Bounding the Diameter of Flip Graphs of Split Chessboards

Michelle Kang

Under the direction of

Tristan Yang
Department of Mathematics
Massachusetts Institute of Technology

Research Science Institute
August 1, 2023

Abstract

This paper presents an investigation into the diameter of flip graphs associated with split chessboards, which are grid-like structures where unit squares can be separated by a diagonal and adjacent sections are of opposite colors. We conduct an exploration of the properties of flip graphs concerning the standard chessboard and identify distinct connected components within. We further establish a compelling link between split chessboards and plabic graphs, shedding light on essential aspects such as the arrangement of colors on the grid, as well as trip permutations within the disk.

Summary

This paper investigates the structure of split chessboards, which are chessboard-like grids where unit squares can be divided into two triangles by a diagonal. We explore the properties of a graph of split chessboards, such as the distance between the farthest two split chessboards, and how specific mutations may affect it. Additionally, we establish a connection between split chessboards and other graphs, providing valuable insights into the grid's color arrangement. Our findings contribute to a better understanding of the complexity of mutations in split chessboards.

1 Introduction

A split chessboard is a chessboard-like grid such that each unit square can be split into two triangles by the diagonal. Sections within the chessboard sharing an edge must be different colors with the grid being colored in two colors in total.

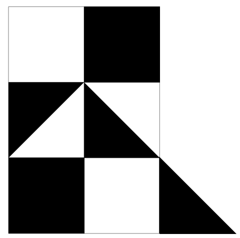

Figure 1: Example split chessboard

Let us define a mutation around a lattice point as follows. Connect the four adjacent lattice points to form four isosceles triangles and swap the colors of each of these triangles from black to white and vice versa. Two possible examples are shown in Figure 2. Applying this mutation on any lattice point in the interior of a split chessboard gives another valid split chessboard without changing the boundary shape.

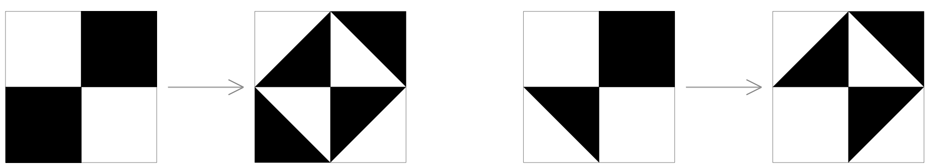

Figure 2: Valid mutations on a 2×2 grid

A flip graph is a graph with vertices being states and edges being transformations. Two vertices are connected by an edge if there is a transformation that changes one to the other. In this case, the vertices are different split chessboards for a specified boundary and the flips correspond to the mutation described above. In Figure 3, we have a connected component of the 2×4 grid. This project aims to study the structure of these flip graphs, specifically the diameter, which is the longest possible distance between any two connected split chessboards in the graph.

In this paper, we determine the diameter of the connected component of a flip graph of a split chessboard containing the standard chessboard and examine the properties of other connected components. We first examine the behavior of these mutations in a standard chessboard, then show that a given algorithm of mutations is the farthest number of mutations from the standard chessboard. We then find that two such chessboards are maximally far apart, giving us the diameter of that connected component. As we are interested in the diameter of the entire flip graph, we examine the properties of other connected components as well.

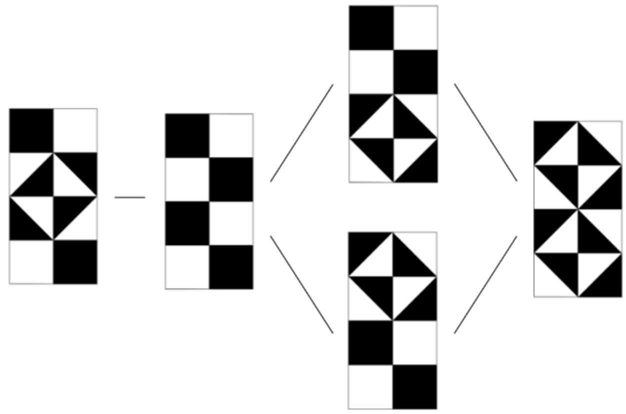

Figure 3: Flip grid for 2×4 grid

Furthermore, we can also form a bijection between split chessboards and planar bicolored graph, also known as plabic graphs. Plabic graphs were first introduced in 2006 by Postnikov [1] in his paper about networks. However, no research has been done on split chessboards.

2 Preliminaries

We first define relevant terms that will be used throughout this paper.

Definition 2.1. The *diameter* of a graph is the largest possible value of the minimum distance between any two vertices in the graph.

Definition 2.2. A *standard chessboard* is a split chessboard with no unit squares split into two separate colors.

Definition 2.3. A *diagonal edge* is the diagonal of a unit square splitting it into two different colors within a split chessboard.

2.1 Plabic graphs

A planar bicolored graph, or a *plabic graph* is a planar undirected graph G drawn within a disk with boundary points lying on the circle incident to exactly one edge, such that each internal vertex is one of two colors as shown in Figure 4a.

We can correspond every chessboard to a plabic graph by placing a colored vertex in every separate section of the chessboard and connecting two vertices if their corresponding sections are adjacent. This graph is clearly on a plane, and edges are non intersecting, so it is planar.

There are also transformations shown in Figure 5 that can be applied to plabic graphs

- (M1) Square move: If the graph contains a square formed by four vertices whose colors alternate as we go around the square, then we can switch colors of these four vertices.

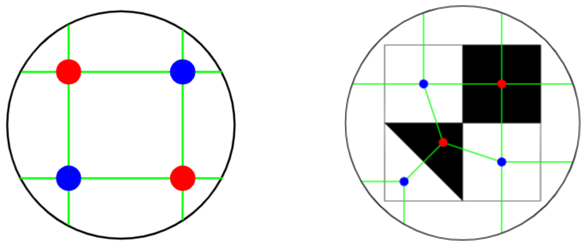

(a) Example plabic graph

(b) Plabic graph of a possible 2×2 grid

Figure 4: Plabic graphs

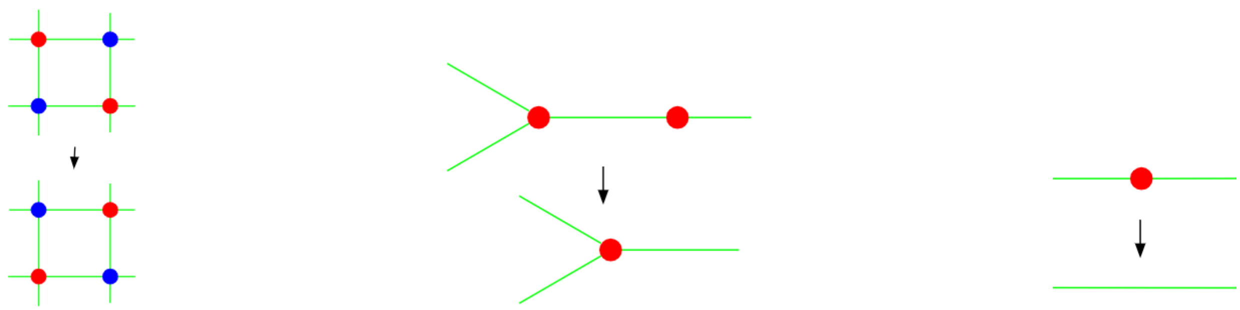

(a) Square move

(b) Unicolored edge contraction

(c) Middle vertex removal

Figure 5: Plabic transformations

- (M2) Unicolored edge contraction/uncontraction: If a graph contains an edge with two vertices of the same color, then we can contract this edge into a single vertex with the same color.
- (M3) Middle vertex insertion/removal: If a graph contains a vertex of degree 2, then we can remove this vertex.

We can correspond each move on the split chessboard with a combination of M1, M2, M3 in plabic graphs as shown in Figure 6.

2.2 Trips in plabic graphs

A **trip** in a plabic graph G is a path which starts at a boundary vertex and proceeds along the edges of G until it reaches another boundary vertex according to the following rules: when a red vertex is reached, the path takes a sharp left turn, and when a blue vertex is reached, the path takes a sharp right turn. An example of this can be seen in Figure 7. A trip permutation is the mapping of each boundary vertex to another in a plabic graph.

Theorem 2.1 ([1]). *Let G and G' be two reduced plabic graphs with the same number of boundary vertices. Then the following claims are equivalent*

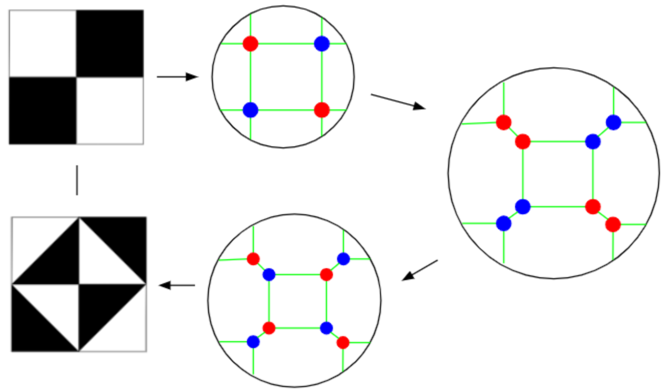

Figure 6: Split chessboard move in plabic transformations

- *G can be obtained from G' by moves M1, M2, M3.*
- *These two graphs have the same trip permutation.*

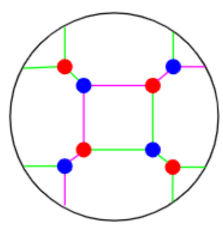

Figure 7: Highlighted trip on a plabic graph

3 Diameter of standard chessboard component

In order to provide intuition on the behavior of these mutations, we start from the smallest case of a $2 \times n$ grid.

Example 3.1. The diameter of the connected component containing the standard chessboard in the flip graph of a $2 \times n$ split chessboard is $n - 1$.

Proof. If a lattice point is mutated, none of the adjacent lattice points can be mutated because this would lead to two intersecting diagonals. Thus, to find the grid farthest from a standard chessboard, we choose to mutate every other lattice point. Because the width of the grid is 2, there is only one column of lattice points within the grid, hence we have two boards that are farthest from the standard chessboard (see Figure 3). Each mutation creates diagonals that can only be removed by reapplying the mutation. Thus, because no lattice point is mutated in both of the farthest boards, one must mutate every lattice point in the interior to transform one to the other. This is $n - 1$ mutations, which gives us the diameter. □

This gives us insight on how these mutations may behave with regards to parity.

Definition 3.1. Let us place the split chessboard on a coordinate plane such that the vertices lie on lattice points. We define *parity* of a vertex to be the parity of the sum of its coordinates. The *parity* of a diagonal edge will be the *opposite* of the parity of the sum of the coordinates of one of its endpoints. This is so that we produce and destroy even diagonal edges by mutating an even vertex, and vice versa.

Remark. In diagrams, even loops will be denoted with purple, and odd loops will be denoted with blue.

Definition 3.2. Let us denote by *loop* a collection of diagonal edges that form a simple closed curve. Note that every loop is necessarily of the same parity.

We can then define the possible states of a split chessboard.

Definition 3.3. Let R be a region bounded by the border of a split chessboard. A *closed diagonal pattern* D in R is a union of diagonal edges in R, such that it can be decomposed as a union of loops intersecting only at lattice points. A *weakly closed diagonal pattern* is a union of diagonal edges in R, such that its union with the boundary of R can be decomposed as a union of simple closed curves intersecting only at lattice points.

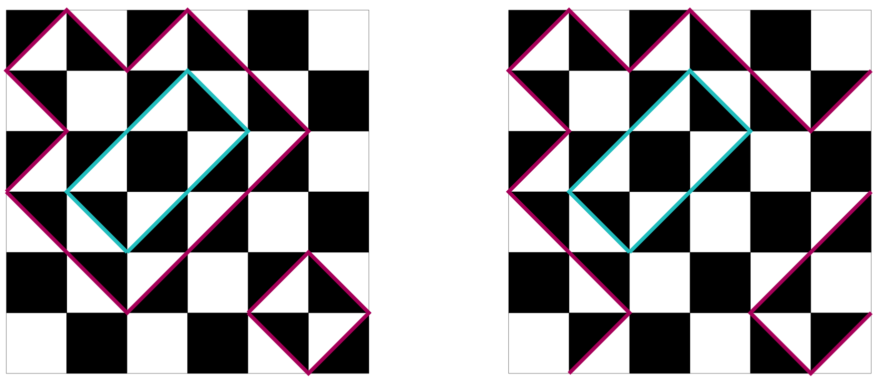

Figure 8: Closed diagonal pattern (left) and weakly closed diagonal pattern (right)

Lemma 3.1. *Let C be a split chessboard obtained through mutations from the standard chessboard C_0. Then the diagonal edges in C form a closed diagonal pattern.*

Proof. Let us prove this by induction. Applying one mutation on a standard grid results in a closed diagonal pattern. Now assuming that we have a closed diagonal pattern, we want to show that applying one more mutation will always create another closed diagonal pattern. If the new mutation creates or destroys a loop sharing no edges with a previously made loop, it clearly remains a closed diagonal pattern. If a mutated edge is shared with a preexisting loop, this edge disappears, and the remaining edges of the new mutation continue the closed curve in possible configurations as shown in Figure 9. □

Lemma 3.2. *Let C be a split chessboard corresponding to a closed diagonal pattern D consisting of loops of the same parity. Then C can be obtained through mutations from the standard chessboard C_0.*

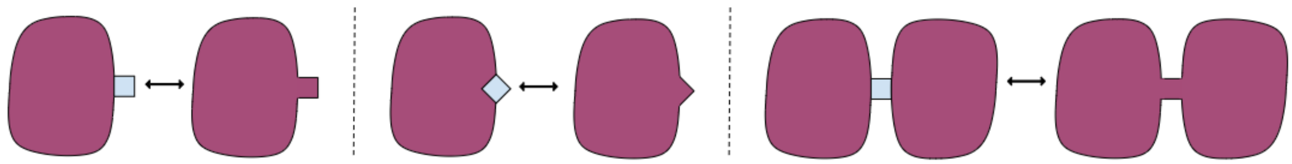

Figure 9: Mutating an adjacent cell to a loop

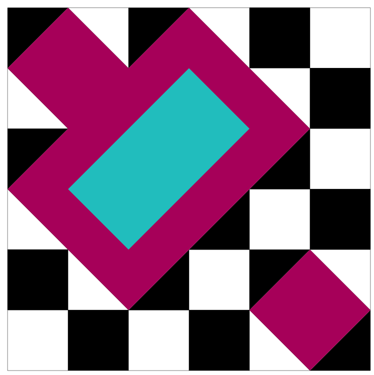

Figure 10: Shaded interior subregions in a closed diagonal pattern

Proof. The closed diagonal pattern partitions the chessboard region into nonintersecting subregions; i.e. connected components of $\mathbb{R}^2 \setminus D$. Consider any path from the boundary of the chessboard to a particular subregion that avoids lattice points; such a path must cross the edges of D either an even or odd number of times. Moreover, note that this parity does not depend on the chosen path. Therefore, we may correspondingly label a subregion as *exterior* or *interior*, respectively, as shown in Figure 10. Let us denote the union of interior subregions by $\text{int}(D)$ By applying the mutation on the vertices in the $\text{int}(D)$ with the same parity, we obtain the chessboard C. □

Let us call the algorithm described in Lemma 3.2 the *loop filling algorithm*. To calculate diameter, we need to calculate the shortest distance between any two split chessboards. Thus, we would like to show that the algorithm in the proof of Lemma 3.2 is the shortest possible path.

Lemma 3.3. *Let C and C_0 be as above. Then the shortest distance between C and C_0 is achieved using the loop filling algorithm.*

Proof. From above, we have a sequence of nested loops of alternating exterior and interior subregions associated to D. By Lemma 3.1, any set of mutations produces a closed diagonal pattern which has an outer boundary. Therefore, no mutations can be made in the outermost exterior ring without later being undone, because the outer boundary of the resulting closed diagonal pattern would then exceed the outer boundary of D. To obtain the boundary of the outermost interior region, each interior vertex in the region must be mutated at least once. Recursively applying this argument subsequent nested region proves the result. □

Though we have found an algorithm to create loops of the same parity, we must translate this to closed diagonal patterns containing loops of both parities. We divide a closed diagonal pattern into layers as follows:

1. The 0th layer is empty
2. The nth layer consists of the union of all loops that are each not contained within any loop of opposite parity after disregarding layers $< n$.

Lemma 3.4. *Let C be a split chessboard corresponding to a closed diagonal pattern. Then C can be obtained through mutations from the standard chessboard C_0.*

Proof. We can split C into layers as seen in Figure 11. Loops in separate layers have no intersections because one must be contained within the other. Then applying Lemma 3.2, we can mutate each of layers in order because they contain loops of the same parity. □

Proposition 3.5. *Given any split chessboard C in the connected component of C_0, the above algorithm achieves the shortest distance between C and C_0.*

Proof. From Lemma 3.2, we must mutate the interior subregions to create loops of the same parity. However, there can be closed diagonal patterns of the opposite parity within the subregions. Hence, layer n must be mutated before layer $n+1$ and the result then follows from applying Lemma 3.3. □

Definition 3.4. The *maximally split chessboards* are the chessboards created from maximizing the size of each nested layer.

We can see the two maximally split chessboards for a 6×6 split chessboard in Figure 11.

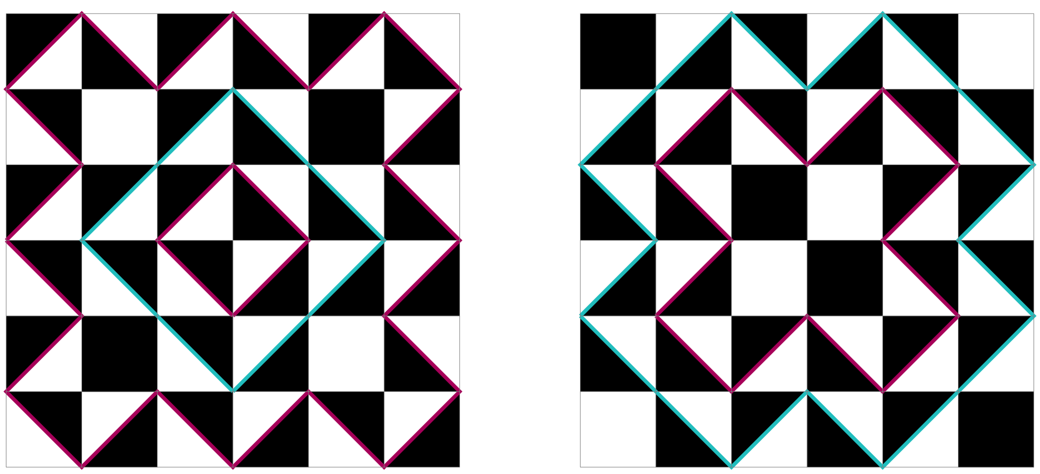

Figure 11: Maximally split chessboards

Lemma 3.6. *The maximally split chessboards M are the split chessboards farthest from the standard chessboard.*

Proof. The outer loop of each of the chessboards is the maximal even or odd interior region that can be created in the first layer. But this also maximizes the potential size of the next innermost layer. By iterating this argument we see that both maximally split chessboards maximize the number of layers and the size of each layer given the outermost loop. □

Theorem 3.7. *The diameter of the component containing the chessboard is the distance between the two maximally split chessboards.*

Proof. As shown in Proposition 3.5, layer n must be mutated before $n+1$. Hence, because the parities of the two maximally split chessboards are different, the path between the two must be mutating layer by layer. For the sake of contradiction, let there be two split chessboards that are farther apart. If d_1 and d_2 are their distances from the standard chessboard, the distance between each other is at most $d_1 + d_2$. However, this is less than the distance between maximally split chessboards, which then must be the diameter. □

4 Other connected components

To find the diameter of the entire flip graph, we need to examine connected components not containing the standard chessboard as well.

Definition 4.1. An anchor is a boundary point where exactly one diagonal edge of the diagonal pattern meets the boundary of the split chessboard.

Theorem 4.1. *Suppose C_1 and C_2 are split chessboards with the same anchor points. Then there is a sequence of mutations connecting C_1 and C_2.*

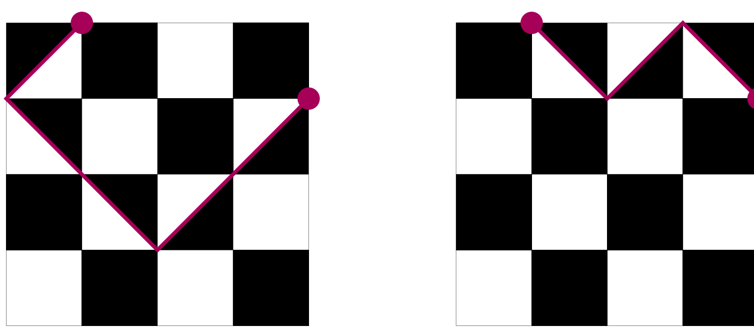

Figure 12: Two connected split chessboards with the same anchors

Let D_1 and D_2 be the weakly closed diagonal patterns corresponding to C_1 and C_2.

Firstly, by applying Lemma 3.4, we may assume without loss of generality that D_1 and D_2 contain no closed curves.

Let us use the notation:

$$D_\cap := D_1 \cap D_2$$
$$D_\cup := D_1 \cup D_2$$
$$D_\ominus := \overline{D_1 \ominus D_2}.$$

Here we are treating everything as subsets of $\mathbb{R}^2$; the overbar refers to topological closure since we want $D_\ominus$ to include the endpoints of its constituent edges. Let us use the notation $(-)^{\text{ev}}$ and $(-)^{\text{od}}$ to refer to operations of taking only the even and odd edges, respectively. Note that $D_\ominus^{\text{ev}}$ and $D_\ominus^{\text{od}}$ are closed diagonal patterns of a single parity, so have corresponding inner regions $\text{int}(D_\ominus^{\text{ev}})$ and $\text{int}(D_\ominus^{\text{od}})$.

Intuitively, we want to mutate the interiors of the "difference" $D_\ominus^{\text{ev}}$ and $D_\ominus^{\text{od}}$ to get from D_1 to D_2. However, it is possible that the interior regions overlap between $D_\ominus^{\text{ev}}$ and $D_\ominus^{\text{od}}$, which means we cannot arbitrarily mutate the interiors in any order. To formalize the restrictions on the mutation order, we construct a colored directed planar graph G as follows:

Construction 4.1.

- For each connected component of $\text{int}(D_\ominus^{\text{ev}}) \setminus D_\cap^{\text{ev}}$, a purple vertex of G.
- For each connected component of $\text{int}(D_\ominus^{\text{od}}) \setminus D_\cap^{\text{od}}$, a blue vertex of G.
- For each intersection between connected components R^{ev} and R^{od}:
  - If an edge e_1 of D_1^{ev} intersects an edge e_2 of D_2^{od} on the boundaries of R^{ev} and R^{od}, respectively, an edge $R^{\text{ev}} \to R^{\text{od}}$.
  - If an edge e_1 of D_1^{od} intersects an edge e_2 of D_2^{ev} on the boundaries of R^{od} and R^{ev}, respectively, an edge $R^{\text{od}} \to R^{\text{ev}}$.

The edges defined above encode the fact that e_1 must be unmutated before e_2 can be mutated—see Figure 13.

A priori, G is a multigraph, although we will soon show that this is not possible.

Given anchor points, we can create a graph as shown above for any two diagonal patterns sharing these anchor points. Proposition 4.2 below implies that we can order mutations to transform one diagonal pattern to the other, thus showing Theorem 4.1

Proposition 4.2. *The vertices of G admit a partial order satisfying $R < R'$ when there is a edge $R \to R'$.*

Proof. Suppose towards contradiction that there exists a cycle $R^1 \to R^2 \to \cdots \to R^n \to R^1$. This corresponds to an annulus-shaped subregion of $\text{int}(D^{\text{ev}}) \cup \text{int}(D^{\text{od}})$. Without loss of generality, we may assume this is the innermost of possibly nested annuli. Then there are no cycles in the center of the annulus, so there exists an ordering in which we can mutate the components of $\text{int}(D_\ominus^{\text{ev}})$ and $\text{int}(D_\ominus^{\text{od}})$ in the center of the annulus to get from D_1 to D_2. Therefore, by performing these mutations first, we can further assume without loss of generality that there are no regions of $\text{int}(D_\ominus^{\text{ev}})$ or $\text{int}(D_\ominus^{\text{od}})$ in the center of the annulus.

Now let us consider paths in $D_\cap$ where by "path" we mean a connected component of $D_\cap \setminus D_\ominus$. By construction, each path ends at either a point of $D_\ominus \cap D_\cap$, or at the chessboard boundary, and paths do not pass through the interior regions $\text{int}(D_\ominus^{\text{ev}}) \cup \text{int}(D_\ominus^{\text{od}})$. Note that there must be exactly one such path emanating from each point of $D_\ominus \cap D_\cap$, because each lattice point must be connected to either two or four diagonal edges—see Figure 14. Moreover, there must be an odd number of such points on the boundary of each R^i on the

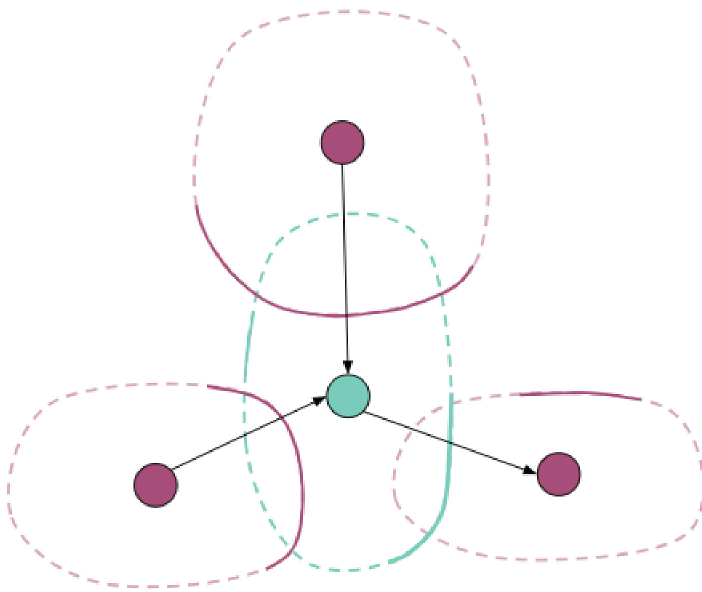

Figure 13: Construction of planar graph according to construction 4.1. Each loop represents a connected component of $\text{int}(D^{\text{ev}})$ or $\text{int}(D^{\text{od}})$, the solid sections are D_1 and the dashed sections are D_2.

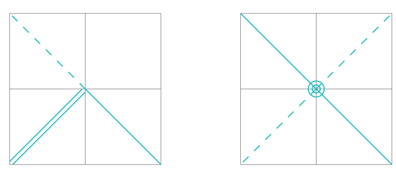

Figure 14: Possible configurations of paths emanating from a point of $D_\ominus \cap D_\cap$. Solid edges are in D_1, dashed edges are only in D_2, and double edges are in $D_\cap$. In the right hand case, we consider the center point to be a length-0 path in $D_\cap$.

inner side of the annulus because overlapping regions must have opposite parity as shown in Figure 15. Then in order for those paths to have endpoints, an odd path must cross an even path. But since these paths consist of edges in both D_1 and D_2, this would imply that an odd diagonal edge overlaps an even diagonal edge a chessboard pattern, which is impossible! Thus we obtain a contradiction. □

Thus, every split chessboard with a set of given anchor points is connected to every other. Therefore, we can study the diameter of such connected components.

5 Conclusion

We establish an algorithm of mutations to find the diameter of the connected component of the flip graph containing the standard chessboard for grids with boundaries containing

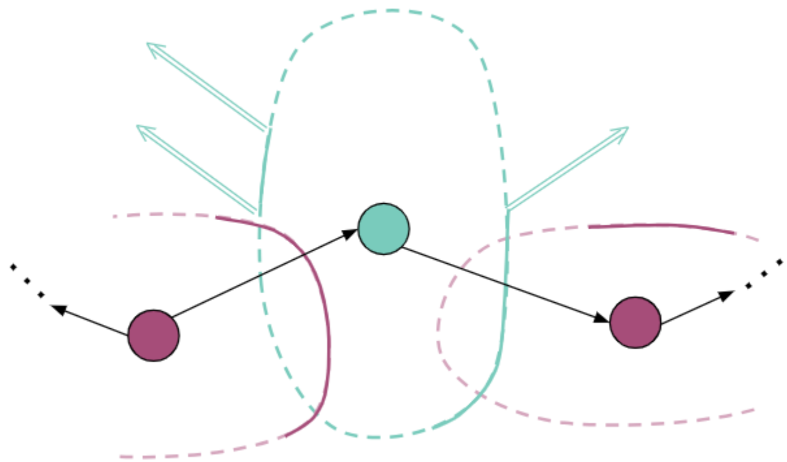

Figure 15: Odd number of paths in $D_\cap$ emanating from a region on one side of the annulus.

only vertical and horizontal edges. We find that each connected component of the flip graphs share the same anchor points.

Further work includes examining the behavior of the diameter in connected components not containing the standard chessboard. Additionally, edge cases should be visited with regards to non vertical and horizontal boundaries or concavities.

6 Acknowledgments

I would like to thank my mentor Tristan Yang for helping and guiding me through the whole project, my head mentor Dr. Tanya Khovanova for giving me great feedback and advice, and my tutor Peter Gaydarov for teaching me how to write and present my research. I would also like to thank my last week TA Allen Lin for giving feedback on my paper as well as my first week TAs AnAn Desimone, Rich Wang, and Sally Zhu for helping with formatting. Furthermore, I would like to thank my friends and family for always supporting me. Lastly, I would like to thank the Center for Education in Excellence, the Research Science Institute, any sponsors, and MIT for giving me this opportunity and assisting with this project.

References

[1] A. Postnikov. Total positivity, grassmannians, and networks, 2006. math/0609764.