

Snake Paths in King and Knight Graphs

Nikolai Beluhov

Bulgaria

Email: nikolai.beluhov@gmail.com

Received: January 4, 2023, Accepted: February 24, 2023, Published: March 24, 2023

The authors: Released under the CC BY-ND license (International 4.0)

ABSTRACT: A snake path in a graph G is a path in G which is also an induced subgraph of G . For all n , we find the greatest length of a snake path in the $n \times n$ king graph and we give a complete description of the paths which attain this greatest length. The even and odd cases behave very differently. We also estimate the greatest length of a snake path or cycle in the $m \times n$ knight graph, for all m and n .

Keywords: Snake; Induced path; King graph; Knight graph; Leaper graph

2020 Mathematics Subject Classification: 05C30; 05C35; 05C38

1. Introduction

Let G be a simple graph. A *snake path* in G is a path in G which is also an induced subgraph of G . Equivalently, a path P in G is a snake path when, for all vertices u and v of P , we have that uv is an edge of G if and only if it is an edge of P . Intuitively, a snake path never comes into contact with itself.

Snake paths are also known as *induced paths* and *chordless paths*.

A *snake cycle* is defined similarly. Just like paths, snake cycles are alternatively called *induced cycles* and *chordless cycles*. Our focus will be mostly on snake paths, though we will touch upon snake cycles, too.

Given a graph G , some of the most natural questions we can ask about its snake paths are as follows:

Question A. *What is the greatest length of a snake path in G ?*

Note that we measure the length of a path by the number of edges that it traverses, rather than the number of vertices that it visits.

The answer to Question A coincides with the greatest diameter of an induced subgraph of G .

The greatest length of a snake path in G is also known as the *induced detour number* of G , and the greatest length of a snake cycle in G as the *induced circumference* of G .

In the special case when G is a hypercube graph, Question A is known as the *snake-in-the-box problem*, and its analog for cycles as the *coil-in-the-box problem*. Both of these problems have been studied extensively.

Question B. *What is the structure of the longest snake paths in G ?*

Question C. *How many longest snake paths are there in G ?*

Note that we formalize paths as subgraphs, rather than as sequences of vertices. Thus we do not distinguish between two paths which are reverses of one another.

We study Questions A–C for certain graphs G associated with chess pieces.

Given a chess piece F and a board A , it is natural to consider the graph whose vertices are the cells of A and whose edges correspond to all possible moves of F on A in the absence of other pieces on A . We proceed to formalize this notion for the king and the knight. Other chess pieces can be handled similarly.

To us, a *cell* is an ordered pair of integers. Or, equivalently, an integer point in the plane.

Let m and n be positive integers. A *board* A of size $m \times n$, with m rows and n columns, is a set of cells of the form $I \times J$, where I and J are integer intervals with $|I| = n$ and $|J| = m$. The *standard board* of size $m \times n$ has $I = [0; n - 1]$ and $J = [0; m - 1]$. Since all boards of the same size are translation copies of one another, sometimes we refer to “the” board of a certain size, meaning the standard board of that size.

Given a set of cells S , we define the *king graph* on S , denoted $\mathcal{G}(\mathfrak{K}, S)$, to be the graph on vertex set S where two distinct cells $a' = (x', y')$ and $a'' = (x'', y'')$ are joined by an edge if and only if $|x' - x''| \leq 1$ and

$|y' - y''| \leq 1$. Since all king graphs on boards of the same size are isomorphic, sometimes we refer to “the” king graph of a certain size, meaning the king graph on the standard board of that size. For convenience, we also use the notation $\mathcal{G}(\text{♔}, m \times n)$ for the king graph of size $m \times n$. Of course, $\mathcal{G}(\text{♔}, m \times n)$ can also be viewed as the strong product of two paths with m and n vertices, respectively.

The *knight graph* on S , denoted $\mathcal{G}(\text{♞}, S)$, is defined similarly, except that the adjacency condition becomes $\{|x' - x''|, |y' - y''|\} = \{1, 2\}$ instead.

Dawson, in problem 187 of [1], considers $\mathcal{G}(\text{♞}, 8 \times 8)$ and presents a snake path of length 31 as well as a snake cycle of length 32.

Knuth, in exercise 172 of [5], discusses the longest snake paths and cycles of various chess piece graphs in the context of algorithmic generation. In particular, he determines that there are 16 essentially distinct snake paths of the greatest length 31 and 6 essentially distinct snake cycles of the greatest length 31 in $\mathcal{G}(\text{♔}, 8 \times 8)$; as well as an essentially unique snake path of the greatest length 33 and 4 essentially distinct snake cycles of the greatest length 32 in $\mathcal{G}(\text{♞}, 8 \times 8)$. (Thus Dawson’s path was not optimal, but his cycle was.)

Our main results are as follows:

Theorem 1.1. *Let n be an even positive integer. Then the greatest length of a snake path in the king graph of size $n \times n$ is $n^2/2 - 1$. Furthermore, when $n \geq 6$, there are exactly $16n$ snake paths which attain this greatest length.*

This count includes rotations and reflections. In Section 3, we will see that with $n \geq 6$ the number of essentially distinct longest snake paths is $2n$ when $n/2$ is even and $2n + 1$ when $n/2$ is odd.

Over the course of the proof of Theorem 1.1, we will also give a complete description of these paths. Roughly speaking, each one of them is shaped like a spiral.

That the total number of paths is given by such a nice formula is most likely only a happy coincidence. Because of the overall structure of the proof, we should expect to see a linear function of n within each parity of $n/2$; however, there is no obvious reason *a priori* to expect these two functions to coincide, or their constant terms to vanish.

Theorem 1.2. *Let n be an odd positive integer. Then the greatest length of a snake path in the king graph of size $n \times n$ is $(n^2 - 1)/2$.*

Remarkably, the odd and even cases behave very differently. Despite the surface similarity between the upper bounds of Theorems 1.1 and 1.2, the former bound is straightforward while the latter one poses considerable difficulties.

In Section 5, we will add to Theorem 1.2 a complete description of the paths which attain the greatest length, stated in Theorem 5.1. However, the description is somewhat complicated and relies on a long series of preceding definitions. Thus we do not reproduce it in the introduction. The gist is that each stamp-folding permutation of $\lceil n/2 \rceil$ elements yields two families of longest snake paths which share the same overall shape but differ from one another by some tiny aberrations.

Theorem 5.1 does not imply an exact answer to Question C. In fact, because of the connection to the stamp-folding problem, it seems unlikely that such an answer would be feasible. Still, the theorem does yield some loose bounds. In particular, we will see (Proposition 5.2) that the logarithm of the number of longest snake paths grows as $\Theta(n)$, in stark contrast to the even case.

Theorems 1.1, 1.2, and 5.1 together completely resolve the questions of the greatest length of a snake path and the structure of the longest snake paths in king graphs on square boards.

Theorem 1.3. *Let m and n be positive integers. Then both the longest snake path and the longest snake cycle in the knight graph of size $m \times n$ are of length $mn/2 + \mathcal{O}(m + n)$.*

Compared to the treatment of Question A in Theorems 1.1 and 1.2, with Theorem 1.3 we do not attempt to obtain an exact answer and are content instead with an asymptotic estimate. On the bright side, this asymptotic estimate applies to cycles as well as paths, and it is valid on all rectangular boards.

It is worth noting that one step in our proof of Theorem 1.3 relies on computer help, and likely cannot be verified manually by a human mathematician. (Our proofs of Theorems 1.1, 1.2, and 5.1 are all human-friendly, though.)

The author obtained Theorems 1.1–7.3 in 2018 after being introduced to the subject by Knuth, in connection with the aforementioned exercise 172. Subsequently, Theorems 1.1, 1.3, and 7.1 were cited in a remark following the exercise’s solution. (Strictly speaking, at that time the author derived Theorem 1.1 in a form referring to the number of essentially distinct paths rather than the total number of paths. This is also how it was stated in [5].)

Then, in 2020, the author proposed Theorem 1.2, appropriately rephrased, as a mathematical olympiad problem for the Cyberspace Mathematical Competition. It was featured as problem 4 on day 1 of the contest. (The CMC was a one-off event intended to approximate the International Mathematical Olympiad. Perhaps the

problem’s difficulty was not an ideal match for the contest’s format; out of 553 participants from 75 countries, only two made substantial progress on it.)

Theorems 1.1–5.1 offer an interesting illustration of how the nature of Questions **A–C** can change when we vary the underlying graph. In the setting of Theorem 1.1, Question **A** is straightforward while Questions **B** and **C** are manageable. For Theorems 1.2 and 5.1, both Questions **A** and **B** become significantly more complicated, while a closed-form answer to Question **C** is likely out of reach. Finally, in the setting of Theorem 1.3, already with Question **A** it seems that an exact answer would be unfeasible, and even for our asymptotic estimate we find ourselves in need of machine help.

2. Preliminaries

Before we continue, let us briefly list some useful notations and observations.

Given two cells $a' = (x', y')$ and $a'' = (x'', y'')$, we write $a' + a''$ for the cell $(x' + x'', y' + y'')$. Given a cell a and a set of cells S , we write $a + S$ for the set of cells $\{a + b \mid b \in S\}$.

A *symmetry* of a board A is the restriction to A of an isometry of the plane which preserves A . Two objects defined with reference to A , such as two sets of cells on A or two graphs on A , are *essentially distinct* (relative to A) when they are distinct under the symmetries of A .

The *grid graph* on a set of cells S , denoted $\mathcal{G}(\square, S)$, is defined similarly to the king and knight graphs on S , except that the adjacency condition becomes $\{|x' - x''|, |y' - y''|\} = \{0, 1\}$ instead. Or, equivalently, cells a' and a'' must be at unit Euclidean distance from one another. Of course, $\mathcal{G}(\square, m \times n)$ can also be viewed as the Cartesian product of two paths with m and n vertices, respectively.

For the edge joining two vertices a and b in a graph G , we write either ab or $a-b$, whichever one reads better in the situation at hand.

We introduce shorthand notation for certain paths within grid and king graphs. Given two cells a and b of a grid or king graph G such that a and b are in the same row or column, we write $a \sim b$ for the path in G connecting a and b whose remaining cells are the ones between a and b in the corresponding row or column. For example, if $a = (x', y)$, $b = (x'', y)$, and $x' \leq x''$, then $a \sim b = (x', y) - (x' + 1, y) - (x' + 2, y) - \dots - (x'', y)$.

Let P be a path in some graph G . Given a vertex a of G , we write $a \in P$ for “ P visits a ”; given an edge e of G , we write $e \in P$ for “ P traverses e ”; and, given a path Q in G , we write $Q \subseteq P$ for “ Q is a subpath of P ”. We will not use these abbreviations too often, but the proofs of Lemmas 3.2 and 5.2 would be cumbersome to state without them.

Suppose, now, that P is a snake path in G . When $ab'b''$ is a three-cycle in G , we have that $a \in P$ implies $b'b'' \notin P$. Similarly, when b, c' , and c'' are three distinct neighbors of a in G , we have that $c'ac'' \subseteq P$ implies $b \notin P$. We will use these simple observations repeatedly throughout the paper.

A subset S of the vertices of G is *k-independent* if, in the induced subgraph of G on vertex set S , every vertex is of degree at most k . When $k = 2$, the induced subgraph itself is a *pseudosnake* of G . Clearly, the number of vertices in the longest snake path or cycle of G cannot exceed the number of vertices in its largest pseudosnake. When G is finite, we define its *pseudosnake density* to be the ratio of the greatest number of vertices in a pseudosnake of G to the total number of vertices in G .

3. King Graphs on Even Boards

In this section, we prove Theorem 1.1. Let n be an even positive integer with $n = 2k$, let A be the standard board of size $n \times n$, and let G be the king graph on A .

For completeness, with $n = 2$ there are 6 longest snake paths, of which 2 are essentially distinct; and with $n = 4$ there are 28 longest snake paths, of which 4 are essentially distinct. From now on, let $n \geq 6$.

The argument we give for the upper bound of Theorem 1.1 is not new; the special case $n = 8$ is in [5], and the general case does not pose any additional difficulties.

Proof of the optimization part of Theorem 1.1. For the upper bound, partition A into k^2 subboards of size 2×2 each. Since a snake path in G can visit at most two cells within each subboard, its length cannot exceed $2k^2 - 1$.




Figure 1




Figure 2

For the lower bound, let P_i be the path $(i, i) \sim (n - i - 2, i) - (n - i - 1, i + 1) \sim (n - i - 1, n - i - 2) - (n - i - 2, n - i - 1) \sim (i + 1, n - i - 1) - (i, n - i - 2) \sim (i, i + 3) - (i + 1, i + 2) - (i + 2, i + 2)$ for all even i with $0 \leq i \leq k - 3$. When k is even and $i = k - 2$, we define P_{k-2} in the same way, except that we stop at cell $(i, n - i - 2) = (k - 2, k)$; and, when k is odd and $i = k - 1$, we also define P_{k-1} to be the path $(k - 1, k - 1) - (k, k)$. Then the concatenation of these paths is a snake path in G of length $n^2/2 - 1$. \square

For example, Figure 1 shows $n = 8$.

The enumeration part of Theorem 1.1 will be somewhat more complicated.

Let B be the standard board of size $k \times k$ and let H be the grid graph on B .

For each cell b of B , let $\Phi(b)$ denote the subboard $2b + [0; 1]^2$ of A , of size 2×2 . (For convenience, if $b = (x, y)$, we also write simply $\Phi(x, y)$.) These subboards, taken over all cells b of B , form a partitioning of A .

Let P be a longest snake path in G . The proof of the optimization part of Theorem 1.1 shows that P visits exactly two cells within each subboard of A of the form $\Phi(b)$.

Suppose that an edge of P joins one cell of $\Phi(b')$ and one cell of $\Phi(b'')$, with $b' \neq b''$. We claim that b' and b'' are then neighbors in H .

Indeed, if not, then b' and b'' must be diagonally adjacent in H , without loss of generality with $b' + (1, 1) = b''$. So P contains a subpath of the form $c' - (2b' + (1, 1)) - 2b'' - c''$, where c' is in $\Phi(b')$ and c'' is in $\Phi(b'')$. Since P is a snake path, it follows that P cannot visit any cells in the set $2(b' + (1, 0)) + \{(0, 0), (0, 1), (1, 1)\}$. Consequently, P visits at most one cell of $\Phi(b' + (1, 0))$, a contradiction. (Figure 2.)

For each edge of P joining one cell of $\Phi(b')$ and one cell of $\Phi(b'')$, with $b' \neq b''$, take the edge $b'b''$ of H . Since P visits $\Phi(b)$ for all b , these edges form a Hamiltonian path in H . Denote this path by ϱ .

From this point on, our plan for the proof will be as follows: First we obtain a complete description of the structure of ϱ . We do this by means of a series of mostly local considerations, starting on the boundary of H and then working our way in. Once we are done, we determine what snake paths P in the original graph G are associated with each path ϱ .

We continue with the details.

We define the i -th frame of B , denoted F_i , to be the subset

$$F_i = [i; k - i - 1]^2 \setminus [i + 1; k - i - 2]^2$$

of B . Thus $F_0, F_1, \dots, F_{\lceil k/2 \rceil - 1}$ form a partitioning of B .

We denote the four corner cells of F_i by $a_i = (i, i)$, $b_i = (k - i - 1, i)$, $c_i = (k - i - 1, k - i - 1)$, and $d_i = (i, k - i - 1)$.




Figure 3

We say that F_i is of type I, II, III, IV, or V (relative to ϱ) when ϱ contains the following subpaths:

For type I, $a_i \sim b_i \sim c_i \sim d_i \sim (a_i + (0, 1)) - (a_i + (1, 1))$. (Figure 3(a).)

For type II, $a_i \sim b_i \sim c_i \sim d_i \sim (a_i + (0, 2)) - (a_i + (1, 2))$ and $(a_i + (0, 1)) - (a_i + (1, 1))$. (Figure 3(b).)

For type III, $a_i \sim b_i \sim c_i \sim (d_i + (1, 0)) - (d_i + (1, -1))$ and $d_i \sim (a_i + (0, 1)) - (a_i + (1, 1))$. (Figure 3(c).)

For type IV, $a_i \sim b_i \sim (c_i + (0, -1)) - (c_i + (-1, -1))$ and $c_i \sim d_i \sim (a_i + (0, 1)) - (a_i + (1, 1))$. (Figure 3(d).)

For type V, $a_i \sim (b_i + (-1, 0)) - (b_i + (-1, 1))$ and $b_i \sim c_i \sim d_i \sim (a_i + (0, 1)) - (a_i + (1, 1))$. (Figure 3(e).)

We say that ϱ itself is of type I when all of $F_0, F_1, \dots, F_{\lceil k/2 \rceil - 2}$ are of type I. Furthermore, let \mathfrak{T} be one of the symbols II, III, IV, and V, and let s be a nonnegative integer with $0 \leq s \leq \lceil k/2 \rceil - 2$. We say that ϱ is of type $\mathfrak{T}(s)$ when all of F_0, F_1, \dots, F_{s-1} are of type I and all of $F_s, F_{s+1}, \dots, F_{\lceil k/2 \rceil - 2}$ are of type \mathfrak{T} . When ϱ is of one of the $4\lceil k/2 \rceil - 3$ types we have just listed, we say that it is *regular*.

For example, Figure 4 shows the unique ϱ of type III(1) when $k = 7$.




Figure 4

Observe that, when k is even, ϱ cannot be of type IV(s), for any s , as $F_{\lceil k/2 \rceil - 2}$ being of type IV prevents ϱ from visiting all cells of $F_{\lceil k/2 \rceil - 1}$. Similarly, when k is odd, ϱ cannot be of type II(s), for any s , as the frame $F_{\lceil k/2 \rceil - 2}$ is too small to be of type II.

In all other cases, if k and the type of ϱ are fixed, there is a unique ϱ of that type, with one exception: When k is even, there are two paths ϱ of type I, differing by just one edge within the innermost frame $F_{\lceil k/2 \rceil - 1}$.

Lemma 3.1. *Suppose that ϱ makes a turn at cell b of B , for concreteness by means of $(b + (0, 1)) - b - (b + (1, 0))$. Then the two cells of $\Phi(b)$ in P are $2b + (0, 1)$ and $2b + (1, 0)$.*

Of course, similar claims hold for the other three possible turns at b as well.

Proof. Suppose, for the sake of contradiction, that $2b + (1, 1) \in P$. Since $(b + (0, 1)) - b - (b + (1, 0)) \subseteq \varrho$, we get that there are two cells a' and a'' with $a' \in \Phi(b + (0, 1))$, $a'' \in \Phi(b + (1, 0))$, and $a' - (2b + (1, 1)) - a'' \subseteq P$. However, then P cannot visit any cells of $\Phi(b)$ other than $2b + (1, 1)$, a contradiction.




Figure 5

Thus $2b + (1, 1) \notin P$. Since $(b + (0, 1)) - b - (b + (1, 0)) \subseteq \varrho$, it follows that $2b + (0, 1) \in P$ and $2b + (1, 0) \in P$, as needed. (Figure 5.) □

Lemma 3.2. *A symmetry of B maps ϱ onto a regular Hamiltonian path in H .*

Proof. First we consider the outermost frame F_0 of B . Observe that all edges of H within F_0 form a cycle E .

Suppose, for the sake of contradiction, that $(0, y) - (0, y + 1) \notin \varrho$ for some y with $2 \leq y \leq k - 4$.




Figure 6

Then P must miss both cells in at least one of the two pairs $\{(0, 2y + 1), (1, 2y + 1)\}$ and $\{(0, 2y + 2), (1, 2y + 2)\}$. Suppose, for concreteness, that this is true of the former pair; the other case is similar. (Figure 6.)

Since P visits two cells in $\Phi(0, y)$, we get that $(0, 2y) \in P$ and $(1, 2y) \in P$. So $(0, 2y - 1) \notin P$ and $(1, 2y - 1) \notin P$. Iterating this sequence of observations, we see that all of $(0, 2y - 2)$, $(1, 2y - 2)$, $(0, 2y - 4)$, and $(1, 2y - 4)$ are in P while all of $(0, 2y - 3)$, $(1, 2y - 3)$, $(0, 2y - 5)$, and $(1, 2y - 5)$ are not. (When $y = 2$, the lattermost couple of cells will lie outside of A . Of course, then they will be outside of P , too.)

Thus all three of $(0, 2y)$, $(0, 2y - 2)$, and $(0, 2y - 4)$ must be endpoints of P , a contradiction.

Consequently, $(0, y) - (0, y + 1) \in \varrho$ for all y with $2 \leq y \leq k - 4$. By symmetry, it follows that ϱ contains all edges of E except for, possibly, $(0, 0) - (0, 1)$, $(0, 1) - (0, 2)$, and their images under the symmetries of B .

On the other hand, at least one edge of E must be outside of ϱ .

If $(0, 0) - (0, 1) \notin \varrho$, then $(0, 0)$ is an endpoint of ϱ and $\Phi(0, 0)$ contains an endpoint of P .

If $(0, 1) - (0, 2) \notin \varrho$, then P must miss both cells in at least one of the two pairs $\{(0, 3), (1, 3)\}$ and $\{(0, 4), (1, 4)\}$. If the latter, then with $k \geq 5$ we arrive at a contradiction as before. If the former, then the

same reasoning as before shows that $\Phi(0, 0)$ and $\Phi(0, 1)$ must each contain an endpoint of P . When $k = 4$, the same conclusion holds up to reflection with respect to the horizontal axis of symmetry of A .

Of course, these observations apply also to all images of the edges $(0, 0) \text{---} (0, 1)$ and $(0, 1) \text{---} (0, 2)$ under the symmetries of B .

Since there are only two endpoints of P , it follows that ϱ cannot omit too many edges of E . We are left to consider the following cases, up to the symmetries of B :

Case 1. The only edge of E outside of ϱ is $(0, 0) \text{---} (0, 1)$. Then F_0 is of type I.

Case 2. The only edge of E outside of ϱ is $(0, 1) \text{---} (0, 2)$, and $k \geq 4$. Then F_0 is of type II.

Case 3. The only edges of E outside of ϱ are $(0, 0) \text{---} (0, 1)$ and $(0, k - 1) \text{---} (1, k - 1)$. Then F_0 is of type III.

Case 4. The only edges of E outside of ϱ are $(0, 0) \text{---} (0, 1)$ and $(k - 1, k - 2) \text{---} (k - 1, k - 1)$. Then F_0 is of type IV.

Case 5. The only edges of E outside of ϱ are $(0, 0) \text{---} (0, 1)$ and $(k - 2, 0) \text{---} (k - 1, 0)$. Then F_0 is of type V.

Cases 6–9. The only edges of E outside of ϱ are $(0, 0) \text{---} (0, 1)$ and one of $(0, 0) \text{---} (1, 0)$, $(0, k - 2) \text{---} (0, k - 1)$, $(k - 2, k - 1) \text{---} (k - 1, k - 1)$, and $(k - 1, 0) \text{---} (k - 1, 1)$. The first one of these cases cannot occur because then cell $(0, 0)$ becomes isolated in ϱ . The other three cannot occur because, in each one of them, part of the edges of E in ϱ form a subpath of ϱ which contains both endpoints of ϱ but does not coincide with ϱ .

With this, we have established that F_0 is of one of our types relative to ϱ , up to the symmetries of B .

We finish the proof by induction on k . Our base cases are $k = 3$ and $k = 4$, when there is nothing left to prove. For the induction step, suppose that $k \geq 5$ and that we have already settled the question on all smaller boards.

Let A^* be the concentric subboard of A of size $(n - 4) \times (n - 4)$ given by $A^* = [2; n - 3]^2$ and let G^* be the king graph on A^* .

When F_0 is of type I, the cells of P outside of A^* form a subpath of P containing one endpoint of P . When F_0 is of one of the remaining four types, the cells of P outside of A^* form two subpaths of P containing the two endpoints of P . Either way, we see that the restriction P^* of P to A^* is a subpath of P . Since P is a snake path in G , also P^* is a snake path in G^* . Furthermore, P^* will be of the greatest possible length within G^* .

Let B^* be the concentric subboard of B of size $(k - 2) \times (k - 2)$ given by $B^* = [1; k - 2]^2 = B \setminus F_0$ and let H^* be the grid graph on B^* . By the same reasoning as above, the restriction ϱ^* of ϱ to B^* is a subpath of ϱ . Furthermore, ϱ^* is a Hamiltonian path in H^* , and ϱ^* and P^* are related in the same way as ϱ and P , in the sense that ϱ^* visits the cells of B^* in the same order as P^* visits their corresponding 2×2 subboards of A^* .

By the induction hypothesis, either ϱ^* or an image of it under a symmetry of B^* must be regular relative to B^* .

We consider five cases for the type of F_0 .

Case 1. F_0 is of type I.

Then $(0, 1) \text{---} (1, 1) \in \varrho$ and $(1, 1)$ is an endpoint of ϱ^* .

Suppose, for the sake of contradiction, that $(1, 1) \text{---} (1, 2) \in \varrho$ and ϱ makes a turn at $(1, 1)$.

If ϱ also makes a turn at $(2, 1)$, then by Lemma 3.1 we get that $(3, 3) \in P$ and $(4, 3) \in P$, in contradiction with $(1, 1) \text{---} (2, 1) \notin \varrho$.

Otherwise, if ϱ does not make a turn at $(2, 1)$, then $(2, 1)$ is an endpoint of ϱ^* . By the induction hypothesis, it follows that F_1 satisfies the reflection, with respect to the line $x = y$, of the conditions defining type II. Thus ϱ makes a turn at $(3, 1)$, and by Lemma 3.1 we get that $(2, 2) \text{---} (3, 3) \text{---} a' \subseteq P$ and $(7, 2) \text{---} (6, 3) \text{---} a'' \subseteq P$ for two cells a' and a'' with $a' \in \Phi(1, 2)$ and $a'' \in \Phi(3, 2)$. Consequently, P cannot visit any cells in $\Phi(2, 1)$, another contradiction.

We conclude that $(1, 1) \text{---} (1, 2) \notin \varrho$ and $(1, 1) \text{---} (2, 1) \in \varrho$. By the induction hypothesis, it follows that ϱ^* is regular. (As in, it is ϱ^* itself that is regular, rather than some image of it under a symmetry of B^* .) Then ϱ is regular as well, and of the same type as ϱ^* .

Case 2. F_0 is of type II.

Then $(0, 1) \text{---} (1, 1) \in \varrho$, $(0, 2) \text{---} (1, 2) \in \varrho$, and $(1, 1)$ and $(1, 2)$ are the two endpoints of ϱ^* . By the induction hypothesis, it follows that ϱ^* is of type II(0), and so is ϱ as well.

Case 3. F_0 is of type III.

Then $(0, 1) \text{---} (1, 1) \in \varrho$, $(1, k - 1) \text{---} (1, k - 2) \in \varrho$, and $(1, 1)$ and $(1, k - 2)$ are the two endpoints of ϱ^* . By the induction hypothesis, it follows that either ϱ^* is of type III(0), or its reflection with respect to the horizontal axis of symmetry of B^* is. (Here, we take into account the fact that types III(0) and V(0) are related by 90° rotation.) If the latter, then ϱ makes a turn at both cells $(1, 1)$ and $(2, 1)$, and we get a contradiction as in Case 1. Thus ϱ^* is of type III(0), and so is ϱ as well.

Case 4. F_0 is of type IV.

Then $(0, 1) \text{---} (1, 1) \in \varrho$, $(k - 1, k - 2) \text{---} (k - 2, k - 2) \in \varrho$, and $(1, 1)$ and $(k - 2, k - 2)$ are the two endpoints of ϱ^* . By the induction hypothesis, it follows that either ϱ^* is of type IV(0), or its reflection with respect to the

line $x = y$ is. From here, the analysis continues as in Case 3, and in the end we obtain that both of ϱ^* and ϱ are of type IV(0).

Case 5. F_0 is of type V.

This case is analogous to Case 3, and in it both of ϱ^* and ϱ are of type V(0).

With this, the induction step is complete. □

Lemma 3.3. *Let ϱ be a regular Hamiltonian path in H . Then the number of longest snake paths P in G associated with ϱ is as shown in Table 1.*

type of ϱ	k even	k odd
I	$8 + 1$	8
II(s)	6	—
III(s)	1	1
IV(s)	—	6
V(s)	1	1

Table 1

These paths are pairwise distinct under the symmetries of A , with the following exceptions: The unique P with ϱ of type III(0) and the unique P with ϱ of type V(0) are related by 90° rotation; and, when k is odd, two pairs of paths P with ϱ of type IV(0) are related by central symmetry.

(The entry of Table 1 for k even and ϱ of type I includes one summand for each regular path ϱ of that type. For the parametrized types, the total count does not depend on the value of the parameter.)

Proof. First we strengthen the claim as follows: Let b be an arbitrary cell of F_0 . Then, for all k with $k \geq 5$:

(i) When ϱ goes straight through b , the two cells of $\Phi(b)$ in P are the ones on the boundary of A , in the union of row 0, row $n - 1$, column 0, and column $n - 1$ of A ; and

(ii) When b is an endpoint of ϱ , by the definitions of our types we get that b is one of $(0, 0)$, $(0, 1)$, $(k - 1, 0)$, $(k - 1, k - 1)$, and $(0, k - 1)$. Then the two cells of $\Phi(b)$ in P are, respectively, the ones in row 0, row 2, column $n - 1$, row $n - 1$, and column 0 of A .

For the proof, we work by induction on k .

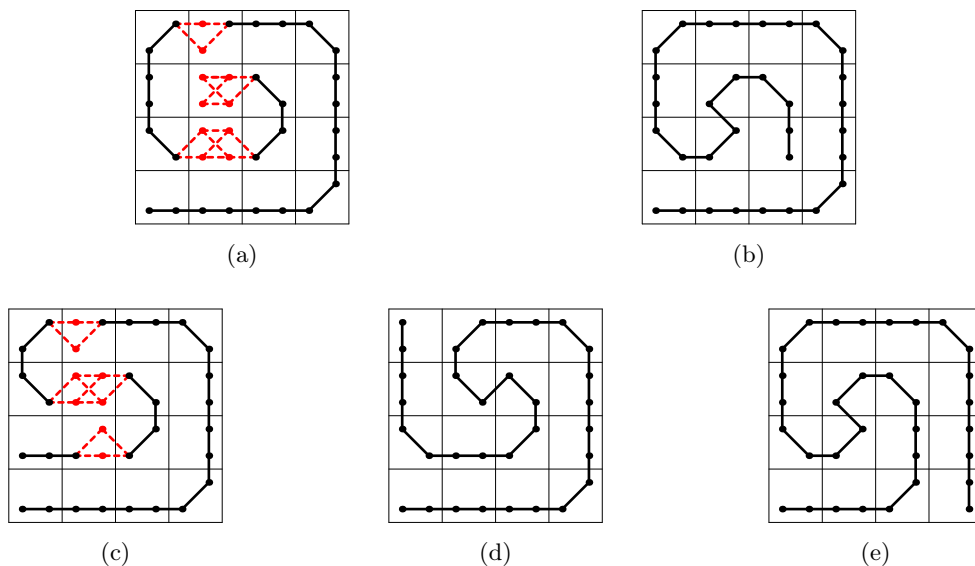


Figure 7

Our base cases are $k = 3$ and $k = 4$, when the strengthening is irrelevant and the original claim follows by direct examination of all cases. These are shown in Figures 7 and 8. For each situation, we depict the cells and edges which are forced to belong to P in black, and the optional cells and edges of P in red. Figures 7(a)–(e) correspond to the two paths ϱ of type I and the paths ϱ of types II(0), III(0), and V(0), respectively, while Figures 8(a)–(d) correspond to the paths ϱ of types I, III(0), IV(0), and V(0), respectively.

For the induction step, suppose that $k \geq 5$ and that we have already settled the question on all smaller boards. Define A^* , B^* , G^* , H^* , P^* , and ϱ^* as in the proof of Lemma 3.2.

We begin with parts (i) and (ii) of the strengthening.




Figure 8

For (i), consider a cell b of F_0 such that ϱ goes straight through b , using two edges on H within F_0 . Let c be the unique neighbor of b in H which belongs to F_1 , let b' and b'' be the two cells of $\Phi(b)$ which are adjacent in G to a cell in $\Phi(c)$, and, conversely, let c' and c'' be the two cells of $\Phi(c)$ which are adjacent in G to a cell in $\Phi(b)$.

Then at least one of c' and c'' must be in P . Indeed, when ϱ^* makes a turn at c , this follows by Lemma 3.1. Otherwise, when ϱ^* goes straight through c using two edges of H within F_1 , or when c is an endpoint of ϱ^* , it follows by the catalog in Figures 7 and 8 when $k = 5$ or $k = 6$, and by the induction hypothesis when $k \geq 7$.

Since $bc \notin \varrho$ and at least one of c' and c'' is in P , we conclude that $b' \notin P$ and $b'' \notin P$. So the other two cells of $\Phi(b)$ must be in P , confirming (i).

For (ii), consider a cell b of F_0 such that b is an endpoint of ϱ . When b is $(0, 0)$, $(0, 1)$, $(k - 1, 0)$, $(k - 1, k - 1)$, or $(0, k - 1)$, define c to be $(0, 1)$, $(0, 2)$, $(k - 2, 0)$, $(k - 1, k - 2)$, or $(1, k - 1)$, respectively, and also define cells b' , b'' , c' , and c'' relative to b and c as in our treatment of (i).

By the definitions of our types, $bc \notin \varrho$. Furthermore, with one exception, to be considered shortly, ϱ makes a turn at c . By Lemma 3.1, it follows that exactly one of c' and c'' is in P . From here, as in our treatment of (i), $bc \notin \varrho$ implies $b' \notin P$ and $b'' \notin P$, and so the other two cells of $\Phi(b)$ must be in P , as needed. The unique exception is when F_0 is of type II, $b = (0, 0)$, and $c = (0, 1)$. However, then $c' \in P$ and $c'' \in P$ by what we just proved applied to the endpoint $(0, 1)$ of ϱ , and once again (ii) is confirmed.

With this, we have established both parts (i) and (ii) of the strengthening.

For the original claim, observe that the type of ϱ determines the type of F_0 uniquely. We are left to show that the type of F_0 , and the subpath P^* , determine P uniquely. We only need to look at the cells of P in $A \setminus A^*$, that is, in the union of the subboards $\Phi(b)$ of A with $b \in F_0$.

Let, then, b be an arbitrary cell of F_0 . When ϱ makes a turn at b , the cells of P in $\Phi(b)$ are uniquely determined by Lemma 3.1. Otherwise, when ϱ goes straight through b , or when b is an endpoint of ϱ , the desired uniqueness follows by parts (i) and (ii) of the strengthening, respectively. The induction step is complete. \square

Proof of Theorem 1.1. We get the number of essentially distinct paths by Lemmas 3.2 and 3.3, summing over all possible types of ϱ and then subtracting out the duplicates specified in the statement of Lemma 3.3.

To convert this into the total number of paths, we analyze symmetries.

When k is even, all of the P 's are asymmetric.

Otherwise, when k is odd, two of the P 's which correspond to the unique ϱ of type IV(0) are preserved under central symmetry, but not under any other symmetries of A ; and all other P 's corresponding to regular ϱ 's are asymmetric. \square

Note that this argument also confirms our remark in the introduction regarding the number of essentially distinct longest snake paths in G .

The proof of Theorem 1.1 shows that, roughly speaking, every longest snake path P in G behaves as follows: Outside of some concentric square subboard of A , it is shaped as a simple spiral; and then, within that subboard, it is shaped as a double spiral instead. One endpoint of P lies on the boundary of A , and the other one lies on the boundary between the two spirals.

4. King Graphs on Odd Boards I

In this section, we prove Theorem 1.2. Let n be an odd positive integer with $n = 2k - 1$, let A be the standard square board of side n , and let G be the king graph on A .

Proof of the lower bound for Theorem 1.2. Consider the set of all cells of A of the form (x, y) with $1 \leq x \leq n - 2$ and y even. The king graph on it is the disjoint union of k paths. To obtain the vertex set of a single snake path in G , we add in also the cells $(0, 0)$; $(0, n - 1)$ if k is even and $(n - 1, n - 1)$ if k is odd; $(0, y)$ in A with $y \equiv 3 \pmod{4}$; and $(n - 1, y)$ in A with $y \equiv 1 \pmod{4}$. \square




Figure 9

For example, Figure 9 shows $n = 9$.

We obtained the upper bound of Theorem 1.1 by summing over some subsets of A such that, for every snake path P , the part of P within each subset must be small. Our approach to the upper bound of Theorem 1.2 will follow a similar strategy, albeit with significant complications. Instead of subsets of A , we sum over subgraphs of G . Furthermore, our notion of smallness will be somewhat unusual: We consider the total number of certain cells and edges of P within each subgraph.

We call a cell (x, y) of A *even* when both of x and y are even, and *odd* when both of x and y are odd. We also call an edge of G *regular* when it is not incident with an odd cell.



Figure 10

Given an even cell $a = (z, z)$ of A with $z \leq k - 2$, we write $\bar{\mathcal{X}}(a)$ for the subgraph of G with vertices $a + [0; 1]^2$ whose edges join a , $a + (0, 1)$, and $a + (1, 0)$ pairwise. Thus $\bar{\mathcal{X}}(a)$ contains four cells, one of which is odd, and three edges, all of which are regular. (Figure 10(a).)

For every symmetry π of A , if $b = \pi(a)$, then we also define $\bar{\mathcal{X}}(b) = \pi(\bar{\mathcal{X}}(a))$. We call each subgraph of G of this form a *little block*.

Given an even cell $a = (x, y)$ of A with $x > y$ and $x + y \leq n - 3$, we write $\bar{\mathcal{X}}(a)$ for the subgraph of G with vertices $a + [-1; 1] \times [0; 2]$ whose edges join a to $a + (-1, 0)$ and $a + (1, 0)$ as well as $a + (0, 1)$ to all elements of the set $a + [-1; 1] \times \{0, 2\}$. Thus $\bar{\mathcal{X}}(a)$ contains nine cells, two of which are odd, and eight edges, all of which are regular. (Figure 10(b).)

For every symmetry π of A , if $b = \pi(a)$, then we also define $\bar{\mathcal{X}}(b) = \pi(\bar{\mathcal{X}}(a))$. We call each subgraph of G of this form a *large block*.




Figure 11

For example, Figure 11 shows $n = 11$ and $n = 13$. (The coloring is only for clarity.)

Observe that every odd cell belongs to at least two blocks and every regular edge belongs to at least one block.

Consider an arbitrary snake path P in G . For every block $\bar{\mathcal{X}}$, let $w_{\text{Cell}}(\bar{\mathcal{X}})$ be the number of odd cells of P in $\bar{\mathcal{X}}$, let $w_{\text{Edge}}(\bar{\mathcal{X}})$ be the number of regular edges of P in $\bar{\mathcal{X}}$, and let $w(\bar{\mathcal{X}}) = w_{\text{Cell}}(\bar{\mathcal{X}}) + w_{\text{Edge}}(\bar{\mathcal{X}})$.

Lemma 4.1. *When $\bar{\mathcal{X}}$ is a little block, $w(\bar{\mathcal{X}}) \leq 1$. Otherwise, when $\bar{\mathcal{X}}$ is a large block, $w(\bar{\mathcal{X}}) \leq 2$.*

Proof. By direct examination of all cases. □

Proof of the upper bound for Theorem 1.2. Let $w_{\text{Cell}}(P)$ be the number of odd cells in P and let $w_{\text{Edge}}(P)$ be the number of regular edges in P . Then the length of P is bounded from above by $2w_{\text{Cell}}(P) + w_{\text{Edge}}(P)$. We sum the inequalities of Lemma 4.1 over all blocks \mathfrak{X} , and we obtain that the latter expression cannot exceed $(n^2 - 1)/2$. \square

5. King Graphs on Odd Boards II

In this section, we give a complete description of the longest snake paths in king graphs on odd square boards. We use the same notations as in Section 4.

We begin with brief overviews of two relevant topics.

Let s be a positive integer and let σ be a permutation of $[0; s - 1]$. For each i with $0 \leq i \leq s - 2$, draw a semicircle with endpoints $(0, \sigma^{-1}(i))$ and $(0, \sigma^{-1}(i + 1))$ which lies on the left of the coordinate axis Oy when i is even and on its right otherwise, when i is odd. The union of all such semicircles is a curve κ in the plane with endpoints $(0, \sigma^{-1}(0))$ and $(0, \sigma^{-1}(s - 1))$. When this curve does not intersect itself, σ is a *stamp-folding permutation*.

For example, Figure 12 shows κ when σ is the permutation 1, 0, 2, 7, 4, 5, 6, 3.

Combinatorially, σ is a stamp-folding permutation if and only if there are no i and j of the same parity with $0 \leq i \leq s - 2$ and $0 \leq j \leq s - 2$ such that exactly one of i and $i + 1$ lies between j and $j + 1$ in σ .




Figure 12




Figure 13

The intuition is as follows: Imagine a paper strip of size $1 \times s$ formed out of s stamps of size 1×1 each. We fold this strip along the perforations between stamps so that all stamps come to lie on top of one another. Then σ is a stamp-folding permutation if and only if it can be obtained as the permutation of the stamps within the resulting stack. The points of κ on the coordinate axis Oy correspond to the s stamps, and the semicircles of κ correspond to the $s - 1$ creases between stamps.

Stamp-folding permutations have been studied extensively.

Consider a Hamiltonian path in the grid graph Γ of size $s \times s$. The smallest number of turns that such a path can make is $2s - 2$. [3] We proceed to review some properties of the paths which attain this minimum. [7]

(The author noticed the connection between stamp-folding permutations and fewest-turn Hamiltonian paths when he obtained Theorems 1.2 and 5.1. Independently, [7] was published before the present work was written.)

Let α be a fewest-turn Hamiltonian path in Γ . We partition α into $2s - 1$ subpaths at the $2s - 2$ cells where it makes a turn. (Each cell with a turn belongs to two subpaths.) We call these subpaths the *segments* of α . Thus in each segment of α either all edges are horizontal or all edges are vertical, and segments of these two types alternate.

Let us call α *mostly-horizontal* when it consists of s horizontal segments and $s - 1$ vertical segments, and *mostly-vertical* otherwise, when it is the other way around. Suppose, for concreteness, that α is mostly-horizontal.

Then every row of the board contains exactly one horizontal segment of α . Orient α arbitrarily, and then number its segments from 0 to $s - 1$ in the order in which they occur along α . When we assign to each row of the board the number of its horizontal segment of α , we obtain a stamp-folding permutation. Conversely, every stamp-folding permutation corresponds in this way to exactly two oriented mostly-horizontal fewest-turn Hamiltonian paths. The two are reflections of one another with respect to the vertical axis of symmetry of the board.

Explicitly, the correspondence is as follows: For each element i of $[0; s - 1]$, let $\omega_{\text{Left}}(i)$ be the number of even nonnegative integers j such that i lies between j and $j + 1$ in σ , and define $\omega_{\text{Right}}(i)$ similarly, but with j odd. In one of the two oriented paths associated with σ , the path's i -th horizontal segment goes from $(s - \omega_{\text{Right}}(i) - 1, \sigma^{-1}(i))$ to $(\omega_{\text{Left}}(i), \sigma^{-1}(i))$ for all even i , and it goes in the opposite direction for all odd i . The second oriented path associated with σ can be obtained by reflection.

For example, Figure 13 shows this for the stamp-folding permutation of Figure 12.

One corollary of this connection is that, for all s with $s \geq 2$, the number of fewest-turn Hamiltonian paths in the grid graph of size $s \times s$ is twice the number of stamp-folding permutations of s elements.

This concludes the two overviews, and we return to our main topic.

Let H be the grid graph of size $k \times k$ and let ρ be a Hamiltonian path in it. (Recall from Section 4 that $k = \lceil n/2 \rceil$.)




Figure 14

We say that the cycle $a'a''b'b''$ in H is *free* (relative to ϱ) when $a'a''$ is an edge of ϱ ; the other three edges of the cycle are outside of ϱ ; both of a' and a'' are cells that ϱ goes straight through (so that, if $a' = (c' + a'')/2$ and $a'' = (a' + c'')/2$, then $c' \sim c'' \subseteq \varrho$); and both of b' and b'' are cells where ϱ makes a turn.

We proceed to associate ϱ with certain paths in G . Figure 14(a) shows one example of a Hamiltonian path ϱ in H , and Figures 14(b)–(d) track the series of definitions given below.

For each edge $a'a''$ of ϱ , we take the path $2a' - (a' + a'') - 2a''$ in G . We denote the concatenation of these paths by $\varphi(\varrho)$. Intuitively, this operation scales ϱ up by a factor of two. (Figure 14(b).)

Observe that the length of $\varphi(\varrho)$ will be twice the length of ϱ , namely $2(k^2 - 1)$.

For each turn $a'ba''$ in ϱ , let us delete the subpath $(a' + b) - 2b - (b + a'')$ from $\varphi(\varrho)$, and let us replace it with the edge $(a' + b) - (b + a'')$. We denote the resulting path in G by $\psi(\varrho)$. Intuitively, this operation smooths down the sharp turns in $\varphi(\varrho)$. (Figure 14(c).)

Let t be the number of turns in ϱ . Then the length of $\psi(\varrho)$ will be $2(k^2 - 1) - t$.

Finally, for each edge $a'a''$ of ϱ which is in a free cycle, either we do nothing or, optionally, we choose one free cycle $a'a''b'b''$ which includes $a'a''$, we set $c = (a' + a'' + b' + b'')/4$, we delete the subpath $2a' - (a' + a'') - 2a''$ from $\psi(\varrho)$, and we replace it with the subpath $2a' - 2c - 2a''$. We call a path in G which can be obtained in this way a *lift* of ϱ . Intuitively, this operation introduces some tiny aberrations in $\psi(\varrho)$. (Figure 14(d).)

In particular, $\psi(\varrho)$ is also a lift of ϱ , namely the one in which we have selected the do-nothing option everywhere.

Observe that every lift of ϱ is a snake path in G . Furthermore, all lifts of ϱ are of the same length, namely $2(k^2 - 1) - t$.

We are ready to state and prove our structure theorem for the longest snake paths in G .

Theorem 5.1. *Let n be an odd positive integer with $n = 2k - 1$. Then every lift of a fewest-turn Hamiltonian path in the grid graph of size $k \times k$ is a longest snake path in the king graph of size $n \times n$. Conversely, every longest snake path in the king graph of size $n \times n$ can be obtained uniquely as a lift of some fewest-turn Hamiltonian path in the grid graph of size $k \times k$.*

Theorem 5.1 yields the following recipe for the generation of all longest snake paths in G : First, we generate all stamp-folding permutations of k elements. Then we convert each stamp-folding permutation into four oriented fewest-turn Hamiltonian paths in H , two mostly-horizontal ones and two mostly-vertical ones. We forget about the orientations and discard the duplicates. Finally, for each fewest-turn Hamiltonian path in H , we identify the corresponding free cycles, and we generate all of its lifts.

(We said in the introduction that each stamp-folding permutation yields two families of longest snake paths. Strictly speaking, the reality is that each permutation yields four families, and each quadruple of families is obtained in this way twice, out of two permutations σ' and σ'' related by $\sigma'(i) + \sigma''(i) = k - 1$ for all i . However, it is straightforward to extract a one-to-two mapping from the two-to-four one.)

Before we go on to the proof of Theorem 5.1, let us briefly discuss the aberrations.

Proposition 5.1. *Let ϱ be a fewest-turn Hamiltonian path in H and let f be the number of its corresponding free cycles. Then ϱ yields exactly 2^f lifts. Furthermore, $f \leq \max\{0, k - 5\}$, and for all k this bound is attained by some ϱ .*

Proof. Suppose, for concreteness, that ρ is mostly-horizontal. Let $a'a''$ be an edge of ρ and let $a'a''b'b''$ be a free cycle which includes $a'a''$.

Suppose, for the sake of contradiction, that edge $a'a''$ is horizontal. Then, since ρ makes a turn at both of b' and b'' , the edges of ρ in the row of b' and b'' cannot form one contiguous subpath of ρ . This contradicts the fact that every row of the board contains exactly one horizontal segment of ρ .

Thus edge $a'a''$ must be vertical.

Suppose now, for the sake of contradiction, that edge $a'a''$ is in a second free cycle $a'a''c'c''$. Since ρ must make a turn at both of c' and c'' as well, it follows that the edges of ρ in the row of b' and c' cannot form one contiguous subpath of ρ . From here, we get a contradiction as before.

Thus no edge of ρ can be in two distinct free cycles. Consequently, in the setting of Theorem 5.1, when we construct a lift of ρ , we never have to choose between two free cycles which include the same edge $a'a''$.

We are left to show that $f \leq \max\{0, k - 5\}$ and the bound is attained. We handle the cases when $k \leq 6$ directly, and from now on we assume that $k \geq 7$.

Each free cycle contains two cells where ρ makes a turn. Conversely, each such cell is in at most one free cycle.

On the other hand, ρ makes $2k - 2$ turns altogether. However, a turn cell in an outermost column of B cannot be in a free cycle. The lowermost turn cell and the topmost turn cell in a non-outermost column of B cannot be in free cycles, either.

When at most one non-outermost column of B contains turns of ρ , since no turn cells outside of that column can be in free cycles, it follows that $2f \leq k - 2$. Then $f \leq k - 5$ by virtue of $k \geq 7$. Otherwise, when at least two non-outermost columns of B contain turns of ρ , it follows that at least eight turn cells are not in free cycles, and so $2f \leq 2k - 10$.

The bound is attained, for example, when ρ corresponds to the stamp-folding permutation $0, 2, 3, 4, \dots, k - 1, 1$. (The path in Figure 14(a) is of this form.) □

In the setting of the proof of Proposition 5.1, our observation that edge $a'a''$ must be vertical allows us to characterize the free cycles corresponding to ρ in terms of the underlying stamp-folding permutation σ , as follows:

Consider the ordered pairs (ε, i) with $\varepsilon \in \{-1, 1\}$ and $0 \leq i \leq k - 2$ such that, in σ , both of $\sigma(i)$ and $\sigma(i + 1)$ lie between $\sigma(i) + \varepsilon(-1)^{\sigma(i)}$ and $\sigma(i + 1) + \varepsilon(-1)^{\sigma(i+1)}$, and, additionally, there is some j with $0 \leq j \leq k - 1$ such that all four of these lie between j and $j + \varepsilon(-1)^j$.

Each such ordered pair yields a free cycle where edge $a'a''$ joins rows i and $i + 1$. When $\varepsilon = 1$, cells a' and a'' are in column $\omega_{\text{Left}}(\sigma(i)) - 1$ and cells b' and b'' are on the right of them. Otherwise, when $\varepsilon = -1$, cells a' and a'' are in column $k - \omega_{\text{Right}}(\sigma(i))$ and cells b' and b'' are on their left. Furthermore, this accounts for all free cycles corresponding to ρ .

Let $\mathcal{N}_{\text{Stamp}}(k)$ be the number of stamp-folding permutations of k elements and let $\mathcal{N}_{\text{King}}(n)$ be the number of longest snake paths in the king graph of size $n \times n$. Proposition 5.1 yields some loose bounds on $\mathcal{N}_{\text{King}}(n)$ in terms of $\mathcal{N}_{\text{Stamp}}(k)$.

Proposition 5.2. *Let n be an odd positive integer with $n = 2k - 1$. Then $\mathcal{N}_{\text{King}}(n)$ lies strictly between $2\mathcal{N}_{\text{Stamp}}(k)$ and $2^{k-4}\mathcal{N}_{\text{Stamp}}(k)$ for all n with $n \geq 11$, and $\mathcal{N}_{\text{King}}(n) = 2\mathcal{N}_{\text{Stamp}}(k)$ when $3 \leq n \leq 9$. Thus, in particular, $\log \mathcal{N}_{\text{King}}(n) = \Theta(n)$.*

Proof. The first part follows by Theorem 5.1 and Proposition 5.1. (Strictly speaking, we should also note that not all fewest-turn Hamiltonian paths in H attain the greatest number of free cycles when $n \geq 11$.) The second part is a corollary of the first part and the well-known asymptotic estimate $\log \mathcal{N}_{\text{Stamp}}(n) = \Theta(n)$. [2] □

We continue with the proof of Theorem 5.1. Let P be a longest snake path in G .

Here is a quick roadmap: Clearly, P must attain exact equality in all inequalities from the proof of Theorem 1.2. We examine all blocks of G from this point of view, one by one, in a certain order, and we see that P must satisfy certain purely local constraints. These constraints allow us to conclude that P must be a lift.

The details, however, are somewhat technical.

We define an even cell a of A to be *nice* (relative to P) when either P visits a ; or, else, P does not visit a but it traverses exactly one edge of G between the four cells in the set $a + \{(1, 0), (0, 1), (-1, 0), (0, -1)\}$. (Note that, for some a , some of these cells might be outside of A .)

Lemma 5.1. *Suppose that all even cells of A are nice and P does not visit any odd cells. Then there is a Hamiltonian path ρ in H such that $P = \psi(\rho)$.*

Proof. Let us call an edge of G *short* when the Euclidean distance between its endpoints is unity, and *long* otherwise, when it is $\sqrt{2}$.

For each long edge $a'a''$ of P , we do the following: Since P does not visit any odd cells, there is a unique even cell b such that both of $a'b$ and ba'' are edges of G . We delete edge $a'a''$ from P , and we replace it with these two edges.

Since all even cells of A are nice and P is a snake path, the result will be a path in G which contains only short edges, which visits all even cells of A , and which does not visit any odd cells. Denote this path by Q .

Let c be an endpoint of Q . Suppose, for the sake of contradiction, that c is not an even cell. Since Q does not visit any odd cells, there are exactly two even cells d' and d'' of A adjacent to c in G . Let cd' be the unique edge of Q incident with c . Observe that c is also an endpoint of P . But then d'' cannot be nice because P is a snake path, and we arrive at a contradiction.

So, to our previous observations about Q , we can add the fact that both of its endpoints are even cells. Consequently, Q is of the form $Q = \varphi(\rho)$ for some Hamiltonian path ρ in H , and $P = \psi(\rho)$, as needed. \square

We define a *rectifiable aberration* in P to be a subpath of P of the form $b'a'ba''b''$ such that a' and a'' are two even cells in the same row or column; the cell $c = (a' + a'')/2$ is a common neighbor of a' and a'' in G ; the cells b' and b'' satisfy $a' = (b' + c)/2$ and $a'' = (c + b'')/2$; and $b \neq c$. (Figure 15.)




Figure 15

To *rectify* a rectifiable aberration, we delete the subpath $a'ba''$ from P , and we replace it with the subpath $a'ca''$. The result will be a new snake path in G of the same length as P . This follows because P being a snake path implies that the only neighbors of c in G which P visits are a' , b , and a'' .

Lemma 5.2. *Suppose that P does not contain any rectifiable aberrations. Then all even cells of A are nice and P does not visit any odd cells.*

Proof. Since P is a longest snake path in G , it must attain exact equality in all inequalities from the proof of Theorem 1.2. Thus:

- (i) A cell of A of the form (z, z) with z odd cannot be in P because it is an odd cell which is in more than two blocks. The same goes for the images of these cells under the symmetries of A ;
- (ii) An edge of G of the form $(z, z + 1) - (z + 1, z)$ with z odd and $z \leq k - 2$ cannot be in P because it is a regular edge which is in more than one block. The same goes for the images of these edges under the symmetries of A ; and
- (iii) Every block \bar{X} must attain exact equality in Lemma 4.1, so that $w(\bar{X}) = 1$ when \bar{X} is a little block and $w(\bar{X}) = 2$ otherwise, when \bar{X} is a large block.

For all odd positive integers s with $1 \leq s \leq n$, we write A_s for the concentric subboard of A of size $s \times s$ given by $A_s = [[\lfloor n/2 \rfloor - \lfloor s/2 \rfloor; \lfloor n/2 \rfloor + \lfloor s/2 \rfloor]^2$.

We will show by induction on s that all even cells of A_s are nice and all odd cells of A_s are outside of P .

Our base case is $s = 1$. Then A_1 consists of a single cell, namely $(k - 1, k - 1)$. Denote this cell by o .

When k is even, o is an odd cell, and $o \notin P$ by (i).




Figure 16

When k is odd, o is an even cell. (Figure 16.) Suppose, for the sake of contradiction, that $k \geq 3$ and o is not nice. Then $o \notin P$. By (i), $o + (1, 1) \notin P$, and similarly for the images of this cell under the symmetries of A . By (ii), $(o + (1, 0)) - (o + (0, 1)) \notin P$, and similarly for the images of this edge under the symmetries of A . Thus (iii) for $\bar{X}(o + (0, -2))$ implies $o + (-1, -2) \in P$, $o + (1, -2) \in P$, and either $o + (0, -1) \in P$ or $o + (0, -2) \in P$; and similarly for the images of this block under the symmetries of P . However, the twelve cells we just concluded must be in P are the vertices of a cycle in G , and we arrive at a contradiction.

This settles the base case.

For the induction step, let $s \geq 3$ and suppose that we have already established the desired result for A_{s-2} . Let $a = (x, y)$ be an arbitrary cell of $A_s \setminus A_{s-2}$. By symmetry, we can assume without loss of generality that $\lfloor n/2 \rfloor - \lfloor s/2 \rfloor \leq x \leq \lfloor n/2 \rfloor$ and $y = \lfloor n/2 \rfloor - \lfloor s/2 \rfloor$.

Figures 17 and 18 show some of the cells, edges, and blocks relevant to our reasoning. Cell a is highlighted in all of them. Note that some blocks which are shown as large in the figures might be little ones in reality. In all such cases, we emphasize this possibility in the text.

Suppose, for the sake of contradiction, that a is an even cell but that it is not nice.



Figure 17

Then $a \notin P$. Furthermore, the odd cells of $\mathfrak{X}(a)$ are not in P , either, by the induction hypothesis.

Case 1. $x = \lfloor n/2 \rfloor - \lfloor s/2 \rfloor$ and $\mathfrak{X}(a)$ is a little block. (Figure 17(a).)

Then (iii) for $\mathfrak{X}(a)$ implies $(a + (1, 0)) - (a + (0, 1)) \in P$. Since a is not nice, at least one more edge of G between the four cells in the set $a + \{(1, 0), (0, 1), (-1, 0), (0, -1)\}$ must be in P . It cannot be $(a + (-1, 0)) - (a + (0, -1))$ because then P would contain the vertices of a cycle in G . The other two subcases are symmetric with respect to the line of unit slope through a , and we assume that $(a + (0, -1)) - (a + (1, 0)) \in P$.

It follows that $a + (2, 0) \notin P$ and $a + (2, 1) \notin P$. Thus no edges of $\mathfrak{X}(a + (2, 0))$ are in P . Since also the odd cells of this block are outside of P by the induction hypothesis, we arrive at a contradiction with (iii). (The conclusion holds regardless of whether the block is a little one or a large one.)

Case 2. $\lfloor n/2 \rfloor - \lfloor s/2 \rfloor < x \leq \lfloor n/2 \rfloor$ and $\mathfrak{X}(a)$ is a large block. (Figure 17(b).)

By (iii) for $\mathfrak{X}(a)$ and the induction hypothesis for $a + (0, 2)$, exactly one of the two edges $(a + (1, 0)) - (a + (0, 1))$ and $(a + (0, 1)) - (a + (-1, 0))$ is in P . The two subcases are analogous, and we assume that the former edge is in P while the latter one is not. As in Case 1, at least one more edge of G between the four cells in the set $a + \{(1, 0), (0, 1), (-1, 0), (0, -1)\}$ must be in P , and it cannot be $(a + (-1, 0)) - (a + (0, -1))$ because then P would contain the vertices of a cycle in G . Thus $(a + (0, -1)) - (a + (1, 0)) \in P$.

From here, we arrive at the exact same contradiction as in Case 1. (Once again, regardless of the type of the block $\mathfrak{X}(a + (2, 0))$.)

We have established that if a is an even cell, then it is nice.

For the second half of the induction step, suppose, for the sake of contradiction, that a is an odd cell with $a \in P$. By (i), it follows that $\lfloor n/2 \rfloor - \lfloor s/2 \rfloor + 2 \leq x \leq \lfloor n/2 \rfloor$.

Case 1. $a + (-2, 0) \in P$. (Figure 18(a).)

By (i), $x \geq \lfloor n/2 \rfloor - \lfloor s/2 \rfloor + 4$. Thus $\mathfrak{X}(a + (-1, 1))$ is a large block.

Note that $a \in P$ implies $(a + (-1, 1)) - (a + (0, 1)) \notin P$ and $a + (-2, 0) \in P$ implies $(a + (-1, 1)) - (a + (-2, 1)) \notin P$. Furthermore, by the induction hypothesis, the odd cells of $\mathfrak{X}(a + (-1, 1))$ are not in P and cell $a + (-1, 3)$ is nice. Then (iii) for $\mathfrak{X}(a + (-1, 1))$ implies that exactly one of the two edges $(a + (0, 1)) - (a + (-1, 2))$ and $(a + (-1, 2)) - (a + (-2, 1))$ is in P . (Here, we take into account the fact that $a \in P$ and $a + (-2, 0) \in P$ together imply $(a + (-1, 1)) - (a + (-1, 2)) \notin P$.) The two subcases are analogous, and we assume that the former edge is in P while the latter one is not.



Figure 18

It follows that $a - (a + (0, 1)) - (a + (-1, 2)) \subseteq P$, $a + (1, 1) \notin P$, and $a + (1, 2) \notin P$. Thus no edges of $\mathfrak{X}(a + (1, 1))$ are in P . Since also the odd cells of this block are outside of P by the induction hypothesis, we arrive at a contradiction with (iii). (The block $\mathfrak{X}(a + (1, 1))$ will always be a large one because of our symmetry-breaking assumption that $x \leq \lfloor n/2 \rfloor$. However, in the analogous subcase when $(a + (0, 1)) - (a + (-1, 2)) \notin P$ and $(a + (-1, 2)) - (a + (-2, 1)) \in P$, the contradiction occurs at block $\mathfrak{X}(a + (-3, 1))$ which could happen to be a little one.)

Case 2. $a + (2, 0) \in P$. This case is analogous to Case 1.

Case 3. $a + (-2, 0) \notin P$ and $a + (2, 0) \notin P$. (Figure 18(b).)

Observe that both of $\mathfrak{X}' = \mathfrak{X}(a + (-1, -1))$ and $\mathfrak{X}'' = a + (1, -1)$ are large blocks.

Since $a \in P$, all edges of \mathcal{X}' other than $(a + (-1, -1)) - (a + (-2, -1))$, $(a + (-2, -1)) - (a + (-1, 0))$, and $(a + (-1, 0)) - (a + (-2, 1))$ are not in P . Then, in light of $a \in P$ and $a + (-2, 0) \notin P$, (iii) for \mathcal{X}' implies that exactly one of these edges is in P . In particular, exactly one cell b' out of the pair $a + \{(-1, -1), (-1, 0)\}$ is in P . Similar reasoning applies to \mathcal{X}'' , and we define b'' analogously.

Suppose, for the sake of contradiction, that $b' = a + (-1, 0)$. Then $ab' \subseteq P$ implies $a + (0, 1) \notin P$ and $a + (-1, 1) \notin P$. Furthermore, the odd cells of $\mathcal{X}(a + (-1, 1))$ are outside of P by the induction hypothesis. When this block is a little one, we arrive at a contradiction with (iii) immediately. Otherwise, when it is a large one, we arrive at a contradiction with (iii) anyway once we take into account the fact that, by the induction hypothesis, both of $a + (-1, 1)$ and $a + (-1, 3)$ are nice.

Consequently, $b' = a + (-1, -1)$. Similarly, $b'' = a + (1, -1)$. By (iii) for \mathcal{X}' and \mathcal{X}'' , it follows that also $(a + (-1, -1)) - (a + (-2, -1)) \in P$ and $(a + (1, -1)) - (a + (2, -1)) \in P$. However, then $(a + (-2, -1)) - (a + (-1, -1)) - a - (a + (1, -1)) - (a + (2, -1))$ becomes a rectifiable aberration in P , and we arrive at a contradiction. (In fact, this is the only place in the proof where we use the constraint that P does not contain any rectifiable aberrations.)

We have established that if a is an odd cell, then it cannot be in P . The induction step is complete. □

We are ready to tackle Theorem 5.1.

Proof of Theorem 5.1. For a start, let us rectify all rectifiable aberrations in P one by one. The result will be a snake path Q in G of the same length as P and without any rectifiable aberrations.

By Lemma 5.2, we get that all even cells are nice relative to Q and Q does not visit any odd cells.

By Lemma 5.1, it follows that there is a Hamiltonian path ϱ in H such that $Q = \psi(\varrho)$.

Let t be the number of turns in ϱ . Then both of P and Q are of length $2(k^2 - 1) - t$.

Since the greatest length of a snake path in G is $(n^2 - 1)/2$ by Theorem 1.2, we conclude that $t = 2k - 2$, and so ϱ is in fact a fewest-turn Hamiltonian path in H .

Finally, in order to transform Q back into P , we must restore the rectifiable aberrations which we removed in the beginning. However, it is straightforward to check that the spots in Q where we can introduce a rectifiable aberration are exactly the ones associated with the free cycles corresponding to ϱ . Therefore, P is a lift of ϱ , as needed.

The reasoning in the last few paragraphs shows also the converse: Every lift of a fewest-turn Hamiltonian path in H is a longest snake path in G . □

6. Knight Graphs

In this section, we prove Theorem 1.3. Let m and n be positive integers, let A be the standard board of size $m \times n$, and let G be the knight graph on A .

We begin with the upper bound.

One natural approach would be as follows: First we find some finite knight graph H with pseudosnake density $1/2$. Then we sum over all translation copies of H contained within our board.

The author was not able to implement this idea in its purest form. Below, we present a slightly more complicated argument which relies on a weighted knight graph instead.

We define a *weighted graph* Γ to consist of a simple graph H and a weighting function w which assigns a nonnegative real weight to each vertex of H . When H is finite, we denote the total weight of all of its vertices by $w(\Gamma)$, and we define the pseudosnake density of Γ to be the ratio of the greatest total weight of the vertices in a pseudosnake of H to $w(\Gamma)$.

Lemma 6.1. *Suppose that there is a weighted knight graph Γ with pseudosnake density τ . Then the number of vertices in a pseudosnake of G cannot exceed $\tau mn + \mathcal{O}(m + n)$.*

Here and in the proof, the implicit constants in the \mathcal{O} -terms depend on Γ .

Proof. Let P be a pseudosnake of G .

Consider all translation copies of Γ that fit within A . There are $mn + \mathcal{O}(m + n)$ of them. For each such copy, the total weight of all cells of P within it cannot exceed $\tau w(\Gamma)$.

Conversely, it is true of all but $\mathcal{O}(m + n)$ cells a of P that a is sufficiently far away from the boundary of A for every translation copy of Γ which contains a to fit within A . For each such cell of P , the sum of its weights over all translation copies of Γ which contain it will be $w(\Gamma)$. □

Proof of the upper bound for Theorem 1.3. By Lemma 6.1, it suffices to exhibit one concrete weighted knight graph with pseudosnake density $1/2$. We claim that the one in Figure 19 works. (The figure shows all cells of the graph together with the weights assigned to them.) It has 68 cells with a total weight of 192.




Figure 19

Our claim would likely be extremely difficult to check by hand. However, it is straightforward to check with the help of a standard constraint satisfaction solver. The author has done this twice, using two different constraint satisfaction frameworks: the Copris package for the Scala programming language and the OR-Tools package for the Python programming language. \square

One might wonder how the weighted knight graph in Figure 19 was found.




Figure 20




Figure 21

For a start, let S_{Tess} be the set of all 16 cells of the form $\varepsilon_1(2, 1) + \varepsilon_2(1, 2) + \varepsilon_3(-1, 2) + \varepsilon_4(-2, 1)$, where $\varepsilon_i \in \{0, 1\}$ for all i . Then the knight graph G_{Tess} on S_{Tess} is isomorphic to the tesseract graph. (Figure 20.)

It is not too difficult to check by hand that the pseudosnake density of the tesseract graph is $9/16$, and that it is attained by an essentially unique pseudosnake. Since $9/16$ is very close to $1/2$ from above, we see that G_{Tess} works almost, but not quite.

We can attempt to fix this by taking the union of several overlapping copies of G_{Tess} . Since G_{Tess} itself just barely manages to push through pseudosnake density $1/2$, we can hope that the interference between its copies will prevent too many of them from doing the same.

We formalize this notion as follows: Let S' and S'' be two nonempty finite sets of cells. We define their *sum*, denoted $S' + S''$, to be the multiset of cells which consists of all cells of the form $a' + a''$ with $a' \in S'$ and $a'' \in S''$, and where the multiplicity of each cell is the number of ways that it can be expressed in this form.

Then, given a multiset of cells \mathcal{S} , we define $\mathcal{G}(\mathbb{Z}, \mathcal{S})$ to be the weighted knight graph on the cells of \mathcal{S} where the weight of each cell is its multiplicity in \mathcal{S} .

For each nonempty finite set of cells S , we can think of the weighted knight graph $\mathcal{G}(\mathbb{Z}, S + S_{\text{Tess}})$ as constructed out of several overlapping copies of G_{Tess} . We experiment with different S , and eventually we strike gold with $S_{\text{Dia}} = [0; 3]^2 \setminus \{0, 3\}^2$. This is the Aztec diamond of order two. (Figure 21.)

We go on to the lower bound.

One natural approach would be as follows: First we find a doubly periodic pseudosnake P_∞ in $\mathcal{G}(\mathbb{Z}, \mathbb{Z}^2)$ with density $1/2$, where furthermore every cell is of degree exactly two and there are no finite cycles. (Here, “doubly periodic” means that there are two linearly independent two-dimensional vectors u and v with $a \in P_\infty \Leftrightarrow a + u \in P_\infty \Leftrightarrow a + v \in P_\infty$ for all cells a of \mathbb{Z}^2 .)

Then, given a board A , we take the restriction P^* of P_∞ to A . Because of the structure of P_∞ , this restriction will be the disjoint union of several paths. We make some modifications near the boundary of A , deleting some cells from P^* and replacing them with new ones, so as to stitch all of these paths together into a single snake path or cycle. Since P_∞ is doubly periodic with density $1/2$, originally P^* will contain $mn/2 + \mathcal{O}(m + n)$ cells. Finally, keeping our modifications close to the boundary of A ensures that they cost us $\mathcal{O}(m + n)$ of these cells altogether.

Finding a suitable P_∞ is straightforward enough. For example, the set of all cells (x, y) with $x \bmod 4 \in \{0, 1\}$ works. It consists of vertical strips of width two spaced two units apart.

However, the second part of our plan runs into significant difficulties. Thus the construction we present below is somewhat more complicated. We divide A into four large regions; we fill up different regions using different

pseudosnakes P_∞ ; and we make stitching-together modifications not only near the boundary of A , but also near the boundaries between regions.

We continue with the details.

We define a *twine* to be a board of height two. When $s \geq 2$, the knight graph on a twine with width s is the disjoint union of four paths, two spanning $\lfloor s/2 \rfloor$ cells each and two spanning $\lceil s/2 \rceil$ cells each.




Figure 22




Figure 23

Consider a twine E with lower left corner cell a . To *tie off* E on the left, we add to it the four cells in the set $a + \{(-2, 1), (-1, 1), (-1, 3), (0, 3)\}$. (Figure 22.) Similarly, to tie off on the right a twine E with lower right corner a , we add to it the reflections of the four cells above with respect to the vertical line through a .

Consider, now, two twines E and F with lower left corners a and b satisfying $a + (1, 4) = b$. To *splice together* E and F on the left, we add to them the ten cells in the set $a + \{(-3, 4), (-3, 5), (-2, 2), (-2, 3), (-2, 6), (-1, 1), (-1, 2), (-1, 6), (0, 5), (0, 6)\}$. (Figure 23.) Similarly, to splice together on the right two twines E and F whose lower right corners a and b satisfy $a + (-1, 4) = b$, we add to them the reflections of the ten cells above with respect to the vertical line through a .

Let k be a positive integer and let $I = [x'; x'']$ be an integer interval with $|I| \geq 8k - 5$. For each i with $0 \leq i \leq k - 1$, if i is even then construct the twine $E_i = [x' + 4i; x'' - 4i] \times [4i; 4i + 1]$, and if i is odd then construct the twine $E_i = [x' + 4i + 3; x'' - 4i + 3] \times [4i; 4i + 1]$. Tie off E_0 on the left; for all i with $0 \leq i \leq k - 2$, splice together E_i and E_{i+1} on the right if i is even, and on the left if i is odd; and, finally, if k is even then tie off E_{k-1} on the left, and if k is odd then tie it off on the right.




Figure 24

We denote the resulting set of cells by $U(k, I)$. For example, Figure 24 shows the knight graph on $U(4, [0; 32])$.

Lemma 6.2. *Suppose that $k \geq 2$ and $|I|$ is odd. Then the knight graph on $U(k, I)$ is a cycle.*

Proof. Denote $H = \mathcal{G}(\hat{\mathcal{L}}, U(k, I))$. It is straightforward to check that all cells of H are of degree two. We are left to verify that H is connected.

Suppose, for the sake of clarity, that k is odd. The opposite case, when it is even, is similar.

Let a_i and b_i be the lower left and lower right corner cells of E_i . Let also a'_i and a''_i be the two cells of E_i adjacent by side to a_i , and define b'_i and b''_i similarly for b_i .

It is straightforward to verify that, since $|I|$ is odd: (a) A path in H connects b'_0 and b''_0 and covers E_0 together with the cells which tie it off; (b) For each i with $1 \leq i \leq k - 2$, two paths in H connect the pairs $\{a'_i, a''_i\}$ and $\{b'_i, b''_i\}$ and cover E_i together with two cells in the adjacent splices; and (c) A path in H connects a'_{k-1} and a''_{k-1} and covers E_{k-1} together with the cells which tie it off.

Finally, the remaining cells of the splices of $U(k, I)$ form additional knight paths in H which connect the pairs $\{b'_i, b''_i\}$ and $\{b'_{i+1}, b''_{i+1}\}$ for all even i with $0 \leq i \leq k - 3$ as well as the pairs $\{a'_i, a''_i\}$ and $\{a'_{i+1}, a''_{i+1}\}$ for all odd i with $1 \leq i \leq k - 2$. \square

Proof of the lower bound for Theorem 1.3. We construct a large snake cycle in G , and for a large snake path we can simply delete one cell from that cycle.

Suppose, without loss of generality, that $m \leq n$. Since we are already willing to accept a tolerance of $\mathcal{O}(m + n)$, we can safely assume that $m = 8k + 14$ for some positive integer k with $k \geq 3$, and that n is even.

Construct $U_I = U(k, [8; n - 12])$. Let also V_I be the reflection of $U(k, [8; m - 12])$ with respect to the line $x = y$. Lastly, let U_{II} and V_{II} be symmetric to U_I and V_I with respect to the center of A .




Figure 25

The knight graph on $U_I \cup V_I \cup U_{II} \cup V_{II}$ is the disjoint union of four cycles. We proceed to stitch these four cycles together into a single longer cycle.

Define $S_{Del} = \{(6, 9), (9, 6)\}$ and $S_{Add} = \{(3, 6), (4, 4), (6, 3), (7, 10), (9, 9), (10, 7)\}$. Delete the two cells of $(4, 4) + S_{Del}$ from U_I and V_I , and replace them with the six cells of $(4, 4) + S_{Add}$. This stitches together the cycles of U_I and V_I . (Figure 25.)

We carry out two more such modifications. For one of them, we reflect the sets S_{Del} and S_{Add} with respect to the vertical axis of symmetry of the board, we delete the two cells in the image of S_{Del} from U_I and V_{II} , and we replace them with the six cells in the image of S_{Add} . This stitches together the cycles of U_I and V_{II} . For the other one, we proceed similarly, except that the reflections are done with respect to the horizontal axis of symmetry of the board. This stitches together the cycles of V_I and U_{II} .




Figure 26

Let W be the final set of cells obtained in this way. For example, Figure 26 shows W in the case when $k = 5$, $m = 54$, and $n = 72$.

Observe that the density of W within A is $1/2$ everywhere except within five strips of bounded width. (Four of these strips surround portions of the interior angle bisectors at the four corners of A , and the fifth one

surrounds a portion of the horizontal axis of symmetry of A . The corresponding mostly hollow areas are clearly visible in Figure 26.) Consequently, the number of cells in W is $mn/2 + \mathcal{O}(m + n)$.

On the other hand, W is the vertex set of a snake cycle in G . □

7. Further Work on King Graphs

In this section, we collect some additional results and open problems on king graphs.

We saw that the behavior of the longest snake paths in $\mathcal{G}(\mathfrak{K}, n \times n)$ depends on the parity of n . For cycles, it appears that there are four classes instead, depending on the value of $n \pmod 4$. The techniques we developed for paths quickly resolve two of them.

Theorem 7.1. *Let n be a positive integer with $n \equiv 0 \pmod 4$ and $n \geq 8$. Then the greatest length of a snake cycle in the king graph of size $n \times n$ is $n^2/2 - 1$. Furthermore, for all such n , there are exactly 48 snake cycles which attain this greatest length. These cycles are all asymmetric, and so six of them are essentially distinct.*

For completeness, there is a unique snake cycle of the greatest length 8 when $n = 4$.

Each one of the cycles of Theorem 7.1 is shaped like a double spiral.

It is curious that the number of longest snake cycles freezes in this way, and the cycles themselves crystallize into a single inflexible structure. A similar phenomenon occurs in the setting of Theorem 7.3.

Proof. Let $n = 2k$ and define A, B, G, H , and Φ as in Section 3. Let also C be a snake cycle in G of length at least $n^2/2 - 1$.

As in Section 3, for each cell b of B at most two cells of $\Phi(b)$ are in C . Let us call b *deficient* when this bound is not attained. Thus there is at most one deficient cell.

Observe that, if an edge of C joins one cell of $\Phi(b')$ and one cell of $\Phi(b'')$, with $b' \neq b''$, then $b'b''$ must still be an edge of H . Otherwise, assuming for concreteness that $b' + (1, 1) = b''$, by an argument similar to the one in Section 3 we see that at least two of the four cells in the set $b' + [0; 1]^2$ must be deficient, a contradiction.

This allows us to define the Hamiltonian cycle ϱ in H relative to C in the same manner as in Section 3.

Define also the cycle E in H as in the proof of Lemma 3.2. Since E is not a Hamiltonian cycle of H when $k \geq 4$, at least one edge of E must be outside of ϱ .

On the other hand, observe that Lemma 3.1 now admits a unique exception: When the turn occurs at a deficient cell. Consequently, every edge of E outside of ϱ must possess a deficient endpoint.

Define A^*, B^*, G^* , and H^* as in Section 3. Let also C^* and ϱ^* be the restrictions of C and ϱ to A^* and B^* , respectively. We obtain that: (a) There is exactly one deficient cell; (b) There is exactly one edge $\beta'\beta''$ of E outside of ϱ ; (c) ϱ^* is a Hamiltonian path in H^* whose endpoints are the two neighbors of β' and β'' in B^* ; and (d) C^* is a snake path in G^* of the greatest possible length whose associated Hamiltonian path in H^* is ϱ^* .

However, the proof of Theorem 1.1 gives us a complete description of all Hamiltonian paths in H^* associated with a longest snake path in G^* . Since the two endpoints of ϱ^* are neighbors in H^* , we conclude that when $n \geq 12$ a symmetry of B^* must map ϱ^* onto the unique regular Hamiltonian path in H^* of type II(0), as defined in Section 3. The rest is straightforward. □

The other class we can tackle without too much extra effort is $n \equiv 3 \pmod 4$. First, though, we need to sort through some preliminaries.

Let s be an even positive integer. Given a permutation σ of $[0; s - 1]$, consider a closed curve in the plane defined in the same way as the curve κ in Section 5, except that one additional arc on the right of the coordinate axis Oy joins points $(0, \sigma^{-1}(s - 1))$ and $(0, \sigma^{-1}(0))$. A curve of this form which does not intersect itself is known as a *closed meander*.

Consider a Hamiltonian cycle in the grid graph Γ of size $s \times s$. The smallest number of turns that such a cycle can make is $2s$. [3] Furthermore, the closed meanders with s arcs and the fewest-turn Hamiltonian cycles in Γ are related in a way analogous to the relation between stamp-folding permutations and fewest-turn Hamiltonian paths.

Observe, lastly, that our definition of a lift in Section 5 works just as well with cycles instead of paths.

Theorem 7.2. *Let n be a positive integer with $n \equiv 3 \pmod 4$ and $n = 2k - 1$. Then the greatest length of a snake cycle in the king graph of size $n \times n$ is $(n^2 - 1)/2$. Every lift of a fewest-turn Hamiltonian cycle in the grid graph of size $k \times k$ is a longest snake cycle in the king graph of size $n \times n$. Conversely, every longest snake cycle in the king graph of size $n \times n$ can be obtained uniquely as a lift of some fewest-turn Hamiltonian cycle in the grid graph of size $k \times k$.*

This time around, our analysis of paths carries over to cycles nearly verbatim.

Proof. The upper bound follows by the same argument as Theorem 1.2. The lower bound is a corollary of the structure description. Finally, the structure description follows by the same argument as Theorem 5.1. □

With the remaining two classes, the main difficulty is this: In both Sections 3 and 5, we introduce the half-sized square board B of side $\lceil n/2 \rceil$ together with its grid graph H , and our reasoning relies heavily on the properties of the Hamiltonian paths of H . For Theorems 7.1 and 7.2, it is the Hamiltonian cycles of H that matter instead. When $n \equiv 1 \pmod{4}$ or $n \equiv 2 \pmod{4}$, however, the side of B is odd and H does not admit a Hamiltonian cycle.

This throws a substantial wrench in the works. While our upper bounds all go through as before, the constructions that support the lower bounds do not, and the gap which opens between the two appears to be difficult to close.

We continue with some tentative remarks.

Let \mathcal{N}_r be the set of all positive integers n such that $n \equiv r \pmod{4}$.

Fix n^* in \mathcal{N}_r with $n^* \geq 7$. For each n in \mathcal{N}_r with $n \geq n^*$, construct the subset D_n of the standard board of size $n \times n$ as follows: Take all cells of the form (x, y) with $x - y \geq 1$, $x + y \leq n - 2$, y even, and $0 \leq y \leq (n - n^*)/2$ together with their images under the symmetries of the board. Delete the cells $(0, 2)$, $(0, 3)$, and $(0, 4)$, and replace them with the cells $(1, 2)$ and $(1, 4)$. Finally, for all even i with $2 \leq i \leq (n - n^*)/2$, delete the three cells in the set $(i, i) + \{(0, 1), (0, 3), (0, 4)\}$, and replace them with the three cells in the set $(i, i) + \{(0, 0), (1, 2), (1, 4)\}$. Note that D_n is the vertex set of a snake path in $\mathcal{G}(\mathfrak{S}, n \times n)$.

We say that \mathcal{N}_r crystallizes at n^* when, for all n in \mathcal{N}_r with $n \geq n^*$ and every longest snake cycle C in $\mathcal{G}(\mathfrak{S}, n \times n)$, there is a symmetry π of the corresponding board such that the set of all cells of $\pi(C)$ outside of the concentric subboard of size $(n^* - 4) \times (n^* - 4)$ coincides with D_n .

Thus, in particular, if \mathcal{N}_r crystallizes, then there are two constants ℓ and μ such that the greatest length of a snake cycle in $\mathcal{G}(\mathfrak{S}, n \times n)$ is $n^2/2 - \ell$ and the number of essentially distinct snake cycles which attain this greatest length is μ for all n in \mathcal{N}_r with $n \geq n^*$. Furthermore, each one of these cycles is asymmetric, and so for all such n the total number of longest snake cycles in $\mathcal{G}(\mathfrak{S}, n \times n)$ is 8μ .

The proof of Theorem 7.1 shows that the class \mathcal{N}_0 crystallizes at $n^* = 12$ with $\ell = 1$ and $\mu = 6$.

The author finds it reasonably plausible that each one of the classes \mathcal{N}_1 and \mathcal{N}_2 might crystallize as well. Experimental data suggests that perhaps the class \mathcal{N}_1 crystallizes at $n^* = 13$ with $\ell = 5/2$ and $\mu = 69$ whereas the class \mathcal{N}_2 crystallizes at $n^* = 14$ with $\ell = 3$ and $\mu = 72$.

The next result might be helpful in the case of the class \mathcal{N}_1 . The two cycles of Theorem 7.3 appear to be related to the longest snake cycles of $\mathcal{G}(\mathfrak{S}, (2k - 1) \times (2k - 1))$ in a way somewhat similar to how the fewest-turn Hamiltonian paths and cycles of grid graphs are related to the paths and cycles of Theorems 5.1 and 7.2.

We define a *near-Hamiltonian cycle* of a graph G to be a cycle in G which visits all but one vertices of G .

Theorem 7.3. *Let k be an odd positive integer with $k \geq 5$. Then the smallest number of turns in a near-Hamiltonian cycle of the grid graph of size $k \times k$ is $2k$. Furthermore, for all such k , there are exactly 16 near-Hamiltonian cycles which attain this smallest number. These cycles are all asymmetric, and so two of them are essentially distinct.*

For completeness, there is a unique near-Hamiltonian cycle with the smallest number of turns, namely four, when $k = 3$.

Just as in Theorem 7.1, the cycles of Theorem 7.3 are shaped like double spirals.

Proof. Let B be the standard board of size $k \times k$, let H be the grid graph on B , and let C be a near-Hamiltonian cycle in H . Denote the unique cell of B which C omits by o .

Suppose, for the sake of contradiction, that there are a row and a column of B without an edge of C . Then the cell at their intersection must be o , and it cannot lie on the boundary of B . Since C visits the neighbors of o in H but the row and column of o do not contain edges of C , it follows that all edges of the cycle $(o + (1, 1)) \sim (o + (-1, 1)) \sim (o + (-1, -1)) \sim (o + (1, -1)) \sim (o + (1, 1))$ must be in C . However, this cycle is not near-Hamiltonian when $k \geq 5$, a contradiction.

Suppose, for concreteness, that every row contains an edge of C . Define the segments of C as in Section 5. Since every row contains a horizontal segment of C , and the endpoints of each such segment are turns, we get that C makes at least $2k$ turns altogether. Suppose, from now on, that this bound is attained and that every row contains exactly one horizontal segment of C . Thus, in particular, o cannot lie in the lowermost or topmost row of B unless it is a corner cell of B .

Suppose, for the sake of contradiction, that the leftmost and rightmost columns of B contain one vertical segment of C each. Then they cannot contain o unless it is a corner cell of B . When, say, $o = (0, 0)$, it follows that C must contain all edges of the cycle $(0, 1) - (1, 1) - (1, 0) \sim (k - 1, 0) \sim (k - 1, k - 1) \sim (0, k - 1) \sim (0, 1)$. Otherwise, when o is not on the boundary of B , it follows that C must contain all edges of the cycle $(0, 0) \sim (k - 1, 0) \sim (k - 1, k - 1) \sim (0, k - 1) \sim (0, 0)$. In both cases, however, the cycle in question is not near-Hamiltonian when $k \geq 5$, a contradiction.

Suppose, for concreteness, that the leftmost column of B contains at least two vertical segments of C . Since the number of vertical segments in C is the same as its number of horizontal segments, and we have already assumed that the latter number equals k , we get that some column u of B does not contain any edges of C .

Consequently, C crosses over u every time when it visits this column. Since the total number of crossings must be even, and u contains an odd number of cells, we obtain that o must be in u .

It follows that there is exactly one column of B without vertical segments of C . (Since each such column must contain o .) Thus the leftmost column of B must contain exactly two vertical segments of C and all other columns except for u must contain exactly one vertical segment of C each. (Since there are a total of k vertical segments in C .)

Let the two vertical segments of C in the leftmost column of B be $(0, 0) \sim (0, w)$ and $(0, w + 1) \sim (0, k - 1)$. Suppose, for concreteness, that $1 \leq w \leq \lfloor k/2 \rfloor - 1$.

From this point on, we establish the identity $w = 1$ and the desired result together, by induction on k . The base case $k = 5$ is straightforward. For the induction step, suppose that $k \geq 7$ and that we have already settled the question on all smaller boards.

Let us delete the subpath $(1, w) - (0, w) \sim (0, 0) \sim (k - 1, 0) \sim (k - 1, k - 1) \sim (0, k - 1) \sim (0, w + 1) - (1, w + 1)$ from C , and let us replace it with the edge $(1, w) - (1, w + 1)$. The result will be a near-Hamiltonian cycle C^* in the grid graph H^* on the concentric subboard B^* of B of size $(k - 2) \times (k - 2)$ given by $B^* = [1; k - 2]^2$.

Since we have deleted at least six turns from C and we have added at most two new ones in their place, C^* can make at most $2k - 4$ turns altogether. Thus our induction hypothesis applies to it, and so in fact C^* makes exactly $2k - 4$ turns, two of which are at cells $(1, w)$ and $(1, w + 1)$ where they form the subpath $(2, w) - (1, w) - (1, w + 1) - (2, w + 1)$. Still by the induction hypothesis, C^* omits exactly one edge of the cycle $(1, 1) \sim (k - 2, 1) \sim (k - 2, k - 2) \sim (1, k - 2) \sim (1, 1)$, and that edge is an image of the edge $(1, 2) - (1, 3)$ under a symmetry of B^* .

We conclude that $w = 1$ and exactly two of the fewest-turn near-Hamiltonian cycles of H^* fit as a suitable C^* . Therefore, there are exactly two essentially distinct fewest-turn near-Hamiltonian cycles in H , both of them asymmetric, and the induction step is complete. \square

8. Further Work on Leaper Graphs

In this section, we collect some additional results and open problems on leaper graphs. (To be defined shortly.)

For the knight, it would be interesting to see a human-friendly proof of the upper bound in Theorem 1.3. Or, if not that, then at least it would be nice to know if there is an unweighted knight graph with pseudosnake density $1/2$ which we could have used instead of the weighted one.

One natural direction of generalization for our results in Section 6 is offered by leapers.

Let p and q be nonnegative integers with $p \leq q$, not both zero. A (p, q) -leaper is a fairy chess piece which moves as a generalized knight, leaping p units away along one coordinate axis and q units away along the other.

Let L be a (p, q) -leaper. The leaper graph of L on a set of cells S , denoted $\mathcal{G}(L, S)$, is defined similarly to the king and knight graphs on S , except that the adjacency condition becomes $\{|x' - x''|, |y' - y''|\} = \{p, q\}$ instead.

Let $d = \gcd(p, q)$. Then the leaper graph of L on the board of size $m \times n$ is the disjoint union of several isomorphic copies of the leaper graphs of a $(p/d, q/d)$ -leaper on the boards of sizes $\lfloor m/d \rfloor \times \lfloor n/d \rfloor$, $\lfloor m/d \rfloor \times \lceil n/d \rceil$, $\lceil m/d \rceil \times \lfloor n/d \rfloor$, and $\lceil m/d \rceil \times \lceil n/d \rceil$, each copy scaled up by a factor of d . Thus we can safely assume that $d = 1$.

For p and q relatively prime, L is known as *free* when $p + q$ is odd and *half-free* when it is even. Briefly, one reason for this distinction is that $\mathcal{G}(L, \mathbb{Z}^2)$ is connected when L is free but consists of two connected components when L is half-free.

A *skew* leaper is one for which p and q are positive and distinct. The only non-skew leapers with relatively prime p and q are the $(0, 1)$ -leaper, known as the *wazir*, and the $(1, 1)$ -leaper, known as the *fers*. Of course, the wazir graph on a board coincides with the grid graph on that board. Furthermore, $\mathcal{G}(\text{Fers}, m \times n)$ can also be viewed as the direct product of two paths with m and n vertices, respectively.

We take a look at the wazir and fers first, and after that we will focus on skew leapers.

Even though Question **A** for grid graphs on rectangular boards is a very natural thing to ask, the only earlier reference for it known to the author as of the time of writing is [6], a puzzle game website where players are invited to construct snake paths in grid graphs on square boards, with longer paths scoring higher.

An asymptotic estimate is straightforward to obtain. The argument we give for the upper bound is not new; it is essentially identical to the argument used in [4] to bound from above the pseudosnake density of $\mathcal{G}(\square, \mathbb{Z}^2)$. (We discuss one natural way to define the pseudosnake density of certain infinite graphs below.) For the lower bound, the general strategy we outlined in Section 6 goes through without a hitch. Once again, [4] contains the same pseudosnake in $\mathcal{G}(\square, \mathbb{Z}^2)$.

Proposition 8.1. *Let m and n be positive integers. Then both the longest snake path and the longest snake cycle in the grid graph of size $m \times n$ are of length $2mn/3 + \mathcal{O}(m + n)$.*

Proof. Let A be the standard board of size $m \times n$ and let G be the grid graph on A .

For the upper bound, let P be a snake path in G ; the argument for cycles is similar.

Let S be the vertex set of P and let T be the complement of S within A . Then nearly every cell of S is adjacent to two cells of T ; each exception is either an endpoint of P or near the boundary of A , and so there are $\mathcal{O}(m+n)$ of them. On the other hand, every cell of T is adjacent to at most four cells of S . Thus $2|S| \leq 4|T| + \mathcal{O}(m+n)$, and so $|S| \leq 2|S \cup T|/3 + \mathcal{O}(m+n)$ as well.

We move on to the lower bound. Let S_∞ be the set of all cells (x, y) in \mathbb{Z}^2 with $x \not\equiv y \pmod{3}$. Then S_∞ induces a pseudosnake P_∞ in $\mathcal{G}(\square, \mathbb{Z}^2)$.

Suppose without loss of generality that $m \geq 10$ and $n \geq 10$, let A^* be the subboard of A given by $A^* = [4; n-5] \times [4; m-5]$, and also let P^* be the restriction of P_∞ to A^* . Then P^* is the disjoint union of several paths, and it is straightforward to add several cells out of $A \setminus A^*$ to P^* so as to stitch these paths together into a single snake path or cycle. \square

The fers can be handled similarly.

Proposition 8.2. *Let m and n be positive integers. Then both the longest snake path and the longest snake cycle in the fers graph of size $m \times n$ are of length $mn/3 + \mathcal{O}(m+n)$.*

The proof is analogous to that of Proposition 8.1, and we omit it.

For wazir and fers graphs on rectangular boards, it might be possible to obtain exact answers to Question A. By way of experimental data, [6] contains a table listing the greatest length of a snake path in the grid graph of size $n \times n$ for all n with $2 \leq n \leq 15$.

We continue on to skew leapers. Suppose, for concreteness, that $p < q$.

It seems highly likely that an exact answer to Question A would be out of reach for skew leapers on arbitrary rectangular boards, or even on arbitrary square boards. For this reason, we propose a weakened version of it.

Question 8.1. *Let L be a skew leaper. What are some interesting lower and upper bounds for the greatest length of a snake path or cycle of L on a given rectangular board?*

Let us pick some of the low-hanging fruit.

When L is half-free, let $\text{Free}(L)$ denote the free (p', q') -leaper with $p' = (q-p)/2$ and $q' = (p+q)/2$.

Suppose, now, that L is a skew free leaper. We will consider this case first, and then for Proposition 8.4 we will reduce the half-free case to the free case using the transformation introduced above.

In all of the following asymptotic estimates, the implicit constants in the \mathcal{O} -terms depend on L .

One construction will be particularly useful to us, and so we introduce special notation for it:

Let P be a pseudosnake in $\mathcal{G}(L, m \times n)$ with vertex set S . Consider the union T of all sets of cells of the form $((n+q)i, (m+q)j) + S$, over all integers i and j . Then the induced subgraph on vertex set T is a pseudosnake in $\mathcal{G}(L, \mathbb{Z}^2)$, as the translation copies of P which this subgraph consists of are too far away from one another to interact in any way. We denote this pseudosnake by $\Upsilon(m \times n, P)$.

Now let τ_n be the pseudosnake density of $\mathcal{G}(L, n \times n)$.

Observe that the sequence $\{\tau_n\}_{n=1}^\infty$ converges. Indeed, let $n^2 \leq N$. Since we can cover the board of size $N \times N$ with $\lceil N/n \rceil^2$ subboards of size $n \times n$, we get that $\tau_N \leq \tau_n + \mathcal{O}(1/n)$. On the other hand, fix a largest pseudosnake P in $\mathcal{G}(L, n \times n)$. Then the restriction of $\Upsilon(n \times n, P)$ to the board of size $N \times N$ is a pseudosnake in $\mathcal{G}(L, N \times N)$, and so $\tau_N \geq \tau_n + \mathcal{O}(1/n)$.

We denote $\tau(L) = \lim_{n \rightarrow \infty} \tau_n$, and we call this the pseudosnake density of $\mathcal{G}(L, \mathbb{Z}^2)$ or, for short, the pseudosnake density of L .

Question 8.2. *Let L be a skew free leaper. What is the pseudosnake density of L ? Or, alternatively, what are some interesting lower and upper bounds for it?*

Consider the four-dimensional infinite grid graph $\mathcal{G}(\square, \mathbb{Z}^4)$. We define its pseudosnake density η similarly to how we defined $\tau(L)$. Observe that η is an absolute constant which does not depend on L . The pseudosnake densities of infinite grid graphs with arbitrarily many dimensions have been studied before. [4]

Proposition 8.3. *For all skew free leapers L , the pseudosnake density of L satisfies $1/2 \leq \tau(L) \leq \eta$.*

Proof. For the lower bound, it suffices to exhibit a doubly periodic pseudosnake in $\mathcal{G}(L, \mathbb{Z}^2)$ with density $1/2$.

Since L is free, exactly one of p and q is even. Denote that even value by r . When $r \equiv 2 \pmod{4}$, let S be the set of all cells (x, y) such that $\lfloor x/2 \rfloor$ is even. Otherwise, when $r \equiv 0 \pmod{4}$, let S be the set of all cells (x, y) such that $\lfloor x/2 \rfloor + y$ is even. Then the induced subgraph of $\mathcal{G}(L, \mathbb{Z}^2)$ on vertex set S works.

For the upper bound, let n be a positive integer, fix a largest pseudosnake P in $\mathcal{G}(L, n \times n)$, and let $Q = \Upsilon(n \times n, P)$. Then Q is a doubly periodic pseudosnake in $\mathcal{G}(L, \mathbb{Z}^2)$ with density $\tau_n + \mathcal{O}(1/n)$.

Consider, now, the induced subgraph R of $\mathcal{G}(\square, \mathbb{Z}^4)$ whose vertex set consists of all four-dimensional integer points (x_1, x_2, x_3, x_4) such that $x_1(p, q) + x_2(q, p) + x_3(-p, q) + x_4(-q, p)$ is a cell of Q . Then R is a quadruply periodic pseudosnake in $\mathcal{G}(\square, \mathbb{Z}^4)$ with the same density as Q . We let n grow without bound, and the conclusion follows. \square

Had it been the case that $\eta = 1/2$, Proposition 8.3 would have resolved Question 8.2 immediately. However, it has been demonstrated that $649/1296 \leq \eta \leq 20/39$, with $649/1296 \approx 0.50077$ and $20/39 \approx 0.51282$. [4] Still, Proposition 8.3 and this result together imply, for all skew free leapers L , that $1/2 \leq \tau(L) \leq 20/39$. Since the gap between these two bounds is rather narrow, and the graph $\mathcal{G}(L, \mathbb{Z}^2)$ is, in some intuitive sense, more crowded than $\mathcal{G}(\square, \mathbb{Z}^4)$, it seems plausible that in fact $\tau(L) = 1/2$ for all skew free leapers L . As we saw in Section 6, this is certainly true of the knight.

The natural connection between pseudosnake density and snake paths and cycles is as follows:

Proposition 8.4. *Let L be a skew leaper and let m and n be positive integers. When L is free, the greatest length of a snake path or cycle of L on the board of size $m \times n$ does not exceed $\tau(L) \cdot mn + \mathcal{O}(m+n)$. Furthermore, when L is half-free, it does not exceed $\tau(L)/2 \cdot mn + \mathcal{O}(m+n)$.*

Proof. Let P be a snake path or cycle in $\mathcal{G}(L, m \times n)$ with s cells.

Suppose first that L is free. Since $\Upsilon(m \times n, P)$ is a doubly periodic pseudosnake in $\mathcal{G}(L, \mathbb{Z}^2)$, its density $s/(m+q)(n+q)$ does not exceed $\tau(L)$.

Suppose, now, that L is half-free. Then $\varepsilon = (x+y) \bmod 2$ is constant over all cells (x, y) of P . Instead of $\Upsilon(m \times n, P)$, take its image under the transformation $(x, y) \rightarrow ((x-y+\varepsilon)/2, (x+y+\varepsilon)/2)$. This is a pseudosnake in $\mathcal{G}(\text{Free}(L), \mathbb{Z}^2)$, and from this point on the argument continues as before. \square

The author finds it reasonably plausible that the upper bounds of Proposition 8.4 might in fact be attained for all skew leapers. As we saw in Section 6, this is indeed the case for the knight.

Note that our construction for the lower bound in the proof of Proposition 8.3 yields a doubly periodic pseudosnake in $\mathcal{G}(L, \mathbb{Z}^2)$ where all cells are of degree exactly two and there are no finite cycles. Thus for free leapers L with $\tau(L) = 1/2$ and their corresponding half-free leapers, this construction might play a role in a proof that the upper bounds of Proposition 8.4 are attained which follows some variant of the strategy we outlined in Section 6.

Acknowledgements

The author would like to thank Professor Donald Knuth for introducing him to the subject of the longest snake paths and cycles in chess piece graphs.

References

- [1] T. Dawson, Échecs Féériques, *L'Échiquier*, volume 2, issue 2, 1930; issue 3, 1931.
- [2] N. Sloane, My Favorite Integer Sequences, *Sequences and Their Applications: Proceedings of SETA '98*, 1999. Revised at <http://neilsloane.com/doc/sg.pdf>.
- [3] G. Jelliss, *Knight's Tour Notes*, website, section Wazir Wanderings, 2001. Internet Archive snapshot at <https://web.archive.org/web/20020206053321/http://www.ktn.freeuk.com/9c.htm>. Revised and collected in G. Jelliss, *Knight's Tour Notes*, twelve-volume monograph, volume 2 (Walker Tours), 2019, <https://www.mayhematics.com/p/p.htm>.
- [4] M. Matamala, E. Prisner, and I. Rapaport, k -Pseudosnakes in Large Grids, *LATIN 2002: Theoretical Informatics*, 2002.
- [5] D. Knuth, *The Art of Computer Programming*, volume 4, pre-fascicle 5c (section 7.2.2.1, Dancing Links), 2019, <https://cs.stanford.edu/~knuth/fasc5c.ps.gz>. Revised and collected in D. Knuth, *The Art of Computer Programming*, volume 4B (Combinatorial Algorithms, Part 2), 2022.
- [6] D. Radcliffe, *Build-a-Snake*, website, 2020, <https://snake.radcliffe.dev/>. Game idea by C. Danielson. Table of optimal path lengths by A. Goupil and A. Howroyd.
- [7] K. Golder, *Minimum Turn Hamiltonian Paths on Rectangular Grids*, thesis presented for the degree of Master of Science, Emporia State University, 2021.