Optimal Results in Staged Self-Assembly of Wang Tiles

Rohil Prasad Jonathan Tidor

January 22, 2013

Abstract

The subject of self-assembly deals with the spontaneous creation of ordered systems from simple units and is most often applied in the field of nanotechnology. The self-assembly model of Winfree describes the assembly of Wang tiles, simulating assembly in real-world systems. We use an extension of this model, known as the staged self-assembly model introduced by Demaine et al. that allows for discrete steps to be implemented and permits more diverse constructions. Under this model, we resolve the problem of constructing segments, creating a method to produce them optimally. Generalizing this construction to squares gives a new flexible method for their construction. Changing a parameter of the model, we explore much simpler constructions of complex monotone shapes. Finally, we present an optimal method to build most arbitrary shapes.

1 Introduction

Self-assembly is the spontaneous creation of complex structures from a collection of basic shapes. Its study is most applicable to the construction of nanoscale objects where direct manipulation is not feasible. Self-assembly can be used in situations ranging from the manufacture of new fabrics to non-invasive surgery.

The self-assembly of Wang tiles was first studied by Winfree who described a model of their behavior and demonstrated that it could be implemented using DNA [1]. Many variations on this model have been studied over the years, with the most pertinent being Demaine et al.'s staged self-assembly model [2]. This model describes the assembly of non-rotatable Wang tiles with a glue on each edge that attaches to other glues of the same type in separate bins and stages.

However, algorithms developed under this model can easily be implemented in a more realistic situation of rotatable tiles and pairs of complementary glues. This model can be applied with a variety of real-world structures, for example, both units of DNA and structures made from two crossed T4 bacteriophage legs have been used to simulate this model of self assembly.

The problems arising in staged self-assembly consist of asking various shapes can be constructed efficiently—with few distinct glues and tiles and in little space and time. In addition, we ask which constructions are optimal—proceeding in as few stages as possible for a given amount of glues, tiles, and bins. A final consideration is the flexibility of a construction. If possible, we seek to make construction possible with variable numbers of glues, tiles and bins.

In Section 3, we first prove a theorem that provides a lower bound for the construction time of a line segment in terms of the numbers of bins and tiles. We demonstrate constructions that meet this bound up to a constant factor in many cases. We start with very small numbers of bins and then work upwards, covering all cases, except for extremely small numbers of tiles. Using the jigsaw method of Demaine et al. [2] we apply this construction to build $n \times n$ squares in Section 4. This construction is still logarithmic, like their original construction, but is now flexible, with the ability to take almost all numbers of bins and tiles.

We then move to construction in temperature $\tau = 2$ in Section 5. We demonstrate two constructions, one for squares and the other for right isosceles triangles that both work in logarithmic time and then generalize the methods used in those constructions to build all monotone shapes. We then introduce the Kolmogorov complexity to show that this construction is optimal. Finally, in Section 6, we return to temperature $\tau = 1$ to demonstrate a construction of radially monotone shapes. We also prove that this construction is optimal using the same Kolmogorov complexity method.

2 Definitions

Theoretical research into self-assembly attempts to create algorithms for the construction of specific shapes or classes of shapes while minimizing the resources used. The building-blocks of these assembly systems are Wang tiles, or simply *tiles*—non-rotatable unit squares with a *glue* on each edge. Each glue has a strength—a non-negative integer that describes the strength of the bond the glue forms with another glue of the same type. The glue with strength zero is the *null glue*, denoted $\emptyset$.

The strength of the bond necessary to hold two tiles together is a parameter of the assembly system known as the *temperature*, denoted τ. Specifically, two tiles will join together along an edge if the glues along that boundary match and have a strength that is at least that of the temperature. Multiple tiles can thus attach together to form large complexes known as *supertiles*. Two supertiles attach in an analogous manner; if the sum of the strengths of the matching pairs of glues along the boundary of the two supertiles is at least the temperature, the two will attach. If every internal edge of a supertile is joined by a matching pair of glue other than the null glue, the supertile is known as *fully connected*. We aim to only construct fully connected shapes as they are sturdier in the real world.

Tiles are divided into separate containers known as *bins*. In each bin, specific types of tiles and supertiles are mixed. An unlimited number of each type of these supertiles combine with each other in all possible ways. If these connections continue on indefinitely, that assembly system is considered not *unique*. We are concerned with the unique case, where all tiles and supertiles will attach a finite number of times before reaching a state where it can no longer attach to any other tile or supertile. Such supertiles are said to be *terminal*.

One *stage* consists of the tiles and supertiles in each bin combining to a terminal state. In the first stage, only individual tiles combine in the bins. Afterwards, the items that assemble in each bin are individual types of tiles as well as the terminally produced contents of any number of the other bins. Note that any bin can either include or not include the supertiles that it terminally produced in the last stage, but it does not ever include non-terminal contents from previous stages.

The specific pattern of mixing is described by a diagram known as the *mix graph*. This graph has a vertex for each pair of bin and stage. We denote mixing from one bin in one stage to another bin in the next stage by an edge between the corresponding vertices. In addition, each vertex is marked by some subset of the tileset which records which types of tiles are added in to each bin at each stage.

It is important to minimize the *glue complexity* of a construction—the number of distinct glues used in total. In the real world, creating many different glues is a technical task whose difficulty increases quickly with more glues as it must be checked that no pair of glues has an unwanted interaction with each other. Since the glue complexity g and the *tile complexity* t, the number of distinct tiles used, are linked by the inequality $\frac{g}{4} \leq t \leq g^4$ we typically are only concerned with minimizing the tile complexity of our assembly systems. In addition, we minimize the *bin complexity* B and the *stage complexity* s, denoting the maximum number of bins used at any stage and the number of total stages, respectively.

3 $1 \times n$ segment assembly

In this section we prove a theorem that provides a lower bound for the stage complexity of the construction of a $1 \times n$ segment in full generality and then demonstrate that this bound may be attained.

We will use the notation $g_1(x)g_2$ to refer to a $1 \times x$ supertile with glue g_1 on its left edge and glue g_2 on its right edge as well as null glues along the top and bottom edges. In addition, $g_1 g_2 = g_1(1)g_2$.

Theorem 3.1. *No two identical tiles may be present in the final $1 \times n$ segment that took the same path through the mix graph.*

Proof. Assume the contrary, that two copies of tile T are present in the final supertile that also took the same path through the mix graph.

We have assumed that the two T's will eventually both become part of a single supertile. Consider the stage before this connection happens. Let the two tiles T have terminally assembled into XTY and $X'TY'$ where X, Y, X', Y' are all (possibly empty) rows of tiles. Then in the next stage XTY and $X'TY'$ combine, along with the row of tiles Z. Without loss of generality, let XTY be to the left of $X'TY'$. Thus the final assembly is $\ldots XTYZX'TY' \ldots$.

However, since the construction takes place all in one row, the two bonding domains on the left and right of a supertile cannot interact. Thus, in addition to XTY and $X'TY'$, $X'TY$ and XTY' will be produced. Then, in the stage where XTY and $X'TY'$ are supposed to combine, this supertile will assemble into the infinitely long segment $\ldots X'TYZX'TYZX'TYZ \ldots$ This is a contradiction, so our original assumption must have been false. □

This theorem manifests itself in two types of lower bounds, one for $B = 1$ and one for $B > 1$. We provide constructions for all values of B, however, very small number of bins require special considerations to be taken into account, which increase the constant factor on the stage complexity and require more tiles.

Theorem 3.2. *Using T tiles and 1 bin, a $1 \times n$ segment may be constructed optimally in $O(\frac{n}{T})$ time.*

Proof. Since we are only using one bin, there are only s paths through the mix graph, one for each stage a tile could have been added at. Thus $sT \geq n$, for a lower bound on the number of stages of $\frac{n}{T}$.

In the rest of this proof we will take all indices modulo $T-2$. Let $t_i = g_i g_{i+1}$ as well as $t_L = \emptyset g_0$ and $t_R = g_{n-2}\emptyset$. Note that this defines exactly T tiles.

A $1 \times n$ segment can be made of the tiles $t_L t_0 t_1 t_2 \cdots t_{n-2} t_{n-3} t_R$. However, if these n tiles were all mixed together, they would connect infinitely.

Instead, we add the first $T-2$ tiles in the first stage, which assemble into a $1 \times (T-2)$ segment. In all other stages we add the next $T-3$ tiles which attach onto the end of the previously assembled supertile until we have formed the full $1 \times n$ segment in $O(\frac{n}{T})$ stages. □

Theorem 3.3. *There is an optimal construction of a $1 \times n$ segment using $B \geq 2$ bins and $T = \Omega(B^2)$ tiles in $O(\log_B \frac{n}{T})$ stages.*

Proof. The number of paths through the mix graph is $B^s + B^{s-1} + B^{s-2} + \ldots + B = \Theta(B^s)$, given s as the number of stages. Then $T \cdot B^s = \Omega(n)$ from Theorem 3.1, or $s = \Omega(\log_B \frac{n}{T})$.

To construct these segments, we first prove the following lemmas.

Lemma 3.4. *A $1 \times (2^n - 1)$ supertile may be constructed using 2 bins, 6 tiles, and n stages.*

Proof. We proceed by induction. In stage 1, we add the tile ab to bin 1 and cd to bin 2.

Assume $g_1(2^k - 1)g_2$ is produced by bin 1 and $g_3(2^k - 1)g_4$ is produced by bin 2 after k stages, where g_1, g_2, g_3, g_4 is some permutation of a, b, c, d. We produce two segments length $2^{k+1} - 1$ in the next stage is follows: Mix both bins into each other and the tile $g_2 g_3$ to bin 1 and $g_4 g_1$ to bin 2. As shown in Figure 1, the supertiles then assemble to form $g_1(2^{k+1} - 1)g_4$ in bin 1 and $g_3(2^{k+1} - 1)g_2$ in bin 2, as desired. □

Lemma 3.5. *A $1 \times n$ supertile may be constructed in 2 bins with $O(1)$ tile complexity and $O(\log n)$ stage complexity.*

Proof. First note that any positive integer n may be written as the sum of distinct integers in the form $2^k + 1$ or 1. This is simply because any n in the interval $2^k + 1 \leq n < 2^{k+1} + 1$ can be transfered into an interval with lower k by subtracting $2^k + 1$. Let $n = \sum_{k=0}^{x} 2^{r_k} + 1$ or

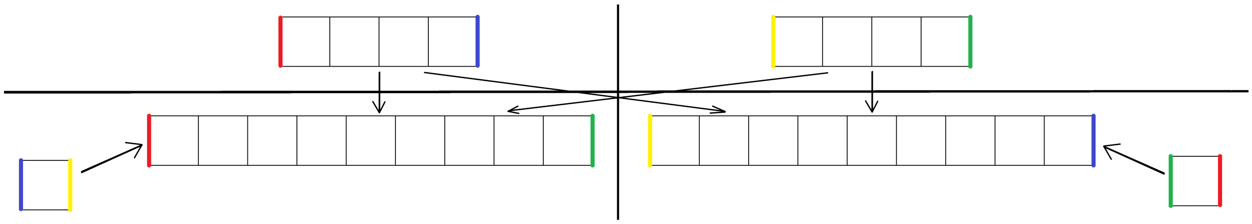

Figure 1: A generic stage in the 2 bin construction of a $1 \times n$ segment. The two segments assembled in the previous stage are combined in two different orders with an addition tile in each bin to create two segments one more than twice as long.

$1 + \sum_{k=0}^{x} 2^{r_k} + 1$, where the r_k are strictly increasing. Define $\{n_i\}$ to be the partial sums; in the first case, $n_i = \sum_{k=0}^{i} 2^{r_k} + 1$ and in the second, the n_i are 1 greater.

In this construction, we build segments length $2^k - 1$ in the same manner as the previous construction and then assemble the $1 \times n$ segment out of those pieces. We store a segment length n_i alongside the rest of the construction that grows sequentially into a segment length n.

In an extra stage after stage r_0, we construct the segment length n_0 by adding in tiles $\emptyset g_1$ and $g_2 e$ to the segment $g_1(2^{r_0} - 1)g_2$ in the first case, or $\emptyset f$, fg_1 and $g_2 e$ in the second case.

After stage r_k, in addition to the segments length $2^{r_k} - 1$, assume $\emptyset(n_{k-1})g$ is produced where g is either equal to e or f. In an additional stage, we attach the segment length $2^{r_k} - 1$ with two tiles to form $\emptyset(n_k)g'$, where g' is whichever of e, f that g was not, as shown in Figure 2.

In this manner, the stored segment grows until it reaches length n after $O(\log n)$ stages.

□

To incorporate more tiles into our assembly, using T^* times as many tiles, we construct T^* segments, each length $\lfloor \frac{n}{T^*} \rfloor$ or $\lceil \frac{n}{T^*} \rceil$, in a proportion such that the total length of all the segments produced is n. These supertiles will each have distinct sets of glues on their edges, except for the glues used to connect the segments that are only added in the last stage.

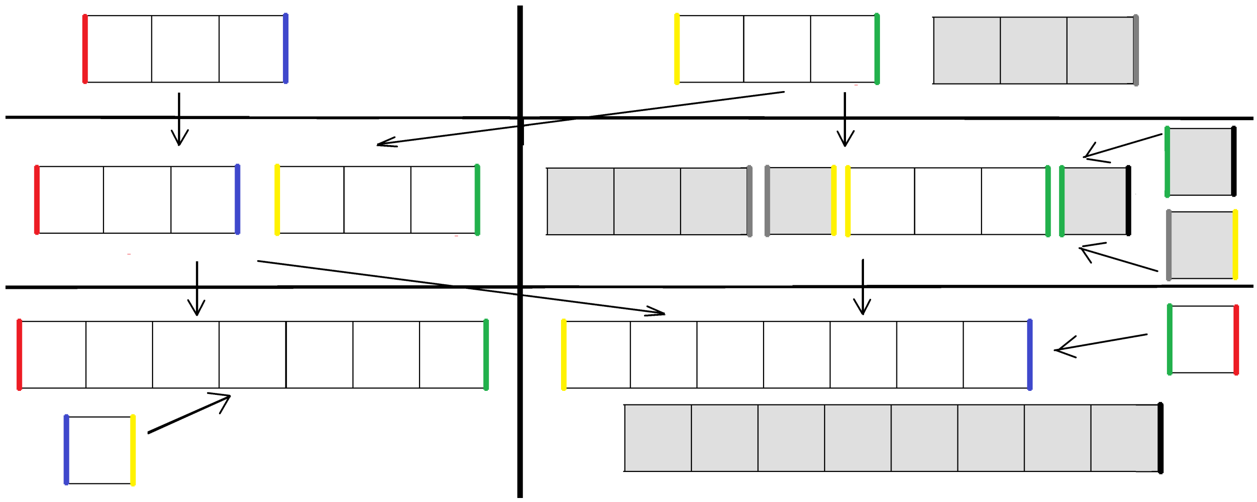

Figure 2: An additional stage in which the stored segment grows from length n_i to n_{i+1} while the rest of the construction proceeds normally.

Thus we can construct each of these segments in parallel since the glues in the different constructions will not interact with each other.

To build these segments we use glues h_x and h_{x+1} on the x^{th} supertile instead of the two null glues that were placed on the ends. The line segment construct proceeds normally until the last tile is added to the end of each of the segments when they all assemble into a long $1 \times n$ segment.

To incorporate more bins into our assembly, we first show that any integer can be written as the sum of integers in the form $B^k + 1$ for nonnegative integer k, or 1 with up to, but not including, B integers of each type. Then we proceed with essentially the same construction as that of Lemma 3.5 except that if tiles $g_i(B^k - 1)h_i$ for $1 \leq i \leq B$, are terminally produced in stage k, then in stage $k+1$, each bin receives all of these supertiles as well as $B-1$ of the B tiles in the form $h_i g_{i+1}$ ($g_{B+1} = g_1$). Since this combines B of the previous tiles at each stages, the total number of stages needed will be on the order of $\log_B \frac{n}{T}$.

Similar to the proof in Lemma 3.5, we simply note that we can always obtain a positive integer less than B^k+1 from a number n in the interval $B^k+1 \leq n < B^{k+1}+1$ by subtracting some number $i(B^k + 1)$ for $1 \leq i < B$ from it. If the remainder is ever reduced to 0, we are done. Otherwise, it will eventually be reduced to less than $B^0 + 1$ (i.e. equal to 1) and we

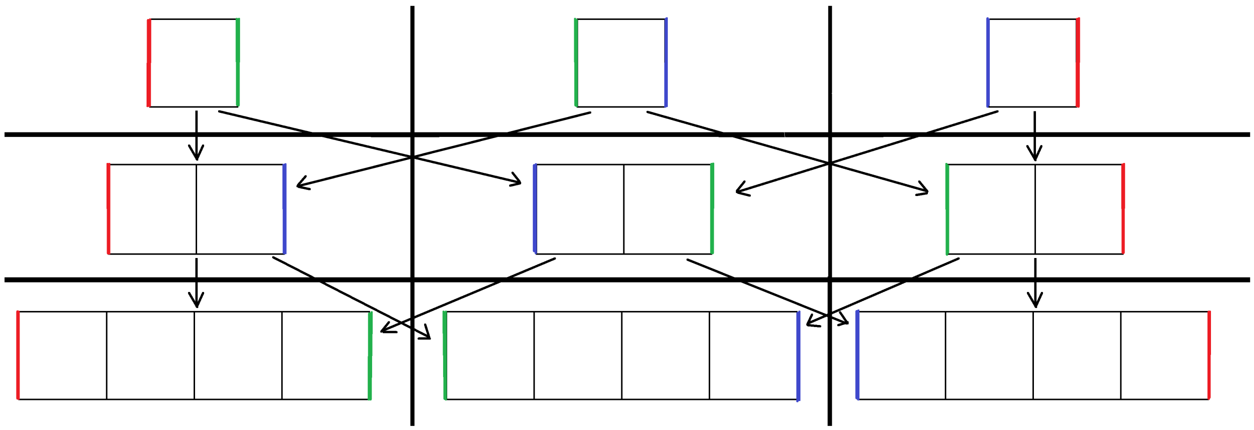

Figure 3: The basic step in the alternate construction of a $1 \times n$ line segment. All bins but one are mixed together in each stage.

finish with a single 1.

Thus we can always assemble a $1 \times n$ supertile in $O(\log_B \frac{n}{T})$ stage complexity as long as we have enough tiles so that $T^* \geq 1$, or equivalently, that $T = \Omega(B^2)$. □

Theorem 3.6. *There is an optimal construction of the $1 \times n$ supertile using $B \geq 4$ bins and $T = \Omega(B)$ tiles in $O(\log_B \frac{n}{T})$ stages.*

Proof. Given $g_i(k)g_{i+1}$ produced terminally in the i^{th} bin ($g_{B+1} = g_1$), a $1 \times (B-1)k$ supertile can be created in the next stage by combining $B-1$ of the produced supertiles in each bin. Starting with $B-1$ tiles in each bin, we produce supertiles length $(B-1)^k$ at the k^{th} stage, as shown in Figure 3.

We introduce one extra bin to create all other lengths. Let $n-2 = \overline{n_k n_{k-1} \ldots n_2 n_1 n_0}_{B-2}$, the base $B-2$ representation of $n-2$. In the final bin in stage $i+1$, add in the n_i segments of length $(B-2)^i$ that have the proper glues so that they will combine with the supertile already in the final bin to form one long segment. In addition, at the first stage, a tile with the null glue on the left edge is added to the final bin and a similar tile with the null glue on the right edge is added in the final stage. This creates a supertile of length $2 + \sum_{i=0}^{k} n_i(B-2)^i = n$.

This construction takes $O(\log_B n)$ stages. Using the same technique as in Theorem 3.3, this number can be reduced to $O(\log_B \frac{n}{T})$ stages as long as there are enough tiles to make one set; in this case, $\Omega(B)$. □

Theorem 3.7. *For $T < B$ tiles, the optimal construction of a $1 \times n$ supertile in $B \geq 4$ bins can be achieved in $O(\log_T n)$ stages.*

Proof. Let g be the glue complexity of the system. Note that T and g are related by $g/4 \leq T \leq g^4$. This implies that $\log g = \Theta(\log T)$.

Define $l(k)$ to be the maximum length of all terminally produced segments after stage k. We will show by contradiction that $\frac{l(k+1)}{l(k)} \leq g$.

Assume that this fraction is greater than g. This means that in stage $k+1$, at least $g+1$ supertiles join together into one long segment. Consider the glues on the left edges of the original supertiles that connect them to another supertile. There are at least g of these and none of them are the null glue, so there are only $g-1$ possibilities for each glue. By the pigeonhole principle, there are two supertiles that share the same glue on their left edge. However, the supertiles between the two copies of this glue will attach to themselves, forming a chain that extends infinitely. Thus this system does not uniquely assemble if $\frac{l(k+1)}{l(k)} > g$.

Therefore, $l(k) = O(g^k)$. Equivalently, to produce a $1 \times n$ supertile, the stage complexity is $\Omega(\log_g n) = \Omega(\log_T n)$.

To achieve this bound, we use the construction in Theorem 3.6: $T-2$ tiles are used in $T-2$ of the bins to assemble segments of length $(T-3)^k$. One more bin is used to hold the segments as they are completed with 2 more tiles to cap off the ends. □

4 Square assembly

The methods used in the construction of a $1 \times n$ segment can be extended to similar constructions in $\tau = 1$ for the $n \times n$ square through the use of the jigsaw method of Demaine et al. [2]. It is currently an open problem to determine whether the bounds that apply to the construction of segments also apply to that of the square in temperature 1.

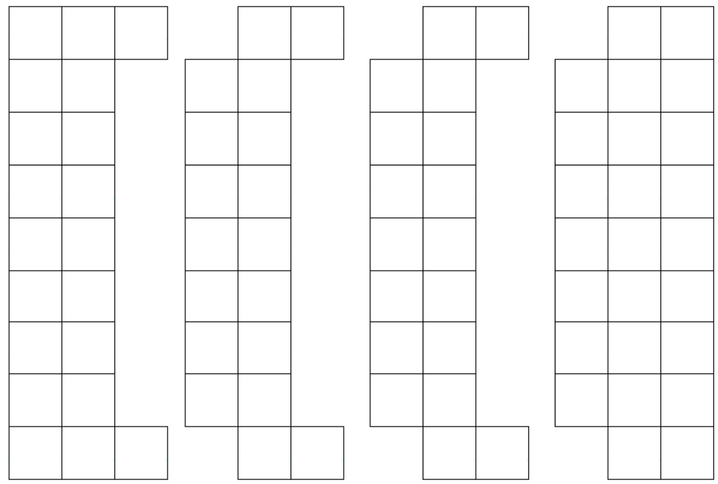

Figure 4: The jigsaw decomposition of a 9×9 square. These pieces can be assembled in a similar manner to the line and due to their shape, assemble with each other in exactly the same fashion as individual tiles building a line.

Theorem 4.1. *An $n \times n$ square may be constructed with tile complexity T and bin complexity B, with $O(\log_B \frac{n}{T})$ stage complexity if $T = \Omega(B)$ and with stage complexity $O(\log_T n)$ for all other T.*

Proof. The jigsaw pieces pictured in Figure 4 can be constructed in an analogous manner to a $1 \times n$ line segment through Theorems 3.6 and 3.7, by first building columns 1, 2, or 3 wide and $n-2$ tall and then adding on the correct top and bottom rows.

Then these jigsaw pieces can be assembled in a square again analogous to the construction of a $1 \times n$ segment, using the jigsaw pieces as the building blocks in the place of tiles. □

5 $\tau = 2$ constructions

Increasing the temperature to 2 allows for the construction of an $n \times n$ square in a much simpler manner and the construction of more complex shapes in a similar manner. This method can be applied to build arbitrary monotone shapes and is optimal in this case.

The following construction is related to one by Rothemund and Winfree [3] but uses ideas

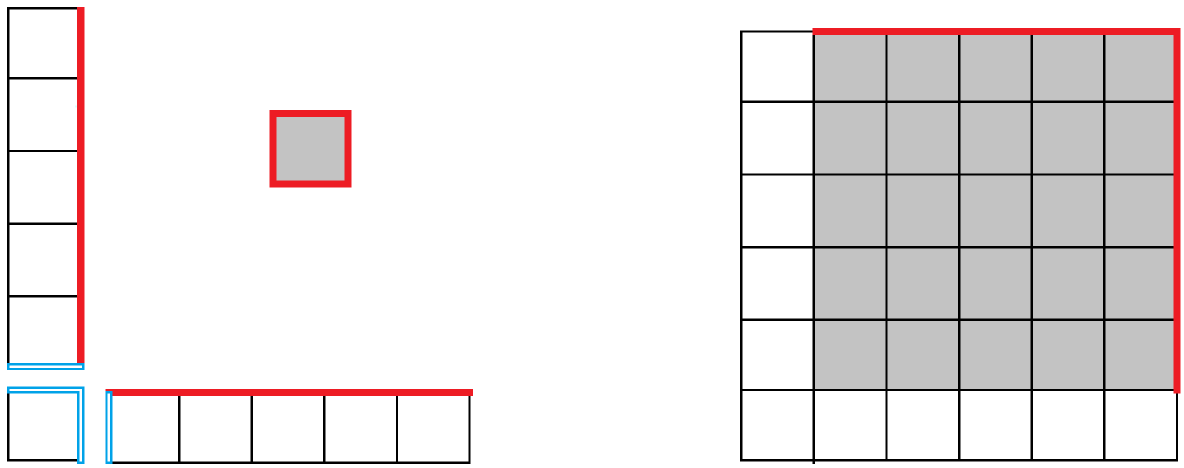

Figure 5: The construction of an $n \times n$ square. The border is constructed using double-strength glues but lined with single-strength glues, allowing it to be completed to a full $n \times n$ square in one stage with the addition of a single tile.

presented above to work with a constant number of tiles and glues.

Theorem 5.1. *An $n \times n$ square may be constructed in temperature $\tau = 2$ using $O(1)$ tiles and glues, 2 bins, and $O(\log n)$ stages.*

Proof. First we construct two $1 \times (n-1)$ segments using double-strength glues such that each has one side lined with a single-strength glue.

Each of the lines can be constructed with $O(1)$ tiles and glues in $O(\log n)$ stages, and they can be constructed in parallel due to their differing orientations. Now, in one additional stage, we add a single tile that connects the two together and another with the same single-strength glue on all sides to the bin where the two lines are constructed. The two lines form two sides of the square, while placing in the final tile fills in the rest, as shown in Figure 5. Since $\tau = 2$, each of these tiles can only connect to an area with two or more sides in the shared border having strength 1 glues. □

Theorem 5.2. *An $n \times n$ isosceles right triangle may be constructed in temperature $\tau = 2$ using $O(1)$ tiles and glues, 2 bins, and $O(\log n)$ stages.*

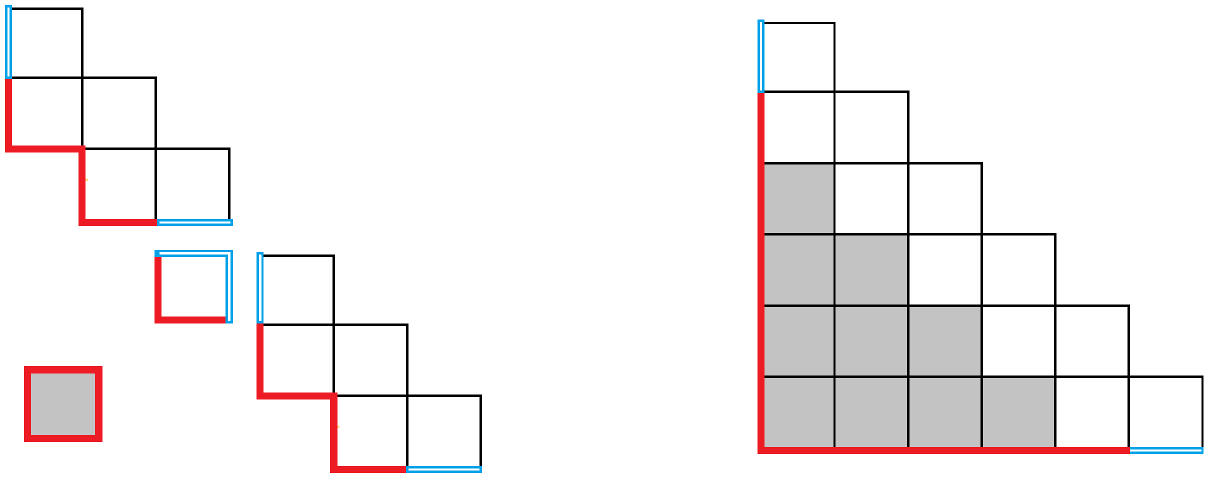

Figure 6: The construction of a $n \times n$ isosceles right triangle in a similar manner to that of the square. Instead of making two segments as the border, a staircase segment is constructed in a manner very similar to that of the segment.

Proof. We will construct the diagonal of the triangle and line it on one side with a single-strength glue as shown in Figure 6.

Assembly of this is in fact similar to the assembly of a line due to the symmetry of the shape. Instead of combining two segments with a tile between them, we instead combine two diagonal segments with a tile between them to produce one of just over twice the length.

Thus the diagonal of the triangle can be constructed with $O(1)$ tiles and glues and $O(\log n)$ stages. Then we place a tile with the same single-strength glue on all sides into the bin, and our right triangle is completed. □

Our constructions for the square and triangle use a border lined with a single-strength glue that we used to fill in the rest of the shape. To construct an arbitrary monotone shape, we first find an appropriate border and then complete it to the full shape.

Theorem 5.3. *Any monotone shape that fits within a square of side length n may be constructed at temperature $\tau = 2$ using $O(n)$ tiles and glues, 2 bins, and $O(\log n)$ stages.*

Proof. Assume our shape is x-monotone, meaning that every column is connected. In this

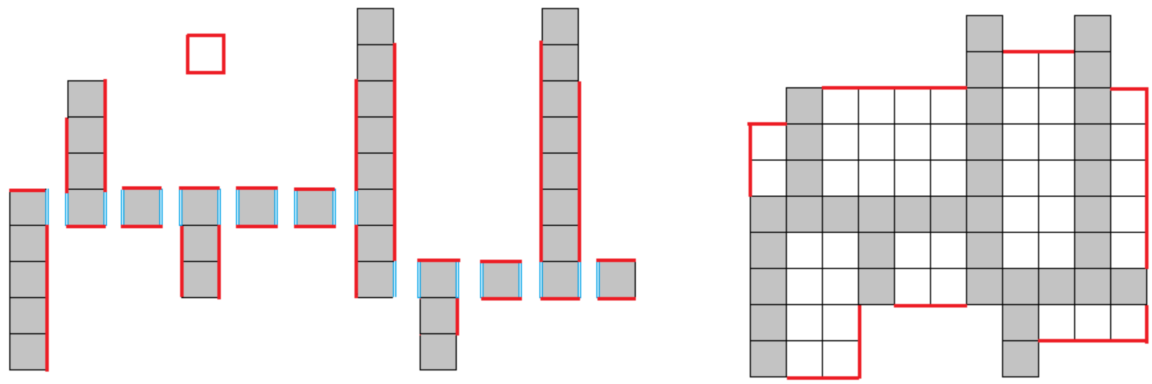

Figure 7: A simplified version of the construction of a monotone shape. The border, shown in gray, can be constructed as a series of segments and allows for the rest of the shape to be filled in in one extra stage.

construction we first choose a border and then fill it in with a small number of tiles in one stage, as shown in Figure 7. We pick a border by choosing columns that are at most two wide at the top and the bottom and then connecting them with a baseline horizontally through the shape.

The top border of our shape is made out of rectangles, and we can chose a set of columns simply by taking the leftmost column in each rectangle (Figure 8b). However, we run into problems if some of the rectangles are only one column wide. If we have chosen a consecutive block of columns, we first pick the columns that are higher than the columns on either side and as well as every other column and discard the rest. If there is more than one peak in a block, the columns we choose will be at most two wide when they join (Figure 8c.

Repeating the above process on the lower border of the shape and then choosing an arbitrary baseline produces a preliminary border (Figure 8d). We further require that the number of columns between any two in our border is either odd or equal to 2. We can modify our border by repeatedly dividing any region that is of even length greater than 2 into smaller pieces (Figure 8e).

We now construct the border out of line segments, lined with single-strength glues as follows: the horizontal segments are lined with n on the top and m on the bottom. The

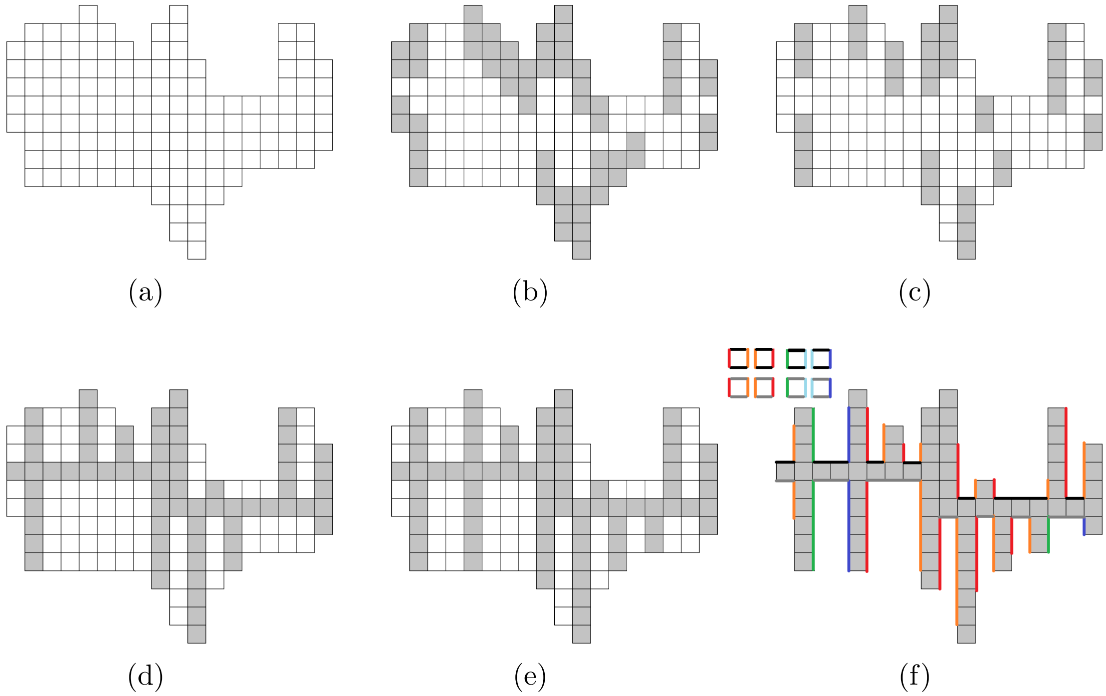

Figure 8: The complete border of a monotone shape. (a) the monotone shape (b) choosing columns preliminarily (c) removing from consecutive blocks (d) adding the baseline (e) removing even sections wider than 2 (f) the proper glues

vertical segments that are on the left edge of a region of odd width are lined with a and on the right with b. The vertical segments on the left edge of a region of width 2 are lined with x and on the right with z (Figure 8f).

We fill in the border with the tiles ab, ba, xy and yz lined with m on the top and bottom to fill in the top half and n on the top and bottom to fill in the bottom half.

There are $O(n)$ rows and columns in the border, where n is the side length of the smallest square that encloses the shape, so assembling all line segments in parallel with distinct tilesets gives the desired complexities. □

Theorem 5.4. *Using $O(n)$ tiles, $O(1)$ bins, any system that assembles arbitrary monotone shapes of diameter n must have stage complexity of at least $\Omega(\log n)$.*

Proof. The number of bits needed to describe a monotone shape of diameter n is $\Omega(n \log n)$. To see this, simply note that that is the number of bits need to describe the top of each of the columns. In actuality, the true number of bits is between $n \log n$ and $2n \log n$. It is a basic fact of the Kolmogorov complexity that the number of bits needed to input into a universal Turing machine that outputs a number k is at least $\log k$ for almost all k. Since assembly is deterministic and computable, an assembly system that produces an arbitrary monotone shape of diameter n essentially is a Turing machine that outputs a number up to $2^{n \log n}$.

The number of bits in the input of this Turing machine is the number of bits needed to describe its mix graph. This is $O(sB^2)$ for the edges of the mix graph and $O(sBT)$ for which tiles are added. With the parameters of this assembly system, we have $s + sn \geq cn \log n$ for some constant c, or that $s = \Omega(\log n)$, as desired. □

6 Assembly of arbitrary shapes

Given any shape S, denote an arbitrary tile in S as its center. The distance between a tile x and the center is denoted by $|x|$ and is defined as the number of squares in the shortest path between it and the center tile. The radius r of S is defined as $\max_{x \in S} |x|$.

Theorem 6.1. *The construction of arbitrary shapes of radius r, with tile complexity $O(r)$ and bin complexity $O(1)$ must have stage complexity $\Omega(r)$.*

Proof. We use the same technique introduced in Theorem 5.4. First we will show that the number of bits needed to determine a shape of radius r is $\Omega(r^2)$. We do so by enumerating a subset of the shapes of radius r for which this is count is still accurate.

Let $f(r)$ be the number of shapes radius r that include all squares on the main diagonals, and also include at least one of every two squares which are diagonally adjacent, and the same distance from the center. We wish to show that $f(r) = 2^{\Omega(r^2)}$.

Clearly $f(1) = 1$. The tiles within distance $r - 1$ from the center of any shape counted in $f(r)$ form a shape counted in $f(r - 1)$. In addition, because of the conditions on the placement of tiles, any tile distance r from the center can be added to the subshape and remain connected. Thus $\frac{f(r)}{f(r-1)}$ is the number of ways to choose a subset of $r-2$ squares that includes at least one of every two consecutive squares, raised to the fourth power. It is well known that this is the F_r for $r \geq 3$. Thus $f(r) = \prod_{i=2}^{r} F_i^4$. Since $\log F_n = \Omega(\log \phi^n) = \Omega(n)$, $\log f(r) = 4 \sum_{i=2}^{r} \Omega(i) = \Omega(r^2)$, as desired.

Now by the Kolmogorov complexity, we know that the number of bits needed to describe the mix graph, $sB^2 + sBT$, is at least r^2. Thus $s + sr \geq cr^2$ for some constant c, which implies that $s = \Omega(r)$, as desired. □

Define a shape S to be *radially monotone* if there exists a choice of center such that $\forall x \in S$, there exists a path from the center to x that consists of tiles in S that are all within $|x|$ of the center. We provide a construction of arbitrary radially monotone shapes with $B = 1$ that runs in $O(r)$ stage complexity.

Theorem 6.2. *Using tile and glue complexities of $O(r)$, there exists a construction of arbitrary radially monotone shapes radius r in 1 bin that has a stage complexity of $O(r)$.*

Proof. We begin with the construction of a very simple shape and then show that with the same parameters that construction can be adapted to all radially monotone shapes. That shape is a diamond of radius r—all tiles around a center of distance at most r.

Lemma 6.3. *A diamond of radius r can be constructed with tile and glue complexities of $O(r)$, 1 bin, and $O(r)$ stages.*

Proof. The shape is built in rings, tiles equidistant from the center, each taking one stage to construct. The first ring is just a single tile, labeled 0 in Figure 9. For clarity, we color this and all other odd numbered rings white and all even numbered rings gray. To construct the n^{th} ring, we add the tiles of the first type in the proper color for $i = 1, 2, \ldots, n-2$, as well as the four tiles of the second type in the proper color for $i = n$. The tiles of the first type make up the body of the shape while the tiles of the second make up the diagonals, as shown in Figure 10. After r stages, a diamond of radius r has been constructed, as desired. □

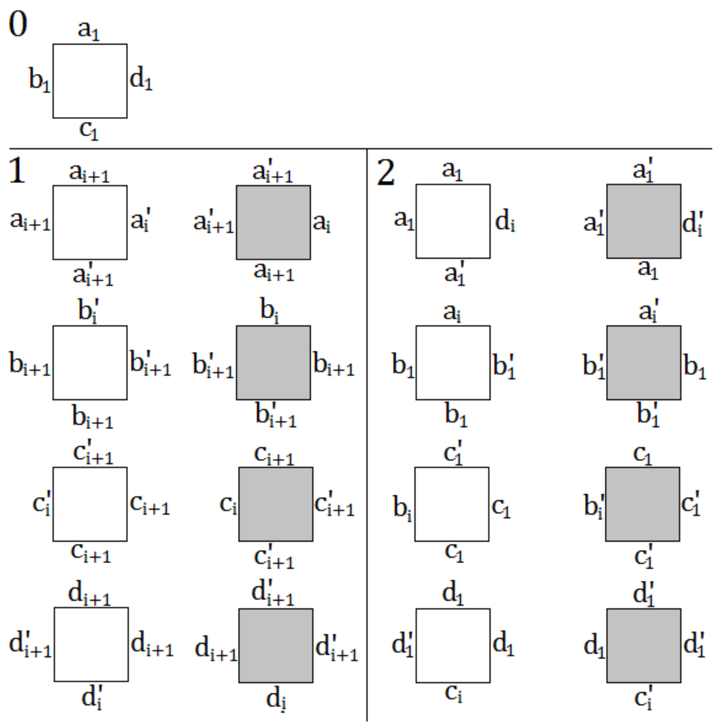

Figure 9: The tiles used in the assembly of arbitrary radially monotone shapes. Tile 0 is the center piece used in the first stage of the construction. Tiles 1 are used in alternating stages with $1 \leq i \leq r-2$ for the body of the shape. Tiles 2 are also used in alternating stages with $2 \leq i \leq r$ for the diagonals of the shape.

The construction of an arbitrary radially monotone shape of radius r proceeds similarly. Instead of using only the tileset shown in Figure 9, we use 15 times as many tiles—including a version of each tile with each proper subset of its glues replaced by the null glue. Now, at each stage, we add in the correct version of each tile required so that there are null glues on all the outside edges and omit tiles which would fill in a square no longer part of the final shape, as shown in Figure 11. Each tile will have at least one tile that it binds to already in the supertile because the shape is radially monotone, and the construction proceeds with the same complexities. □

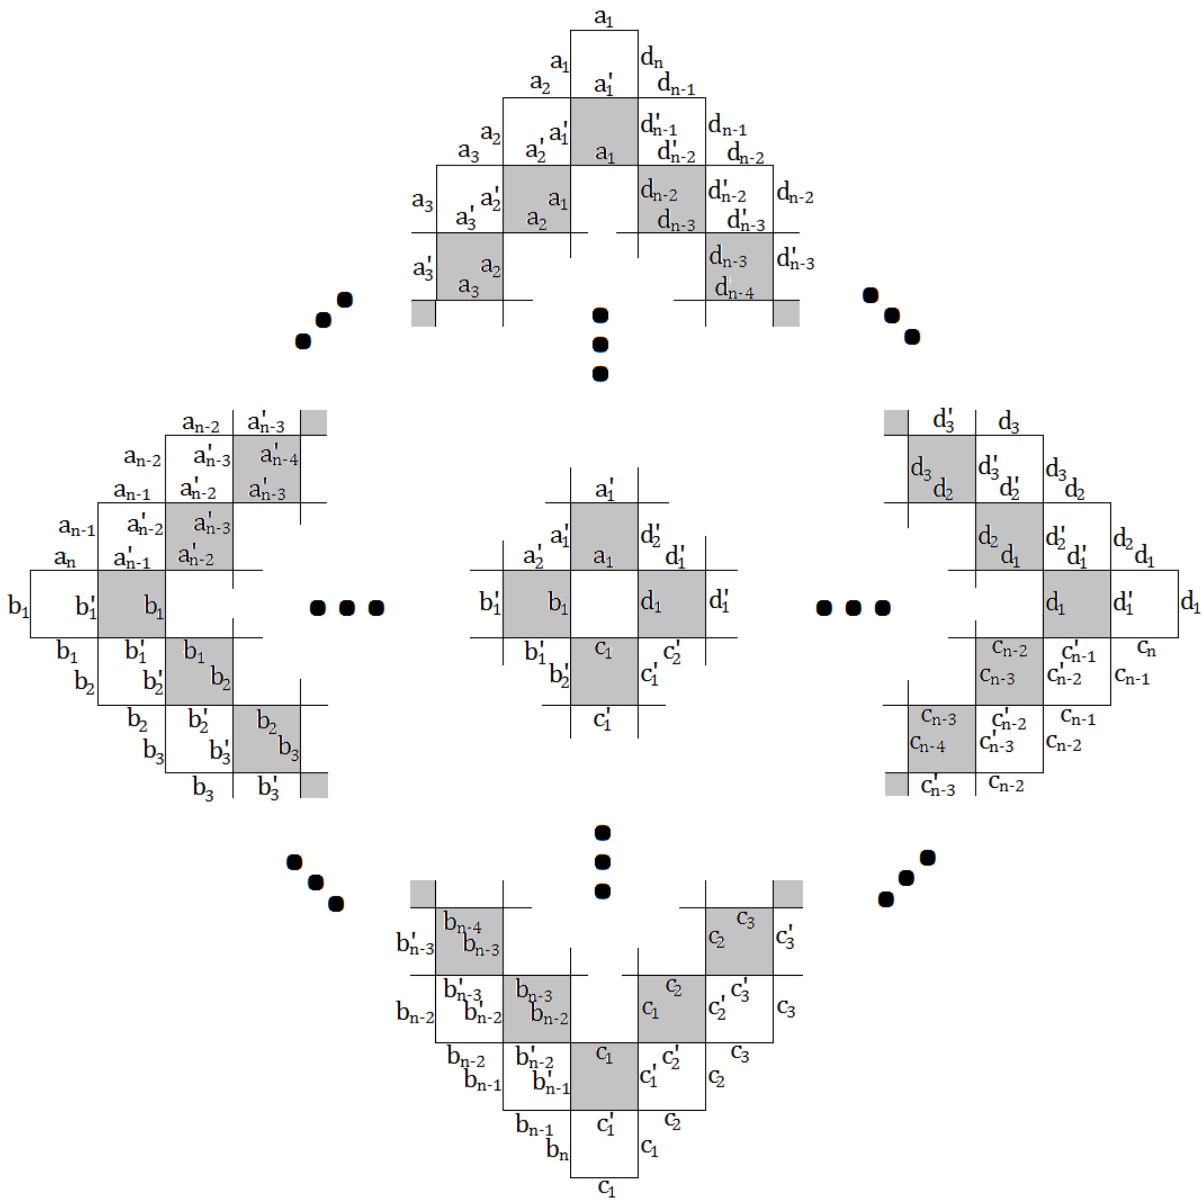

Figure 10: Assembly of the full n^{th} ring in the construction of arbitrary radially monotone shapes. Each ring after the first is made up of $4(n-1)$ distinct tiles, $4(n-2)$ of type 1 and 4 of type 2. All rings of the same parity use the same set of tiles of type 1.

7 Summary and Open Problems

For $1 \times n$ line segments, we have demonstrated an optimal and flexible construction that can use nearly any desired number of bins and tiles. In addition, we have created optimal con-

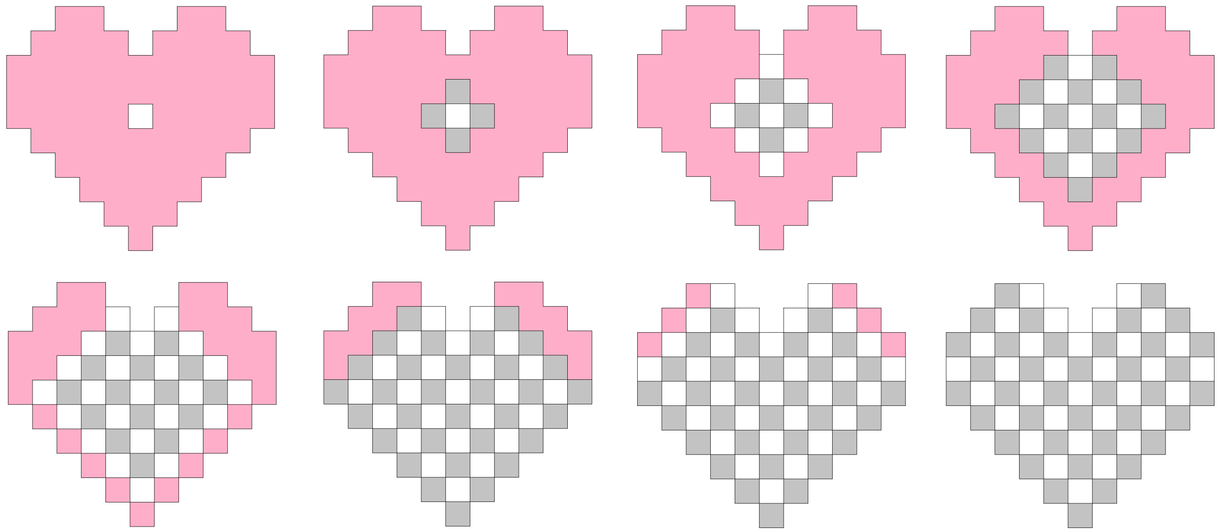

Figure 11: Successive stages in the assembly of a radially monotone shape. Tiles that would not be inside the final figure are simply not added and tiles that are on the edge are replaced by versions with the null glue so that no unwanted attachment occurs.

struction of monotone and radially monotone shapes that sacrifice constant tile complexities for the ability to construct broader classes of shapes.

We also have a flexible construction of squares, though it is an open problem whether this construction is optimal in temperature 1 or if faster constructions exist. In addition, making flexible constructions of more shapes, namely monotone and radially monotone is worth considering. Finally, it would be interesting to determine other classes of shapes that can be constructed quickly as well as constructions for the non-radially monotone shapes.

Two useful generalizations of this model exist which have not been studied deeply to the best of our knowledge. The first is construction in three or more dimensions which could prove useful in the real world. The second is a probabilistic version of this model which is more applicable to the real world by not assuming that all assembly completely occurs, and instead leaves some fraction of non-terminal structures present.

Acknowledgments. We would like to thank Jesse Geneson, without whom this research would have been impossible as well as the the MIT PRIMES Program through which this research was done.

References

[1] Erik Winfree, “Algorithmic Self-Assembly of DNA”, In Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy, California Institute of Technology, Pasadena, California, 1998.

[2] Erik D. Demaine, Martin L. Demaine, Sándor P. Fekete, Mashhood Ishaque, Eynat Rafalin, Robert T. Schweller, and Diane L. Souvaine, “Staged Self-Assembly: Nanomanufacture of Arbitrary Shapes with O(1) Glues”, *Natural Computing*, volume 7, number 3, September 2008, pages 347370. Special issue of selected papers from the 13th International Meeting on DNA Computing, 2007.

[3] Paul W. K. Rothemund and Erik Winfree. The program-size complexity of self-assembled squares. In Proceedings of the 32nd Annual ACM Symposium on Theory of Computing, pages 459–468, 2000.