

The non-cooperative tile assembly model is not intrinsically universal or capable of bounded Turing machine simulation

Pierre-Étienne Meunier
Inria
pierre-etienne.meunier@inria.fr*

Damien Woods
Inria
damien.woods@inria.fr†

Abstract

The field of algorithmic self-assembly is concerned with the computational and expressive power of nanoscale self-assembling molecular systems. In the well-studied cooperative, or temperature 2, abstract tile assembly model it is known that there is a tile set to simulate any Turing machine and an intrinsically universal tile set that simulates the shapes and dynamics of any instance of the model, up to spatial rescaling. It has been an open question as to whether the seemingly simpler noncooperative, or temperature 1, model is capable of such behaviour. Here we show that this is not the case, by showing that there is no tile set in the noncooperative model that is intrinsically universal, nor one capable of time-bounded Turing machine simulation within a bounded region of the plane.

Although the noncooperative model intuitively seems to lack the complexity and power of the cooperative model it has been exceedingly hard to prove this. One reason is that there have been few tools to analyse the structure of complicated paths in the plane. This paper provides a number of such tools. A second reason is that almost every obvious and small generalisation to the model (e.g. allowing error, 3D, non-square tiles, signals/wires on tiles, tiles that repel each other, parallel synchronous growth) endows it with great computational, and sometimes simulation, power. Our main results show that all of these generalisations provably increase computational and/or simulation power. Our results hold for both deterministic and nondeterministic noncooperative systems. Our first main result stands in stark contrast with the fact that for both the cooperative tile assembly model, and for 3D noncooperative tile assembly, there are respective intrinsically universal tilesets. Our second main result gives a new technique (reduction to simulation) for proving negative results about computation in tile assembly.

*This work was carried out while at Inria, Paris, France, and the Department of Computer Science, Aalto University, Finland, and Aix Marseille Université, CNRS, LIF UMR 7279, 13288, Marseille, France, and LIAFA UMR 7089, Paris 7, France, and California Institute of Technology, Pasadena, CA 91125, USA. Supported in part by National Science Foundation Grant CCF-1219274.

†This work was carried out while at Inria, Paris, France, as well as California Institute of Technology, Pasadena, CA 91125, USA, and during a brief visit to LIAFA (UMR 7089), Paris 7, France. Supported by National Science Foundation grants CCF-1219274, 0832824 (The Molecular Programming Project), CCF-1219274, and CCF-1162589, USA, a visiting professor award from Paris 7, and by INRIA.

1 Introduction

The design and laboratory fabrication of nanoscale molecular systems that implement sophisticated computation is a goal held by many. If we are to have such an engineering discipline that exploits the idea that molecules can compute, then we need a firm foundation of the kind of computational theory that is relevant to such systems. The field of algorithmic tile assembly provides one such theoretical framework targetted specifically at molecular self-assembling systems. One of most well-studied models of computation for molecular self-assembly systems is the abstract tile assembly model, put forward by Winfree [26]. The model describes crystal-like growth process where, starting from a small connected arrangement of square tiles, called a seed assembly, a growth process takes place where other unit-size square tiles stick to the ever-larger growing assembly. Local rules specify which tiles can stick at each location along the boundary of the assembly. Growth happens asynchronously and in parallel; the model is a kind of asynchronous nondeterministic cellular automaton. Winfree [26] showed that the model can simulate Turing machines, Winfree and Rothemund showed that it can efficiently self-assemble squares [23,24], and Winfree and Soloveichik [25] used bounded-space simulation of time-/space-bounded Turing machines to exhibit for each finite connected shape a Kolmogorov-efficient tile set that assembles a scaled version of that shape. Recently, it has been shown that there is even a single *intrinsically universal* tile set set that faithfully simulates the geometry (shapes) and dynamics of any instance of the model, up to rescaling [8].

These results were all shown for the so-called *cooperative* (or *temperature 2*) model, where tiles bind to the growing assembly if they, or at least some of them, bind on two or more sides. This provides a kind of “context sensitivity” in the growth process. What happens if we allow *noncooperative* (or *temperature 1*) growth where tiles bind if they match on at least one side? Growth like this looks like growing and branching tips in 2D. Tendrils snake out from the seed, possibly crashing into each other, and more often than not they seem to merely form simple structures (cycles and/or repeated path segments), and certainly not the kind of structures needed for computation. Putting proofs behind this intuition has been a challenge and the literature has seen a number of unproven conjectures about the limitations of temperature 1. In this paper, we settle two such questions.

Our first main result is on the topic of simulation in tile assembly. As noted, it has been shown that there is an intrinsically universal tile set for the cooperative model; that is, a tile set is capable of simulating any instance of the cooperative model [8]. More precisely, there is a tile set U that given as input (encoded as a seed assembly) any instance \mathcal{T} of the tile assembly model, tiles from U self-assembly to simulate the geometry (shapes) and dynamics of \mathcal{T} perfectly, modulo a spatial rescaling. By spatial rescaling we mean that each unit-sized square tile in \mathcal{T} is simulated by an $m \times m$ square block of tiles over U . The result is a kind of completeness result for the abstract tile assembly model: the tile set U is “hard” for all tile assembly systems in the sense it is able to capture all possible production and dynamics of all systems, and of course every instantiation of U is itself also a valid tile assembly system. Since then, it has been shown [20] that the noncooperative tile assembly model can not simulate the cooperative model but it was left open (Conjecture 1.4 [20]) whether the noncooperative model can simulate itself. So although the noncooperative model is weak, perhaps it is just strong enough for self-simulation? In other words, is there a noncooperative tile set that is “hard” for the noncooperative model? We answer this conjecture showing that there is no such intrinsically universal tile set for the noncooperative model.

Our second main result is on computation in the noncooperative model. We show, that it is impossible to simulate a time-bounded Turing machine in a bounded rectangular region of the plane in the noncooperative model (see Theorem 1.2 for the formal statement). Indeed, this result implies that the noncooperative model can not simulate Turing machines using any method with a geometry remotely similar to any of the known ways to simulate Turing machines in any known tile assembly model [3, 5, 6, 11, 14, 15, 21–27].¹ It is important to note that the negative result about simulation in ref [20] does not say anything about computation in the model; in fact that particular negative result also holds in the 3D noncooperative model, despite the fact that model can simulate Turing machines. It is also important to point out that many generalisations (enumerated below) of the classical 2D noncooperative model can indeed carry out “bounded” simulation of Turing machines; thus our result formally separates these generalised noncooperative models from the classical 2D model.

New tools and future work on noncooperative tile assembly. Besides showing limitations on noncooperative growth in terms of simulation and computation power, we contend that this paper brings some new techniques to the table. Generally in tile assembly systems, in order to carry out nontrivial computation for finite or infinite shape-building one often has the goal of building structures that (a) are large but (b) not too large (e.g. neither hardcoding a small shape nor filling the entire plane could be reasonably regarded as interesting computation—the interesting *algorithmic* stuff lies in-between). In this paper we provide two tools to analyse, and prove negative results on, building such shapes in the noncooperative model. The first is a method to show that any *any* path of tiles P that travels a long enough horizontal distance while staying above some horizontal line can be either pumped forever or else blocked by growing something else. Hence if P was supposed to form part of some interesting shape, then our first tool (Lemma 5.10) makes it so that we can use P to make another path that either goes outside the shape (P is pumpable), or else prevents P from growing to completion (P is blocked). This contrasts with previous works, e.g. [10, 18], by *using* non-pumpable paths to prove that other assemblies are also possible. In fact one of the main new ideas in our work is to prove strong properties on non-pumpable paths.

Our second tool (Theorem 6.1) builds on this to simultaneously block multiple paths, despite the fact they may interact with each other in very complicated ways. More precisely, given a set of paths of tiles, we define a total ordering on those paths so that we can iteratively apply the first tool to infinitely pump and/or block all of the paths. See Section 4 for a proof overview. Another contribution of this work that might prove useful in the future is a set of definitions (Section 2.3) and lemmas (Section 5.1) that capture a number of basic properties about producible paths at temperature 1, which in turn allow us to frequently use reasoning that is at the abstraction level of *paths in the plane* as opposed to the more low-level of individual tiles and glues. Our conventions and tricks for reasoning about paths of tiles via embeddings in \mathbb{R}^2 could be applied to a variety of models. Together this collection of tools allow us to disrupt any attempt to build shapes of a certain kind, and they work whether or not nondeterminism is deployed as a tool by the ill-fated programmer. We hope these ideas may find use independently of the two main problems we solve here.

The main goal of this paper is to give a negative result on a model that has resisted multiple attempts thus far. Nevertheless, we believe our work to open up new directions for research on the noncooperative model. For example, there are a large number of papers on temperature 1

¹I.e. by simulating a time $t(n)$ and space $s(n)$ -bounded Turing machine in a $O(f(t(n))) \times O(g(s(n)))$ region for finite functions f and g .

models that are generalizations of the classical temperature 1 model that we study (references below). Since those models achieve Turing universality, naïve application of our techniques to those models is *provably impossible*. But often we care more about shape-building than computation and our techniques give a method to edit producible shapes, hence we ask: Can our techniques be generalised to show limitations to the classes of shapes producible in those models? Another question: Is there a non-trivial hierarchy of simulation power within the noncooperative model? We leave this as an open research direction to further clarify and investigate the power of noncooperative self-assembly,² that would certainly require new techniques beyond what we’ve seen to date and beyond what appears in this paper.

Intrinsic universality, and simulation between tile assembly systems, is giving rise to a kind of complexity theory for comparing models of self-assembly [28]. It is interesting to note that in this setting sometimes it is possible to prove negative results on the *simulation power* of models that are already known to be Turing universal [7, 15, 20]. Here we show that one can obtain a negative result on *Turing machine-style computation itself*, via a negative result on simulation between tile assembly systems. Hence we show (for the first time) that simulation between tile assembly systems is a new method to obtain negative results on Turing computation in tile assembly.

Relationship to other work. A large number of papers have conjectured or discussed that in one sense or another, sophisticated computation such as Turing machine simulation or building shapes with few tile types is impossible in the noncooperative model [1, 4, 5, 10, 18, 22–24]. Our Theorem 1.2 implies that any claimed simulator of Turing machines by noncooperative (temperature 1) systems would have to look very different from the known methods for cooperative abstract tile assembly model [24–26] and its generalizations such as the two-handed [3, 7] or polygon [6, 13] models, as well as all known generalizations of the noncooperative model (3D [5], allowing error [5, 23], negative glues [22], staged and stepwise assembly [1], active signals [16, 21], polyominoes [11, 15] and polygons [14]).

Rothmund and Winfree [24] gave the first negative result on 2D temperature 1 systems: building an $N \times N$ square requires N^2 tile types if we insist that the square is *fully connected*. They conjecture this holds in the absence of that assumption. Mañuch, Stacho, and Stoll [18] show that 2D temperature 1 systems *without mismatches* require at least $2N - 1$ tile types to uniquely self-assemble $N \times N$ squares. Doty, Patitz and Summers [10] conjecture that every 2D noncooperative system is *pumpable* meaning, roughly speaking, that every sufficiently long path has a segment that can be producibly repeated infinitely often. Their paper shows that if this conjecture holds then certain forms of computation (e.g. infinite computation) are impossible at temperature 1 in 2D. We remark that proving that conjecture would not imply our main results which are concerned with bounded (finite) computation, nor do our results imply that temperature 1 systems are pumpable. Also here our negative results do not make any assumptions about pumpability nor mismatches. As already noted, it has been shown [20] that the noncooperative tile assembly model can not simulate the cooperative model, here we answer the main open question from that paper (Conjecture 1.4 [20]).

Meunier [19] gives positive results for 2D noncooperative systems. First, by showing the existence of relatively simple noncooperative tile assembly systems that always build finite assemblies that contain a path where at least one tile type is repeated. A second, more general, construction gives for each $n \in \mathbb{N}$ a noncooperative tile assembly system with n tile types that

²The question is not without merit, as recently it was shown that the two-handed, or hierarchical, model of self-assembly has an infinite set of hierarchies with each level in the hierarchy more power than the one below [7].

builds terminal assemblies of height $2n - o(n)$. So although general-purpose computation seems impossible in any reasonable form at temperature 1, we know that one form of algorithmic self-assembly is possible, namely building long paths by re-using tile types.

One of the main reasons one simulates Turing machines with tile assembly systems is to build shapes. Theorem 1.2 shows that none of the standard ways to make shapes in models that are generalisations of the noncooperative model can possibly work in the noncooperative model itself. This gives one formal sense in which shape building via computation is impossible at temperature 1.

Beyond self-assembly, the combinatorics of self-avoiding walks in the plane, first introduced by Flory in 1953 [12] in the context of polymer chemistry, has provided long-standing open problems attracting attention from mathematicians and computer scientists [2, 17]. Our setting and results can be interpreted as memory-bounded versions of this topic: indeed, noncooperative self-assembly is exactly the process of building self-avoiding paths in \mathbb{Z}^2 , but with a *memory* encoded by tile types. Our techniques could be applied to that domain to shed an algorithmic light on the problem of counting or sampling self-avoiding walks.

1.1 Results

We give an overview of our two main results, although a number of notions have yet to be formally defined (see Section 2 for definitions). Our first main result shows that the noncooperative abstract tile assembly model is not intrinsically universal:

Theorem 1.1. *The noncooperative abstract tile assembly model is not intrinsically universal. In other words, there is no tileset U that at temperature 1 simulates all noncooperative tile assembly systems.*

The intuition behind the proof is given in Section 4, and the proof is given in Sections 5 and 6.

Our second main result, that is almost immediate from our main theorem, shows that temperature 1 systems are severely limited in their ability to simulate Turing machines. The standard published methods to simulate Turing machines in 2D in the abstract tile assembly model and its generalizations [3, 5, 6, 11, 14, 15, 21–27], are (or can be easily modified to be) such that simulation of a $s(n)$ space bounded, and $t(n)$ time bounded Turing machine M can be achieved in a $O(s(n)) \times O(t(n))$ rectangle with (a) a seed assembly (encoding M, x) contained in the leftmost $O(1)$ columns, (b) an output assembly (encoding the output of M on input x) that includes a unique tile type appearing on the rightmost column, and (c) no tile ever goes outside this rectangle. The following theorem states, in a formal way, that simulating Turing machines in a bounded rectangular region without error is impossible for the 2D noncooperative abstract tile assembly model, for deterministic or even nondeterministic tile assembly systems. In the theorem statement it is important to note that the “bounding function” B_M is arbitrary in the sense that it allows a potential simulator tile assembly system to use much more space than the actual running time or space usage of the Turing machine M ; this generality serves to strengthen the theorem statement (e.g. bounded Turing machine simulation is impossible even if we allow the tile assembly system to use, say, exponential, doubly exponential, or any finite spatial scaling).

Theorem 1.2. *Let $t : \mathbb{N} \rightarrow \mathbb{N}$, $s : \mathbb{N} \rightarrow \mathbb{N}$ and let $B_M : \mathbb{N} \rightarrow \mathbb{N}$ such that $\forall n \in \mathbb{N}$, $B_M(n) \geq s(n)$. Let M be any Turing machine that halts on all inputs $x \in \{0, 1\}^*$ in time $t(|x|)$ using space $s(|x|)$. There is no pair (V, B_M) where V is a tileset and B_M is a function such that for all $x \in \{0, 1\}^*$, $|x| = n$, there is a seed assembly $\sigma_{M,x}$ and tile assembly system $\mathcal{V}_x = (V, \sigma_{M,x}, 1)$ such that:*

1. $\text{dom}(\sigma_{M,x}) \subseteq \{0, 1, \dots, B_M(n) - 1\} \times \{0, 1, \dots, B_M(n) - 1\}$
2. for all $\alpha \in \mathcal{A}_{\square}[\mathcal{V}]$, $\text{dom}(\alpha) \subseteq \{0, 1, \dots, t(n)B_M(n) - 1\} \times \{0, 1, \dots, B_M(n) - 1\}$, $\text{dom}(\alpha) \cap (\{b + 1, b + 2, \dots, b + B_M(n) - 1\} \times \{0, 1, \dots, B_M(n) - 1\}) \neq \emptyset$ where $b = B_M(n)(t(n) - 1)$ and α places at least one occurrence of a special tile type $H \in V$ on the rightmost column, and nowhere else, of $\text{dom}(\alpha)$ if and only if M accepts x .

The formalism simply states that there is no tile set V , such that when V is instantiated as a noncooperative (temperature 1) tile assembly system $\mathcal{V}_x = (V, \sigma_{M,x}, 1)$, with an input seed assembly $\sigma_{M,x}$ (that somehow encodes a Turing machine M and its input x), then \mathcal{V}_x simulates M on x within a finite rectangular region, writing a yes/no answer anywhere on the rightmost column of tiles. Since the “bounding function” B_M in the theorem statement can be arbitrarily large, the theorem holds even if we allow the noncooperative system to use an *arbitrarily large*, but finite, rectangular bounding box for the simulation. Section 4 gives an intuitive overview of the proof, and the actual proof is given in Section 7.

2 Definitions and preliminaries

Let \mathbb{Z} be the integers, $\mathbb{Z}^+ = \{1, 2, 3, \dots\}$ and $\mathbb{N} = \{0, 1, 2, 3, \dots\}$.

When referring to the relative placements of positions in the grid graph of \mathbb{Z}^2 , or in the plane \mathbb{R}^2 , we say that a position $P = (x_P, y_P)$ is *to the right of* (respectively, *to the left of*, *above*, *below*) of another position $Q = (x_Q, y_Q)$ if $x_P \geq x_Q$ (respectively $x_P \leq x_Q$, $y_P \geq y_Q$, $y_P \leq y_Q$). This definition should not be confused with the definitions of right and left turns, nor with the definition of right-hand side and left-hand side, all of which are defined below.

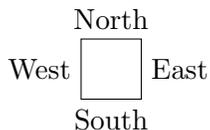
Moreover, unless stated otherwise, vectors of \mathbb{Z}^2 and \mathbb{R}^2 are *column vectors*, i.e. $\vec{u} = \begin{pmatrix} x_u \\ y_u \end{pmatrix}$.

2.1 Abstract tile assembly model

The abstract tile assembly was introduced by Winfree [26]. In this paper we study a restriction of the abstract tile assembly model called the temperature 1 abstract tile assembly model, or noncooperative abstract tile assembly model. For definitions of the full model, as well as intuitive explanations, see for example [23, 24].

A *tile type* is a unit square with four sides, each consisting of a glue *type* and a nonnegative integer *strength*. Let T be a finite set of tile types. In all sets of tile types used in this paper, we assume the existence of a well-defined total ordering that we call the *canonical ordering*.

The sides of a tile type are respectively called north, east, south, and west, as shown in the following picture:



An *assembly* is a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T$ where T is a set of tile types and the domain of α (denoted $\text{dom}(\alpha)$) is connected.³ A *tile* is a pair $((x, y), t) \in \mathbb{Z}^2 \times T$ where (x, y) is a *position*

³Intuitively, an assembly is a positioning of unit-sized tiles, each from some set of tile types T , so that their centers are placed on (some of) the elements of the discrete plane \mathbb{Z}^2 and such that those elements of \mathbb{Z}^2 form a connected set of points.

and t is a tile type. Hence the elements of an assembly are tiles. We let \mathcal{A}^T denote the set of all assemblies over the set of tile types T . In this paper, two tile types in an assembly are said to *bind* (or *interact*, or are *stably attached*), if the glue types on their abutting sides are equal, and have strength ≥ 1 . An assembly α induces a weighted *binding graph* $G_\alpha = (V, E)$, where $V = \text{dom}(\alpha)$, and there is an edge $\{a, b\} \in E$ if and only if a and b interact, and this edge is weighted by the glue strength of that interaction. The assembly is said to be τ -stable if every cut of G has weight at least τ .

A *tile assembly system* is a triple $\mathcal{T} = (T, \sigma, \tau)$, where T is a finite set of tile types, σ is a τ -stable assembly called the *seed*, and $\tau \in \mathbb{N}$ is the *temperature*. Throughout this paper, $\tau = 1$.

Given two τ -stable assemblies α and β , we say that α is a *subassembly* of β , and write $\alpha \sqsubseteq \beta$, if $\text{dom}(\alpha) \subseteq \text{dom}(\beta)$ and for all $p \in \text{dom}(\alpha)$, $\alpha(p) = \beta(p)$. We also write $\alpha \rightarrow_1^T \beta$ if we can obtain β from α by the binding of a single tile type, that is: $\alpha \sqsubseteq \beta$, $|\text{dom}(\beta) \setminus \text{dom}(\alpha)| = 1$ and the tile type at the position $\text{dom}(\beta) \setminus \text{dom}(\alpha)$ stably binds to α at that position. We say that γ is *producible* from α , and write $\alpha \rightarrow^T \gamma$ if there is a (possibly empty) sequence $\alpha_1, \alpha_2, \dots, \alpha_n$ where $n \in \mathbb{N} \cup \{\infty\}$, $\alpha = \alpha_1$ and $\alpha_n = \gamma$, such that $\alpha_1 \rightarrow_1^T \alpha_2 \rightarrow_1^T \dots \rightarrow_1^T \alpha_n$. A sequence of $n \in \mathbb{N}^+ \cup \{\infty\}$ assemblies $\alpha_0, \alpha_1, \dots$ over \mathcal{A}^T is a *\mathcal{T} -assembly sequence* if, for all $1 \leq i < n$, $\alpha_{i-1} \rightarrow_1^T \alpha_i$.

The set of *productions*, or *producible assemblies*, of a tile assembly system $\mathcal{T} = (T, \sigma, \tau)$ is the set of all assemblies producible from the seed assembly σ and is written $\mathcal{A}[\mathcal{T}]$. An assembly α is called *terminal* if there is no β such that $\alpha \rightarrow_1^T \beta$. The set of all terminal assemblies of \mathcal{T} is denoted $\mathcal{A}_\square[\mathcal{T}]$.

As mentioned, in this paper $\tau = 1$. Also throughout this paper, we make the simplifying assumption that all glue types have strength 0 or 1: it is not difficult to see that this assumption does not change the behavior of the model (if a glue type g has strength $s_g \geq 1$, in the $\tau = 1$ model then a tile with glue type g binds to a matching glue type on an assembly border irrespective of the exact value of s_g).

2.2 Simulation between tile assembly systems and intrinsic universality

To state our main result, we must formally define what it means for one tile assembly system to “simulate” another. A number of definitions of simulation have been put forward for various self-assembly models [6–9, 11, 20], here and in Appendix A we use those from [20].

Let T be a tile set, and let $m \in \mathbb{Z}^+$. An *m -block supertile* over T is a partial function $\alpha : \mathbb{Z}_m^2 \dashrightarrow T$, 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 $(x, y) \in \mathbb{Z}^2$, define $\alpha_{(x,y)}^m$ to be the m -block supertile defined by $\alpha_{(x,y)}^m(x', y') = \alpha(mx+x', my+y')$ for all $x', y' \in \{0, 1, \dots, m-1\}$. For some tile set S , 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*⁴ $R^* : \mathcal{A}^S \rightarrow \mathcal{A}^T$ such that $R^*(\alpha') = \alpha$ if and only if $\alpha(A) = R(\alpha'_A)$ for all $A \in \mathbb{Z}^2$. 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 α'_A , $A + \vec{u} \in \text{dom}(\alpha)$ for some $\vec{u} \in \mathbb{Z}^2$ such that $\|\vec{u}\|_2 \leq 1$. In other words, α' may have tiles on supertile blocks representing empty

⁴Note that R^* is a total function since every assembly of S represents *some* assembly of T ; the functions R and α are partial to allow undefined points to represent empty space.

space in α , but only if that position is adjacent to a tile in α . We call such growth “around the edges” of α' *fuzz* and thus restrict it to be adjacent to only valid supertiles, but not diagonally adjacent (i.e. we do not permit *diagonal fuzz*).

Below, let $\mathcal{T} = (T, \sigma_T, \tau_T)$ be a tile assembly system, let $\mathcal{S} = (S, \sigma_S, \tau_S)$ be a tile assembly system, and let R be an m -block representation function $R : B_m^S \rightarrow T$.

Definition 2.1. We say that \mathcal{S} and \mathcal{T} have *equivalent terminal shapes* (under R) if $\{\text{dom}(R^*(\alpha)) \mid \alpha \in \mathcal{A}_\square[\mathcal{S}]\} = \{\text{dom}(\beta) \mid \beta \in \mathcal{A}_\square[\mathcal{T}_n]\}$.

Our main negative result on simulation (Theorem 1.1) shows that any claimed intrinsically universal noncooperative tileset U does not satisfy Definition 2.1 when used to simulate certain noncooperative tile assembly systems. Intrinsically universal tilesets must satisfy Definition 2.1 (see Observation A.2) and moreover must satisfy a significantly stronger set of definitions than Definition 2.1; such stronger definitions are given in Appendix A.

2.3 Paths and non-cooperative self-assembly

This definition sections introduces quite a number of key definitions and concepts that will be used extensively throughout the paper.

Let $n \in \mathbb{N}$ and let T be a set of tile types. As already defined in Section 2.1, a *tile* is a pair $((x, y), t)$ where $(x, y) \in \mathbb{Z}^2$ is a position and $t \in T$ is a tile type.

Intuitively, a path is a finite or one-way-infinite simple (non-self-intersecting) sequence of tiles placed on points of \mathbb{Z}^2 so that each tile in the sequence interacts with the previous one, or more precisely:

Definition 2.2 (Path). A *path* is a (finite or infinite) sequence $P = P_0P_1P_2 \dots$ of tiles $P_i = ((x_i, y_i), t_i) \in \mathbb{Z}^2 \times T$, such that:

- for all P_j and P_{j+1} defined on P it is the case that t_j and t_{j+1} interact, and
- for all P_j, P_k such that $j \neq k$ it is the case that $(x_j, y_j) \neq (x_k, y_k)$.

Whenever P is finite, i.e. $P = P_0P_1P_2 \dots P_{n-1}$ for some n , n is termed the *length* of P . By definition, paths are simple (or self-avoiding), and this fact will be repeatedly used through the paper. A *position* of P is an element of \mathbb{Z}^2 that appears in P (and therefore appears exactly once), and an *index* i of P is simply an integer in $\{0, 1, \dots, n-1\}$. For a path $P = P_0 \dots P_i P_{i+1} \dots P_j \dots$, we define the notation $P_{i,i+1,\dots,j} = P_i P_{i+1} \dots P_j$, i.e. “the subpath of P between indices i and j , inclusive”.

Although a path is not an assembly, we know that each adjacent pair of tiles in the path sequence interact implying that the set of path positions forms a connected set in \mathbb{Z}^2 and hence every path uniquely represents an assembly containing exactly the tiles of the path, more formally: For a path $P = P_0P_1P_2 \dots$ we define the set of tiles $\text{asm}(P) = \{P_0, P_1, P_2, \dots\}$ which we observe is an assembly⁵ and we call $\text{asm}(P)$ a *path assembly*. A path P is said to be *producible* by some tile assembly system $\mathcal{T} = (T, \sigma, 1)$ if the assembly $(\text{asm}(P) \cup \sigma) \in \mathcal{A}[\mathcal{T}]$ is producible, and we call such a P a *producible path*. We define

$$\mathbf{P}[\mathcal{T}] = \{P \mid P \text{ is a path and } (\text{asm}(P) \cup \sigma) \in \mathcal{A}[\mathcal{T}]\}$$

⁵I.e. $\text{asm}(P)$ is a partial function from \mathbb{Z}^2 to tile types that is defined on a connected set.

to be the set of producible paths of \mathcal{T} .⁶

For any path $P = P_0P_1P_2, \dots$ and integer $i \geq 0$, we write $\text{pos}(P_i) \in \mathbb{Z}^2$, or $(x_{P_i}, y_{P_i}) \in \mathbb{Z}^2$, for the position of P_i and $\text{type}(P_i)$ for the tile type of P_i . Hence if $P_i = ((x_i, y_i), t_i)$ then $\text{pos}(P_i) = (x_{P_i}, y_{P_i}) = (x_i, y_i)$ and $\text{type}(P_i) = t_i$.

If two paths, or two assemblies, or a path and an assembly, share a common position we say they *intersect* at that position. Furthermore, we say that two paths, or two assemblies, or a path and an assembly, *agree* on a position if they both place the same tile type at that position and *conflict* if they place a different tile type at that position.

Note that, since the domain of a producible assembly is a connected set in \mathbb{Z}^2 , and since in an assembly sequence of some TAS $\mathcal{T} = (T, \sigma, 1)$ each tile binding event $\beta_i \rightarrow_1^{\mathcal{T}} \beta_{i+1}$ adds a single node v to the binding graph G_{β_i} of β_i to give a new binding graph $G_{\beta_{i+1}}$, and adds at least one weight-1 edge joining v to the subgraph $G_{\beta_i} \in G_{\beta_{i+1}}$, then for any tile $((x, y), t) \in \alpha$ in a producible assembly $\alpha \in \mathcal{A}[\mathcal{T}]$, there is a edge-path (sequence of edges) in the binding graph of α from σ to $((x, y), t)$. From there, the following important fact about temperature 1 tile assembly is straightforward to see.

Observation 2.3. *Let $\mathcal{T} = (T, \sigma, 1)$ be a tile assembly system and let $\alpha \in \mathcal{A}[\mathcal{T}]$. For any tile $((x, y), t) \in \alpha$ there is a producible path $P \in \mathbf{P}[\mathcal{T}]$ that for some $i \in \mathbb{N}$ contains $P_i = ((x, y), t)$.*

For $A, B \in \mathbb{Z}^2$, we define $\overrightarrow{AB} = B - A$ to be the vector from A to B , and for two tiles $P_i = ((x_i, y_i), t_i)$ and $P_j = ((x_j, y_j), t_j)$ we define $\overrightarrow{P_iP_j} = \text{pos}(P_j) - \text{pos}(P_i)$ to mean the vector from $\text{pos}(P_i) = (x_i, y_i)$ to $\text{pos}(P_j) = (x_j, y_j)$. The translation of a path P by a vector $\vec{v} \in \mathbb{Z}^2$, written $P + \vec{v}$, is the path Q where and for all indices i of P , $\text{pos}(Q_i) = \text{pos}(P_i) + \vec{v}$ and $\text{type}(Q_i) = \text{type}(P_i)$. As a convenient notation, for a path PQ composed of subpaths P and Q , when we write $PQ + \vec{v}$ we mean $(PQ) + \vec{v}$ (i.e. the translation of all of PQ by $+\vec{v}$). The translation of a path P by a vector $\vec{v} \in \mathbb{Z}^2$, written $P + \vec{v}$, is the path Q where and for all indices i of P ,

The translation of an assembly α by a vector \vec{v} , written $\alpha + \vec{v}$, is the assembly β defined on the set $\text{dom}(\alpha) + \vec{v}$ as $\beta(x, y) = \alpha((x, y) - \vec{v})$ where $(x, y) \in \mathbb{Z}^2$. A *column* $x \in \mathbb{Z}$ is the set of all points of \mathbb{Z}^2 with x-coordinate x , and a *row* $y \in \mathbb{Z}$ is the set of all points of \mathbb{Z}^2 with y-coordinate y .

Next, for a path P and two indices i, j on P , we will define a (not necessarily producible) sequence called the pumping of P between i and j .

Definition 2.4 (pumping of P between i and j). Let $\mathcal{T} = (T, \sigma, 1)$ be a tile assembly system and $P \in \mathbf{P}[\mathcal{T}]$. We say that the “*pumping of P between i and j* ” is the sequence \bar{q} of elements from $\mathbb{Z}^2 \times T$ defined by:

$$\bar{q}_k = \begin{cases} P_k & \text{for } 0 \leq k \leq i \\ P_{i+1+((k-i-1) \bmod (j-i))} + \lfloor (k-i-1)/(j-i) \rfloor \overrightarrow{P_iP_j} & \text{for } i < k, \end{cases}$$

Hence, intuitively, \bar{q} has two parts. It begins with a finite sequence $P_{0,1,\dots,i}$. Then appended to that, there is an infinite sequence where the tile types appear with positions at regular intervals in the plane. We formalize the latter intuition in the following Lemma:

⁶Intuitively, although producible paths are not assemblies, any producible path P has the nice property that it encodes an unambiguous description of how to grow $\text{asm}(P)$ from the seed σ , in (P) path order, to produce the assembly $\sigma \cup \text{asm}(P)$.

Lemma 2.5. *Let P be a path with tiles from some tiling set T , $i < j$ be two integers, and q be the pumping of P between i and j . Then for all integer $k \geq i$, $q_{k+(j-i)} = q_k + \overrightarrow{P_i P_j}$.*

Proof. By the definition of q :

$$\begin{aligned}
q_{k+(j-i)} &= P_{i+1+((k+(j-i)-i-1) \bmod (j-i))} + \left\lfloor \frac{k+(j-i)-i-1}{j-i} \right\rfloor \overrightarrow{P_i P_j} \\
&= P_{i+1+((k-i-1) \bmod (j-i))} + \left\lfloor \frac{k-i-1}{j-i} + 1 \right\rfloor \overrightarrow{P_i P_j} \\
&= P_{i+1+((k-i-1) \bmod (j-i))} + \left(\left\lfloor \frac{k-i-1}{j-i} \right\rfloor + 1 \right) \overrightarrow{P_i P_j} \\
&= q_k + \overrightarrow{P_i P_j}
\end{aligned}$$

□

The following definition gives the notions of pumpable and finitely pumpable that are used in our proofs. It is followed by a less formal but more intuitive description.

Definition 2.6 (Pumpable). Let $\mathcal{T} = (T, \sigma, 1)$ be a tile assembly system. We say that a producible path $P \in \mathbf{P}[\mathcal{T}]$, is *infinitely pumpable*, or simply *pumpable*, if there are two integers $i < j$ such that the pumping of P between i and j is a producible (infinite) path, i.e. $\bar{q} \in \mathbf{P}[\mathcal{T}]$.

In other, more intuitive, words, a producible path $P \in \mathbf{P}[\mathcal{T}]$ is *infinitely pumpable*, or simply *pumpable*, if there is a producible infinite assembly $\alpha \in \mathcal{A}[\mathcal{T}]$ and two indices $i < j$ on P , such that α contains exactly σ , then $\text{asm}(P_{0,1,\dots,j})$, and then infinitely many occurrences of the “pumpable segment” $\text{asm}(P_{i,i+1,\dots,j-1})$ each translated by successive positive integer multiples $\{1, 2, 3, \dots\}$ of $\overrightarrow{P_i P_j}$, where these occurrences do not intersect σ , $\text{asm}(P_{0,1,\dots,j-1})$ or themselves, each tile along this path assembly is bound to the previous, and α contains no other tiles.⁷

For all i such that both P_i and P_{i+1} are defined (i.e. for all $i \in \mathbb{N}$ if P is infinite, and for all $i < |P| - 1$ otherwise), we define the “*output side of P_i* ” to be the side of $\text{type}(P_i)$ adjacent to $\text{type}(P_{i+1})$, and for all $i > 0$, we define the “*input side of P_i* ” to be the side of $\text{type}(P_i)$ adjacent to $\text{type}(P_{i-1})$. The sides of $\text{type}(P_i)$ that are neither output sides nor input sides of P_i are said to be *free*, as are the glues of those sides.⁸

Let $P = P_0 P_1 \dots$. For $i > 0$, we say that a *right turn* (respectively *left turn*) from P at index i is a path with prefix $P_0 P_1 \dots P_i x$ for some $x \in \mathbb{Z}^2 \times T$ adjacent to P_i such that orientated in the direction $\overrightarrow{P_{i-1} P_i}$, $\text{pos}(x)$ is clockwise (respectively anti-clockwise) from P_{i+1} . More formally, let $\vec{u} = \overrightarrow{P_i P_{i-1}}$ (the unit column vector from $\text{pos}(P_i)$ to $\text{pos}(P_{i-1})$), let $\rho = \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$, and let $\tau = (\rho \cdot \vec{u}, \rho \cdot \rho \cdot \vec{u}, \rho \cdot \rho \cdot \rho \cdot \vec{u})$, then we say that $P_0 P_1 \dots P_i x$ is a right turn from $P_0 P_1 \dots P_i P_{i+1}$ if $\overrightarrow{P_i x}$ appears after $\overrightarrow{P_i P_{i+1}}$ in τ .

⁷We remark that this definition of “infinitely pumpable” intentionally excludes pumping that intersects with and *agrees with* the seed, $P_{0,1,\dots,j}$ or some translated (pumped) segment.

⁸By this definition of input and output sides the first tile of a path does not have an input side, and the last one does not have an output side. We also remark that this definition of input/output sides is defined relative to a specific *path*, and is *not* a property of the tiles themselves; moreover, the tiles, including the first and last tiles, may have other glue types, i.e. *free* glue types, not used by the path. Despite the fact free sides may have strength 1 glue types we typically ignore this in our analysis of paths—this is because our proofs typically analyse paths one at a time and thus require us to consider only the non-free tiles sides that actually bind the tiles along the path assembly and thus don’t require us to make statements about free sides.

We define $\text{glue}(P_i P_{i+1}) = (g, i)$ [i.e. (g, i) is a pair of the form (glue type, path index)] where g is the shared glue type between consecutive tiles P_i and P_{i+1} on the path P . When we say “glue” in the context of a path, we mean a pair of the form (glue type, path index). We define $\text{type}(\text{glue}(P_i P_{i+1})) = g$ to denote the glue type of $\text{glue}(P_i P_{i+1})$, we write $\text{pos}(\text{glue}(P_i P_{i+1})) = (\text{pos}(P_i), \text{pos}(P_{i+1}))$ (the “position of glue $\text{glue}(P_i P_{i+1})$ ”) to denote the edge (of the grid graph of \mathbb{Z}^2) of $\text{glue}(P_i P_{i+1})$, oriented from $\text{pos}(P_i)$ to $\text{pos}(P_{i+1})$. Moreover, for $A, B \in \mathbb{R}^2$ we define $\text{mid}(A, B) = A + \frac{1}{2}\overrightarrow{AB} \in \mathbb{R}^2$ to be the midpoint of the line segment $[A, B] \subsetneq \mathbb{R}^2$, and for a pair of tiles P_i, P_j we define $\text{mid}(P_i, P_j)$ to be the midpoint of the line segment $[\text{pos}(P_i), \text{pos}(P_j)] \subsetneq \mathbb{R}^2$.

Definition 2.7 (The right priority path of a set of paths). Let P and Q , where $P \neq Q$, be two paths with $\text{pos}(P_0) = \text{pos}(Q_0)$ and $\text{pos}(P_1) = \text{pos}(Q_1)$. Let i be the smallest index such that $i \geq 0$ and $P_i \neq Q_i$. We say that P is the *right priority path* of P and Q if either (a) $P_{0,1,\dots,i}$ is a right turn from Q or (b) $\text{pos}(P_i) = \text{pos}(Q_i)$ and the type of P_i is smaller than the type of Q_i in the canonical ordering of tile types.

For sets of paths, we extend this definition as follows: let $p_0 \in \mathbb{Z}^2, p_1 \in \mathbb{Z}^2$ be two adjacent positions. If S is a set of paths such that for all $P \in S$, $P_0 = p_0$ and $P_1 = p_1$, we call the *right-priority path of S* the path that is right-priority path of all other paths in S .

The *left priority path of a set of paths* is defined symmetrically: swap left for right in Definition 2.7.

2.3.1 Curves: embedding paths in \mathbb{R}^2

A *curve*, or a *curve in \mathbb{R}^2* , is defined to be a continuous function $f : [0, 1] \rightarrow \mathbb{R}^2$. We say that f is continuous at some $x_0 \in \text{dom}(f)$ if $\forall \varepsilon, \exists \eta, \forall x, |x - x_0| \leq \eta \Rightarrow \|f(x) - f(x_0)\|_2 \leq \varepsilon$, where for all $a, b \in [0, 1]$, $\|(a, b)\|_2 = \sqrt{a^2 + b^2}$, and we say that f is continuous if and only if f is continuous at all $x_0 \in [0, 1]$.

Intuitively, we will define the concatenation of a finite sequence of curves f_0, f_1, \dots, f_{k-1} to be a function $F : [0, 1] \rightarrow \mathbb{R}^2$ that for each $i \in \{0, 1, \dots, k-1\}$ represents f_i by rescaling the domain of f_i to be in the interval $[\frac{i}{k}, \frac{i+1}{k}]$. Thus F is defined on $[0, 1]$ and has range $\bigcup_{i=0}^{k-1} (\text{range}(f_i))$. This is defined as follows:

Definition 2.8 (Concatenation of curves in \mathbb{R}^2). Given a finite sequence of $k \in \mathbb{N}$ curves f_0, f_1, \dots, f_{k-1} in \mathbb{R}^2 their *concatenation* is the function $F : [0, 1] \rightarrow \mathbb{R}^2$ defined for all i such that $0 \leq i < k$ and all $x \in [\frac{i}{k}, \frac{i+1}{k}]$ as $F(x) = f_i(xk - i)$.

For example, Figure 2.1(c) shows the concatenation of two curves: the curve in Figure 2.1(b) and a unit-length vertical line segment.

The following observation states that the concatenation F of k continuous functions f_0, f_1, \dots, f_{k-1} , that have the property $f_i(1) = f_{i+1}(0)$ for $0 \leq i < k-1$, is itself a continuous function and although the proof is straightforward, it is worth explicitly stating since it is used extensively in this paper:

Observation 2.9. *Let f_0, f_1, \dots, f_{k-1} be a finite sequence of curves in \mathbb{R}^2 that have the property that for all $i \in \{0, 1, \dots, k-2\}$, $f_i(1) = f_{i+1}(0)$ and let F be the concatenation of f_0, f_1, \dots, f_{k-1} . Then F is a curve in \mathbb{R}^2 .*

Proof. First note that for each $i \in \{0, 1, \dots, k-1\}$, f_i is a continuous function and that rescaling (shrinking) the domain of f_i from $[0, 1]$ to $[\frac{i}{k}, \frac{i+1}{k}]$ preserves continuity. Secondly, since for each $i \in \{0, 1, \dots, k-2\}$, $f_i(1) = f_{i+1}(0)$ and F contains $\frac{i+1}{k} \rightarrow f_i(1)$, the function F is continuous

on the (“double-length”) interval $[\frac{i}{k}, \frac{i+2}{k}]$. Since this holds for all such i , F is continuous on its entire domain $[0, 1]$, and thus is a curve in \mathbb{R}^2 . \square

Observation 2.10. *Let f_0, f_1, \dots, f_{k-1} be a finite sequence of finite-length simple curves in \mathbb{R}^2 such that for all $i \in \{0, 1, \dots, k-2\}$, $f_i(1) = f_{i+1}(0)$, also $f_{k-1}(1) = f_0(0)$, and those k nonempty intersections between f_0, f_1, \dots, f_{k-1} are the only nonempty intersections between them. Let F be the concatenation of f_0, f_1, \dots, f_{k-1} . Then F is a finite-length closed simple curve in \mathbb{R}^2 .*

Proof. The hypotheses of Observation 2.9 are satisfied hence F is a curve. F is composed of a finite set of k finite length component curves so F is of finite length. F is closed because for $i \in \{0, 1, \dots, k-2\}$, $f_i(1) = f_{i+1}(0)$, and $f_{k-1}(1) = f_0(0)$, and F is simple because those are the only nonempty intersections between f_0, f_1, \dots, f_{k-1} . \square

Also, we will sometimes need other curves that are not defined by paths:

Definition 2.11 (Line segment). Let $A, B \in \mathbb{R}^2$. The *line segment from A to B* , which we write $[A, B]$, is the curve defined for all $x \in [0, 1]$ by $f(x) = A + x\overrightarrow{AB}$.

Definition 2.12. For any path P we define \mathfrak{E}_P to be the *canonical embedding* of P where $\mathfrak{E}_P : [0, 1] \rightarrow \mathbb{R}^2$, such that for all s such that $0 \leq s < 1$

$$\mathfrak{E}_P(s) = \text{pos}(P_{\lfloor s \cdot (|P|-1) \rfloor}) + (s \cdot (|P| - 1) - \lfloor s \cdot (|P| - 1) \rfloor) \overrightarrow{P_{\lfloor s \cdot (|P|-1) \rfloor} P_{\lfloor s \cdot (|P|-1) \rfloor + 1}}$$

and

$$\mathfrak{E}_P(1) = \text{pos}(P_{|P|-1}).$$

Note that by Definition 2.12, the canonical embedding of a path is a curve, i.e. the canonical embedding is a continuous function from $[0, 1]$ to \mathbb{R}^2 . Figure 2.1(a) shows an example path P and Figure 2.1(b) shows its canonical embedding \mathfrak{E}_P .

This paper frequently uses the Jordan curve theorem, which is a statement about curves in \mathbb{R}^2 : any simple closed (and hence finite) curve in \mathbb{R}^2 partitions \mathbb{R}^2 into exactly two connected components, a bounded one and an unbounded one.

In our proofs, we will often reason about right turns and left turns from a curve, and also about on which side of a closed simple curve is the bounded connected component. Since all of the closed simple curves we will define will be simple finite polygons, their left-hand side and right-hand side can be defined by taking any point A on a segment of the polygonal curve c , not at a corner, and reasoning as follows. Since c is locally a straight line around A , c is differentiable at A . Also c has a direction (from domain element 0 to domain element 1). The left-hand side of c is therefore the connected component to the left of A when orientated in the direction from 0 to 1 along c , and the right-hand side of c is the connected component to the right of A . By defining curves within a very small distance of c , we can show that the left-hand side of c is connected, and the right-hand side of c is also connected. For example, Figure 2.1(c) shows such a polygonal closed simple curve c , with its left-hand side highlighted in grey.

3 A family of tile assembly systems \mathcal{T}_N

Definition 3.1 defines a (very simple) infinite family of noncooperative tile assembly systems $\{\mathcal{T}_N \mid N \in \mathbb{Z}^+\}$. The proof of our main theorem shows that there is no tile set U that for all N simulates \mathcal{T}_N .

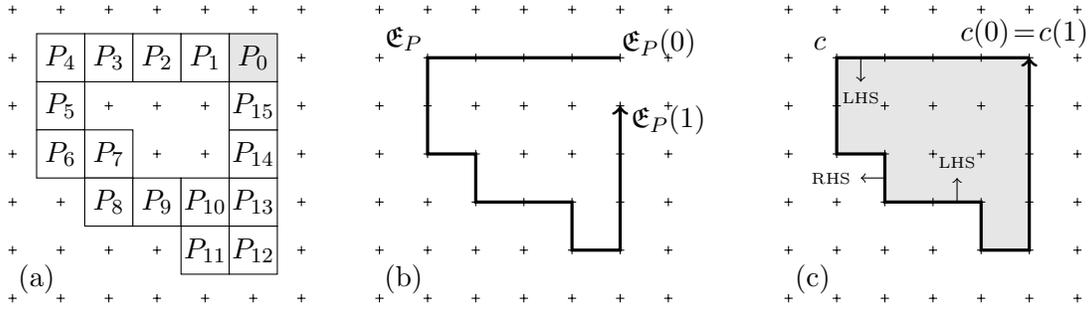


Figure 2.1: Crosses denote points in \mathbb{Z}^2 . (a) A path $P = P_0, P_1, \dots, P_{15}$. (b) The curve \mathfrak{E}_P , the canonical embedding of P in \mathbb{R}^2 . (c) The curve c defined as the concatenation of \mathfrak{E}_P and the unit-length line segment $[\text{pos}(P_{15}), \text{pos}(P_0)] = [\mathfrak{E}_P(1), \mathfrak{E}_P(0)]$. Two endpoints of c are identical, $c(0) = c(1)$, moreover c is a closed simple curve in \mathbb{R}^2 . A single thick arrowhead on c indicates the direction of c and the position of c 's identical endpoints. ‘By the Jordan curve theorem, such a closed simple curve in \mathbb{R}^2 partitions the plane into two connected components exactly one of which is ‘bounded’ (has finite area). The ‘right-hand side’ (RHS) and ‘left-hand side’ (LHS) of c are indicated by small thin arrows with the LHS being the bounded component highlighted in grey.

Definition 3.1. For each $N \in \mathbb{Z}^+$, let $\mathcal{T}_N = (T_N, \sigma_N, 1)$ be the tile assembly system that assembles the infinite assembly shown in Figure 3.1.

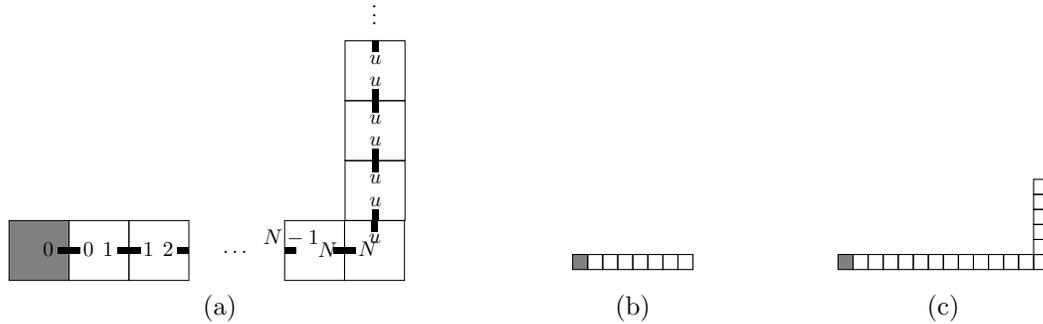


Figure 3.1: (a) The “flipped-L” TAS $\mathcal{T}_N = (T_N, \sigma_N, 1)$, that deterministically assembles a single, infinite assembly. T contains $N + 3$ tile types as shown, and the grey tile is the seed σ_N which is placed at the origin $(0, 0)$. \mathcal{T}_N grows from the seed, distance $N + 1$ to the east, and then grows infinitely to the north. Hence \mathcal{T}_N builds an infinite path assembly. Productions of an example such tile assembly system for $N = 12$ are shown: (b) \mathcal{T}_{12} after 7 tile additions, and (c) \mathcal{T}_{12} after 18 tile additions.

4 Intuition behind the proofs of Theorems 1.1 and 1.2

We begin with a description of the high-level intuition behind the proof of our main result, Theorem 1.1. One of the main difficulties of this result is that for *any* finite number of non-cooperative tile assembly systems, there is in fact a single non-cooperative simulator for all of them: simply let the tiles of the simulator be the disjoint union of all tilesets of the simulated systems. Moreover, it is known [20] that in 3D, there is a tiling set, operating at temperature 1 (i.e. noncooperative), that simulates all non-cooperative tile assembly systems. Hence our proof

is going to crucially make use of the fact that any claimed simulator tileset is of finite size, and must work in the plane.

First, we assume, for the sake of contradiction, that there is a single tileset U , that simulates all noncooperative (or temperature 1) tile assembly systems. Hence, in particular, the tileset U simulates the class of systems $\{\mathcal{T}_N \mid N \in \mathbb{Z}^+\}$ described in Section 3 and shown in Figure 3.1. Hence, for all \mathcal{T}_N , there is a tile assembly system $\mathcal{U}_{\mathcal{T}_N} = (U, \sigma_{\mathcal{T}_N}, 1)$ and scale factor $m \in \mathbb{Z}^+$ such that $\mathcal{U}_{\mathcal{T}_N}$ simulates \mathcal{T}_N .⁹ In particular, by the definition of simulation (Section 2.2), and in particular by Definition 2.1, this implies that for each terminal producible assembly α of \mathcal{T}_N there is a producible assembly $\alpha' \in \mathcal{U}_{\mathcal{T}_N}$ that represents the shape of α , and vice-versa. Figure 4.1 shows what such a simulation should look like.

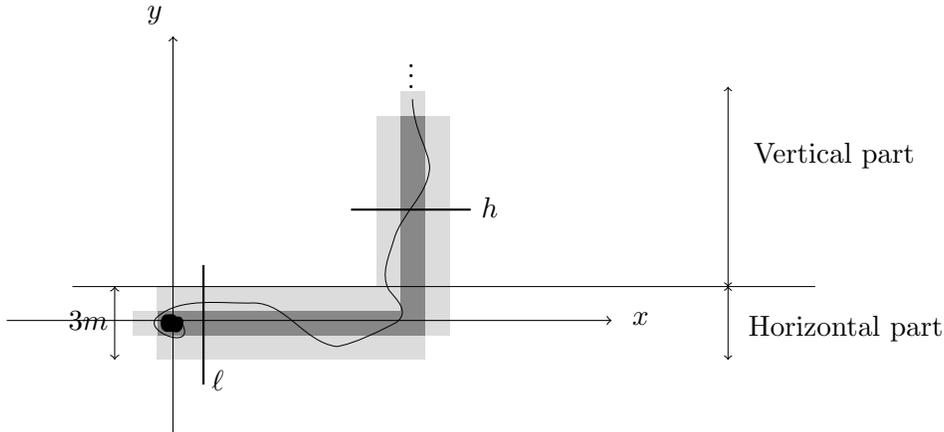


Figure 4.1: An example (claimed) simulation of some \mathcal{T}_N , $N \in \mathbb{Z}^+$, by $\mathcal{U}_{\mathcal{T}_N} = (U, \sigma_{\mathcal{T}_N}, 1)$. The seed assembly of $\mathcal{U}_{\mathcal{T}_N}$ is shown in black on the left. The simulator $\mathcal{U}_{\mathcal{T}_N}$ is free to place tiles anywhere in the *simulation zone*, defined to be the union of the dark gray (tile-representing supertiles) and light gray (fuzz) regions, and an example valid path of tiles (that has a h -successful prefix) is shown in those regions as a thin black curve. The scale factor is m and $\mathcal{U}_{\mathcal{T}_N}$ places $N + 2$ horizontal tiles, hence the width of the dark gray region is $m(N + 2)$. Our main result is proven by showing that any simulator that claims to produce a valid terminal assembly must also produce one with an *incorrect shape*—one that either (a) places tiles outside of the simulation zone (e.g. because some path is infinitely pumpable to the right or can be modified to grow upwards at an incorrect location that is outside of the vertical part), or (b) is finite (e.g. does not grow infinitely upwards). The seed supertile contains the seed (in black) and the origin $(0, 0)$. The lines ℓ (at x -coordinate $|U|(3m + 1) + m + 1$) and h (at y -coordinate $10m$) are used in many of the proofs.

The proof is then broken into two stages. First, in Section 5, we consider any path that can grow in the claimed simulator long enough so it places at least one tile on a horizontal line at some height h . An example such path is shown in Figure 4.1. For any such path, we let P denote its shortest prefix that contains exactly one tile at height $h = 10m$ and where that tile is P 's last tile. Any assemblable path of this form is called “ h -successful” (see Definition 5.7). We show that any h -successful path P can be modified in two different ways: (1) P is modified so that it grows at an invalid position, either by (1.1) *infinite pumping*, by which we mean P can be modified to give another path P' which grows to form an assembly that is infinitely long horizontally to the right, and hence is not a simulation of the “flipped-L” shaped \mathcal{T}_N , or modifying it so that (1.2) the vertical arm grows displaced to the left or to the right (and hence is

⁹Later in the paper we drop the \mathcal{T}_N subscripts from $\mathcal{U}_{\mathcal{T}_N}$, $\sigma_{\mathcal{T}_N}$ and simply say that $\mathcal{U} = (U, \sigma, 1)$ simulates \mathcal{T}_N .

grown in the wrong position) or (1.3) grows something in the wrong position by finite pumping (repeating a path segment that gets blocked). (2) P is blocked, by which we mean *another* assembly (a path R) can be grown that blocks P (forcing P to be of finite length) and thus the simulator can not rely on the growth of path P to obtain a valid simulation. (1) and (2) together show that no *single path* can carry out a valid simulation, which is stated formally in Theorem 5.11. This can be regarded as our first technical tool for handling long paths.

However this leaves open the possibility that many paths could simultaneously and non-deterministically grow and interact in a way that carries out a valid simulation. In particular, in (2) above, if we first block a path P^1 using another path R^1 , and then attempt to block another path P^2 using a path R^2 , then R^2 may itself get blocked by R^1 , hence P^2 could possibly “escape” to become h -successful and carry out the simulation. In the second part of the proof, in Section 6, we consider the ways that h -successful paths P , and the paths R we use to block them, can interact. Based on this, we provide an explicit growth order for paths, such that if we apply (2) to each such path in turn, we guarantee that the claimed simulator fails. This can be regarded as our second technical tool: a method to control the interactions of multiple long paths.

The proof of our second main result Theorem 1.2 is very short. We note that, for the sake of contradiction, if for each time-bounded Turing machine M there is a noncooperative tile set V that simulates M using arbitrarily large 2D space, bounded by a rectangle, then V could be easily modified to build a family of assemblies that are of the same (scaled) shape as systems \mathcal{T}_N , thus contradicting our result that there is no such tile set. Hence we give a reduction from a Turing machine prediction problem to the problem of simulating the class of all \mathcal{T}_N systems, a new technique for proving negative results about computation in tile assembly.

5 Pumping or blocking any sufficiently wide and tall path

Suppose, for the sake of contradiction, that there is a tile set U , such that for $N = 10|U|$, there is a scale factor $m \in \mathbb{Z}^+$ and a seed σ for which $\mathcal{U} = (U, \sigma, 1)$ simulates \mathcal{T}_N at scale factor m . Then it must be the case that there is a path P such that the simulator produces the assembly $\text{asm}(P)$ and $\text{asm}(P)$ grows up to meet a line at height $h = 10m$ above the horizontal arm of the assembly (See Figure 4.1). In this section, we prove that we can use the existence of such a “ h -successful” path P to force the simulator to produce another assembly $\alpha \in \mathcal{A}[\mathcal{U}]$ that either (I) illegally places tiles outside of the simulation zone or else (II) is finite and blocks P from growing (i.e. $\alpha \cup \text{asm}(P)$ is not producible from the assembly $\alpha \in \mathcal{A}[\mathcal{U}]$). Case (I) contradicts that \mathcal{U} simulates \mathcal{T} and we are done with the proof in that case. Case (II) gives a way to block the arbitrary single path P , but may not prevent other paths from growing and this is handled later in Section 6.

Since showing impossibility of simulating a single tile assembly system \mathcal{T}_N is sufficient to prove our result, we will set $N = 10|U|$ in the rest of the paper, and call σ the seed and m the scale factor at which our claimed simulator $\mathcal{U} = (U, \sigma, 1)$ simulates $\mathcal{T}_N = \mathcal{T}_{10|U|}$.

5.1 Glue visibility: definitions and basic results about V_P^+ and V_P^-

We begin this section with a definition that will be used in every proof in the rest of the paper, and is illustrated in Figure 5.1. This notion of visibility is *from the south*.

Definition 5.1 (Visible glue (from the south)). Let P be a path and let $i, j \in \{0, 1, \dots, |P| - 2\}$.

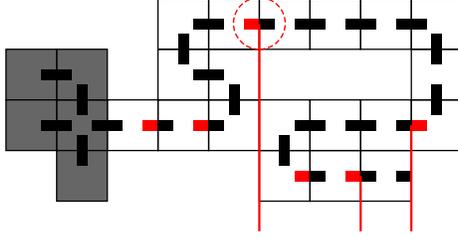


Figure 5.1: A short path with six glues that are visible (from the south). Growth of this path begins from the seed which is shown in grey. Glues here are coloured not by their type, but by their visibility (and direction). Five red-black glues are in V_P^+ and one black-red glue (to the right) is in V_P^- giving a total of six glues in V_P . None of the other glues (black-black, or black-) are visible. Three red rays to the south testify to the visibility of three of the glues, and are called the “visibility rays” of these glues.

Let $\text{glue}(P_i P_{i+1}) = (g, (x, y), d)$ where g is the interacting glue type between P_i and P_{i+1} , $(x, y) = \text{mid}(P_i, P_{i+1})$ and $d = \overrightarrow{P_i P_{i+1}}$ ¹⁰. We say that $\text{glue}(P_i P_{i+1})$ is *visible relative to* P if (a) there is a vertical ray r in \mathbb{R}^2 that starts from $\text{mid}(P_i, P_{i+1}) \in \mathbb{R}^2$ and goes infinitely to the south and (b) for all $j \neq i$, r does not contain $\text{mid}(P_j, P_{j+1}) \in \mathbb{R}^2$.

We write V_P for the set of (glue type, path index) pairs of P that are visible relative to some path P . We define V_P^+ (respectively, V_P^-) to be the set of those (glue type, path index) pairs of V that are on *east* (respectively, *west*) of output sides of tiles of P . Since, on a path P , each tile has exactly one output side, it follows that (V_P^+, V_P^-) is a partition of V_P .

For a path P and a vertical line $l = \{(x, y) \mid y \in \mathbb{R}\}$ at some position $x \in \{\frac{x'}{2} \mid x' \in \mathbb{Z}\}$ such that $l \cap \mathfrak{E}_P \neq \emptyset$, we write “ P ’s visible glue on l ” to mean the unique glue $\text{glue}(P_i P_{i+1})$ such that $\text{mid}(P_i, P_{i+1}) \in \mathbb{R}^2$ is on l and for all j , such that $0 \leq j \leq |P| - 2$, $j \neq i$ it is not the case that $\text{mid}(P_j, P_{j+1})$ has a smaller y -coordinate on l than $\text{mid}(P_i, P_{i+1})$. Moreover “the position of the visible glue of P on l ” is the point $\text{mid}(P_i, P_{i+1})$. For a glue $\text{glue}(P_i P_{i+1})$ that is visible relative to P we write “the position of $\text{glue}(P_i P_{i+1})$ ” to mean the point $\text{mid}(P_i, P_{i+1}) \in \mathbb{R}^2$. For a glue $\text{glue}(P_i P_{i+1})$ that is visible relative to P we write “the visibility ray of $\text{glue}(P_i P_{i+1})$ ” to mean the infinite ray starting at the point $\text{mid}(P_i, P_{i+1}) \in \mathbb{R}^2$ and going vertically to the south.

It is important to note that “visible glue” is defined relative to a particular path. Hence, we often say that a glue of P is “visible *relative to* P ”, although when P is clear from the context, we may simply say that a glue is “visible”. Figure 5.1 gives examples of glues on a path that are and are not visible. Intuitively, note that although a visibility ray starts at the “position” of a visible glue on the path P , the ray is permitted to “touch” other free glues (Section 2.3) on tiles of P (recall free glues are by definition not on the path since they are not input/output glues along the path, hence they are not visible nor can they prevent some other glue from being visible). See Figure 5.1 for examples.

In the following lemma we show that for a path P that has both V_P^+ glues and V_P^- glues, the V_P^+ glues are positioned to the right of its V_P^- glues. The intuition behind the proof is as follows: suppose otherwise, then draw a finite length curve c in \mathbb{R}^2 that runs from a V_P^+ glue along the positions of P to a V_P^- glue, then includes segments of the visibility rays (to the south) of these two glues and finally includes a horizontal line that lies far below P and runs between those two visibility rays. It turns out that c is simple and closed and thus cuts the plane into an

¹⁰Note that since paths are defined to be simple, for any $(x, y) \in \mathbb{Z}^2$, there is at most one index i on P such that (x, y) is the position of $\text{glue}(P_i P_{i+1})$.

unbounded component and a bounded component \mathcal{C} . It turns out that P must both go inside \mathcal{C} for a time and then leave \mathcal{C} , but since c was defined using curves that P never crosses we get a contradiction. This rough intuition is more rigorously formalised in the proof.

Lemma 5.2 (V_P^+ glues are to the right of V_P^- glues). *Let $P \in \mathbf{P}[\mathcal{U}]$ be an assemblable path of any tile assembly system $\mathcal{U} = (U, \sigma, 1)$, and let i and j be two indices such that $\text{glue}(P_i P_{i+1}) \in V_P^+$, $\text{glue}(P_j P_{j+1}) \in V_P^-$. If P has at least one tile P_k where $k > \max(i, j)$ and $\text{pos}(P_k)$ is strictly to the right of, or strictly above, all tiles of $P_{\min(i,j), \min(i,j)+1, \dots, \max(i,j)}$, then $x_{P_i} > x_{P_j}$. In other words, the glues of V_P^+ are all to the right of the glues of V_P^- .*

Proof. We have argued above that (V_P^+, V_P^-) is a partition of V_P . Hence since $\text{glue}(P_i P_{i+1}) \in V_P^+$, $\text{glue}(P_j P_{j+1}) \in V_P^-$ it is the case that $i \neq j$ which in turns implies $x_{P_i} \neq x_{P_j}$. Assume, for the sake of contradiction, that there are two integers i and j that satisfy the lemma hypotheses (in particular that $\text{glue}(P_i P_{i+1}) \in V_P^+$, $\text{glue}(P_j P_{j+1}) \in V_P^-$) but where $x_{P_i} < x_{P_j}$.

Since $\text{glue}(P_i P_{i+1}) \in V_P^+$, there is a vertical ray l_i that starts from the point $\text{mid}(P_i, P_{i+1}) \in \mathbb{R}^2$, goes infinitely to the south, and does not contain any point of the canonical embedding \mathfrak{E}_P of P in \mathbb{R}^2 besides $\text{mid}(P_i, P_{i+1})$. Likewise, let l_j be the vertical ray to the south starting from the point $\text{mid}(P_j, P_{j+1})$, and observe that since $\text{glue}(P_j P_{j+1}) \in V_P^-$, then l_j does not contain any point of \mathfrak{E}_P besides $\text{mid}(P_j, P_{j+1})$. We will use l_i and l_j to define the three line segments s_i^{\leftarrow} , s_j and $s_{i,j}$. First let y_0 be a y -coordinate below all of $P_{i, i+1, \dots, j}$, for instance

$$y_0 = \min\{y_{P_i}, y_{P_{i+1}}, \dots, y_{P_j}\} - 10$$

We then define:

$$s_i = [\text{mid}(P_i, P_{i+1}), (X(\text{mid}(P_i, P_{i+1})), y_0)] \subsetneq l_i \quad (1)$$

where $X(x, y) = x$ and we note that $s_i \subsetneq l_i$ since $(X(\text{mid}(P_i, P_{i+1})), y_0)$ is directly to the south of $\text{mid}(P_i, P_{i+1})$. Let s_i^{\leftarrow} be the “reverse direction” of the line segment s_i , more precisely:

$$s_i^{\leftarrow} = [(X(\text{mid}(P_i, P_{i+1})), y_0), \text{mid}(P_i, P_{i+1})] \quad (2)$$

Also, define the line segment

$$s_j = [\text{mid}(P_j, P_{j+1}), (X(\text{mid}(P_j, P_{j+1})), y_0)] \subsetneq l_j \quad (3)$$

And the line segment

$$s_{j,i} = [(X(\text{mid}(P_j, P_{j+1})), y_0), (X(\text{mid}(P_i, P_{i+1})), y_0)] \quad (4)$$

There are two (almost identical) cases, (a) and (b).

Claim (a): $i < j$. We let c be the *concatenation* (see Definition 2.8) of the following six curves:

$$\begin{aligned} & s_i^{\leftarrow} \\ & [\text{mid}(P_i, P_{i+1}), \text{pos}(P_{i+1})] \\ & \mathfrak{E}_{P_{i+1, i+2, \dots, j}} \\ & [\text{pos}(P_j), \text{mid}(P_j, P_{j+1})] \\ & s_j \\ & s_{j,i} \end{aligned}$$

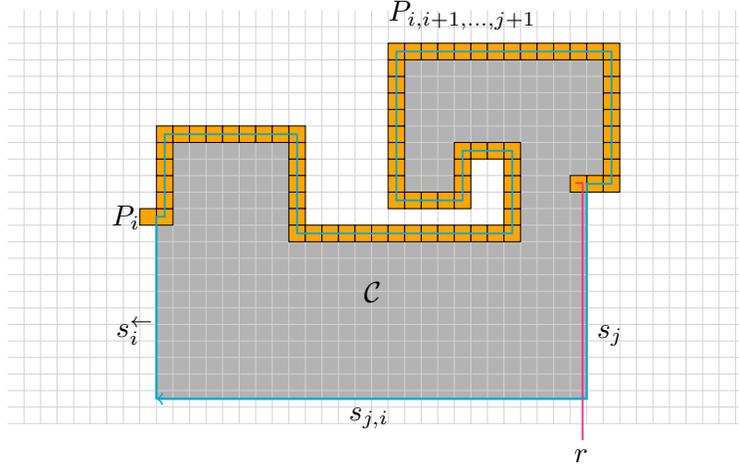


Figure 5.2: An example illustrating the proof of Lemma 5.2, Case (a): $i < j$. In this figure, $\text{glue}(P_i P_{i+1}) \in V_P^+$, $\text{glue}(P_j P_{j+1}) \in V_P^-$, and $\text{pos}(P_i)$ is to the left of $\text{pos}(P_j)$. The canonical embedding of the path $P_{i+1,i+2,\dots,j}$ in \mathbb{R}^2 and five line segments are used to define a closed simple curve c (light blue) and thus a bounded connected component \mathcal{C} of \mathbb{R}^2 (two of the line segments are unlabelled in this figure). Then, the ray r is used to verify that $\text{pos}(P_{j+1})$ is inside \mathcal{C} . Thus $P_{j+1,j+2,\dots}$ “grows inside” \mathcal{C} , but must leave \mathcal{C} at some point leading to a contradiction as P can not cross any of the components that define c .

By Observation 2.10, c is a finite closed simple curve¹¹ and thus defines a bounded connected component \mathcal{C} . (See Figure 5.2 for an example.)

We claim that $\text{pos}(P_{j+1})$ is inside \mathcal{C} . First, note that $\text{pos}(P_{j+1})$ is to the left of $\text{pos}(P_j)$ (because $\text{glue}(P_j P_{j+1}) \in V_P^-$), and therefore, $\text{pos}(P_{j+1})$ is to the left of l_j . But since $x_{P_i} < x_{P_j}$ by assumption, $X(\text{mid}(P_i, P_{i+1},)) \neq X(\text{mid}(P_j, P_{j+1},))$ by visibility, and P_{j+1} is unit distance to the left of P_j , then $\text{pos}(P_{j+1})$ is in fact between l_i and l_j (i.e. to the right of l_i and to the left of l_j). Secondly, $\text{pos}(P_{j+1})$ is above the horizontal line $s_{j,i}$. Consider the ray r at x-coordinate $X(\text{pos}(P_{j+1})) + 0.25$ that comes from the south and stops at position $p = (X(\text{pos}(P_{j+1})) + 0.25, Y(\text{pos}(P_{j+1})))$. Observe that r crosses c at the segment $s_{j,i}$ exactly once, and crosses c nowhere else, and that due to the visibility of $\text{glue}(P_j P_{j+1})$ we get that r does not intersect $\mathfrak{E}_{P_{i+1,i+2,\dots,j}}$, and that by its definition r is positioned away from the other four components of c . Furthermore, since c does not cross the short line segment $[\text{pos}(P_{j+1}), p]$ then starting at the point p , one can walk (westwards) along the segment $[\text{pos}(P_{j+1}), p]$ to the point $\text{pos}(P_{j+1})$, without crossing c . Hence $\text{pos}(P_{j+1})$ is inside \mathcal{C} as claimed.

Since, from the lemma statement, P has at least one tile P_k after P_j (i.e. $k > j$) positioned to the right of, or above, all tiles of $P_{i,i+1,\dots,j}$, then P has tiles positioned outside of \mathcal{C} after P_j . But since P is a path, it does not cross itself. Therefore, $\mathfrak{E}_{P_{j+1,j+2,\dots,|P|-1}}$ must cross at least one of l_i or l_j contradicting that both $\text{glue}(P_i P_{i+1})$ and $\text{glue}(P_j P_{j+1})$ are visible. Thus $x_{P_i} > x_{P_j}$.

¹¹To see this one needs to check that the components of c satisfy the hypotheses of Observation 2.10. Less formally but more intuitively, it can be seen that c is of finite length because its components are, also c is closed as each of the components are closed curves and their endpoints are pairwise equal in such a way to satisfy closure, and finally c is simple since the components are simple and their only intersection is at their endpoints in the order they are given.

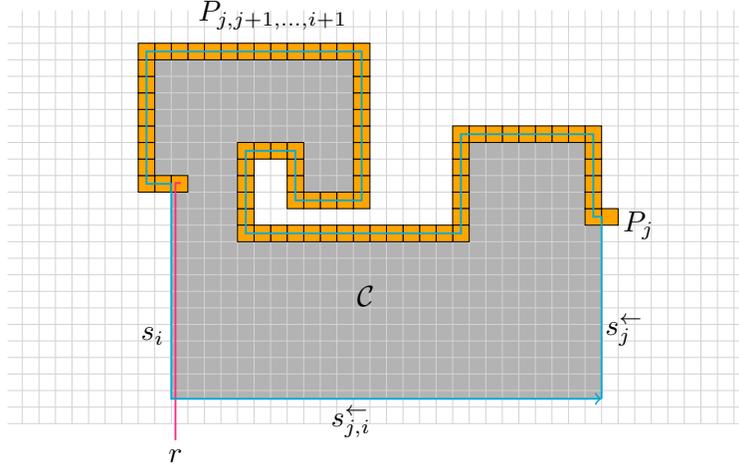


Figure 5.3: An example illustrating the proof of Lemma 5.2, Case (b): $j < i$. In this figure, $\text{glue}(P_i P_{i+1}) \in V_P^+$, $\text{glue}(P_j P_{j+1}) \in V_P^-$, and $\text{pos}(P_i)$ is to the left of $\text{pos}(P_j)$ (i.e. $P_{j,j+1,\dots}$ begins at the rightmost orange tile). The canonical embedding of the path $P_{j+1,j+2,\dots,i}$ in \mathbb{R}^2 and five line segments are used to define a closed simple curve c (light blue) and thus a bounded connected component \mathcal{C} of \mathbb{R}^2 (two of the line segments are unlabelled in this figure). Then, the ray r is used to verify that $\text{pos}(P_{i+1})$ is inside \mathcal{C} . Thus $P_{i+1,i+2,\dots}$ “grows inside” \mathcal{C} , but must leave \mathcal{C} at some point leading to a contradiction as P can not cross any of the components that define c .

Case (b): $j < i$. We define the curve c as the concatenation (see Definition 2.8) of the following six curves:

$$\begin{aligned}
& s_j^{\leftarrow} \\
& [\text{mid}(P_j, P_{j+1}), \text{pos}(P_{j+1})] \\
& \mathfrak{E}_{P_{j+1,j+2,\dots,i}} \\
& [\text{pos}(P_i), \text{mid}(P_i, P_{i+1})] \\
& s_i \\
& s_{j,i}^{\leftarrow}
\end{aligned}$$

By Observation 2.10, c is a finite closed simple curve¹¹, and thus defines a bounded connected component \mathcal{C} . (See an example in Figure 5.3.)

By a similar¹² argument as Case (a) (where $i < j$), $\text{pos}(P_{i+1})$ is inside \mathcal{C} . Since P has at least one tile P_k after P_i (i.e. $k > i$) to the right of or above $P_{j,j+1,\dots,i}$, then $P_{i+1,i+2,\dots,|P|-1}$ has tiles positioned outside \mathcal{C} . But since P is a path, it does not cross itself. Therefore, $\mathfrak{E}_{P_{i+1,i+2,\dots,|P|-1}}$ must cross at least one of l_i or l_j contradicting that both $\text{glue}(P_i P_{i+1})$ and $\text{glue}(P_j P_{j+1})$ are visible. Thus $x_{P_i} > x_{P_j}$. \square

Lemma 5.2 was our first statement proven using visibility, and the following lemma (5.3)

¹²Specifically, consider the ray r with x-coordinate $X(\text{pos}(P_{i+1}) - 0.25)$ that comes from the south and stops at position $p = (X(\text{pos}(P_{i+1}) - 0.25), Y(\text{pos}(P_{i+1})))$. Observe that r crosses c exactly once at the segment $s_{j,i}^{\leftarrow}$, that due to the visibility of $\text{glue}(P_i P_{i+1})$ we get that r does not intersect $\mathfrak{E}_{P_{j+1,j+2,\dots,i}}$, and that by its definition r is positioned away from the other four components of c . Furthermore, since c does not cross the short line segment $[p, \text{pos}(P_{i+1})]$ then starting at the point p one can walk (eastwards) along the segment $[p, \text{pos}(P_{i+1})]$ to the point $\text{pos}(P_{i+1})$, without crossing c . Thus $\text{pos}(P_{i+1})$ is inside \mathcal{C} .

similarly exploits the technique of using rays given by visibility, an embedded path, and other line segments to enclose a connected component of the plane, enabling us to reason about how visible glues are organised along a path. Together Lemmas 5.2 and 5.3 and Corollary 5.4 give properties of how tiles of V_P^+ and V_P^- are arranged in the plane. This is formalized in greater detail in Lemma 5.8.

Lemma 5.3 (V_P^+ glue order preserves path order). *Let $P \in \mathbf{P}[\mathcal{U}]$ be a path producible by any tile assembly system $\mathcal{U} = (U, \sigma, 1)$, and let i, j be such that $\text{glue}(P_i P_{i+1}) \in V_P^+$, $\text{glue}(P_j P_{j+1}) \in V_P^+$ and $x_{P_i} < x_{P_j}$. If P has at least one tile P_k after P_i and P_j (i.e. $i < k, j < k$), where $\text{pos}(P_k)$ is to the right of, or above, all tiles of $P_{\min(i,j), \min(i,j)+1, \dots, \max(i,j)}$, then $i < j$.*

Proof. Define s_i as in Equation (1), s_i^{\leftarrow} as in Equation (2), s_j as in Equation (3) and $s_{j,i}$ as in Equation (4).

First assume, for the sake of contradiction, that $j < i$. We let c be the *concatenation* (see Definition 2.8) of the following six curves:

$$\begin{aligned} & s_i^{\leftarrow} \\ & [\text{mid}(P_i, P_{i+1}), \text{pos}(P_i)] \\ & \mathfrak{C}_{P_{j+1}, j+2, \dots, i}^{\leftarrow} \\ & [\text{pos}(P_{j+1}), \text{mid}(P_j, P_{j+1})] \\ & s_j \\ & s_{j,i} \end{aligned}$$

By Observation 2.10, c is a simple closed curve, hence partitions \mathbb{R}^2 into two connected components, exactly one of which is bounded. Let \mathcal{C} be that bounded connected component. By a similar argument¹³ using ray r as in the proof of Lemma 5.2, $P_{i+1, i+2, \dots, |P|-1}$ starts inside \mathcal{C} . However, since $P_{i+1, i+2, \dots, |P|-1}$ has at least one tile to the right of, or above $P_{j, j+1, \dots, i}$, then $P_{i+1, i+2, \dots, |P|-1}$ cannot be entirely inside \mathcal{C} . Therefore, P needs to cross the border of \mathcal{C} , contradicting either the fact that P is simple, or that $\text{glue}(P_i P_{i+1})$ or $\text{glue}(P_j P_{j+1})$ are visible. \square

By flipping the use of “+” and “-”, and “left” and “right”, in the statement of the previous lemma, we immediately get the following corollary:

Corollary 5.4 (V_P^- glue order preserves path order). *Let $P \in \mathbf{P}[\mathcal{U}]$ be a path producible by any tile assembly system $\mathcal{U} = (U, \sigma, 1)$, and let i, j be such that $\text{glue}(P_i P_{i+1}) \in V_P^-$, $\text{glue}(P_j P_{j+1}) \in V_P^-$ and $x_{P_i} > x_{P_j}$. If P has at least one tile P_k after P_i and P_j (i.e. $i < k, j < k$), where $\text{pos}(P_k)$ is to the right of, or above all tiles of $P_{\min(i,j), \min(i,j)+1, \dots, \max(i,j)}$, then $i < j$.*

In the proof of Lemma 5.10, we will attempt to pump a path segment. We will make use of the following lemma stating that the V_P^+ glues on the prefix remain visible even if that prefix is pumped. In other words *visibility survives pumping*.

Lemma 5.5 (Visibility survives pumping for V_P^+). *Let $P \in \mathbf{P}[\mathcal{U}]$ be a path producible by any tile assembly system $\mathcal{U} = (U, \sigma, 1)$, and i, j be two integers such that $i < j$, $\text{glue}(P_i P_{i+1}) \in V_P^+$, $\text{glue}(P_j P_{j+1}) \in V_P^+$, and $\text{type}(\text{glue}(P_i P_{i+1})) = \text{type}(\text{glue}(P_j P_{j+1}))$. Let \bar{q} be the pumping of P*

¹³Here, r comes from the south with x-coordinate $X(\text{pos}(P_{i+1})) - 0.25$, crosses $s_{j,i}$ (thus entering \mathcal{C}), ends at y-coordinate $Y(\text{pos}(P_{i+1}))$, and then we walk from that end point to $\text{pos}(P_{i+1})$ staying inside \mathcal{C} .

between i and j (as defined in Definition 2.4), and let Q be the maximal prefix of \bar{q} that is an assemblable path.

Then $V_{P_{0,1,\dots,j}}^+ \subseteq V_Q^+$. Intuitively, this means that all “+” glues visible relative to $P_{0,1,\dots,j}$ are also visible relative to Q (note that Q contains $P_{0,1,\dots,j}$ as a prefix).

Proof. Setup. Let l_i be the visibility ray of $\text{glue}(P_i P_{i+1})$ and l_j be the visibility ray of $\text{glue}(P_j P_{j+1})$. Now let P_{left} be a leftmost tile of $\sigma \cup \text{asm}(P_{0,1,\dots,j})$, and let l_{left} be the horizontal ray to the west starting from $\text{pos}(P_{\text{left}})$. Moreover, let P_{bottom} be a lowest tile of $P_{0,1,\dots,j}$.

We now define two helper points of \mathbb{R}^2 , far enough from $P_{0,1,\dots,j}$:

$$\begin{aligned} A &= P_{\text{left}} + \begin{pmatrix} -10 \\ 0 \end{pmatrix} \\ B &= \begin{pmatrix} X(\text{mid}(P_j, P_{j+1})) \\ Y(\text{pos}(P_{\text{bottom}})) - 10 \end{pmatrix} \end{aligned}$$

where $X(x, y) = x$, $Y(x, y) = y$.

Intuitively, A has the same y -coordinate as, and is 10 units to the left of P_{left} , and B has the same x -coordinate as the glue between P_j and P_{j+1} , and is 10 units below the lowest point of $P_{0,1,\dots,j}$.¹⁴

Define $\mathcal{Q}_A = \{(x, y) \in \mathbb{R}^2 \mid x \leq x_A \wedge y \leq y_A \text{ where } A = (x_A, y_A)\}$ (the quarter-plane below and to the left of A) and $\mathcal{Q}_B = \{(x, y) \in \mathbb{R}^2 \mid x \leq x_B \wedge y \leq y_B \text{ where } B = (x_B, y_B)\}$ (the quarter-plane below and to the left of B). Also, let $s_j = [\text{pos}(P_j), \text{mid}(P_j, P_{j+1})]$ be the half-unit horizontal segment of \mathbb{R}^2 from $\text{pos}(P_j)$ to $\text{mid}(P_j, P_{j+1})$.

Moreover, let c be the curve that is the concatenation of $[(x_A, y_B), A]$, $[A, \text{pos}(P_{\text{left}})]$, $\mathfrak{E}_{P_{\text{left}}, \text{left}+1, \dots, j}$, s_j , $[\text{mid}(P_j, P_{j+1}), B]$, $[B, (x_A, y_B)]$. By Observation 2.10, c is a simple closed curve, hence defines a bounded connected component \mathcal{C} of \mathbb{R}^2 .

Now, let $C_j = \mathcal{Q}_A \cup \mathcal{Q}_B \cup \mathcal{C}$, which is also a single connected component of \mathbb{R}^2 , because (x_A, y_B) is in all components of the union $((x_A, y_B) \in \mathcal{Q}_A \cap \mathcal{Q}_B \cap \mathcal{C})$ and each component is connected.

Proof argument. We claim that for all $k < j$, $l_k \subset C_j$ since $l_k \subset \mathcal{Q}_B \cup \mathcal{C}$ which can be seen as follows: (i) $l_k^B = \{(x, y) \mid (x, y) \in l_k \wedge y \leq B\}$ is contained in \mathcal{Q}_B since l_k is to the left of l_j (by the contrapositive of Lemma 5.3) and l_k^B reaches infinitely far to the south as does \mathcal{Q}_B ; and (ii) by the definition of visibility the segment $l_k \setminus l_k^B$ is contained in \mathcal{C}_j .

Now suppose, for the sake of contradiction, that $\mathfrak{E}_{Q_{j+1, j+2, \dots}}$, the canonical embedding of the path $Q_{j+1, j+2, \dots}$ in \mathbb{R}^2 , intersects l_k . Then $\mathfrak{E}_{Q_{j+1, j+2, \dots}}$ needs to enter C_j , because $\text{pos}(Q_{j+1})$ is outside C_j (by the contrapositive of Lemma 5.3). This crossing can happen at only four different parts of the border of C_j :

- l_{left} , but this is impossible since $\text{pos}(P_{\text{left}})$ is a leftmost point of $\sigma \cup \text{asm}(P_{0,1,\dots,j})$ and $\overrightarrow{P_i P_j}$ (as a vector) has a strictly positive x -coordinate (by the contrapositive of Lemma 5.3).
- $P_{\text{left}, \text{left}+1, \dots, j}$ or s_j , but this would contradict the fact that Q is simple.
- l_j . We claim that this is impossible: assume, for the sake of contradiction, that it is not, and let $k_0 > j$ be the smallest integer such that $Q_{j+1, j+2, \dots, k_0}$ intersects l_j . Moreover, $k_0 - (j - i) > j$ because $Q_{i, i+1, \dots, j}$ does not intersect its own visibility rays (except at their endpoints). Therefore, $Q_{j+1, j+2, \dots, k_0 - (j - i)}$ would also intersect l_i , hence $Q_{j+1, j+2, \dots, k_0 - (j - i)}$ also enters C_j ,

¹⁴The choice of 10 units is arbitrary here, we simply need to define a curve with segments that are strictly to the left of $\sigma \cup \text{asm}(P_{0,1,\dots,j})$ and below $P_{0,1,\dots,j}$.

contradicting the assumption that k_0 is the smallest integer such that $Q_{j,j+1,\dots,k_0}$ intersects l_j (since in the previous bullet points we have shown that $Q_{j+1,j+2,\dots,k_0-(j-i)}$ cannot cross other parts of the border of C_j).

Therefore, $\mathfrak{E}_{Q_{j+1,j+2,\dots}}$ does not enter C_j , which is a contradiction, which in turn contradicts our assumption that $\mathfrak{E}_{Q_{j+1,j+2,\dots}}$ intersects l_k . Hence the canonical embedding of Q does not cross any of the visibility rays defining $V_{P_{0,1,\dots,j}}^+$, thus $V_{P_{0,1,\dots,j}}^+ \subseteq V_Q^+$. \square

By flipping the use of “+” and “-” in the statement of the previous lemma, and “left” and “right” in the proof, we get the following corollary:

Corollary 5.6 (Visibility survives pumping for V_P^-). *Let $P \in \mathbf{P}[\mathcal{U}]$ be a path producible by any tile assembly system $\mathcal{U} = (U, \sigma, 1)$, and i, j be two integers such that $i < j$, $\text{glue}(P_i P_{i+1}) \in V_P^-$ and $\text{glue}(P_j P_{j+1}) \in V_P^-$, and $\text{type}(\text{glue}(P_i P_{i+1})) = \text{type}(\text{glue}(P_j P_{j+1}))$. Let \bar{q} be the pumping of P between i and j , and let Q be the maximal prefix of \bar{q} that is an assemblable path.*

Then $V_{P_{0,1,\dots,j}}^- \subseteq V_Q^-$. Intuitively, this means that all “-” glues visible relative to $P_{0,1,\dots,j}$ are also visible relative to Q (note that Q contains $P_{0,1,\dots,j}$ as a prefix).

5.2 Blocking any h -successful path

Keeping in mind that the seed assembly supertile of \mathcal{U} includes the origin $(0, 0) \in \mathbb{Z}^2$, for the rest of the paper fix a horizontal line at height $h = 10m$ above the origin.

Definition 5.7 (The set of h -successful paths of \mathcal{U}). The set of h -successful paths $\mathbb{P}_{\mathcal{U}}$ of \mathcal{U} is defined as:

$$\mathbb{P}_{\mathcal{U}} = \{P \mid P \in \mathbf{P}[\mathcal{U}] \text{ and } P \text{ contains exactly one tile at height } h = 10m, \text{ its last tile}\}$$

where $\mathbf{P}[\mathcal{U}]$ is the set of producible paths of \mathcal{U} (defined in Section 2.3).

Note that *any* claimed successful simulation by \mathcal{U} of \mathcal{T}_n (defined in Section 3) must exhibit at least one path that has a h -successful prefix P . When we write “ P is a h -successful path” we mean $P \in \mathbb{P}_{\mathcal{U}}$. The set of h -successful paths $\mathbb{P}_{\mathcal{U}}$ is finite because the tileset U is finite and the area of the simulation zone below the horizontal line at height h is finite.

We define a “*nowhere- h -successful path*” to be a path that has no tile at height h . In other words, a nowhere- h -successful path has no h -successful prefixes.

5.2.1 Visibility setup

We begin with the following lemma, which has a straightforward proof and is used merely to define i, j and ℓ , which are used extensively in later proofs. Recall that \mathcal{U} is a tile assembly system with tile set U simulating $\mathcal{T}_{10|U|}$ at scale factor m .

Lemma 5.8. *Let $P \in \mathbb{P}_{\mathcal{U}}$ be a h -successful path, and let ℓ be a vertical line in \mathbb{R}^2 with x -coordinate $|U|(3m+1) + m + 1.5$. If the visible glue placed by P on ℓ is a V_P^+ glue (respectively a V_P^- glue), then there exist $i < j \in \mathbb{N}$ that satisfy all of the following properties:*

1. $\text{pos}(P_i)$ is to the right (respectively to the left) of ℓ (i.e. $x_{P_i} > \ell$, respectively $x_{P_i} < \ell$)
2. $\text{pos}(P_j)$ is horizontal distance at least $3m$ from $\text{pos}(P_i)$ (i.e. $|x_{P_i} - x_{P_j}| \geq 3m$)
3. $\text{pos}(P_j)$ is to the right (respectively to the left) of $\text{pos}(P_i)$

4. $\text{glue}(P_i P_{i+1}) \in V_P^+$ and $\text{glue}(P_j P_{j+1}) \in V_P^+$ (respectively $\text{glue}(P_i P_{i+1}) \in V_P^-$ and $\text{glue}(P_j P_{j+1}) \in V_P^-$)
5. $\text{type}(\text{glue}(P_i P_{i+1})) = \text{type}(\text{glue}(P_j P_{j+1}))$
6. $\text{pos}(P_i)$ and $\text{pos}(P_j)$ are within vertical distance $3m$ (i.e. $|y_{P_j} - y_{P_i}| \leq 3m$)

Proof. Firstly, the path P is of width (horizontal extent) $\geq 8|U|m + 2$ (as we are simulating $\mathcal{T}_{10|U|}$ at scale factor m , see Figure 4.1). Secondly, since P is h -successful, P crosses ℓ , and since ℓ is positioned at distance $\geq |U|(3m + 1) + 1$ to the right¹⁵ of the seed σ of \mathcal{U} , there are at least $|U|(3m + 1) + 1$ visible glues to the left of ℓ .

Moreover:

- If the visible glue placed by P on ℓ is in V_P^+ , then since P has at least one tile in the rightmost $2m$ positions of the simulation zone, P has at least $|U|(3m + 1) + 1$ visible glues to the right of ℓ . By Lemma 5.2 all of these glues are in V_P^+ .
- Else, since ℓ is at horizontal distance at least $|U|(3m + 1) + 1.5$ from the rightmost tile of σ , P has at least $|U|(3m + 1) + 1$ visible glues to the left of ℓ . By Lemma 5.2 all of these glues are in V_P^- .

Thus, if we look at the first $|U|(3m + 1) + 1$ visible glues of P that are immediately to the right (respectively to the left) of ℓ , by the pigeonhole principle, at least one of their glue types appears at least $3m + 1$ times. Since each x-coordinate has exactly one visible tile, we can find two V_P^+ (respectively V_P^-) glues, $\text{glue}(P_i P_{i+1})$ and $\text{glue}(P_j P_{j+1})$ for some $i < j$, with the same type, that are at least horizontal distance $3m$ away from each other, which shows Conclusions 1, 2, 4 and 5 of this lemma. Taking the contrapositive of Lemma 5.3 (in that lemma letting P_k be the tile of P at height h), we get that $x_{P_i} < x_{P_j}$ (respectively, of Corollary 5.4, that $x_{P_j} < x_{P_i}$), which shows Conclusion 3.

Finally, since the region we chose to apply the pigeonhole argument is located immediately to the right of ℓ (respectively, left) and is of width merely $|U|(3m + 1) + 1$ neither $\text{glue}(P_i P_{i+1})$ nor $\text{glue}(P_j P_{j+1})$ are in the “vertical part” of the simulation zone (Figure 4.1), and since the “horizontal part” of the simulation zone is of height $3m$, this proves Conclusion 6. \square

5.2.2 Blocking any h -successful path by growing a branch from it

In this section we give Lemma 5.10 which is the first main tool used in this paper. We also give Theorem 5.11 whose short proof gives a method to block any h -successful path. We begin with the definition of an *enclosing branch*, which is a path D branching from P , and enclosing a connected component of \mathbb{R}^2 . The enclosing branch achieves this in one of two ways: (1) either by intersecting $\sigma \cup P$ (see Figure 5.4(Left)) and hence the enclosure is bordered by P , the enclosing branch and possibly σ , or (2) by placing a new visible glue on ℓ (see Figure 5.4(Right)) and hence the enclosure is bordered by P , the enclosing branch, a segment of ℓ and possibly σ .

Definition 5.9. [Enclosing branch for a path] Let $P \in \mathbb{P}_{\mathcal{U}}$ be a h -successful path and for any ℓ that satisfies the hypotheses of Lemma 5.8: Let $k \in \{0, 1, \dots, |P| - 1\}$ be such that $P_{0,1,\dots,k}$ includes $\text{glue}(P_\ell P_{\ell+1})$, P 's visible glue on ℓ . We call a path D an *enclosing branch* for P at k if $P_{0,1,\dots,k} D_{0,1,\dots,|D|-2} \in \mathbf{P}[\mathcal{U}]$ is an assemblable path, $\text{pos}(P_{k+1}) = \text{pos}(D_0)$, $\text{glue}(P_k P_{k+1}) = \text{glue}(P_k D_0)$ and is visible relative to $P_{0,1,\dots,k} D_{0,1,\dots,|D|-2} \in \mathbf{P}[\mathcal{U}]$, and at least one of the following is the case:

¹⁵Note that the $m \times m$ seed supertile region contains the point $(0, 0) \in \mathbb{Z}^2$.

1. at $\text{pos}(D_{|D|-1})$ the path D intersects $\sigma \cup \text{asm}(P_{0,1,\dots,k})$,¹⁶ or
2. $\text{mid}(D_{|D|-2}, D_{|D|-1}) \in \mathbb{R}^2$ is on ℓ and strictly below the points in the set $\mathfrak{E}_P \cap \ell$.¹⁷

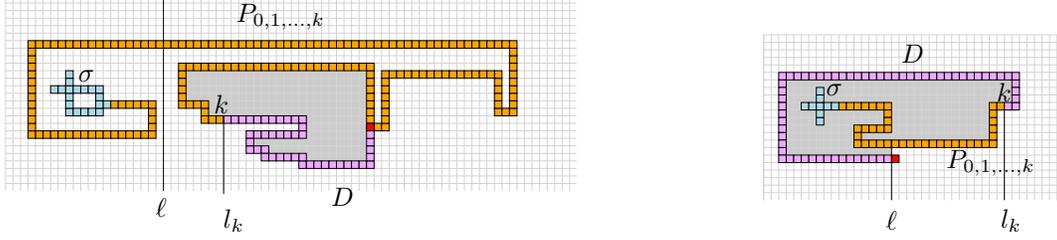


Figure 5.4: Enclosing branches that enclose the part of the plane shaded grey. **Left:** Example of an enclosing branch D for P at i for Definition 5.9(1). D is the pink and red path, the last tile of D , which intersects P , is drawn in red (hence PD is not a path because its last tile intersects an earlier part of itself). **Right:** Example of an enclosing branch D for P at i for Definition 5.9(2). D is the pink and red path, the last tile of D (immediately after D places a new visible glue on ℓ) is drawn in red. Note that Figure 5.13 shows a concrete construction of such an enclosing branch.

Before stating Lemma 5.10, we give an intuitive, although imprecise, *summary* of its statement, of *how* we obtain such a statement and *why* we would want such a statement. Let P be any h -successful path, which implies that P is long enough to repeat some glue type many times. This implies one of two things (which are the two conclusions of the lemma):

1. There is a path, assemblable by \mathcal{U} , that places tiles outside of the simulation zone. If the proof yields this conclusion, then such a path is found either by repeating one part of P enough times (this is point P1 in the proof), or else by deriving a h -successful path P' from P , and “shortcutting” a part of P' , so as to translate a suffix of P' by a large enough vector. In both cases, this contradicts that \mathcal{U} simulates $\mathcal{T}_{10|U|}$, which gives the proof of our main theorem.
2. Else, we can grow a path of the form $P_{0,1,\dots,k}D_{0,1,\dots,|D|-2}$, for some $k \in \{i, j\}$, which has the following properties:
 - $\text{glue}(P_k P_{k+1})$, which is visible relative to P , is also visible relative to D , and
 - any h -successful path that turns to the right from $D_{0,1,\dots}$ must hide the visibility of $\text{glue}(P_k P_{k+1})$.

First we observe that P must be blocked by such a $P_{0,1,\dots,k}D_{0,1,\dots}$ (i.e. P cannot grow from an assembly that already contains $P_{0,1,\dots,k}D_{0,1,\dots}$), because we can show that P turns to the right from $D_{0,1,\dots,a}$ ¹⁸, but this causes a contradiction as P can not hide its own visible glue. This is one useful fact from Conclusion 2.

The second useful fact (and the main reason for the particular form of Conclusion 2), is that we have found a way to consume a finite resource: visible glues (note that the entire set of h -successful paths is finite and each such path has a finite set of visible glues). Then, later in the proof of Theorem 1.1, we will use this “consumption of visible glues” property to show

¹⁶and therefore D cannot grow completely from $\sigma \cup \text{asm}(P_{0,1,\dots,k})$ without intersecting $\sigma \cup \text{asm}(P_{0,1,\dots,k})$

¹⁷And therefore D cannot grow completely from $\sigma \cup \text{asm}(P_{0,1,\dots,k})$ without placing a glue on ℓ that is lower than P ’s visible glue on ℓ .

¹⁸From the definition of enclosing branch P places a tile at $\text{pos}(D_0)$ therefore either P branches to the right (see proof of Theorem 5.11) from D , or if it happens that $\text{type}(P_{k+1}) \neq \text{type}(D_0)$ then P is blocked by D .

that the “supply of visible tiles” must decrease with each new “branch” to the right starting from $P_{0,1,\dots,k}D_{0,1,\dots,|D|-2}$; and the particular form of Conclusion 2 allows us to do this.

The previous intuitive overview of the lemma statement is somewhat incomplete, but may serve as a guide. All the steps of the *proof* are given in Figure 5.5, with a running pictorial example, and short summaries of each step.

Lemma 5.10 (Pumping or enclosing any h -successful path). *Let $P \in \mathbb{P}_{\mathcal{U}}$ be a h -successful path, let i, j, ℓ be as in Lemma 5.8 with P 's visible glue on ℓ being in V_P^+ (respectively, V_P^-). At least one of the following holds:*

1. *There is an assemblable path (i.e. from $\mathbf{P}[\mathcal{U}]$) that places tiles outside of the simulation zone of \mathcal{U} , contradicting that \mathcal{U} simulates $\mathcal{T}_{10|U|}$.*
2. (i) *There is an enclosing branch $D = D_{0,1,\dots,|D|-1}$ for P at some $k \in \{i, j\}$, such that $P_{0,1,\dots,k}D_{0,1,\dots,|D|-2} \in \mathbf{P}[\mathcal{U}]$ is an assemblable nowhere- h -successful path that has the same visible glue on ℓ as P .*
(ii) *Furthermore, for all paths R such that for some $a \leq |D| - 2$, $P_{0,1,\dots,k}D_{0,1,\dots,a}R$ is h -successful and turns right (respectively, left) from $P_{0,1,\dots,k}D_{0,1,\dots,|D|-2}$, at least one of the following is the case:*
 - *glue($P_k P_{k+1}$) is not visible relative to $P_{0,1,\dots,k}D_{0,1,\dots,a}R$, or*
 - *R has a lower visible glue on ℓ than P .*

Proof. A tree representing the proof is shown in Figure 5.5. The proof performs a depth-first, left-first, search of the tree. Points in the tree and proof are numbered P1, P3, . . . , P9 correspondingly.

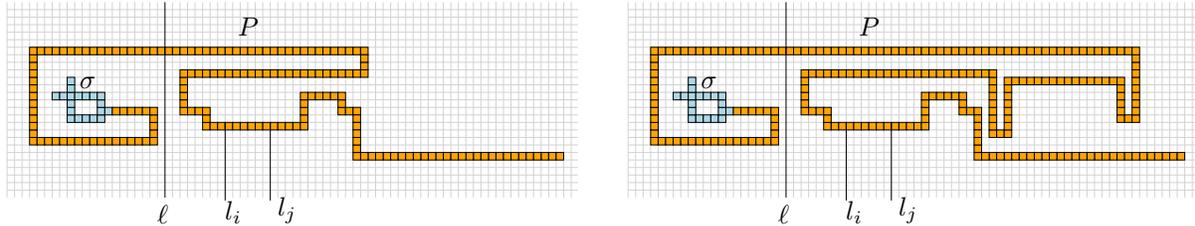


Figure 5.6: Two example h -successful paths. For brevity, their suffixes that grow to height h are not shown. Here, and in all figures in the paper, l_i and l_j denote the two visibility rays of the glues $\text{glue}(P_i P_{i+1})$ and $\text{glue}(P_j P_{j+1})$ respectively. **Left:** A path P where the segment $P_{i+1,i+2,\dots,j}$ can be (infinitely) pumped (see Figure 5.7). **Right:** A path P where the segment $P_{i+1,i+2,\dots,j}$ can not be (infinitely) pumped (see Figure 5.9).

Intuitively, let \bar{q} be the sequence composed of $P_{0,1,\dots,i}$ followed by the “infinite pumping” of the segment $P_{i+1,i+2,\dots,j}$. Formally, let \bar{q} be the pumping of P between i and j (Definition 2.4), and let $Q \in \mathbf{P}[\mathcal{U}]$ be the longest prefix of \bar{q} that (a) is a path (i.e. non-self intersecting), (b) does not intersect σ , and (c) is assemblable from the seed σ of \mathcal{U} . Q is either infinite or finite.

P1 If Q is infinite (as in the example in Figure 5.7): Then since $\overrightarrow{P_i P_j}$ has a nonzero horizontal component (by Lemma 5.8 the horizontal distance between $\text{glue}(P_i P_{i+1})$ and $\text{glue}(P_j P_{j+1})$ is $\geq 3m$), this infinite assembly reaches infinitely to the right for glue g in the lemma statement being in V_P^+ (respectively, to the left for g being in V_P^-), and thus places tiles outside of the simulation zone, giving Conclusion 1.

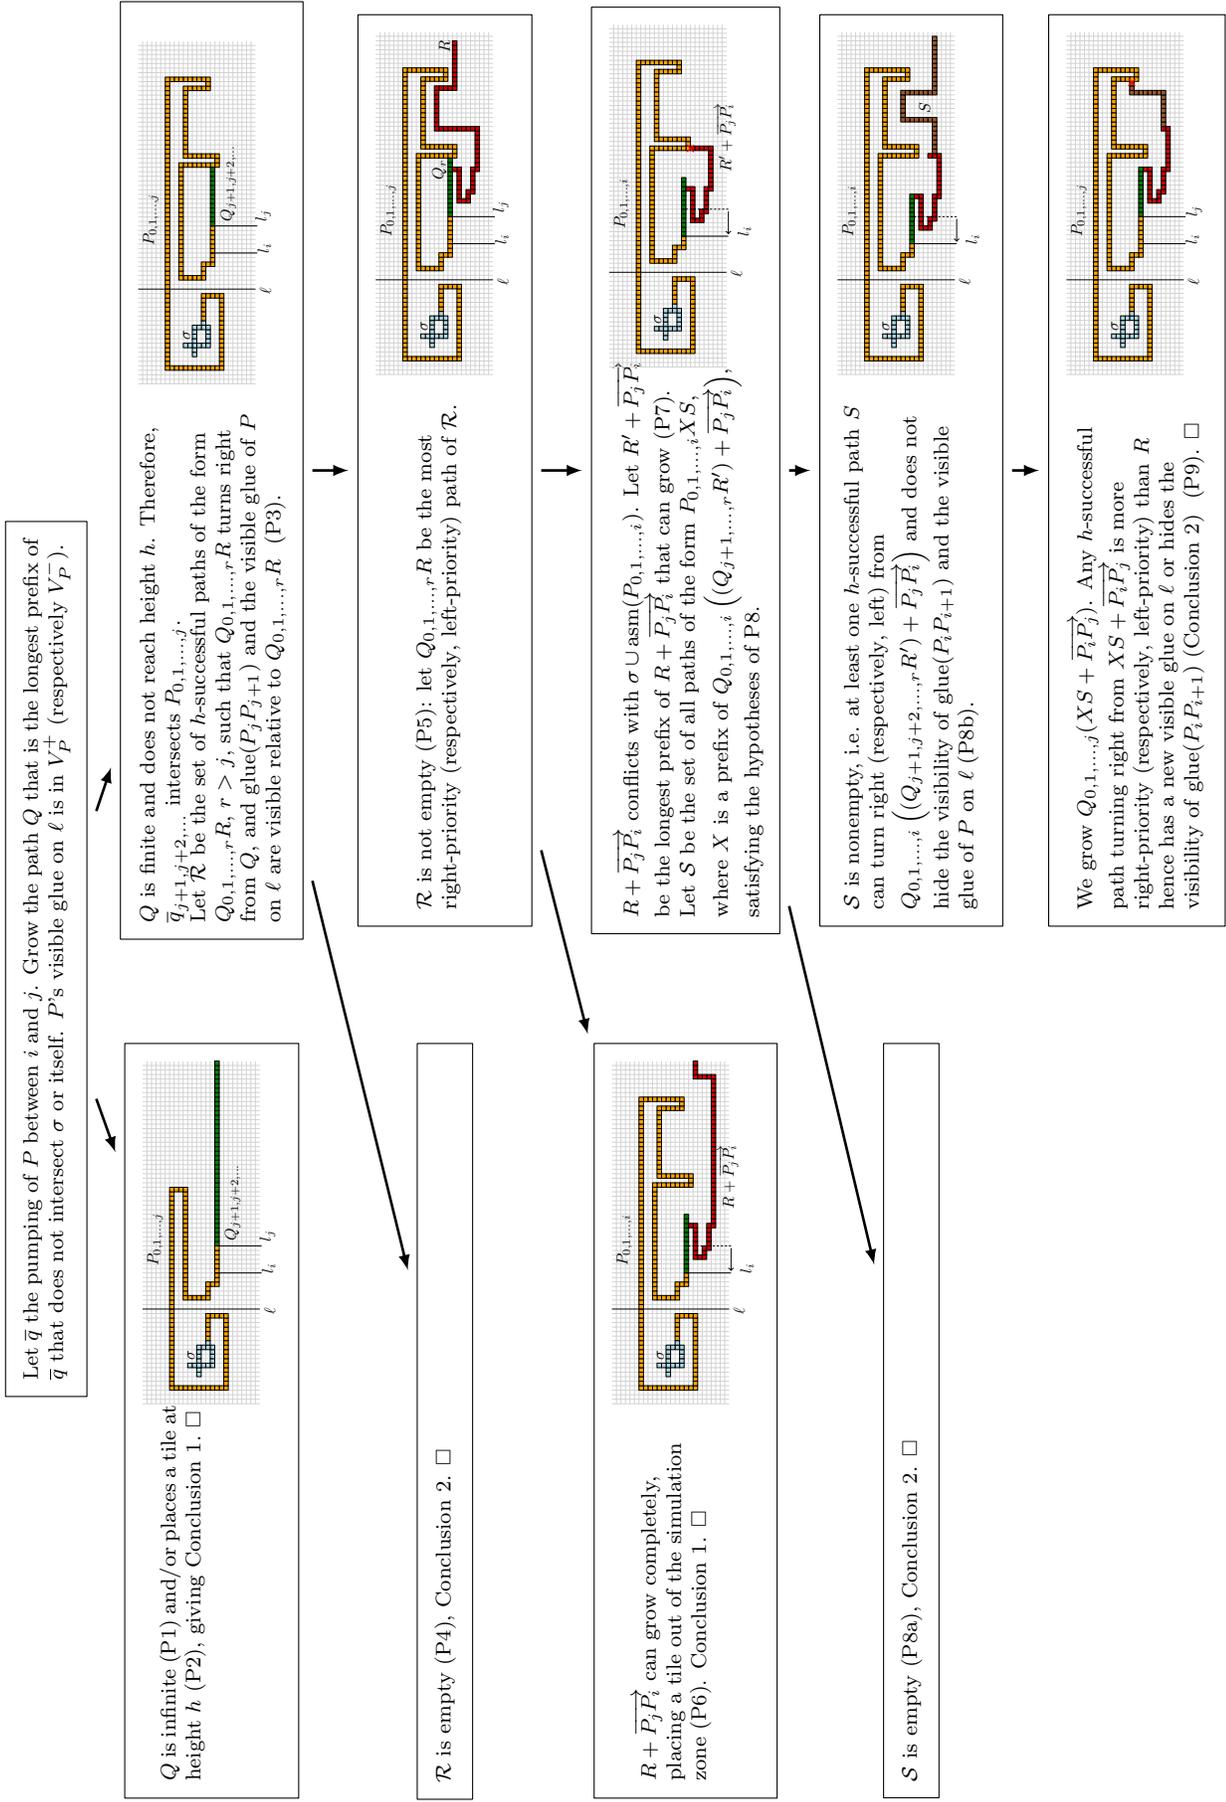


Figure 5.5: A summary of the major points in the proof of Lemma 5.10. The proof essentially performs a depth-first, left-first, search of this tree with each leaf ending in one of the lemma conclusions. Points in the tree and proof are numbered P1, P3, ..., P9 correspondingly.

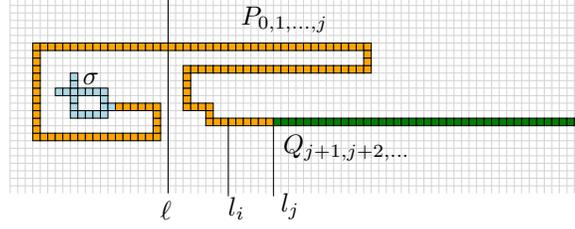


Figure 5.7: Case P1: Q is an infinite path because the sequence \bar{q} does not intersect $P_{0,1,\dots,j}$ nor σ . Note that Q is defined so that it shares a prefix with P : i.e. $P_{0,1,\dots,j} = Q_{0,1,\dots,j}$.

P2 If some prefix Q' of Q is h -successful we claim that we get Conclusion 1 of the lemma statement. To see this, note that $P_{0,1,\dots,j}$ is nowhere- h -successful. Therefore, if Q' were h -successful, Q would place a tile $Q_s = Q_{|Q'|}$ (where $|Q'| \in \mathbb{N}$ is the length of Q') on the horizontal line $y = h$, which means that Q would also have a tile at position $\text{pos}(Q_s - \overrightarrow{P_i P_j})$ (i.e. positioned one iteration of the “pumping” earlier, see Figure 5.8). But $\text{pos}(Q_s - \overrightarrow{P_i P_j})$ is outside of the simulation zone: $\text{pos}(Q_s)$ is in the simulation zone on line h , the “vertical part” of the simulation zone (that intersects line h and for distance $7m$ below h) is of width $3m$, $\overrightarrow{P_i P_j}$ has horizontal length $\geq 3m$ and vertical length $\leq 3m$ (Lemma 5.8). This immediately gives Conclusion 1.

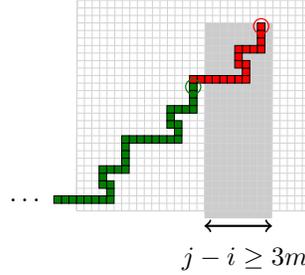


Figure 5.8: Case P2: Q (in green and red in the picture) is h -successful. In this case, the last tile $Q_{|Q|-1}$ of Q (circled in red) is at height h . Since $|x_{P_j} - x_{P_i}| \geq 3m$, and because $Q_{i,i+1,\dots}$ is periodic, Q also contains tile $Q_{|Q|-1} - \overrightarrow{P_j P_i}$ (in red in the picture), which is outside of the simulation zone (in gray).

P3 Else Q is finite, and nowhere- h -successful. We assert the claim that $\bar{q}_{|Q|}$ intersects $\sigma \cup \text{asm}(P_{0,1,\dots,j})$, which we will next prove.¹⁹

If $\bar{q}_{|Q|}$ intersects σ , we are immediately done. Else, \bar{q} intersects itself. Let $a < |Q|$ be the *smallest* integer such that $\text{pos}(\bar{q}_a) = \text{pos}(\bar{q}_{|Q|})$. We will show that $a \leq j$. Assume, for the sake of contradiction that $a > j$.

Note that $|Q| - (j - i) \geq a - (j - i) \geq i + 1$. Therefore, we can apply Lemma 2.5 to show that $Q_{a-(j-i)} = Q_a - \overrightarrow{P_i P_j}$ and $Q_{|Q|-(j-i)} = \bar{q}_{|Q|} - \overrightarrow{P_i P_j}$. But then since $\text{pos}(Q_a) = \text{pos}(\bar{q}_a) = \text{pos}(\bar{q}_{|Q|})$, and since $a > j$,²⁰ we get that $\text{pos}(Q_{a-(j-i)}) = \text{pos}(Q_a) - \overrightarrow{P_i P_j} = \text{pos}(\bar{q}_{|Q|}) - \overrightarrow{P_i P_j} =$

¹⁹Note that we already know that $\bar{q}_{|Q|}$ intersects σ or $\bar{q}_{0,1,\dots,|Q|-1}$. Intuitively, if $\bar{q}_{j+1,j+2,\dots}$ intersects itself, then since $\bar{q}_{i,i+1,\dots}$ is periodic and $i < j$, we can find another self-intersection $j - i$ indices earlier along \bar{q} . Applying this argument repeatedly will show an intersection with $P_{i,i+1,\dots,j}$.

²⁰intuitively, we are in the ≥ 2 nd iteration of pumping

$\text{pos}(Q_{|Q|-(j-i)})$. Therefore, Q intersects itself which is a contradiction. Hence $a \leq j$ and thus $\bar{q}_{|Q|}$ intersects $\sigma \cup \text{asm}(P_{0,1,\dots,j})$ as claimed. An example of this case is shown in Figure 5.9.

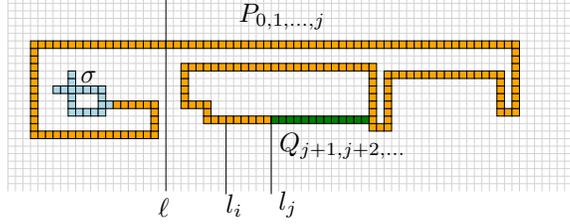


Figure 5.9: Case P3: \bar{q} conflicts with $P_{0,1,\dots,j}$. The figure shows σ and the path Q . Note that Q is defined so that it shares a prefix with P : i.e. $P_{0,1,\dots,j} = Q_{0,1,\dots,j}$.

P4 Let \mathcal{R} be the set of all h -successful paths of the form $Q_{0,1,\dots,r}R \in \mathbf{P}[\mathcal{U}]$, for any r such that $j < r \leq |Q| - 1$ and any path R , such that all of the following hold:

- (i) $Q_{0,1,\dots,r}R$ turns right (respectively, left) from Q ,
- (ii) $\text{glue}(P_j P_{j+1}) = \text{glue}(Q_j Q_{j+1})$ is visible relative to $Q_{0,1,\dots,r}R$,
- (iii) $Q_{0,1,\dots,r}R$ has the same visible glue as P on ℓ .

If \mathcal{R} is empty²¹, then we claim that we get Conclusion 2 of the lemma statement, with $k = j$ and $D = \bar{q}_{j+1,j+2,\dots,|Q|}$. Indeed, D is an enclosing branch for P at j with the *same visible glue* on ℓ as P . Moreover, by Lemma 5.5 (respectively, Corollary 5.6), Q does not hide the visibility of $\text{glue}(P_i P_{i+1})$ or $\text{glue}(P_j P_{j+1})$, hence D satisfies Conclusion 2(i). Then, Conclusion 2(ii) is immediately satisfied since \mathcal{R} is empty (intuitively, \mathcal{R} is the set of paths that do not meet that conclusion).

P5 **Else, \mathcal{R} is not empty.** Until the end of this proof, let r and R be such that $Q_{0,1,\dots,r}R$ is the most right-priority (respectively, left-priority) of \mathcal{R} . Note that since all paths of \mathcal{R} are h -successful, the last tile of $R_{|R|-1}$ is the only tile of R at height h . Such a path is shown in Figure 5.10. Notice that, because R is the most right-priority path of \mathcal{R} , R does not turn left from Q (before turning right from Q), or else we could find a h -successful path turning right earlier than R .

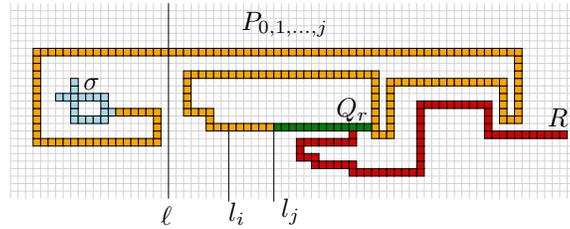


Figure 5.10: Case P5: there is at least one path in \mathcal{R} . Let $Q_{0,1,\dots,r}R \in \mathcal{R}$ be the most right-priority one.

We grow $\sigma \cup \text{asm}(P_{0,1,\dots,i}) = \sigma \cup \text{asm}(Q_{0,1,\dots,i})$, and then the maximal assemblable prefix of the following “translated path”: $(Q_{j+1,j+2,\dots,r}R) + \overrightarrow{P_j P_i}$. In other words, starting from σ , we grow the following path: $Q_{0,1,\dots,i}((Q_{j+1,j+2,\dots,r}R) + \overrightarrow{P_j P_i})$ (Figure 5.11 below highlights this

²¹which includes in particular the case where no h -successful path can branch from Q

“backwards translation”).

P6 If all of $(Q_{j+1,j+2,\dots,r}R) + \overrightarrow{P_j P_i}$ is assemblable from $\sigma \cup \text{asm}(P_{0,1,\dots,i})$, as in the example in Figure 5.11, then we claim that $\text{pos}(R_{|R|-1} + \overrightarrow{P_j P_i})$ is outside the simulation zone, yielding Conclusion 1. To see this first note that since R is h -successful the tile $R_{|R|-1}$ (i.e. untranslated) is on the horizontal line $y = h$ and thus in the “vertical part” of the simulation zone (see Figure 4.1) which is of width $3m$. But, by Lemma 5.8, $\overrightarrow{P_j P_i}$ has horizontal length $\geq 3m$ and vertical length $\leq 3m$. This means that $\text{pos}(R_{|R|-1} + \overrightarrow{P_j P_i})$ is outside of the simulation zone, giving Conclusion 1.

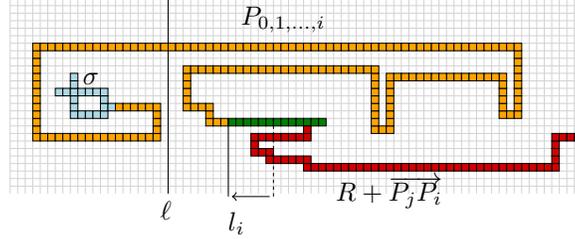


Figure 5.11: Case P6: $(Q_{j+1,j+2,\dots,r}R) + \overrightarrow{P_j P_i}$ is assemblable from $\sigma \cup \text{asm}(P_{0,1,\dots,i})$.

P7 Else not all of $(Q_{j+1,j+2,\dots,r}R) + \overrightarrow{P_j P_i}$ is assemblable from $\sigma \cup \text{asm}(P_{0,1,\dots,i})$. Let R' be the longest prefix of R such that $R' + \overrightarrow{P_j P_i}$ does not intersect $\sigma \cup \text{asm}(P_{0,1,\dots,i})$ and $R' + \overrightarrow{P_j P_i}$ does not have any visible glue on ℓ below the visible glue of P on ℓ . (See Figure 5.12 for an example where the longest assemblable prefix of $R + \overrightarrow{P_j P_i}$ does not place any new visible glue on ℓ , and Figure 5.13 for an example where it does.)

In this case, we make the more specific claim that $R + \overrightarrow{P_j P_i}$ conflicts with $\sigma \cup \text{asm}(P_{0,1,\dots,i})$: indeed, by its definition in P5 above, R does not conflict with $Q_{0,1,\dots,r}$, and hence does not conflict with $Q_{j+1,j+2,\dots,r}$, hence $R + \overrightarrow{P_j P_i}$ does not conflict with $Q_{j+1,j+2,\dots,r} + \overrightarrow{P_j P_i} = Q_{i+1}Q_{i+2} \dots Q_{r-(j-i)}$. The only part of the assembly $\sigma \cup \text{asm}(P_{0,1,\dots,i}) \cup (\text{asm}(Q_{j+1,j+2,\dots,r} + \overrightarrow{P_j P_i}))$ that $R + \overrightarrow{P_j P_i}$ can conflict with is therefore $\sigma \cup \text{asm}(P_{0,1,\dots,i})$.

Observe that $(Q_{j+1,j+2,\dots,r}R_{0,1,\dots,|R'|}) + \overrightarrow{P_j P_i}$ is an enclosing branch for P at i (Definition 5.9²²).

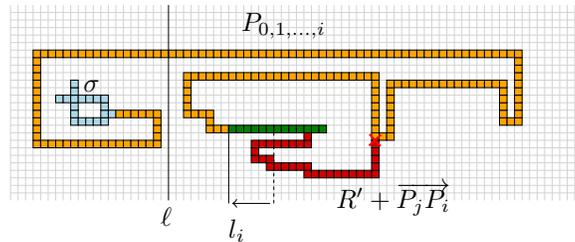


Figure 5.12: Case P7: Not all of $(Q_{j+1,j+2,\dots,r}R) + \overrightarrow{P_j P_i}$ is assemblable from $\sigma \cup \text{asm}(P_{0,1,\dots,i})$.

P8 We now consider the set \mathcal{S} of h -successful paths of the form $P_{0,1,\dots,i}XS$ for some prefix X of $(Q_{j+1,j+2,\dots,r}R') + \overrightarrow{P_j P_i}$ and some path S , such that all of the following hold:

²²Note that both of the cases (1) and (2) of Definition 5.9 can happen here.

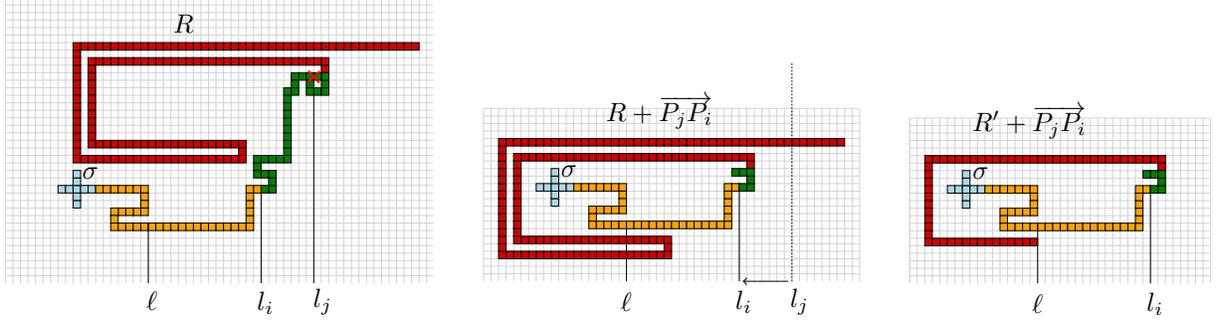


Figure 5.13: **Left:** an example path R from P7 (the red cross is an intersection of Q with itself). **Centre:** Example for P7, i.e. $R + \overrightarrow{P_j P_i}$ places a glue on ℓ below the visible glue of P on ℓ . **Right:** $R' + \overrightarrow{P_j P_i}$ is the largest prefix of $R + \overrightarrow{P_j P_i}$ that does not place a visible glue on ℓ below the visible glue of P on ℓ .

- $P_{0,1,\dots,i}XS$ turns right (respectively, left) from $P_{0,1,\dots,i}((Q_{j+1,j+2,\dots,r}R') + \overrightarrow{P_j P_i})$,
- $\text{glue}(P_i P_{i+1})$ is visible relative to $P_{0,1,\dots,i}XS$, and
- the visible glue of P on ℓ is visible relative to $P_{0,1,\dots,i}XS$.

There are two cases:

- If \mathcal{S} is empty this gives Conclusion 2 with $D = Q_{j+1,j+2,\dots,r}R_{0,1,\dots,|R'|} + \overrightarrow{P_j P_i}$ and $k = i$. In particular, as noted in case P7, D is an enclosing branch for P at i , which satisfies Conclusion 2(i). Moreover, the fact that \mathcal{S} is empty immediately shows Conclusion 2(ii) (intuitively, \mathcal{S} is the set of paths that do not meet that conclusion).
- Else, there is at least one path $P_{0,1,\dots,i}XS$ in \mathcal{S} . See Figure 5.14 for two examples.

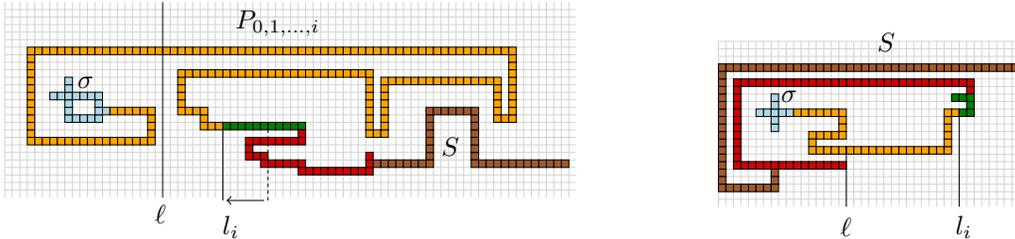


Figure 5.14: Two different example paths S that illustrate Case P8b: At least one h -successful path XS can turn right from $(Q_{j+1,j+2,\dots,r}R') + \overrightarrow{P_j P_i}$ without hiding the visibility of $\text{glue}(P_i P_{i+1})$ nor of the visible glue of P on ℓ .

P9 The only unresolved case after step P8 is therefore case P8b, which we now reason about with the goal of obtaining Conclusion 2.

We next grow the longest assemblable prefix of the “forward translated” segment $XS + \overrightarrow{P_i P_j} = (XS) + \overrightarrow{P_i P_j}$ and use this to show that in all remaining cases we get Conclusion 2. See Figure 5.15 for an example.

For notation, let $X' = X + \overrightarrow{P_i P_j}$ and $S' = S + \overrightarrow{P_i P_j}$. Suppose that the visible glue of P on ℓ is in V_P^+ (respectively, in V_P^-).

From P8b, S does not hide the visibility of $\text{glue}(P_i P_{i+1})$. Therefore, S' does not place a glue directly below $\text{mid}(P_j, P_{j+1})$ either. Notice that $P_{0,1,\dots,j}X'$ is assemblable (as it is a prefix of $P_{0,1,\dots,j}Q_{j+1,\dots,r}R$), and let s be the largest integer such that $P_{0,1,\dots,j}X'S'_{0,1,\dots,s-1}$ is

Proof. We apply Lemma 5.10. If we get Conclusion 1, we are immediately done (we get α in the statement). Else, let i, j, ℓ be as defined in Lemma 5.8, let $k \in \{i, j\}$ and let D be the enclosing branch constructed in Conclusion 2 of Lemma 5.10, assuming P places a V_P^+ (respectively, V_P^-) glue on ℓ . We begin by defining a connected component \mathcal{C} . There are two cases:

- If D intersects $\sigma \cup \text{asm}(P_{0,1,\dots,k})$ let \bar{P} be any path from $\text{pos}(D_{|D|-1})$ (the first such intersection) to $\text{pos}(P_k)$ in the grid graph of $\sigma \cup \text{asm}(P_{0,1,\dots,k})$. Let \mathcal{C} be the bounded connected component of \mathbb{R}^2 enclosed by the concatenation of $\mathfrak{E}_{\bar{P}}$ and \mathfrak{E}_D (the canonical embedding of the paths \bar{P} and D , respectively).
- Else, $D = D_{0,1,\dots,|D|-1}$ places a glue, denoted $\text{glue}(D_{|D|-2} D_{|D|-1})$, on ℓ below the visible glue of P on ℓ , and D does not intersect $\sigma \cup \text{asm}(P_{0,1,\dots,k})$. In this case, we let c be the concatenation of the following four curves (where $\text{glue}(P_\ell P_{\ell+1})$ is the visible glue of P on ℓ):

$$\begin{aligned} & [\text{mid}(P_\ell, P_{\ell+1}), \text{pos}(P_{\ell+1})] \\ & \mathfrak{E}_{P_{\ell+1},\dots,k D_{0,1,\dots,|D|-2}} \\ & [\text{pos}(D_{|D|-2}), \text{mid}(D_{|D|-2}, D_{|D|-1})] \\ & [\text{mid}(D_{|D|-2}, D_{|D|-1}), \text{mid}(P_\ell, P_{\ell+1})] \end{aligned}$$

By Observation 2.10, c is a finite closed simple curve and thus defines a bounded connected component \mathcal{C} .

In both cases, by the fact that D is an enclosing branch, $\text{glue}(P_k D_0)$ is visible relative to $P_{0,1,\dots,k} D_{0,1,\dots,|D|-2}$, which implies that the left-hand side (respectively, right-hand side) of $\text{glue}(P_k D_0)$ is *inside* \mathcal{C} (when walking in the direction from P_k to D_0).

Also, in both cases, $P_{0,1,\dots,k} D_{0,1,\dots,|D|-2}$ is not h -successful (by Lemma 5.10). Therefore, if P can still grow to be h -successful after $P_{0,1,\dots,k} D_{0,1,\dots,|D|-2}$ is grown, then $P_{k,k+1,\dots}$ turns right (respectively, left) from $P_k D_{0,1,\dots,|D|-2}$, and thus R from the statement of Lemma 5.10 is a suffix of P . But this implies that P places a lower visible glue than $\text{glue}(P_\ell P_{\ell+1})$ (on ℓ) and/or $\text{glue}(P_k D_0)$ (on the visibility ray l_k), which contradicts the visibility of $\text{glue}(P_\ell P_{\ell+1})$ and/or $\text{glue}(P_k P_{k+1}) = \text{glue}(P_k D_0)$ relative to P .

Therefore, P conflicts with $P_{0,1,\dots,k} D_{0,1,\dots,|D|-2}$ and thus P cannot grow to be h -successful from the assembly $\sigma \cup \text{asm}(P_{0,1,\dots,k} D_{0,1,\dots,|D|-2})$. \square

6 Blocking all paths

We restate our main theorem here:

Theorem 1.1. The noncooperative abstract tile assembly model is not intrinsically universal. In other words, there is no tileset U that at temperature 1 simulates all noncooperative tile assembly systems.

This result is an immediate corollary of Theorem 6.1 below. Intuitively, Theorem 6.1 states that there is no tile set that, at temperature 1, produces (or simulates) the “shapes” of all \mathcal{T}_N systems,²⁴ even if the simulator is allowed to use spatial rescaling. Thus Definition 2.1 is violated which immediately implies (via Observation A.2) that there is no tile set that, at temperature 1, simulates the productions of all \mathcal{T}_N systems (thus contradicting Definition A.1, “equivalent productions”), which in turn contradicts Definition A.6 (“intrinsically universal”, at temperature 1), giving Theorem 1.1.

²⁴The class of “flipped-L” tile assembly systems $\{\mathcal{T}_N \mid N \in \mathbb{Z}^+\}$ were defined earlier in Section 3.

Theorem 6.1. *There is no tileset U , scale factor $m \in \mathbb{Z}^+$, seed σ and m -block supertile representation function R_m such that for all $N \geq 10|U|$, $\text{dom}(R_m^*(\sigma)) = \text{dom}(\sigma_N)$ and $\{\text{dom}(R_m^*(\alpha)) \mid \alpha \in \mathcal{A}_\square[\mathcal{U}]\} = \{\text{dom}(\beta) \mid \beta \in \mathcal{A}_\square[\mathcal{T}_N]\}$ where $\mathcal{U} = (U, \sigma, 1)$ and $\mathcal{T}_N = (T_N, \sigma_N, 1)$.*

Proof. Assume, for the sake of contradiction, that there is a tileset U such that for $N = 10|U|$, there is an integer m , a seed assembly $\sigma \in \text{asm}(\mathcal{U})$, and an m -block representation function R_m such that the terminal assemblies of $\mathcal{U} = (U, \sigma, 1)$ map cleanly to the terminal assemblies of $\mathcal{T}_{10|U|}$ under R_m , where $\mathcal{T}_{10|U|}$ is the flipped-L tile assembly system defined in Definition 3.1.

We will show that \mathcal{U} also produces terminal assemblies mapping to non-terminal or non-producible assemblies of $\mathcal{T}_{10|U|}$ under R_m . More specifically, we will show that either some of the assemblies of \mathcal{U} map cleanly to *non-producible* assemblies of $\mathcal{T}_{10|U|}$ under R_m , or else we will construct one *producible* assembly $\alpha \in \mathcal{A}[\mathcal{U}]$ conflicting with all h -successful paths of \mathcal{U} . This will then conclude the proof since α grows into a *terminal* assembly, i.e. $\alpha \rightarrow^{\mathcal{U}} \alpha'$ where $\alpha' \in \mathcal{A}_\square[\mathcal{U}]$, that does not map cleanly to a *terminal* assembly of $\mathcal{A}_\square[\mathcal{T}_{10|U|}]$ under R_m (since all tiles of α' are below the horizontal line at height h).

Blocking h -successful paths individually For the remainder of the proof, let ℓ be a vertical (glue) line at x-coordinate $|U|(3m+1) + m + 1.5$ (in other words, at distance $\geq |U|(3m+1) + 1$ to the right of the rightmost tile of σ), as defined by Lemma 5.8. We apply Lemma 5.10 on each h -successful path P , individually²⁵. For each such P , Lemma 5.10 has one of two conclusions, numbered Conclusion 1 and Conclusion 2. If we get Conclusion 1 for any of the h -successful paths, we can conclude the proof immediately, because that conclusion shows that it is possible to grow a path from σ that places tiles outside of the simulation zone of \mathcal{U} , contradicting that \mathcal{U} simulates $\mathcal{T}_{10|U|}$, and hence assemblies of \mathcal{U} do not simulate the shape of $\mathcal{T}_{10|U|}$ and we are done with the proof of Theorem 6.1.

Therefore, in the rest of this proof, we assume that for *all* h -successful paths of \mathcal{U} we get Conclusion 2 of Lemma 5.10. That conclusion gives, for each h -successful path P , a nowhere- h -successful enclosing branch D for P at some integer k_P .

If it were the case that the entire set of these enclosing branches could grow together in the same assembly, we would immediately be done: indeed, the union of the seed with all of these enclosing branches (and their prefixes from P) would be an assembly conflicting with all h -successful paths of \mathcal{U} (implying in particular that this union does not contain any h -successful path).

The rest of the proof deals with the situation where this is not the case, i.e. at least one (and possibly very many) enclosing branches D from Lemma 5.10 conflict with other paths or with other enclosing branches, and thus not all enclosing branches D can grow completely together in the same assembly.

Path order We will build an assembly that does not reach height h and that blocks all of the paths from the set of h -successful paths $\mathbb{P}_{\mathcal{U}}$ of \mathcal{U} . Recall that the set of h -successful paths of \mathcal{U} is finite. In order to block them all, we will tackle h -successful paths in a specific order, called the “path order,” defined as follows. Let \prec be the *path order* relation on the set $\mathbb{P}_{\mathcal{U}}$ of h -successful paths of \mathcal{U} where for $P, Q \in \mathbb{P}_{\mathcal{U}}$ with $P \neq Q$ we say that $P \prec Q$ if and only if at least one of (A) or (B) holds:

²⁵By “individually” we mean that we are currently merely looking at the case where we grow each path separately: of course it may be the case that not all of these paths can be simultaneously grown as they may conflict with each other—the main point of this proof is to handle this.

- (A) the visible glue of P on ℓ is strictly higher than the visible glue of Q on ℓ , or
- (B) the visible glues $\text{glue}(P_p P_{p+1})$ and $\text{glue}(Q_q Q_{q+1})$ of P and Q on ℓ are at the same position²⁶ and one of the following holds:
- $\text{glue}(P_p P_{p+1}) \in V_P^+$ and $\text{glue}(Q_q Q_{q+1}) \in V_Q^-$, or
 - $\text{glue}(P_p P_{p+1}) \in V_P^+$, $\text{glue}(Q_q Q_{q+1}) \in V_Q^+$, and $Q_{q,q+1,\dots,|Q|-1}$ is the right-priority path of $P_{p,p+1,\dots,|P|-1}$ and $Q_{q,q+1,\dots,|Q|-1}$, or
 - $\text{glue}(P_p P_{p+1}) \in V_P^-$, $\text{glue}(Q_q Q_{q+1}) \in V_Q^-$, and $Q_{q,q+1,\dots,|Q|-1}$ is the left-priority path of $P_{p,p+1,\dots,|P|-1}$ and $Q_{q,q+1,\dots,|Q|-1}$, or
 - Else, notice that P and Q share their suffix from their visible glue on ℓ onwards until their last tile at height h (because none of these suffixes is the right-priority or left-priority one). Then P is the right priority path of P and Q if P and Q share two consecutive tiles $P_a P_{a+1} = Q_b Q_{b+1}$ before disagreeing (note that $P \neq Q$), and if they do not share such a pair then P is the lexicographically first path of P and Q if we describe both using some canonical encoding of P and Q as two binary strings.

Here is an intuitive description of the path order: we first consider paths by the height of their visible glue on ℓ (highest first), and then if both visible glues on ℓ are in the same direction, we first consider the most right-priority of P and Q after they cross ℓ if these glues are in V_P^+ and V_Q^+ , or the most left-priority if these glues are in V_P^- and V_Q^- , and if the glues are at the same position with different +/- orientations, the one with a “+” visible glue on ℓ comes first. Finally if P and Q happen to agree (are equal) on their suffix from their visible glue on ℓ onwards, then we (arbitrarily) choose the right priority path (note that in this latter case all of the differences between P and Q must be before their respective visible glues on ℓ).

Note that the relation \prec is a total order on the set \mathbb{P}_U of h -successful paths, since the last case of the definition of \prec covers all remaining cases using right-priority, and right-priority is itself a total order. Also, recall that the set of h -successful paths is a finite number (see Section 5.2), and let H be that number.

Thus let $P^0 \prec P^1 \prec P^2 \prec \dots \prec P^{H-1}$ be the list of all h -successful paths according to path order (so that no path has a higher visible glue on ℓ than P^0).

Enclosing branch D^n . For each path P^n , applying Lemma 5.10 gives an index k_n and an “enclosing branch D^n for P^n at k_n ” such that $P_{0,1,\dots,k_n}^n D_{0,1,\dots,|D|-2}^n$ conflicts with P^n (by Theorem 5.11). Let

$$E^n \stackrel{\text{def}}{=} P_{0,1,\dots,k_n}^n D_{0,1,\dots,|D|-2}^n$$

The “path blocking” assembly α_n . Let the notation $\alpha \rightsquigarrow P$ denote the path that is the longest assemblable prefix of P that can be grown from the assembly α .²⁷

We define an assembly α_n which has a special form (composed of σ and assemblable prefixes of E^k grown in path, i.e. k , order), to be used in our induction hypothesis:

$$\alpha_n = \sigma \cup \left(\bigcup_{k=0}^n \text{asm}(F^k) \right)$$

where $F^0 = \sigma \rightsquigarrow E^0 = E_{0,1,\dots,|E^0|-1}^0$ and for all $k \geq 1$, $F^k = \alpha_{k-1} \rightsquigarrow E^k$.

²⁶I.e. $\text{mid}(P_p, P_{p+1}) = \text{mid}(Q_q, Q_{q+1})$.

²⁷Observe that if α is producible by some tile assembly system then for all paths P it is (trivially) the case that $\alpha \cup \text{asm}(\alpha \rightsquigarrow P)$ is an assembly producible by that same tile assembly system.

Claim: for all $n \geq 0$, $\alpha_n \in \mathcal{A}[\mathcal{U}]$. To see this claim note that:

- First, α_0 is producible: indeed, $F^0 = E^0 = P_{0,1,\dots,k_0}^0 D_{0,1,\dots,|D^0|-2}^0$ is a producible path of \mathcal{U} , by Lemma 5.10. Therefore, $\alpha_0 = \sigma \cup \text{asm}(F^0)$ is producible.
- Then, assuming α_n is producible, i.e. $\alpha_n \in \mathcal{A}[\mathcal{U}]$, remember that $\alpha_n \rightsquigarrow E^{n+1}$ is the maximal prefix of E^{n+1} that can grow from α_n . Therefore, $\alpha_n \rightarrow^{\mathcal{U}} (\alpha_n \cup \text{asm}(\alpha_n \rightsquigarrow F^{n+1})) = \alpha_{n+1}$, and therefore $\alpha_{n+1} \in \mathcal{A}[\mathcal{U}]$.

Hence $\alpha_n \in \mathcal{A}[\mathcal{U}]$ as claimed.

To conclude the proof we will consider the assembly²⁸ $\alpha_{H-1} \in \mathcal{A}[\mathcal{U}]$ which we claim has the (as yet unproven) property that all producible h -successful paths conflict with it. Then allowing tiles to attach to α_{H-1} will eventually yield²⁹ a terminal assembly $\alpha \in \mathcal{A}_{\square}[\mathcal{U}]$ with no tiles at height h and thus no tiles above height h , which contradicts that \mathcal{U} simulates (the shape of) $\mathcal{T}_{10|U|}$. We will use induction to show that all producible h -successful paths conflict with α_{H-1} .

Induction hypothesis: All of the paths P^0, P^1, \dots, P^n conflict with α_n .

Some intuition and implications of our induction hypothesis: The induction hypothesis implies that for $\forall k \leq n$, $\alpha_n \rightsquigarrow P^k$ is not h -successful. To see this note that since for all k the last tile of P^k is its only tile at height h , and the induction hypothesis implies that P^k can not grow from α_n to be h -successful. Also, no tile of α_n reaches height h (because E^k is constructed via Lemma 5.10) which implies that none of the F^k (of which α_n is composed) are h -successful. Finally, since there are a finite number H of h -successful paths the induction exhausts those H paths in H steps, and thus yields an assembly α_{H-1} which is a finite union of finite (path) assemblies, and thus α_{H-1} is a finite producible assembly that blocks all h -successful paths.

Initial step of induction (P^0 and α_0). At the initial step of the induction, we apply Lemma 5.10 to P^0 , to obtain an enclosing branch $P_{0,1,\dots,k_0}^0 D_{0,1,\dots,|D^0|-1}^0$. This proves our induction hypothesis for the initial step: by Theorem 5.11, P^0 conflicts with α_0 (i.e. P^0 cannot grow to be h -successful from α_0), and we have already defined $\alpha_0 = \sigma \cup \text{asm}(F^0)$ where $F^0 = \sigma \rightsquigarrow E^0 = E_{0,1,\dots,|E^0|-2}^0 = P_{0,1,\dots,k_0}^0 D_{0,1,\dots,|D^0|-2}^0$.

Inductive step (P^{n+1} and α_{n+1}). The remainder of the proof is concerned with the inductive step. For any $n \geq 0$ suppose the induction hypothesis holds,³⁰ i.e. all of first $n+1$ paths P^0, P^1, \dots, P^n conflict with α_n .

We recall that $F^{n+1} = \alpha_n \rightsquigarrow E^{n+1}$ is the maximal prefix of E^{n+1} that can grow from α_n , and that $\alpha_{n+1} = \alpha_n \cup \text{asm}(F^{n+1})$. If P^{n+1} conflicts with α_n , then we are immediately done with the induction step for α_{n+1} , because this proves that P^{n+1} cannot grow from α_{n+1} . Hence from now we will assume that P^{n+1} does not conflict with α_n .

If P^{n+1} conflicts with F^{n+1} , then we are immediately done with the induction step for α_{n+1} , because this proves that P^{n+1} cannot grow from α_{n+1} . Otherwise P^{n+1} does not conflict with F^{n+1} . This implies that F^{n+1} is a *strict* prefix of E^{n+1} (otherwise we would contradict

²⁸Recall that H is the number of h -successful paths of \mathcal{U} .

²⁹Growth can only happen within the finite area simulation zone below height h so must eventually stop.

³⁰Recall that α_n contains only the seed σ and assembled paths F^0, F^1, \dots, F^n (i.e. $\alpha_n = \sigma \cup (\bigcup_{k=0}^n \text{asm}(F^k))$) that are respective prefixes of E^0, E^1, \dots, E^n , none of which are h -successful.

Theorem 5.11) and therefore E^{n+1} (and in particular D^{n+1}) conflicts with F^k for some $k \in \{0, 1, \dots, n\}$. We will reason about this F^k .

We next split the inductive step into three cases³¹ each of which will be concluded independently:

- (Case 1) P^{n+1} and F^k share the position of their visible glue on ℓ , one of these glues is a V^+ glue, the other one is a V^- glue.
- (Case 2) P^{n+1} and F^k do not share the position of their visible glue on ℓ (which includes the case where F^k does not reach ℓ).
- (Case 3) P^{n+1} and F^k share the position of their visible glue on ℓ , and either both are V^+ glues, or both are V^- glues.

In all three cases, let F^k 's visible glue on ℓ be denoted $\text{glue}(F_f^k F_{f+1}^k)$ and let P^{n+1} 's visible glue on ℓ be denoted $\text{glue}(P_p^{n+1} P_{p+1}^{n+1})$.

Case 1: P^{n+1} and F^k share the position of their visible glue on ℓ , one of these glues is a V^+ glue, the other one is a V^- glue. Exactly one of $\text{glue}(F_f^k F_{f+1}^k)$ and $\text{glue}(P_p^{n+1} P_{p+1}^{n+1})$ is in V^+ and the other is in V^- . (See Figure 6.1 for an example.) At the beginning of the inductive step, we assumed that P^{n+1} does not conflict with α_n , hence P^{n+1} and F^k agree on all points where they intersect. Moreover, since P^{n+1} and F^k share their visible glue on ℓ (i.e. $\text{mid}(P_p^{n+1}, P_{p+1}^{n+1}) = \text{mid}(F_f^k, F_{f+1}^k)$, i.e. their visible glues on ℓ are at the same position) we know they agree on at least two tiles each, specifically $F_f^k = P_{p+1}^{n+1}$ and $F_{f+1}^k = P_p^{n+1}$. Now let b be the largest integer such that there is an integer $a \leq f$ where $F_a^k = P_b^{n+1}$. Since $\text{mid}(P_p^{n+1}, P_{p+1}^{n+1}) = \text{mid}(F_f^k, F_{f+1}^k)$ we know that $p+1 \leq b$, also $b \leq |P^{n+1}| - 2$ since P^{n+1} is h -successful and E^k is not. Consider the sequence $Q = F_{0,1,\dots,a}^k P_{b+1,b+2,\dots,|P|-1}^{n+1}$. First note that the positions of Q form a connected sequence of positions in \mathbb{Z}^2 : this follows from the fact that the positions of $F_{0,1,\dots,a}^k$ are connected, the positions of $P_{b+1,b+2,\dots,|P|-1}^{n+1}$ are connected, and that $F_a^k = P_b^{n+1}$. Also, we claim that Q is simple: to see this, note that (i) $F_{0,1,\dots,a}^k$ is simple, (ii) $P_{b+1,b+2,\dots,|P|-1}^{n+1}$ is simple, and finally that (iii) $P_{b+1,b+2,\dots,|P|-1}^{n+1}$ does not intersect $F_{0,1,\dots,a}^k$ (by definition of b). Since Q has a connected simple set of positions, Q is a path. Furthermore it is the case that $Q \in \mathbf{P}[\mathcal{U}]$, which follows immediately from the following facts: Q is a path, $F_{0,1,\dots,a}^k \in \mathbf{P}[\mathcal{U}]$, $P^{n+1} \in \mathbf{P}[\mathcal{U}]$ and $F_a^k = P_b^{n+1}$.

Next we claim that $Q \prec P^{n+1}$. First, we know that, since $b \geq p+1$ all of the glues that $P_{b+1,b+2,\dots}^{n+1}$ places on ℓ are at height strictly higher than the height of $\text{glue}(P_p^{n+1} P_{p+1}^{n+1})$ on ℓ which is P^{n+1} 's visible (i.e. lowest) glue on ℓ . Second, since $a \leq f$ we know that all of the glues that $F_{0,1,\dots,f}^k$ places on ℓ are at height strictly higher than the height of $\text{glue}(F_f^k F_{f+1}^k)$ on ℓ which is F^k 's visible glue on ℓ which is at the same height as P^{n+1} 's visible glue on ℓ . Since $Q = F_{0,1,\dots,a}^k P_{b+1,b+2,\dots}^{n+1}$, then Q 's visible glue on ℓ is strictly higher than the visible glue of P^{n+1} on ℓ . Thus $Q \prec P^{n+1}$.

Since P^{n+1} is h -successful, and since P^{n+1} and Q share a nonempty suffix $P_{b+1,b+2,\dots,|P^{n+1}|-1}^{n+1}$, this implies that Q is also h -successful. Moreover, no strict prefix of Q is h -successful, because F^k is not h -successful and by the definition of h -successful no strict prefix of P^{n+1} is h -successful. But since $Q \prec P^{n+1}$, this means that Q satisfies the induction hypothesis, meaning that Q conflicts with α_n . Since $\text{asm}(F_{0,1,\dots,a}^k)$ is a subassembly of α_n then the prefix $F_{0,1,\dots,a}^k$ of Q does

³¹For the sake of proof simplicity, we present them in the order in which we handle these cases.

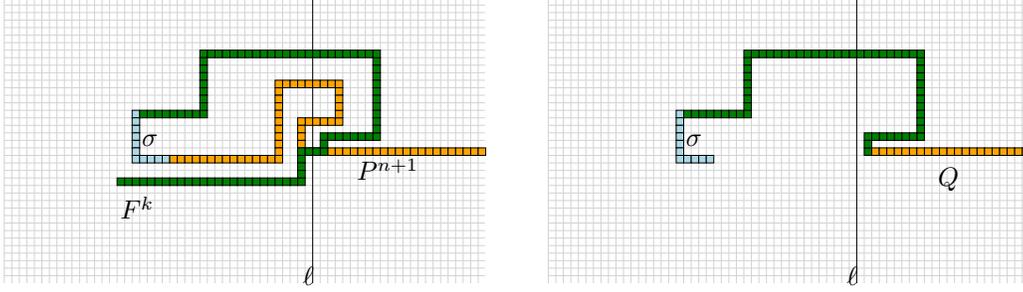


Figure 6.1: **Left:** An example of Case 1, where F^k places a V^- glue on ℓ . **Right:** The path $Q = F_{0,1,\dots,a}^k P_{b+1,b+2,\dots}^{n+1}$ which has the property that it comes earlier in path order than P^{n+1} , i.e. $Q \prec P^{n+1}$, since the visible glue of Q on ℓ is higher than the visible glue of P^{n+1} on ℓ .

not conflict with α_n , which in turn implies that the suffix $P_{b+1,b+2,\dots,|P^{n+1}|-1}^{n+1}$ of Q conflicts with α_n , which implies that P^{n+1} conflicts with α_n , satisfying the induction hypothesis.³²

Case 2: P^{n+1} and F^k do not share the position of their visible glue on ℓ (which includes the case where F^k does not reach ℓ). Moreover, the glue placed by P^{n+1} on ℓ is a $V_{P^{n+1}}^+$ glue (respectively a $V_{P^{n+1}}^-$ glue). We assumed that P^{n+1} does not conflict with α_n , hence P^{n+1} does not conflict with F^k . We first show that $P_{p+1,p+2,\dots,|P^{n+1}|-1}^{n+1}$ intersects and agrees with F^k , and then use an argument similar to Case 1 above:

- Assume, for the sake of contradiction, that $P_{p+1,p+2,\dots,|P^{n+1}|-1}^{n+1}$ does not intersect F^k . Let $a \geq p$ be the smallest³³ integer such that $\text{pos}(E_a^{n+1}) = \text{pos}(F_b^k)$ for some b . We are going to define a closed connected component in which P^{n+1} starts to grow. First note that F^k is connected, connected to σ , and E^{n+1} is connected to σ . Therefore, $\sigma \cup \text{asm}(F^k) \cup \text{asm}(P_{0,1,\dots,p}^{n+1})$ contains at least one path from F_b^k to P_p^{n+1} (note that $P_{0,1,\dots,p}^{n+1}$ is a prefix of E^{n+1}). Let \mathcal{P} be any shortest such path.

Let then c be the closed curve defined by the concatenation of $c^k = \mathfrak{E}_{\mathcal{P}}$ and $c^{n+1} = \mathfrak{E}_{E_{p,p+1,\dots,a}^{n+1}}$.

Curve c is simple because c^{n+1} and c^k only intersect at their endpoints because a was chosen to be the smallest integer ($\geq p$) such that $\text{pos}(E_a^{n+1}) = \text{pos}(F_b^k)$ and because \mathcal{P} is a shortest path.

Therefore, by the Jordan Curve Theorem, c encloses a single bounded connected component \mathcal{C} of \mathbb{R}^2 . (This connected component is shown in gray in the example in Figure 6.2.)

Now, $\text{mid}(P_p^{n+1}, P_{p+1}^{n+1})$, the position of the visible glue $\text{glue}(P_p^{n+1} P_{p+1}^{n+1})$ of P^{n+1} on ℓ , is on curve c . Moreover, since no other point of c intersects ℓ at the height of, or below $\text{mid}(P_p^{n+1}, P_{p+1}^{n+1})$, then $\text{mid}(P_p^{n+1}, P_{p+1}^{n+1})$ is the unique lowest intersection of c and ℓ .

Then, since $\text{glue}(P_p^{n+1} P_{p+1}^{n+1})$ is in $V_{P^{n+1}}^+$ (respectively in $V_{P^{n+1}}^-$), the left-hand side (respectively right-hand side) of c is inside \mathcal{C} . Therefore, since P^{n+1} is h -successful and places tiles (with positions) on c , $P_{p,p+1,\dots,|P|-1}^{n+1}$ needs to turn from c (because all points of c are below height h). However, by Lemma 5.10, P^{n+1} cannot turn right (respectively, left) from D^{n+1} , hence from E^{n+1} ; if it did P^{n+1} would hide at least one of its own visible glues, which is

³²Recall that we have already defined $F^{n+1} = \alpha_n \rightsquigarrow E^{n+1}$ and $\alpha_{n+1} = \alpha_n \cup \text{asm}(F^{n+1})$.

³³There is at least one such integer since we know that E^{n+1} conflicts with F^k , and we know that this conflict happens after (in E^{n+1} order) the visible glue ($\text{glue}(P_p^{n+1} P_{p+1}^{n+1})$) of P^{n+1} and E^{n+1} .

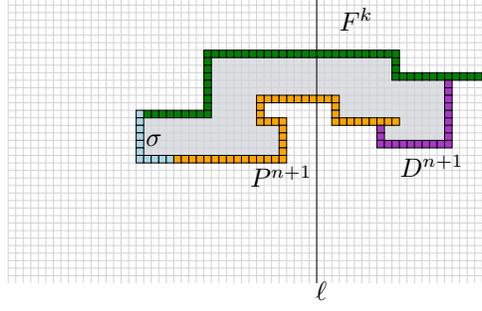


Figure 6.2: An example of Case 2, where P^{n+1} and F^k both place a V^+ glue on ℓ , at different heights.

impossible. Therefore, $P_{p,p+1,\dots,|P^{n+1}|-1}^{n+1}$ must turn from, and thus intersect, other parts of c , i.e. σ or F^k , which is a contradiction. Hence $P_{p,p+1,\dots,|P^{n+1}|-1}^{n+1}$ intersects F^k .

- We have shown that $P_{p+1,p+2,\dots,|P^{n+1}|-1}^{n+1}$ intersects (and agrees with) F^k at least once. In fact all such intersections are agreements because P^{n+1} does not conflict with F^k . Let $a \geq p+1$ be the largest integer such that $P_a^{n+1} = F_b^k$ for some integer b . We claim that $Q = F_{0,1,\dots,b}^k P_{a+1,a+2,\dots}^{n+1}$ is a path: indeed, $P_{a+1,a+2,\dots}^{n+1}$ does not intersect $F_{0,1,\dots,b}^k$ by the definition of a , and Q is connected. Moreover, $Q \in \mathbf{P}[\mathcal{U}]$. Furthermore, $Q \prec P^{n+1}$, because since $a > p$ and P^{n+1} is simple, and since the visible glue of F^k on ℓ is not shared with that of P^{n+1} , all glues of Q on ℓ are strictly higher than $\text{glue}(P_p^{n+1} P_{p+1}^{n+1})$. Therefore, by the induction hypothesis, Q conflicts with α_n , and hence α_{n+1} , which means that $P_{a+1,a+2,\dots}^{n+1}$ conflicts with α_{n+1} (since $F_{0,1,\dots,b}^k$ does not conflict with α_{n+1}).

Case 3: P^{n+1} and F^k share the position of their visible glue on ℓ , and either both are V^+ glues, or both are V^- glues. In this case, because $P^k \prec P^{n+1}$, and P^{n+1} and F^k , and hence P^k , share their visible glue at the same height on ℓ , we know by the definition of \prec that either:

- $P_{p,p+1,\dots}^{n+1}$ is more right-priority (respectively, left-priority) than $P_{f,f+1,\dots}^k$ if $\text{glue}(P_p^{n+1} P_{p+1}^{n+1})$, $\text{glue}(P_f^k P_{f+1}^k)$ are both V^+ glues (respectively, V^- glues), where $\text{glue}(P_p^{n+1} P_{p+1}^{n+1})$ is the visible glue of P^{n+1} on ℓ , and $\text{glue}(P_f^k P_{f+1}^k)$ is the visible glue of P^k on ℓ .
- $P_{p,p+1,\dots}^{n+1} = P_{f,f+1,\dots}^k$.

However, in the second case, since α_n conflicts with P^k (by the induction hypothesis, since $k \leq n$), and F^k places the visible glue of P^k on ℓ , then α_n conflicts with $P_{f,f+1,\dots}^k$, hence α_n also conflicts with $P_{p,p+1,\dots}^{n+1}$, and we are done with Case 3 by simply letting $\alpha_{n+1} = \alpha_n \cup \text{asm}(\alpha_n \rightsquigarrow E^{n+1})$.

To conclude this proof, we will therefore handle the first case, i.e. the case where $P_{p,p+1,\dots}^{n+1}$ is more right-priority (respectively, left-priority) than $P_{f,f+1,\dots}^k$.

We assumed that P^{n+1} does not conflict with α_n , hence in particular $P_{p,p+1,\dots}^{n+1}$ does not conflict with $F_{f,f+1,\dots}^k$, and does not conflict with $P_{f,f+1,\dots}^k$. Let $q \geq p$ and $g \geq f$ be the smallest integers such that $P_q^{n+1} \neq F_g^k$. An example of this situation is shown in Figure 6.3.

The argument follows along the same lines as Case 2 (building a closed connected component in which P^{n+1} starts to grow), but requires a new technique to identify the inside and outside

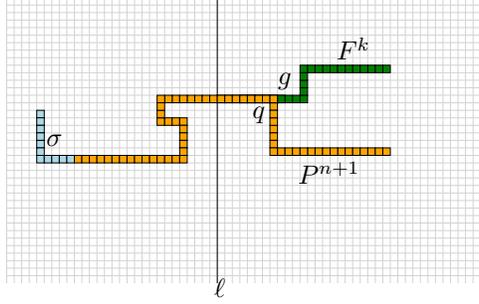


Figure 6.3: An example of Case 3, where P^{n+1} and F^k (and by implication P^k) both place a V^+ glue on ℓ , at the same height.

of that connected component.

- Assume, for the sake of contradiction, that $P_{q,q+1,\dots}^{n+1}$ does not intersect F^k .

We will now describe a closed connected component inside which a suffix of P^{n+1} starts to grow.

We first introduce a new variant of embedding of paths into \mathbb{R}^2 , which we call the *nano-embedding of a path* P , denoted \mathfrak{N}_P . This is illustrated in Figure 6.4 and defined as follows. For a path P consider its canonical embedding \mathfrak{E}_P . Then, define \mathfrak{N}_P to be the curve in \mathbb{R}^2 where all of the points of \mathfrak{N}_P are at L_∞ distance exactly 0.25 from their closest point on \mathfrak{E}_P , and are positioned on the right (respectively, left) hand side of \mathfrak{E}_P as we walk along \mathfrak{E}_P from $\mathfrak{E}_P(0)$ to $\mathfrak{E}_P(1)$. For tiles on P with input side being their west side, we show in Figure 6.4(top) all three cases of nano-embeddings. Other cases where the input side is north, east or south are rotations of these three cases. Special cases for start and end tiles of a path are illustrated in Figure 6.4.

Since E^{n+1} intersects F^k (because in particular, E^{n+1} conflicts with F^k), let $b > q$ be the smallest integer such that $\text{pos}(E_b^{n+1}) = \text{pos}(F_d^k)$ for some integer d .

We now define a simple closed curve c inside which a suffix of P^{n+1} starts to grow: let c be the concatenation of $\mathfrak{N}_{F_{p,p+1,\dots,b-1}^{n+1}}$, then a length < 1 line segment from the final point of $\mathfrak{N}_{F_{p,p+1,\dots,b-1}^{n+1}}$ to $\text{pos}(F_d^k)$, then $\mathfrak{E}_{F_{f,f+1,\dots,d}^k}^{\leftarrow}$, and finally a line segment of length 0.25 from $\text{pos}(F_f^k)$ to $\mathfrak{N}_{F_{p,p+1,\dots,b-1}^{n+1}}(0)$, which is the first point of $\mathfrak{N}_{F_{p,p+1,\dots,b-1}^{n+1}}$. (Figure 6.5 shows an example c .)

We claim that c is simple: indeed, since F^k only turns left from F^{n+1} , and since $\mathfrak{N}_{F^{n+1}}$ stays immediately to the right of F^{n+1} , the four curves used to construct c intersect each other only at the last and first endpoints of each pair of consecutive curves. Notice that c is also closed. Therefore, by the Jordan Curve Theorem, c encloses a bounded connected component \mathcal{C} of \mathbb{R}^2 . Moreover, at the visible glue of P^{n+1} and P^k on ℓ , the nano-embedding $\mathfrak{N}_{P^{n+1}}$ of P^{n+1} is below $\mathfrak{E}_{P^{n+1}}$, and since P^{n+1} places a $V_{P^{n+1}}^+$ (respectively, $V_{P^{n+1}}^-$) glue on ℓ , the left-hand (respectively, right-hand) side of c is the inside of \mathcal{C} .

Finally, since P^{n+1} does not turn right from F^{n+1} (otherwise, by Lemma 5.10 P^{n+1} would hide the visibility of at least one of its own glues, which is impossible), a suffix of P^{n+1} starts inside \mathcal{C} or on c . But since no point of \mathcal{C} is at or above height h , and P^{n+1} is h -successful, the last tile of P^{n+1} is positioned outside of \mathcal{C} , which can happen in only two different ways: either P^{n+1} turns right from $F_{p,p+1,\dots,b-1}^{n+1}$ (contradicting Lemma 5.10), or else $\mathfrak{E}_{P_{q,q+1,\dots}^{n+1}}$ intersects

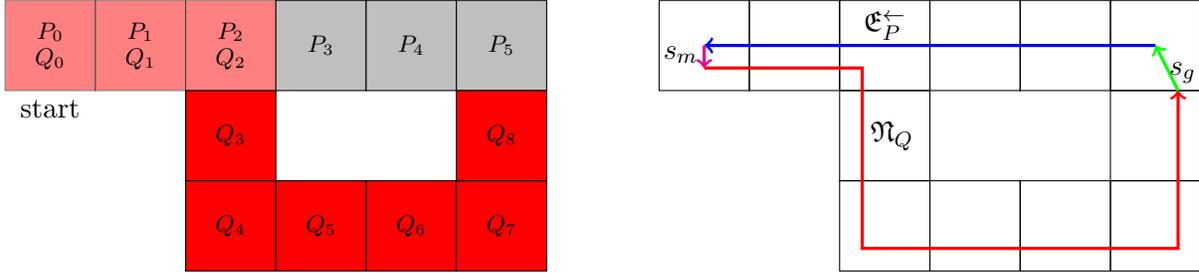
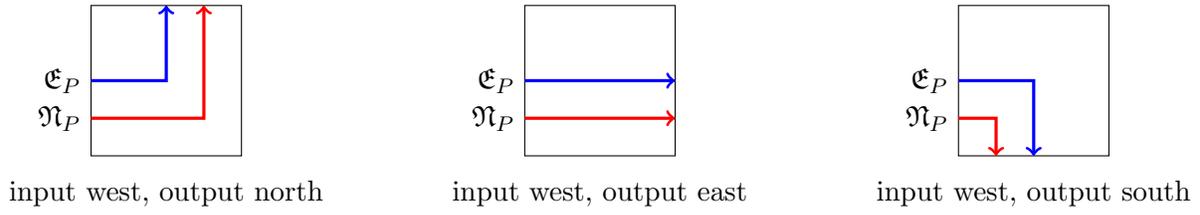


Figure 6.4: Nano-embedding and canonical embedding of a path in \mathbb{R}^2 . Top: Three tiles on some path P that have their input side as their west side, and their output sides as north, east and south respectively, as indicated. In each of the three cases, the canonical embedding \mathcal{E}_P of P is shown in blue, and the nano-embedding \mathfrak{N}_P is shown in red. Each point in \mathfrak{N}_P is L_∞ distance exactly 0.25 from its closest point in \mathcal{E}_P . Rotating these diagrams by 90° , 180° and 270° give the other 9 cases needed to define the nano-embedding of any tile on a path, except for the first and last tile. Bottom left: two paths that start at the common tile $P_0 = Q_0$; path P is shown in pink and grey, Q is shown in pink and red. Bottom right: example showing how we combine the canonical embedding \mathcal{E}_P of the path P and the nano-embedding \mathfrak{N}_Q of the path Q to make a simple closed curve c in \mathbb{R}^2 . The start tile on the left is a special case (in our construction it is always the case that start tile of a nano-embedding has its output side on the east). For the start tile a short vertical magenta segment s_m is drawn so that it ends at the start point of \mathfrak{N}_Q . We follow \mathfrak{N}_Q until it ends (“in” the last tile of Q), then draw a short green segment s_g to the position (“center”) of the last tile of P . From there the reverse of \mathcal{E}_P , denoted \mathcal{E}_P^\leftarrow , traces backwards through the positions of tiles of P to the start point of the magenta segment. The resulting curve c is the concatenation of the curves $s_m, \mathfrak{N}_Q, s_g, \mathcal{E}_P^\leftarrow$ and is a simple closed curve in \mathbb{R}^2 .

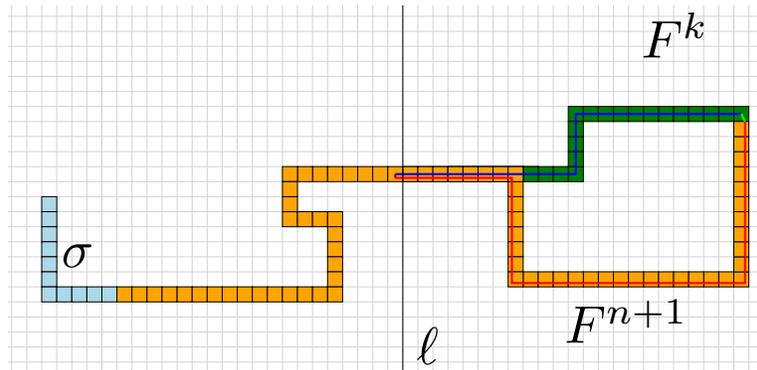


Figure 6.5: An example of case 3.1, showing the paths F^{n+1} and F^k . The nano-embedding of F^{n+1} is shown as the red curve, and the canonical embedding of F^k is shown as the blue curve. Together with the two small (length ≤ 1) magenta and green segments, these four curves form a simple closed curve c in \mathbb{R}^2 .

$\mathfrak{C}_{F_{g,g+1,\dots,d}^k}^{\leftarrow}$ which is also a contradiction.

- Therefore, $P_{q,q+1,\dots}^{n+1}$ intersects F^k . Moreover, that intersection is necessarily an agreement. Let therefore q' be the largest integer such that $P_{q'}^{n+1} = F_{g'}^k$ for some $g' \geq 0$ (notice that $q' \geq q$), and let $Q = F_{0,1,\dots,g'}^k P_{q'+1,q'+2,\dots,|P^{n+1}|-1}^{n+1}$. Note that Q is connected, and that $P_{q'+1,q'+2,\dots}^{n+1}$ does not intersect $F_{0,1,\dots,g'}^k$ (because of our choice of q'), and that the tiles $F_{g'}^k$ and $P_{q'+1}^{n+1}$ bind (since $P_{q'}^{n+1} = F_{g'}^k$). Therefore, Q is an assemblable path in \mathcal{U} .
 - If $g' \geq f$, then Q and P^{n+1} have the same visible glue on ℓ (this visible glue is on $F_{0,1,\dots,g'}^k$), and the first difference between $Q_{f,f+1,\dots}$ and $P_{p,p+1,\dots}^{n+1}$ is a right turn of $P_{p,p+1,\dots}^{n+1}$ from $Q_{f,f+1,\dots}$, meaning that $Q_{f,f+1,\dots}$ is less right-priority than $P_{p,p+1,\dots}^{n+1}$. Therefore, $Q \prec P^{n+1}$.
 - Else, $g' < f$, and hence the visible glue of Q on ℓ is not the same as the visible glue of P^{n+1} on ℓ . Therefore, the visible glue of Q on ℓ is strictly higher than that of P^{n+1} . This means that $Q \prec P^{n+1}$.

In both cases, $Q \prec P^{n+1}$, hence Q conflicts with α_n by the induction hypothesis. Recall that $Q = F_{0,1,\dots,g'}^k P_{q'+1,q'+2,\dots,|P^{n+1}|-1}^{n+1}$, and that $\text{asm}(F^k) \sqsubseteq \alpha_n$, and therefore $P_{q'+1,q'+2,\dots,|P^{n+1}|-1}^{n+1}$ conflicts with α_n .

Hence P^{n+1} conflicts with α_n which proves the induction hypothesis for α_{n+1} . □

7 Noncooperative tile assembly: Impossibility of bounded Turing machine simulation

We begin by restating Theorem 1.2. Note that the “bounding function” B_M in the statement is an arbitrary upperbound on the space usage of the Turing machine M as we wish to allow any claimed temperature 1 simulator of Turing machines to be arbitrarily 2D-space-inefficient in its attempt to do so.

Theorem 1.2. Let $t : \mathbb{N} \rightarrow \mathbb{N}$, $s : \mathbb{N} \rightarrow \mathbb{N}$ and let $B_M : \mathbb{N} \rightarrow \mathbb{N}$ such that $\forall n \in \mathbb{N}$, $B_M(n) \geq s(n)$. Let M be any Turing machine that halts on all inputs $x \in \{0, 1\}^*$ in time $t(|x|)$ using space $s(|x|)$. There is no pair (V, B_M) where V is a tiling set and B_M is a function such that for all $x \in \{0, 1\}^*$, $|x| = n$, there is a seed assembly $\sigma_{M,x}$ and tile assembly system $\mathcal{V}_x = (V, \sigma_{M,x}, 1)$ such that:

1. $\text{dom}(\sigma_{M,x}) \subseteq \{0, 1, \dots, B_M(n) - 1\} \times \{0, 1, \dots, B_M(n) - 1\}$
2. for all $\alpha \in \mathcal{A}_{\square}[\mathcal{V}]$, $\text{dom}(\alpha) \subseteq \{0, 1, \dots, t(n)B_M(n) - 1\} \times \{0, 1, \dots, B_M(n) - 1\}$, $\text{dom}(\alpha) \cap (\{b + 1, b + 2, \dots, b + B_M(n) - 1\} \times \{0, 1, \dots, B_M(n) - 1\}) \neq \emptyset$ where $b = B_M(n)(t(n) - 1)$ and α places at least one occurrence of a special tile type $H \in V$ on the rightmost column, and nowhere else, of $\text{dom}(\alpha)$ if and only if M accepts x .

Proof. Intuitively, the proof proceeds by supposing for the sake of contradiction that there is such a tiling set V and then modifying V to get another tile set that that can be instantiated as an infinite set of tile assembly systems each one of which produces terminal assemblies that have the same scaled (simulated) shape as some system \mathcal{T}_N defined in Section 3. But this violates Theorem 6.1, giving a contradiction. We argue this as follows.

Let M be a Turing machine with input alphabet $\{1\}^*$ that accepts all of its inputs $x \in \{1\}^n$ using space $s(n)$ and time $t(n)$. So suppose for the sake of contradiction that there is a tiling set V

that simulates M on all inputs using some “bounding function” B_M as described in the theorem statement.

We will modify the tile set V . Since the tile type H is on the rightmost vertical column of every *terminal* assembly α , H 's east glue type g_E is either (a) of strength 0, or else (b) of strength 1 and matches no west glue in the tile set V . On tile type H we replace g_E with a new glue type g'_E that is of strength 1 and where g'_E appears on no other tile type of V . We also add two new tile types t_1, t_2 to V : the west side of t_1 has the glue type g'_E and so binds to the east side of H , and the south side of t_2 binds to the north side of t_1 , and the south side of t_2 binds to the north side of itself. When H appears in some assembly α it is always possible to bind a tile of type t_1 to the east side of H in α (by hypothesis H is placed in the rightmost column of α hence there is always sufficient (unit) space to the right of H to place t_1). There may be a number of places where tiles of type t_1 can bind (each to the east of a tile of type H), nevertheless every terminal assembly will have an infinite vertical line of t_2 tiles growing to the north of some instance of t_1 (i.e. since H is in the rightmost column, there can be nothing to the north of a tile of type t_1 , other than possibly another tile of type t_1 , and hence there is nothing to stop some tile of type t_1 growing an the infinite vertical line of t_2 tiles to its north).

For each x , $|x| = n$, using the modified tileset V , the system $\mathcal{V}_x = (V, \sigma_{M,x}, 1)$ builds an assembly that has (roughly) the same rescaled shape as $\mathcal{T}_{B_M(n)/B_M(n)t(n)} = \mathcal{T}_{t(n)}$ (defined in Section 3), but with some spatial rescaling (by a factor of $B_M(n)$). Let x be any input such that, with $|x| = n$, $t(n) \geq 10|V|$ (e.g. choosing $|x| \geq 10|V|$ does the trick). Finally, setting $m = B_M(n)$, let R_m be the m -block supertile representation function that is undefined on empty m -blocks and maps nonempty m -blocks to the tile $\sigma \in T$.³⁴ Then $\text{dom}(R_m^*(\sigma_{\mathcal{V}_x})) = \text{dom}(\sigma_{\mathcal{T}_n})$ and $\{\text{dom}(R_m^*(\alpha)) \mid \alpha \in \mathcal{A}_{\square}[\mathcal{V}_x]\} = \{\text{dom}(\beta) \mid \beta \in \mathcal{A}_{\square}[\mathcal{T}_n]\}$ for $n \geq 10|V|$. Hence our modified V violates Theorem 6.1 where in the theorem statement we set $U = V$, $m = B_M$, $\sigma = \sigma_{M,x}$ and $R_m = R_{B_M}$. \square

Acknowledgements We thank Damien Regnault, Matthew Patitz, Trent Rogers and Andrew Winslow for important technical feedback and the following people for interesting discussions that helped to improve our thinking: Nicolas Schabanel, Dave Doty, Robert Schweller and Jacob Hendricks. We also thank for Nicolas Schabanel for hosting both authors at LIAFA (Paris 7, France) for the months of September 2014 and May 2015. A special thanks to our wives Elisa and Beverley.

A Additional simulation definitions

This appendix continues the definitions in Section 2.2 and were used in previous work on intrinsic universality [6–9, 11, 20].

Definition A.1. We say that \mathcal{S} and \mathcal{T} have *equivalent productions* (under R), and we write $\mathcal{S} \Leftrightarrow \mathcal{T}$ if the following conditions hold:

1. $\{R^*(\alpha') \mid \alpha' \in \mathcal{A}[\mathcal{S}]\} = \mathcal{A}[\mathcal{T}]$.
2. $\{R^*(\alpha') \mid \alpha' \in \mathcal{A}_{\square}[\mathcal{S}]\} = \mathcal{A}_{\square}[\mathcal{T}]$.

³⁴ m -block supertile representation functions were defined in Section 2.2. Since the proof of Theorem 1.2 reasons merely about the *shapes* of terminal assemblies, we do not even require that R_m sometimes maps different nonempty m -blocks to different tile types of T . In other words, having R_m map nonempty m -blocks to some tile type of T (here $\sigma \in T$) is sufficient to reason about the shape of assemblies under R_m .

3. For all $\alpha' \in \mathcal{A}[\mathcal{S}]$, α' maps cleanly to $R^*(\alpha')$.

Observation A.2. *If \mathcal{S} and \mathcal{T} that have equivalent productions (they satisfy Definition A.1) then they have equivalent shapes (they satisfy Definition 2.1).*

The following two definitions (for simulation *dynamics*) can be safely ignored by the reader and are included only for the sake of completeness.

Definition A.3. We say that \mathcal{T} *follows* \mathcal{S} (under R), and we write $\mathcal{T} \dashv_R \mathcal{S}$ if $\alpha' \rightarrow^{\mathcal{S}} \beta'$, for some $\alpha', \beta' \in \mathcal{A}[\mathcal{S}]$, implies that $R^*(\alpha') \rightarrow^{\mathcal{T}} R^*(\beta')$.

Definition A.4. We say that \mathcal{S} *models* \mathcal{T} (under R), and we write $\mathcal{S} \models_R \mathcal{T}$, if for every $\alpha \in \mathcal{A}[\mathcal{T}]$, there exists $\Pi \subset \mathcal{A}[\mathcal{S}]$ where $R^*(\alpha') = \alpha$ for all $\alpha' \in \Pi$, such that, for every $\beta \in \mathcal{A}[\mathcal{T}]$ where $\alpha \rightarrow^{\mathcal{T}} \beta$, (1) for every $\alpha' \in \Pi$ there exists $\beta' \in \mathcal{A}[\mathcal{S}]$ where $R^*(\beta') = \beta$ and $\alpha' \rightarrow^{\mathcal{S}} \beta'$, and (2) for every $\alpha'' \in \mathcal{A}[\mathcal{S}]$ where $\alpha'' \rightarrow^{\mathcal{S}} \beta'$, $\beta' \in \mathcal{A}[\mathcal{S}]$, $R^*(\alpha'') = \alpha$, and $R^*(\beta') = \beta$, there exists $\alpha' \in \Pi$ such that $\alpha' \rightarrow^{\mathcal{S}} \alpha''$.

The previous definition essentially specifies that whenever \mathcal{S} simulates an assembly $\alpha \in \mathcal{A}[\mathcal{T}]$, there must be at least one valid growth path in \mathcal{S} for each of the possible next steps \mathcal{T} could make from α .

Definition A.5. We say that \mathcal{S} *simulates* \mathcal{T} (under R) if $\mathcal{S} \Leftrightarrow_R \mathcal{T}$ (equivalent productions), $\mathcal{T} \dashv_R \mathcal{S}$ and $\mathcal{S} \models_R \mathcal{T}$ (equivalent dynamics).

A.1 Intrinsic universality

Now that we have a formal definition of what it means for one tile assembly system to simulate another, we can proceed to formally define the concept of intrinsic universality, i.e. when there is one general-purpose tile set that can be appropriately programmed to simulate any other tile system from a specified class of tile assembly systems. Let REPR denote the set of all supertile representation functions (i.e. m -block supertile representation functions for all $m \in \mathbb{Z}^+$). Define \mathfrak{C} to be a class of tile assembly systems, \mathfrak{C}_1 to be the class of all temperature 1 tile assembly systems, and let U be a tileset.

Definition A.6. We say U is *intrinsically universal* for \mathfrak{C} at temperature $\tau' \in \mathbb{Z}^+$ if there are functions $\mathcal{R} : \mathfrak{C} \rightarrow \text{REPR}$ and $S : \mathfrak{C} \rightarrow \mathcal{A}_{<\infty}^U$ 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, $R = \mathcal{R}(\mathcal{T})$ is a representation function that interprets assemblies of $\mathcal{U}_{\mathcal{T}}$ as assemblies of \mathcal{T} , and $\sigma_{\mathcal{T}} = S(\mathcal{T})$ is the seed assembly used to program tiles from U to represent the seed assembly of \mathcal{T} . In this paper, we disprove the existence of an intrinsically universal tileset for \mathfrak{C}_1 (the set of all temperature 1 tile assembly systems) at temperature $\tau' = 1$.

References

- [1] B. Behsaz, J. Mañuch, and L. Stacho. Turing universality of step-wise and stage assembly at temperature 1. In *DNA18: Proc. of International Meeting on DNA Computing and Molecular Programming*, volume 7433 of *LNCS*, pages 1–11. Springer, 2012.

- [2] M. Bousquet-Mélou. Families of prudent self-avoiding walks. *J. Comb. Theory, Ser. A*, 117(3):313–344, 2010.
- [3] S. Cannon, E. D. Demaine, M. L. Demaine, S. Eisenstat, M. J. Patitz, R. Schweller, S. M. Summers, and A. Winslow. Two hands are better than one (up to constant factors). In *STACS: Proceedings of the Thirtieth International Symposium on Theoretical Aspects of Computer Science*, pages 172–184. LIPIcs, 2013. Arxiv preprint: 1201.1650.
- [4] H. Chandran, N. Gopalkrishnan, and J. Reif. Tile complexity of approximate squares. *Algorithmica*, 66(1):1–17, 2013.
- [5] M. Cook, Y. Fu, and R. T. Schweller. Temperature 1 self-assembly: deterministic assembly in 3D and probabilistic assembly in 2D. In *SODA: Proceedings of the 22nd Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 570–589, 2011. Arxiv preprint: arXiv:0912.0027.
- [6] E. D. Demaine, M. L. Demaine, S. P. Fekete, M. J. Patitz, R. T. Schweller, A. Winslow, and D. Woods. One tile to rule them all: Simulating any tile assembly system with a single universal tile. In *ICALP: Proceedings of the 41st International Colloquium on Automata, Languages, and Programming*, volume 8572 of *LNCS*, pages 368–379. Springer, 2014. Arxiv preprint: arXiv:1212.4756.
- [7] E. D. Demaine, M. J. Patitz, T. A. Rogers, R. T. Schweller, S. M. Summers, and D. Woods. The two-handed tile assembly model is not intrinsically universal. In *ICALP: Proceedings of the 40th International Colloquium on Automata, Languages, and Programming*, volume 7965 of *LNCS*, pages 400–412. Springer, July 2013. Arxiv preprint: arXiv:1306.6710.
- [8] D. Doty, J. H. Lutz, M. J. Patitz, R. T. Schweller, S. M. Summers, and D. Woods. The tile assembly model is intrinsically universal. In *FOCS: Proceedings of the 53rd Annual IEEE Symposium on Foundations of Computer Science*, pages 439–446. IEEE, Oct. 2012. Arxiv preprint: arXiv:1111.3097.
- [9] D. Doty, J. H. Lutz, M. J. Patitz, S. M. Summers, and D. Woods. Intrinsic universality in self-assembly. In *STACS: Proceedings of the 27th International Symposium on Theoretical Aspects of Computer Science*, pages 275–286, 2009. Arxiv preprint: arXiv:1001.0208.
- [10] D. Doty, M. J. Patitz, and S. M. Summers. Limitations of self-assembly at temperature 1. *Theoretical Computer Science*, 412(1–2):145–158, 2011. Arxiv preprint: arXiv:0906.3251.
- [11] S. P. Fekete, J. Hendricks, M. J. Patitz, T. A. Rogers, and R. T. Schweller. Universal computation with arbitrary polyomino tiles in non-cooperative self-assembly. In *SODA: ACM-SIAM Symposium on Discrete Algorithms*, pages 148–167. SIAM, 2015.
- [12] P. J. Flory. *Principles of Polymer Chemistry*. Cornell University Press, 1953.
- [13] B. Fu, M. J. Patitz, R. T. Schweller, and R. Sheline. Self-assembly with geometric tiles. In *ICALP: Proceedings of the 39th International Colloquium on Automata, Languages, and Programming*, volume 7391 of *LNCS*, pages 714–725. Springer, 2012.
- [14] O. Gilbert, J. Hendricks, M. J. Patitz, and T. A. Rogers. Computing in continuous space with self-assembling polygonal tiles. In *SODA: ACM-SIAM Symposium on Discrete Algorithms*, pages 937–956. SIAM, 2016. Arxiv preprint: arXiv:1503.00327.

- [15] J. Hendricks, M. J. Patitz, T. A. Rogers, and S. M. Summers. The power of duples (in self-assembly): It's not so hip to be square. In *COCOON: Proceedings of 20th International Computing and Combinatorics Conference*, pages 215–226, 2014. Arxiv preprint: [arXiv:1402.4515](https://arxiv.org/abs/1402.4515).
- [16] N. Jonoska and D. Karpenko. Active tile self-assembly, part 1: Universality at temperature 1. *Int. J. Found. Comput. Sci.*, 25(2):141–164, 2014.
- [17] D. E. Knuth. Mathematics and computer science: coping with finiteness. *Mathematics: people, problems, results*, 2, 1984.
- [18] J. Mañuch, L. Stacho, and C. Stoll. Two lower bounds for self-assemblies at temperature 1. *Journal of Computational Biology*, 17(6):841–852, 2010.
- [19] P.-É. Meunier. Non-cooperative algorithms in self-assembly. In *UCNC: Unconventional Computation and Natural Computation*, volume 9252 of *LNCS*, pages 263–276. Springer, 2015.
- [20] P.-É. Meunier, M. J. Patitz, S. M. Summers, G. Theyssier, A. Winslow, and D. Woods. Intrinsic universality in tile self-assembly requires cooperation. In *SODA: Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*, pages 752–771, 2014. Arxiv preprint: [arXiv:1304.1679](https://arxiv.org/abs/1304.1679).
- [21] J. E. Padilla, M. J. Patitz, R. T. Schweller, N. C. Seeman, S. M. Summers, and X. Zhong. Asynchronous signal passing for tile self-assembly: Fuel efficient computation and efficient assembly of shapes. *International Journal of Foundations of Computer Science*, 25(4):459–488, 2014. Arxiv preprint: [arxiv:1202.5012](https://arxiv.org/abs/1202.5012).
- [22] M. J. Patitz, R. T. Schweller, and S. M. Summers. Exact shapes and Turing universality at temperature 1 with a single negative glue. In *DNA 17: Proceedings of the Seventeenth International Conference on DNA Computing and Molecular Programming*, LNCS, pages 175–189. Springer, Sept. 2011. Arxiv preprint: [arXiv:1105.1215](https://arxiv.org/abs/1105.1215).
- [23] P. W. K. Rothmund. *Theory and Experiments in Algorithmic Self-Assembly*. PhD thesis, University of Southern California, December 2001.
- [24] P. W. K. Rothmund and E. Winfree. The program-size complexity of self-assembled squares (extended abstract). In *STOC: Proceedings of the thirty-second annual ACM Symposium on Theory of Computing*, pages 459–468, Portland, Oregon, United States, 2000. ACM.
- [25] D. Soloveichik and E. Winfree. Complexity of self-assembled shapes. *SIAM Journal on Computing*, 36(6):1544–1569, 2007.
- [26] E. Winfree. *Algorithmic Self-Assembly of DNA*. PhD thesis, California Institute of Technology, June 1998.
- [27] E. Winfree. Simulations of computing by self-assembly. Technical Report Caltech CS TR:1998.22, California Institute of Technology, 1998.
- [28] D. Woods. Intrinsic universality and the computational power of self-assembly. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 373(2046), 2015. [dx.doi.org/10.1098/rsta.2014.0214](https://doi.org/10.1098/rsta.2014.0214).