Crossing Number Bounds in Knot Mosaics

Hugh Howards, Andrew Kobin

October 17, 2014

Abstract

Knot mosaics are used to model physical quantum states. The mosaic number of a knot is the smallest integer m such that the knot can be represented as a knot m-mosaic. In this paper we establish an upper bound for the crossing number of a knot in terms of the mosaic number. Given an m-mosaic and any knot K that is represented on the mosaic, its crossing number c is bounded above by $(m-2)^2 - 2$ if m is odd, and by $(m-2)^2 - (m-3)$ if m is even. In the process we develop a useful new tool called the dual of the mosaic.

1 Introduction

In [7], Lomonaco and Kauffman introduce a standard system of knot mosaics as a model of physical quantum states. In this paper we introduce a new tool for analyzing mosaics, a dual to the mosaic, denoted D, together with T an ordered triple associated to D. Mosaics contain 5 distinct tiles, up to rotation, and all 11 orientations are shown below. We label the tiles with roman numerals for the 5 types and when applicable the letters a though d for the distinct rotations of those types. We also introduce a type 0 tile which consists of a square with a dot in the center. Type 0 tiles are not a part of a mosaic, and have not previously been used in the literature but will be used when we define the dual.

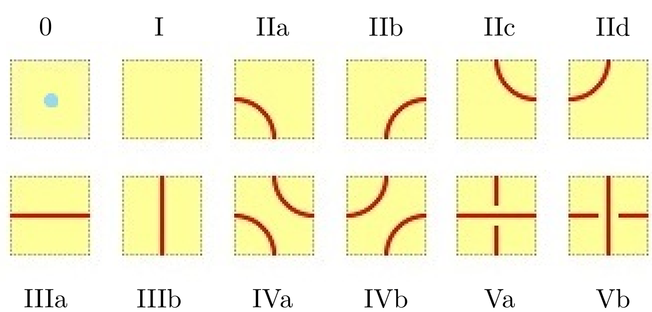

For a positive integer n, define an *n-mosaic* M_n as an $n \times n$ matrix composed of mosaic tiles. When specifying a given tile, we model subscripts after the entries in a matrix so the tile $R_{i,j}$ will refer to the tile in the i^{th} row and j^{th} column, where columns are counted from the left and rows are counted **from the bottom** (we diverge from matrix notation slightly here to allow row numbers to reflect a height function). If we think of M_n as a square disk, the $4n-4$ tiles that intersect ∂M_n will be called *boundary tiles* and the other $(n-2)^2$ tiles will be referred to as the *interior* of the mosaic and denoted S. See Figure 2.3 for a depiction of S and the boundary tiles in a mosaic. We will often be focused on S.

A *link n-mosaic* is an n-mosaic with all of its tiles suitably connected so that after all the tiles are placed on the mosaic, the result is a projection of a link. In such a mosaic, tiles of types II, III, IV, and V must have adjacent tiles that are also of one of these four types to extend the arcs started on those tiles.

A *knot n-mosaic* is a link n-mosaic that corresponds to a projection of a one component link (a knot). Thus every knot mosaic is a link mosaic, but not every link mosaic is a knot mosaic. Define *the mosaic number* of a knot (or link) to be the smallest natural number m such that the knot (link) is representable as a knot (link) m-mosaic.

An important motivation for studying knot mosaics is the Lomonaco-Kauffman Conjecture, which states that knot mosaic theory is equivalent to classic (tame) knot theory. This was proven by Kuriya in [5], so as a result, we can treat knot equivalence and crossing number of knot mosaics in the usual way.

We introduce the concept of a *dual* for a mosaic in this paper. Given a mosaic M_n for link L (or knot K), we define the dual on S, the inner tiles of M_n. The dual D is obtained by replacing the tiles on S in the following manner: every type V is replaced with a type I, type IV with type 0, type IIIa with type IIIb, type IIIb with type IIIa, type IIa with type IIc, type IIc with type IIa, type IIb with type IId, type IId with type IIb, and finally type I with type IV (although this means the dual is not uniquely defined, either type IVa or IVb is fine). Intuitively, the dual is like a complement of the link in the mosaic in the sense that the union of the dual and the mosaic will intersect the boundary of each tile in S exactly four times - once on each of the tile's edges. Most importantly, although the definition of the dual is not quite unique, there is an inverse function associated to it that is unique (reverse the orders in the definition above), so the portion of the mosaic contained in S may be deduced from D.

The dual consists of loops, type 0 tiles, and arcs with both endpoints on ∂S that we will call edges. The term arc will be used to refer to a subset of a loop or an edge. *The constant* $|D|$ is defined to be the total number of tiles in D (excluding the blank type I tiles, of course). The tile $T_{i,j}$ refers to the dual tile in the i^{th} row and j^{th} column. **Throughout the paper $R_{i,j}$ refers to the knot and $T_{i,j}$ refers to the dual.** Although $T_{i,j}$ is defined by looking at $R_{i,j}$, we will be focused on the dual in most of our arguments so we will almost always refer directly to $T_{i,j}$.

We next define an ordered triple T, which is computed from the dual of a given mosaic. Let D' be the complement in D of all the type 0 tiles in D. Let D'' be the set of all tiles $T_{i,j}$ in D that are of type IV. Let $l = |D|$, $l' = |D'|$ and $l'' = |D''|$. We define the ordered triple $T = (l, l', l'')$. Notice that although the dual is not unique for a given mosaic, the only choices were which of the two type IV tiles to pick and so any choice of dual for this mosaic will give the same ordered triple T.

Of all possible ways to embed a specific knot K on an $n \times n$ mosaic, let M_n be a mosaic that minimizes the ordered triple T lexicographically (alphabetically). For example if K can be built with with a dual yielding $T_1 = (7, 4, 2)$ or with a dual yielding $T_2 = (8, 0, 0)$ we pick the first embedding since $T_1 < T_2$ lexicographically. We say in such a case that T is *minimal*.

Note that D' consists of a set of loops and properly embedded edges on a square disk S. Another way to think of D' is this: all type 0 tiles in the dual correspond to smoothing one crossing from a mosaic that contained one more crossing. Then D' is just the dual of the knot or link mosaic which has none of these crossings smoothed.

We begin by introducing a theorem (Theorem 3.1 in [3]) which is the natural starting place for a discussion of the relation between crossing number and mosaic number.

Theorem 1.1 (Upper Bound on Crossing Number)**.** *[3] Given an m-mosaic, if a knot is representable on the mosaic then its crossing number is bounded above by: $c \leq (m-2)^2$*

Using this bound, one can quickly show the mosaic numbers for many knots. However, it is clear that this bound is of limited use when it comes to determining the mosaic number of more complex knots. In this paper, we sharpen the upper bound on crossing number by proving

Theorem 7.2. Given an m-mosaic and any knot K that is projected onto the mosaic, the crossing number c of K is bounded above by the following:

$$c \leq \begin{cases} (m-2)^2 - 2 & \text{if } m = 2k+1 \\ (m-2)^2 - (m-3) & \text{if } m = 2k. \end{cases}$$

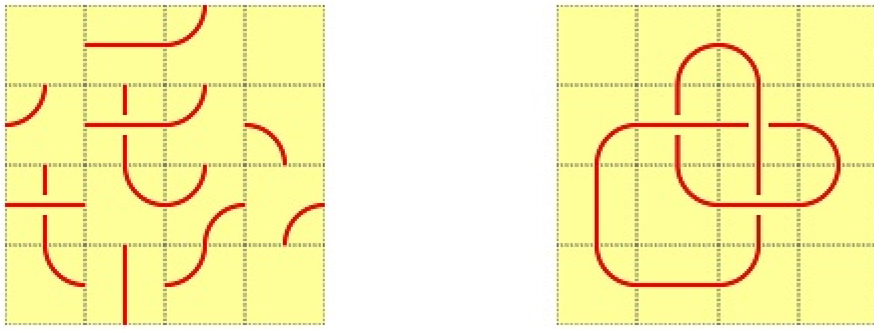

Figure 1.1: Pictured above are two 4×4 mosaics, but only the one on the right is suitably connected to yield the projection of a knot.

2 Saturation and duals

A mosaic is said to be *saturated* if every tile of S, the interior of the mosaic, is a crossing tile. In this case $D = \emptyset$. Conversely, if a link mosaic has a nonempty dual, it is not saturated. Theorem 7.2 stated above shows that the even and odd knot mosaic boards have radically different properties regarding how "nearly saturated" they can be.

Lemma 2.1. *In a link mosaic, boundary tiles cannot be crossing tiles. Therefore all crossings of a link must fit on S, the interior of the mosaic.*

Proof. For all tiles to be suitably connected, every edge containing a strand must be adjacent to another edge containing a strand, and for crossing tiles on the boundary, this is not the case (for example, see Fig. 1.1). Therefore boundary tiles cannot be crossing tiles. □

Theorem 2.2. *A saturated mosaic cannot contain the projection of a knot that achieves its crossing number.*

Proof. Start by filling S with type V tiles. The proof will not depend on if we choose type Va or Vb so at this stage we have not lost generality of the argument no matter which type V tiles we choose. Now we notice that since edges intersecting ∂S can only connect to one of its two adjacent edges in ∂S there are only two choices of how to connect up the strands through the boundary tiles to get a knot or link. The vertical strand in $R_{2,3}$, for example, must either connect to the vertical strand in $R_{2,2}$ as it does in Figure 2.6 or $R_{2,4}$ as in Figure 2.5. In the first case this means tile $R_{1,3}$ is a type IId and $R_{1,2}$ is type IIc and in the second case tile $R_{1,3}$ is a type IIc and $R_{1,4}$ is a type IId, again both matching the examples in Figures 2.6 and 2.5 respectively. On a saturated board once this single choice has been made the rest of the choices on the outside are uniquely determined.

Suppose K is a knot such that its mosaic representation on M_n is saturated. If the mosaic (shown in Fig. 2.5) contains any crossings easily removed by a type I Reidemeister move (nugatory crossings), the number of crossings in the mosaic exceeds the crossing number of K.

If n is even, then one of the two ways of connecting up through the boundary tiles results in nugatory crossings in each of the four corners showing that crossing number is not achieved. Thus if it achieves its crossing number in this projection it must be attached in the other way, but then as Fig. 2.6 illustrates, we actually have a link of $n - 2$ components, and not a single knot as we would like.

Now suppose n is odd. Here either way we choose the tiles on the outside we get exactly two nugatory crossings in corners that can be removed with type I Reidemeister moves. As shown in Fig. 2.7 selecting not to have a nugatory crossing in the northwest corner forces nugatory crossings in the northeast and southwest corners. This shows that an odd, saturated mosaic has

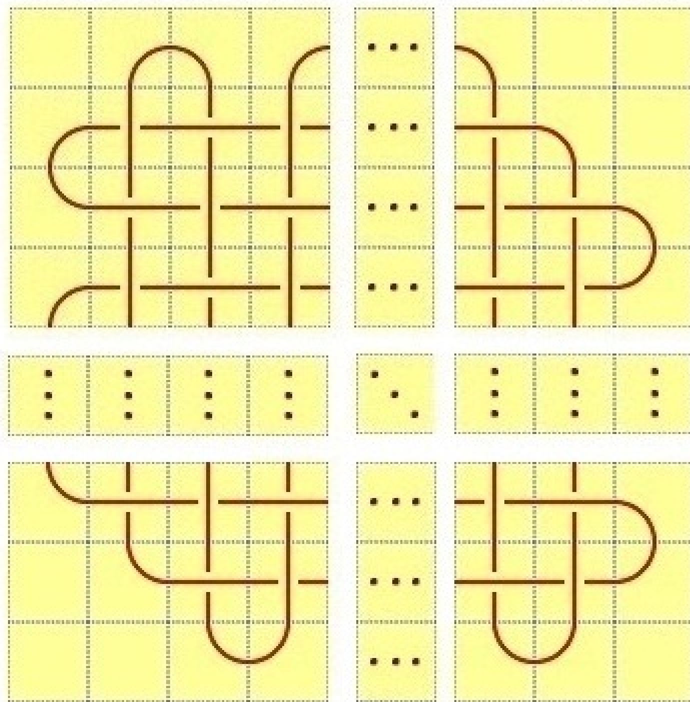

Figure 2.1: An odd mosaic with near-saturation of degree 2.

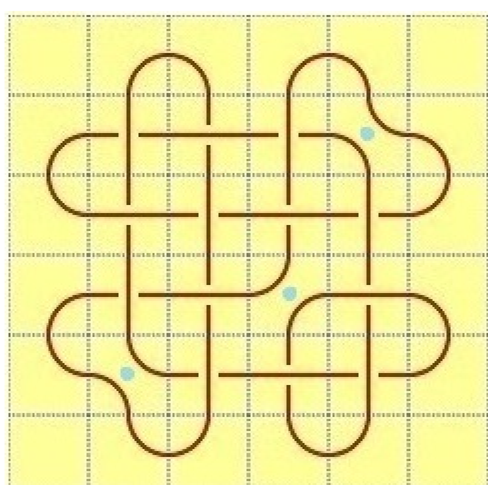

Figure 2.2: There are three type 0 tiles in this dual on M_6. One could think of the mosaic as a saturated mosaic with three crossings smoothed reducing the number of components in the link from 4 to 1.

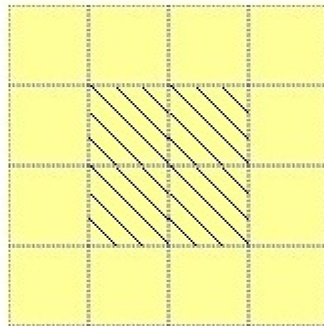

Figure 2.3: In a 4×4 mosaic the the interior S is the shaded 2×2 sub-mosaic.

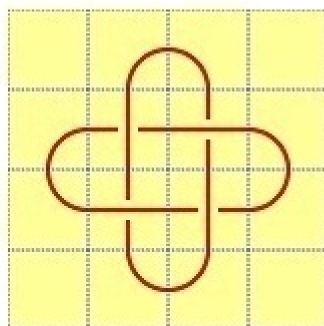

Figure 2.4: A projection of Solomon's link on a saturated mosaic (M_4).

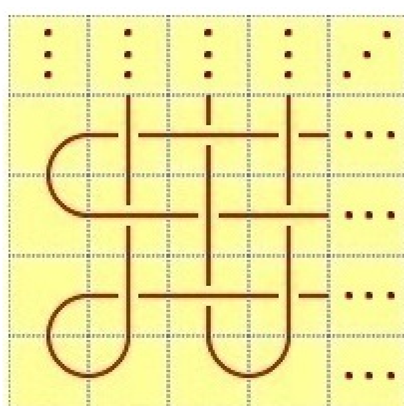

Figure 2.5: A nugatory crossing in the corner.

two nugatory crossings in the corners. Therefore crossing number is not achieved. In all cases, a saturated mosaic cannot contain the projection of a knot that achieves its crossing number. $\square$

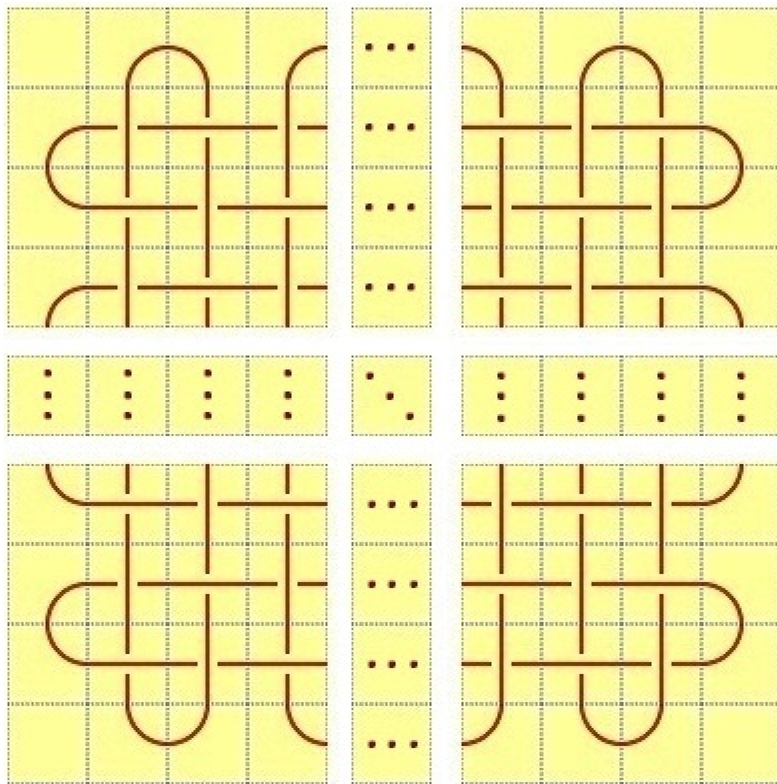

Figure 2.6: A link of $n-2$ components.

We next consider mosaics with duals consisting of a single tile, that is $|D| = 1$. Almost every knot will fail to achieve its crossing number on such mosaics, and the following lemma precedes a general theorem for mosaics with $|D| = 1$.

Lemma 2.3. *The trefoil knot has mosaic number 4.*

Proof. The trefoil knot has crossing number 3. Since a 3×3 board can only support one crossing, we must have at least a 4×4 board to have a non-trivial knot. Indeed, the mosaic on the right in Figure 1.1 shows that the trefoil knot may be embedded on M_4 and achieve its crossing number of 3. Therefore the trefoil knot has mosaic number 4. $\square$

Theorem 2.4. *Given a knot K with crossing number c, suppose its mosaic number m is odd. Then $c \leq (m-2)^2 - 2$.*

Proof. We showed above that an odd mosaic represents a knot with crossing number at most $(n-2)^2 - 1$. A mosaic with two interior tiles that are not crossing tiles (type V) will have $c \leq (m-2)^2 - 2$ so we are left to focus on the case of exactly one interior tile that is not a type V tile. If the knot K is achieved with only one interior tile that is not a crossing tile, then this means $|D| = 1$. This can only happen if D is either a single type 0 tile or if D is a single type II tile in one of the corners of S. If D is not in one of the corners of S, then M_n has two opposite corners with crossings that can be reduced by a type I Reidemeister move just as in a saturated odd board. Placing the dual in one of the corners can at most remove one of these nugatory crossings. Thus even though this embedding of K has $(n-2)^2 - 1$ crossings, K does have an embedding with $(n-2)^2 - 2$ crossings or less, bounding the crossing number from above. $\square$

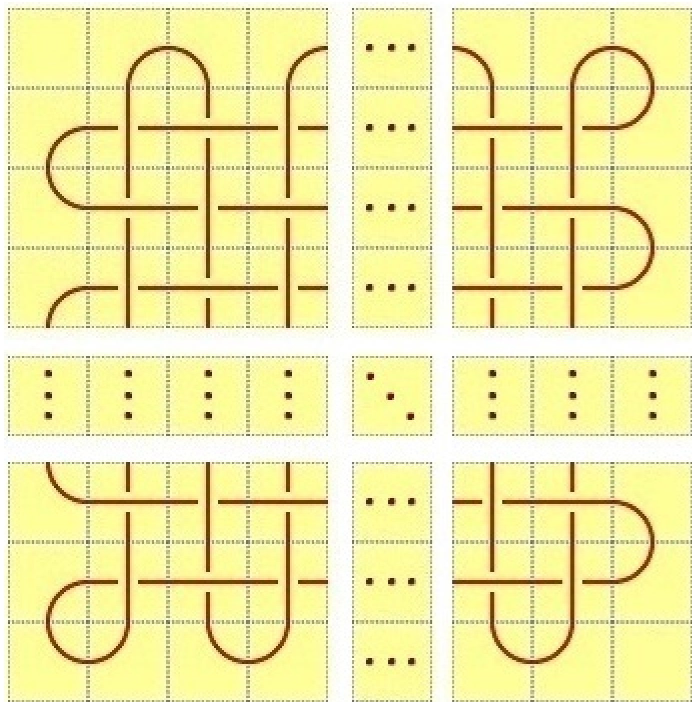

Figure 2.7: Odd mosaic with corners easily removed by type I Reidemeister moves..

Theorem 2.4 establishes the first part of the main theorem in this paper (Theorem 7.2). Moreover, an odd mosaic that is saturated and with the type V crossings chosen to give an alternating knot achieves this bound, so the bound is sharp. See Figure 2.1 for the same knot simplified via two type I Reidemeister moves to show that the knot achieves its crossing number.

Once the even case is established, it will follow that the trefoil is the only knot which can be constructed on a mosaic whose dual consists of a single tile. For the rest of the paper we focus on the even mosaic case M_{2k} where we assume T is minimal with respect to all $2k \times 2k$ knot mosaics giving knot K.

3 Loops in the dual

We now construct an argument showing that we can assume D contains no loops while keeping T minimal.

Lemma 3.1. *Let M_n be a mosaic for a knot K for which T is minimal and the number of loops in dual D is minimized over all such possible mosaics and for which $|D| \leq n - 4$. Let $\{c_1, c_2, \ldots, c_k\}$ be the set of loops in the dual. Then if the set of loops is not empty, at least one of the c_i's contains a type II tile.*

Proof. Since each loop has at least 4 corners, the only way for a loop to avoid a type II tile is if each corner is part of a type IV tile. If none of the corners are type II then the loop has at least 4 type IV corners that meet other components of the dual. If one of these corners is type IVa replace it with one that is type IVb; if not do the opposite. This swap yields the connect sum of the loop in the dual with another component of the dual, decreasing the number of loops in the dual by one. Since it does not change T and since it still yields a dual for K (the type IV corners are chosen arbitrarily) we see that the original dual did not meet the hypothesis of the lemma, a contradiction. Thus the corners of each loop may be assumed to be type II, an even stronger conclusion than the one type II tile in the lemma. $\square$

Lemma 3.2. *If some loop in the dual contains a type II tile then $T = (l, l', l'')$ is not minimal.*

Proof. If there is such a loop, call it c_1. Any time c_1 crosses the knot, K, we may dictate that it goes under K. By virtue of the definition of the dual, c_1 never crosses itself. Now remove c_1 from the dual and add it instead to the knot, replacing the knot mosaic with a link mosaic containing K and an unknot. Next place a type V tile into the mosaic (type V in the link, not the dual) where the type II tile of c_1 had been. Because c_1 was entirely below K, and c_1 had no crossings with itself, we have just taken the connect sum of K with an unknot. Thus we have a new mosaic representing K, but $|D|$ has dropped contradicting the minimality of the ordered triple T in the original mosaic. □

In the proof above we found an unknot in the dual that contained a type II tile, swapped it out of the dual and into the mosaic and changed the type II tile to a type V, yielding a connect sum of K with an unknot resulting in a new embedding of K and lowering T. We will repeat this process multiple times in different contexts throughout the paper and we call the process *corner conversion.*

Lemmas 3.1 and 3.2 imply

Corollary 3.3. *Let M_n be a mosaic for knot K with $|D| \leq n - 4$ and for which T achieves its minimum over all such mosaics. We may then choose M_n so that simultaneously D contains no loops and T is minimal.*

We now want to look at a particular class of loops. These are the shortest possible loops: ones of length 4 (contained in exactly 4 tiles) coming from a combination of 4 tiles of types II and IV. We call these short loops *bubbles.* We see a bubble in each of the pictures in Figure 3.1.

The proof above showed that the dual may be chosen to contain no loops if T is minimal, but it did not show that a dual containing loops could not also be minimal. Later we may start with a dual that contains no loops and use moves that create bubbles, which we then want to argue is impossible, so we need a stronger result saying that if T is minimal the dual cannot contain bubbles. In the argument it will be useful to have the following lemma that gives us some flexibility in where a bubble might be positioned.

Lemma 3.4. *Given a knot mosaic M_n for knot K with dual D and ordered triple T, if M_n contains any $2 \times n$ or $n \times 2$ subset of tiles, $n \geq 2$, that consists of exclusively type IV tiles, then we may pick any 2×2 subset of these tiles and form a (possibly) new dual for M_n and K in which there is a bubble in the 2×2 subset so that T is unchanged for the new dual.*

Proof. The proof is easy. Pick each of the four tiles to be IVa or IVb so that they have a bubble within them. Leave the other type IV tiles unchanged. Since swapping type IV tiles in the dual doesn't affect K or T, the new dual has shifted the bubble to the desired location and satisfies our requirements on K and T. □

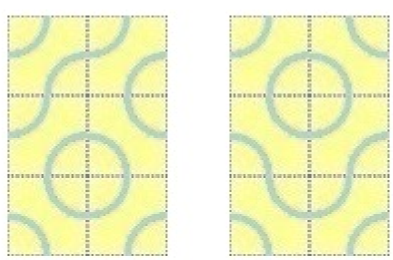

Figure 3.1: A bubble percolates up through a 3×2 set of type IV tiles.

This process allows us to shift the location of an existing bubble through nearby type IV tiles. We call this process of shifting a bubble to a new location *percolation.* A 3×2 example is shown in Figure 3.1.

Theorem 3.5. *Let M_n be a mosaic for knot K for which T is minimal. D cannot contain a bubble.*

Proof. If we ever have a bubble which contains a type II tile, then we can do a corner conversion as we did in Lemma 3.2, yielding a new embedding of K but lowering T. This contradicts the fact that T was minimal. Thus we proceed with the argument under the assumption that the bubble is entirely contained in type IV tiles.

Next we show the dual is not minimal if the knot intersects either a row or a column of S containing the bubble. In this case, the ability to rotate the mosaic assures us that we may assume that the intersection is in a column above the bubble. If there are any type IV tiles below the knot in those columns, but above the bubble we use Lemma 3.4 to shift the bubble up to the four tiles directly below the knot.

Explicitly, if K intersects $T_{i+1,s} \cup T_{i+1,s+1}$ we let the bubble be contained in tiles $T_{i-1,s}$, $T_{i,s}$, $T_{i-1,s+1}$, and $T_{i,s+1}$. Without loss of generality let K intersect $T_{i+1,s}$ (and possibly $T_{i+1,s+1}$, too). Because it is directly above a type IV tile, but it contains part of the knot, $T_{i+1,s}$ is either type IIb, IIa, or IIIb (the knot cannot intersect the bottome edge of the tile). $T_{i+1,s+1}$ is also above a type IV tile and must pair with $T_{i+1,s}$. This means $T_{i+1,s+1}$ must be a IIa tile or a type IV tile if $T_{i+1,s}$ is type IIb, and $T_{i+1,s+1}$ must be IIb tile or a type IIIb if $T_{i+1,s}$ is IIa or IIIb.

We show moves in Figure 3.2 for six possible combinations that allow us to connect sum the bubble with the knot and reduce T contradicting minimality – in the case of a type IV in $T_{i+1,s+1}$ we show only type IVb case since the IVa case is nearly identical.

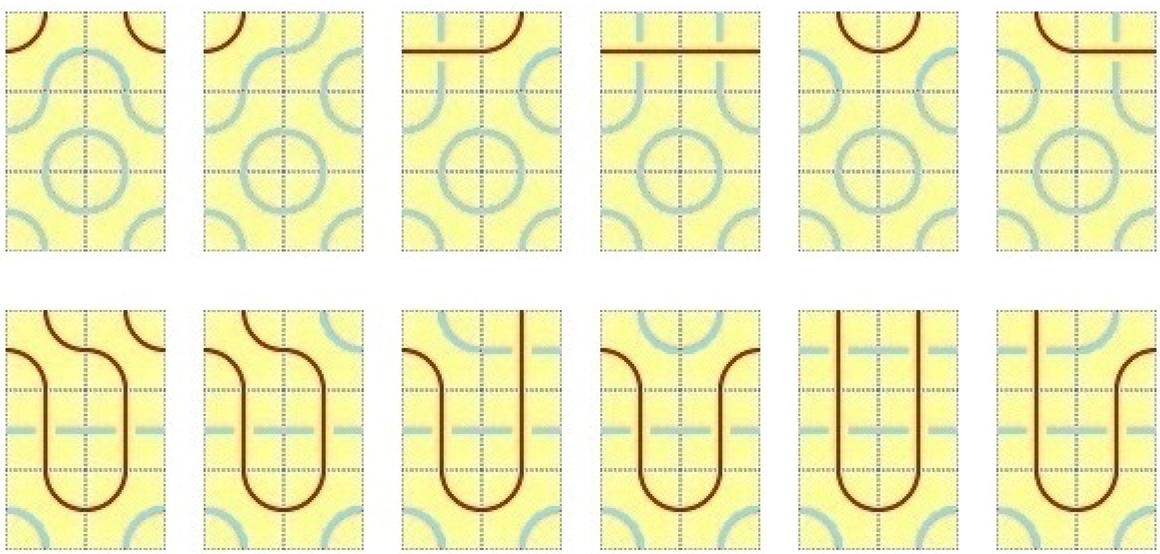

Figure 3.2: A bubble can never appear in the dual when T is minimal. Here we see that if a tile above the bubble contains an arc of K, there is always an embedding of K that decreases T and 'bursts' the bubble.

Finally we are left with the case where neither the columns nor the rows of S containing the bubble intersect the knot. Thus they are exclusively full of type IV tiles. Since K exists and in any interesting case has at least one crossing, we know that some row of S intersects K. (Of course the unknot fits on a 2×2 board with $S = \emptyset$ so we are only interested in knots with positive crossing number.)

By Lemma 3.4 we can percolate the bubble within the columns containing it to make it intersect the row that already contains part of K. Now as before we have not changed K or T. We have, however, reduced to the previous case, which shows T can be reduced without changing K, contradicting minimality. $\square$

4 Edges in the dual

Since we now know that we can get rid of loops in a dual without increasing T we turn our attention to edges.

Lemma 4.1. *If $|D| \leq n - 4$ then no edge in D runs from one side S to the opposite side.*

Proof. The proof is trivial since such an edge must be of length at least $n - 2$. $\square$

This implies that if e is an edge of the dual, then the endpoints will have to either be on adjacent sides, as in edges e_1, e_2 and e_3 in Figure 6.3, or on the same side of S, i.e. starting and ending in the exact same row or column as in edges f_1, f_2, f_3 and f_4 in the same figure. If the endpoints are on the same side as each other we call it an *XX-edge* and if adjacent sides we call it an *XY-edge*.

Both XX and XY-edges cut S into two disks. The smaller side is considered the outside of the edge. In topology we may not have a metric so we often avoid talking about the smaller part of a disk, but a mosaic as an $n \times n$ subset of the plane has a natural metric on it so we are free to use the term.

Because the edges in the dual are relatively short, if e is an XX-edge then the boundary of the outside (smaller) disk consists of e together with part of one side of S. Likewise if e is an XY-edge the outside consists of e together with parts of two sides of S. Thus for any arc e we have a notion of *outside*. An edge $e \subset D'$ is called *outermost* in D' if there are no edges outside of e on S in D'.

Note that if D contains no loops – which Corollary 3.3 allows us to assume – and $D' \neq \emptyset$, so there is at least one edge, then there must be an outermost edge e. Also note that our definition is not quite the same as the traditional definition of outermost arcs on disks in topology. If e is outermost in our context it is outermost in the traditional definition, but not every traditionally outermost arc is outermost in our definition because it might be outermost, but on the wrong side (the side of its larger disk). An edge can still be outermost even if there is a type 0 tile of D outside of it.

5 Reduction moves

We establish a set of moves that when applied to the arcs of the dual will reduce the ordered triple T without changing the isotopy class of the knot. One primary use of the moves is to lower an arc a of the dual that represents a local maximum (possibly after rotating the mosaic). This will eventually lead to the conclusion first that no such moves can be made to the edges of D, and then after a further argument that D contains no arcs at all if T is minimal. We define the moves starting with the more elementary moves.

5.1 The type IV moves: bubble release and XX-through-XY moves

In the proof of Lemma 3.1 we swapped type IV tiles to reduce the number of loops in a dual by taking the connect sum of a loop with another part of the dual. We now consider the inverse operation on the dual when it would create a bubble. We define a bubble release move when we swap a type IVa tile for IVb or vice-versa to yield a bubble without altering K or T. Such a move is pictured in Figure 5.1.

We know, however, by Theorem 3.5 that a minimal dual can never contain a bubble and T is unchanged by a bubble release move so we see immediately the following lemma.

Lemma 5.1. *If T is minimal then D cannot contain tiles on which we can perform a bubble release move.*

The other type IV move is the XX-through-XY move. Let D contain e_1 an XY-edge that has e_2 an XX-edge outside of it such that the two edges share a type IV tile. Switching the tile from IVa to IVb or vice-versa will, of course, have no effect on T or K, but will replace e_1 and

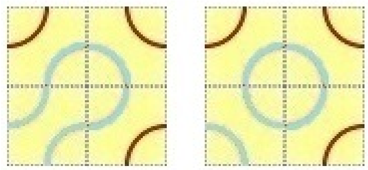

Figure 5.1: Neither the embedding of the knot nor T are altered when we break a bubble off of an arc of the dual using the bubble release move. Note that the move is identical if any of the type II dual tiles are swapped for type IV tiles.

e_2 with a new XY-edge and a new XX-edge. Call the XY-edge e_1' and the XX-edge e_2'. The move reduces the overall number of XX-edges outside of XY-edges. In particular at the very least e_2' is not outside of e_1' and e_1' has fewer XX-edges outside of it than e_1 did.

Iterating this process will eventually terminate since the number of type IV tiles is bounded by $|D|$.

We define a knot mosaic together with a dual D and associated ordered triple T to be *a minimal embedding* for a knot K if T is minimal, D contains no loops and in D no XX-through-XY moves are possible. Corollary 3.3 together with the process we have just described assures that every knot K that has a knot mosaic with $|D| \leq n-4$ has a minimal embedding. We call the dual in a minimal embedding a *minimal dual.*

5.2 Corner-corner moves

We describe this move in terms of an arc that acts as a local maximum for an edge and is moved downward. By symmetry, we can rotate the mosaic any multiple of 90 degrees or reflect along a horizontal or vertical line so the move is equally valid if the arc is a minimum and is moved upward, or one that is concave right and is moved to the right or concave left and is moved left.

Corner-corner move: Let e be an edge in D that intersects row i in an arc a that represents a local maximum for e. A local maximum must run directly across row $i, i > 2$, from a type IIb (or IVb) tile in column s to a type IIa (or IVa) tile in column t with $s < t$ as in Figure 5.2. We want to reduce the ordered triple $T = (l, l', l'')$ by moving part of the knot up across a and shorten a by moving it down. Figure 5.3 shows the basic move.

We pay close attention to any portion of the dual in tiles $T_{i-1,w}$ with $s \leq w \leq t$ (row $i-1$ directly under a). Since each of the tiles of a of the form $T_{i,w}$ with $s < w < t$ consists exclusively of type IIIa tiles, clearly those tiles of the form $T_{i-1,w}$ cannot ever be type IIc, IId IIIb, IV or V.

Certainly $T_{i-1,w}$ can be a Type 0 tile as shown, together with the corresponding corner-corner move in, Figure 5.3. On the other hand, if there is a tile in the dual $T_{i-1,w}$ with $s < w < t$ that is type IIa, IIb, or IIIa then the corner-corner move is undefined on a. Such examples are seen in the nested arcs in Figure 6.3. Another obstruction to the definition we can encounter is that if $T_{i-1,t}$ is type IId, IVa, or IVb. Symmetrically it is also undefined if the dual tile $T_{i-1,s}$ is type IIc, IVa, or IVb. We see arcs of this form in Figure 5.4. We will never need to use the corner-corner move in any of the undefined contexts, so the lack of definition here will not be a problem.

We now focus on the definition of the move in the situations where it can be applied. The move at its core just takes an arc a that is a local maximum for the dual and pushes it down one row when there is nothing from the dual already below it to block it. The exact prescription is given in two parts. For each w, $s < w < t$ we switch $T_{i-1,w}$ with $T_{i,w}$. This tells us how we apply the move to tiles that are between the corners of the arc, but not in the corners themselves. We now specify the move on the two corner tiles and the two tiles directly below them ($T_{i,s}$, $T_{i-1,s}$, $T_{i,t}$ and $T_{i-1,t}$). The corner tiles $T_{i,s}$ and $T_{i,t}$ are either type II tiles or type IV. When the move is defined tile $T_{i-1,s}$ must be type IIIb tile or type IIc. On the other corner, $T_{i-1,t}$ is either type

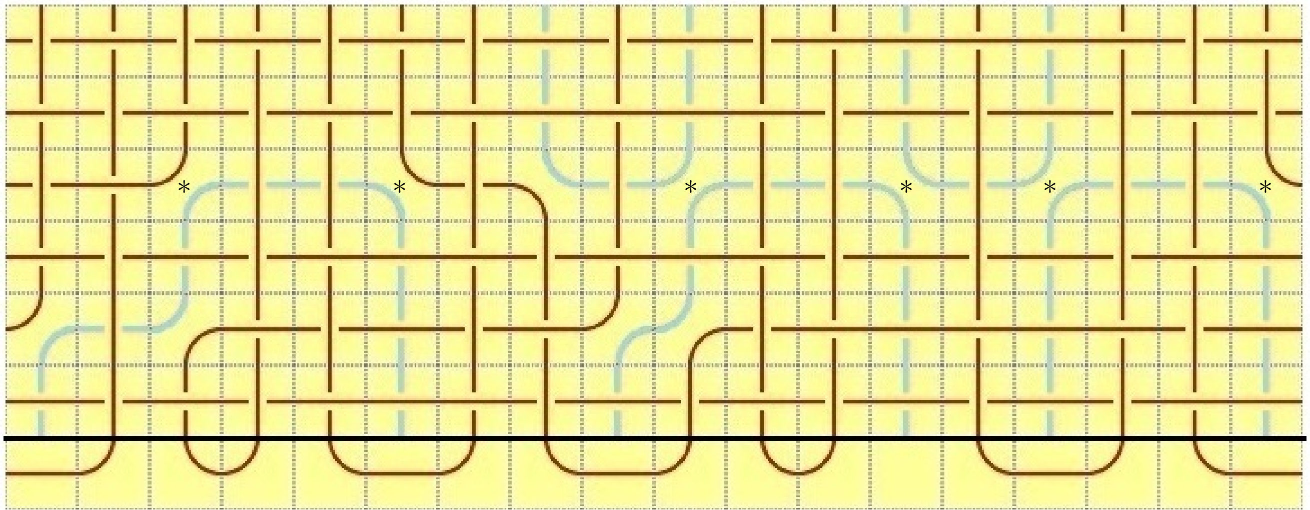

Figure 5.2: Each * denotes one of four types of corners possible in a corner-corner arc (two up to reflective symmetry).

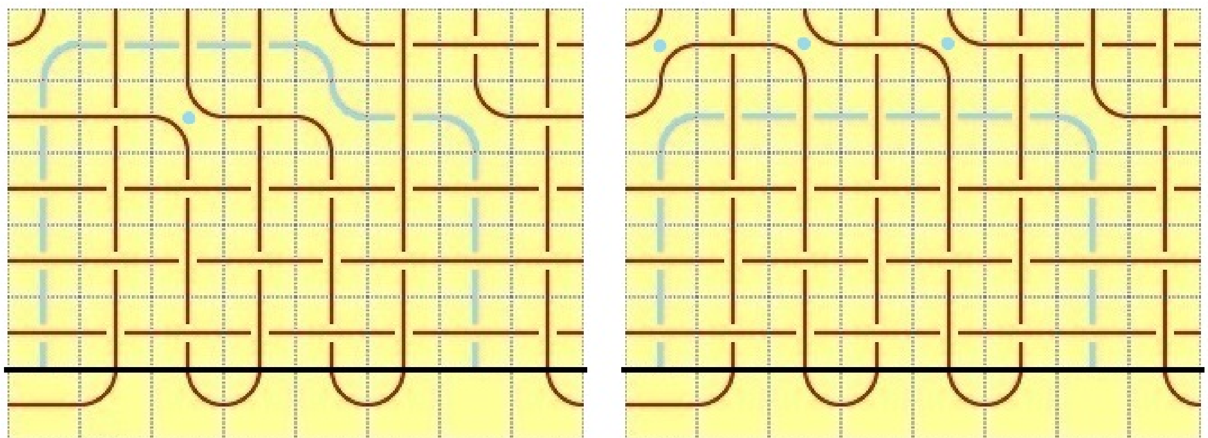

Figure 5.3: We see a basic corner-corner move.

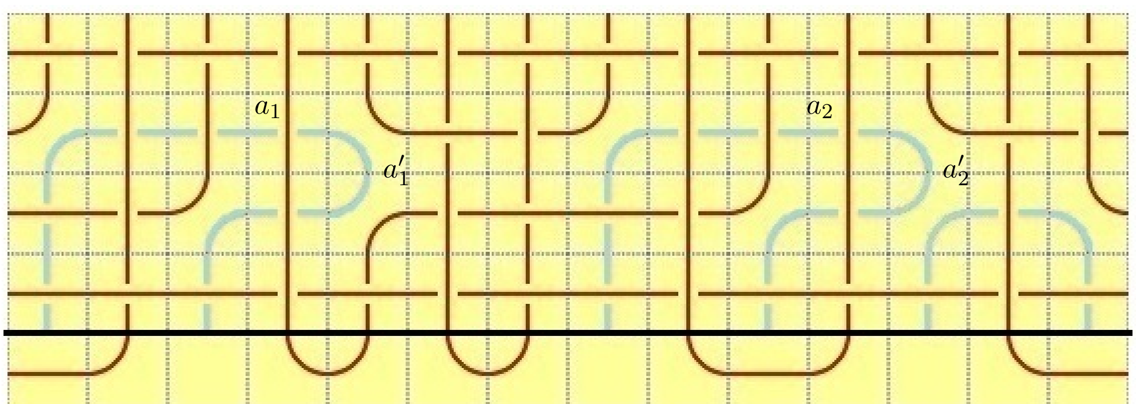

Figure 5.4: The corner-corner move is not defined on the arcs a_1 and a_2 at the top of the edges because of the bottom tile in a_1' and a_2' blocking the move, but this is not a problem because it is defined on the two-tile arcs a_1' and a_2'.

IIIb or IId. As mentioned earlier, the move is not defined if either of the tiles below the corners are type IV; we address this situation later.

The swap for a typical situation is pictured in Figure 5.3. If the tile $T_{i,s}$ or $T_{i,t}$ is type II we replace it with a type 0 tile. If it is type IVa it is replaced with a type IIc tile. A type IVb is replaced with type IId. If the tile $T_{i-1,s}$ is type IIIb it is replaced by a type IIb tile. If it is type IId, it is replaced by IIIa. If $T_{i-1,t}$ is type IIIb then it is replaced by type IIa. If it is IIc, it is replaced by type IIIa.

Lemma 5.2. *A corner-corner move causes a planar isotopy of K and reduces T. Hence there cannot be an arc on which a corner-corner move can be applied in a mosaic that minimizes T.*

Proof. The lemma follows directly from the definition of the move. The tiles between columns s and t swap places, but pairwise remain identical and thus cannot change T. As seen in the figures no matter which configuration appears in column s and t, the ordered triple T decreases in these columns. Specifically, the contributions to $|D|$ remains the same, but the contribution in column s to either $|D'|$ or $|D''|$ is reduced by one and the same is true in column t. $\square$

We now turn our attention to the two cases in which the corner-corner move was not defined to see that neither of these is a problem. The following lemma states that the first one can never occur in a reduced dual.

Lemma 5.3. *A corner-corner arc a with $T_{i-1,t}$ either type IId or type IVb or with $T_{i-1,s}$ type IIc or IVa cannot occur if T is minimal.*

Proof. Given a in the dual running from $T_{i,s}$ to $T_{i,t}$ if $T_{i-1,t}$ in the dual is either type IVb (meaning $R_{i-1,t}$ is type I) or type IId then the corner-corner move is not defined on a. Let the portion of e in tiles $T_{i,t}$ and $T_{i-1,t}$ be called a'. Arcs a'_1 and a'_2 in Figure 5.4 are examples of such arcs. If $t = s+1$ and $T_{i,s}$ is type IV then we are in a situation such as Figure 5.1, but this is impossible since T is minimal and the existence of a bubble release move would contradict minimality. Given the structure of a corner-corner arc a together with adjacent two-tile corner-corner arc a', this is the only case in which the corner-corner move is not defined on a' Therefore we push it to the left so in all other cases a corner-corner move can be applied to a' reducing T just as it can be to a'_1 and a'_2 in Figure 5.4. By Corollary 5.2 we know that the move cannot happen if T is minimal, so $T_{i-1,t}$ cannot be type IId or IVb. The analogous argument holds by reflective symmetry if $T_{i-1,s}$ is type IIc or IVa. $\square$

Thus it is not a problem that the corner-corner move was not defined in this context. We are left only with the following situations in which the corner-corner move was not defined. We could have a corner-corner arc a in the dual running across row i from $T_{i,s}$ to $T_{i,t}$ and if $t > s+1$ we have a dual tile $T_{i-1,w}$ with $s < w < t$ that is type IIa, IIb, or IIIa. If $t = s+1$ then $T_{i,s}$ is type IVb and $T_{i-1,t}$ is type IVa; $T_{i,s}$ may also be type IVb and $T_{i-1,t}$ may be type IVa if $t > s+1$, too, of course.

Lemma 5.4. *If T is minimal, then the only way we can have a corner-corner arc a in row i is if a is part of a nested series of corner-corner arcs $\{a_2, a_3 \dots a_{i-1}\}$ with each a_j contained in row j for $2 \le j \le i-1$.*

Proof. We have seen already that the definition of the dual dramatically limits the choices for tiles beneath a in row $i-1$ so the arc of the dual containing $T_{i-1,w}$ must be a corner-corner arc a_{i-1} from $T_{i-1,s'}$ to $T_{i-1,t'}$ for some s' and t' with $s \le s' < t' \le t$. Iterating the process we either find a corner-corner arc that does not have a corner-corner arc below it in some row j with $2 < j \le i$ contradicting minimality or there are nested corner-corner arcs extending in every row from i down to 2 as in the edges e_1, e_2 and e_3 in Figure 6.3. $\square$

5.3 Corner-edge moves

We again for simplicity choose to describe this move as it moves an arc a of the dual down, but as before, symmetrical moves to the right, left, or up are all valid by rotations or reflections of the mosaic. The move is very similar to the corner-corner move as are the arguments about it.

Our goal in applying the corner-edge move is to reduce T, and we will always do any available corner-corner moves before doing any corner-edge moves, so we need not worry about defining the corner-edge move on an edge for which a corner-corner move is possible.

Corner-edge move: Let e be an edge of the dual that intersects row $i > 2$ in an arc a running directly across i in columns 2 through t and turning down in column $t, t \geq 2$. More precisely e intersects row i in an arc a such that each tile in the dual $T_{i,w}$, $w < t$ is a type IIIa tile and tile $T_{i,t}$ is a type IIa or IVa tile as in Figures 5.5 and 5.6, respectively.

As in the corner-corner move, we pay close attention to any portion of the dual in tiles $T_{i-1,w}$ with $2 \leq w < t$ (row $i - 1$ directly under a). Again $T_{i-1,w}$ cannot ever be type IIc, IId IIIb, IV or V, but can be a type 0 tile without causing any problems.

As before, if there is a tile in the dual $T_{i-1,w}$ with $s < w < t$ that is type IIa, IIb, or IIIa then the corner-edge move is undefined on a. Such examples are seen in the nested edges f_1, f_2, f_3 and f_4 in Figure 6.3.

We may have an obstruction where $t = 2$, $T_{i-1,t}$ is type IId or type IVb and $T_{i,t} \cup T_{i-1,t}$ forms a two tile outermost XX-edge, but we never apply a corner-edge move in this context so we do not mind this obstruction. With this exception we do not encounter an obstruction to the definition where $T_{i-1,t}$ is either type IId or type IVa because it would lead to a reduction via a corner-corner move of tiles $T_{i,t} \cup T_{i-1,t}$ to the left which already contradicts the minimality of T for the dual.

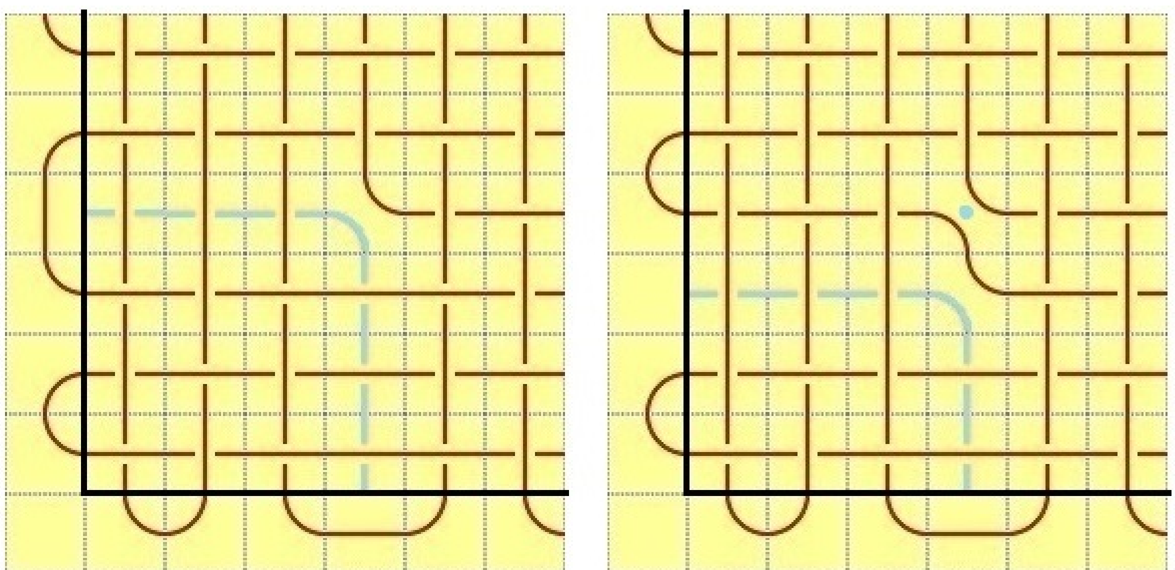

Figure 5.5: The corner-edge move may have a type II tile in its corner.

We again reduce T by moving part of the knot across the dual. Figures 5.5, 5.6, and 5.7 show the basic move.

Because there can be no corner-corner moves in a minimal dual and there are never type V tiles in a dual, the tile $T_{i-1,t}$ must be either a type IIIb, type IIc or type IVa. If $t = 2$ and $T_{i-1,t}$ is type IVa then the corner-edge move is not defined, but again this is an obstruction that we do not mind as we will never need to apply it in this context. Instead examine the case where $t > 2$. If $T_{i-1,t}$ is type IVa we could swap the type IVa dual tile with a type IVb tile, replacing dual D by dual B without affecting the knot. Since the only thing we have changed to go from

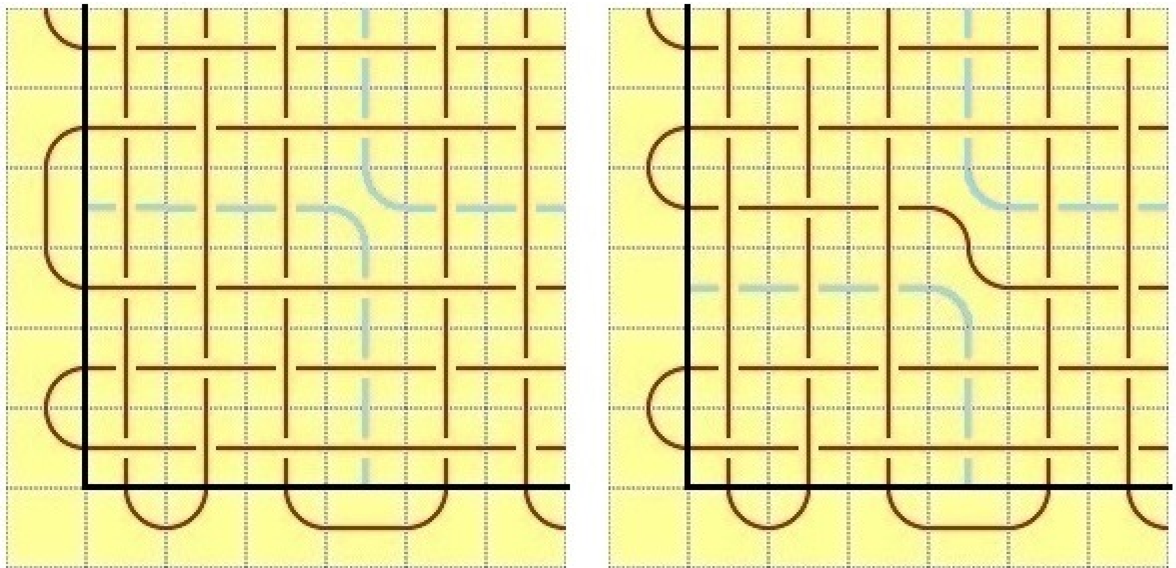

Figure 5.6: Alternatively, the corner-edge move may have a type IV tile in its corner. The resulting move is only slightly different.

D to B is one type IV tile for another, the ordered triple T_D for D is clearly identical to the ordered triple T_B for B. T_B is reduced by a corner-corner move, showing it was not minimal for the knot K whose dual is B (and D) and therefore T_D also was not minimal for K. Since we always choose our embedding of K so that T is minimal we may assume that tile $T_{i-1,t}$ is not type IVa when $t > 2$.

We are now left with the possibilities that $T_{i-1,t}$ must be type IIIb (Figure 5.5) or IIc as depicted in Figure 5.7 and we define the corner-edge move accordingly. The exact prescription for the move is that for each $w < s$ we switch tile $T_{i-1,w}$ with tile $T_{i,w}$. $T_{i-1,t}$ and tile $T_{i,t}$ are treated exactly as they were in the corner-corner move: if $T_{i,t}$ is type II we replace it with a type 0 tile. If it is type IVa it is replaced with a type IIc tile. If $T_{i-1,t}$ is type IIIb then it is replaced by type IIa. If it is IIc, it is replaced by type IIIa. Typical corner-edge moves are depicted in Figures 5.5 through 5.7.

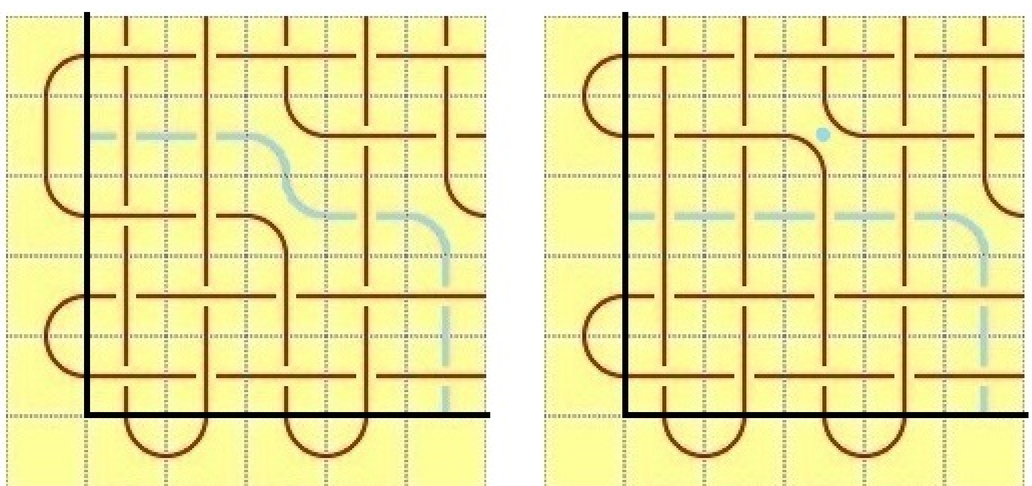

Figure 5.7: $T_{i-1,t}$ may be a type II tile oriented as pictured instead of a type III tile. The resulting move still reduces T.

Again this move was described in terms of a row, but it can be rotated or reflected to move corner to edge row arcs up and down and corner to edge column arcs right and left.

Lemma 5.5. *A corner-edge move causes only a planar isotopy of K and reduces T. Therefore in a minimal mosaic, there cannot be an arc on which a corner-edge move may be applied.*

Proof. The lemma follows directly from the definition of the move. The argument is analogous to in the corner-corner move. □

We now have the moves defined and in the next section will turn our focus to the XX-edges (edges with both endpoints on the same edge), showing that they cannot exist without contradicting minimality. Then once we know there are no edges of this type we will eliminate XY-edges, too.

6 Reduction steps towards the main theorem

Lemma 6.1. *If $E = \{e_1, e_2, \ldots, e_n\}$ is the set of all edges in D and $|D| \leq n - 4$ then there is some e_i containing a type II tile in the dual. If the XX-edges do not share a type IV tile with the XY-edges then at least one of the XX-edges contains a type II tile or the set is empty. The same is true for the XY-edges.*

Proof. Each XX-edge has at least two "outside" corners and each XY-edge must have at least one. These corners must either be type II tiles or Type IV tiles where they meet another edge of the dual. To avoid any type II tiles, the dual would have to stretch from one side of S to the opposite side, but this would mean $|D| \geq n - 2$ violating the hypothesis of the lemma. □

Lemmas 6.2 through 6.6 put together will show that if e is an XX-edge with both endpoints on the bottom in a minimal dual then e contains exactly one corner-corner arc, and that arc can only be concave down. We note that as always, symmetrical arguments can be made by rotation and reflection for edges with endpoints on the other sides of S.

Lemma 6.2. *If e is an XX-edge in D with both endpoints on the bottom side of S or an XY-edge with one end point on the bottom of S and the other on the left side and $|D| \leq n - 4$ and T is minimal then e cannot contain a corner-corner arc a that is concave up. By symmetry this also means the XY-edge cannot have a corner-corner arc that is concave right.*

Proof. By Lemma 5.4, T can be reduced via a corner-corner move applied to a unless there is a nested set of corner-corner arcs inside of a including one in each of the rows of S above a. This, however, cannot happen since it would imply that there are dual tiles in every row of S, contradicting $|D| \leq n - 4$. □

Lemma 6.3. *Suppose $|D| \leq n - 4$ and T minimal. Let e be an XX-edge in D with both endpoints on the bottom edge of S or an XY-edge with one end point on the bottom of S. If a is a corner-corner arc of e in row i that is concave down, then a is the only corner-corner arc on e that is concave down.*

Proof. A second concave down corner-corner arc would require a concave up corner-corner arc between the two: an edge in the plane with endpoints at the same height may not have two local maxima without a local minimum. We know we cannot have a concave up corner-corner arc on e by Lemma 6.2. □

Lemma 6.4. *Let e be an XX-edge or XY-edge in D where $|D| \leq n - 4$ and T is minimal. Then e cannot contain a corner-corner arc concave to the left (representing a maximum in the direction right) and also a corner-corner arc concave right.*

Proof. Each corner-corner arc would need nested corner-corner arcs going all the way to the edge of S. This would, of course, require at least one tile in each column of S, contradicting the fact that $|D| \leq n-4$. □

We have now established several lemmas that work for both XX-edges and XY-edges. The next few lemmas will be just concerned with XX-edges. After establishing further structure on the XX-edges we will be able to return to the XY-edges and deal with them more efficiently.

Lemma 6.5. *Let e be an XX-edge in D with both endpoints on the bottom edge of S with $|D| \leq n-4$ and T minimal. Let the left endpoint of e be in column s and the right endpoint in column t, $s < t$. We cannot encounter a corner-corner arc concave to the left intersecting column w for $w < t$. The same is true for corner-corner arcs concave right in columns w with $s < w$*

Proof. If a is such a corner-corner arc in e, to connect with $T_{2,t}$ in the first case and $T_{2,s}$ in the second, e would need to turn around via a corner-corner arc concave in the opposite direction contradicting Lemma 6.4. □

Lemma 6.6. *Let e be an XX-edge in a minimal dual D with both endpoints on the bottom edge of S with $|D| \leq n-4$ and T minimal. Then e cannot have any corner-corner arcs concave to the left or right.*

Proof. We show the proof for corner-corner arcs concave to the right since the proof to the left is identical up to symmetry. By Lemma 6.5 the concave right corner-corner arc a must start and end in column w, $w \leq s$. However by Lemma 5.4, a must have nested corner-corner arcs inside of it extending all the way to the right side of S. Since a is on the left side of e, this implies that at least one of the nested corner-corner arcs for a is also in e. For e to contain two corner-corner arcs that are concave to the right it must also contain a corner-corner arc concave to the left between them. This contradicts Lemma 6.4 . Thus there were no corner-corner arcs concave to the right. □

These lemmas imply

Corollary 6.7. *Let e be an XX-edge in minimal dual D with both endpoints on the bottom edge of S with $|D| \leq n-4$ and T minimal. Let the left endpoint of e be in column s and the right endpoint in column t, $s < t$. Let row i contain the maximum of e. Then e is strictly contained between columns s and t (inclusive) and below row i (inclusive).*

We now apply this result to outermost XX-edges to build an argument that they must consist of only two tiles exemplified by edge e_1 in Figure 6.3.

Lemma 6.8. *If e is an outermost XX-edge in minimal dual D with $|D| \leq n-4$ then e consists of two adjacent tiles in the second layer. Each of the tiles is either type II or type IV.*

Proof. Without loss of generality let both endpoints of e be on the bottom edge of S. First we argue that if e is an outermost edge then e is totally contained in row 2 and consists of $T_{2,s}$ a type IIb or IVb tile, $T_{2,t}$ a type IIa or IVa tile and type IIIa tiles $T_{2,w}$ for $s < w < t$ (if $t = s+1$ then this last set of tiles is not used). Up to rotation, h_1 and e_1 are examples of such arcs in Figure 6.3. We then strengthen the result to show that e not only is in the second row, but it contains only two tiles.

Let a be the corner-corner arc of e in row i, the highest row that contains a tile of e. If $i = 2$ we are done with the first step of the proof. If $i > 2$ then the row $i+1$ just above $a \subset e$ cannot contain any portion of e since it is above the global maximum of e and this dictates that the tiles in row $i+1$ are inside (not outside) of e. In turn this means that the row $i-1$ just below a must either contain part of e or be outside of e. Since T is minimal, there must be a corner-corner arc $a' \subset D$ below a or we could do a corner-corner move on a pushing it down and reducing T. Then a' cannot be part of an edge other than e since that would imply it was outside of e and e is outermost. Thus $a' \subset e$, but this contradicts Lemma 6.3. This implies that a is in row 2 and this can only happen if $e = a$ and consists of $T_{2,s}$ a type IIb tile, $T_{2,t}$ a type IIa tile and $T_{2,w}$ type IIIa tiles for $s < w < t$.

Now we know that e is entirely contained in row 2. If e contains more than 2 tiles, then the leftmost tile of e can be moved to the right using a corner-edge move to reduce T contradicting minimality. Thus e must be just 2 tiles long and we are only left with the desired type of outermost XX-edges (see arc e_1 in Figure 6.3). □

Lemma 6.9. *If e is an XX-edge in a minimal dual D with both endpoints on the bottom edge of S and $|D| \leq n-4$ and the maximum of e occurs in column i, $i > 2$ then there is a set of nested edges $\{e_2, e_3 \dots e_{i-1}\}$ outside of e with the maximum of each e_j in row j.*

Proof. We must have nested corner-corner arcs outside of e in each row and each edge has only one corner-corner arc. □

Lemma 6.10. *If e is an XX-edge in minimal dual D with both endpoints on the bottom edge of S and $|D| \leq n-4$ and the left endpoint of e occurs in column s at $T_{2,s}$, and e is not outermost, then there is a nested set of edges $\{e_{s+1}, e_{s+2} \dots e_k\}$ outside of e with the left endpoint of e_j in column j for each j, $s+1 \leq j \leq k$ and e_k an outermost edge in D.*

Proof. Because the XX-edges contain no corner-corner arcs that are concave left or right we can always use a corner-edge move to reduce T unless there is an edge with its endpoint in the adjacent column blocking the move. □

Lemma 6.11. *If e is an XX-edge in D and $|D| \leq n-4$ and all the edges outside of e are nested with each other, and e contains a type II tile, then D is not minimal.*

Proof. Examine $K \subset M_n$. If the arcs of K contained in S are connected using the tiles under e in row 1, as it does in the top mosaic in Figure 6.1, then the edge in row 1 also runs under an outermost edge e' outside of e ($e = e'$ if e is outermost). We can add e' to K, connecting it up to the original knot as in the bottom picture in Figure 6.1, and obtain a knot isotopic to K, but we have reduced the ordered triple contradicting minimality.

If K does not use the tiles in row 1 under e, then K is disjoint from all of the tiles in row 1 between the endpoints of e. If necessary change crossings between the knot and dual – but not the knot with itself, thus not changing the knot at all – to make sure e always goes under K, and connect e to itself through row 2 giving a loop. Remove the loop from the dual and add it to the mosaic creating a link mosaic with two components, K and an unknot. This takes us to the middle picture in Figure 6.2. Then use a corner conversion by placing a type V crossing tile where the type II tile of e had been going from the middle to the bottom picture in Figure 6.2. As in the proof of Lemma 3.2, a corner conversion takes the connect sum of K with an unknot giving another version of K on an $n \times n$ mosaic, but with a reduced ordered triple contradicting the minimality of T. □

Lemma 6.12. *If e is an XX-edge in a minimal dual D and $|D| \leq n-4$ then all the edges outside of e are nested with each other.*

Proof. If not, then examine the outermost edge e which has edges outside of it which are not nested with each other. Let e_1 and e_2 be the innermost non-nested edges outside of e (so e_1 is not outside of e_2 and vice-versa and e_1 and e_2 are just outside of e). None of the corners at the top of $e_1 \cup e_2$ can be type II or T is not minimal by the previous lemma, but if they are all type IV tiles then e has a corner-corner arc that is concave up contradicting Lemma 6.2. □

These lemmas imply

Theorem 6.13. *D contains no XX-edges containing a type II tile if D is chosen minimally.*

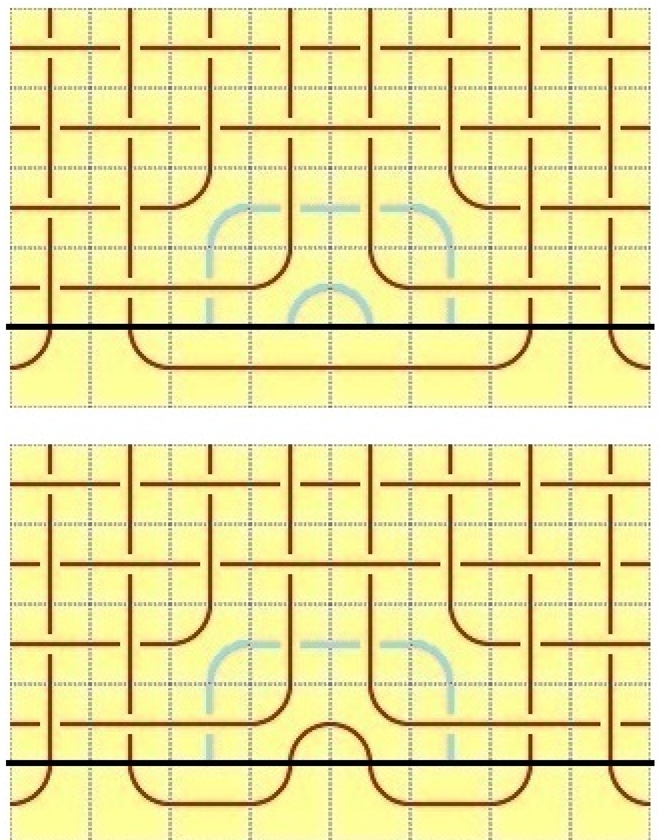

Figure 6.1: If K passes outside of an outermost arc of the dual, the move pictured shows the dual is not reduced.

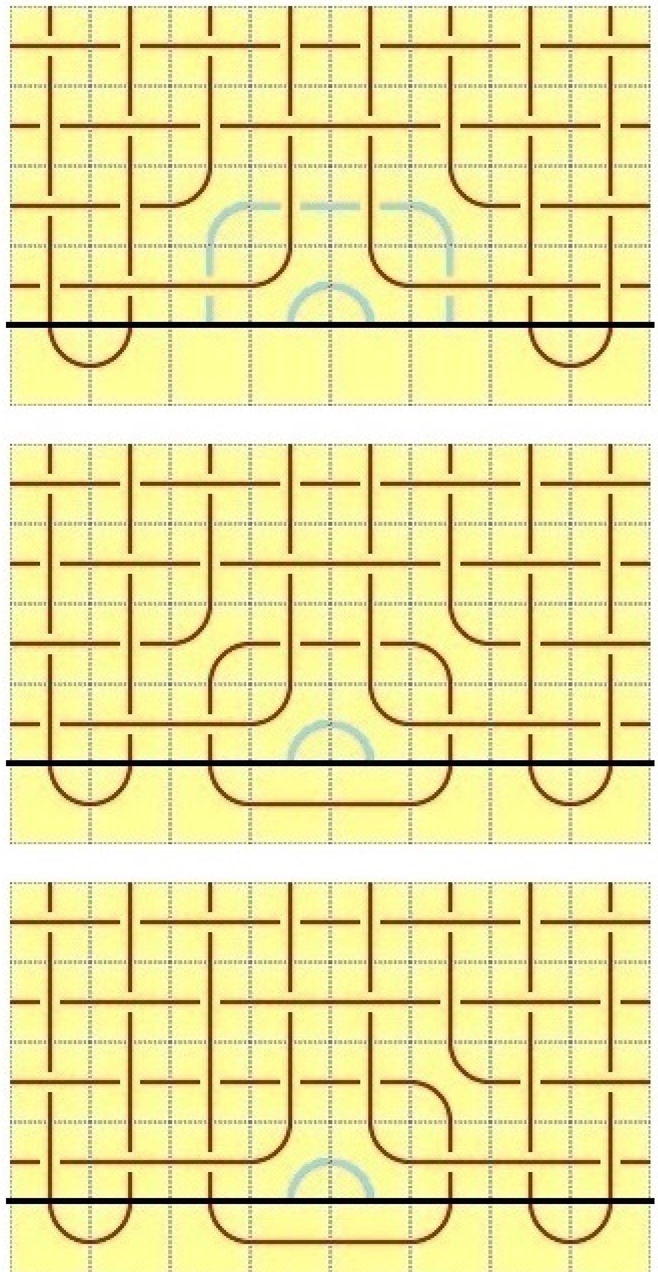

Figure 6.2: If a nested XX-edge in the dual has a type II tile in the dual and K does not pass outside the edge, we can alter the dual reducing T.

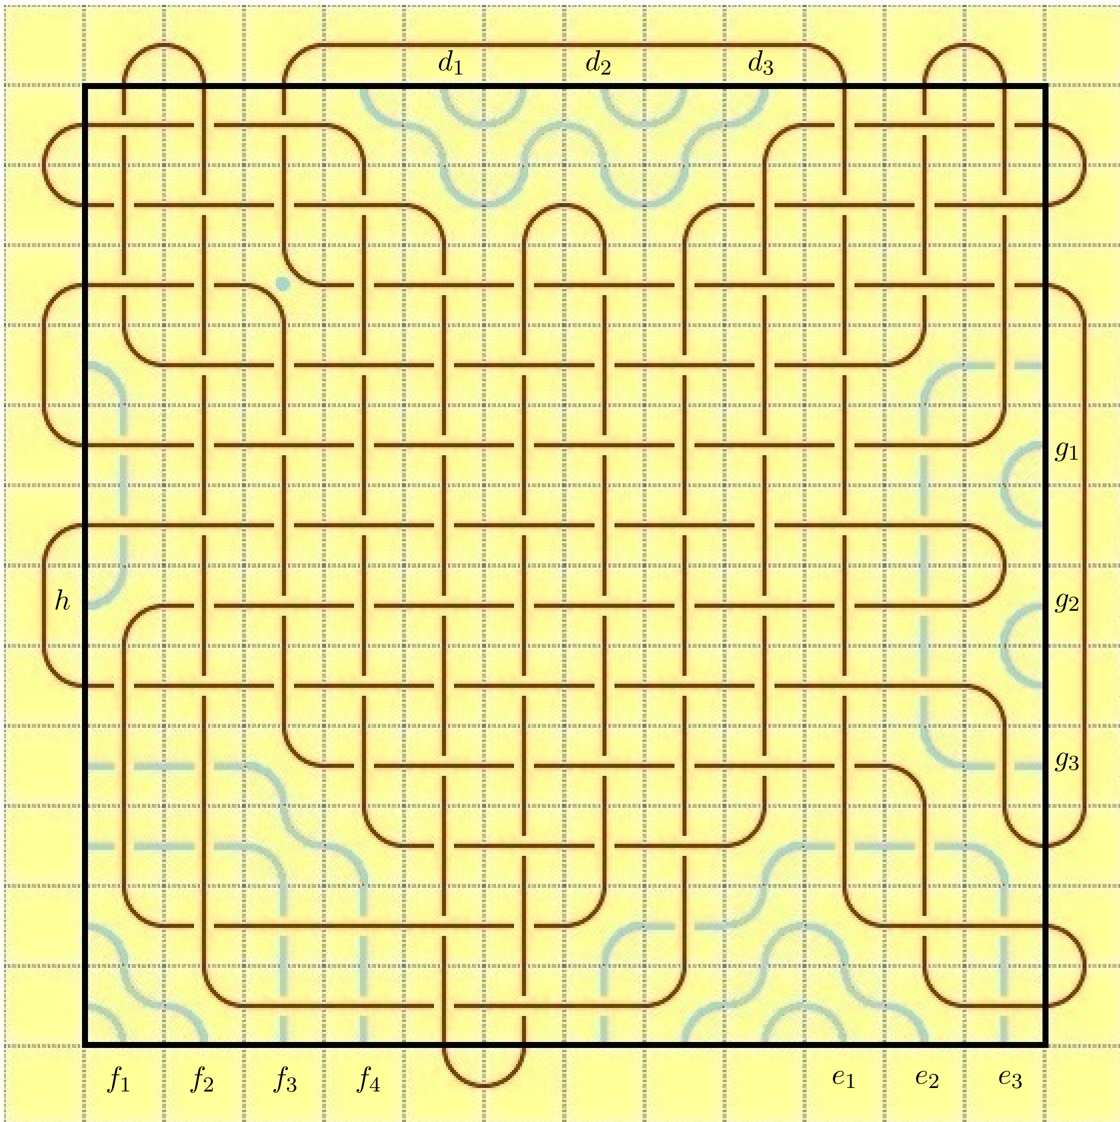

Figure 6.3: An example of a knot mosaic and its dual.

Now that we know that all the corners in an XX-edge are type IV we exploit this fact to get rid of all XX-edges.

Lemma 6.14. *Let e be an XX-edge with both end points on the bottom of S for a minimal dual D with $|D| \leq n-4$. Then e cannot simultaneously share a type IV tile with an XY-edge that has its end point on the right side of S and share a type IV with another XY-edge that has its end point on the left side of S.*

Proof. This is clear because D would stretch all the way across S dictating $|D| \geq n-2$. $\square$

Theorem 6.15. *If $|D| \leq n-4$ and D is chosen minimally, then D contains no XX-edges.*

Proof. If D does contain an XX-edge then there is at least one that is not outside of any of the other XX-edges – the innermost edge from any of the nested sequences would suffice. Call that edge e. Without loss of generality let e have both end points on the bottom of S, specifically in tiles $T_{2,s}$ and $T_{2,t}$ with $s < t$. We know e has no type II tiles by Theorem 6.13, but it must contain two type IV tiles at its maximum. Let f be an edge that meets e in one of these type IV tiles. Note that the tile coming from a maximum for e implies that f is not outside of e. Since e has endpoints on the bottom of S and $|D| \leq n-4$ we know f cannot have an end point on the top edge of S. If there is a second edge g that shares the other type IV tile from the maximum of e, Lemma 6.14 says that it cannot be the case that one of these edges had an end point on the right side of S and the other on the left side so without loss of generality we may assume that the end points of f and g are contained in at most the left and bottom sides of S.

We next argue that both end points of f (or g) cannot just be on the bottom of S. If f has both end points on the bottom it is by definition an XX-edge, but recall that we chose e so it was not outside of any XX-edges so we know that e is not outside of f. As a result both end points of f have to be contained in rows that are on the same side of rows containing the end points of e. In particular f must have end points in tiles $T_{2,u}$ and $T_{2,v}$ with $u < v$ since they are on the bottom of S, but then since the two edges are not nested and do not intersect we must have $u < v < s < t$ or $s < t < u < v$. Either way by Corollary 6.7 f never intersects any of the columns between s and t and e never leaves these columns so they cannot contain a common type IV tile. This implies that f does not have both end points on the bottom of S. Now without loss of generality f (and any edge sharing a type IV tile with e) either has both end points on the left side of S if it is an XX-edge or one on the left side of S and the other on the bottom side if it is an XY-edge. In the latter case, since e is not outside of f we also know that the end point on the bottom of S is in some tile $T_{2,u}$ with $u < s < t$.

Now examine again the top right type IV tile from the maximum of e and the edge f that shares this tile. The portion of f within this tile looks like a IIc tile. Because both of its end points are to the left of e and it does not intersect e we know f must contain a concave up corner-corner arc to contain this tile. This, however, contradicts Lemma 6.2.

Thus e cannot exist so there are no XX-edges in the dual. $\square$

Now we know by Lemma 6.1 that if there are any edges there must be an XY-edge, say e_1, containing a type II tile. From here we proceed by eliminating XY-edges.

Theorem 6.16. *If $|D| \leq n-4$ and D is chosen minimally, then D contains no XY-edges.*

Proof. We know all edges must be XY-edges, and that if the set of edges is nonempty then at least one of the XY-edges contains a type II tile by Lemma 6.1.

Let e be such an XY-edge. We can rotate the entire mosaic if necessary until it encloses the bottom left corner, so let e have one endpoint in tile $T_{s,2}$ on the left edge of S and the other in $T_{2,t}$ on the bottom edge of S enclosing the bottom left corner on its outside. First observe that e cannot contain a corner-corner arc. Any corner-corner arc requires a set of nested corner-corner arcs terminating in an XX-edge on the boundary of S. D, however, contains no XX-edges so this is impossible.

The lack of any corner-corner arcs implies that e is completely contained in the rows below row s (inclusive) and to the left of column t (inclusive). In turn this implies that if $s > 2$ then there must be a set of nested XY-edges $\{e_2, e_3 \dots e_w \dots e_{s-1}\}$ for each w, $2 \leq w < s$. If not we could apply a corner-edge move to reduce T. The analogous result holds for the endpoints on the bottom of S with respect to the columns.

Now we close the argument in the same manner as before. The outermost XY-edge must be a type II tile or one of the arcs from a type IV tile in $T_{2,2}$. The next most outermost has one endpoint in tile $T_{2,3}$ and the other in $T_{3,2}$ etc. Examine $K \subset M_n$. If K runs through row 1 and column 1 past the ends of these arcs, then we add the outermost XY-edge to K and connect it up to give a planar isotopy of K and reducing T, contradicting minimality. If not, recall that e contains a type II tile. Since K does not go through the tiles in row 1 or column 1 outside of e we can turn e into an unknot by hooking the endpoints of e to itself through these tiles. There is no effect on K if we assume that e always passes under K. As in previous theorems we use a corner conversion to replace the hypothesized type II tile from e with a type V tile to connect sum the new unknot with K, yielding another embedding of a knot isotopic to K for which T has decreased contradicting minimality. Thus there can be no XY-edges. $\square$

Since the dual contains no XX-edges, no XY-edges and no loops, we may now conclude the following Corollary.

Corollary 6.17. *If $M_n = M_{2k}$ is an even knot mosaic yielding knot K, with minimal dual D and $|D| \leq n-4$, then $l' = |D'| = 0$. Therefore we may assume D consists exclusively of type 0 tiles.*

7 Proof of the main result for even boards

We use Corollary 6.17 to show our main theorem below.

Theorem 7.1. *If $M_n = M_{2k}$ is an even knot mosaic yielding knot K, then the crossing number of K is less than $(n-2)^2 - (n-4)$.*

Proof. A mosaic with $|D| = l$ and $D' = \emptyset$ has l type 0 tiles in the dual and nothing else. It therefore is obtained from a saturated mosaic by smoothing l crossings. In the language of tiles, we are replacing l type V tiles in the link with l type IV tiles. Each time this is done the number of components in the mosaic changes by at most one.

If $l > n-4$ then K has at most $(n-2)^2 - (n-5)$ crossings failing to exceed our bound. As we saw in the section on saturated mosaics, a saturated even mosaic has $n-2$ or $n-3$ components depending on how the arcs contained in S are connected in the boundary tiles of the mosaic. Since we are only smoothing l crossings, we see that if $l < n-4$ this leaves at least 2 components and we did not really have a knot mosaic. In this context we insist that we are left with a knot, that $l \leq n-4$ and also that $l \geq n-4$ so it must be that $l = n-4$.

If the saturated mosaic starts with $n-2$ components, smoothing $n-4$ crossings still leaves at least 2 components, so to avoid a contradiction the original saturated mosaic must have been connected to yield $n-3$ components. This, however, can only happen when we also have nugatory crossings in each of the four corners of the mosaic. If the $n-4$ type 0 tiles yield a knot then none of them are in a corner of S as smoothing one of these crossings fails to lower the number of components in the link. This means that the knot that results from smoothing $n-4$ crossings will have $(n-2)^2-(n-4)$ crossings, but it also still has the 4 trivial loops in the corners that can be removed with type I Reidemeister moves. Thus K could be embedded with 4 fewer crossings, showing its crossing number is at most $(n-2)^2-(n-4)-4 < (n-2)^2-(n-4)$ so this knot does not exceed our bound on crossing number on even mosaics. Therefore a knot mosaic on an even board cannot have crossing number greater than or equal to $(n-2)^2 - (n-4)$. $\square$

Note that the trefoil establishes that this bound is sharp since it is achievable on M_4 showing a knot of crossing number 3 can be built on a 4×4 board, but our bound says that we cannot have a knot of crossing number 4 on such a board.

The primary goal of this paper is to refine existing upper bounds for crossing number. Theorems 2.4 and 7.1 together establish the following upper bound for crossing number given mosaic number.

Theorem 7.2 (New Upper Bound for Crossing Number)**.** *Given an m-mosaic and any knot K that is projected onto the mosaic, the crossing number c of K is bounded above by the following:*

$$c \leq \begin{cases} (m-2)^2 - 2 & \text{if } m = 2k+1 \\ (m-2)^2 - (m-3) & \text{if } m = 2k. \end{cases}$$

8 Lower bound for mosaic number

At the beginning of this paper we used Theorem 1.1 to relate crossing number and mosaic number. In a similar fashion, Theorem 7.2 may be used to bound a knot's mosaic number from below. First we define

$$B_1 = \sqrt{2+c} + 2$$
$$B_2 = \frac{5 + \sqrt{4c-3}}{2}$$

As a corollary to Theorem 7.2, we have

Corollary 8.1 (New Lower Bound for Mosaic Number)**.** *Let K be a knot with crossing number c and mosaic number m. Then $m \geq \min\{B_1, B_2\}$.*

This will prove useful in future computations of mosaic number. For now, we will briefly explore the behavior of B_1 and B_2. It is easy to see that B_1 and B_2 are asymptotic, as

$$\lim_{c\to\infty} \frac{B_1}{B_2} = 1.$$

Informally, this means that B_1 and B_2, as functions of c, grow at relatively the same rate. A stronger result is that the difference $|B_1 - B_2|$ is bounded by $\frac{1}{2}$, and although this difference is always increasing, it turns out that

$$\lim_{c\to\infty} |B_1 - B_2| = \tfrac{1}{2}.$$

This result may be somewhat surprising: our work has shown that the even and odd cases require different approaches, but in reality the estimates for each are actually quite similar and the predictive power of one never strays too far from the other.

9 Future research directions

Our research has provoked several questions about knot mosaics which are left open to further investigation.

Question. *How does mosaic number behave for the connect sum of knots?*

Question. *Are there any knots whose mosaic number is 2 greater than the number predicted by Theorem 7.2?*

Question. *What are the mosaic numbers for all knots of 10 crossings or fewer?*

It is also natural to look at a more general class of mosaics where instead of insisting the board be $n \times n$ we allow it to be $n \times m$. One might then define the rectangular mosaic number in terms of the number of tiles in the mosaic or perhaps even better the number of tiles on its interior (of course if we wanted to stay more parallel to the current definition we could use the square root of the number of tiles). Mosaics that need not be square should allow for more efficient embeddings especially in the case of knots that are not prime.

Question. *How does crossing number relate to rectangular mosaic number?*

The authors in [6] establish an upper bound on the mosaic number of a knot using arc index, and use this to prove stronger bounds for several classes of knots. Their results, together with Theorem 7.2 and Corollary 8.1 of this paper, provide a clearer picture of the relation between the mosaic number and crossing number of a knot. However, in general a more precise formula that relates the mosaic number and crossing number of a knot would be desirable. The mosaic number of some relatively simple knots remains unknown.

Conjecture. 6_1 *and* 6_3 *have mosaic number 6.*

Note a 5×5 board is simple enough that it can only support 9 crossings. It is likely that one could build these two knots on a 6×6 board and then simply analyze all possible 5×5 boards to establish the answer to the conjecture above, but for knots with higher crossing numbers the complexity of their mosaic representations increases rapidly making a case by case analysis impractical.

Next recall the class of knots first introduced in Section 2, that is, the alternating class of knots that fit on M_{2k+1} with $|D| = 2$. We know several things about these knots already, namely:

- There is an odd knot mosaic (e.g. M_5, M_7, etc.) with a choice of dual such that $|D| = 2$ for each of these knots.
- The crossing number is given by: $c = (2k-1)^2 - 2$, where $k \geq 1$. In particular, this shows that the bound in Theorem 2.4 is sharp.
- 7_4 is the smallest knot of this class; the next would have 23 crossings.

The 7_4 knot shows up in many artifacts in Asian culture where it is called the endless knot, perhaps due to its exquisite symmetry. For this reason, we shall refer to the family of knots that embeds on M_{2k+1} with $|D| = 2$ as E_{2k+1}, with E being an homage to the endless knot.

A knot is *invertible* if it can be deformed to itself, reversing orientation along the way. All knots with 7 crossings or less are known to be invertible [2], so in particular the 7_4 knot is invertible. Given the precise crossing structure of the E_{2k+1} knots and the symmetries of the mosaic board itself, one can easily prove that each E_{2k+1} knot is invertible.

An important invariant in knot theory is chirality. A knot is *amphichiral* if it is ambient isotopic to its mirror image; otherwise, the knot is *chiral* [1]. Knot polynomials are a useful tool in determining chirality: in [4], a knot is shown to be chiral if its Jones polynomial is not palindromic, i.e. if $V(K,t) \neq V(K,t^{-1})$. One may compute the Jones polynomial for the 7_4 knot – the smallest knot in E_{2k+1} – and see that it is not palindromic, indicating that 7_4 is in fact a chiral knot.

Conjecture. E_{2k+1} *knots are chiral.*

This conjecture is supported by the crossing pattern of these knots, but as it would be difficult to calculate knot polynomials even for the 23 crossing endless knot, not to mention larger crossing numbers, the conjecture remains open to other methods. Nevertheless, one potential implication is that for each M_n where n is odd, there would be two endless knots, a left-handed and right-handed version, that are representable on M_n.

Finally, we suggest an extension of knot mosaics to three dimensions. Using cubic blocks rather than square tiles, we can represent a knot in its three dimensional form without imposing crossings onto a 2D representation, while still maintaining a degree of rigidity. Define a *standard cube* as an analog of a mosaic tile: a cube contains 0, 1, 2 or 3 strands, and each face of the cube intersects at most 1 strand in the center of the face. In addition, for each strand within a cube there is at most one critical point in any direction on the interior arc of the strand. Define an *n-cubic knot* as an $n \times n \times n$ array of suitably connected standard cubes. Furthermore, define the *grid number* of a knot (analog of mosaic number) to be the smallest natural number g such that the knot is representable as a g-cubic knot. Note that mosaic number is a (bad) upper bound for grid number. Further research questions on the topic of cubic knots are proposed below.

Question. *For a knot K with mosaic number m and grid number g, $g \leq m$. Find a sharper upper bound for g.*

The authors would like to thank Sam Lomonaco and Lou Kauffman for use of many of the figures in this paper, as well as Joe Paat and Lew Ludwig for inspirational discussions resulting from their paper with Erica Evans [8].

References

[1] Adams, Colin C. *The Knot Book: An Elementary Introduction to the Mathematical Theory of Knots.* Henry Holt and Company, New York (2002).

[2] Burde, G., and Zieschang, H. *Knots,* 2^{nd} *rev. ed.* Berlin: de Gruyter (2002).

[3] Garduño, Irina T. Virtual Mosaic Knots. *Rose Hulman Undergraduate Mathematics Journal.* **10** 2009, 2.

[4] Jones, Vaughan F.R. A polynomial invariant for knots via von Neumann algebra. *Bulletin American Mathematical Society.* vol. 12 (1985). pp. 103-111.

[5] Kuriya, Takahito. On a Lomonaco-Kauffman Conjecture. `arXiv:0811.0710v3 [math.GT]` (2008).

[6] Lee, Hwa Jeong, et al. Upper Bound on Mosaic Number of Knots and Links. `arXiv:1301.6041v1 [math.GT]` (2013).

[7] Lomonaco, Samuel J., and Kauffman, Louis H. Quantum Knots and Mosaics. *Journal of Quantum Information Processing.* vol. 7, nos. 2-3 (2008). pp. 85-115.

[8] Ludwig, Lewis, Evans, Erica, and Paat, Joseph. An infinite family of knots whose mosaic number is realized in non-reduced projections. *Journal of Knot Theory and Its Ramifications* vol. 22, Issue 7 (2013).