Journal of Applied and Computational Topology (2020) 4:263–280
https://doi.org/10.1007/s41468-019-00043-w

Discrete configuration spaces of squares and hexagons

Hannah Alpert[1]

Received: 2 August 2017 / Accepted: 21 September 2019 / Published online: 28 September 2019
© Springer Nature Switzerland AG 2019

Abstract
A fixed fraction of a large box is occupied by tiny spheres; how difficult is it for those spheres to move around in the space available in the box, and how does that difficulty depend on the chosen density? We consider a discrete version of this problem, a generalization of the familiar fifteen-piece sliding puzzle on the 4 by 4 square grid. On larger grids with more pieces and more holes, asymptotically how fast can we move the puzzle into the solved state? We consider these questions for both sliding squares and sliding hexagons, to model very densely packed disks.

Keywords Configuration space · Sliding block puzzle · Parallel sorting · Grid graph · Hard sphere model

Mathematics Subject Classification 05C85 · 55R80 · 68W10

1 Introduction

A substance such as water can be a solid, liquid, or gas depending on the temperature and pressure. Theoretically we should be able to predict which phase we get, just from the attracting and repelling forces between individual molecules. Questions like this are part of statistical mechanics, which relates the large-scale properties of a substance to the properties of the tiny particles it is made of. In this paper we study questions related to the hard spheres model. The hard spheres model looks at the water molecules, pretends that they are identical spheres, and forgets the attracting and repelling forces between them, so that water now looks the same as any other substance. The only rule is that the spherical molecules must be disjoint. So, we study the set of all ways to arrange our collection of small disjoint spherical molecules in some container, such as a large cube. These arrangements are called configurations, and the set of them is called the configuration space. The hope is that the geometry

✉ Hannah Alpert
hcalpert@math.ubc.ca

1 University of British Columbia, 1984 Mathematics Road, Vancouver, BC, Canada

and topology of the configuration space already contains enough information to tell us about solid/liquid/gas phase transitions. Diaconis gives a quick overview of this program in Diaconis (2009), and Löwen gives a more detailed survey in Löwen (2000) from the physics perspective.

We don't care so much about the configuration space of a specific number of spheres of a specific size. What matters more is how much space the spheres cover as a fraction of the container's volume. This is called the packing fraction. We care most about the configuration spaces with a given packing fraction in which each sphere is negligibly small compared to the container. No one has mathematically defined how to take the limit of configuration spaces as the sphere size goes to zero, but nevertheless this limiting configuration space is called the thermodynamic limit. It is not so farfetched to want to define such a limit, because it has been done for other similar problems, such as in Shnirelman's work (1987). There he models the configuration space of a liquid as the space of volume-preserving diffeomorphisms of a 3–dimensional cube. He shows that this liquid configuration space is the limit of other configuration spaces by proving that the liquid configuration space with a particular metric is well approximated by the space of permutations of tiny cubes in the same 3–dimensional cube. To apply Shnirelman's framework to the hard spheres model, one would hope to come up with a collection of self-maps of the cube that can be said to be the limit of configuration spaces of hard spheres with fixed packing fraction in the cube; perhaps someone will do so in the future.

Researchers have found ways to study the hard spheres model even without defining the thermodynamic limit. Some have used computational experiments to study the configuration spaces in terms of their volume, which roughly means the probability that the spherical molecules will be disjoint if you place them independently at random positions in the container. The volume of configuration spaces is physically relevant because it is closely related to the Helmholtz free energy, which is a more commonly studied property of physical systems. These experiments have found that the volume seems to have a phase change at a packing fraction (density) of approximately .71 (see Löwen 2000). Here the definition of phase change is that if we consider the volume as a function of the packing fraction, then that function is not analytic. Another direction of current study is the Topological Hypothesis, which is a conjecture saying that phase change corresponds to a change in topology of the surfaces of constant potential in the configuration space, for a particular notion of potential; see (Carlsson et al. 2012) for a brief overview. Other topological exploration of sphere and disk configuration spaces can be found in papers such as Deeley (2011), Baryshnikov et al. (2013), and Alpert (2017).

In this paper, we study the diameter of the configuration spaces: given two configurations of n labeled unit spheres in an m by m by m cube, how fast can the spheres move from one configuration to the other if they are not allowed to intersect, leave the cube, or exceed a speed of 1? Actually, we consider only the 2–dimensional version of this question. In addition to its relevance to the hard spheres model in physics, the 2–dimensional version is a problem in motion planning, which is part of robotics. Instead of spheres we have n robots, each one a unit disk, and when we talk about a path through the configuration space, that means a way for the robots to move simultaneously without hitting each other. In the robotics context, the amount of time taken by the robots on a particular path through the configuration space is called the makespan. This terminology is useful because researchers in robotics also consider the computation time needed to plan how the robots will move. In the present paper we do not consider computation time, so all time estimates concern the makespan.

We would like to understand the diameter of the configuration space of disks in a box when the packing fraction is large. The maximum possible packing fraction would be if the disks were arranged in a hexagonal grid. If the packing fraction is very close to the maximum, how difficult is it to move the disks from one arbitrary configuration to another? One flaw in this question is that the configuration space is often not connected, partly because the disks can get stuck: there are arbitrarily sparse configurations of disks in which the disks cannot move [see Böröczky (1964) and Kahle (2012)]. If the configuration space is not connected, then its diameter is not well-defined. To simplify the question, we replace the disks by regular hexagons that can slide but are not allowed to turn. Then the densest packing is still a hexagonal grid—the tiling of the plane by hexagons—and the hexagons can still get stuck in some ways but the geometry is more predictable than that of disks. We simplify further by requiring that the hexagons stay at grid positions on the hexagonal lattice except when sliding between adjacent grid positions; the precise definition appears in Sect. 4. The resulting configuration space is still not connected but turns out to have a large connected component, and we estimate the diameter of this component.

Less unusual than sliding hexagons are sliding squares, as in the familiar fifteen-piece sliding puzzle. In the square version, there are labeled square tiles at all but a few grid positions, and any tile adjacent to a hole can slide into that hole. There is no way for sliding squares to get stuck, so the problem is simpler than that of hexagons. In this paper we address the square problem first as preparation for the hexagon problem.

This paper addresses questions similar to those in the paper (Demaine et al. 2018) but differs in a few relevant ways. That paper, like this one, is about taking two configurations of disks and estimating the minimum makespan of a path between them in the configuration space, and it also works mainly with a discrete version in which the robots are confined to a square grid. The discrete results in that paper are stronger than those in the present paper, in the sense that instead of just bounding the diameter of the configuration space, for each pair of configurations the authors look at the distance between the starting and ending positions of each robot, and they bound the makespan between those two configurations in terms of that distance. I wish that I had known about their fine work as I worked on this project, but their paper had not appeared when I first wrote a draft of this paper.

However, there are significant geometric differences between the discrete setups in the two papers. In the present paper, the hexagonal or square robots nearly tile their container, and they touch along edges as they slide into the empty spaces. In Demaine et al. (2018), the authors imagine the robots to be disks smaller than the grid squares, so that even with one robot at every grid position, the robots have enough room to move simultaneously along grid lines. So, they do not consider what may happen when the packing fraction is very close to the maximum possible. In fact, packing fraction is not a consideration in their paper because their robots move in the plane, not confined to a bounded region—there is no container. In addition to considering the packing fraction, the present paper differs from Demaine et al. (2018) by introducing the sliding hexagon question. As far as I know, this is the first time anyone has studied a hexagonal analogue of the square fifteen-piece sliding puzzle. The motion-planning paper (Chinta et al. 2018) considers robots on a hexagonal grid, but there as with the square grid there is enough space to have one unit-disk robot at each grid position and still be able to permute the robots. In the present paper, the hexagonal robots cannot move without some empty grid positions, and the availability of those empty positions is determined by the packing fraction.

In Sect. 2 we set up the definitions necessary for the main theorem for squares (Theorem 2), which we prove in Sect. 3. Sections 4 and 5 address the analogous result, the main theorem for hexagons (Theorem 7). In Sect. 4 we define the discrete configuration spaces of hexagons and prove basic connectivity properties. And, in Sect. 5 we prove the hole-moving lemma for hexagons (Lemma 10), which serves to reduce the hexagon problem to the square problem, completing the proof of the main theorem for hexagons.

2 Square configuration spaces

In this section we define the discrete configuration spaces of squares in order to set up the main theorem for squares (Theorem 2). Let $\widetilde{\mathrm{sq}}(n; m)$ denote the set of ways to arrange n pieces numbered 1 through n in distinct positions on the m by m square grid (so we have $n \leq m^2$). We refer to this as a ***labeled configuration space***. We let the ***unlabeled configuration space*** $\mathrm{sq}(n; m)$ be the set of ways to arrange n indistinguishable pieces in distinct positions on the m by m square grid. Sometimes instead of an m by m square grid we want an m_1 by m_2 rectangular subset of the square grid, and use the notation $\widetilde{\mathrm{sq}}(n; m_1, m_2)$ and $\mathrm{sq}(n; m_1, m_2)$ for the corresponding configuration spaces.

These configuration spaces can be viewed as metric spaces. First we give a naive choice of metric and show in Proposition 1 that we can easily estimate, up to a constant factor, the diameter of the resulting metric space. Then we give a more physically relevant metric that we use for the remainder of the paper.

Naively, we can say that two configurations are adjacent, or have distance 1, if they differ by moving one piece into a hole (i.e., an unoccupied position) adjacent to it. We refer to the metric generated by this relation as the L^1 ***intrinsic metric***. It is always greater than or equal to the L^1 ***extrinsic metric***, under which the distance between two configurations is the sum over all pieces of the distance along the grid between the two positions of that piece. (This definition is for the labeled configuration space; for the unlabeled space, take the minimum distance achieved over all labelings of the two configurations.) That is, the L^1 extrinsic metric would be the minimum number of moves between the two configurations if multiple pieces were allowed to occupy the same position.

We would like to estimate the diameter of the labeled configuration space $\widetilde{\mathrm{sq}}(n; m)$. As m grows, does the L^1 intrinsic diameter grow a lot faster than the L^1 extrinsic diameter? The following proposition shows that the answer is no.

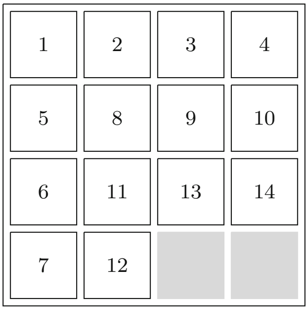

Fig. 1 The home configuration for $n = 14$ pieces on a board of side $m = 4$. Because we can move the pieces home in numerical order in $O(nm)$ moves, the L^1 intrinsic diameter of the configuration space $\widetilde{\mathrm{sq}}(n; m)$ is $O(nm)$

Proposition 1 *The labeled configuration space $\widetilde{\mathrm{sq}}(n; m)$, with $n \le m^2 - 2$ so that the space is connected, has L^1 intrinsic diameter that grows as $\Theta(nm)$, matching the L^1 extrinsic diameter up to a constant factor.*

Proof We check first that the L^1 extrinsic diameter grows as $\Omega(nm)$, and then that the L^1 intrinsic diameter grows as $O(nm)$. For the extrinsic: take any configuration, and permute the pieces by shifting the top $\lceil \frac{m}{2} \rceil$ rows down past the remaining $\lfloor \frac{m}{2} \rfloor$ rows. Each piece travels distance $\lfloor \frac{m}{2} \rfloor$ or $\lceil \frac{m}{2} \rceil$, so the L^1 extrinsic distance between the two configurations is at least $\lfloor \frac{m}{2} \rfloor \cdot n$.

To check that the L^1 intrinsic diameter grows as $O(nm)$, we designate one home configuration and sketch an $O(nm)$—step algorithm for getting home. The home configuration is as follows, shown in Fig. 1: numbers 1 through m go in order along the top row, then $m + 1$ through $2m - 1$ down the left column, then the remainder of the second row, then the remainder of the second column, and so on until the pieces run out. We use induction on m. It takes $O(m)$ moves to move a hole to be adjacent to piece 1 and then move piece 1 home, then another $O(m)$ moves to bring piece 2 home, and so on, to fill the top row and left column with the correct pieces. Extra care is needed at the ends of the row and column, but that adjustment takes only a constant number of moves. Then we apply the inductive hypothesis to the remaining grid; having at least two holes guarantees that in the base case of the 2 by 2 grid we can get to the home configuration. □

Having quickly resolved the question of L^1 intrinsic diameter, we do not pursue it any further. In the remainder of the paper, we allow arbitrarily many independent moves to occur simultaneously. This new construction is more reasonable from a physics perspective, if we imagine the pieces to be independent particles that might randomly wander into neighboring holes. We say that the L^∞ ***intrinsic metric*** is generated by the relation that two configurations are adjacent, and have distance 1, if they differ by moving any number of pieces into holes adjacent to them. (Note that we do not permit sliding a piece into a position that is simultaneously being vacated; it must really be a hole.) The L^∞ ***extrinsic metric*** is the maximum over all pieces of the distance along the grid between the positions of that piece in the two configurations. If the number of holes is bounded, then the L^∞ intrinsic metric and the L^1 intrinsic metric differ by at most a constant factor equal to the number of holes, because at most that many pieces can move at a time. In the remainder of the paper, we fix the proportion of positions that contain pieces, so that as m grows the number of holes grows. The main theorem for squares (Theorem 2), proved in the next section, says that for a fixed ratio of pieces to holes, the L^∞ intrinsic and extrinsic diameters match up to a constant factor.

3 Main theorem for squares

In this section we prove the main theorem for squares (Theorem 2), except for the proof of the hole-moving lemma for squares (Lemma 5). The proof of Lemma 5 serves as a template for the proof of the analogous lemma about hexagons, so it appears in Sect. 5 with the corresponding hexagon lemmas.

Theorem 2 (Main theorem for squares) *Fix $\alpha \in (0, 1)$. Then for m a power of 2, and $1 \le n \le \alpha m^2$, the labeled configuration space $\widetilde{\mathrm{sq}}(n; m)$ has L^∞ intrinsic diameter that grows as $\Theta(m)$, matching the L^∞ extrinsic diameter up to a constant factor depending only on α.*

The hypothesis that m is a power of 2 is only for convenience. The result should be true for arbitrary m, but that would require tedious modifications to the proof, including modifying the cited result Theorem 3 which in its original paper is only proved for m a power of 2.

The proof of the theorem is based on the solution to the following simpler problem. Suppose that we have an m by m grid with pieces numbered 1 through m^2 in distinct positions. We can swap a pair of adjacent pieces, and in fact any disjoint set of such swaps can happen simultaneously. How long does it take to sort the pieces? Taking the L^∞ extrinsic diameter gives a lower bound of $\Omega(m)$ time steps; can we sort in $\Theta(m)$ time steps?

An affirmative answer was found by Thompson and Kung in 1977, and appears below as Theorem 3. Furthermore, the sequence of comparisons in their algorithm does not depend on the starting permutation; such an algorithm is called a ***sorting network*** on m^2 numbered pieces. More specifically, by a ***comparison step*** we mean an ordered pair of positions adjacent in the grid, indicating that the pieces in those positions should be swapped if necessary to match the order specified. By a ***routing step*** we mean an unordered pair of positions adjacent in the grid, indicating that the pieces in those positions should be swapped without comparing them. The algorithm is a sequence of time steps, each consisting of a collection of disjoint comparison steps or routing steps. The pieces end up in ***snake-like row major order***, which means that the numbering goes across the top row to the right, across the second row to the left, and so on, alternating directions.

Theorem 3 [Thompson and Kung (1977)] *Consider the m by m grid graph, with pieces numbered 1 through m^2 distributed one per vertex. For m a power of 2, there is a sorting-network algorithm to sort the pieces into snake-like row major order, using $\Theta(m)$ time steps of disjoint comparisons or routings along grid edges.*

The proof of the main theorem for squares (Theorem 2) is based on this theorem of Thompson and Kung, but in our setup there may be many fewer holes than there are pieces, so we can't swap as many pairs at a time. Very roughly, the strategy is to group together several pieces that are near one hole, and to think of applying Thompson and Kung's algorithm to the groups of pieces.

To be more precise, given a sorting network on n numbered pieces, we can apply it instead to n buckets each containing an unordered set of k numbered pieces, as follows. At each comparison step, we take the $2k$ pieces in the two buckets and redistribute so that the least k pieces are in one bucket and the greatest k pieces are in the other. Then the greater of the two buckets takes the role of the greater of the two pieces in the original sorting network, and the lesser bucket takes the role of the lesser piece. I would not be surprised if the following lemma is well-known, but the proof is included for completeness.

Lemma 4 (Bucket lemma) *Given a sorting network on n pieces, if we apply it instead to n buckets of k pieces, then it sorts the pieces: the least k pieces end up in the least bucket, the next k pieces in the next bucket, and so on.*

Proof We use the zero-one principle for sorting networks, explained in Theorem Z of Section 5.3.4 of Knuth (1973), which says that it suffices to show that if all pieces have labels 0 and 1 instead of distinct labels, then the algorithm sorts them into non-decreasing order.

We use induction on the number of buckets in the starting configuration that contain both 0 and 1 (mixed buckets). The base case is if there is at most one mixed bucket; in this case the sorting proceeds as if each bucket were a 0 piece, a 1 piece, or in the case of the mixed bucket, a piece labeled $\frac{1}{2}$. These buckets are sorted correctly by the sorting network.

Suppose that there are at least two mixed buckets in the starting configuration. There must be, at some time during the sorting, a comparison step involving two mixed buckets; we select a first such comparison step. After this step, at most one of the two buckets is a mixed bucket, and the sorting proceeds the same as if these two buckets had come into this comparison step in the same state that they leave it. That is, we can follow these two buckets backward in time to construct a different starting configuration with one mixed bucket fewer than our original starting configuration, but with the same ending configuration. By the inductive hypothesis, this ending configuration must be the sorted configuration. □

The remaining ingredient needed to prove the main theorem for squares (Theorem 2) is the following lemma, for which the proof appears in Sect. 5 along with the corresponding proof for hexagons.

Lemma 5 (Hole-moving lemma for squares) *For m a power of 2, and arbitrary $n \leq m^2$, the unlabeled configuration space $\mathrm{sq}(n; m)$ has L^∞ intrinsic diameter that grows as $O(m)$ independent of n.*

Using this lemma we can finish the proof of the main theorem for squares.

Proof of main theorem for squares (Theorem 2) The L^∞ extrinsic diameter of the configuration space $\widetilde{\mathrm{sq}}(n; m)$ is $\Theta(m)$ because the furthest that a piece may need to move is between two opposite corners of the grid. Thus, we need to provide an algorithm for sorting $\widetilde{\mathrm{sq}}(n; m)$ to some home configuration in $O(m)$ time steps. Roughly, for some choice of r we view the m by m grid as an $\frac{m}{r}$ by $\frac{m}{r}$ grid of r by r blocks, and each block plays the role of a bucket in the sorting network on the $\frac{m}{r}$ by $\frac{m}{r}$ grid.

More specifically, we fix r a power of 2 large enough that $\alpha \leq \frac{r^2-1}{r^2}$. This means that for m sufficiently large, if we view the m by m grid as an $\frac{m}{r}$ by $\frac{m}{r}$ grid of r by r blocks, then we have enough holes to put at least one in each block. Note that the L^∞ intrinsic diameter is monotonic in n, since we can always delete a piece from any sequence of moves. So it suffices to provide a sorting algorithm for the case where $n = \left(\frac{m}{r}\right)^2 (r^2 - 1)$ and we have the same number of holes as blocks.

The sorting algorithm goes as follows. First we use the hole-moving lemma (Lemma 5) to move the holes so that each block has one hole. Then we apply Thompson and Kung's sorting algorithm (Theorem 3), using the blocks as the buckets from Lemma 4. That is, each pair of blocks forms an r by $2r$ or $2r$ by r grid with two holes, so it can be sorted so that the lesser half of the pieces end up in the first block in order, and the greater half end up in the second block in order; the time required is a function only of r (and hence of α but not m). All together, by Theorem 3 the sorting takes $O\left(\frac{m}{r}\right) = O(m)$ time steps.

At the end of this process, the blocks are in snake-like row major order and the pieces in each block are sorted—say, also in snake-like row major order with the hole at the end. Thus, it takes $O(m)$ time steps to get from an arbitrary configuration to this home configuration. $\square$

4 Sorting hexagons with six holes

In the hexagonal sliding puzzle, instead of numbered square tiles we use numbered hexagon tiles. Geometrically if a hole is surrounded by pieces, then none of the pieces can move into the hole. But if two holes are adjacent, then a piece adjacent to both holes can slide into either hole, as shown in Fig. 2. We can ask the same asymptotic questions about the diameter of this hexagon puzzle as for the square puzzle. For our growing boards it would be most natural to use a hexagon shape, but we use a parallelogram shape so that the techniques from the square case carry over better. By an m_1 ***by*** m_2 ***parallelogram*** we mean m_1 rows each of m_2 hexagons, each row half a step to the right of the previous row.

We let $\widetilde{\mathrm{hex}}(n; m_1, m_2)$ and $\mathrm{hex}(n; m_1, m_2)$ denote the labeled and unlabeled configuration spaces of n hexagon pieces on an m_1 by m_2 parallelogram board, and let $\widetilde{\mathrm{hex}}(n; m)$ and $\mathrm{hex}(n; m)$ denote those on an m by m parallelogram board. We allow a piece to move into either of two adjacent holes, and require that if multiple pieces move simultaneously, the pairs of holes that they use must be disjoint. If we can make a set of simultaneous moves between two configurations in this way, we say that the configurations are adjacent and have L^∞ intrinsic distance 1; the L^∞ ***intrinsic metric*** is generated by these pairs. The L^∞ ***extrinsic metric*** is, as with squares, the maximum over all pieces of the number of steps needed to move the piece on an empty board between its two positions.

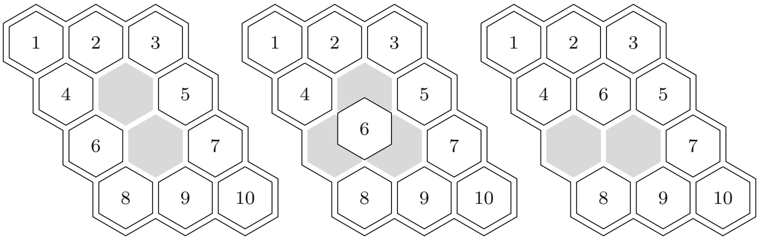

Fig. 2 In the configuration space $\widetilde{\text{hex}}(10; 4, 3)$ a piece adjacent to two adjacent holes can move into either of those positions. The intermediate configuration, with the piece balanced between three positions, is not part of the configuration space

We would like to prove a theorem for hexagons analogous to the main theorem for squares (Theorem 2). However, unlike the square configuration spaces with at least two holes, the space $\widetilde{\text{hex}}(n; m)$ is not always connected. Therefore it does not make sense a priori to ask about its diameter. For instance, any configuration where all holes are isolated is not adjacent to any other configuration. However, the next proposition states that as long as there are at least six holes, this is the only obstruction to connectivity, so each hexagon configuration space with at least six holes consists of a large connected component and possibly some isolated configurations. Just as sorting a square grid with two holes is an important step in the main theorem for squares (Theorem 2), this proposition is an important step in the main theorem for hexagons (Theorem 7), which concerns the diameter of the large connected component of $\widetilde{\text{hex}}(n; m)$.

Proposition 6 *For $m \geq 5$ and $n \leq m^2-6$, in the labeled configuration space $\widetilde{\text{hex}}(n; m)$ there is a large connected component containing every configuration for which not all the holes are isolated.*

Proof It suffices to prove the statement when there are exactly six holes. We start with an arbitrary configuration in which some two holes are adjacent, and show that we can move the holes into a 3 by 2 parallelogram in the upper-left corner, and sort the pieces into an arbitrary order.

First we move the holes. From the point of view of the holes, a hole is permitted to move one step whenever there is another hole to which it remains adjacent. So, we can take the two adjacent holes in our starting configuration and walk them over to a third hole, then move all three—staying connected—into the upper left corner. Then, depositing one hole in the corner, we can walk the remaining two over to retrieve the next hole, and so on. In this way we move the holes to form a 3 by 2 parallelogram in the upper-left corner.

Next we want to show that we can use this block of holes to permute the pieces arbitrarily. We start by showing this for a 3 by 4 board: the holes are in the first two columns, and pieces labeled 1 through 6 are in the remaining two columns, say in column-major order. We claim that we can transpose pieces 1 and 2, pieces 2 and 3, pieces 3 and 4, pieces 4 and 5, or pieces 5 and 6. By composing these transpositions we can achieve any permutation. The method for making the transpositions is as follows, shown in Fig. 3. It is not hard to see how to transpose pieces 1 and 2, or 2 and 3, while keeping pieces 4, 5, and 6 in place. Similarly, to transpose pieces 4 and 5, or 5 and 6,

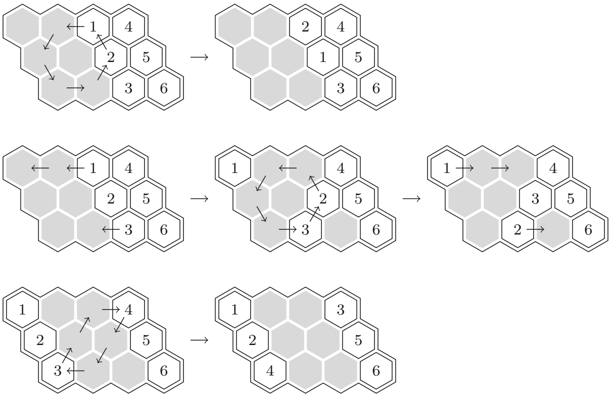

Fig. 3 In this 3 by 4 parallelogram with 6 pieces, we can swap any two consecutively numbered pieces; swaps (12), (23), and (34) are shown

we start by shifting all pieces left by two columns. It remains to see how to transpose pieces 3 and 4. There is enough room to do this if we start with pieces 1, 2, and 3 shifted to the leftmost column and pieces 4, 5, and 6 in the rightmost column.

This same method shows that we can sort a 3 by m parallelogram for arbitrary m, with the first two columns all holes. For a parallelogram with more rows, we show that we can transpose any two pieces that are in the same row and adjacent columns, or that are in the same column and adjacent rows; these transpositions are more than enough to generate all permutations. Given such a pair of adjacent pieces, if they are both in the first three rows, then we already know what to do. Otherwise, we shift the block of holes to the right until its pair of columns do not contain either of the pieces we want to swap; if those two pieces are both in column at least 3, we do not move the holes at all, and if not, then both pieces are in column at most 3, so we may move the holes to columns 4 and 5—this is why we have assumed $m \geq 5$. Then we shift the block of holes down until the pair of pieces we want to transpose is in the three rows that have holes in them. We use the 3 by m case to achieve the desired transposition, then undo the hole-moving to return the other pieces to where they were. By composing such transpositions, we can sort the pieces into any order.

Note that the proposition is still true for the cases $m = 3$ and $m = 4$, which can be done by hand, but the proof is not included here. □

Having established the existence of a large connected component, we are ready to state the main theorem for hexagons.

Theorem 7 (Main theorem for hexagons) *Fix* $\alpha \in (0, 1)$. *Then for* m *a power of* 2, *and* $1 \leq n \leq \alpha m^2$, *the large connected component of the labeled configuration space* $\widetilde{\text{hex}}(n; m)$ *has* L^∞ *intrinsic diameter that grows as* $\Theta(m)$, *matching the* L^∞ *extrinsic diameter up to a constant factor depending only on* α.

The proof is completely analogous to the proof of the main theorem for squares (Theorem 2). However, the hole-moving step, in which the holes are moved from their arbitrary starting positions to be evenly distributed over the board, is trickier for hexagons. The hole-moving lemma (Lemma 10) is the main goal of Sect. 5, which then ends with the completed proof of the main theorem for hexagons.

5 Hole-moving lemmas

In this section we prove the hole-moving lemma for squares (Lemma 5), the hole-moving lemma for hexagons (Lemma 10), and the main theorem for hexagons (Theorem 7). First we prove the hole-moving lemma for squares as a template for the hole-moving lemma for hexagons. (There are simpler proofs for the square version, but we give a proof here that can be modified for hexagons.) Then we prove the hexagon version and finish the main theorem for hexagons.

Here we repeat the statement of the hole-moving lemma for squares.

Lemma 5 (Hole-moving lemma for squares) *For m a power of 2, and arbitrary $n \leq m^2$, the unlabeled configuration space $\mathrm{sq}(n; m)$ has L^∞ intrinsic diameter that grows as $O(m)$ independent of n.*

The proof relies on the hole-pushing sublemma (Sublemma 8) and the hole-turning sublemma (Sublemma 9).

Sublemma 8 (Hole-pushing sublemma for squares) *For arbitrary $n \leq m$, in the one-row unlabeled configuration space $\mathrm{sq}(n; 1, m)$ the L^∞ intrinsic distance between the configuration with all holes on the left and the configuration with all holes on the right is bounded above by m.*

Proof Start with all holes on the left, and number the pieces 1 through n from left to right. We can move piece k to the left at time steps k through $k + m - n - 1$ to get to the configuration with all holes on the right. $\square$

Sublemma 9 (Hole-turning sublemma for squares) *For m_1 and m_2 powers of 2, and arbitrary $n \leq m_1 m_2$, in the unlabeled configuration space $\mathrm{sq}(m_1 m_2 - n; m_1, m_2)$ the configuration with holes in the first n positions in row-major order and the configuration with holes in the first n positions in column-major order have L^∞ intrinsic distance that grows as $O(m_1 + m_2)$ independent of n.*

Proof Because we may swap the roles of the rows and the columns, we may assume without loss of generality that $m_1 \leq m_2$; in particular, we know that m_1 and m_2 are powers of 2, so m_1 divides m_2. We start in column-major order and give an algorithm for changing to row-major order. We use induction on $m_1 + m_2$. We use different strategies for when the holes make up at most half the grid and for when they make up more than half the grid. Of course we could just swap the roles of pieces and holes, but this trick will not work in the hole-turning sublemma for hexagons (Sublemma 12).

We divide the m_1 by m_2 board into quadrants of size $\frac{m_1}{2}$ by $\frac{m_2}{2}$. In the case where there are $n \leq \frac{1}{2} m_1 m_2$ holes, we use at most m_2 time steps to push the lower-left-quadrant holes into the lower-right quadrant, $\frac{m_2}{2}$ steps to the right. Then we use the inductive hypothesis to change each quadrant into row-major order (really only the upper-left and lower-right quadrants because the other two have no holes). Then we

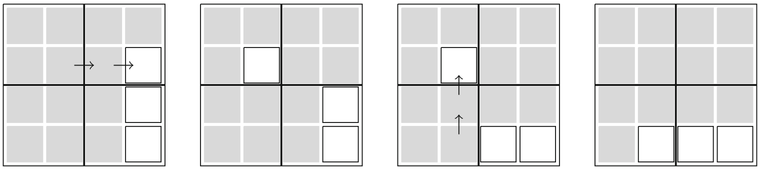

Fig. 4 The hole-turning method is to start in column-major order, push some holes right, turn each quadrant, and push all holes up into row-major order. (The holes are shown shaded gray)

use at most m_1 time steps to push the lower-right-quadrant holes into the upper-right quadrant, $\frac{m_1}{2}$ steps up. At this stage, using the fact that m_1 divides m_2 we observe that there is some k such that each of the two upper quadrants consists of k full rows of holes and then one more row that is empty or partial or full. We push this last partial row of holes from the upper-right quadrant to the upper-left quadrant, using at most m_2 time steps, to finish putting the whole board of holes into row-major order.

If it takes at most $C \cdot \left(\frac{m_1}{2} + \frac{m_2}{2}\right)$ time steps to turn each quadrant, then the number of time steps needed for the whole process is at most

$$m_2 + C \cdot \left(\frac{m_1}{2} + \frac{m_2}{2}\right) + m_1 + m_2 \leq \left(2 + \frac{1}{2}C\right) \cdot (m_1 + m_2),$$

which is at most $C \cdot (m_1 + m_2)$ as long as we choose $C \geq 4$.

The case where there are $n > \frac{1}{2}m_1m_2$ holes is slightly more complicated, and is depicted in Fig. 4. First we swap the configurations in the upper two quadrants, in the following way. Suppose the starting configuration has k holes in the upper-right quadrant, in column-major order, and $\frac{1}{4}m_1m_2$ holes filling the upper-left quadrant. We fix the first k holes (in column-major order) in the upper-left quadrant; there are $\frac{1}{4}m_1m_2$ other holes in the upper half of the board, and we push these to the right until they fill the upper-right quadrant. This takes at most m_2 time steps. Then we use the inductive hypothesis to change each quadrant to row-major order. Then we push all holes upward as far as possible, using at most m_1 time steps, and as above we use at most m_2 time steps to push the last partial row of holes from right to left. The result is row-major order, and again the total time is at most $\left(2 + \frac{1}{2}C\right) \cdot (m_1 + m_2) \leq C \cdot (m_1 + m_2)$ steps. □

Proof of hole-moving lemma for squares (Lemma 5) We start with an arbitrary configuration of holes, and give an algorithm for changing to the configuration where all the holes precede all the pieces in row-major order. The process is depicted in Fig. 5. We use induction on m. Using the inductive hypothesis (with columns and rows swapped), we put each $\frac{m}{2}$ by $\frac{m}{2}$ quadrant into column-major order. In column-major order there may be several full columns of holes followed by at most one partial column, which we call the short column. We combine the two upper quadrants and combine the two lower quadrants by working separately on the upper and lower half-boards as follows. In each half-board, push down the holes in the short column of the right quadrant, and then push all holes left. At this stage, in each half-board there may be several full

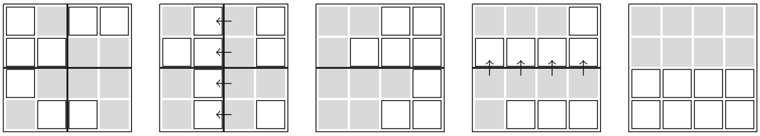

Fig. 5 The hole-moving method is to put each quadrant into column-major order, merge top quadrants and merge bottom quadrants, turn top and bottom half-boards, and merge top and bottom into row-major order. (The holes are shown shaded gray)

columns of holes followed by at most one partial column. If the combined number of holes we had in the short columns of the left and right quadrants is less than the height of the half-board (namely, $\frac{m}{2}$), then this partial column consists of an interval of holes at the top and an interval of holes at the bottom; if the combined number of short-column holes is greater than the height of the half-board, then the partial column consists of an interval of holes somewhere in the middle of the column. In either case we refer to this combined partial column as a short column and push up all the holes in it to the top of the half-board.

Having combined the two upper quadrants and the two lower quadrants, we apply the hole-turning sublemma (Sublemma 9) to the upper and lower half-boards to convert them into row-major order. Then we combine: push right the holes in the short row of the lower half, then push all holes up, then push left the holes in the short row. If it takes time at most $C \cdot \frac{m}{2}$ to put each quadrant into column-major order, and time at most $C' \cdot m$ for the other steps, then the total time is at most $C \cdot m$ as long as we choose $C \geq 2C'$. □

Next we prove the hole-moving lemma for hexagons.

Lemma 10 (Hole-moving lemma for hexagons) *For m a power of 2, and arbitrary $n \leq m^2$, the large connected component of the unlabeled configuration space* $\mathrm{hex}(n; m)$ *has L^∞ intrinsic diameter that grows as $O(m)$ independent of n.*

The proof again relies on a hole-pushing sublemma (Sublemma 11) and a hole-turning sublemma (Sublemma 12). For hexagons the hole-pushing sublemma is more involved than what is needed for squares. It says that we can push the holes in two half-boards together.

Sublemma 11 (Hole-pushing sublemma for hexagons) *For m_1 and m_2 divisible by 2, and n satisfying $3 \leq n \leq m_1 m_2$, consider the unlabeled configuration space* $\mathrm{hex}(m_1 m_2 - n; m_1, m_2)$.

1. *Consider the configurations in which the holes in the left m_1 by $\frac{m_2}{2}$ half-board occupy the first positions in column-major order and the remaining holes occupy the first positions in column-major order of the right half-board. The L^∞ intrinsic distance from any of these configurations to the one where the holes occupy the first n positions in column-major order has an $O(m_1 + m_2)$ upper bound independent of n.*
2. *Consider the configurations in which the holes in the left and right half-boards occupy the first positions in row-major order, and there is only one row of the full*

board that includes both holes and pieces. The L^∞ intrinsic distance from any of these configurations to the one where the holes occupy the first n positions in row-major order has an $O(m_2)$ upper bound independent of n.

Proof For part (1), we modify the strategy for combining quadrants or half-boards of squares described in the proof of Lemma 5: push down the short column of holes in the right half-board, push all holes to the left, then push up the short column. First we consider the case where the right half-board has at least one full column of holes. In this case pushing down the short column proceeds just as in the hole-pushing sublemma for squares (Sublemma 8). To push left, we divide the board into 2 by m_2 ribbons. Within each ribbon, we pair up the holes by column and walk each pair to the left along the ribbon, with each pair playing the role of a square hole in the hole-pushing sublemma for squares; if there is only one hole in the right-most column, we group it with the previous column and move that group of 3 to the left together, somewhat slower than the groups of 2 but still in $O(m_2)$ steps. After pushing left, we push up the short column again as in the hole-pushing sublemma for squares.

Next we consider the case where the right half-board has less than a full column of holes. If there is just one hole in the right half, we send a pair of holes from the left half to walk over and retrieve it. Otherwise, to push that column down, we need to use the column to the right of it as well, with a zigzag motion as in Fig. 6. We group the holes into adjacent pairs and triples. We walk each group one position to the right, so that the entire column of holes is shifted right, and then walk each grouping one position down and to the left; this process pushes the column of holes down one space from its original position. Using this strategy, we push the column of holes all the way down. In the same groupings, we push all the way to the left. Then we push the short column up, either directly or in a zigzag depending on whether there are holes to the left of it or not.

For part (2), if there is at least one row of all holes, then we can push the holes (or pieces) in the partial row just as in the hole-pushing sublemma for squares (Sublemma 8), taking at most m_2 time steps. If there is no row of all holes, then the strategy is to consider the 2 by m_2 ribbon consisting of the first two rows (the partial first row, and the all-pieces second row), pair up adjacent holes in the right half-board, and walk each pair to the left along the ribbon. One possible difficulty is if there is only one hole in the right half-board; in this case, because $n \geq 3$ there are at least two adjacent holes in the left half-board, and we can walk them over to the lone hole using $O(m_2)$ steps before starting the process of moving the right holes to the left. The other possible difficulty is if there is an odd number of holes in the right half-board; in this case, as above we divide the holes into pairs and perhaps one triple, and move the triple along the ribbon as one group. In any of these cases, it takes $O(m_2)$ time steps to move the holes from the right half of the first row to the left half. □

Sublemma 12 (Hole-turning sublemma for hexagons) *For m_1 and m_2 powers of 2, and arbitrary $n \leq m_1 m_2$, in the unlabeled configuration space* $\mathrm{hex}(m_1 m_2 - n; m_1, m_2)$ *the configuration with holes in the first n positions in row-major order and the configuration with holes in the first n positions in column-major order have L^∞ intrinsic distance that grows as $O(m_1 + m_2)$ independent of n.*

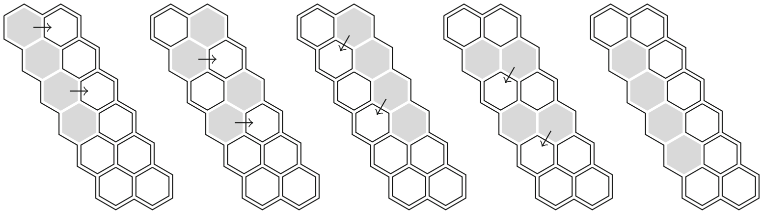

Fig. 6 Zigzagging between two columns, an arbitrarily long column of holes can advance up to m_1 positions downward in $O(m_1)$ time steps. (The holes are shown shaded gray)

Proof We modify the proof of the hole-turning sublemma for squares (Sublemma 9). There, the method consists of four steps: horizontal pushing, turning each quadrant, vertical pushing, and horizontally pushing the last partial row. At the end of each of these steps, the goal configuration of hexagon holes exactly corresponds to the goal configuration of square holes in the corresponding step of the hole-turning sublemma for squares (with one exception, addressed below), but in order to get to that goal configuration we need to apply the hole-pushing sublemma for hexagons (Sublemma 11).

For the first horizontal pushing step, we observe that getting from the goal configuration back to the original configuration is the situation addressed by the first part of the hole-pushing sublemma for hexagons, so we can find the necessary steps and then follow them in reverse to get from the original configuration to the goal configuration. For the vertical pushing step we apply hole-pushing as in the first part of the hole-pushing sublemma, but with the roles of rows and columns swapped. And, for the final horizontal pushing step we apply the second part of the hole-pushing sublemma.

The one exception arises when we want to apply the hole-pushing sublemma for hexagons but there are fewer than 3 holes in the relevant half-board. This situation occurs if we have $n = \frac{m_1}{2} + 1$ or $n = \frac{m_1}{2} + 2$ so that we start with just over half a column of holes. (In all other situations with too few holes, there is a uniform bound on n, so we know we can move the holes in $O(m_1 + m_2)$ time steps.) In this case, we can swap the configurations in the upper half and the lower half of the board by applying the second part of the hole-pushing sublemma for hexagons (Sublemma 11) in reverse and with the roles of rows and columns swapped, so that the first column has one or two holes at the top and $\frac{m_1}{2}$ holes at the bottom. After this modification we can proceed as in the usual case: push the lower-left quadrant holes to the lower-right quadrant, turn each quadrant, push the lower-right quadrant holes to the upper-right quadrant, and push the partial row of holes to the left. □

Proof of hole-moving lemma for hexagons (Lemma 10) We modify the proof of the hole-moving lemma for squares (Lemma 5). There, the method consists of four steps: moving each quadrant to column-major order, horizontal pushing, turning upper and lower half-boards, and vertical pushing. For hexagons we need to add a preliminary step of getting one pair of adjacent holes into each quadrant. To do this, we walk our original pair over to a third hole, and carry it over to a fourth hole, and so on, until we have four pairs of adjacent holes. (If the total number of holes is less than 8, we can carry them all directly to row-major order and be done.) Then we move one pair of holes to each quadrant; this whole process can be done in $O(m)$ time steps. For the remaining steps of the proof for squares, we apply the hexagonal analogues: the inductive hypothesis, hole-pushing (Sublemma 11), hole-turning (Sublemma 12), and hole-pushing again. The result is an algorithm to move the holes to precede all the pieces in row-major order, using $O(m)$ time steps.

Having proved the hole-moving lemma, we are ready to prove the main theorem for hexagons.

Proof of main theorem for hexagons (Theorem 7) The proof is very similar to the proof of Theorem 2. As with the case of squares, the L^∞ extrinsic diameter is $\Theta(m)$. To show the L^∞ intrinsic diameter of the large component is $O(m)$, we fix r a power of 2 large enough that $\alpha \leq \frac{r^2-3}{r^2}$, and view the m by m board as an $\frac{m}{r}$ by $\frac{m}{r}$ grid of r by r blocks. We may assume that the number of holes is exactly three times the number of blocks.

First we use the hole-moving lemma (Lemma 10) to move the holes so that each block has three holes, adjacent to each other. Then we apply Thompson and Kung's sorting algorithm (Theorem 3), using Proposition 6 to sort the pair of adjacent blocks in each comparison step. This process allows us to sort the pieces in $O(m)$ time steps so that the blocks are in snake-like row major order and the pieces in each block are in snake-like row major order with the three holes at the end. $\square$

6 Conclusion

When exploring configuration spaces such as these, one goal is to find evidence of a phase transition: does the diameter of $\widetilde{\mathrm{sq}}(n; m)$ or $\widetilde{\mathrm{hex}}(n; m)$ grow differently when the density $\frac{n}{m^2}$ is low, compared to when the density is high? The main results of this paper state that for each density $\alpha \in (0, 1)$ the quantities $\frac{1}{m} \operatorname{diam}(\widetilde{\mathrm{sq}}(\alpha m^2; m))$ and $\frac{1}{m} \operatorname{diam}(\widetilde{\mathrm{hex}}(\alpha m^2; m))$ are bounded as m goes to infinity. If each one converges to a limit, then that limit is a function of α, and for any value of α at which this function or its derivatives are discontinuous, we can say that the diameter has a phase transition at that density α. One could speculate on which densities are good candidates for a phase transition, such as the density at which there exist configurations where every piece is adjacent to a hole. It would also be useful to estimate the limits of $\frac{1}{m} \operatorname{diam}(\widetilde{\mathrm{sq}}(\alpha m^2; m))$ and $\frac{1}{m} \operatorname{diam}(\widetilde{\mathrm{hex}}(\alpha m^2; m))$, if they exist, in terms of α. Put differently, can we estimate the diameter of $\widetilde{\mathrm{sq}}(n; m)$ or $\widetilde{\mathrm{hex}}(n; m)$ as a function of both n and m, up to a constant factor?

The purpose of studying the square and hexagon configuration spaces was to gather clues about disk configuration spaces. Do the main theorems of this paper have an analogue that describes the diameters of disk configuration spaces or their connected components? And, for any further results on whether the square and hexagon problems show a phase transition in diameter as we vary density, do those statements apply to disks as well?

This work suggests many other questions. For instance, does the main theorem for squares (Theorem 2) also apply to square grids in higher dimensions? More generally, to what extent does Thompson and Kung's result (Theorem 3) generalize to arbitrary graphs? Can we guarantee a sorting time in terms of the diameter, say, and some other properties of the graph? Specifically, given a graph on n vertices, we can construct another graph on the $n!$ permutations of the vertices. Two permutations are adjacent in the new graph if there exists a set of disjoint edges in the original graph such that we can get from one permutation to the other by swapping the two ends of each of the chosen edges. Can we estimate, up to a constant factor, the diameter of the new graph?

Finally, we could ask about properties of the configuration spaces other than diameter. The notion of adjacency used in defining the L^∞ intrinsic metric turns the square and hexagon configuration spaces into graphs. Instead of diameter, we could look at other graph-theoretic properties such as expansion. Another approach is to add higher-dimensional cells so as to approximate the disk configuration spaces topologically, and keep track of topological invariants such as Betti numbers; Matthew Kahle and Robert MacPherson (personal communication) have constructed such a cell complex for the square configuration spaces. Any of these directions for inquiry might give new understanding of how the configuration spaces change shape according to the shape of the board, the shape of the pieces, and the density of the pieces on the board.

Acknowledgements I would like to thank Matt Kahle for suggesting the problem of estimating the diameters of square sliding-piece puzzles, Larry Guth and Jeremy Mason for additional helpful discussions, the editors Yuliy Baryshnikov and Shmuel Weinberger for pointing me toward references (Shnirelman 1987) and Demaine et al. (2018), and the referees for close reading that uncovered some mathematical glitches. I was supported by the Institute for Computational and Experimental Research in Mathematics (ICERM) while working on this project, and by the National Science Foundation under Award No. DMS 1802914 during the paper revisions.

Compliance with ethical standards

Conflict of interest The author states that there is no conflict of interest.

References

Alpert, H.: Restricting cohomology classes to disk and segment configuration spaces. Topol. Appl. **230**, 51–76 (2017)

Baryshnikov, Y., Bubenik, P., Kahle, M.: Min-type Morse theory for configuration spaces of hard spheres. Int. Math. Res. Not. **2014**(9), 2577–2592 (2013). https://doi.org/10.1093/imrn/rnt012

Böröczky, K.: Über stabile Kreis- und Kugelsysteme. Ann. Univ. Sci. Budapest. Eötvös Sect. Math. **7**, 79–82 (1964)

Carlsson, G., Gorham, J., Kahle, M., Mason, J.: Computational topology for configuration spaces of hard disks. Phys. Rev. E **85**, 011–303 (2012)

Chinta, R., Han, S.D., Yu, J.: Coordinating the motion of labeled discs with optimality guarantees under extreme density (2018)

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

Demaine, E.D., Fekete, S.P., Keldenich, P., Scheffer, C., Meijer, H.: Coordinated Motion Planning: Reconfiguring a Swarm of Labeled Robots with Bounded Stretch. In: 34th International Symposium on Computational Geometry (SoCG 2018), Leibniz International Proceedings in Informatics (LIPIcs). **99**, 29:1–29:15 (2018)

Diaconis, P.: The Markov chain Monte Carlo revolution. Bull. Am. Math. Soc. **46**(2), 179–205 (2009)

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

Knuth, D.E.: The art of computer programming. Volume 3. Addison-Wesley Publishing Company, Reading, Mass.-London-Don Mills, Ont. Sorting and searching, Addison-Wesley Series in Computer Science and Information Processing (1973)

Löwen, H.: Fun with hard spheres. In: Statistical Physics and Spatial Statistics (Wuppertal, 1999), Lecture Notes in Physics, **554**, pp. 295–331. Springer, Berlin (2000)

Shnirelman, A.I.: On the geometry of the group of diffeomorphisms and the dynamics of an ideal incompressible fluid. Math. USSR-Sb. **56**(1), 79–105 (1987)

Thompson, C.D., Kung, H.T.: Sorting on a mesh-connected parallel computer. Commun. ACM **20**(4), 263–271 (1977)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.