

The tile assembly model is intrinsically universal

David Doty* Jack H. Lutz† Matthew J. Patitz‡ Robert T. Schweller§
Scott M. Summers¶ Damien Woods||

Abstract

We prove that the abstract Tile Assembly Model (aTAM) of nanoscale self-assembly is *intrinsically* universal. This means that there is a single tile assembly system \mathcal{U} that, with proper initialization, simulates any tile assembly system \mathcal{T} . The simulation is “intrinsic” in the sense that the self-assembly process carried out by \mathcal{U} is exactly that carried out by \mathcal{T} , with each tile of \mathcal{T} represented by an $m \times m$ “supertile” of \mathcal{U} . Our construction works for the full aTAM at any temperature, and it faithfully simulates the deterministic or nondeterministic behavior of each \mathcal{T} .

Our construction succeeds by solving an analog of the cell differentiation problem in developmental biology: Each supertile of \mathcal{U} , starting with those in the seed assembly, carries the “genome” of the simulated system \mathcal{T} . At each location of a potential supertile in the self-assembly of \mathcal{U} , a decision is made whether and how to express this genome, i.e., whether to generate a supertile and, if so, which tile of \mathcal{T} it will represent. This decision must be achieved using asynchronous communication under incomplete information, but it achieves the correct global outcome(s).

*Computing and Mathematical Sciences, California Institute of Technology, Pasadena, CA 91125, USA. ddoty@caltech.edu. This author’s research was supported by a Computing Innovation Fellowship under NSF grant 1019343.

†Computer Science, Iowa State University, Ames, IA 50011 USA. lutz@cs.iastate.edu. This author’s research was supported by NSF grants 0652569 and 1143830. Part of this work was done during a sabbatical at Caltech and the Isaac Newton Institute for Mathematical Sciences at the University of Cambridge.

‡Computer Science, University of Texas–Pan American, Edinburg, TX, 78539, USA. mpatitz@cs.panam.edu. This author’s research was supported in part by NSF grant CCF-1117672.

§Computer Science, University of Texas–Pan American, Edinburg, TX, 78539, USA. schweller@cs.panam.edu. This author’s research was supported in part by NSF grant CCF-1117672.

¶Computer Science and Software Engineering, University of Wisconsin–Platteville, Platteville, WI 53818, USA. summerss@uwplatt.edu.

||Computer Science, California Institute of Technology, Pasadena, CA 91125, USA. woods@caltech.edu. This author’s research was supported by NSF grant 0832824, the Molecular Programming Project.

1 Introduction

Structural DNA nanotechnology, pioneered by Seeman in the 1980s [44], exploits the information-processing capabilities of nucleic acids to engineer complex structures and devices at the nanoscale. This is a very “hands-off” sort of engineering: The right molecules are placed in solution, and the structures and devices self-assemble spontaneously according to the principles of chemical kinetics. Controlling such self-assembly processes is an enormous technical challenge, but impressive progress has already been made. Regular arrays [49], polyhedra [27], fractal structures [25, 42], maps of the world [43], curved three-dimensional vases [26], DNA tweezers [50], logic circuits [38], neural networks [39], and molecular robots [34] are a few of the nanoscale objects that have self-assembled in successful laboratory experiments. Motivating future applications include smaller, faster, more energy-efficient computer chips, single-molecule detection, and in-cell diagnosis and treatment of disease.

Theoretical computer science became involved with structural DNA nanotechnology just before the turn of this century. In his 1998 Ph.D. thesis, Winfree introduced a mathematical model of DNA tile self-assembly and proved that this model is Turing-universal, i.e., that it can simulate any Turing machine [48]. This implies that nanoscale self-assembly can be algorithmically directed, and that extremely complex structures and devices can in principle be engineered by self-assembly. Rothemund and Winfree [41] subsequently refined this model slightly, formulating the *abstract Tile Assembly Model (aTAM)*. The (two-dimensional) aTAM is an idealized model of error-free self-assembly in two dimensions that has been extensively investigated [1–8, 12–15, 18–20, 23, 28, 29, 35, 45, 46] and is the subject of this paper.

Very briefly, a *tile* in the aTAM is a unit square with a kind and strength of “glue” on each of its sides. A *tile assembly system* \mathcal{T} consists of a finite collection T of tile types (with infinitely many tiles of each type in T available), a *seed assembly* σ consisting of one or more tiles of types in T , and a *temperature* τ . Self-assembly proceeds from the seed assembly σ , with tiles of types in T successively and nondeterministically attaching themselves to the existing assembly. Two tiles placed next to each other *interact* if the glues on their abutting sides match, and a tile *binds* to an assembly if the total strength on all of its interacting sides is at least τ . A more complete description of the aTAM appears in Section 2.

Our topic is the *intrinsic universality* of the abstract Tile Assembly Model. We now explain what this means, starting with what it does *not* mean. By Winfree’s above-mentioned result, there is a tile assembly system \mathcal{U} that simulates a universal Turing machine. This universal Turing machine, and hence \mathcal{U} , can simulate any tile assembly system \mathcal{T} (in fact, there are various aTAM software simulators available, e.g., [37]). But this is only a *computational* simulation. It *tells* us what \mathcal{T} does, but it does not actually *do* what \mathcal{T} does. The task of a Turing machine is to perform a computation, and a universal Turing machine performs the same computation as a machine that it simulates. The task of a tile assembly system is to perform the process of self-assembly, so a universal tile assembly system should *perform* the same self-assembly process as a tile assembly system that it simulates. This is what is meant by intrinsic universality.

This paper proves that the abstract Tile Assembly Model is intrinsically universal.

This means that there is a single tile set U that, with proper initialization (calling the initialized system \mathcal{U}), simulates any tile assembly system \mathcal{T} . The simulation is “intrinsic” in the sense that the self-assembly process carried out by \mathcal{U} is exactly that carried out by \mathcal{T} , with each tile of \mathcal{T} represented by an $m \times m$ “supertile” of \mathcal{U} . Our construction works for the full aTAM at any temperature (the simulating system \mathcal{U} uses temperature 2), and it faithfully simulates the

deterministic or nondeterministic behavior of each \mathcal{T} . This notion was studied by Ollinger [36] and others [11, 16, 24, 30–32] in the context of cellular automata and Wang tiling,

Our construction succeeds by solving an analog of the cell differentiation problem in developmental biology: Each supertile of \mathcal{U} , starting with those in the seed assembly, carries a complete encoding of the simulated system \mathcal{T} (the “genome” of \mathcal{T}) along each of its sides, which we call “supersides”. (This genome accounts for most of the m tiles of \mathcal{U} that appear on each superside. Additional tiles along the superside identify the glue of the simulated tile of \mathcal{T} and support a variety of communication mechanisms.) At each location of a potential supertile a decision is made whether and how to express this genome, i.e., whether to generate a supertile and, if so, which tile of \mathcal{T} it will represent. (This latter choice will be nondeterministic precisely insofar as \mathcal{T} is nondeterministic at this location.) This decision depends on very limited local information, but it achieves the correct global outcome(s). The self-assembly of \mathcal{U} is thus a “developmental process” in which “supertile differentiation” is governed by local communication, while the “genome” is passed intact from supertile to supertile.

Our construction uses three basic interacting primitives to carry out the asynchronous communication under imperfect information needed for supertile differentiation and genome copying. These mechanisms are called frames, crawlers, and probes. The *frame* of a potential supertile consists of four layers of tiles just inside each extant superside. This frame is used for communication with adjacent supersides, which may or may not exist. Much of its function is achieved by a symmetry-breaking “competition” at each corner. Our construction uses many types of *crawlers*, which are messengers that copy and carry various pieces of information from place to place in the supertile. The *probes* of a superside are used to communicate with the opposite superside, which may or may not exist. The challenge is to program all this activity without ever blocking a path that may later be needed for intra-supertile communication. This summary is greatly oversimplified. An overview of our construction is presented in Section 4, and the full construction is presented in Sections 5–10. We have used various symmetries to simplify the construction and its presentation, and we hope that the sheer number of cases does not obscure the underlying elegance of the machinery.

Our result shows that the aTAM is universal for itself, without recourse to indirect simulations by Turing machines or other models that obscure important properties of the model. For example, our result shows that the tile assembly model is able to simulate local interactions between tiles, nondeterminism, and tile growth processes in general, all on a global scale. Thus our intrinsically universal tile set captures, in a well defined way, all properties of any tile assembly system.

Intrinsic universality, with its precise notion of “simulate”, has applications to the theory of self-assembly. Firstly, a useful type of simulation is where one shows that for all aTAM systems \mathcal{T} there exists a tile assembly system \mathcal{T}' in some *other* self-assembly model that simulates \mathcal{T} . This style of $\forall \mathcal{T}, \exists \mathcal{T}'$ -simulation has been used previously [5, 9, 17] and is useful when comparing the power of tile assembly models. However, combining such a simulation statement with the statement of our main result gives the immediate corollary that there is a *single* set of tiles U in the other model that, when appropriately seeded, simulates *any* aTAM tile assembly system \mathcal{T} . Hence our main result automatically shows the existence of a single, very powerful, tile set in the other model, a seemingly strong statement. Secondly, and more speculatively, our result opens the possibility for new research directions in self-assembly. For example, taken together with the result in [22], we now know of two classes of tile assembly systems that exhibit intrinsic universality: the full aTAM (our main result), and the more restricted locally consistent systems [22]. This gives a kind of closure property for these classes of systems. In the field of cellular automata, the notion of

intrinsic universality has led to the development of formal tools to classify models of computation in terms of their ability to simulate each other [16]. The intrinsically universal cellular automata sit at the top of this “quasi-order”. As an example of a concrete application of this work, the notion of intrinsic universality has been used [10, 11] to show that various elementary cellular automata are strictly less powerful than others. Specifically, it was shown that the communication complexity of those systems is too low for them to exhibit intrinsic universality, and so there is a wide range of behaviors they can never achieve. Such statements crucially make use of the fact that intrinsic universality uses a tight notion of “simulate”. In tile self-assembly, we currently have very few tools by which to compare the abilities of models; the main comparisons essentially boil down to comparing tile complexity, or establishing whether or not the system can simulate Turing machines and thus make arbitrary computable shapes. Both comparisons, especially the latter, are necessarily rather coarse for comparing the expressibility of models, and we hope that our result, and the notion of “simulate” that we use, can inspire the development of work that elucidates a fine-grained structure for self-assembly.

We conclude this introduction with a brief discussion of related work. The most recent precursor is [22], in which some of the present authors showed that a restricted submodel of the aTAM is intrinsically universal. This was an extensive, computationally expressive submodel of the aTAM, but its provisos (temperature 2, no glue mismatches, and no binding strengths exceeding the temperature) were artificially restrictive, awkward to justify on molecular grounds, and inescapable from the standpoint of that paper’s proof technique. Our approach here is perforce completely different. Both papers code the simulated system’s genome along the supersides, but the resemblance ends there. The frames, crawlers, and probes that we use here are new. (The “probe-like” structures in [22] are too primitive to work for simulating the full aTAM.)

As noted in [22], constructions of Soloveichik and Winfree [45] and Demaine, Demaine, Fekete, Ishaque, Rafalin, Schweller, and Souvaine [18] can be used to achieve versions of intrinsic universality for tile assembly at temperature 1, but this appears to be a severe restriction. Additionally, the latter paper uses a generalized version of the aTAM (i.e., “hierarchical” self-assembly or the “two-handed” aTAM) that has a mechanism for long-range communication that is lacking in the standard aTAM and that obviates the need for the distributed communication mechanisms we employ to build supertiles. Also discussed in [22] are studies of universality in Wang tiling [47] such as those by Lafitte and Weiss [30–32]. While these studies are very significant in the contexts of mathematical logic and computability theory, they are concerned with the *existence* of tilings with no mismatches, and not with any *process* of self-assembly. In particular, most attempts to adapt the constructions of Wang tiling studies (such as those in [30–32]) to self-assembly result in a tile assembly system in which many junk assemblies are formed due to incorrect nondeterministic choices being made that arrest any further growth and/or result in assemblies that are inconsistent with the desired output assembly. We therefore require novel techniques to ensure that the only produced assemblies are those that represent the intended result or valid partial progress toward it. Furthermore, techniques used in constructing intrinsically universal cellular automata do not carry over to the aTAM as the models have fundamental differences; in particular, when a tile is placed it remains in-place forever, whereas cellular automata cells can be reused indefinitely. In fact, many of the challenging issues in proving our result are related to the fact that tiles, once placed, can block each other and, of course, that self-assembly is a highly asynchronous and nondeterministic process.

2 Abstract Tile Assembly Model

This section gives a brief informal sketch of the abstract Tile Assembly Model (aTAM). See Section A for a formal definition of the aTAM.

A *tile type* is a unit square with four sides, each consisting of a *glue label* (often represented as a finite string) and a nonnegative integer *strength*. We assume a finite set T of tile types, but an infinite number of copies of each tile type, each copy referred to as a *tile*. An *assembly* (a.k.a., *supertile*) is a positioning of tiles on the integer lattice \mathbb{Z}^2 ; i.e., a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T$. Let \mathcal{A}^T denote the set of all assemblies of tiles from T , and let $\mathcal{A}_{<\infty}^T$ denote the set of finite assemblies of tiles from T . Write $\alpha \sqsubseteq \beta$ to denote that α is a *subassembly* of β , which means that $\text{dom } \alpha \subseteq \text{dom } \beta$ and $\alpha(p) = \beta(p)$ for all points $p \in \text{dom } \alpha$. Two adjacent tiles in an assembly *interact* if the glue labels on their abutting sides are equal and have positive strength. Each assembly induces a *binding graph*, a grid graph whose vertices are tiles, with an edge between two tiles if they interact. The assembly is τ -*stable* if every cut of its binding graph has strength at least τ , where the weight of an edge is the strength of the glue it represents. That is, the assembly is stable if at least energy τ is required to separate the assembly into two parts.

A *tile assembly system* (TAS) is a triple $\mathcal{T} = (T, \sigma, \tau)$, where T is a finite set of tile types, $\sigma : \mathbb{Z}^2 \dashrightarrow T$ is a finite, τ -stable *seed assembly*, and τ is the *temperature*. An assembly α is *producible* if either $\alpha = \sigma$ or if β is a producible assembly and α can be obtained from β by the stable binding of a single tile. In this case write $\beta \rightarrow_1^{\mathcal{T}} \alpha$ (α is producible from β by the attachment of one tile), and write $\beta \rightarrow^{\mathcal{T}} \alpha$ if $\beta \rightarrow_1^{*\mathcal{T}} \alpha$ (α is producible from β by the attachment of zero or more tiles). When \mathcal{T} is clear from context, we may write \rightarrow_1 and \rightarrow instead. An assembly is *terminal* if no tile can be τ -stably attached to it. Let $\mathcal{A}[\mathcal{T}]$ be the set of producible assemblies of \mathcal{T} , and let $\mathcal{A}_{\square}[\mathcal{T}] \subseteq \mathcal{A}[\mathcal{T}]$ be the set of producible, terminal assemblies of \mathcal{T} . A TAS \mathcal{T} is *directed* (a.k.a., *deterministic, confluent*) if $|\mathcal{A}_{\square}[\mathcal{T}]| = 1$.

We make the following assumptions that do not affect the fundamental capabilities of the model, but which will simplify our main construction. Since the behavior of a TAS $\mathcal{T} = (T, \sigma, \tau)$ is unchanged if every glue with strength greater than τ is changed to have strength exactly τ , we assume henceforth that all glue strengths are in the set $\{0, 1, \dots, \tau\}$. We assume that glue labels are never shared between glues of unequal strength.

3 Main result

To state our main result, we must formally define what it means for one TAS to “simulate” another. We focus in particular on a sort of “direct simulation” via block replacement ($m \times m$ blocks of tiles in the simulating system represent single tiles in the simulated system). The intuitive goal of the following definition is identical to that in [22], and corrects some subtle errors there.

Let $m \in \mathbb{Z}^+$. An *m -block supertile* over tile set T is a partial function $\alpha : \mathbb{Z}_m \times \mathbb{Z}_m \dashrightarrow T$, where $\mathbb{Z}_m = \{0, 1, \dots, m-1\}$. Let B_m^T be the set of all m -block supertiles over T . The m -block with no domain is said to be *empty*. For a general assembly $\alpha : \mathbb{Z}^2 \dashrightarrow T$ and $x, y \in \mathbb{Z}$, define $\alpha_{x,y}^m$ to be the m -block supertile defined by $\alpha_{x,y}^m(i, j) = \alpha(mx + i, my + j)$ for $0 \leq i, j < m$. A partial function $R : B_m^S \dashrightarrow T$ is said to be a *valid m -block supertile representation* from S to T if for any $\alpha, \beta \in B_m^S$ such that $\alpha \sqsubseteq \beta$ and $\alpha \in \text{dom } R$, then $R(\alpha) = R(\beta)$.

For a given valid m -block supertile representation function R from tile set S to tile set T , define the *assembly representation function* $R^* : \mathcal{A}^S \rightarrow \mathcal{A}^T$ such that $R^*(\alpha') = \alpha$ if and only if

$\alpha(x, y) = R(\alpha'_{x,y})$ for all $x, y \in \mathbb{Z}$.¹ For an assembly $\alpha' \in \mathcal{A}^S$ such that $R(\alpha') = \alpha$, α' is said to map *cleanly* to $\alpha \in \mathcal{A}^T$ under R^* if for all non empty blocks $\alpha'_{x,y}$, $(x+u, y+v) \in \text{dom } \alpha$ for some $u, v \in \{-1, 0, 1\}$, or if α' has at most one non-empty m -block $\alpha'_{0,0}$. In other words, α' may have tiles on supertile blocks representing empty space in α , but only if that position is adjacent to a tile in α .

A TAS $\mathcal{S} = (S, \sigma_S, \tau_S)$ *simulates* a TAS $\mathcal{T} = (T, \sigma_T, \tau_T)$ at scale $m \in \mathbb{Z}^+$ if there exists an m -block representation $R : B_m^S \rightarrow T$ such that the following hold:

1. Equivalent Production.

- (a) $\{R^*(\alpha') \mid \alpha' \in \mathcal{A}[\mathcal{S}]\} = \mathcal{A}[\mathcal{T}]$.
- (b) For all $\alpha' \in \mathcal{A}[\mathcal{S}]$, α' maps cleanly to $R^*(\alpha')$.

2. Equivalent Dynamics.

- (a) If $\alpha \rightarrow^{\mathcal{T}} \beta$ for some $\alpha, \beta \in \mathcal{A}[\mathcal{T}]$, then for all α' such that $R^*(\alpha') = \alpha$, $\alpha' \rightarrow^{\mathcal{S}} \beta'$ for some $\beta' \in \mathcal{A}[\mathcal{S}]$ with $R^*(\beta') = \beta$.
- (b) If $\alpha' \rightarrow^{\mathcal{S}} \beta'$ for some $\alpha', \beta' \in \mathcal{A}[\mathcal{S}]$, then $R^*(\alpha') \rightarrow^{\mathcal{T}} R^*(\beta')$.

Let REPR denote the set of all supertile representation functions (i.e., m -block supertile representation functions for some $m \in \mathbb{Z}^+$). Let \mathfrak{C} be a class of tile assembly systems, and let U be a tile set.² Note that every element of \mathfrak{C} , REPR, and $\mathcal{A}_{<\infty}^U$ is a finite object, hence can be represented in a suitable format for computation in some formal system such as Turing machines. We say U is *intrinsically universal* for \mathfrak{C} if there are computable functions $\mathcal{R} : \mathfrak{C} \rightarrow \text{REPR}$ and $S : \mathfrak{C} \rightarrow \mathcal{A}_{<\infty}^U$ and $\tau' \in \mathbb{Z}^+$ such that, for each $\mathcal{T} = (T, \sigma, \tau) \in \mathfrak{C}$, there is a constant $m \in \mathbb{N}$ such that, letting $R = \mathcal{R}(\mathcal{T})$, $\sigma_{\mathcal{T}} = S(\mathcal{T})$, and $\mathcal{U}_{\mathcal{T}} = (U, \sigma_{\mathcal{T}}, \tau')$, $\mathcal{U}_{\mathcal{T}}$ simulates \mathcal{T} at scale m and using supertile representation function R . That is, $\mathcal{R}(\mathcal{T})$ outputs a representation function that interprets assemblies of $\mathcal{U}_{\mathcal{T}}$ as assemblies of \mathcal{T} , and $S(\mathcal{T})$ outputs the seed assembly used to program tiles from U to represent the seed assembly of \mathcal{T} .

Our main theorem states that there is a single tile set capable of simulating any tile assembly system.

Theorem 3.1. *There is a tile set U that is intrinsically universal for the class of all tile assembly systems.*

The rest of the paper is devoted to describing the construction of U and justifying its correctness. Throughout this paper, $\mathcal{T} = (T, \sigma, \tau)$ will denote an arbitrary TAS being simulated. Let $g \in \mathbb{Z}^+$ denote the number of different glues in T ; note that $g = O(|T|)$. In our main construction, we achieve scale factor $m = O(g^4 \log g)$; an interesting open question is decreasing this scale factor or proving a nontrivial lower bound on it.

¹Note that R^* is a total function since every assembly of S represents *some* assembly of T ; the other functions such as R and α are partial to allow undefined points to represent empty space.

²TAS's having tile set U are not necessarily elements of \mathfrak{C} , although this will be true in our main theorem since \mathfrak{C} will be the set of all TAS's.

4 High-level description of construction

In this section we sketch an intuitive overview of the construction. The full construction, including detailed figures, is contained in Sections 5–10, in the Technical Appendix. Let $\mathcal{T} = (T, \sigma, \tau)$ be a TAS being simulated by $\mathcal{U} = (U, \sigma_{\mathcal{T}}, 2)$, where U is the universal tile set and $\sigma_{\mathcal{T}}$ is the appropriate seed assembly for U to simulate \mathcal{T} . The seed assembly $\sigma_{\mathcal{T}}$ encodes information about the glues from \mathcal{T} that are on the perimeter of σ , with each exposed tile-side of σ encoded as a “superside” as shown in Figure 2. In particular, glues are simply encoded as binary strings of length $O(\log |T|)$. (Glue strengths are not explicitly encoded since their effect on binding is implicitly accounted for by other parts of the design.) Most importantly, each of these supersides, as well as each superside of all subsequently grown supertiles, encodes information about the entire TAS \mathcal{T} . This information is like the “genome” of the system that is transported to each supertile of the assembly in order to help direct its growth based on the contents of T .

4.1 The fundamental problem of simulating arbitrary tile systems

The basic problem faced by any superside adjacent to an empty supertile is this: the superside must determine what other superside(s) are adjacent to the same empty supertile, what glue(s) are on those sides, whether those glues are part of a tile type $t \in T$ and whether they have enough strength to bind t (and to choose among multiple tile types if more than one match the glues), and if so, the supersides on the remaining sides of t must be constructed (i.e., placed as “output” on empty supersides). This must be done in concert with other supersides that will be attempting the same thing, possibly “unaware” of each others’ presence, and it must be done without prior knowledge of which other supersides will eventually arrive and the order and timing of their arrival.

To illustrate the nontriviality of this problem, consider the following scenario illustrated in Figure 1a. Two supertiles arrive at positions that are north and south of an empty supertile position, with the east and west positions being unoccupied. The south superside has no choice but to attempt to “contact” the north superside, for it may be the case that their glues match that of some tile type $t \in T$, in which case the west and east supersides representing the sides of t must be put in place. But suppose that although there is a north superside, the glue it represents is not shared with the south glue on any tile type in T , or perhaps their combined strength is less than τ . (See Figure 1a.) Intuitively, it seems that to determine this, the north and south supersides must connect, in order to bring their glues together and do a computation/lookup to find that no tile type in T shares them. But once they have connected, the west and east sides of the supertile are now sealed off from each other.

Suppose that at a later time, a superside arrives on the west, and its encoded glue *is* shared with the west glue on some tile type $t \in T$ (with a north glue mismatching that of the supertile already present there; see Figure 1b). This means that t ’s east glue must now be represented by constructing an east output superside; however, this information cannot be communicated from the west side of the supertile because the previous attempt to connect the south and north has created a barrier between east and west.

Note that such problems do not appear in Wang tiling constructions where nondeterminism could be used to simply guess what the other sides are.

This is not the only potential pitfall to be faced, but it illustrates one difficulty of coordinating interaction between multiple supersides in the absence of knowledge about which supersides will



(a) The north side got to talk with the south side, and the consensus is that there is no tile whose north side is 'N' and whose south side is 'S'. Better luck next time!

(b) Uh oh! There seems to be no way to “communicate” from the west side to the east that the ‘WE’ tile should be represented here.

Figure 1: How should supertiles communicate across a gap without “cutting the supertile in two”?

eventually arrive, the order in which they will arrive, and what glues they will represent.³

4.2 Basic protocol

Here we give a high-level overview of our construction.

4.2.1 Frame

Each superside fills in a 4-layer “frame” in the supertile before doing anything else. The purpose of the frame is to give each superside as much information as possible about the other available supersides, to help coordinate their interaction. Each superside attempts to “become an input superside” of the supertile by competing to place a single tile at a particular position on its left end, and another on its right end. It is competing with a (potential) adjacent superside near their common corner; for example, the south superside competes on its left end with the west superside (at the southwest corner) and on its right end with the east superside (at the southeast corner). (The canonical definitions of the “left” and “right” end of each superside are given in Figure 25.) Therefore there are four competitions, one at each corner, and for each corner there is both a winner superside and a loser superside. The “loser” may simply be a superside that is not present and never will be, or it may be a superside that is present but lost the competition because it arrived later. Corner tiles initiate the growth of the first (outermost) layer of the frame, and it is by the initiation from a corner that it lost that a losing side gains the information that there was an adjacent side to compete with. For corners that it won, it cannot know whether an adjacent superside is present and lost, or whether there is simply no adjacent superside.⁴

³These problems would be easier (if cumbersome) to overcome by growing in three dimensions, but achieving a planar construction is nontrivial. Furthermore, since two is the standard number of dimensions in tile assembly, a planar construction is required if we want our result to be most applicable to other (future) results in the abstract tile assembly model, as well as in the wide spectrum of other tile assembly models.

⁴Parallel programmers may be reminded of a similar phenomenon: a thread locking a mutex does not know whether other threads will eventually attempt to access it, but a thread encountering a locked mutex knows for sure that

The first layer of the frame grows from the corners of the superside to its center, at which point the entire superside knows whether it won or lost each corner. The subsequent layers of the frame allow the pattern of wins and losses for each side to be propagated among adjacent sides in a well defined way. The careful design of the frame and the algorithm for passing this information allows the large set of possible scenarios (of all subsets of sides which may be present in all possible orderings) to be condensed into a much smaller set of equivalent classes of scenarios which can then be properly handled by the next portions of the supertile to grow (as necessary) from the frame: “crawlers” which grow along adjacent supersides, and “probes” which grow across the centers of supertile spaces attempting to communicate with opposite supersides (if they exist). It is notable that, as the frame grows, it propagates all information from the superside of the adjacent supertile (which is acting as an output side for that supertile that initiates the formation of this potential input side for a new supertile), including the encoding of the glue on that superside and the encoding of the system \mathcal{T} (i.e. the “genome”), into the interior of the new supertile. Figure 14 shows all possible combinations of input sides, together with all possible win-lose scenarios at each corner (rotational symmetries omitted). Section 7 details the algorithm used to grow the frame to achieve this gathering of information.

4.2.2 Crawlers and lookup tables

Once the frame has formed, information about the glues represented by the supersides, the simulated tile system \mathcal{T} , and the win/loss status of each side (possibly along with information about additional sides) is presented on the inside of the frame. At this point, pairs of adjacent edges (i.e. those sharing a corner) may initiate the growth of a “crawler” component. (The determination of whether or not a particular pair will do so is discussed later.) At a high level, when a crawler is initiated it contains the encoding of the glue for one superside, and as it forms it grows across the adjacent superside, gathering the information of the glue on that side as it grows across it. Next, it grows across the encoding of a “tile lookup table” which is an encoding of the tile set T that allows it to determine if the glues that it has collected so far match a tile in $t \in T$ and have sufficient strength for t to bind. If so, then this supertile should simulate t and the crawler’s job becomes to grow around the remaining sides of the supertile and to create output supersides for those sides which aren’t already occupied by input supersides. See Figure 5 for an example of this scenario.

Of course, this is an ideal scenario; what could go wrong? Perhaps the glues do not match a tile type. Perhaps another superside arrived while we were attempting to output on that side. Perhaps there are only supersides on the south and north, and they must reach across the gap of the empty supertile between them to cooperate and place east and west output supersides. More generally, a crawler crawls around a supertile “collecting” input glues. It starts in an **unfilled** state until a tile lookup reveals that it has collected glues with sufficient strength to place an output tile type; at this point the crawler enters a **full** state. However, it has not yet committed to creating an output tile type because there may be other crawlers present that are also **full**. Some symmetry-breaking is used to determine when a **full** crawler changes to an **output** state and takes responsibility for determining the output tile type and placing output supersides. Our main goal in justifying the correctness of the construction is proving that if subsets of supersides represent glues that are sufficient to bind a tile, then eventually exactly one crawler will enter the **output** state and decide the output tile type t (or two crawlers in the special case where they originate from meeting probes

another thread is currently accessing it.

and are guaranteed to make the same output decision).

4.2.3 More general crawler protocol

The more general protocol followed by crawlers is this. Whenever two supersides “connect”, at a corner as in Figure 5, or by reaching across the gap as in Figure 6, they (sometimes, depending on information supplied by the frame) initiate a crawler that first combines their glues and does a lookup to see if a tile matches these glues. Crawlers always move counterclockwise around a supertile. Figure 14 shows green arrows to indicate where crawlers are initiated. The general rule is: *Initiate a crawler when two sides meet, unless we have enough information from the frame to see that another crawler will be on its way from another corner.* Note that sometimes two crawlers are initiated because the “later” crawler (the crawler in the more counterclockwise direction) does not “know” (based on only its two adjacent sides) about the first crawler. If the lookup is successful, the crawler becomes **full** and will attempt to place output supersides if there are potentially empty supersides. On each potential output superside, the crawler first “tests” to see if an output side is already present, only outputting if necessary. If the lookup is unsuccessful (i.e., the glues available to the crawler were not sufficient to bind a tile), the **unfilled** crawler crawls to the edge of the supertile to wait for a potential new input superside to arrive. If this superside ever does arrive, it will initiate its own crawler that will combine the information from the first crawler (and its two glues), to see if all *three* glues are sufficient by performing a new lookup. This new crawler will follow the same protocol.

4.2.4 Multiple crawlers

As previously mentioned (and as shown in the scenarios of Figure 14 where there is more than one green arrow), there are some situations in which two crawlers may be initiated and begin growth. In such a situation, a crawler c_1 may arrive at a side to find that another crawler c_2 has already begun growth from there; if so, c_1 crawls over the “back” of c_2 (see Figure 7) to see if c_2 became **full** (i.e., had a successful table lookup). If c_2 is **full**, c_1 stops, allowing c_2 to take responsibility for outputting, as in Figure 7. Otherwise, c_1 does its own lookup using all glues (whatever glues that c_1 has already collected before encountering c_2 , plus the new glue on the side that initiated the growth of c_2). It is possible for c_1 to receive this additional information from c_2 because crawlers pass all collected and computed information up through themselves. This is necessary for supporting such “piggybacking” crawlers, as well as making the necessary information available when it becomes time to create output supersides. In this case c_1 may overtake c_2 to place output, as in Figure 11.

In cases where all four supersides are present, although the **output** crawler does not need to deposit output supersides, it must still decide on an output tile type so that the representation function can uniquely decode which tile type is represented by the supertile. In this case, it may be the case that two crawlers exist but one of them does not run into the other. However, we still require symmetry-breaking so that only one of them changes to the **output** state. In this case, once a crawler has encountered the fourth superside (which happens after it has traversed the full length of two supersides; see Figure 12), it has complete information (gathered from the frame) about the win-loss configuration of Figure 14 and therefore knows whether another crawler was independently initiated. In this case, a precedence ranking on corners that initiate crawlers (NW > SW > SE > NE) is used to determine whether to transition to the **output** state or to simply die (in effect, letting the other, higher-precedence crawler become the unique **output** crawler).

4.2.5 Probes

“Probes” are used for communication across a gap between two potential input supersides on opposite sides (i.e. north and south or east and west) when necessary. Suppose the south superside needs to communicate with the north superside. Recall that the south superside’s frame either won or lost on each end of the superside. If either end lost, then this means there is a supertile adjacent to both south and north (west if south lost in the southwest corner, and east if south lost in the southeast corner). Therefore there is no need for probes, since crawlers will eventually connect the south glue with the north glue. Only if the south superside is “win-win” does it send probes to potentially connect with the north superside. Since all supersides follow this rule, this ensures that at most two sides ever grow probes, and if so, then they are opposite sides (since adjacent sides cannot both be win-win). This ensures that orthogonal probes cannot grow and interfere with each other.

Probes are specifically designed so that they “close the gap” (connect two sides of the supertile) if and only if their supersides represent glues that map to a tile type $t \in T$ and have sufficient strength for t to bind. The probes grow from a region on the superside known as the “probe region”. Each glue in T has its own unique subregion in the probe region (see Figure 24). The supersides do not grow probes symmetrically: north and east grow probes in one way, and west and south grow them in a complementary way. Suppose the glue on the north is n and the glue on the south is s . The north superside will grow a probe in the subregion associated with n . For every glue g that has the property that there is some tile type $t \in T$ with g on the north, s on the south, and g and s have combined strength at least τ , the south superside will grow a probe in the subregion associated with that g . If no tile type matches glue s on the south and n on the north (or if n and s have insufficient combined strength), but both north and south probes form, they will be guaranteed to leave sufficiently wide gaps for crawlers, which may be initiated and arrive later, to make their way around and between the probes (see Figure 10). This is because each probe subregion is at least $\Omega(|T|)$ tiles from its adjacent probe subregions, but crawlers are only $O(\log |T|)$ tiles wide. If the probes do meet in the middle of the supertile (indicating that a tile type $t \in T$ matches the north/south glues and has sufficient strength to bind), they initiate their own crawlers (see Figure 6) which can place the output supersides representing the west and east sides of t .

4.2.6 Simulation of nondeterministic tile systems

In each of these cases, if the simulated system \mathcal{T} is nondeterministic, there may be more than one tile type that matches a given set of input glues. To handle this scenario, a “random number” is produced through nondeterministic attachment of tile types to a special “random number selector component” and used as an index to select one of the possible tile types. (A similar mechanism was used in [22].) It is crucial that if two probes cut off two sides of a supertile from each other, each side’s crawlers must use the same random number to select the tile type to output, or else they may choose differently and place output glues that are not consistent with any single tile type in T . This is why probes generate a random number and advertise it to each side of the probe. However, if probes do not meet, then eventually a single crawler will be responsible for choosing an output tile type, so it is sufficient for the crawler to generate a random number just before it begins a tile lookup.

Sections 5–10 describe the details of the full construction.

5 Supertile layout

The layout of completed supertiles is heavily influenced by the number of supersides that need to cooperate to produce output supersides. First we describe the layout and encoding of supertile sides (or supersides). We then describe supertile layout, beginning with the simplest case: one-sided binding. The goal in this section is to convey the high-level idea about how frames, probes and crawlers work in a way to share information appropriately so that the required information can move around the simulated supertile. The low-level details about *how* frames, probes and crawlers interact is given in Sections 7, 8 and 9. Section 6 builds on these descriptions in order to argue the correctness of our construction.

5.1 Superside layout




Figure 2: Superside layout for a south superside. The superside encodes a *glue* on the left and right, and encodes the entire simulated tile assembly system in both copies of its *tile lookup table*. The *probe region* is where probes grow from, and the *probe table* is used to compute which probes should grow from the probe region. The *blank region* is of the same length as the probe region and probe table, and is simply used to maintain symmetry.

A side of a supertile is called a superside and consists of a linear sequence of $O(g^4 \log g)$ tiles that encode the simulated tile assembly system \mathcal{T} , as well one side of an encoded tile t . This superside information is embedded in the frame, which in turn is described in Section 7.

Figure 2 shows the layout for a south superside (i.e., a superside south of an empty supertile, which is the *north* output superside of the supertile to the south of the empty supertile). Starting from the west, we have a 4-tile wide frame region. Next, the glue area encodes *two* copies (side-by-side, in a linear sequence) of t 's south glue. Then, there is a tile-lookup table encoded as a sequence of $O(g^4 \log g)$ tiles. Next, we have a blank region, probe region and probe lookup table. The blank region is simply used to preserve symmetry in the construction, the probe region is where probes grow and the probe table is used to compute the positions of probes in yet-to-be-produced output sides (see Section 8). Finally we have another copy of the tile lookup table, two copies of the glue and strength, and another width-4 frame region. This redundancy (two copies of each region) is used in the construction to enable tile lookups on both ends of the probe region and to facilitate copying of encoded superside information to new output supersides (which is initiated from the right end of an input superside).

Finally, if a superside represents a strength- τ glue, then this bit is encoded into each tile type in the superside. This is useful for implementing some rules regarding when certain crawlers grow or change state, etc.

5.2 One-sided binding

Here we describe the simulation of the binding to a tile type t that binds on its south side with a strength- τ bond. We assume there are no other input sides in this scenario, which corresponds to Figure 14(1.1). The order of growth is described at a high level in Figure 3. Growth begins with




Figure 3: One-sided binding by a strength- τ south superside. An (orange) crawler is initiated (from the south probe) and outputs to the east, north and west. In this and similar figures, dark gray represents the frame, and dotted lines represent boundaries between supertiles. Therefore, supersides on the “inside” of the dotted lines are output supersides, and supersides on the “outside” of the dotted lines are input supersides (and are always adjacent to a frame on the inside of the dotted lines).

the frame (Figure 15), which results in a “win-win” (WW) for this superside since we are assuming there are no competing east/west supersides. As the south frame completes its final (fourth) row, a single *strength- τ probe* grows from a specific location in the probe region. The probe, shown in blue in Figure 3, is of width $O(\log g)$, and grows to the center of the supertile (supersides are of odd length) using a binary counter. The purpose of this probe is to compete with a potential opposite strength- τ superside that also wishes to place output supersides (in the scenario we are considering there is no such opposite superside, but the south superside does not “know” this so must grow a probe anyway). The probe claims the center tile position by placing a blue tile there, which initiates an orange crawler. The crawler picks up the south glue information g_S and a random number r from the side of the probe (see Section 8), and provides these as input to a tile lookup table (see

Section 10), which decides if the supertile should produce output sides. The random number r is used to simulate nondeterminism (i.e. if t is one of a set S of valid tiles that could be simulated). In the tile lookup table, the south glue g_S is used to find the set of valid tiles S , and r is used to choose t from S . Crawlers can be in one of 4 states `{unfilled, full, output, dead}`, described in detail in Section 9.3. All crawlers start in state `unfilled`.

After the tile-lookup the orange crawler changes to state `full` (since a single strength- τ glue is always sufficient to bind a tile), and its goal is to produce three output supersides. On the one hand, in the absence of other input supersides, the orange crawler succeeds at its goal and triggers the growth pattern depicted in light blue in Figure 3. This growth pattern serves to output properly encoded supersides to all other sides, and includes computing the correct output glues, strengths and probe regions, as well as copying the remaining superside information. The computation of correct output glues, strengths and probe regions is discussed in Sections 9 and 8. The crawler changes from `full` to `output` upon detecting that the east superside is empty (for details see Section 9.4). On the other hand, if there are extra *non-contributing input supersides* (e.g. we are simulating mismatches) the outputting crawler is designed to detect this and not output on those supersides. This behavior is described next.

5.2.1 One-sided binding with non-contributing input sides

We now deal with one-sided binding in the presence of *non-contributing input supersides*.⁵ The first scenario we consider, shown in Figure 4, is one-sided binding where two *opposite* strength- τ supersides compete to see which should place output supersides. The two input supersides are opposite, and there are no adjacent supersides, therefore both supersides have win-win frames (Figure 14(2.1)) and so they both grow probes. Within a supertile, probes meet (and block the path of crawlers) only if we are simulating either (a) two-sided opposite binding for matching opposite sides or (b) two opposite strength- τ sides. In all other cases probes do not meet and crawlers can crawl over them (see Section 8). In Figure 4, the two strength- τ probes meet, and actually compete to see which will fill out the supertile: south wins this competition by being the first to claim the center tile position. When two strength- τ probes meet, they trigger the growth of two crawlers: one east crawler and one west crawler (both begin in state `unfilled`). Outputting to the east side proceeds as before. Output to the north by the east-side crawler is prevented by the east crawler carrying out a test in the north-east corner and detecting the presence of the north frame (see Figure 27(b) for details). This crawler also detects that the competing input superside to the north is of strength τ (each frame tile encodes a single bit that flags whether or not a superside is of strength τ). This information tells the east crawler to stop growing (it now knows that there will be a north probe blocking the path to the west). Outputting to the west is handled by the west crawler. As shown in Figure 4, the west crawler copies the south glue information g_S and a random number r from the south (blue) probe. The west crawler uses this information to do a tile lookup along the north superside. Since both the east and west crawlers use the same inputs (g_S, r) , after the tile lookup they both encode the same output tile type t .

In one-sided binding (i.e. where a single input superside of strength- τ determines the encoded tile type), for all other scenarios of non-contributing input sides, opposite probes either do not grow at all, or if they grow they do not meet. In this case a single outputting crawler moves counter-clockwise around the supertile. When the outputting crawler arrives at a corner it detects whether

⁵*Non-contributing input supersides* are input supersides that are not used to simulate tile binding. Mismatching supersides are one such an example.




Figure 4: One-sided binding by a strength- τ south (S) superside, in the presence of a non-contributing strength- τ north superside. The east (orange) crawler is initiated from the south probe and outputs to the east. The west (also in orange) crawler is initiated from cooperation between the north (red) and south (blue) probes and outputs to the west.

or not there is a non-contributing input side, see Section 9.4 for details. If not, the outputting crawler produces an output side using the growth pattern shown in Figure 3. Otherwise, if there is a non-contributing input side, the outputting crawler detects this and simply crawls across the non-contributing input side.

5.3 Two-sided binding

We now describe the simulation of two-sided binding between adjacent supersides. Figure 5 illustrates this for south and west cooperating adjacent supersides (the south is win-win, corresponding to Figure 14(2.3)). The orange crawler is initiated by cooperation between the south and west frames, in the south-west corner. This crawler picks up the west and south glue information (g_W, g_S) while in state `unfilled`, and grows towards the first (leftmost) tile lookup table on the




Figure 5: Two-sided binding for two adjacent sides on west and south. An (orange) crawler is initiated from the south-west corner and outputs to the east and north.

south superside. Upon entering the tile lookup table a random number r is generated, then the tile lookup is performed using (g_W, g_S, r) as input. The crawler transitions to state **full** upon completing the tile lookup. In the south-east corner, the **full** crawler detects the absence of an east superside, transitions to state **output**, and produces the east and north output supersides using the growth pattern shown in Figure 5.

In Figure 5, as in many others, we show that there are probes that do not meet their complementary probes, to illustrate that the probes grow because the south superside, being win-win, does not “know” that eventually the west superside will arrive to initiate a crawler. However, since the north superside does not arrive, the west superside must be win on its left (north) end (i.e., the frame configuration matches in Figure 14(2.3)), so any north side cannot be win-win and therefore cannot grow any probes. Therefore the crawler (having knowledge of the frame configuration on the west superside) safely proceeds across the probe region since it knows that no north probes will be present.




Figure 6: Two-sided binding for two opposite sides on north and south. North and south grow probes (red and blue), which meet and initiate two (orange) crawlers. These crawlers output to the east and west respectively.

We next describe the simulation of two-sided binding of a tile type t that binds with two opposite sides. Figure 6 shows an example of this cooperation “across the gap”. Both north and south are win-win, corresponding to Figure 14(2.1). Here, probes are used to establish cooperation between the two opposite sides. Since we are dealing with two matching opposite sides, the north and south probes meet (see Observation 8.1). In Figure 6 the north and south probes cooperate to initiate both an east and a west crawler, both shown in orange and both are initiated in state **unfilled**. Each crawler picks up both the north and south glues (g_N, g_S) from the two probes, as well as a random number r from the south (blue) probe only. The crawlers crawl counterclockwise across the south and north supersides to their respective tile lookup tables, where they both supply the same input (g_N, g_S, r), and hence both produce the (same) simulated tile type t as output. The crawlers transition to state **full**, detect the absence of the east and west supersides, transition to state **output**, and then produce the east and west output supersides using the growth patterns

shown in Figure 6. So although the probe blocked the path across the center of the supertile, we still managed to share all the information required for the simulation of consistent binding of a single tile type.

Note that this situation of across the gap cooperative binding and the previously described situations of one-sided binding in which the strength- τ probes meet are the only two situations in which two physically separate crawlers are both in state `output`. This is safe to do since each crawler is initiated from the same probe and uses the same information to determine an output tile type (hence make the same decision). In all other situations described subsequently, at most a single crawler is ever in state `output`, to ensure that a consistent decision is made regarding the output tile type (if one exists).

In the absence of extra non-contributing input sides, all forms of two-sided binding are simulated using the above two techniques, or their rotations.

5.3.1 Two-sided binding with non-contributing input sides

Up to rotation, there are four cases for two-sided binding in the presence of additional non-contributing input sides.

Case (i): Two-sided adjacent binding with one adjacent non-contributing input supertile (three supertiles total). Figure 7 gives an example where two adjacent sides (west and south) cooperate to simulate a tile binding, in the presence of a north non-contributing input side. The win-lose configuration is given in Figure 14(3.2). As per Figure 14(3.2), the orange crawler in Figure 7 is initiated in the south-west corner in state `unfilled`, transitions to state `full` (encodes a simulated tile type $t \in T$) after the first tile lookup, and proceeds to output on the east (as happened in Figure 5). In the meantime a north supertile arrives. The outputting orange crawler detects the presence of the north supertile at the north-east corner, and then stops growth as it knows there are no further supertiles to be output. Although the simulation of the binding of tile type t is essentially complete, at the north-west corner there is insufficient information to know this fact. So a green crawler gets initiated at the north-east corner by cooperation between north and west. The green crawler does its first tile lookup table on the west, and in the absence of knowledge about the orange crawler, green makes a decision about whether it should remain in state `unfilled` or transition to `full`. However, after green goes through its first tile lookup table on the south it detects the state (`full`) of the orange crawler directly to its south. Green now knows that orange has already claimed responsibility for encoding the simulated tile type t , therefore green enters state `dead` and stops growth.

For the same win/lose scenario for these three frames (north, south, west), another possible order of growth is where the green crawler arrives at the south-west corner before the orange crawler is initiated. In this case there is no orange crawler, the green crawler transitions to state `full` upon collecting the glue from the south supertile, then state `output` upon successfully reaching the east supertile and proceeds to output the east supertile.

For the win/lose configuration in Figure 14(3.1) (where one of the sides is a non-contributing input side) the orange and green crawlers act similarly as in Figure 14(3.2) just described, the only difference being the presence/absence of non-blocking probes on the various sides. For the win/lose configurations in Figures 14(3.3) and 14(3.4), the situation is somewhat simpler to describe as only one crawler is created; we omit the details.

Case (ii): Two-sided adjacent binding with two adjacent non-contributing input sides (four supertiles total). This is similar to Case (i), the only difference being that the `full` crawler detects




Figure 7: Two-sided binding for two adjacent sides on west and south, with an additional non-contributing input side to the north. An (orange) crawler is initiated from the south-west corner and outputs to the east. West and north also initiate cooperation by producing a (green) crawler, which halts when it detects that orange is already outputting.

the presence of both non-contributing input sides and stops growing, and thus does not output any supersides. However, immediately before stopping growth the **full** crawler transitions to state **output** and grows a single column of $O(\log g)$ tiles that encode the simulated tile type $t \in T$ (for details, see Figure 27(c)). In order to determine the encoded tile type from T , the representation function R checks these $O(\log g)$ tiles.

Case (iii): Two sided opposite binding, with one adjacent non-contributing input side (three supersides total). Figure 8 shows a simulation of Figure 14(3.4), where the west is a non-contributing input side. Probe locations within the probe region are such that probes meet if we are simulating two matching opposite sides (Observation 8.1), as is the case here. In Figure 8 the probes meet and cooperate (as was described above for Figure 6). In the meantime a green crawler is initiated from the north-west corner. After green does its second tile lookup on the south side it halts as by




Figure 8: Two-sided binding for two opposite win-win sides on north and south, with an additional non-contributing input side to the west. North and south grow probes (red and blue), which meet and initiate (orange) crawlers. The east crawler outputs a superside to the east, while the west crawler is prevented from outputting by the presence of a non-contributing input superside to the west. West and north initiate cooperation by producing the green crawler, which halts after it knows that the center of the supertile is blocked by probes.

this time it has sufficient information (i.e. north glue, south glue, and the win/lose configuration of the three frames) to know that north and south match, and will grow probes that meet and block the path to the east. The orange crawler outputs to the west.

For the same win/lose scenario for these three frames (north, south, west), another possible order of growth is where the west orange crawler arrives at the north-west corner before the green crawler is initiated. In this case there is no green crawler. The west orange crawler detects the presence of the frame of the west superside, transitions to state **output**, and stops growth.

If the non-contributing input side is instead on the west, a rotated version of the previous simulation is used.

Case (iv): Two sided opposite binding, with two adjacent non-contributing input sides (four supersides total). This win/lose scenario is shown in Figure 14(4.1). This is similar to Case (iii), the only differences being (a) that up to four crawlers are created (two corresponding to the green arrows in Figure 14(4.1) and two originating from the meeting probes) and (b) the *two* outputting crawlers (those coming from the probes) detect the presence of both non-contributing (lose-lose) input sides, transition to state *output*, and stop growth.

For cases (i)-(iv) above we have omitted description of their rotations, and we did not explicitly list all win/lose frame configurations. However, the same style of argument applies for all these scenarios.

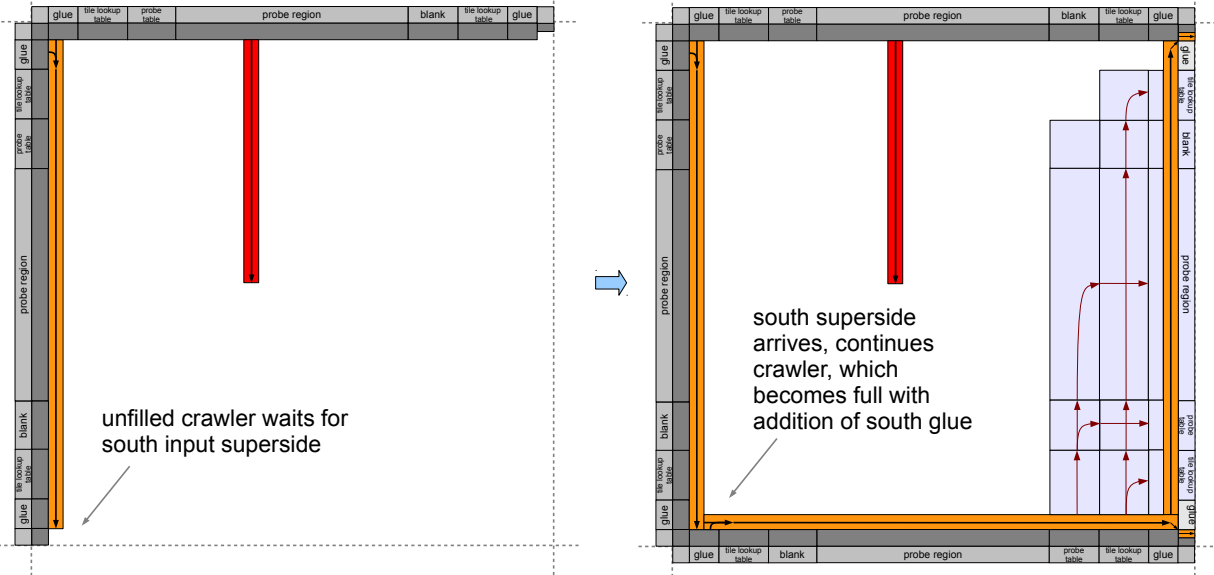


Figure 9: Three-sided binding with input supersides on the north, west and south. An (orange) crawler is initiated from the north-west corner but has insufficient glue information to output so waits at the south-west corner. When the south superside arrives, the crawler continues, which after a table lookup becomes full, so it has sufficient glue information to simulate the placement of a tile type and thus outputs on the east.

5.4 Three-sided binding

Figure 9 shows a three-sided binding scenario with win/lose configuration given in Figure 14(3.1). An (orange) crawler is initiated from the north-west corner but has insufficient glue information to output, so it waits at the south-west corner in state *unfilled*. Eventually the south superside appears. The orange crawler cooperates with the south, gathers the south glue, does a tile lookup and transitions to state *full*. The crawler outputs to the east. The purpose of this example is to illustrate the asynchronous nature of superside cooperation: where appropriate, crawlers must wait for potential input supersides.

Figure 10 illustrates north, west and south cooperating supersides (north and south are both win-win, corresponding to Figure 14(3.4). The south-west corner does not initiate a crawler: as can be seen in Figure 14(3.4) the frames at this corner know that there will be a crawler coming




Figure 10: Three-sided binding with input supersides on the north, west and south. A single (orange) crawler is initiated from the north-west corner and outputs to the east. North and south grow probes, that do not meet.

from north-west corner. The north-west corner initiates an (orange) crawler that travels counter-clockwise gathering glues from the three input supersides. After completing its second tile lookup (on the left of the south superside) the orange crawler transitions to state *full*. Since north and south are win-win they both grow probes. However in this case the north and south supersides have insufficient strength to match, so, by Observation 8.1, their probes do not meet. The orange crawler crawls along the south superside, outputs on the east, and then halts.

Another very similar case, shown in Figure 14(3.3) is where north is win-win and west and south are both lose-win. This case is handled exactly as in the previous one, the only difference being that there is no south probe present.

Another class of cases is where we have three-sided binding, with up to *two crawlers*. This occurs with frame configurations given by Figures 14(3.1) and 14(3.2) and an example is illustrated in Figure 11. In Figure 11 the south frame is win-win and grows probes (i.e. Figure 14(3.2)). An




Figure 11: Three-sided binding with input supersides on the north, west and south. An (orange) crawler is initiated from the south-west corner but has insufficient glue information to output so waits at the south-east corner. A (green) crawler is initiated from the north-west corner, and has sufficient glue information to simulate the placement of a tile type and thus outputs on the east.

(orange) crawler is initiated in the south-west corner, however after going through the first tile lookup (on the left side of the south superside) it has insufficient strength from the south and west glues to simulate binding of a tile type. The orange crawler (in state **unfilled**) waits at the south-east corner (waiting to cooperate with a possible east superside). In the meantime a (green) crawler is initiated in the north-west corner, by the time it does a tile lookup on the south (using glue information from north, west and south) it transitions to state **full** as it has sufficient glues and strength to simulate binding of a tile type. The green crawler crawls over the orange crawler and outputs on the east side (see Figure 29(b) for details).

For the same win/lose scenario for these three input sides, another possible order of growth is where the green crawler arrives to the south-west corner before the orange crawler is initiated. In this case the green crawler simply transitions to **full**, after doing a tile lookup on the south, and

proceeds to output on the east.

The same arguments work for all rotations of these 3-sided binding cases.

5.4.1 Three-sided binding with a non-contributing input side

Three-sided binding in the presence of an additional non-contributing input side is conceptually identical to the case of four-sided binding, the only difference being the number of glues that contribute in a table lookup, and the fact that one crawler will try to output a superside (and fail). Therefore we omit explicit discussion of these cases, since four-sided binding is discussed in Section 5.5.

5.5 Four-sided binding

In terms of win/lose configurations, there are six 4-sided cases given in Figure 14(4.1)–(4.6). When simulating four-sided binding, it is not necessary to produce output supersides, so crawlers use the state `output` solely to define the simulated tile type (i.e. to communicate to the representation function R).

One of the simplest cases to consider is shown in Figure 12. Here the win/lose configuration is either of Figures 14(4.5) or 14(4.6). No probes are created, and a single crawler gathers all four glues, at which point it transitions to state `output`, and lays down a row of $O(\log g)$ tiles that describe the simulated tile type t . If there is no matching tile type, it transitions to state `dead`. The win/lose configuration of Figure 14(4.3) is handled in a very similar way, there is a (useless) probe and a single crawler that actually does the work.

The remaining three win/lose configurations are given in Figure 14(4.1), (4.2) and (4.4). Up to two crawlers are created. For four-sided binding, our construction is designed so that (a) if there is exactly one crawler, it simulates binding as required, and (b) if there are two crawlers then at some point both crawlers become aware of each others' presence, and one transitions to state `dead` while the other transitions to state `output`. The latter crawler encodes the output tile type. Rather than giving an exhaustive analysis at this point, we describe an example in Figure 13, which has win/lose configuration shown in Figure 14(4.1). Since this is four-sided binding, the opposite probes do not meet (Observation 8.1) in Figure 13. Two crawlers are initiated, as shown in Figure 14(4.1). The orange crawler is unfilled by the time it has gathered 3 glues, when it reaches the south-east corner it detects the green crawler (i.e. orange fails to claim point of competition POC1), reads the east glue, and does a tile lookup. There is a crawler precedence (based on the initiation corner: $NW > SW > SE > NE$, see Section 9.3), so since orange has highest precedence it transitions to state `full`, then immediately transitions to state `output`. When the other (green) crawler detects the presence of the orange crawler (in the north-west corner) it stops growth as it has lower precedence.

The general rule here is that if a crawler learns the win/lose configuration of all 4 supersides (including the fact that there *are* four input supersides), and from that learns that there is/will be a second crawler with higher precedence, then it stops growth. If it learns that *it* is the crawler with highest precedence, then it outputs a tile type from T .

6 Correctness of construction

As a reminder of notation, let g be the number of glues in the tile set T of the simulated tile assembly system \mathcal{T} .




Figure 12: The win-loss configuration forms a cycle (Figure 14(4.5) or (4.6)). Since all four sides of the frame have complete knowledge of the win-loss configuration, we can break symmetry and begin a crawler only from the NW corner. Its only job is to gather all glues (which it has after visiting the east side), and transition to `output` if there is a tile type matching these glues.

Our argument for correctness of an assembly sequence of \mathcal{U} simulating an assembly sequence of \mathcal{T} is based on showing that supertiles (a) decode input supersides to produce correct output supersides, and (b) correctly encode a tile type from T . Then, via the representation function R , this implies that an entire assembly sequence from \mathcal{T} is simulated correctly. The argument is based on an analysis of crawler states and on the growth patterns within a supertile.

Section 9.3 defines the set of crawler states: $\{\text{unfilled}, \text{full}, \text{output}, \text{dead}\}$, and describes how crawlers transition between states. In particular, the crawler state diagram in Figure 26 is used extensively in this section. A supertile that encodes a tile type t , does so explicitly in any column of any crawler that is in state `output`. Therefore the representation function R decodes a supertile by checking $O(\log g)$ tiles. More precisely, R checks a constant number, independent of the simulated tile assembly system \mathcal{T} , of regions within the supertile, each of size $O(\log g)$, in order




Figure 13: Two crawlers are initiated when four sides are present. After crawling for two sides, each learns that there are four sides, as well as the entire supertile’s win-loss configuration. At this point, the crawler with higher precedence (orange) continues, while the lower-precedence crawler (green) halts. This is because if the four sides have sufficient strength to bind, only one crawler should transition to `output`; the orange crawler does this after reading the lookup table on the east side, transitioning to `full` and then immediately to `output` (since it knows there are no empty sides necessary to output to).

to determine t .⁶

We consider three cases that represent three different types of supertile formation, and argue that in each case the correct tile type t is simulated, otherwise no tile type is simulated.

Case 1. Claim: If a supertile simulates binding across the gap⁷, there are exactly two crawlers in the supertile that transition to the state `output` and both encode the same tile type $t \in T$.

Proof sketch: By Observation 8.1, simulating binding “across the gap” involves either (1a) two

⁶Thus R is a very simple representation function in the sense that it has very low computational complexity: using reasonable (standard) encodings of tiles as bit strings, it is contained in FAC^0 (the function analog of the constant depth unbounded-fanin Boolean circuit complexity class AC^0).

⁷The term *binding across the gap* refers to binding that is initiated by two probes that meet, see Observation 8.1.

probes, each representing strength $< \tau$ glues, meeting to cooperate, or (1b) two strength- τ probes that meet (one of which wins the competition between them). In the case of (1a), immediately after the probes meet, two crawlers in state `unfilled` are initiated that both go on to transition to state `full` after their first tile table lookup (see Figures 6 and 8). Since both crawlers used the same input glues and random number (gathered from the south or west probe) in their tile lookup, when they enter the state `full` both encode the same tile type t . If there are no other (i.e. non-contributing) input supersides, these crawlers transition to state `output` and output the two relevant encoded output sides of tile t (e.g. Figure 6). If there is one or two non-contributing input sides, then when a `full` crawler detects this (by failing to claim a specific “point of competition”, called POC2, because of the presence of a frame, as shown in Figure 27(c)) it transitions to state `output`, produces a single column of $O(\log g)$ tiles that encode $t \in T$ and immediately stops growing (see Figures 8 and 27(c) for details).

The case of (1b) is analogous: the two strength- τ probes connect in the middle. The one that wins the competition initiates a crawler independently of whether the losing probe is present. The arrival of the second probe initiates a second crawler, which in analogy to case (1a) uses the same glue information and random number as the first crawler, guaranteeing that it enters state `full`, and then `output` encoding the same output tile type as the first crawler.

For both (1a) and (1b), besides probe-initiated crawlers, the only other crawlers that get initiated are due to the arrival of third or fourth input supersides (whose frames are necessarily lose-lose). Each such superside initiates at most one such new crawler. As usual, the crawler crawls from the left end to the right end of the lose-lose superside. It meets the win-win superside, then does a tile lookup on the win-win superside (the green crawler in Figure 8 is an example). Immediately after the tile lookup the crawler enters state `dead`, as it has enough glue information to decide that both probes meet up ahead (and that there are other crawlers whose job it is to encode the tile type). This state change is defined in Figure 26, and illustrated in Figure 8.

Case 2. Claim: If the supertile simulates binding that is not “across the gap”, there is exactly one crawler in the supertile that transitions to state `output` (and thus the supertile uniquely encodes some tile type $t \in T$).

Proof sketch: As defined in Figure 26, transitioning to the state `output` can happen in one of two ways (2a) or (2b), both from state `full`:

(2a) A crawler c_1^{full} in state `full` claims POC2, i.e., it succeeds in constructing an output superside. Figure 27(a) shows a `full` crawler claiming POC2. There are three subcases to consider, (2a.1), (2a.2) and (2a.3):

(2a.1) The crawler c_1^{full} came from the previous corner in the clockwise direction. From Figure 14, at most one other crawler c_2 can be initiated. If c_2 exists, it came from the corner previous to the corner that initiated c_1^{full} (because one superside is empty). If c_2 is initiated it will eventually piggyback on the `full` crawler c_1^{full} and then stop growth in state `dead` and so c_2 never gets to state `output` (Figure 26 defines this behavior of c_2 , and an example is shown in Figure 7).

(2a.2) The crawler c_1^{full} came from two corners previous: again there will be no other crawlers that enter state `full`, because either c_1^{full} piggybacks on an unfilled crawler c_2 and thus c_1^{full} completely covers c_2 (since c_1^{full} claims POC2, see Figures 29(b) and 11), or else there are no other crawlers at all. (If c_1^{full} piggybacks on a full crawler then it will enter state `dead` before reaching the empty superside, so it could not have claimed POC2.)

(2a.3) The crawler came from a strength- τ probe. The crawler enters state `output` upon claiming POC2, and proceeds to output. It also outputs supersides for any other supersides where it can

claim POC2. Any other crawler that is initiated comes from a previous corner and stops growth when it meets the strength- τ superside (Figure 26).

(2b) There are four supersides, and there is no other crawler with higher precedence that will eventually transition to **output**:

As Figure 14 shows, the frame initiates at most two crawlers. **One crawler:** If a crawler gathers four glues directly from four frames (claiming POC1 at each frame), then it knows that there are, and will be, no other crawlers. If it is **full** at this point it immediately transitions to state **output**, grows a single column of $O(\log g)$ tiles encoding the tile type t , and then stops growth. An example is shown in Figure 12. If it is not **full** at this point, it simply enters state **dead** and stops growth (i.e. the supertile encodes no tile from T).

Two Crawlers: If there are four input sides and two crawlers (see Figure 14), there is a simple algorithm that causes one crawler to enter the **dead** state, and the other crawler (after collecting all 4 glues) to enter the **output** state. The basic rule is: if a crawler c_1 is in state **full** and learns (from reading win/lose frame information) that it is the highest precedence crawler, or that a higher precedence crawler would have piggybacked on c_1 when c_1 was full, then it immediately transitions to **output**. Note, that the necessary information not contained in the frame (location of the lookup table where c_1 transitioned to **full**) can be stored in a constant number of bits by c_1 . This rule is defined in Figure 26.

Case 3. Claim: If the supertile should not simulate any tile type, then there are no crawlers in state **output** within that supertile location in a producible terminal assembly.

Proof sketch: We argue that no crawler transitions to state **full**, which implies that no crawler transitions to state **output**, for all cases where we should not simulate placement of a tile type from T . If there are no input supersides, then no growth occurs within the supertile. If there is one input superside that has strength $< \tau$, no crawlers are initiated (Section 9.1). If there are exactly 2 opposite supersides that do not represent glues (of sufficient strength) on any tile type $t \in T$, then their probes do not meet (see Observation 8.1), and thus no crawlers are initiated. If there are exactly 2 adjacent input supersides, that do not represent any tile type $t \in T$, or have insufficient strength for binding, then an **unfilled** crawler c is initiated at their corner. Crawler c goes through a tile lookup, remains in state **unfilled**, crawls to the right end of a superside, and waits (forever) in state **unfilled**, for a third input superside that never arrives (an example of an **unfilled** waiting crawler is shown on the left side of Figure 9). If there are exactly 3 adjacent input supersides that do not represent glues (of sufficient strength) on any tile type $t \in T$, then first, we know that no probes meet (and so no crawlers are generated from probes). We also know that either one or two crawlers are initiated from the two corners, which remain in state **unfilled**, forever waiting for an input superside that never arrives (Figure 28(b)). If there are exactly 4 input supersides that do not represent glues (of sufficient strength) on any tile type $t \in T$, then either one or two crawlers are created and that will transition to state **dead** after gathering 4 glues and doing a tile lookup.

7 The Frame

The *frame* is the portion of a supertile consisting of the outermost 4 rows on supersides that serve as potential input supersides for the creation of that supertile. The growth of any particular superside of a frame is dependent upon there being an adjacent supertile that has placed an “output” superside adjacent to this newly forming supertile (which uses it as an “input” superside), and therefore the




Figure 14: Win-Lose scenarios for all possible combinations of input supersides (with rotationally equivalent scenarios omitted). Green arrows show corners from which crawlers will begin and the directions in which they will grow.

frame for a particular supertile may consist of any non-empty subset of such input supersides.

The growth of the frame executes a sort of distributed algorithm in which supersides either grow or wait for adjacent supersides to grow before growing, so as to incorporate that superside’s information into their own. The basic frame formation algorithm is this: if “I” (a superside) win on both ends, then I fill in four layers that encode this information (see Figure 15). If I lose on either end, then this means there is a superside there that won, so I wait for it to fill in its four layers before forming my own four layers. This is because I want to incorporate all the information that the adjacent winning superside is able to gather into my own frame layer. This idea almost works, except that two cases (4.5 and 4.6) in Figure 14 have a Dining Philosopher’s deadlock problem: all supersides lost exactly one end, so all of them wait forever. Most of the complexity of the frame algorithm (and the need for four layers rather than just two) stem from the need to handle these cases.

The fix for these cases, at an intuitive level, works roughly as follows. The south, north, and east supersides work as described. The west superside is special. If west loses on both ends, then the configuration cannot be 4.5 or 4.6 in Figure 14, so it is safe for both ends to wait since frames of adjacent supersides will eventually arrive. Otherwise, if the west superside loses on exactly one end, then it sends a row of tiles (designated the “*WAI* signal”, standing for “Waiting on Additional Information”) towards the winning end. The *WAI* signal propagates around the edge of the supertile within the frame. For concreteness assume the losing end of the west superside is right (south) and the winning end is left (north).

Two scenarios are possible: either the win-loss configuration is 4.5 (in our concrete example) in Figure 14, or not. If so, then the *WAI* signal eventually traverses the entire supertile and reaches back to the west superside (on its losing end) to serve as the signal that the west superside’s losing end was waiting for. At this point, the entire frame can fill in with the complete information describing the win-loss configuration (hence each superside’s frame will have complete common knowledge of all supersides).

If the win-loss configuration is not 4.5 in Figure 14, then at least one superside to the right of west’s losing end must win on its right end. Therefore this superside is guaranteed to complete its frame, and all supersides between west and this superside (which all lost on their right end, including west) will receive the signal they are waiting for. In this scenario, the *WAI* signal is not needed, and it will “bounce” off of the first superside it encounters that wins on its right (guaranteed to exist as we just observed), without getting all the way around the supertile (which may not even be possible if there is a missing superside). But it is acceptable for the *WAI* signal not to make it all the way around, since in this scenario a different signal (that of the superside that won on its right end) will reach the losing end of the west superside. Therefore in either scenario, west’s losing end is guaranteed to receive the signal it is waiting for, preventing deadlock.

We now formally describe the frame construction algorithm. The first (outermost) layer of the frame serves merely as a “spacer” row and grows each superside (N, E, S, and W) as 2 independent pieces, with each beginning assembly from a position immediately adjacent to a corner of the square space to be occupied by the supertile forming. Each such half-row grows from its starting location (which is initiated by a strength 2 bond from the output superside of the input supertile) toward the center of its superside, stopping one location short of the center. This growth utilizes cooperation with the input supertile, thus transferring information about that input glue, etc., into the frame, and also prevents any frame growth along supersides that don’t have inputs. This transfer of information from the input sides is continued through all levels of the frame, resulting




Figure 15: The growth sequence of a WW superside of a frame for arbitrary direction D . Layers cooperate via strength 1 glues on North and South supersides except where bonds are explicitly shown with black tabs. (a) Layer 1 grows from each end towards the center, stopping with a one-tile gap. (b) The superside wins both competition corners, then grows toward the center. (c) Layer 2 completes via a single tile that binds via cooperation between the two halves, and exposes a double strength bond upward. Note that this tile contains the information that this superside was WW and passes that along. (d) Layer 3 grows by first attaching a tile to the double strength bond in the center of Layer 2, then grows outward to both supersides. (e) Layer 4 grows in the same way as Layer 3. The leftmost and rightmost tiles of Layer 4 expose single strength glues on their tops that can initiate frame growth for the losing supersides to the left and right (if they exist) and also pass along information that superside D is WW . Furthermore, only if D is South (or North), then the second leftmost (rightmost) tile also exposes such a glue.

in the interior of the frame having all of the necessary information in its exposed glues so that the next modules of the construction can form. In the few cases where frame tiles don't require cooperation on their outer edges for assembly, those positions are intentionally left empty by the input supertile in terms of output information. See Figure 20-1 for an example of the formation of the first layer of a frame for a supertile with 4 input supertiles.

The first tiles that must be placed for any superside of Layer 2 of a frame are *competition* tiles (pictured as the dark grey tiles in Figure 20). Competition tiles bind with strength 2 bonds to their input tiles from Layer 1, which are in turn similarly bound to their input supertiles. Thus, clearly for each competition tile position only one adjacent supertile can serve as the input that initiates its assembly. In the case where there are two previously assembled supertiles at positions necessary to initiate the placement of a particular competition tile, the actual competition tile that is placed represents a nondeterministic "competition" determining which one of those supertiles is actually bound to (with the other having a glue mismatch). We say that the supertile that is bound to *wins* the competition (often denoted simply by a W), and the other *loses* (L). When talking about the W/L pair for a particular superside, we denote it as WW , WL , LW , or LL with the first of the pair referring to the LEFT end and the second to the RIGHT end.

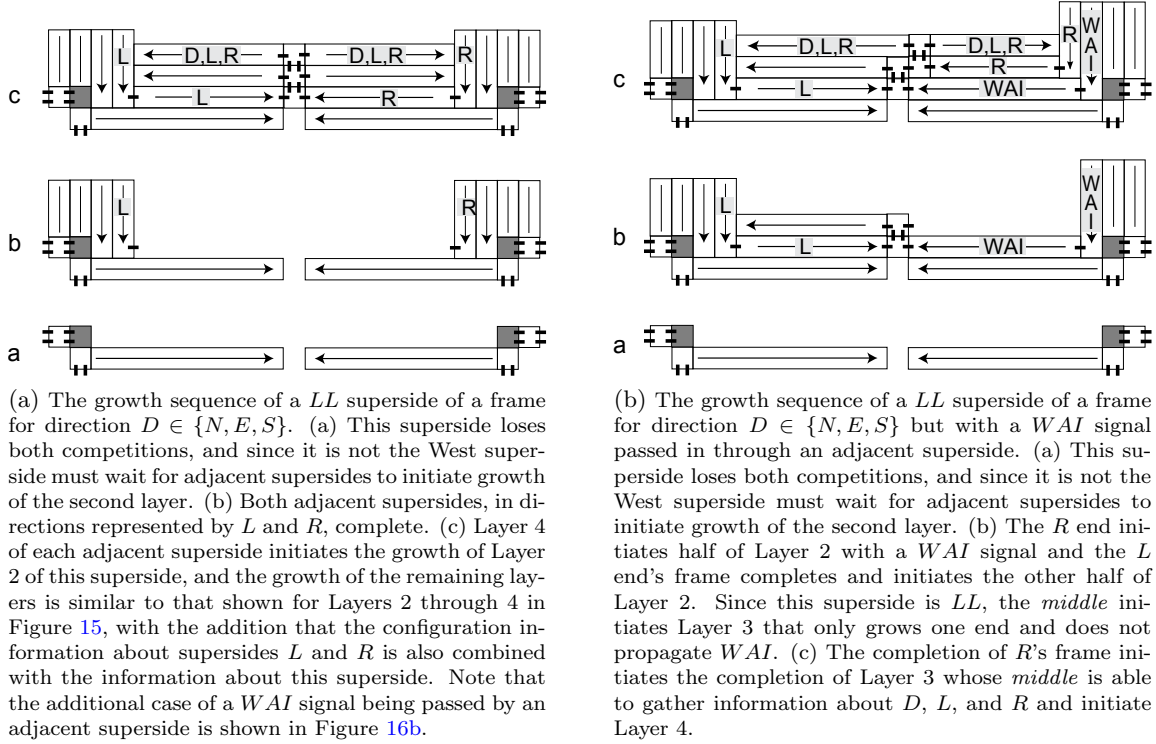


Figure 16: Examples of LL superside frame growth for supersides other than the West and beginning after the formation of Layer 1 that is shown in Figure 15 and is the same for all configurations.

Along with the subset of neighboring supertiles present, the pattern of W and L positions along with some nondeterminism based on asynchronous timing determine how the rest of a frame assembles. The entire algorithm for the growth of all supersides follows. Also, Figure 20 shows one possible growth sequence of a frame for a supertile with 4 potential input supersides, while Figures 21 and 22 show completed frames for supertiles with 4 potential input supersides that have differing W/L patterns. The full set of potential combinations of input supersides and W/L configurations (excluding rotational duplicates) can be seen in Figure 14.

Following is the *frame formation algorithm*, which specifies the growth of the 4 layers of the frame. Note that there are unique tile types for each superside and for each piece of information being passed such as W/L pairs for (possibly multiple) sides. Throughout the following algorithm, let $D \in \{N, S, E, W\}$ and $S \in \{W, L\}$. Each layer is broken into 3 components, the *halves*, which are the portions including a corner and extending toward the center but stopping one tile short of the center position, and the *middle*, which is the single tile position between 2 halves. Note that in some cases the *middle* may be one position away from the actual center of the row and thus one *half* is correspondingly one tile longer and the other is one shorter, and due to the fact that L -*halves* begin growth from portions of adjacent frame sides that have completed, they are shorter than W -*halves*. We say S -*half* to refer to a *half* that is on the S superside of a layer.

Note that scenarios 4.5 and 4.6 create conditions of potential deadlock in which all supersides have the same W/L configuration. To “break symmetry” and avoid deadlock, we designate the

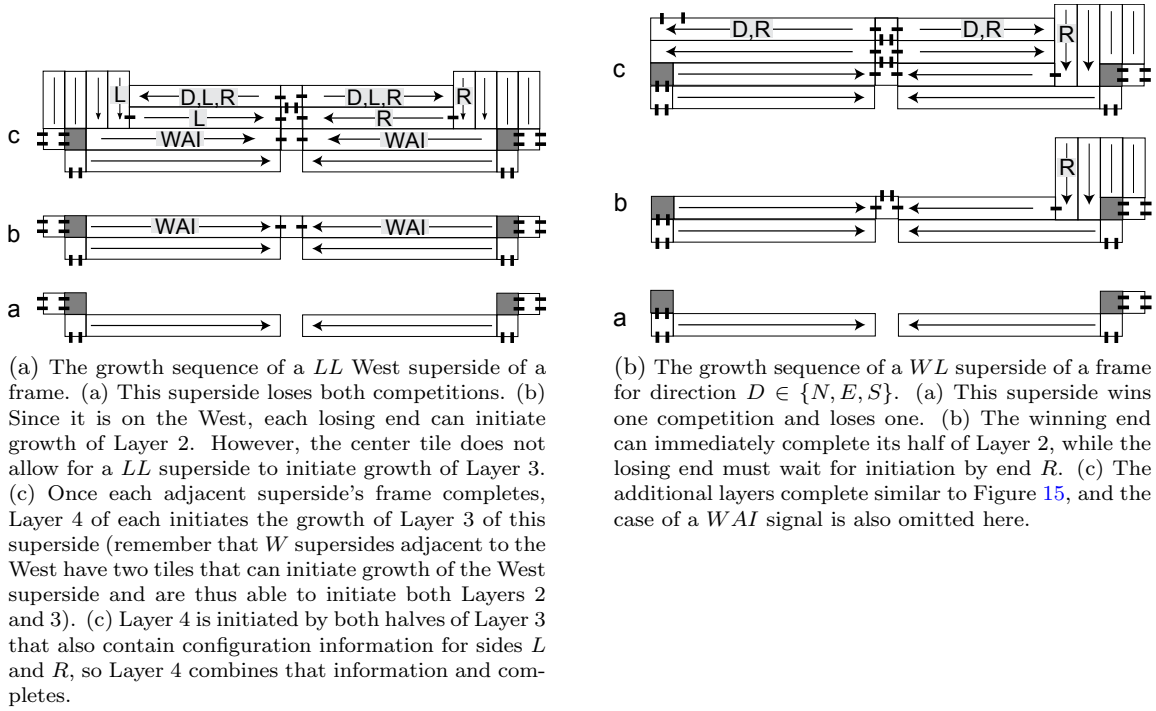


Figure 17: Another *LL* superside example and a *WL* superside example.

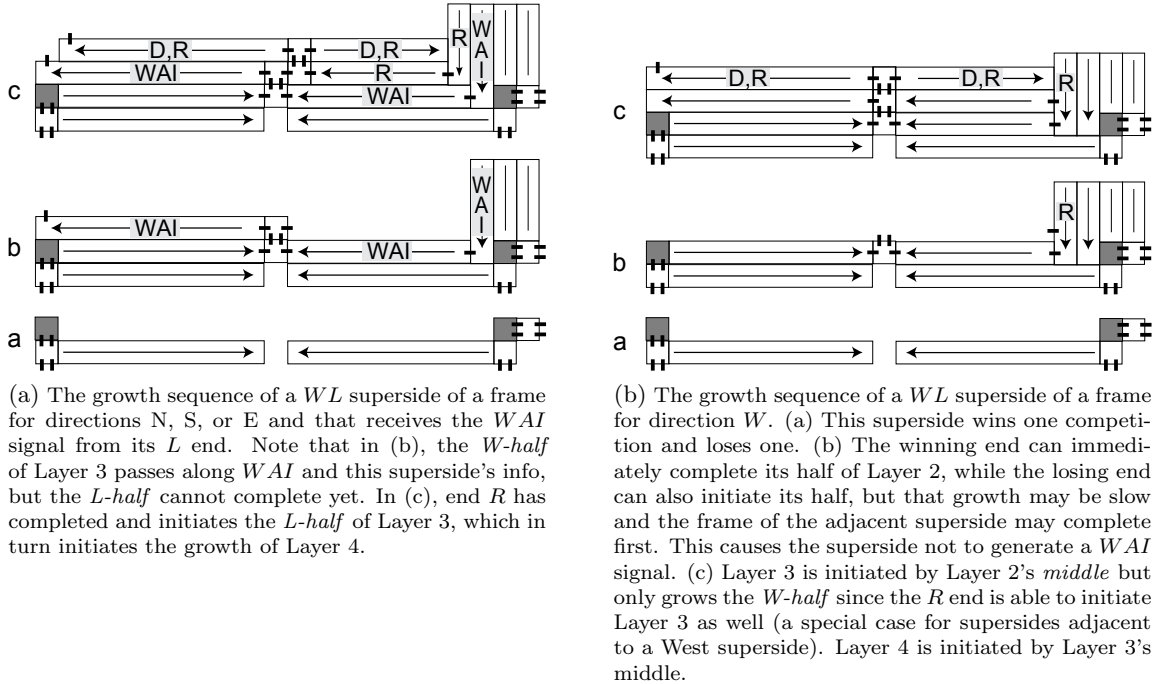



Figure 18: Various examples of *WL* superside growth for each direction.



(a) The growth sequence of a WL West superside of a frame when the WAI signal does not make it all the way around the frame (i.e. not scenarios 4.5 or 4.6). (a) This superside wins one competition and loses one. (b) The winning end can complete its half of Layer 2, and since this is the West superside, so can the losing end. However, when they meet in the center the halves of Layer 2 initiate the formation of a Layer 3 that only grows half of the layer towards the W end and sending the WAI signal. (c) The WAI signal row can initiate growth of the frame for the adjacent superside. Once end R 's frame completes, it initiates growth of the right half of Layer 3 that grows toward the center where it is able to cooperate with the L end to begin formation of Layer 4 that has information about this end and (at least) end R . Note that the left end of the completed frame can also initiate the remaining growth of end L 's frame.

(b) The growth sequence of a WL West superside of a frame in a scenario where the WAI signal makes it all the way around (which must be either scenario 4.5 or 4.6). Note that in (b), the W -half of Layer 3 passes along WAI and this superside's info, but the L -half cannot complete yet. (c) The WAI signal has made it completely around the frame and gathered the configuration information about every superside. This is used to complete Layer 3 and initiate and complete Layer 4, and is passed to the superside adjacent to the W , allowing it and subsequently every other superside to complete.

Figure 19: Examples of WL superside frame growth for the West superside when the L end is “fast” and initiated from the competition tile.

West superside as a special case superside to which slightly different rules of frame growth apply, as shown below.

1. Layer 1 growth (for all D): Both *halves* grow toward the center, and there is no tile for the *middle*.
2. Layer 2 growth:
 - (a) For $D \in \{N, S, E\}$:
 - i. Each W -half: Initiated by winning competition tile, grows completely.
 - ii. Each L -half: Initiated by the frame of the adjacent superside:
 - A. Adjacent superside completes: Complete standard L -half grows
 - B. Adjacent superside passes WAI signal: Complete L -half grows that passes WAI signal to *middle*
 - (b) For West:
 - i. Each W -half: Initiated by winning competition tile, grows completely





Figure 20: The phases of frame growth for case 4.2 from Figure 14.



(a) The full frame produced from the example in Figure 20. Note that the North (L/L) superside does not actually need to pass the *WAI* signal along in this configuration.

(b) The full frame produced from a different ordering of tile placements in the configuration of Figure 20. Namely, the West superside's South half doesn't begin layer 1 until after the South superside has completed. This provides the West superside with the necessary information to make the *WAI* signal unnecessary.

Figure 21: Examples of fully formed frames.




Figure 22: The full frame produced for configuration 4.5 of Figure 14, which shows the necessity of the *WAI* signal, namely to break the “deadlock” situation where all sides are W/L (from LEFT to RIGHT).

- ii. Each “Fast” *L-half*: Initiated by competition tile’s losing end, grows completely, and initiates *WAI* signal
 - iii. Each “Slow” *L-half*: Initiated by completed frame of adjacent superside, behaves same as non-West superside. This can occur for scenarios other than 4.5 and 4.6, depending on timing.
- (c) For all *D*:
- i. *Middle*:
 - A. *WW* superside: Initiates Layer 3 and passes superside info (i.e. *WW*)
 - B. *LL* superside: Does not initiate Layer 3 and thus doesn’t propagate (possible) *WAI* signal
 - C. *WL* or *LW* superside, and *WAI* signal from *L*: Initiates Layer 3 and passes *WAI* along with superside info
 - D. *WL* or *LW*, no *WAI* signal: Initiates Layer 3 and passes superside info (i.e. *WL* or *LW*)
3. Layer 3 growth:
- (a) *WW* superside (all *D*): Initiated by Layer 2 *middle*, both halves grow and Layer 4 is initiated from *middle*
 - (b) *LL* superside:
 - i. For $D \in \{N, S, E\}$:
 - A. Layer 2 doesn’t pass *WAI*: Initiated by Layer 2 *middle*, both halves grow and Layer 4 is initiated from *middle*
 - B. Layer 2 passes *WAI*: *W-half* initiated by Layer 2 and *WAI* is discarded, *L-half* initiated by completed frame of adjacent superside and causes placement of a “second *middle*” that initiates Layer 4
 - ii. For West superside: Both halves initiated by completed frames of adjacent sides, *middle* initiates Layer 4
 - (c) *WL* and *LW* sides:
 - i. For $D \in \{N, S, E\}$:
 - A. Layer 2 doesn’t pass *WAI*: Initiated by Layer 2 *middle*, both halves grow and Layer 4 is initiated from *middle*
 - B. Layer 2 passes *WAI*: *W-half* initiated by Layer 2 and *WAI* is propagated, *L-half* initiated by completed frame of adjacent superside and causes placement of a “second *middle*” that initiates Layer 4
 - ii. For West superside:
 - A. With “fast” *L-half*: *L-half* initiated by *WAI* from adjacent superside (if *WAI* makes it all the way around the frame) or by completed frame of adjacent superside and causes placement of a “second *middle*” that initiates Layer 4, *W-half* initiated by Layer 2 and passes *WAI*
 - B. With “slow” *L-half*: *L-half* initiated by completed frame of adjacent superside, *W-half* initiated by Layer 2 (with no need to create *WAI* since superside adjacent to *L* was able to complete)

4. Layer 4 growth:

- (a) *WW*, *WL* or *LW* superside: Initiated by *middle* of Layer 3 and both halves complete while copying information about all known sides from previous layer. Each end tile exposes a glue that can initiate frame growth for an adjacent superside (if it exists) while passing all known information about any sides, and if the adjacent superside is West, the tile next to the end tile also exposes that glue.
- (b) *LL* superside: Initiated by *middle* of Layer 3 and both halves complete while copying information about all known sides from previous layer.

8 Probes




Figure 23: A probe that grows from a west superside. The probe grows to length $\ell = \lceil (\text{superside length})/2 \rceil$, while copying the west glue from its base and generating a random number. Both are advertised as shown. Not shown: the probe length ℓ is encoded in the base of the probe as a binary sequence of $O(\log \ell)$ tiles.




Figure 24: Use of the probe table by a crawler to generate a probe region. This figure shows a crawler reading the probe table (bottom row) in order to determine which positions in the probe region have a probe, if the superside is south or west. North and east supersides have simpler probe regions that simply grow one probe in the position corresponding to the glue on that superside. The probe table also has a position for a potential probe for a strength- τ probe (not shown for brevity), which is grown if and only if the glue has strength τ .

In the case that a superside wins on both ends, it attempts to contact the opposite superside via probes. The superside could have a glue that is either strength $< \tau$ or strength $= \tau$. First

we discuss the case that the glue is strength $< \tau$. The design of probes ensures that probes from opposite sides meet in the middle (closing off the other two sides of the supertile from each other) if and only if their supersides represent glues that occur on a tile type and have sufficient strength to bind it. The probes grow from a region on the superside known as the “probe region”. Each glue in T has its own unique position in the probe region; see Figure 24. The sides are not symmetric: north and east grow probes in one way, and west and south grow them in a complementary way. Suppose the glue on the north is n and the glue on the south is s . North will grow a probe in the subregion associated with n . South will grow a probe in every subregion associated with a glue g that has the property that there is some tile type $t \in T$ with g on the north, s on the south, and g and s have combined strength at least τ . Therefore, if no tile type matches glue s on the south and n on the north (or matches it but n and s have insufficient strength), then the north and south probes leave sufficiently wide gaps for crawlers to later make their way around the probes (see Figure 10). Otherwise, the probes meet in the middle of the supertile and initiate their own crawlers (see Figure 6).

When two probes are complementary, they meet as shown in Figure 6. The probes will stop at a common row in the middle of the supertile, with one probe staggered so that it is one tile away from the other. Once both probes are complete, a single tile binds between the probes. This initiates two crawlers, one that is intended to output the west superside (in this example), and the other that is intended to output the east superside. The probes also advertise the glue they represent on their “tip” (north edge if a southern probe, east edge if a western probe, etc.) facing one side of the supertile, and on their side facing the other side of the supertile. This glue information is picked up by the crawler in order to combine the two glues from each probe for the later tile lookup.

The south and west probes each generate a random number $r \in \{1, \dots, |T|\}$, which is used by the crawlers to choose nondeterministically from among different potential output tiles during the tile lookup conducted later at the base of the probe. Since we need the two crawlers to output supersides that are consistent with a single tile type, the same random number is used, advertised on each side of the south probe (or west probe, in a rotated case of Figure 6). This ensures that each crawler will pick the same output tile type when they do their lookups.

The other reason to grow a probe is that the superside (again, south superside for concreteness) represents a strength- τ glue on a win-win frame. In this case, there is a potential problem if the north superside also represents a strength- τ glue on a win-win frame. If they both began to output supersides as in Figure 3, then if they represent different tile types, they could output glues to the east and west that are inconsistent (not actually shared by any tile type in T). Since we want only a single crawler to be responsible for outputting supersides to ensure consistency, we test for the presence of an opposite-side strength- τ superside (with a win-win frame configuration) using a probe. The purpose of these probes is not to cooperate, but merely to claim the “right” to fill out the supertile.

In Figure 4, the south superside has “won” and therefore begins filling out the east side. Unlike in Figure 3, however, before it can output on the north superside, a north input superside arrives, sending its own probe to attempt to claim the right to output. It loses, but the attempt cuts off the east from the west. Therefore, the probes use a similar cooperation as in the previous case to initiate a second crawler from the middle of the supertile, growing counterclockwise to eventually output on the west superside. Similarly to the “across-the-gap” cooperation case described first, if there is more than one tile type that shares the south glue, then the random number used to choose one of them must be the same as that used by the crawler that output on the east. Therefore the

probe advertises its random number on two sides, as in Figure 23.

It is clear by the design of the probe table that the space between each adjacent probe is at least $\Omega(|T|)$, so that crawlers (even two crawlers, one on top of the other) have sufficient space to crawl around probes in the case that the probes do not meet in the middle. This is because the width of any crawler is $O(\log |T|)$.

The following observation is clear from the probe design described above.

Observation 8.1. *Within a supertile, probes meet only if there are two opposite win-win sides and if we are simulating either (a) two-sided opposite binding for matching opposite sides or (b) two opposite strength- τ sides.*

9 Crawlers

Conceptually, *crawlers* are the components resembling thick bands of rows of tiles which assemble counterclockwise around the inside edges of a supertile, gathering glue information from input supertiles via frame sides or probes. Each time a crawler grows along a sufficient portion of a side of a frame to read the glue information being output by the adjacent supertile, it combines that glue information with any it may already have acquired. It then performs a lookup on the tile lookup table that encodes the simulated tile set to see if the currently acquired set of input glues represent a sufficient set of glues to match a tile type from the simulated set for the forming supertile to simulate. If so, the crawler then performs the duty of outputting the necessary output glue information along all available (non-input) sides of the supertile. If not, the crawler continues along an adjacent input side if it exists, or waits at a special location until and unless another input side arrives.

9.1 Crawler initiation

A crawler can be initiated in either the corner of a frame or near the tip of a probe. We first discuss those initiated by the corners of frames.

Figure 25 shows the basic set of conventions used when talking about supertile side orientation, defining a fixed LEFT and RIGHT corner for each side and showing the counterclockwise direction of growth of crawlers. The purpose of the frame is to gather and process the information about existing input sides in order to create initiation points for crawlers at exactly the corners specified by green arrows in the various input side and frame configurations shown in Figure 14. For shorthand, when referring to one of the configurations of Figure 14 (and its rotations), we will simply refer to it as “case $x.y$ ” where $x.y$ refers to the number in the bottom left corner of the box depicting that configuration. The set of cases shown in Figure 14 represent the full set of possible input side combinations and W/L configurations for each such combination (with cases which are rotationally equivalent omitted).

The following algorithm is implemented during frame formation, and determines whether or not the growth of a new crawler is initiated from each side, specifically the RIGHT corner of a given side. Crawler initiation occurs in corners created by two sides of a frame, by the placement of an initial crawler tile which uses cooperation between the two sides. Thus, each frame side must “agree” that a crawler should be initiated in that corner based on the information about its W/L configuration and that of any sides it has gathered. The frame formation algorithm was specifically




Figure 25: Basic conventions: LEFT and RIGHT ends as defined for each superside (North, East, South, and West), and counterclockwise direction of growth of crawlers.

designed to ensure that the gathered information is sufficient to make the correct decisions for every corner of every case.

The general rule is that if the LEFT of a side wins (is W), then the RIGHT corner adjacent to that side will initiate a crawler. The algorithm highlights the exceptions to that rule, which are the situations in which that rule is not followed and the crawler that would “normally” be created from a corner is prevented from forming, usually by a tile from one of the frame sides not providing the glue necessary for cooperation.

The RIGHT side corner of the $D \in \{N, E, S, W\}$ side initiates a crawler if and only if:

1. D is WW or WL , unless:
 - (a) It’s case 4.2 (in any rotation) and D is WW . In this case, the LL side contains information about the configurations of every side, and thus will prevent the initiation of the crawler that would normally form in the WW - LL corner.
 - (b) It’s case 4.5 and D is South, East, or West (all rotations of this case are the same). Therefore, in case 4.5 only the North side initiates a crawler (in the Northwest corner). All sides are WL , and the configurations of every side and thus the fact that case 4.5 is encountered is known by each side of the frame (due to the transmission of the WAI signal as discussed in the frame formation algorithm), and thus each side can place the correct tiles adjacent to the corners to ensure that only the North crawler begins.
2. It’s case 4.6 and D is North. Case 4.6 is the opposite of case 4.5 since all sides are LW . However, for the same general reason as in 1b, it can be guaranteed that exactly the North side will initiate a crawler.

In addition to initiation by frame corners, crawler growth can begin on probes. In the first such case, a strength- τ probe (i.e. a probe representing a glue of strength τ), if that probe wins its point of competition, it will initiate a crawler which will acquire the glue information and random number from that probe, then grow down the counterclockwise side of the probe and then around the frame. See Figure 3 for an example.

In the second case of crawler initiation by probes, *two* crawlers are initiated by a pair of matching probes which fully grow from opposite sides and meet to cooperate in the middle of the supertile, as

shown in Figure 6. In this case, one crawler forms on each side of the pair of now connected probes, each gathering and combining the glue information from both probes and the random number from one (by convention, the North or West). They then both grow down the counterclockwise sides of the probes adjacent to them and then counterclockwise around the supertile sides.

9.2 Multiple crawlers

As seen in Figure 14, no case of frame growth results in more than 2 crawlers being initiated by the frame. However, the cases in which 2 crawlers can arise allow for the possibility of one crawler needing to “piggyback” on the other. At a high level, if crawler B catches up with crawler A , it will orient itself so that it can assemble on A ’s “back”, or along the surface facing the interior of the supertile. Additionally, a crawler initiated by either a strength- τ probe or by the meeting of two matching probes can be forced to piggyback on another crawler. Crawler A will follow crawler B until it is forced to halt (for instance, when A becomes **full**), or until A halts and B can continue and begin outputting. Sections 9.3 and 9.4 present the details of how these state transitions occur.

Also, as shown in Figure 13, it is possible for two crawlers to independently gather enough glue information to decide the output tile type; in particular, this happens in some four-sided binding cases in which there are no empty output sides for one crawling to begin outputting on. In these cases, when a crawler reaches the fourth side, it now has the frame information of all four sides. Using this information, it can determine whether there are other crawlers. If there is another crawler, a precedence ranking on corners is used to break symmetry; one crawler “dies” while the other continues gathering glue information to determine the output tile type. The (arbitrary) precedence ranking we choose is $NW > SW > SE > NE$. As shown in Figure 13, this results in the orange crawler continuing on to determine the output tile type, while the green crawler halts.

9.3 Crawler states

Crawlers are $O(\log g)$ tiles tall. This allows a crawler to encode up to 4 glues, along with a random number n (generated by a component to be described) where $0 \leq n < \lceil \log g \rceil$

Every column of each crawler can be thought of as being in one of four logical states: **{unfilled, full, output, dead}**, and we consider the state of the latest column of a crawler to form to be that crawler’s “current” state. The crawler state is denoted by labels on each tile and this is what is used by the supertile replacement function R (see Section 3) to map the supertile to a tile type in T . The state transition diagram for crawlers is shown in Figure 26, and details of those transitions follow. Additionally, the growth of crawlers into and around frame corners is depicted in Figures 28-27. These figures show how crawlers interact with the frame as well as potential piggybacking crawlers.

When a crawler is initiated and begins growth, it is considered to be **unfilled**. It can transition to **full** only after it has gathered glues whose strengths sum to at least τ and which are consistent with glues of a tile type from T (which may already be the case when it is initiated) **and** it has completed a tile table lookup with those glues. Once **full**, every column of the crawler contains the full definition of the tile which was selected by the lookup (i.e. the encoding of the glues of its 4 sides). If a **full** crawler encounters an empty side to which it can output, it switches its state to **output**. The other cases where it can transition from **full** to **output** involve situations where all 4 adjacent supertiles formed first and provide potential input sides, and will be discussed in greater detail in Section 9.4.

In some situations, it is necessary for a crawler to permanently cease growth. This is accomplished by a transition to state **dead**. Following is the list of scenarios in which a crawler transitions to **dead**:

1. It is piggybacking on a crawler and grows to a point where the crawler beneath it transitioned to **full**
2. It grows onto a side which is WW and has enough glue information to know that the opposite side is also WW and that the probes from those sides will meet and thus determine the supertile's identity. For examples, see cases 3.4 and 4.1. In both cases, crawlers will grow along an LL side and then turn onto a WW side. Since they were initiated by a WW side, they will have information about both WW sides and be able to determine whether or not the probes from those sides will meet (which implies they have enough strength and match a tile type from T). This is performed by simultaneously performing dual lookups on the tile lookup table: one with all 3 currently collected glues and one with the 2 glues from opposite sides.
3. It arrives at a side which is WW and whose glue strength is τ , which means that either that side or the opposite side which may also have a τ -strength glue will initiate the necessary crawlers to complete the supertile.
4. It gathers 4 potential input glues but there is no matching tile from T
5. If a crawler gathers 4 potential input glues and is piggybacking on a crawler with a higher *precedence*, where the precedence order is $NW > SW > SE > NE$. This situation can occur in cases 4.1, 4.2, and 4.4. By one crawler halting it is ensured that in these cases exactly one crawler changes state to **output** and thus determines the supertile's identity.

9.4 Outputting crawlers

In this section, we discuss how a **full** crawler can change states to **output** and thus determine the identity of the forming supertile. In all situations other than those in which crawlers are initiated by matching probes meeting in the center of the supertile (which is possible in cases 2.1, 3.4, and 4.1), exactly one crawler changes state to **output**. In the cases of matching probes there will be exactly 2 such crawlers, but they will have both shared the output of a single random number selection and are thus guaranteed to make the same selection of supertile identity. Also note that, due to the location of the probes, those two crawlers will be completely isolated from each other and thus unable to encounter or crawl over each other.

The transition from **full** to **output** (note that only a **full** crawler can transition **output**) can occur in 3 general situations:

1. When a **full** crawler reaches a frame corner where the adjacent side has not formed, signifying that there may not be an adjacent supertile there and thus that side should serve as an output side of the forming supertile. If the crawler is able to win POC2 (as shown in Figure 29), it transitions to **output** and begins the process of outputting the superside (described below).
2. In the case where the crawler has failed to win all POC2's due to the existence of all 4 frame sides, and it determines that no other crawler will have higher precedence (See item 5 in Section 9.3 for more details.)

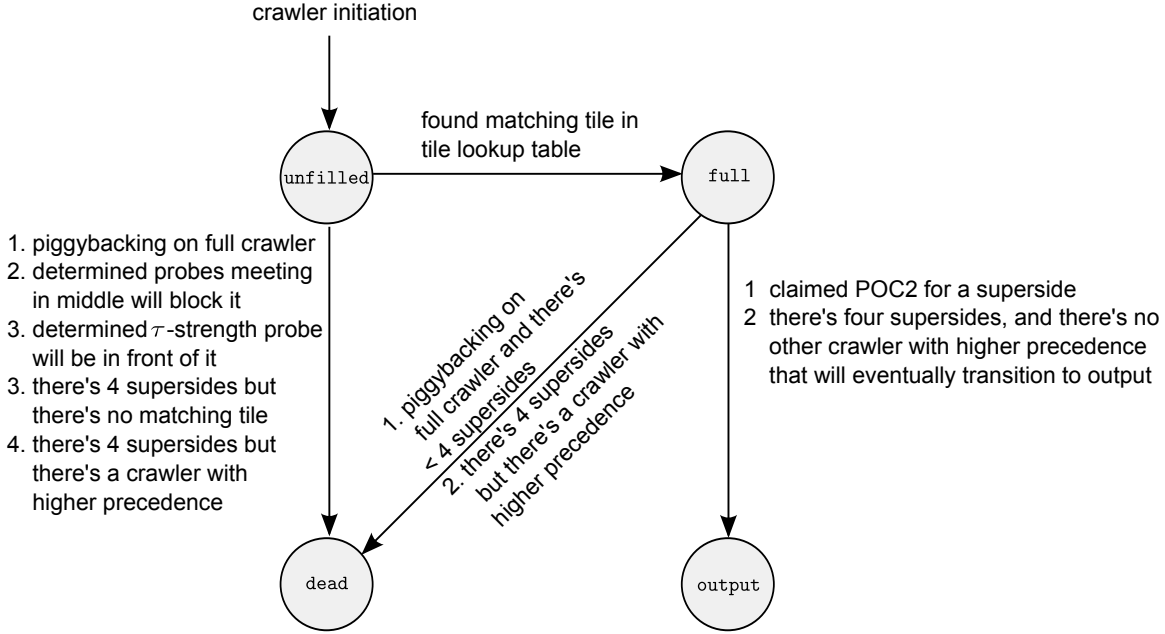


Figure 26: Definition of state transitions for crawlers. A state change occurs if any of the edge labels are satisfied. A crawler’s state determines whether it has encountered sufficient glue strength to output a tile type, and also whether it has actually taken responsibility for determining the output tile type (which it will not if another crawler exists that takes responsibility).




Figure 27: Three possible scenarios for a `full/output` crawler meeting a frame. (a) A crawler in either of states `full/output` that wants to produce an east output superside arrives at an empty east superside, (i.e. it claims POC1 and POC2, see Figure 30a for more details) and initiates outputting on the east. If the crawler was previously in state `full` it transitions to state `output`. Once it has output the east glue, it must pause until additional output information arrives from its left. (b) A crawler in either of states `full/output` that wants to produce an east output superside arrives at a non-empty east superside (i.e. it claims POC1, fails to claim POC2), which triggers growth of the crawler to the north. The crawler does not change state. It continues on its journey to find an empty north superside to place output. (c) A crawler in state `full` that wants to place output at a single superside to the east (after this it has no more output supersides to produce). However, the crawler arrives at a non-empty east superside (i.e. it claims POC1, fails to claim POC2). The crawler transitions to state `output` and outputs a single column of $O(\log g)$ tiles.




Figure 28: Three scenarios of crawlers arriving at frame corners. (a) A single, unfilled crawler arrives at the location which is the width of one crawler away from the corner. It then grows a single row of tiles, each of which cooperates with the frame below in order to pass the glue information from the frame to its top side, toward “point of competition 1” (POC1). The crawler loses POC1 as the right side of the frame is fully formed and initiates growth of a new crawler (following the crawler initiation protocol in Section 9.1). The newly formed crawler provides the cooperation necessary for the continuation of the first crawler. (b) An unfilled crawler and a piggybacking unfilled crawler near the corner. Since there are already 2 crawlers, it is guaranteed that no other crawler can be initiated and therefore the bottom crawler will win POC1. At this point, since both crawlers are unfilled the growth of both will pause until and unless the frame’s right side completes. (c) If and when the frame’s right side arrives, the 2 positions immediately above POC1 will be tiled utilizing cooperation from below and the right, ensuring that there is another input side for the crawler to collect (possibly allowing it to become full) and initiating the growth of the rows back toward the crawler which will result in its rotation and growth upward. (The need for 2 such rows is shown in Figure 29.) Note that the growth of the bottom crawler in this scenario is the same whether or not there is a piggybacking crawler.

3. In the case of 4-sided binding (where all 4 potential input sides actually serve as input sides to a simulated tile which would bind on all sides), once a crawler has acquired glue information about all 4 sides and determines that no other crawler will have a higher precedence

Recall that once a crawler transitions to `output`, every column then contains the full definition of the tile selected by the lookup along with a special marker to denote that that crawler is now in state `output` and contains the definitive output of the tile lookup and therefore the identify of the forming supertile. In fact, in situations where 4 potential input supersides exist, once the correct crawler transitions to `output`, it simply “prints” the definition of the selected tile on a single column of tiles and ceases growth. This provides the representation function with the necessary information to map the supertile to a tile from T .

When a crawler wins POC2 for a side, it begins the process of creating an output superside. A portion of this process is depicted in Figure 29(c). In order to do this, it begins the assembly of a row of tiles, known as the *output initiation row*, which grow back along the top surface of the crawler until reaching the representation of the tile lookup table which was passed upward through the growing crawler. Once arriving there, this row provides the cooperation necessary for the tile lookup and probe tables to grow upward and rotate into positions necessary to grow toward the side being output. Upon arriving next to the outputting crawler, which can’t continue growth until those tables arrive, they provide the cooperation necessary for the crawler to advance and also pass




Figure 29: Two additional scenarios of crawlers arriving at frame corners. (a) The bottom crawler (which in this scenario can only be unfilled because if it was full the top crawler would have halted) arrives near the corner and wins POC1 (which it must be able to do since there are already 2 crawlers). The piggybacking crawler is full and begins a path towards POC2 which, if won, would allow it to begin outputting. If the path is blocked at any point by the arrival of the frame’s right side, the completion of that frame side will provide the necessary cooperation for the third row which allows both crawlers to rotate and continue upward. (b) If the top crawler’s path to POC2 succeeds in winning POC2, it is guaranteed to be able to grow a return path which is unblocked and can initiate the transition of the top crawler to outputting. (See Figure 30b for more details.) (c) A crawler which wins POC2 and begins outputting sends an “output initiation row” across its top, back toward the center. When this row reaches the representation of the tile lookup table which was passed up through the crawler previously, it initiates growth that brings the tile lookup table and probe table upward. The definitions of those tables rotate as necessary to allow them to grow toward the side(s) being output. Once they arrive in position next to the outputting crawler, they provide the cooperation necessary for the crawler to continue growing, and also outputting the values of those tables. See Figure 3 for more details on the paths of growth of those tables.

the table encodings through to the output superside.

Note that the growth of the output initiation row is guaranteed never to be blocked by the presence of a piggybacking crawler (although climbing over it wouldn’t necessarily be problematic), since the only time that a piggybacking crawler continues growth along the back of a full crawler is in situations where there are 4 potential input supersides, and thus no output initiation row is possible. In other piggybacking scenarios, the bottom crawler transitions to full on the Left end of a superside, which is therefore where the piggybacking crawler would decide to transition to dead and cease growth, while a possible output initiation row would be initiated from the Right end and not need to grow back all the way to the halted piggybacking crawler.

10 The tile lookup table

The tile lookup table is composed of a row of tiles that encode T , in a way to be discussed, and leave space necessary for particular computations during the lookup process. This table is used by a crawler that grows across its surface reading the information encoded there and utilizing information about the supertile’s input glues (those it has already gathered) to find a table entry corresponding to a tile type $t \in T$ that has those same glues. If such an entry is found, the crawler extracts and encodes the definition of t and switches its state to full. If no such entry is found, the crawler preserves its encoding of collected input glues and remains in state unfilled.

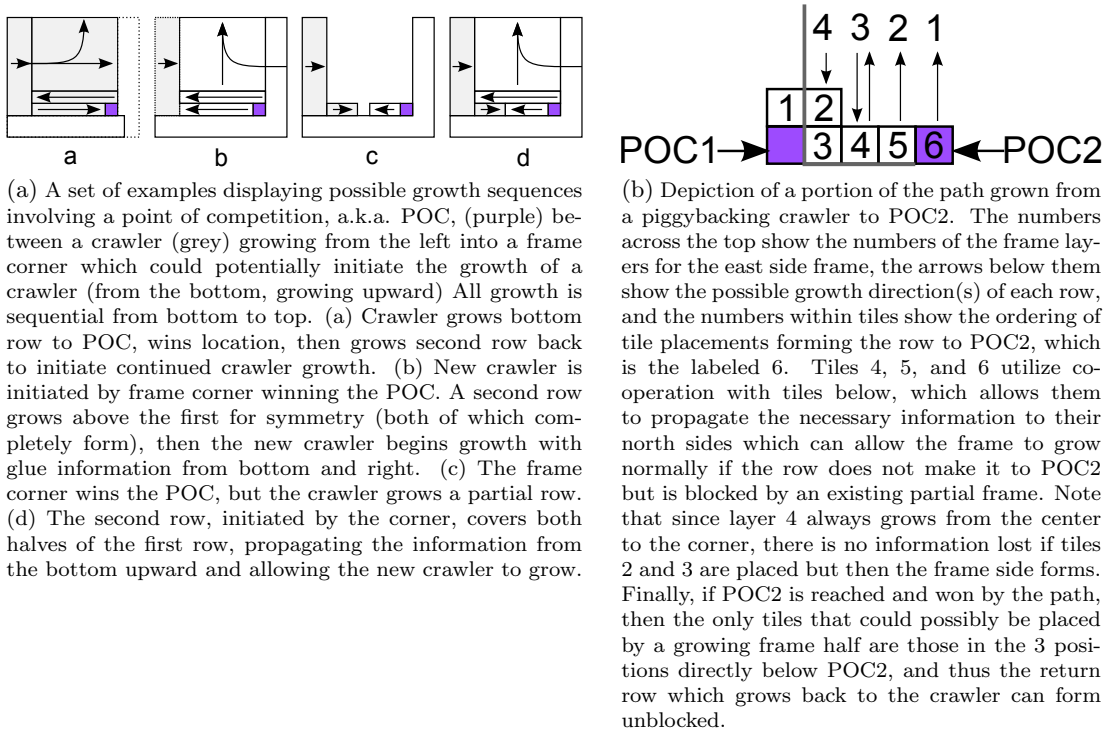


Figure 30: Points of competition between crawlers and frames.

10.1 Tile lookup table structure

Let G be the set of all glues on tiles in T plus the *null* glue, and $g = |G|$ be the count of those glues. Assign the elements of G a fixed ordering with the *null* glue at index 0, and for $x \in G$ let $\text{bin}(x)$ be the binary encoding of x 's index in G , padded to $\lceil \log g \rceil$ bits. Let $\text{bin}(x)_i$ for $0 \leq i < \lceil \log g \rceil$ represent the i th bit of $\text{bin}(x)$. An *address* A is a string of hexadecimal characters of length $\lceil \log g \rceil$, which is a combination of $\text{bin}(a)$, $\text{bin}(b)$, $\text{bin}(c)$, $\text{bin}(d)$ for $a, b, c, d \in G$ such that the i th position of A , A_i for $0 \leq i < \lceil \log g \rceil$, is the hexadecimal value of the bit string $\text{bin}(a)_i \text{bin}(b)_i \text{bin}(c)_i \text{bin}(d)_i$. We can denote such an address as $A_{a,b,c,d}$ to show the glues composing it and their ordering. By convention, the glues composing an address are assumed to be those of the sides of a tile ordered North, East, South, West. The absence of a glue (representing a missing side; note that this is different from the *null* glue) is represented by a null glue when doing a table lookup.

There are g^4 possible addresses for tiles in T (i.e. every possible combination and ordering of four glues from G). The purpose of the entries in the tile lookup table is to list the set of tiles from T that are consistent with each possible set of input sides to the supertile whose glue strengths sum to at least τ_T . Therefore, it is possible for each tile to map to several addresses. In fact, for each $t \in T$, for every subset of glues on t whose strengths sum to at least τ_T , t is *validly addressed* by every address that is composed of at least those glues on the sides on which they appear on t .

These addresses are “accessed” by crawlers by means of counters; if a tuple of glues (binary strings) whose concatenation represents an integer k , then by counting (once per entry) from k down to 0, a crawler can determine the location of the entry relevant to it. Most crawlers do only a single lookup, but one situation requires doing two lookups at the same time, illustrated in

Figure 8. In this case, the green crawler must do a standard look up corresponding to its three glues from the north, west, and south. However, it must simultaneously perform a lookup using only the north and south glues, to check whether the north and south probes will meet (halting if so). These two lookups can be performed simultaneously by running two independent counters in parallel as the crawler moves over the lookup table.

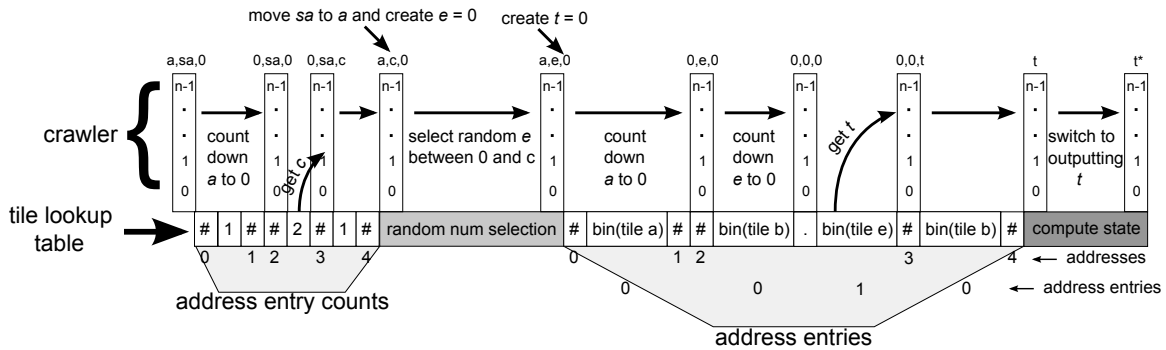


Figure 31: The tile lookup table (bottom row) and depiction of the front column of a crawler growing over it and performing a tile lookup.

The general structure of the tile lookup table for a trivial example is shown in Figure 31 as the bottom row of tiles. The main sections of the table are the “address entry counts”, “random number selection”, “address entries”, and “compute state”.

Address entry counts This section (the leftmost) of the tile lookup table consists of g^4 tiles with “#” symbols on their tops - one for every address, representing those addresses in order from 0 to $g^4 - 1$. For each address A , if n tiles from T are validly addressed by A , then the binary encoding of n is represented by the $\lceil \log n \rceil$ tiles immediately to the right of the A^{th} #, which is followed immediately by the $(A + 1)^{\text{th}}$ #. If $n = 0$ there are no tiles between those #'s.

Random number selection This section of the tile lookup table conveys no information on the top of tiles, but instead consists of a series of tiles that provide the spacing necessary for the crawler to perform the function of random number generation. Since there can be multiple entries for the address being looked up and exactly one of them must be selected, and it is also important to fairly choose between them, it is necessary to (approximately) uniformly select a random number to perform that selection. Since there is the requirement that the space used for random number generation be bounded (so it can't require assembly growth that collides with other portions of the supertile or grows outside of the $m \times m$ square), it is impossible to guarantee a perfectly uniform distribution for the random number selection (see [21] for details). However, in [21] the authors presented a number of modular tile sets that can be used to select random numbers while providing varying degrees of uniformity (with tradeoffs in the amount of space required), and for the random number selection component of the crawler it is possible to choose any of those tile sets. The result of the computation is the selection of a value e where $0 \leq e < n$, representing the address entry to be selected. Note that if a crawler is initiated by a probe, rather than the corner of a frame, the random number selection is performed by the probe so that the same random selection can be utilized by both of the crawlers that will grow into opposite sides of the supertile. In this case, a

random number r where $0 \leq r < g$ is contained within the crawler before it reaches the tile lookup table, and in this section of the table the computation that is performed is $e = r \% n$. Note that since $r < g$, $O(\log g)$ bits are required to represent r .

Address entries Similar to the “address entry counts” section, this section (the rightmost) of the tile lookup table consists of one tile representing a “#” symbol for every address, representing those addresses in order from 0 to $g^4 - 1$. For each address A , if n tiles from T are validly addressed by A , then there are n entries, each separated by a tile with a “.” symbol. Each entry consists of the address (which fully specifies the tile type) of one of those validly addressed tiles.

Compute state This section of the tile lookup table conveys no information on the top of its tiles, but instead consists of a sequence of spacer tiles that provide enough spacing for the crawler to perform the function of computing its (possibly new) state after having performed a tile lookup. The full set of state changes that can occur are shown in Section 9.3, and this computation is performed by a (constant sized) modular tile set that utilizes information gathered by the crawler related to: the number and W/L configuration of potential input sides, the result of the tile table lookup, and the crawler’s initiation number (as defined in Section 9.3). The state change is realized by the columns of the crawler forming from tiles specific to the new state, and in some cases a crawler may effectively transition from **unfilled** to **output** by making two quick transitions: first forming a column in state **full**, then the next in **output**.

10.2 Example

Figure 31 shows an example of a crawler growing across the top of a tile lookup table from left to right, with each vertical column representing the front of the crawler at a particular point. The initial column of the crawler contains two copies of the address for the combination of input sides that have been collected by the crawler, represented as a and sa , with the n hexadecimal value positions of each represented by n positions from bottom to top with the least significant position at the bottom. As the crawler grows to the right, across the “address entry counts” section, the value of a is decremented each time a position representing a “#” is encountered. Once a reaches 0, the correct address location has been reached and the following string of bits (which represents the number of entries for that address) is rotated upward into the crawler (occupying the position shown as c that was initially filled with 0’s). The crawler then grows to the section for random number selection where the value of sa is moved to the position for a and a position for number e is initialized to 0 before the crawler performs a random number selection between 0 and c . Next, the crawler once again counts a down to 0. At that point, it counts e down to 0, after which it has arrived at the appropriate entry of the correct address. The definition of the matching tile t is rotated upward into the crawler and the crawler grows toward the section for state computation. Finally, following the rules of crawler state transition as shown in Section 9.3, it can potentially change state, which consists of utilizing different sets of tiles marked with those states (if it continues growth forward).

10.3 Scale

The size of the tile lookup table dominates the size of a supertile and is therefore the main factor in the scaling factor m . There are g^4 addresses, falling into each of the following categories:

1. 1 address has 4 *null* glues and contains 0 entries
2. $4(g-1)$ addresses have exactly 1 non-*null* glue, and each such address can contain at most g^3 entries, totalling $4g^4 - 4g^3$ entries
3. $6(g-1)^2$ addresses have exactly 2 non-*null* glues, and each such address can contain at most g^2 entries, totalling $6g^4 - 12g^3 + 6g^2$ entries
4. $4(g-1)^3$ addresses have exactly 3 non-*null* glues, and each such address can contain at most g entries, totalling $4g^4 - 12g^3 + 12g^2 - 4g$ entries
5. $(g-1)^4$ addresses have 4 non-*null* glues, and each such address can contain exactly 1 entry, totalling $g^4 - 4g^3 + 6g^2 - 4g + 1$ entries

This yields a maximum possible $15g^4 - 32g^3 + 24g^2 - 8g + 1$ entries. Each entry in the “address entries” section requires $\log g$ space, for a total space requirement of $O(g^4 \log g)$, and since this section is the largest in the tile lookup table, the overall size is $O(g^4 \log g)$. In order to allow for crawlers to grow over the top of a tile lookup table and read the encoded information while also performing computations based on that information, it is necessary for the columns of a crawler to be able to grow in a zig-zag manner, growing one column from the bottom to top, then the next from the top to bottom. This means that only the tiles of every other column cooperate with those of the lookup table and thus read its information. Therefore, the tile lookup table actually consists of “spacer” tiles in every other position and is thus twice as wide as previously mentioned, but that doubling factor does not change the $O(g^4 \log g)$ bound.

A Abstract Tile Assembly Model

This section gives a terse definition of the abstract Tile Assembly Model (aTAM, [48]). This is not a tutorial; for readers unfamiliar with the aTAM, [41] gives an excellent introduction to the model.

Fix an alphabet Σ . Σ^* is the set of finite strings over Σ . \mathbb{Z} , \mathbb{Z}^+ , and \mathbb{N} denote the set of integers, positive integers, and nonnegative integers, respectively. Given $A \subseteq \mathbb{Z}^2$, the *full grid graph* of A is the undirected graph $G_A^f = (V, E)$, where $V = A$, and for all $u, v \in V$, $\{u, v\} \in E \iff \|u-v\|_2 = 1$; i.e., iff u and v are adjacent on the integer Cartesian plane.

A *tile type* is a tuple $t \in (\Sigma^* \times \mathbb{N})^4$; i.e., a unit square with four sides listed in some standardized order, each side having a *glue* $g \in \Sigma^* \times \mathbb{N}$ consisting of a finite string *label* and nonnegative integer *strength*. We assume a finite set T of tile types, but an infinite number of copies of each tile type, each copy referred to as a *tile*. An *assembly* is a nonempty connected arrangement of tiles on the integer lattice \mathbb{Z}^2 , i.e., a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T$ such that $G_{\text{dom } \alpha}^f$ is connected and $\text{dom } \alpha \neq \emptyset$. Let \mathcal{A}^T denote the set of all assemblies of tiles from T , and let $\mathcal{A}_{<\infty}^T$ denote the set of finite assemblies of tiles from T . The *shape* $S_\alpha \subseteq \mathbb{Z}^2$ of α is $\text{dom } \alpha$. Two adjacent tiles in an assembly *interact* if the glues on their abutting sides are equal (in both label and strength) and have positive strength. Each assembly α induces a *binding graph* G_α^b , a grid graph whose vertices are positions occupied by tiles, with an edge between two vertices if the tiles at those vertices interact.⁸ Given $\tau \in \mathbb{Z}^+$, α is τ -*stable* if every cut of G_α^b has weight at least τ , where the weight of an edge is the strength of the glue it represents. That is, α is τ -stable if at least energy τ is

⁸For $G_{S_\alpha}^f = (V_{S_\alpha}, E_{S_\alpha})$ and $G_\alpha^b = (V_\alpha, E_\alpha)$, G_α^b is a spanning subgraph of $G_{S_\alpha}^f$: $V_\alpha = V_{S_\alpha}$ and $E_\alpha \subseteq E_{S_\alpha}$.

required to separate α into two parts. When τ is clear from context, we say α is *stable*. Given two assemblies $\alpha, \beta : \mathbb{Z}^2 \dashrightarrow T$, we say α is a *subassembly* of β , and we write $\alpha \sqsubseteq \beta$, if $S_\alpha \subseteq S_\beta$ and, for all points $p \in S_\alpha$, $\alpha(p) = \beta(p)$.

A *tile assembly system* (TAS) is a triple $\mathcal{T} = (T, \sigma, \tau)$, where T is a finite set of tile types, $\sigma : \mathbb{Z}^2 \dashrightarrow T$ is the finite, τ -stable *seed assembly*, and $\tau \in \mathbb{Z}^+$ is the *temperature*. Given two τ -stable assemblies $\alpha, \beta : \mathbb{Z}^2 \dashrightarrow T$, we write $\alpha \rightarrow_1^\mathcal{T} \beta$ if $\alpha \sqsubseteq \beta$ and $|S_\beta \setminus S_\alpha| = 1$. In this case we say α *\mathcal{T} -produces β in one step*.⁹ If $\alpha \rightarrow_1^\mathcal{T} \beta$, $S_\beta \setminus S_\alpha = \{p\}$, and $t = \beta(p)$, we write $\beta = \alpha + (p \mapsto t)$. The *\mathcal{T} -frontier* of α is the set $\partial^\mathcal{T}\alpha = \bigcup_{\alpha \rightarrow_1^\mathcal{T} \beta} S_\beta \setminus S_\alpha$, the set of empty locations at which a tile could stably attach to α .

A sequence of $k \in \mathbb{Z}^+ \cup \{\infty\}$ assemblies $\alpha_0, \alpha_1, \dots$ is a *\mathcal{T} -assembly sequence* if, for all $1 \leq i < k$, $\alpha_{i-1} \rightarrow_1^\mathcal{T} \alpha_i$. We write $\alpha \rightarrow^\mathcal{T} \beta$, and we say α *\mathcal{T} -produces β* (in 0 or more steps) if there is a \mathcal{T} -assembly sequence $\alpha_0, \alpha_1, \dots$ of length $k = |S_\beta \setminus S_\alpha| + 1$ such that 1) $\alpha = \alpha_0$, 2) $S_\beta = \bigcup_{0 \leq i < k} S_{\alpha_i}$, and 3) for all $0 \leq i < k$, $\alpha_i \sqsubseteq \beta$. If k is finite then it is routine to verify that $\beta = \alpha_{k-1}$.¹⁰ We say α is *\mathcal{T} -producible* if $\sigma \rightarrow^\mathcal{T} \alpha$, and we write $\mathcal{A}[\mathcal{T}]$ to denote the set of \mathcal{T} -producible assemblies. The relation $\rightarrow^\mathcal{T}$ is a partial order on $\mathcal{A}[\mathcal{T}]$ [33, 40]. A \mathcal{T} -assembly sequence $\alpha_0, \alpha_1, \dots$ is *fair* if, for all i and all $p \in \partial^\mathcal{T}\alpha_i$, there exists j such that $\alpha_j(p)$ is defined; i.e., no frontier location is “starved”.

An assembly α is *\mathcal{T} -terminal* if α is τ -stable and $\partial^\mathcal{T}\alpha = \emptyset$. We write $\mathcal{A}_\square[\mathcal{T}] \subseteq \mathcal{A}[\mathcal{T}]$ to denote the set of \mathcal{T} -producible, \mathcal{T} -terminal assemblies. A TAS \mathcal{T} is *directed* (a.k.a., *deterministic, confluent*) if the poset $(\mathcal{A}[\mathcal{T}], \rightarrow^\mathcal{T})$ is directed; i.e., if for each $\alpha, \beta \in \mathcal{A}[\mathcal{T}]$, there exists $\gamma \in \mathcal{A}[\mathcal{T}]$ such that $\alpha \rightarrow^\mathcal{T} \gamma$ and $\beta \rightarrow^\mathcal{T} \gamma$.¹¹

When \mathcal{T} is clear from context, we may omit \mathcal{T} from the notation above and instead write \rightarrow_1 , \rightarrow , $\partial\alpha$, *frontier*, *assembly sequence*, *produces*, *producible*, and *terminal*. We make the following assumptions that do not affect the fundamental capabilities of the model, but which will simplify our main construction. Since the behavior of a TAS $\mathcal{T} = (T, \sigma, \tau)$ is unchanged if every glue with strength greater than τ is changed to have strength exactly τ , we assume henceforth that all glue strengths are in the set $\{0, 1, \dots, \tau\}$. We assume that glue labels are never shared between glues of unequal strength.

Acknowledgement

We would like to thank Matthew Cook for pointing out a simplification to our construction.

References

- [1] Zachary Abel, Nadia Benbernou, Mirela Damian, Erik D. Demaine, Martin Demaine, Robin Flatland, Scott Kominers, and Robert Schweller, *Shape replication through self-assembly and*

⁹Intuitively $\alpha \rightarrow_1^\mathcal{T} \beta$ means that α can grow into β by the addition of a single tile; the fact that we require both α and β to be τ -stable implies in particular that the new tile is able to bind to α with strength at least τ . It is easy to check that had we instead required only α to be τ -stable, and required that the cut of β separating α from the new tile has strength at least τ , then this implies that β is also τ -stable.

¹⁰If we had defined the relation $\rightarrow^\mathcal{T}$ based on only finite assembly sequences, then $\rightarrow^\mathcal{T}$ would be simply the reflexive, transitive closure $(\rightarrow_1^\mathcal{T})^*$ of $\rightarrow_1^\mathcal{T}$. But this would mean that no infinite assembly could be produced from a finite assembly, even though there is a well-defined, unique “limit assembly” of every infinite assembly sequence.

¹¹The following two convenient characterizations of “directed” are routine to verify. \mathcal{T} is directed if and only if $|\mathcal{A}_\square[\mathcal{T}]| = 1$. \mathcal{T} is *not* directed if and only if there exist $\alpha, \beta \in \mathcal{A}[\mathcal{T}]$ and $p \in S_\alpha \cap S_\beta$ such that $\alpha(p) \neq \beta(p)$.

- RNase enzymes*, SODA 2010: Proceedings of the Twenty-first Annual ACM-SIAM Symposium on Discrete Algorithms (Austin, Texas), Society for Industrial and Applied Mathematics, 2010.
- [2] Leonard M. Adleman, Qi Cheng, Ashish Goel, and Ming-Deh Huang, *Running time and program size for self-assembled squares*, STOC 2001: Proceedings of the thirty-third annual ACM Symposium on Theory of Computing (Hersonissos, Greece), ACM, 2001, pp. 740–748.
- [3] Leonard M. Adleman, Qi Cheng, Ashish Goel, Ming-Deh Huang, and Hal Wasserman, *Linear self-assemblies: Equilibria, entropy and convergence rates*, In Sixth International Conference on Difference Equations and Applications, Taylor and Francis, 2001.
- [4] Leonard M. Adleman, Qi Cheng, Ashish Goel, Ming-Deh A. Huang, David Kempe, Pablo Moisset de Espanés, and Paul W. K. Rothmund, *Combinatorial optimization problems in self-assembly*, STOC 2002: Proceedings of the Thirty-Fourth Annual ACM Symposium on Theory of Computing, 2002, pp. 23–32.
- [5] Leonard M. Adleman, Jarkko Kari, Lila Kari, Dustin Reishus, and Petr Sosík, *The undecidability of the infinite ribbon problem: Implications for computing by self-assembly*, SIAM Journal on Computing **38** (2009), no. 6, 2356–2381, Preliminary version appeared in FOCS 2002.
- [6] Gagan Aggarwal, Qi Cheng, Michael H. Goldwasser, Ming-Yang Kao, Pablo Moisset de Espanés, and Robert T. Schweller, *Complexities for generalized models of self-assembly*, SIAM Journal on Computing **34** (2005), 1493–1515, Preliminary version appeared in SODA 2004.
- [7] Florent Becker, Ivan Rapaport, and Eric Rémila, *Self-assembling classes of shapes with a minimum number of tiles, and in optimal time*, FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science, 2006, pp. 45–56.
- [8] Nathaniel Bryans, Ehsan Chiniforooshan, David Doty, Lila Kari, and Shinnosuke Seki, *The power of nondeterminism in self-assembly*, SODA 2011: Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2011, pp. 590–602.
- [9] Sarah Cannon, Erik D. Demaine, Martin L. Demaine, Sarah Eisenstat, Matthew J. Patitz, Robert Schweller, Scott M. Summers, and Andrew Winslow, *Two hands are better than one (up to constant factors)*, Arxiv preprint arXiv:1201.1650 (2012).
- [10] Eric Goles Ch., Cedric Little, and Ivan Rapaport, *Understanding a non-trivial cellular automaton by finding its simplest underlying communication protocol*, ISAAC 2008, The 19th International Symposium on Algorithms and Complexity, Lecture Notes in Computer Science, vol. 5369, 2008, pp. 592–604.
- [11] Eric Goles Ch., Pierre-Etienne Meunier, Ivan Rapaport, and Guillaume Theyssier, *Communication complexity and intrinsic universality in cellular automata*, Theor. Comput. Sci. **412** (2011), no. 1-2, 2–21.
- [12] Harish Chandran, Nikhil Gopalkrishnan, and John H. Reif, *The tile complexity of linear assemblies*, ICALP 2009: 36th International Colloquium on Automata, Languages and Programming, vol. 5555, Springer, 2009, pp. 235–253.

- [13] Ho-Lin Chen and David Doty, *Parallelism and time in hierarchical self-assembly*, SODA 2012: Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2012.
- [14] Ho-Lin Chen, David Doty, and Shinnosuke Seki, *Program size and temperature in self-assembly*, ISAAC 2011: Proceedings of the 22nd International Symposium on Algorithms and Computation, Lecture Notes in Computer Science, vol. 7074, Springer-Verlag, 2011, pp. 445–453.
- [15] Matthew Cook, Yunhui Fu, and Robert T. Schweller, *Temperature 1 self-assembly: Deterministic assembly in 3d and probabilistic assembly in 2d*, SODA, 2011, pp. 570–589.
- [16] Marianne Delorme, Jacques Mazoyer, Nicolas Ollinger, and Guillaume Theyssier, *Bulking II: Classifications of cellular automata*, Theor. Comput. Sci. **412** (2011), no. 30, 3881–3905.
- [17] Erik D. Demaine, Martin L. Demaine, Sandor Fekete, Matthew J. Patitz, Robert Schweller, Andrew Winslow, and Damien Woods, *One tile to simulate any tile system*, Preprint (2012).
- [18] 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 **7** (2008), no. 3, 347–370, Preliminary version appeared in DNA 13.
- [19] Erik D. Demaine, Matthew J. Patitz, Robert T. Schweller, and Scott M. Summers, *Self-assembly of arbitrary shapes using RNase enzymes: Meeting the Kolmogorov bound with small scale factor*, STACS 2011: Proceedings of the 28th International Symposium on Theoretical Aspects of Computer Science, 2011.
- [20] David Doty, *Randomized self-assembly for exact shapes*, SIAM Journal on Computing **39** (2010), no. 8, 3521–3552, Preliminary version appeared in FOCS 2009.
- [21] David Doty, Jack H. Lutz, Matthew J. Patitz, Scott M. Summers, and Damien Woods, *Random number selection in self-assembly*, UC 2009: Proceedings of The Eighth International Conference on Unconventional Computation, Lecture Notes in Computer Science, vol. 5715, Springer, 2009, pp. 143–157.
- [22] ———, *Intrinsic universality in self-assembly*, STACS 2010: Proceedings of The Twenty-Seventh International Symposium on Theoretical Aspects of Computer Science, Leibniz International Proceedings in Informatics (LIPIcs), vol. 5, 2010, pp. 275–286.
- [23] David Doty, Matthew J. Patitz, Dustin Reishus, Robert T. Schweller, and Scott M. Summers, *Strong fault-tolerance for self-assembly with fuzzy temperature*, FOCS 2010: Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science, IEEE, 2010, pp. 417–426.
- [24] B. Durand and Zs. Róka, *The game of life: universality revisited*, Cellular Automata (M. Delorme and J. Mazoyer, eds.), Kluwer, 1999.
- [25] Kenichi Fujibayashi, Rizal Hariadi, Sung Ha Park, Erik Winfree, and Satoshi Murata, *Toward reliable algorithmic self-assembly of DNA tiles: A fixed-width cellular automaton pattern*, Nano Letters **8** (2007), no. 7, 1791–1797.
- [26] D. Han, S. Pal, J. Nangreave, Z. Deng, Y. Liu, and H. Yan, *DNA origami with complex curvatures in three-dimensional space*, Science **332** (2011), no. 6027, 342.

- [27] Y. He, T. Ye, M. Su, C. Zhang, A.E. Ribbe, W. Jiang, and C. Mao, *Hierarchical self-assembly of DNA into symmetric supramolecular polyhedra*, Nature **452** (2008), no. 7184, 198–201.
- [28] Ming-Yang Kao and Robert T. Schweller, *Reducing tile complexity for self-assembly through temperature programming*, SODA 2006: Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, 2006, pp. 571–580.
- [29] ———, *Randomized self-assembly for approximate shapes*, ICALP 2008: International Colloquium on Automata, Languages, and Programming (Luca Aceto, Ivan Damgård, Leslie Ann Goldberg, Magnús M. Halldrsson, Anna Ingólfssdóttir, and Igor Walukiewicz, eds.), Lecture Notes in Computer Science, vol. 5125, Springer, 2008, pp. 370–384.
- [30] Grégory Lafitte and Michael Weiss, *Universal tilings*, STACS 2007, 24th Annual Symposium on Theoretical Aspects of Computer Science, Aachen, Germany, February 22-24, 2007, Proceedings (Wolfgang Thomas and Pascal Weil, eds.), Lecture Notes in Computer Science, vol. 4393, Springer, 2007, pp. 367–380.
- [31] ———, *Simulations between tilings*, Tech. report, University of Athens, 2008.
- [32] ———, *An almost totally universal tile set*, Theory and Applications of Models of Computation, 6th Annual Conference, TAMC 2009, Changsha, China, May 18-22, 2009. Proceedings (Jianer Chen and S. Barry Cooper, eds.), Lecture Notes in Computer Science, vol. 5532, Springer, 2009, pp. 271–280.
- [33] James I. Lathrop, Jack H. Lutz, and Scott M. Summers, *Strict self-assembly of discrete Sierpinski triangles*, Theoretical Computer Science **410** (2009), 384–405, Preliminary version appeared in CiE 2007.
- [34] Kyle Lund, Anthony T. Manzo, Nadine Dabby, Nicole Micholotti, Alexander Johnson-Buck, Jeanetter Nangreave, Steven Taylor, Renjun Pei, Milan N. Stojanovic, Nils G. Walter, Erik Winfree, and Hao Yan, *Molecular robots guided by prescriptive landscapes*, Nature **465** (2010), 206–210.
- [35] Ján Maňuch, Ladislav Stacho, and Christine Stoll, *Step-assembly with a constant number of tile types*, ISAAC 2009: Proceedings of the 20th International Symposium on Algorithms and Computation (Berlin, Heidelberg), Springer-Verlag, 2009, pp. 954–963.
- [36] Nicolas Ollinger, *Intrinsically universal cellular automata*, CSP, 2008, pp. 199–204.
- [37] Matthew J. Patitz, *Simulation of self-assembly in the abstract tile assembly model with ISU TAS*, FNANO 2009: 6th Annual Conference on Foundations of Nanoscience: Self-Assembled Architectures and Devices (Snowbird, Utah, USA, April 20-24 2009), Sciencetechnica, 2009, pp. 209–219.
- [38] L. Qian and E. Winfree, *Scaling up digital circuit computation with DNA strand displacement cascades*, Science **332** (2011), no. 6034, 1196.
- [39] L. Qian, E. Winfree, and J. Bruck, *Neural network computation with dna strand displacement cascades*, Nature **475** (2011), no. 7356, 368–372.

- [40] Paul W. K. Rothmund, *Theory and experiments in algorithmic self-assembly*, Ph.D. thesis, University of Southern California, December 2001.
- [41] Paul W. K. Rothmund and Erik Winfree, *The program-size complexity of self-assembled squares (extended abstract)*, STOC 2000: Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, 2000, pp. 459–468.
- [42] Paul W.K. Rothmund, Nick Papadakis, and Erik Winfree, *Algorithmic self-assembly of DNA Sierpinski triangles*, PLoS Biology **2** (2004), no. 12, 2041–2053.
- [43] PW Rothmund, *Folding DNA to create nanoscale shapes and patterns*, Nature **440** (2006), no. 7082, 297–302.
- [44] Nadrian C. Seeman, *Nucleic-acid junctions and lattices*, Journal of Theoretical Biology **99** (1982), 237–247.
- [45] David Soloveichik and Erik Winfree, *Complexity of self-assembled shapes*, SIAM Journal on Computing **36** (2007), no. 6, 1544–1569, Preliminary version appeared in DNA 10.
- [46] Scott M. Summers, *Reducing tile complexity for the self-assembly of scaled shapes through temperature programming*, Algorithmica, to appear.
- [47] Hao Wang, *Proving theorems by pattern recognition – II*, The Bell System Technical Journal **XL** (1961), no. 1, 1–41.
- [48] Erik Winfree, *Algorithmic self-assembly of DNA*, Ph.D. thesis, California Institute of Technology, June 1998.
- [49] Erik Winfree, Furong Liu, Lisa A. Wenzler, and Nadrian C. Seeman, *Design and self-assembly of two-dimensional DNA crystals.*, Nature **394** (1998), no. 6693, 539–44.
- [50] B. Yurke, A.J. Turberfield, A.P. Mills Jr, F.C. Simmel, and J.L. Neumann, *A DNA-fuelled molecular machine made of DNA*, Nature **406** (2000), no. 6796, 605–608.