. *The cohomology ring of the colored braid group*, pages 183–186. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.
- [3] Yuliy Baryshnikov, Peter Bubenik, and Matthew Kahle. Min-type Morse theory for configuration spaces of hard spheres. *Int. Math. Res. Not. IMRN*, 9:2577–2592, 2014.
- [4] Ulrich Bauer, Herbert Edelsbrunner, Grzegorz Jablonski, and Marian Mrozek. Čech–Delaunay gradient flow and homology inference for self-maps. arXiv:1709.04068, 2017.

TABLE 1. Betti numbers of $\mathcal{C}(n, w)$ for small n and w . Bold font indicates that homology is in the “liquid regime”.

n	w	β_0	β_1	β_2	β_3	β_4	β_5	β_6	β_7
2	1	2	0	0	0	0	0	0	0
2	2	1	1	0	0	0	0	0	0
3	1	6	0	0	0	0	0	0	0
3	2	1	7	0	0	0	0	0	0
3	3	1	3	2	0	0	0	0	0
4	1	24	0	0	0	0	0	0	0
4	2	1	31	6	0	0	0	0	0
4	3	1	6	29	0	0	0	0	0
4	4	1	6	11	6	0	0	0	0
5	1	120	0	0	0	0	0	0	0
5	2	1	111	110	0	0	0	0	0
5	3	1	10	169	40	0	0	0	0
5	4	1	10	35	146	0	0	0	0
5	5	1	10	35	50	24	0	0	0
6	1	720	0	0	0	0	0	0	0
6	2	1	351	1160	90	0	0	0	0
6	3	1	15	714	780	80	0	0	0
6	4	1	15	85	1066	275	0	0	0
6	5	1	15	85	225	875	0	0	0
6	6	1	15	85	225	274	120	0	0
7	1	5040	0	0	0	0	0	0	0
7	2	1	1023	9212	3150	0	0	0	0
7	3	1	21	2568	6468	3920	0	0	0
7	4	1	21	175	5272	5957	840	0	0
7	5	1	21	175	735	7678	2058	0	0
7	6	1	21	175	735	1624	6084	0	0
7	7	1	21	175	735	1624	1764	720	0
8	1	40320	0	0	0	0	0	0	0
8	2	1	2815	61194	60900	2520	0	0	0
8	3	1	28	8385	37464	76146	6720	0	0
8	4	1	28	322	21477	54910	36239	2520	0
8	5	1	28	322	1960	43728	49959	7896	0
8	6	1	28	322	1960	6769	62525	17101	0
8	7	1	28	322	1960	6769	13132	48348	0
8	8	1	28	322	1960	6769	13132	13068	5040

- [5] Ulrich Bauer, Michael Kerber, Jan Reininghaus, and Hubert Wagner. PHAT – Persistent Homology Algorithms Toolbox. *Journal of Symbolic Computation*, 78:76–90, January 2017. Software available at <https://bitbucket.org/phant-code/phant>.
- [6] Anders Björner. Subspace arrangements. In *First European Congress of Mathematics, Vol. I (Paris, 1992)*, volume 119 of *Progr. Math.*, pages 321–370. Birkhäuser, Basel, 1994.
- [7] Anders Björner and László Lovász. Linear decision trees, subspace arrangements and Möbius functions. *J. Amer. Math. Soc.*, 7(3):677–706, 1994.

- [8] Anders Björner, László Lovász, and A Yao. Linear decision trees: volume estimates and topological bounds. In *Proc. 24th ACM Symp. on Theory of Computing*, volume 170, page 177. Citeseer, 1992.
- [9] Anders Björner and Volkmar Welker. The homology of “ k -equal” manifolds and related partition lattices. *Adv. Math.*, 110(2):277–313, 1995.
- [10] Pavle V. M. Blagojević and Günter M. Ziegler. Convex equipartitions via equivariant obstruction theory. *Israel J. Math.*, 200(1):49–77, 2014.
- [11] Gunnar Carlsson, Jackson Gorham, Matthew Kahle, and Jeremy Mason. Computational topology for configuration spaces of hard disks. *Phys. Rev. E*, 85:011303, Jan 2012.
- [12] Michael W. Davis. The geometry and topology of Coxeter groups. In *Introduction to modern mathematics*, volume 33 of *Adv. Lect. Math. (ALM)*, pages 129–142. Int. Press, Somerville, MA, 2015.
- [13] Kenneth Deeley. Configuration spaces of thick particles on a metric graph. *Algebr. Geom. Topol.*, 11(4):1861–1892, 2011.
- [14] Persi Diaconis. The Markov chain Monte Carlo revolution. *Bull. Amer. Math. Soc. (N.S.)*, 46(2):179–205, 2009.
- [15] Steve Ferry, Konstantin Mischaikow, and Vidit Nanda. Reconstructing functions from random samples. *J. Comput. Dyn.*, 1(2):233–248, 2014.
- [16] Robin Forman. A user’s guide to discrete Morse theory. *Sém. Lothar. Combin.*, 48:Art. B48c, 35, 2002.
- [17] Mark Goresky and Robert MacPherson. *Stratified Morse theory*, volume 14 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1988.
- [18] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.
- [19] Matthew Kahle. Sparse locally-jammed disk packings. *Ann. Comb.*, 16(4):773–780, 2012.
- [20] Mikhail Khovanov. Real $K(\pi, 1)$ arrangements from finite root systems. *Math. Res. Lett.*, 3(2):261–274, 1996.
- [21] Rob Kusner, Wöden Kusner, Jeffrey C. Lagarias, and Senya Shlosman. Configuration spaces of equal spheres touching a given sphere: The twelve spheres problem. In *New Trends in Intuitive Geometry, Bolyai Society Mathematical Studies*, pages 219–277. Springer-Verlag GMBH, Germany, 2018.
- [22] Davide Lofano and Giovanni Paolini. Euclidean matchings and minimality of hyperplane arrangements. arXiv:1809.02476, 2018.
- [23] Francesca Mori and Mario Salvetti. (Discrete) Morse theory on configuration spaces. *Math. Res. Lett.*, 18(1):39–57, 2011.
- [24] M. Salvetti. Topology of the complement of real hyperplanes in \mathbf{C}^N . *Invent. Math.*, 88(3):603–618, 1987.
- [25] Mario Salvetti and Simona Settepanella. Combinatorial Morse theory and minimality of hyperplane arrangements. *Geom. Topol.*, 11:1733–1766, 2007.
- [26] Dev P. Sinha. The (non-equivariant) homology of the little disks operad. In *OPERADS 2009*, volume 26 of *Sémin. Congr.*, pages 253–279. Soc. Math. France, Paris, 2013.
- [27] Richard P. Stanley. *Enumerative combinatorics. Volume 1*, volume 49 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, second edition, 2012.