

Space-Efficient Prime Knot 7-Mosaics

Aaron Heap  and Natalie LaCourt

SUNY Geneseo, Department of Mathematics, 1 College Circle, Geneseo, NY 14454, USA

* Correspondence: heap@geneseo.edu

Received: 28 February 2020; Accepted: 30 March 2020; Published: 5 April 2020



Abstract: The concepts of tile number and space-efficiency for knot mosaics were first explored by Heap and Knowles in 2018, where they determined the possible tile numbers and space-efficient layouts for every prime knot with mosaic number 6 or less. In this paper, we extend those results to prime knots with mosaic number 7. Specifically, we find the possible values for the number of non-blank tiles used in a space-efficient 7×7 mosaic of a prime knot are 27, 29, 31, 32, 34, 36, 37, 39, and 41. We also provide the possible layouts for the mosaics that lead to these values. Finally, we determine which prime knots can be placed within the first of these layouts, resulting in a list of knots with mosaic number 7 and tile number 27.

Keywords: knot theory; knot; mosaic; tile number; space efficient; mosaic number; quantum knot

1. Introduction

Knot mosaics were first introduced by Lomonaco and Kauffman in [1] as a basic building block of blueprints for constructing an actual physical quantum system, with a mosaic knot representing a quantum knot. The mosaic system they developed consisted of creating a square array of tiles selected from the list of tiles given in Figure 1. These mosaic tiles are identified, respectively, as $T_0, T_1, T_2, \dots, T_{10}$.




Figure 1. Tiles used for constructing mosaic knots.

The first mosaic tile, T_0 , is a blank tile, and the remaining mosaic tiles, referred to as *non-blank tiles*, depict pieces of curves that will be used to construct knots or links when appropriately connected. These non-blank tiles consist of single arcs, horizontal or vertical line segments, double arcs, and over/under knot projection crossings. A *connection point* of a tile is an endpoint of a curve drawn on the tile. A tile is *suitably connected* if each of its connection points touches a connection point of an adjacent tile.

Definition 1. An $n \times n$ array of suitably connected tiles is called an $n \times n$ knot mosaic, or n -mosaic.

Note that an n -mosaic could represent a knot or a link, as illustrated in Figure 2. The first two mosaics depicted are 4-mosaics, and the third one is a 5-mosaic. In this paper, we will be working only with knots, not links.




Figure 2. Examples of knot mosaics.

In addition to the original eleven tiles $T_0 - T_{10}$, we will also make use of *nondeterministic tiles*, such as those in Figure 3, when there are multiple options for a specific tile location. For example, if a tile location must contain a crossing tile T_9 or T_{10} but we have not yet chosen which, we will use the nondeterministic crossing tile, shown as the first tile in Figure 3. Similarly, if we know that a tile location must have four connection points but we do not know if the tile is a double arc tile (T_7 or T_8) or a crossing tile (T_9 or T_{10}), we will indicate this with a tile that has four connection points, as seen in the second tile of Figure 3. If the tile contains dashed lines or arcs, these will indicate the options for that tile. The third tile in Figure 3 could be a horizontal segment T_5 or a single arc T_2 .



Figure 3. Examples of nondeterministic tiles.

There are a few knot invariants of primary importance in this paper. The *crossing number* of a knot is the least number of crossings in any projection of the knot. The remaining invariants are directly related to knot mosaics. The first is the mosaic number of a knot, introduced in [1].

Definition 2. The mosaic number of a knot K is defined to be the smallest integer m for which K can be represented on an m -mosaic.

The next knot invariant is the tile number of a knot, introduced by Lee, Ludwig, Paat, and Peiffer in [2] and first explored by Heap and Knowles in [3].

Definition 3. The tile number of a knot K is the smallest number of non-blank tiles needed to construct K on any size mosaic.

We say that a knot mosaic is *minimal* if it is a realization of the mosaic number of the knot depicted on it. That is, if a knot with mosaic number m is depicted on an m -mosaic, then that mosaic is a minimal knot mosaic. It turns out that the tile number of a knot may not be realizable on a minimal mosaic. This fact was discovered by Heap and Knowles in [4], where it was shown that the knot 9_{10} has mosaic number 6 and tile number 27, but that on a 6-mosaic 32 non-blank tiles were required. The tile number 27 was only achievable on a larger mosaic. Because of this, it is also of some interest to know how many non-blank tiles are necessary to depict a knot on a minimal mosaic, which is known as the minimal mosaic tile number of a knot, first introduced in [3].

Definition 4. Let m be the mosaic number of K . The minimal mosaic tile number of K is the smallest number of non-blank tiles needed to construct K on an m -mosaic.

So the knot 9_{10} has mosaic number 6, tile number 27, and minimal mosaic tile number 32, with the tile number achieved on a 7-mosaic. 9_{10} is the simplest knot for which the tile number and minimal mosaic tile number are not equal. In this paper, we give more examples of this phenomenon.

As we work with knot mosaics, we can move parts of the knot around within the mosaic via *mosaic planar isotopy moves* to obtain a different knot mosaic that depicts the same knot. Two knot mosaic diagrams are of the *same knot type* (or *equivalent*) if we can change one to the other via a sequence of

these mosaic planar isotopy moves. An examples of a mosaic planar isotopy move is given in Figure 4, which is equivalent to a Reidemeister Type I move. If we have a mosaic that has one of these 2×2 submosaics within it, then that submosaic can be replaced by either of the other two without changing the knot type of the depicted knot. While these moves are technically tile replacements within the mosaic, they are analogous to the planar isotopy moves used to deform standard knot diagrams. A more complete list of these moves are given and discussed in [1,5]. We will make significant use of these moves throughout this paper, as we attempt to construct knot mosaics that use the least number of non-blank tiles.



Figure 4. Example of a mosaic planar isotopy move.

A knot mosaic is called *reduced* if there are no *reducible* crossings in the knot mosaic diagram. A crossing in a knot diagram is reducible if there is a circle in the projection plane that meets the diagram transversely at the crossing but does not meet the diagram at any other point. Removable crossings are unnecessary and can be easily removed by twisting. One such reducible crossing is given in the first 2×2 submosaic of Figure 4. Another example is given in Figure 5, together with a transverse circle, where the crossing in the fourth row and third column is reducible. If we want to create knot mosaics efficiently, using the least number of non-blank tiles necessary, we will want to avoid these reducible crossings.




Figure 5. Reducing a reducible knot mosaic.

Definition 5. A knot n -mosaic is *space-efficient* if it is reduced and the number of non-blank tiles is as small as possible on an n -mosaic without changing the knot type of the depicted knot.

The number of non-blank tiles in a knot mosaic that is space-efficient cannot be decreased through a sequence of mosaic planar isotopy moves. In Figure 6, the two knot mosaics depict the same knot (the 5_1 knot). However, the first knot mosaic uses nineteen non-blank tiles and the second knot mosaic uses only seventeen. In fact, seventeen is the minimum number of non-blank tiles possible to create this knot on a 5-mosaic. Therefore, the second mosaic is space-efficient, but the first one is not.




Figure 6. Space-inefficient and space-efficient mosaics of the 5_1 knot.

In [3], the possible layouts for space-efficient n -mosaics, together with the possible values of the minimal mosaic tile numbers and tile numbers, are given for all $n \leq 6$. In the supplement to [4], we are provided with a table of knot mosaics that includes space-efficient mosaics for all prime knots with mosaic number 6 or less. In each of these prime knot mosaics, either the tile number or minimal mosaic tile number is realized. In this paper, we expand upon these ideas to include 7-mosaics.

For a quality introduction to knot mosaics, we refer the reader to [2]. For more details related to traditional knot theory, we refer the reader to [6] by Adams. We also point out that throughout this paper we make use of KnotScape [7], created by Thistlethwaite and Hoste, to verify that a given knot mosaic represents the specific knot we claim it does. Finally, special thanks are due to James Canning, who was kind enough and brilliant enough to create for us a program that automated the process of creating the mosaics in the Table of Mosaics of Section 4.

2. Space-Efficient 7-Mosaics

In [3], the authors determined the bounds for the number of non-blank tiles needed for a space-efficient n -mosaic ($n \geq 4$) depicting a prime knot and in which either every row or every column of the mosaic is occupied. The lower bound is $5n - 8$. If n is even, then the upper bound is $n^2 - 4$. If n is odd, then the upper bound is $n^2 - 8$. Therefore, in the specific case of $n = 7$, if t is the number of non-blank tiles used in the mosaic, then $27 \leq t \leq 41$. In [3], the authors provided a conjecture for the possible values for the number of non-blank tiles used in the mosaic, and we provide a confirmation of that conjecture now (we note that the original conjecture included the extra possibility of 40 non-blank tiles, but the layout that resulted in that option was not space-efficient).

Theorem 1. *If we have a space-efficient 7-mosaic of a prime knot K for which either every column or every row is occupied, then the only possible values for the number of non-blank tiles used in the mosaic are 27, 29, 31, 32, 34, 36, 37, 39, and 41. Furthermore, any such mosaic of K is equivalent (up to symmetry) to one of the mosaics Figure 7.*




Figure 7. Possible layouts for a space-efficient 7-mosaic.

The layouts given in Figure 7 are listed in order of the number of non-blank tiles used, with that number displayed in the upper-right corner. As we can see from the layouts, other than the non-blank tiles that form the outer edges of the knot, the remaining non-blank tiles each have four connection points. As a consequence, the use of the horizontal and vertical segment tiles, T_5 and T_6 , are not necessary for space-efficient mosaics. This does not mean that a mosaic that makes use of these tiles is not space-efficient. However, for any space-efficient mosaic that makes use of them, there is a planar isotopy move that can be applied that removes the line segment tiles without changing the number of non-blank tiles. For the sake of clarity, we postpone the proof of this theorem until Section 3.

Now that we know the space-efficient layouts of prime knot 7-mosaics, we seek to determine what prime knots actually fit within these. Naively, there are four tile options for each remaining nondeterministic tile, which means that there are $4^{13} = 67,108,864$ options for the first layout and 4^{23}

(over 70.3 trillion) options for the last layout. Of course these options can be greatly reduced using symmetry and space-efficiency. Also, we know from [4] that every prime knot with crossing number 8 or less can fit on a 6-mosaic and has tile number 27 or less. Therefore, we only seek to find prime knots with crossing number 9 or more. For the first layout, symmetry, space-efficiency, and restricting to nine or more crossing tiles, we can reduce the number of options to the low thousands, a huge improvement over 67 million. Because of the vast number of options, especially for the larger layouts, we limit ourselves in this paper to just the three simplest layouts, the ones that use 27 non-blank tiles.

Using the techniques of [4], we fill portions of these layouts with the 3×3 building blocks shown in Figure 8, which have either two, three, or four crossing tiles (a similarly filled block with only one crossing tile is not space-efficient). For example, in the first layout with 27 non-blank tiles, the upper-left 3×3 submosaic can be filled with any of the options given in Figure 8. The lower-right 3×3 submosaic can be filled with a rotation of these. Then we populate the rest of the layout with either double-arc tiles or crossing tiles, making sure that we use at least nine crossing tiles.




Figure 8. 3×3 building blocks.

Once the mosaics are completely filled within the specified layout, we eliminate any links and composite knots, any duplicate layouts that are equivalent to others via obvious mosaic planar isotopy moves, and any mosaics for which the tile number can easily be reduced by a simple mosaic planar isotopy. Finally, we use KnotScape to determine what knots are depicted in the mosaic by choosing the crossings so that they are alternating, as well as all possible non-alternating combinations. Doing this for all three layouts with 27 non-blank tiles, we find all knots with mosaic number 7 and tile number 27. We also find prime knots with mosaic number 6 and minimal mosaic tile number 32 but whose tile number 27 is only realized on mosaics of size 7 or larger. The following theorems summarize the results. When listing prime knots with crossing number 10 or less, we will use the Alexander-Briggs notation, matching Rolfsen's table of knots in *Knots and Links* [8]. For knots with crossing number 11 or higher, we use the Dowker–Thistlethwaite name of the knot. See information about this naming at *KnotInfo* [9].

Theorem 2. *The following prime knots have mosaic number 6, minimal mosaic tile number 32, and tile number 27 realized on a space-efficient 7-mosaic: 9_{10} , 10_{11} , 10_{20} , 10_{21} , and $11a_{341}$, and their tile numbers are realized in the mosaics of Figure 9.*




Figure 9. Knots with mosaic number 6 whose minimal mosaic tile number is greater than their tile number.

Proof. Each of the knots listed were found to have mosaic number 6 and minimal mosaic tile number 32 in [4]. In Figure 9, we give space-efficient 7-mosaics with tile number 27 for each of these knots. So, in order to create these knots on a 6-mosaic, the fewest non-blank tiles possible is 32, while it is possible to depict these knots using only 27 non-blank tiles on a 7-mosaic. \square

Theorem 3. *The following prime knots have mosaic number 7 and tile number 27:*

- (a) 9_6 , 9_{15} , 9_{18} ,

- (b) $10_5, 10_6, 10_7, 10_8, 10_9, 10_{10}, 10_{13}, 10_{14}, 10_{15}, 10_{16}, 10_{17}, 10_{18}, 10_{19}, 10_{24}, 10_{25}, 10_{26}, 10_{29}, 10_{30}, 10_{31}, 10_{32}, 10_{33}, 10_{35}, 10_{36}, 10_{38}, 10_{39},$
- (c) $11a_{90}, 11a_{93}, 11a_{119}, 11a_{145}, 11a_{180}, 11a_{184}, 11a_{185}, 11a_{192}, 11a_{203}, 11a_{205}, 11a_{210}, 11a_{226}, 11a_{306}, 11a_{307}, 11a_{308}, 11a_{309}, 11a_{311}, 11a_{333}, 11a_{336}, 11a_{337}, 11a_{363},$
- (d) $12a_{541}, 12a_{601}, 12a_{1024}, 12a_{1034}, 12a_{1126},$ and
- (e) $13a_{4304}.$

Proof. We have given space-efficient mosaics with tile number 27 for each of these knots in the table of knots in Section 4. Since the mosaic number of each of these knots is 7, we know that they cannot have a tile number smaller than 27. \square

Finally, we point out that all of the knot mosaics for the knots listed in Theorem 3 come from the first layout given in Theorem 1 (Figure 7). Neither of the next two layouts, also with 27 non-blank tiles, resulted in any knots that the first layout did not. This is analogous to what we see with 6-mosaics, where there are two space-efficient layouts with 22 non-blank tiles, but the second layout led to the same results as the first. See [4] for more information. Based on these early results and preliminary results using the remaining layouts, we propose the following general conjecture.

Conjecture. Any two space-efficient layouts for an n -mosaic with the same number of non-blank tiles used lead to the same collection of prime knots.

3. Useful Observations and the Proof of Theorem 1

In this section, we will prove Theorem 1. As we progress toward this goal, we first provide some useful terms and observations that we will make use of as we attempt to create space-efficient mosaics and count the number of non-blank tiles needed to create them.

Suppose there are two adjacent single arc tiles, $T_1, T_2, T_3,$ or T_4 , that share a connection point, and the other connection points enter the same adjacent row or column. The four options are shown in Figure 10, and we will refer to these collectively as *caps* and individually as *top caps*, *right caps*, *bottom caps*, and *left caps*, respectively. We will encounter mosaics that have pieces that are similar to caps but with line segment tiles, T_5 or T_6 , between the arc tiles or pieces that can be easily changed to these via planar isotopy moves. We will call these *reducible caps* and define them to be a collection of suitably connected tiles with no crossing tiles that are, or can be changed via planar isotopy moves to, two single arc tiles that are connected by line segment tiles and whose other connection points enter the same adjacent row or column. Essentially, a reducible cap can meander around but can easily be reduced to a reducible cap that does not meander, such as the last example in Figure 10, which could be simplified to the fifth example.



Figure 10. Caps and reducible caps.

Equipped with this terminology, we consider the following observations that will assist us in counting the minimum number of non-blank tiles necessary to create knot mosaics. More general versions of the first three observations were proven in [3], and we include them here, without proof, specifically applied to prime knots. The fourth observation is a mosaic planar isotopy that we make use of several times that can reduce the number of non-blank tiles used in a mosaic. The fifth observation is simply a reminder of how to recognize when a mosaic does not represent a prime knot.

This first observation tells us that we can create all of our space-efficient knot mosaics without using the corner tile locations, the first and last tiles in the first row and last row of the mosaic. Because the outer rows and columns need not be occupied, we may assume that the first tile and the

last tile in the first occupied row and column is a blank tile, and similarly for the last occupied row and column.

Observation 1 (Lemma 3.1, [3]). *We can assume that the corner tiles of any space-efficient knot mosaic are blank T_0 tiles. Furthermore, for a space-efficient knot mosaic, the first and last tile location of the first and last occupied row and column are blank (or can be made blank via a planar isotopy move that does not change the tile number).*

Observation 2 (Lemma 3.3, [3]). *Suppose we have a space-efficient knot mosaic. If there is a cap or reducible cap in any row (or column), then the two tiles that share connection points with the cap must have four connection points. Examples are shown in Figure 11.*



Figure 11. Tiles that share connection points with caps must have four connection points.

Observation 3 (Lemma 3.6, [3]). *Suppose we have a space-efficient 7-mosaic. Then the first occupied row of the mosaic can be simplified so that the only non-blank tiles form either one or two top caps. Similarly, the last occupied row is made up of bottom caps, and the first and last occupied columns are made up of left caps and right caps, respectively. Although it is not explicitly stated in [3], the proof there shows that every reducible cap can be simplified to a cap without increasing the number of non-blank tiles.*

Observation 4. *The mosaic planar isotopy move given in Figure 12 does not change the knot type of the knot mosaic. The two tiles with four connection points rotate into the tile positions to the right. The number of non-blank tiles in the resulting mosaic, after applying this isotopy, is always less than or equal to the number of non-blank tiles in the original mosaic.*



Figure 12. This planar isotopy move does not increase the number of non-blank tiles in a mosaic.

Observation 5. *If a knot mosaic can be separated into two nontrivial, space-efficient pieces connected by exactly two connection points, then the knot mosaic represents a composite knot. An example of this can be seen in Figure 5.*

As we have previously noted, the use of the horizontal and vertical segment tiles, T_5 and T_6 , are not necessary for create space-efficient mosaics. Once we show this fact, proving Theorem 1 is simple. Observation 3 tells us that T_5 and T_6 tiles are not needed in the first and last occupied rows and columns. Next, we show that they are not needed in the second and penultimate occupied rows and columns.

Lemma 1. *Suppose we have a space-efficient 7-mosaic of a prime knot. Then the mosaic can be simplified so that there are no horizontal or vertical line segment tiles in the second occupied row/column and the next-to-last occupied row/column.*

The proof of this lemma is quite simple but long because of the accompanying figures that depict the various cases. However, knowing that we do not need to use horizontal and vertical segment tiles in the first two and last two rows and columns greatly reduces the possibilities for the layouts of a space-efficient 7-mosaic.

Proof. We prove the lemma for the second occupied row, and the other options follow by rotational symmetry. By Observation 3, the first occupied row only contains one or two top caps, with the rest of the positions filled by blank tiles. By Observation 2, the tile locations directly below these caps

must have four connection points, and this also prevents any vertical segment tiles from occurring in the second row. The only options (up to symmetry) for the first two occupied rows are as depicted in Figure 13. We assume that the first occupied row of the mosaic is actually the top row since, if it is not, we can simply shift all of the tiles of the mosaic upward. In each case below, we will reach a contradiction when we assume that there is a horizontal segment tile in the second row. In most cases, we construct partial mosaics for every possible tile choice until it is obvious that the mosaic is not space-efficient or does not represent a prime knot. We always base our choices on the observations stated above, knowing that the tiles sharing connection points with any cap must have four connection points and avoiding the corner tiles, reducible caps, links, composite knots, and reducible knots.



Figure 13. The first occupied row can only have one or two caps.

Case 1: First, we consider the case where there are two consecutive top caps. The first five tile positions in the second row are determined by Observation 2, which also prevents the sixth tile position from being a horizontal segment tile. Otherwise, there would be a cap in the seventh row, which would necessitate a tile with four connection points in the sixth position. Thus the sixth tile must be a single arc tile T_1 , and the final tile position must be blank.

Case 2: The next case has two top caps with a blank tile in between them. In this case, the second row is completely determined and must have a horizontal segment tile as seen in Figure 13. Our claim is that any completion of this mosaic will not be space-efficient or will not represent a prime knot. To see this, we will examine the remaining rows. There is actually only one possibility for the third row as well, which can be seen in the first mosaic of Figure 14. The first and last tiles in the third row must complete a left and right cap, respectively, and the second and sixth tiles must have four connection points. The third and fifth tile positions in the third row must also have four connection points. Otherwise, they would be single arc tiles, and according to Observation 5, any resulting space-efficient mosaic would not represent a prime knot.

We now consider the remaining positions for the middle column. The tile in the fourth position can either be blank, another horizontal tile, or a single arc tile T_1 or T_2 (because of symmetry, the T_1 and T_2 cases are equivalent). If the entire middle column has only blank or horizontal segment tiles, then the mosaic is not space-efficient, as we can collapse this column. That is, there must eventually be a T_1 (or T_2) single arc tile. This tile clearly cannot be in the sixth row (by Observation 2) or the seventh row. Thus, the only options are as in the second and third mosaic of Figure 14.




Figure 14. The first three rows are determined and there must be a single arc tile in the fourth or fifth position of the middle column.

Suppose the tile in the middle position of the fourth row is the single arc tile T_1 , as in the second mosaic of Figure 14. Because of the locations of connection points and sides of tiles with no connection points, the options are limited. For example, the fifth position in the fourth row can only be a vertical segment tile T_6 or single arc tile T_3 . The vertical segment leads to the first partial mosaic in Figure 15. If the T_3 tile is used, the tile in the sixth position could be a single arc tile T_4 or a tile with four connection points. The first option is the second partial mosaic in Figure 15, and the second option leads to the remaining partial mosaics of Figure 15. All of these depend on the tile in the fifth position of the fifth row, which can be a horizontal tile or a single arc tile T_2 . The last three of these options depend on the third tile position in the fourth row, which can only be a single arc tile T_3 or a tile with four connection points. Filling in a few more tiles around these make it is easy to see that none of

these are space-efficient. For the first three, simply slide the right-most tiles to the left by collapsing the horizontal tiles. For the fourth, fifth, and sixth mosaics, slide the upper, left quadrant to the right by collapsing the horizontal segments.




Figure 15. If the tile in the fourth row, fourth column is a T_1 tile, then the mosaic is not space-efficient.

Finally, suppose the tile in the middle position of the fourth row is either blank or a horizontal tile, and the tile in the fifth row, fourth column is the single arc tile T_1 , as in the third mosaic of Figure 14. There are only seven ways to complete the tile positions to the right of these, and we provide them in Figure 16. In each case, it is easy to see that the resulting mosaics are not space-efficient. In the first six cases, we simply slide the upper, right quadrant of the mosaics to the left by collapsing all of the horizontal segment tiles. In the case of the seventh mosaic, we can lower the tile number of the mosaic using the planar isotopy given in Observation 4.




Figure 16. If the tile in the fourth row, fourth column is blank or a horizontal segment, then the mosaic is not space-efficient.

Case 3: Now let us consider the case where there is only one top cap in the first occupied row, and it is located in the first two tile positions after the corner tile, as in Figure 17. Then the first tile in the second occupied row must be a single arc tile T_2 , followed by two tiles with four connection points. There must also be a single arc tile T_1 in this row, but this T_1 tile cannot be part of a right cap (that is, the tile below it is not a T_4 tile). To see this, assume the T_1 tile is part of a right cap. If the T_1 tile is in the fourth tile position of the second row, then, using Observations 2 and 5, it is easy to see that the knot mosaic is either not prime or not space-efficient. If the T_1 tile is in the fifth, sixth, or seventh position, then Observation 2 says the preceding tile position must have four connection points, which contradicts the fact that the first row only has a single top cap.




Figure 17. The first case with only one top cap.

We now examine the third row. The first tile must complete the left cap with a single arc tile T_3 , and the second tile position must have four connection points. The third position must also have four connection points. Otherwise, this tile position would be a single arc tile T_4 , and the mosaic would either not be space-efficient or would not represent a prime knot (Observation 5).

For the sake of contradiction, suppose there are horizontal segment tiles in the second row. There clearly cannot be three horizontal segment tiles because this forces the arc tile T_1 into the seventh position in this row, which is necessarily part of a right cap, and we have already ruled this out.

If there is only one horizontal segment in the second row, then it must be in the fourth tile position, and the fifth tile position is the arc tile T_1 . We know this is not part of a right cap, and we look at the tiles below the horizontal segment. Directly below the horizontal segment must be a single arc tile T_1 or another horizontal segment tile. The T_1 option, shown in the first mosaic of Figure 18, is easily simplified to reduce the number of non-blank tiles. For the horizontal segment option, the options are similar to those in Case 2. The tile below it can only be a blank, horizontal segment, or single arc (T_1 or T_2) tile. Eventually there must be a T_1 or T_2 tile to avoid the possibility of just collapsing the entire column, and this must occur in either the fourth row or the fifth row. With each of these options in mind, the remaining fourteen partially completed mosaics shown in Figure 18 cover all possibilities in this scenario, and it is again easy to see that none of them are space-efficient.




Figure 18. Possible configurations with a horizontal segment in the fourth position of the second row.

If there are two horizontal line segments in the second row, they must be in the fourth and fifth tile positions, and the sixth tile position is the single arc tile T_1 . In the third row, the tile in the fourth position must be a single arc tile T_1 or a horizontal segment tile. If it is the T_1 tile, then the fifth position must be blank or the single arc tile T_2 . In either case, shown in the first two mosaics of Figure 19, it is easy to see that we can eliminate the horizontal segments in the second row without increasing the number of non-blank tiles.




Figure 19. Possible configurations of the first three rows when the second row has two horizontal line segments.

If the tile in the fourth position of the third row is a horizontal segment tile, there are only two ways to complete the third row, and they are depicted in the last two mosaics of Figure 19. The first one is easily seen to reduce to only one horizontal segment in the second row, which was covered above. The second possibility is also simple to eliminate after we examine a couple of tiles in the fourth row of the mosaic and perform a simple planar isotopy move. The sixth and seventh positions of the fourth row must complete the right cap, and there is an equivalent mosaic, with fewer non-blank tiles, in which the tiles of this right cap are moved into the position of a top cap and the horizontal segments are removed, as shown in Figure 20.



Figure 20. Simplifying the mosaic when the second and third row have two horizontal line segments.

Case 4: Now we consider the final case, shown in Figure 21, where there is a single top cap, and it is located in the third and fourth tile positions of the first occupied row. We will assume that the first column is occupied, otherwise a shift of the mosaic to the left would reduce this to Case 3.



Figure 21. The second case with only one top cap.

It is easy to see that a horizontal segment tile is not allowed in the second or sixth position of the second row, as this would necessarily violate Observation 2. Thus, if there is a horizontal segment tile, the only possibility is for it to be in the fifth position. Just as in previous cases, there must eventually be a single arc tile, T_1 or T_2 , in the fifth column below the horizontal segment. There are twenty possibilities, and they are given in Figure 22. Each mosaic is either not space-efficient or does not represent a prime knot.




Figure 22. Possible configurations when there is a single top cap occupying the third and fourth position of the first row and a horizontal segment tile in the fifth position of the second row.

Having completed all four cases, we have shown that any space-efficient 7-mosaic of a prime knot can be simplified so that the second occupied row does not have a horizontal segment tile. \square

Now that we have shown that all space-efficient 7-mosaics of prime knots can be created without horizontal or vertical segment tiles in the first two and last two rows and columns, we are ready to turn our focus to proving Theorem 1.

Proof of Theorem 1. We assume we have a space-efficient 7-mosaic of a prime knot for which every row is occupied. The case where every column is occupied is equivalent. Our goal is to show that the only possible layouts for the mosaic are those given in the statement of the theorem or that it is equivalent to one of them via a planar isotopy that does not increase the number of non-blank tiles. The resulting number of non-blank tiles follows immediately.

The first and seventh rows of the mosaic have either 2 or 4 non-blank tiles (one or two caps by Observation 3). In either case, we have either 4 or 6 non-blank tiles in the second and sixth rows. In the proof of Lemma 1, we show that there are only three options for the first row. These lead to only four possibilities for the first two rows (up to symmetry), and these are given in Figure 23.



Figure 23. The first two rows of a space-efficient 7-mosaic.

Assuming the first column is occupied, the options for the first two rows extend to the first two columns. After removing any duplicates that are equivalent up to symmetry, there are nine options for the first two rows and columns of the mosaic. See Figure 24.




Figure 24. The first two rows and columns of a space-efficient 7-mosaic.

Next, we consider the bottom two rows and the two rightmost columns. We have assumed that all rows are occupied but not necessarily every column. However, in order to avoid composite knots or space-inefficiency, there must be at least four connection points between any two rows (Observation 5), except between the first two and last two rows. Removing options that are equivalent to others via symmetry, we find 20 possibilities for the two outermost rows and columns. See Figure 25.

Now that we have a complete set of possibilities for the outer shell of a space-efficient 7-mosaic, we want to fill in the inner 3×3 block. Some of that can be accomplished using Observation 2. Using Observation 5, we need at least four connection points between each of the middle three rows and columns, and, for example, in the upper, left 3×3 block of the first mosaic in Figure 25, the tile in the third row and third column cannot be a single arc tile T_4 . Therefore, the tile in that location must have four connection points.

These simple observations are applied as we determine the options for filling in the inner 3×3 blocks of the mosaics in Figure 25. For some, there is only one way to complete the mosaic, such as the first, second, third, fourth, eighth, ninth, tenth, and thirteenth mosaic. All of these, except the second one, lead to one of the desired layouts given in the statement of the theorem. The others have more than one possible way to complete them and require further consideration.




Figure 25. The first and last two rows and columns of a space-efficient 7-mosaic.

The second outer shell in Figure 25 must be completed with horizontal or vertical segment tiles, as seen in the first mosaic of Figure 26. The resulting mosaic is clearly not space-efficient, as we can reduce the number of non-blank tiles with a planar isotopy that shifts the upper, left 4×4 block to the right (or down).




Figure 26. The second outer shell completed is not space-efficient, and the fifth outer shell can be completed in three ways.

For each of the remaining outer shells, if we fill the remaining tile positions of the inner 3×3 block, when possible, with tiles that have four connection points, we end up with the remaining layouts given in the statement of the theorem. If we assume that at least one of those tile positions does not have four connection points, there are several possibilities to consider, which we look at next, but each one is either not space-efficient or can be changed to one of the desired layouts without changing the number of non-blank tiles. For example, if we complete the fifth outer shell with tiles that have four connection points when possible, the result is the second mosaic in Figure 26, which is the third layout given in the statement of the theorem. However, there are two alternative completions, given in the third and fourth mosaics of Figure 26. In both cases, the vertical segment tiles can be altered by a planar isotopy that changes it to the second mosaic without changing the number of non-blank tiles.

The sixth, seventh, eleventh, twelfth, fifteenth, eighteenth, and twentieth outer shells each have one alternative completion, which are shown in Figure 27. None of these are space-efficient as each can be simplified to one of the layouts given in the statement of the theorem using the planar isotopy shown in Observation 4.




Figure 27. The sixth, seventh, eleventh, twelfth, fifteenth, eighteenth, and twentieth outer shells completed in ways other than those given in Theorem 1.

For the fourteenth and seventeenth outer shells, there are two alternative completions, given in Figure 28. In each case, the vertical segment tiles can be removed by a planar isotopy without changing the number of non-blank tiles.




Figure 28. The alternative completions of the fourteenth and seventeenth outer shells can be simplified to not use the vertical segment tiles.

The sixteenth outer shell has six alternative completions, and they are given in Figure 29. The first two are easily simplified, reducing the number of non-blank tiles, using the isotopy of Observation 4. The next three are easily seen to be equivalent to the first two. For the fifth one, just slide the center tile to the right, and the resulting mosaic is identical to the fourth. The sixth alternative can be simplified, as shown, by rotating the tile in the fourth row and second column, which has four connection points, into the fourth row, third column.




Figure 29. The sixteenth outer shell completed in ways other than those given in Theorem 1.

Finally, we encounter six alternative completions of the nineteenth outer shell, see in Figure 30. Each of these are handled in ways similar to the sixteenth outer shell.




Figure 30. The nineteenth outer shell completed in ways other than those given in Theorem 1.

This completes our discussion of filling in the outer shells given in Figure 25. We have exhausted all possibilities of completing a space-efficient 7-mosaic in which all rows are occupied, arriving at those listed in the statement of the theorem. \square

4. Mosaics for Theorem 3

In this section we include the knot mosaics for each of the prime knots listed in Theorem 3. The mosaics given below in Figures 31 and 32 constitute the proof for that theorem. For each knot mosaic, the mosaic number, minimal mosaic tile number, and tile number of the knot are realized, but the crossing number may not be realized. If the knot mosaic is marked with an asterisk (*) then the given mosaic has more crossing tiles than the crossing number for the represented knot, but it is the minimum number of crossing tiles needed in order for the tile number to be realized.




Figure 31. Mosaics of knots listed in Theorem 3.

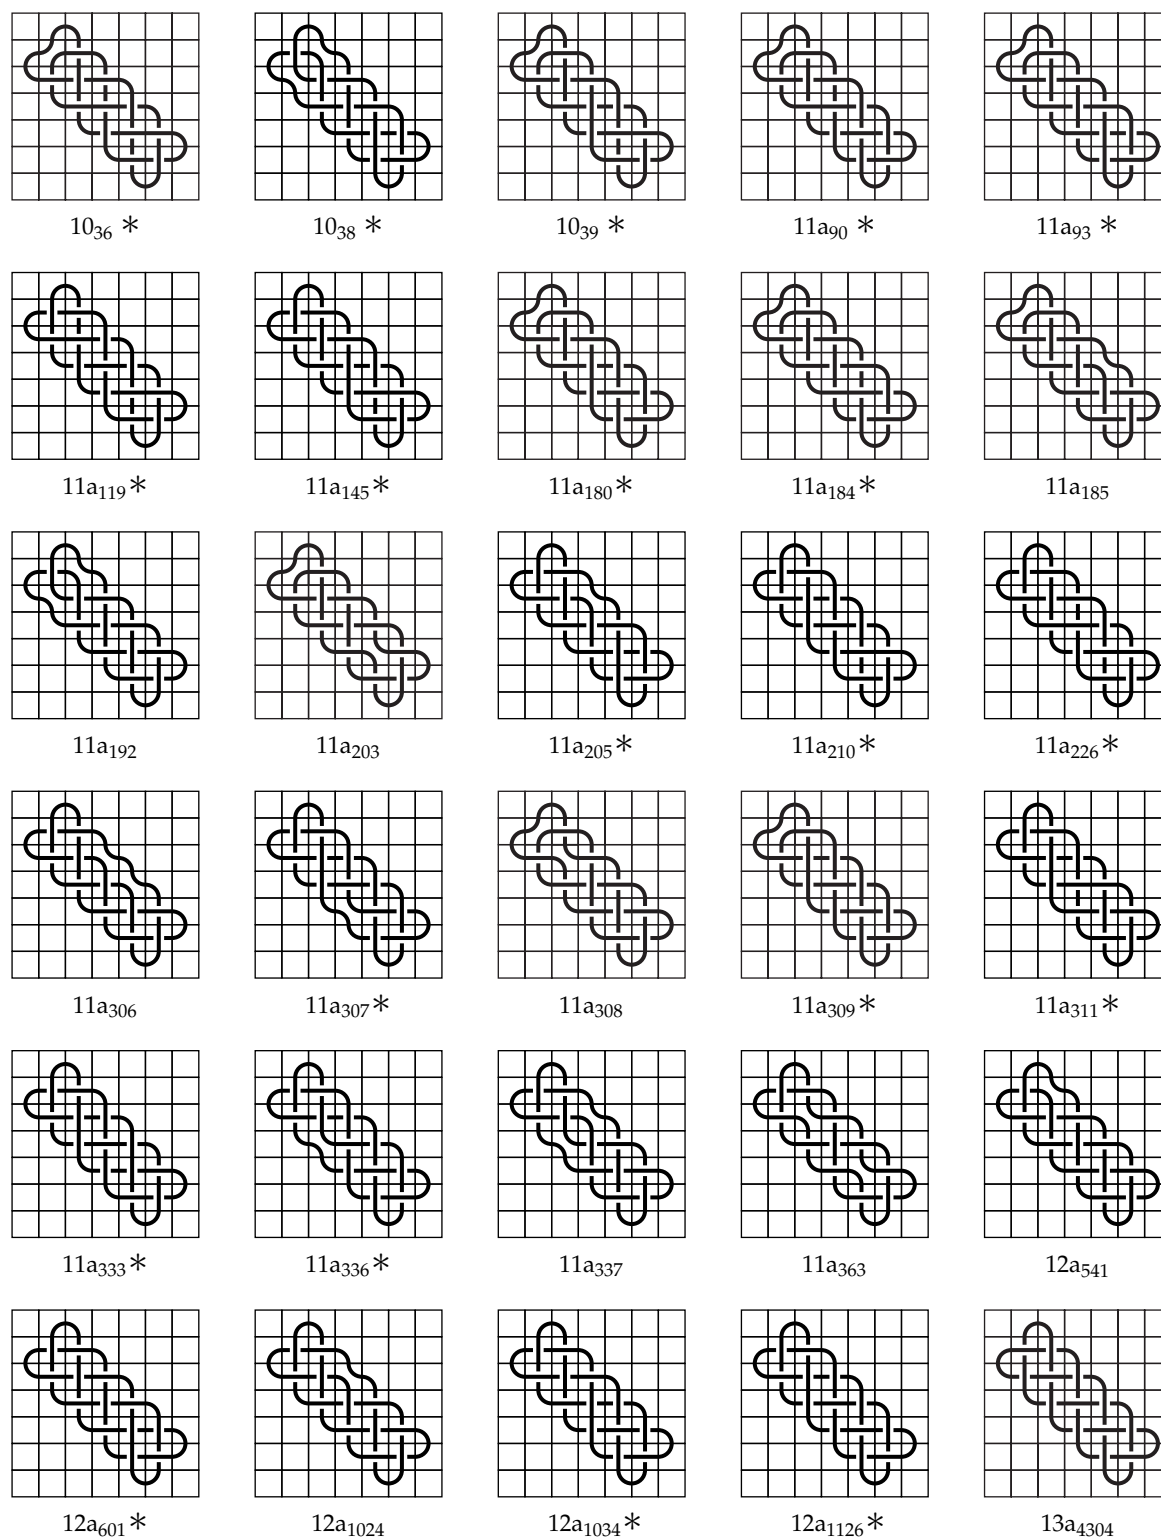


Figure 32. Continuation of mosaics of knots listed in Theorem 3.

Author Contributions: Conceptualization, A.H.; Formal analysis, N.L.; Investigation, A.H. and N.L.; Methodology, A.H.; Project administration, A.H.; Supervision, A.H.; Validation, A.H. and N.L.; Visualization, N.L.; Writing original draft, A.H. and N.L.; Writing review and editing, A.H. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Lomonaco, S.J.; Kauffman, L.H. Quantum Knots and Mosaics. *Quantum Inf. Process.* **2008**, *7*, 85–115. [CrossRef]
2. Lee, H.J.; Ludwig, L.D.; Paat, J.S.; Peiffer, A. Knot Mosaic Tabulation. *Involve* **2018**, *11*, 13–26. [CrossRef]
3. Heap, A.; Knowles, D. Tile Number and Space-Efficient Knot Mosaics. *J. Knot Theory Ramif.* **2018**, *27*. [CrossRef]
4. Heap, A.; Knowles, D. Space-Efficient Knot Mosaics for Prime Knots with Mosaic Number 6. *Involve* **2019**, *12*. [CrossRef]
5. Kuriya, T.; Shehab, O. The Lomonaco-Kauffman Conjecture. *J. Knot Theory Ramif.* **2014**, *23*. [CrossRef]
6. Adams, C. *The Knot Book: An Elementary Introduction to the Mathematical Theory of Knots*; American Mathematical Society: Providence, RI, USA, 2004.
7. Thistlethwaite, M.; Hoste, J. KnotScape. Available online: <http://www.math.utk.edu/~morwen/knotscape.html> (accessed on 1 September 2016)
8. Rolfsen, D. *Knots and Links*; Publish or Perish Press: Berkeley, CA, USA, 1976.
9. Cha, J.C.; Livingston, C. KnotInfo: Table of Knot Invariants. Available online: <http://www.indiana.edu/~knotinfo> (accessed on 1 June 2019).



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).