CONFIGURATION SPACES OF DISKS IN AN INFINITE STRIP

HANNAH ALPERT, MATTHEW KAHLE, AND ROBERT MACPHERSON

Abstract. We study the topology of the configuration space $\mathcal{C}(n,w)$ 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 inclusion i into the configuration space of n points in the plane $\mathcal{C}(n,\mathbb{R}^2)$ induces an isomorphism on homology $i_*: H_j[\mathcal{C}(n,w)] \to H_j[\mathcal{C}(n,\mathbb{R}^2)]$. The Betti numbers of $\mathcal{C}(n,\mathbb{R}^2)$ were computed by Arnold [1], and so as a corollary of the isomorphism, if w and j are fixed then $\beta_j[\mathcal{C}(n,w)]$ is a polynomial of degree $2j$ in n.

On the other hand, we show that w and j are fixed and $2 \leq w \leq j+1$, then $\beta_j[\mathcal{C}(n,w)]$ grows exponentially fast 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 "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 are interested here in the configuration spaces $\mathcal{C}(n,w)$ of n hard disks, of unit diameter, in an infinite strip of width w.

More precisely, for non-negative integers n, w we define

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

One could define these configuration spaces with strict inequalities instead, giving an open manifold instead of a closed semi-algebraic set. Up to homotopy equivalence, it makes no difference. Similarly, if w is not an integer then the configuration space $\mathcal{C}(n,w)$ deformation retracts onto $\mathcal{C}(n,\lfloor w \rfloor)$, so there is no loss of generality in assuming that w is an integer.

Our main result gives sharp estimates for the Betti numbers, as the number of disks tends to infinity. We use the notation $f = \Theta(g)$ to indicate that there exist positive constants c_1, c_2 such that

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

M.K. thanks IAS for hosting him as a member in 2010–11. He gratefully acknowledges the support of a Sloan Research Fellowship and an NSF-CAREER award.

All three authors thank ICERM for hosting them during the special thematic semester on Topology in Motion in Fall 2016.

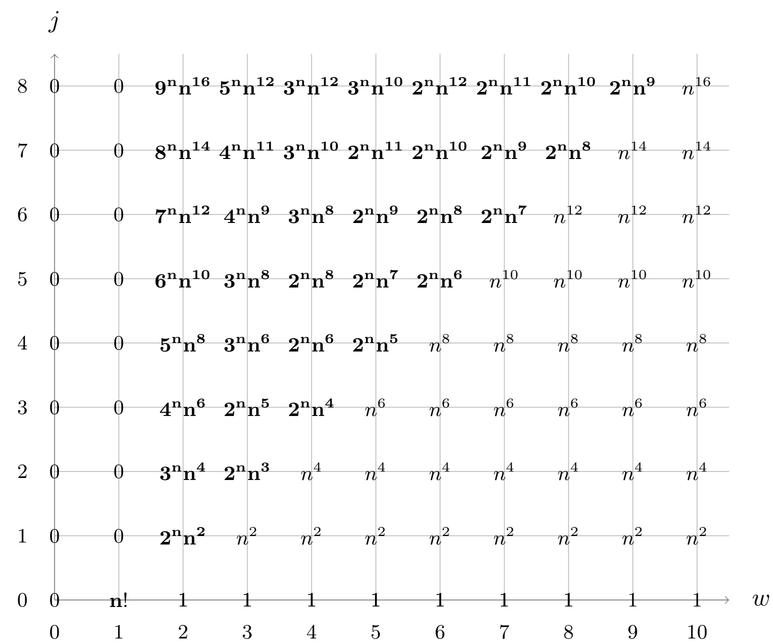

FIGURE 1. Theorem 1.1 describes the rate of growth of $\beta_j[\mathcal{C}(n,w)]$, for fixed j and w, as $n \to \infty$. The results are up to a constant factor, e.g. $\beta_8[\mathcal{C}(n,3)] = \Theta\left(5^n n^{12}\right)$.

for all sufficiently large n. Whenever there is asymptotic notation in the following, the implied constants depend on j and w but not on n.

Theorem 1.1.

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

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

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

$$\beta_j\left[\mathcal{C}(n,w)\right] = \Theta\left(n^{2j}\right).$$

(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*

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

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

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

In principle, one could compute homology of $\mathcal{C}(n,w)$ exactly. For example, we believe that $\mathcal{C}(n,w)$ is homotopy equivalent to the complement of a certain subspace arrangement that we do not describe here. Then, one could use essentially combinatorial formulas [14] to derive a recursive formula for $\beta_j[\mathcal{C}(n,w)]$. We do not take such an approach here, and we believe that the asymptotic approach we take may be in some ways more meaningful.

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 [12], and the review of the physics literature in Carlsson et al. in [9].

The topology of configuration spaces of particles with thickness has been studied earlier, for example in [2], [11], and [18], but not much seems to be known. Some of this past work is also inspired in part by applications to engineering, particularly motion planning for robots.

Inspired by the statement of Theorem 1.1, we introduce the notion of "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 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, as in an ideal gas. Arnold [1] showed that the Poincaré polynomial of $\mathcal{C}(n, \mathbb{R}^2)$ is given by

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

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

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

For a self-contained and readable overview of the homology and cohomology of $\mathcal{C}(n, \mathbb{R}^2)$, see Sinha [23].

One can use a standard recursive formula for Stirling numbers to write $\left[\begin{smallmatrix} n \\ n-j \end{smallmatrix}\right]$ as a polynomial in n of degree $2j$. For example, 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 were 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.

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

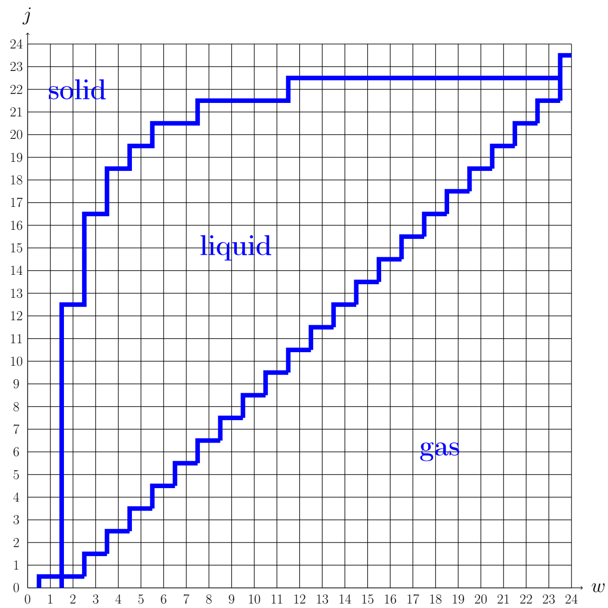

FIGURE 2. Theorem 1.2 describes the shapes of the homological solid, liquid, and gas regimes for every n. We illustrate the case $n = 24$.

apply 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 only asymptotically. The following describes the shapes of the regimes for every n. We note 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 following hold for every* $n \geq 2$.

(1) *If* $w \geq 2$ *and* $0 \leq j \leq w - 2$, *then the inclusion map* $i : \mathcal{C}(n, w) \to \mathcal{C}(n, \mathbb{R}^2)$ *induces an isomorphism on homology*
$$i_* : \ H_j[\mathcal{C}(n, w)] \to H_j[\mathcal{C}(n, \mathbb{R}^2)].$$
If $w \geq n$, *then* $\mathcal{C}(n, w)$ *is homotopy equivalent to* $\mathcal{C}(n, \mathbb{R}^2)$.

(2) *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* $i : \mathcal{C}(n, w) \to \mathcal{C}(n, \mathbb{R}^2)$ *does not induce an isomorphism on homology*
$$i_* : \ H_j[\mathcal{C}(n, w)] \to H_j[\mathcal{C}(n, \mathbb{R}^2)].$$

(3) *If either* $w = 0$, *or* $w \geq 1$ *and* $j \geq n - \lceil n/w \rceil + 1$, *then*
$$H_j[\mathcal{C}(n, w)] = 0.$$

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, and then the number of d-cells is an upper bound on the Betti number β_d.

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

2.1. **poset**(n) **and cell**(n)**.** 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.* The Salvetti complex and related constructions have appeared implicitly or explicitly many times—see Section 3 of [8] for a brief review of the literature. The complex was apparently first described explicitly by Salvetti in [21].

Definition 2.1. The poset $\text{poset}(n)$ has as it underlying set $\mathcal{A}(n)$, which we define as follows. Elements of $\mathcal{A}(n)$ we call "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.

We call a part of the permutation between two bars a *block.* 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, $(1 \mid 3 \mid 2)$, $(31 \mid 2)$, and (321) 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

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

it means that each c_i is a piece of the permutation, separated from the rest by bars. Forgetting the order of permutation elements within a block, 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 [8], usually called the Salvetti complex. It was shown in [21] that

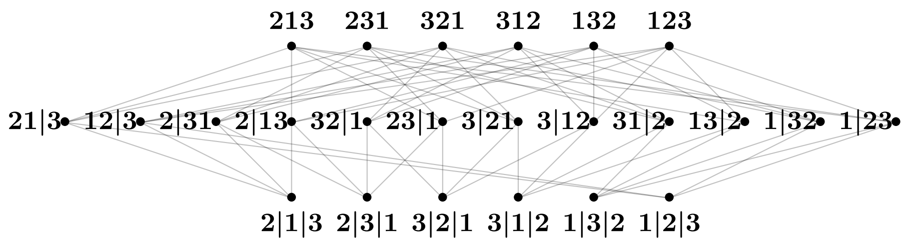

FIGURE 3. The Hasse diagram of poset(3). This is the face poset of the Salvetti complex for the configuration space of 3 points in the plane.

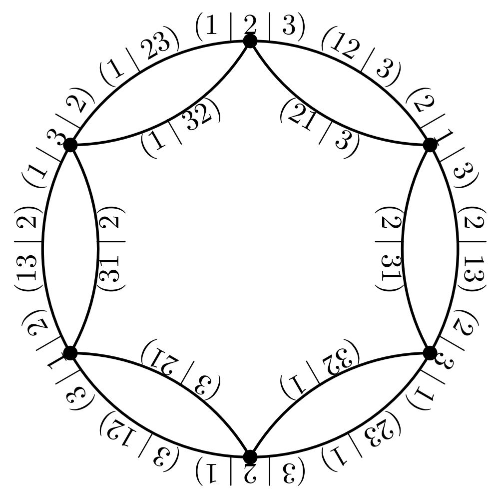

FIGURE 4. The cell complex cell(3, 2).

cell(n) is homotopy equivalent to the configuration space of points in the plane $\mathcal{C}(n, \mathbb{R}^2)$.

The cell complex 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 $\alpha = (c_1 \mid c_2 \mid \cdots \mid c_m)$ with m blocks, then the cell has dimension $j = n - m$.

We will be mostly concerned with certain subcomplexes of cell(n), described as follows.

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

The cell complex cell(3, 2) is illustrated in Figure 4.

We define the closely related "configuration space of vertical line segments" as follows.

Definition 2.3. Assume that $0 < \epsilon < 1$. The configuration space of vertical line segments of unit length in the strip of width $w + \epsilon$ is defined by

$$\mathcal{C}_I(n, w+\epsilon) = \{(x_1, y_1, \ldots, x_n, y_n) \mid 1/2 < y_i < w - 1/2 + \epsilon \text{ for every } 1 \le i \le n, \\ \text{and } \; x_i = x_j \implies |y_i - y_j| > 1 \text{ for every } 1 \le i < j \le n.\}$$

We prove in Section 3 that

$$\mathcal{C}_I(n, w+\epsilon) \simeq \mathcal{C}(n, w),$$

as one step in the proof of the main homotopy equivalence $\mathrm{cell}(n, w) \simeq \mathcal{C}(n, w)$. We also use $\mathcal{C}_I(n, w+\epsilon)$ in the proof of lower bounds in Section 4. It is convenient, for example, that $\mathcal{C}_I(n, w+\epsilon)$ is an open subset of $\mathbb{R}^{2n}$ so an open manifold.

3. Homotopy equivalence

In this section we prove the main homotopy equivalence $\mathcal{C}(n, w) \simeq \mathrm{cell}(n, w)$, and list a few of the immediate consequences.

Theorem 3.1. *For every n, w, we have $\mathcal{C}(n, w) \simeq cell(n, w)$.*

3.1. **Proof of Theorem 3.1.** First we show the main step, that $\mathcal{C}_I(n, w+\epsilon) \simeq \mathrm{cell}(n, w)$.

Lemma 3.2. *For every $n, w \ge 1$ and $0 < \epsilon < 1$, we have the homotopy equivalence*

$$\mathcal{C}_I(n, w+\epsilon) \simeq cell(n, w).$$

Our strategy will be to use the nerve theorem. We will consider the nerve of the following open cover.

Definition 3.3. Given a symbol $\alpha \in \mathrm{poset}(n, w)$, we define an open set U_α as follows. Write α in block notation $\alpha = (c_1 \mid c_2 \mid \cdots \mid c_m)$. We define

$$U_\alpha = \{(x_1, y_1, x_2, y_2, \ldots, x_n, y_n) \in \mathbb{R}^{2n} : \text{ the following conditions are met}\}.$$

- Whenever σ_k and σ_l are in different blocks and $k < l$, we have $x_k < x_l$.
- Whenever σ_k and σ_l are in the same block and $k < l$, we have $y_k > y_l$.
- If σ_k and σ_l are in the same block, and $\sigma_{k'}$ and $\sigma_{l'}$ are in different blocks, then

  $$|x_k - x_l| < |x_{k'} - x_{l'}|.$$

  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-ccordinate.

Every U_α is convex, hence all the nonempty intersections of these sets is contractible. In the following, we check that $\{U_\alpha\}$ form an open cover of $\mathcal{C}_I(n, w)$, and that the intersection

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

is nonempty if and only if the symbols

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

form a chain in $\mathrm{poset}(n, w)$. Then the nerve theorem gives the homotopy equivalence.

Proof of Lemma 3.2. Given a point $p = \mathcal{C}_I(n,w)$, we first describe an algorithm for finding $\mathcal{A}_p$, the subset of symbols $\alpha \in \mathcal{C}(n,w)$ such that $p \in U_\alpha$.

As a preliminary, we define the poset of ordered partitions part(n). An element of 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 $[n]$. This is very similar to an element of $\mathcal{A}(n)$, where we forget the order within a block (replacing an ordered subset of $[n]$ by a subset).

The partial order on 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 by their union at the same place in the order. We remark that part(n) is somewhat similar to poset $\mathcal{A}(n)$, but in part(n) we forget the order of the elements within a block.

Now let a point

$$p = (x_1, y_1, x_2, y_2, \dots, x_n, y_n) \in \mathcal{C}_I(n,w)$$

be given.

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

For every real number $\rho \geq 0$, there is a unique partition

$$\pi(\rho) = (S_1(\rho) \mid S_2(\rho) \mid \cdots \mid S_{m(\rho)})$$

such that if x_i and x_j are consecutive x values and $i \in S_k(\rho)$, $j \in S_\ell(\rho)$, then

- if $x_j - x_i \leq \rho$ then $k = \ell$, i.e. i and j lie in the same set of the partition $\pi(\rho)$, and
- if $x_j - x_i > \rho$ then $k + 1 = \ell$, i.e. i and j lie in consecutive sets of the partition $\pi(\rho)$.

Let $K_i(\rho)$ be the closed interval $[x_i - \rho/2, x_i + \rho/2]$ of length ρ centered on x_i and let $K(\rho)$ be the union of all the intervals $K_i(\rho)$. The set $K(\rho)$ will itself be a union of intervals of various lengths, namely its connected components. We cluster the integers $[n]$ according to which connected component of K they lie in. The set $S_j(\rho)$ is the subset of $[n]$ whose x coordinates are in the j-th connected component of $K(\rho)$, counting from the lower end.

When $\rho = 0$, i and j lie in the same cluster only when $x_i = x_j$. When ρ is very large, there is only one cluster $S_1 = [n]$. In general, as ρ increases from 0, 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 $\pi_1, \pi_2, \dots$ This sequence is clearly a chain in the poset of ordered partitions, part(n).

Step 2 lifts a subsequence of the chain $\pi_1, \pi_2, \dots$ in the poset of ordered partitions part(n) to a chain in the poset of symbols poset(n). This step uses the y coordinates of p but not the x coordinates.

Given a partition $\pi_i = (S_1, S_2, \dots)$, for each set S_j, order the elements of S_j in such a way that if k and ℓ are elements of S_j that are consecutive in the ordering, then $y_k + 1 < y_\ell$. If for some S_j in the ordered partition π_i this can't be done, then discard π_i and exclude it from further consideration.

Since a symbol in $\mathcal{A}_n$ is just an ordered partition $(S_1, S_2, \ldots)$ of $[n]$ together with an ordering of the elements of each set S_j, we now have a list of symbols in $\mathcal{A}_n$. Using the fact that $\pi_1, \pi_2, \ldots$ is a chain in part(n), it is easy to see that our list of symbols is a chain in poset(n). It actually lies in the sub-poset poset(n, w) because the y_i values are restricted to lie between $1/2$ and $w + \epsilon - 1/2$.

Step 3 extracts a subchain using the x_i values. From the list of symbols α coming from Step 2, throw out those that don't satisfy $p \in \mathcal{A}_\alpha$. The result will still be a chain, since a subsequence of a chain is a chain.

From the definition of U_α it is not hard to check that the resulting chain is $\mathcal{A}_p$.

The algorithm always produces at least A_1, so the set $\mathcal{A}_p$ is always nonempty, and therefore the set U_α form an open cover of $\mathcal{C}_I(n, w)$.

Each U_α is convex for every α, so every U_α is contractible and since the intersection of convex sets is convex, every nonempty intersection is contractible.

Finally, we check that an intersection

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

is nonempty if and only if the symbols

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

form a chain in poset(n, w).

First of all, if the intersection is nonempty then the algorithm produces a chain of partitions that include all of $\alpha_1, \ldots, \alpha_k$, and a subposet of a chain is a chain.

Now suppose we have a chain $\alpha_1 < \alpha_2 < \cdots < \alpha_m$ in poset(n, w). We produce a point p such that $p \in U_{\alpha_i}$ for every i. We may assume without loss of generality that the chain is maximal.

None of the blocks in any symbol have width more than w. So we can choose y-coordinates based on the symbol α_m as follows. If σ_k and σ_l are in the same block and $k < l$ then choose $y_k > y_l + 1$, then there is enough room vertically in $\mathcal{C}_I(n, w)$ to do this.

For x-coordinates, we first look at the symbol α_1, which by maximality of the chain has all blocks of size one. Our first restriction is that we choose x-coordinates of p such that $\sigma_k < \sigma_l$ in α_1 then $x_k < x_l$.

Finally, we cluster x-coordinates so that as one ascends the chain, the blocks merge in the correct order. For example, let $0 < \lambda_1 < \lambda_2 < \cdots < \lambda_m$ be any increasing sequence of real numbers. If going from α_i to α_{i+1} merges blocks c_{k_i} and c_{k_i+1}, our only other restriction is simply to make sure that

$$\max\left\{|x_a - x_b| \;:\; \sigma_a \in c_{k_i},\, \sigma_b \in c_{k_{i+1}}\right\} = \lambda_i,$$

and

$$\max\{|x_a - x_b| \;:\; \sigma_a \text{ and } \sigma_b \text{ are in the same block in } c_{k_i}\} < \lambda_i.$$

We conclude that the nerve of the cover $\{U_\alpha\}$ is isomorphic to the order complex $\Delta\,(\text{poset}(n, w))$. So by the Nerve Theorem, $\mathcal{C}_I(n, w + \epsilon) \simeq \text{cell}(n, w)$.

□

Next, we give a homotopy equivalence between the closed configuration space of disks and the configuration space of intervals.

Definition 3.4. Let $\bar{\mathcal{C}}(n,w)$ denote the closed configuration space of disks.

We first comment that the closed configuration space of disks is homotopy equivalent to a slightly larger open configuration space.

Lemma 3.5. *For integers $n, w \geq 1$ and $0 < \epsilon < 1$, we have a homotopy equivalence between the configuration space of intervals and the configuration space of disks*

$$\mathcal{C}_I(n, w+\epsilon) \simeq \mathcal{C}(n,w).$$

Proof. For $x = (x_1, y_1, \ldots, x_n, y_n)$, set

$$R = \max\left\{ \max_{i,j} \sqrt{\frac{1-(y_i-y_j)^2}{(x_i-x_j)^2}}, 1 \right\}.$$

Here the maximum is taken over all pairs (i,j) such that $x_i \neq x_j$.

Now we dilate by a factor of R in the horizontal direction. In other words, we define

$$D(x_1, y_1, x_2, y_2, \ldots, x_n, y_n) = (Rx_1, y_1, Rx_2, y_2, \ldots, Rx_n, y_n).$$

This map D gives a deformation retraction onto the smaller configuration space

$$\begin{aligned}\mathcal{C}'(n, w+\epsilon) = \{(x_1, y_1, \ldots, x_n, y_n) \mid\ & 1/2 < y_i < w - 1/2 + \epsilon \text{ for every } 1 \leq i \leq n, \\ & (x_i - x_j)^2 + (y_i - y_j)^2 \geq 1 \text{ for every } 1 \leq i < j \leq n, \\ & \text{and } x_i = x_j \implies |y_i - y_j| > 1 \text{ for every } 1 \leq i < j \leq n.\}\end{aligned}$$

Next, define

$$\begin{aligned}\mathcal{C}''(n, w+\epsilon) = \{(x_1, y_1, \ldots, x_n, y_n) \mid\ & 0 \leq x_i \leq 2n \text{ for every } 1 \leq i \leq n, \\ & 1/2 < y_i < w - 1/2 + \epsilon \text{ for every } 1 \leq i \leq n, \\ & (x_i - x_j)^2 + (y_i - y_j)^2 \geq 1 \text{ for every } 1 \leq i < j \leq n, \\ & \text{and } \ x_i = x_j \implies |y_i - y_j| > 1 \text{ for every } 1 \leq i < j \leq n.\}\end{aligned}$$

It is clear that $\mathcal{C}'(n, w+\epsilon)$ deformation retracts onto $\mathcal{C}''(n, w+\epsilon)$. We note that $\mathcal{C}''$ is essentially the configuration space of n disks of radius $r = 1/2$ in a rectangle of dimensions $(2n+1) \times (w+\epsilon)$. The min-type Morse theory techniques in [2] give deformation retractions from both $\mathcal{C}''(n, w+\epsilon)$ and $\mathcal{C}(n,w)$ to the configuration space $\mathcal{C}'''(n,w)$ of n disks of radius $1/2$ in a rectangle of dimensions $(2n+1) \times w$ defined by

$$\begin{aligned}\mathcal{C}'''(n, w) = \{(x_1, y_1, \ldots, x_n, y_n) \mid\ & 0 \leq x_i \leq n \text{ for every } 1 \leq i \leq n, \\ & 1/2 \leq y_i \leq w - 1/2 \text{ for every } 1 \leq i \leq n, \\ & \text{and } (x_i - x_j)^2 + (y_i - y_j)^2 \geq 1 \text{ for every } 1 \leq i < j \leq n.\}\end{aligned}$$

This step is where we use the hypothesis that $\epsilon < 1$.

Putting it together, we have that $\mathcal{C}_I(n, w+\epsilon)$ deformation retracts to $\mathcal{C}'(n,w)$, $\mathcal{C}''(n,w)$, and $\mathcal{C}'''(n,w)$. We also have that $\mathcal{C}(n,w)$ deformation retracts to $\mathcal{C}'''(n,w)$, so we have a homotopy equivalence between $\mathcal{C}_I(n, w+\epsilon)$ and $\mathcal{C}(n,w)$, as desired. □

3.2. **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.

Proof of Part (1) of Theorem 1.1. We show that if $w \geq j+2$ and i is the inclusion map

$$i : \mathcal{C}(n,w) \hookrightarrow \mathcal{C}(n,\mathbb{R}^2),$$

then the induced map on homology

$$i_* : H_j(\mathcal{C}(n,w)) \to H_j(\mathcal{C}(n,\mathbb{R}^2))$$

is an isomorphism.

Note that every cell in cell(n) but not in 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 $(w-1)$-skeleton of cell(n,w) is the same as the $(w-1)$-skeleton of cell(n), which is homotopy equivalent to $\mathcal{C}(n,\mathbb{R}^2)$. The homology in degrees $\leq w-2$ only depend on the $(w-1)$-skeleton. □

Another consequence of the homotopy equivalence is that cell$(n,2)$ is an Eilenberg–Maclane space.

Theorem 3.6. *The cubical complex 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 [15]. The only thing to check is that the link of every vertex in cell$(n,2)$ is a "flag" simplicial complex. A complex is flag if it has no empty triangles or higher-dimensional empty simplices. A precise statement and complete proof of Gromov's criterion can be found in Appendix I.6 of Davis's book [10]. □

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 C is an open d-dimensional manifold, with submanifolds $Z_1, Z_2, \ldots, Z_k$ and $Z_1^*, Z_2^*, \ldots, Z_k^*$ 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 $\dim H_j(C,\mathbb{R}) \geq k$.

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

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

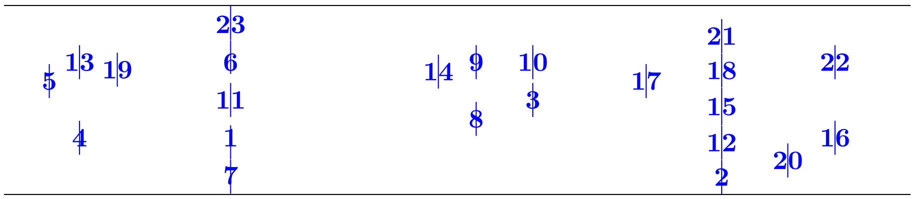

FIGURE 5. A point on the 34-dimensional cocycle Z^*_α, where $\alpha \in \mathcal{A}(23, 5)$ is the special symbol

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

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

□

Rather than work directly with $\mathcal{C}(n, w)$, it is convenient to work here in the configuration space of hard vertical segments, i.e. setting the ambient manifold $M = \mathcal{C}_I(n, w)$. We can make this substitution since in Section 3 it is shown that $\mathcal{C}_I(n, w)$ is homotopy equivalent to $\mathcal{C}(n, w)$.

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 $(c_1 \mid c_2 \mid \cdots \mid c_m)$ 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
(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}.

We emphasize that which symbols are special depends on n, j, w, but for the sake of simplicity in notation we assume that these parameters are always known.

Definition 4.3. Given a special symbol α, we define a closed submanifold Z^*_α in $\mathcal{C}_I(n, w)$ as follows.

(1) If σ_k and σ_l are in the same block and $k < l$, then $x_k = x_l$ and $y_k \geq y_l$.
(2) If σ_k and σ_l are in different blocks and $k < l$, and either σ_k or σ_l is in a block of width w, then $x_k < x_l$.

The only thing that we need to check is that every Z^*_α is closed in $\mathcal{C}_I(n, w)$. All the inequalities that define the submanifold are clearly closed, except possible the condition $x_i < x_j$. However this inequality is also actually closed. Indeed, since the assumption is that at least one of the two adjacent blocks is maximal, there is not room for another vertical interval to merge . So one could replace this inequality by $x_i \leq x_j$.

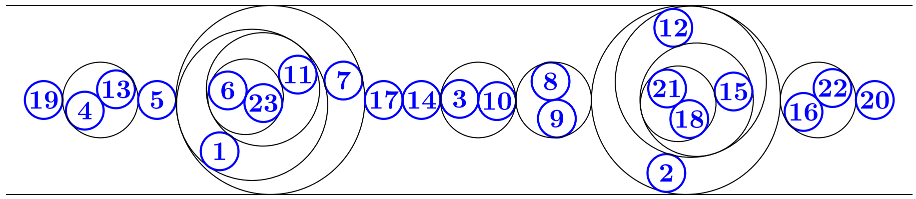

FIGURE 6. A point on the 12-dimensional cycle Z_α.

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

Proposition 4.4. *Given a special symbol α, there exists a cycle Z_α represented by an embedded j-dimensional torus, such that whenever $\alpha' \neq \alpha$ we have that $Z_\alpha \cap Z^*_{\alpha'} = \emptyset$, and such that Z_α intersects Z^*_α transversely in a point for every α.*

Proof. 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 describe a point $(x_1, y_1, \dots, x_n, y_n) \in \mathbb{R}^{2n}$, as follows.

Read the symbol $\alpha = (c_1|c_2|\dots|c_m)$ from left to right, one block at a time. Let $w(c_i)$ denote the width of c_i. Since α is a special symbol, blocks in α are all of width 1, 2, or w.

Let

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

This variable tells us how far to horizontally shift the center of the torus for the next block.

Let

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

This is a counter which tells us which angle we are on.

(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 $w(c_i) = 2$, i.e. $c_i = |\sigma_m \sigma_{m+1}|$, then set

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

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

(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 $k = 1, 2, \ldots, w$, let

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

and

$$(u_k, v_k) = (u_{k-1}, v_{k-1}) - \frac{1}{2}(\cos\theta_d, \sin\theta_d).$$

The point of (3) is to rotate the first element around the second, and rotate the first two elements around the third, and so on; see Figure 6. A block of width w contributes a $(w-1)$-dimensional torus.

This construction embeds a j-dimensional torus in the closedconfiguration space of disks $\mathcal{C}(n, w)$. To obtain a j-dimensional torus in the configuration space of intervals $\mathcal{C}_I(n, w+\epsilon)$, dilate by a factor of $1+\epsilon/2w$ and translate vertically by $w/2$.

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

Suppose that

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

Define an equivalence relation on $[n]$ by setting $k \sim l$ if $x_k = x_l$. By the definition of cycle Z_α , if $k \sim l$ then σ_k and σ_l are in the same block of α. By the definition of cocycle Z^*_α, if σ_k and σ_l are in the same block, then $k \sim l$. So then if $p \in Z_\alpha \cap Z^*_{\alpha'}$, if σ_k and σ_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 σ_k and σ_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 the cocycle $Z^*_{\alpha'}$ it corresponds to the element at the top of the column (i.e. has the largest y-coordinate).

The only point on cycle Z_α with all elements in every block vertically aligned, and the largest, first element of block on top has all elements of block in vertical order. The maximum element in the block is on top by assumption. Then since the elements are vertically aligned and in a 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 a disk of diameter k, then the kth element of the block must be immediately below the $(k-1)$st element.

We assumed that $p \in Z_\alpha \cap Z^*_{\alpha'}$, and have shown that then $\alpha = \alpha'$. On the other hand, it is easy to see from the definitions that there is a single point of intersection $Z_\alpha \cap Z^*_\alpha$.

The only thing left to vertify is that in this case the intersection is transverse. Since Z_a 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 $\mathcal{C}_I(n, w+\epsilon)$ is the direct sum of the tangent spaces to Z_α and Z^*_α. This is a routine verification, and we omit the details. □

Finally, we are ready to prove lower bounds.

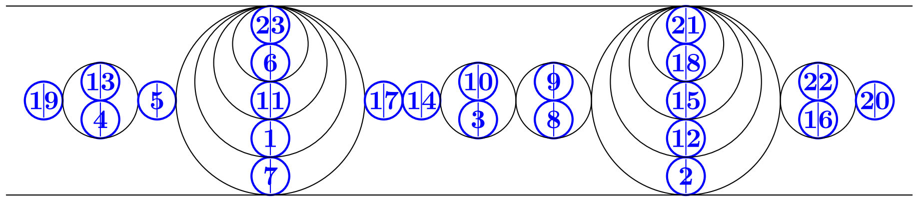

FIGURE 7. The single point of transverse intersection $Z_\alpha \cap Z_\alpha^*$.

Proof of lower bounds for part (2) of Theorem 1.1. If we verify the conditions of Lemma 4.1 we will have shown that if $n \geq qw + 2r$ (i.e. for sufficiently large n) then

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

This counts the number of special symbols in $\mathcal{A}(n, w)$. The mulitinomial 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 $n \to \infty$, then we write the simpler asymptotic expression

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

□

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

5. The solid, liquid, and gas regimes for finite 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) This is the same as the proof of (1) of Theorem 1.1, in Section 3.2.

(2) 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 an extra-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 $i : \mathcal{C}(n,w) \to \mathcal{C}(n,\mathbb{R}^2)$ does not induce an isomorphism on homology

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

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.

It seems clear that $i_*(Z_\alpha)$ and $i_*(Z_{\alpha'})$ are homologous in $\mathcal{C}(n,\mathbb{R}^2)$, so in other words $Z_\alpha - Z_{\alpha'}$ is in the kernel of i_*.

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

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

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 $n - \left\lceil \frac{n}{w} \right\rceil$, since the minimum number of blocks is $\left\lceil \frac{n}{w} \right\rceil$. So if $j \geq n - \left\lceil \frac{n}{j} \right\rceil + 1$, then there are no j-dimensional cells, in which case there is no nontrivial j-dimensional homology.

□

6. Discrete Morse theory

A *discrete vector field* V on a regular CW complex X is a collection of pairs of faces $[\alpha, \beta]$ where n 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 $[\alpha_1, \beta_1], [\alpha_2, \beta_2], \ldots, [\alpha_r, \beta_r]$ 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 [13] 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. 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 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 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 kth block is a follower, and their combined number of elements is at most w; or
(2) The kth 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 kth 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 kth 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 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 $v(v(\alpha)) = \alpha$. □

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

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

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 kth 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 kth 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 $\text{key}(\alpha_{i+1})$. 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. □

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. □

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 $\mathrm{cell}(n, w)$, 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 $\mathrm{code}_S(\alpha)$. 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 $\mathrm{code}_S(\alpha)$, the function code_S is injective. □

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 $j = q(w-1)+2r$, 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

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

completing the proof of Theorem 1.1. □

8. Comments

(1) The definitions of homological solid, liquid, and gas make sense even for 0th homology. The homological solid-liquid phase transition for 0th homology is the "sphere packing" problem. The largest radius spheres that will fit in the region corresponds to where the configuration space goes from empty to nonempty.

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, Diaconis writes about it in the context of ergodicity, a requirement for being able to effectively sample a configuration space by making small random movements of disks in his survey article [12]. See also [16] for discussion of the connectivity threshold.

(2) 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.

(3) Discrete Morse theory has been studied on the Salvetti complex before. For a more geometric approach to discrete gradients on cell(n), see [22], [20], and [19]. We do not know whether the techniques from these papers can improve the upper bounds on $\beta_j[\mathcal{C}(n, w)]$, or even produce perfect discrete Morse functions or minimal CW complexes for $\mathcal{C}(n, w)$.

(4) A related family of spaces is the "no k-equal space" studied in [5, 6, 7]. In particular, there is a natural map $\mathcal{C}(n, w) \to M^{\mathbb{R}}_{n,w+1}$ by projecting onto the x-coordinates. We do not know much about the induced map on homology in general. We point out here a coincidence we notice in our 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 Björner and Welker's paper [7], it seems possible that $H_1(\mathcal{C}(n, 2))$ is isomorphic to $H_1(M^{\mathbb{R}}_{n+1,3})$—at least the Betti numbers seem to be equal. Is this true?

We emphasize that it is a configuration space of n points in the first space and $n + 1$ points in the second space, so at least on the surface of things 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.

We showed that $\mathcal{C}(n, 2)$ is a $K(\pi, 1)$ in Section 3.2. The question of whether $M^{\mathbb{R}}_{n,3}$ is a $K(\pi, 1)$ was asked by Björner [4] and answered affirmatively by Khovanov [17]. Khovanov describes this as a real analogue of the fact that $M^{\mathbb{C}}_{n,2}$ (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 $\beta_j\,[\mathrm{cell}(n, w)]$ (for homology with $\mathbb{Z}/2$ coefficients) for small n using the software PHAT [3]. These appear in Table 1. For a point of reference, we note that cell(8) has over 5 million cells.

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

REFERENCES

[1] Vladimir I. Arnold. *The cohomology ring of the colored braid group*, pages 183–186. Springer Berlin Heidelberg, Berlin, Heidelberg, 2014.

[2] 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.

[3] 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/phat-code/phat`.

[4] 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.

[5] Anders Björner and László Lovász. Linear decision trees, subspace arrangements and Möbius functions. *J. Amer. Math. Soc.*, 7(3):677–706, 1994.

[6] 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.

[7] Anders Björner and Volkmar Welker. The homology of "k-equal" manifolds and related partition lattices. *Adv. Math.*, 110(2):277–313, 1995.

[8] Pavle V. M. Blagojević and Günter M. Ziegler. Convex equipartitions via equivariant obstruction theory. *Israel J. Math.*, 200(1):49–77, 2014.

[9] Gunnar Carlsson, Jackson Gorham, Matthew Kahle, and Jeremy Mason. Computational topology for configuration spaces of hard disks. *Phys. Rev. E*, 85:011303, Jan 2012.

[10] 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.

[11] Kenneth Deeley. Configuration spaces of thick particles on a metric graph. *Algebr. Geom. Topol.*, 11(4):1861–1892, 2011.

[12] Persi Diaconis. The Markov chain Monte Carlo revolution. *Bull. Amer. Math. Soc. (N.S.)*, 46(2):179–205, 2009.

[13] Robin Forman. A user's guide to discrete Morse theory. *Sém. Lothar. Combin.*, 48:Art. B48c, 35, 2002.

[14] 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.

[15] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[16] Matthew Kahle. Sparse locally-jammed disk packings. *Ann. Comb.*, 16(4):773–780, 2012.

[17] Mikhail Khovanov. Real $K(\pi, 1)$ arrangements from finite root systems. *Math. Res. Lett.*, 3(2):261–274, 1996.

[18] 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.

[19] Davide Lofano and Giovanni Paolini. Euclidean matchings and minimality of hyperplane arrangements. arXiv:1809.02476, 2018.

[20] Francesca Mori and Mario Salvetti. (Discrete) Morse theory on configuration spaces. *Math. Res. Lett.*, 18(1):39–57, 2011.

[21] M. Salvetti. Topology of the complement of real hyperplanes in $\mathbf{C}^N$. *Invent. Math.*, 88(3):603–618, 1987.

[22] Mario Salvetti and Simona Settepanella. Combinatorial Morse theory and minimality of hyperplane arrangements. *Geom. Topol.*, 11:1733–1766, 2007.

[23] 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.