Negative Interactions in Irreversible Self-assembly*

David Doty[1], Lila Kari[1], and Benoît Masson[2]

[1] U. of West. Ontario, Dept. of Computer Science, London, Canada
{ddoty,lila}@csd.uwo.ca
[2] IRISA (INRIA), Campus de Beaulieu, Rennes, France
benoit.masson@irisa.fr

Abstract. This paper explores the use of negative (i.e., repulsive) interactions in the abstract Tile Assembly Model defined by Winfree. Winfree in his Ph.D. thesis postulated negative interactions to be physically plausible, and Reif, Sahu, and Yin studied them in the context of *reversible* attachment operations. We investigate the power of negative interactions with *irreversible* attachments, and we achieve two main results. Our first result is an impossibility theorem: after t steps of assembly, $\Omega(t)$ tiles will be forever bound to an assembly, unable to detach. Thus negative glue strengths do not afford unlimited power to reuse tiles. Our second result is a positive one: we construct a set of tiles that can simulate an s-space-bounded, t-time-bounded Turing machine, while ensuring that no intermediate assembly grows larger than $O(s)$, rather than $O(s \cdot t)$ as required by the standard Turing machine simulation with tiles.

1 Introduction

Tile-based self-assembly is a model of "algorithmic crystal growth" in which square "tiles" represent molecules that bind to each other via highly-specific bonds on their four sides, driven by random mixing in solution but constrained by the local binding rules of the tile bonds. Erik Winfree [10], based on experimental work of Seeman [7], modified Wang's mathematical model of tiling [9] to add a physically plausible mechanism for growth through time. Winfree defined a model of tile-based self-assembly known as the abstract Tile Assembly Model (aTAM). The fundamental components of this model are un-rotatable, but translatable square "tile types" whose sides are labeled with "glues" representing binding sites. Two tiles that are placed next to each other are attracted with strength determined by the glues where they abut, and in the aTAM, a tile *binds* to an assembly if it is attracted on all sides with total strength at least

* This research was supported in part by Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant R2824A01 and the Canada Research Chair Award in Biocomputing to Lila Kari, and by the NSF Computing Innovation Fellowship grant to David Doty.

Y. Sakakibara and Y. Mi (Eds.): DNA 16, LNCS 6518, pp. 37–48, 2011.

a certain threshold value τ.[1] Assembly begins from a "seed" tile and progresses until no more tiles may attach.

We study a variant of this model in which glue strengths are allowed to be negative as well as positive. This leads to the situation in which a stable assembly may become unstable through the addition of a tile that, while binding strongly enough to the assembly to remain attached itself, exerts a repulsive force on a neighboring tile, which is sufficiently strong to detach some portion of the assembly. This is formally modeled by allowing an assembly to break into two parts any time that the two parts have total connection strength less than τ (i.e., if there is a cut of the interaction graph of strength less than τ). Negative glue strengths were discussed as a plausible mechanism in Winfree's thesis [10], and explored theoretically in a more general model of *graph-based self-assembly* by Reif, Sahu and Yin [4]. We compare the results of [4] to the present paper in more detail later in this section.

This paper has two main contributions, an impossibility result and a positive result. The impossibility result is that under the irreversible model, negative glue strengths are not sufficient to achieve perfect reuse of tiles as in [4]. It is tempting to believe that with negative glue strengths, the monotonic growth of the aTAM could be overcome to such a degree that a bounded set of tiles could be reused for arbitrarily long computations,[2] hence implementing the observation that "you can reuse space but you can't reuse time". Alas, you cannot reuse space (tiles) too much with irreversible reactions. We show that under the irreversible model of tile assembly, even with negative glue strengths, many tiles will be forever bound to an assembly, unable to detach. In fact, this number is linear in the number of assembly operations, so that after t operations, $\Omega(t)$ tiles will be permanently bound to some assembly.

The positive result is a construction attempting to make do with this limitation. For concreteness, our construction shows how to simulate a single-tape Turing machine. But the idea applies to the iterated computation of any function f that can be "computed with constant height" by a tile assembly system (a formal definition is given in Section 4). The function f_M mapping the configuration of a Turing machine M to its next configuration is an example of one such function. Other examples include the incrementing or decrementing of a counter, or the selection of a uniformly distributed random number from a finite set $\{1, 2, \ldots, n\}$ using flips of a fair coin via von Neumann's *rejection method*, as shown in [2].

Our construction achieves the following property: if the Turing machine M being simulated on input x (with $n = |x|$) has space bound $s(n)$ and time bound $t(n)$, then $O(t(n) \cdot s(n))$ tiles (meaning total count of tiles, which is greater than

[1] The threshold τ models the temperature at which insufficiently strong chemical bonds will break, such as those formed by Watson-Crick complementarity in DNA-based implementations of tiles.

[2] Subject, of course, to computational complexity constraints such as $\mathsf{DTIME}(t(n)) \subseteq \mathsf{DSPACE}(2^{t(n)})$, based on the observation that configurations cannot repeat during the course of a halting computation.

the number of unique tile types), mixed in solution, will simulate the computation of M on input x, and no intermediate assembly will grow to size larger than $O(s(n))$. The impossibility result can be interpreted to imply that external energy must be supplied to break bonds between tiles if we wish to reuse them for computation. If we wish to limit the volume of a solution, and therefore the number of molecules it can contain (by the finite density constraint, see [8]) to $O(s(n))$, then we cannot allow intermediate assemblies to grow larger than this value. Of course, by the impossibility result, many more than $s(n)$ different such assemblies will form if $t(n) \gg s(n)$ (for instance, when simulating a linear-space, cubic-time computation). With a mechanism to "vacuum" away junk assemblies and supply the external energy needed to break them up (a mechanism not modeled in the aTAM), these tiles could be reused, bringing down the required number of tiles from $O(t(n) \cdot s(n))$ to $O(s(n))$.

The main difference between [4] and the present paper is that [4] employs *reversible* reactions, and the present paper employs *irreversible* reactions.[3] Within the aTAM, the main difference between our model and [4] amounts to a difference in the definition of a legal attachment operation. In [4], the authors define a tile attachment to be legal if the tile attaches with strength $\tau - 1$ (in fact, they define it a bit differently but restricting attention to our construction and that of [4], this definition is equivalent). This is a phenomenon not modeled by the aTAM, but it is physically plausible to suppose that it occurs, though with less frequency than strength τ attachments (see the kinetic TAM of [10]). Therefore the tile may detach after attaching since it is held with insufficient strength. But, if it first causes another tile or group of tiles to be bound with total strength less than τ, then those tiles may also fall off, possibly resulting in stabilization of the original attachment. In the present paper, we define attachments to be legal only if they have strength at least τ, whereas detachments may only happen between assemblies attached with strength at most $\tau - 1$. This difference implies that our impossibility result does not apply to [4], which can be considered an advantage of reversible interactions. But this advantage does not come without disadvantages: due to the second law of thermodynamics, their construction must necessarily be implemented as an unbiased random walk with equal rates of forward and reverse reaction, lest the entropy of the system increase with time if one direction is more favorable. Therefore their construction takes expected time n^2 to go forward n steps.

We should also note that although [4] uses a more general model of graph-based self-assembly, this does not imply that their construction of an assembly system simulating a space-bounded Turing machine simulation is a stronger result than our construction. The more general model affords more power to aid in a construction, such as allowing non-planar interactions, in addition to the extra power of reversible interactions. Therefore, we emphasize that our positive construction is not merely a specialization of the construction of [4] to grid graphs. The construction of [4] is inherently non-planar and reversible. A major

[3] [4] also uses a more general graph-based model of self-assembly, but this difference is less crucial than the reversibility issue.

source of the effort in designing our construction was getting it to work in the plane and use irreversible attachments. Many similar (and simpler) constructions that superficially appear to do the same thing as our construction do not actually work in our model, as they introduce not only a desired cut of strength less than τ, but also some undesired cuts of strength less than τ, which, if detached, will ruin the construction.

For color figures, see `http://www.csd.uwo.ca/~ddoty/papers/niisa.pdf`.

2 Abstract Tile Assembly Model

This section gives a brief definition of the abstract Tile Assembly Model (aTAM, [10]) with negative glue strengths. This not a tutorial on the aTAM; for readers unfamiliar with the model, please see [5] for an excellent introduction.

$\mathbb{Z}$ and $\mathbb{Z}^+$ denote the set of integers and positive integers, respectively. Let G be a finite alphabet of *glues*. A *tile type* is a tuple $t \in G^4$, i.e., a unit square with a glue on each side. Associated with the tile types is a *glue strength function* $str : G \times G \to \mathbb{Z}$ that indicates, given two glues g_1 and g_2, the strength $str(g_1, g_2)$ with which they interact. We assume a finite set T of tile types, but an infinite number of copies of each tile type, each copy referred to as a *tile*. Let $G(T)$ denote the set of all glues of tile types in T. An *assembly* (a.k.a., *supertile*) is a positioning of tiles on the integer lattice $\mathbb{Z}^2$ (i.e., a partial function $\alpha : \mathbb{Z}^2 \dashrightarrow T$, where $\dashrightarrow$ denotes that the function is partial). Each assembly induces a *binding graph*, a grid graph whose vertices are tiles, with an edge between two tiles if they are adjacent (i.e., are Euclidean distance 1 apart).[4] The assembly is *τ-stable*, or simply *stable* if τ is understood from context, if every cut of its binding graph has weight (strength) at least τ, where the weight of an edge is the strength of the glue it represents. That is, the assembly is stable if at least energy τ is required to separate the assembly into two parts. In this paper, where not stated otherwise, we assume that $\tau = 2$.

A *tile assembly system* (TAS) is a 4-tuple $\mathcal{T} = (T, str, \sigma, \tau)$, where T is a finite set of tile types, $str : G(T) \times G(T) \to \mathbb{Z}$ is the *glue strength function*, $\sigma : \mathbb{Z}^2 \dashrightarrow T$ is the finite and τ-stable *seed assembly*, and $\tau \in \mathbb{Z}^+$ is the *temperature*. Given a TAS $\mathcal{T} = (T, str, \sigma, \tau)$, an assembly α is *producible* if either (base case) $\alpha = \sigma$, or (recursive case 1) α results from the τ-stable attachment of a single tile to a producible assembly ("τ-stable attachment" meaning that the cut separating the tile from the rest of the assembly has strength $\geq \tau$), or (recursive case 2) α consists of one side of a cut of strength $< \tau$ of a producible assembly. Note in particular that a producible assembly need not be stable, but may be stabilized by attachments before it can break apart. An assembly α is *terminal* if α is τ-stable and no tile can be τ-stably attached to α. Let $B \subseteq T$ be a set of "black"

[4] Previous papers model the binding graph as having edges only between tiles that interact with positive strength. In the present paper, the presence of negative glue strengths means that we must consider every possible interaction between adjacent tiles, whether positive, negative, or 0.

tile types. $\mathcal{T}$ is *B-directed* (a.k.a., *B-deterministic*, *B-confluent*) if it has exactly one terminal, producible assembly containing one or more tiles from B.[5]

To define reversible assembly at temperature $\tau = 2$ (as in [4]), it suffices to define attachment events with strength threshold $\tau - 1 = 1$, rather than strength threshold $\tau = 2$. This behavior is illustrated on Fig. 1(a), and can be compared with our new notion, whose evolution is shown on Fig. 1(b).

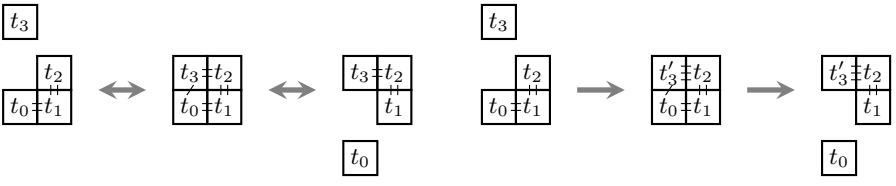

(a) Reversible model, as defined in [4]. (b) Irreversible model.

Fig. 1. Two different implementations of negative interactions at temperature 2. The slanted bonds represent a strength of -1. In the reversible model, the tile t_3 can attach with a total strength of 1 (one bond of strength 2 and one of strength -1) and hence is unstable, while with our definition, t'_3 is attached with a total strength of 2 and forces t_0 to detach.

3 Limitation of Tile Reuse with Irreversible Reactions

If α is an assembly, define $\Phi(\alpha)$, the *(negative) free energy* of α, to be the sum of all glue strengths between adjacent tiles in the assembly.[6] In particular, an assembly consisting of a single tile has free energy 0. If S is a multiset of assemblies (such as that produced by a TAS with negative glue strengths, considering even the "junk" assemblies that are discarded after a cut), define the (negative) free energy of S to be the sum of the free energies of each assembly in S, denoted $\Phi(S)$. Note that even postulating an infinite count of tiles, after a finite number of operations, only finitely many assemblies in S consist of more than one tile, and each of these is a finite assembly. Therefore $\Phi(S) < \infty$ for any multiset S of assemblies producible by a TAS, even in the case that $|S| = \infty$ (such as the initial multiset consisting of a countably infinite number of copies of each individual tile type).

[5] We define this notion of B-directedness but do not henceforth discuss it explicitly, since our construction simulates a general "computation", and B would depend on the goals of the computation being simulated. In our example construction in Section 4 of simulating steps of a Turing machine, B could, for instance, consist of the tile types that represent a halting state, so that only a terminal assembly representing the configuration of a halted Turing machine would be considered the result.

[6] The standard definition of free energy is the negative of this quantity, but as in [5] we use its negation so that the quantity will be positive for stable assemblies. Intuitively, it is the energy *required* to separate α into individual tiles, whereas the standard definition is the energy *released* by such a separation.

When we discuss the "number of steps" for the assembly process of a TAS, we mean the total number of attachment and detachment operations that have been applied so far. We do not claim that this is a proper model of "running time", but it is convenient to think of attachment and detachment events as discrete and equally-spaced steps, even though they may happen in parallel or with interval times governed by a continuous distribution.

Theorem 1. *Let $\mathcal{T}$ be a TAS, and let S be a multiset of assemblies producible by $\mathcal{T}$ after $t \in \mathbb{N}$ steps. Then $\Phi(S) \geq t/2$.*

A proof of Theorem 1 will appear in a full version of this extended abstract.

Since the glue strengths are bounded above by some constant s, an immediate consequence of Theorem 1 is that after t steps, at least $t/(2s)$ sides of tiles are bound. With the finite tile count assumption, once t is sufficiently large that $t/(2s)$ exceeds the total number of sides available (i.e., 4 times the total number of tiles in solution), no more sides are available for binding, and self-assembly grinds to a halt. This is the sense in which a finite number of tiles cannot be reused indefinitely.

There is a natural thermodynamic interpretation of Theorem 1: work done by tiles on tiles, in an irreversible manner, increases the entropy of the system by the second law of thermodynamics, thus decreasing the potential energy available to do more work. Therefore, any potential energy stored in the unattached glues is eventually permanently used up if external energy is not supplied to break these bonds. In our main construction, many junk assemblies are created that are no longer useful once the tiles in them have been used once. Theorem 1 tells us that no amount of cleverness will allow us to break up those assemblies and reuse the tiles solely through design of tile types with negative glues; some external force must be supplied to break them apart using a mechanism not modeled in the aTAM.

Of course, Theorem 1, interpreted in light of the molecular interactions that are being modeled by the aTAM, should not be surprising to any physicist. But we believe it is important to formally establish the truth of such a statement within the model. One develops more confidence in a model of reality when it tells us something already known about reality (e.g., the Positive Mass Theorem [6]).

Theorem 1 does not apply to the negative glue strength construction of Reif, Sahu, and Yin [4], because their model allows *reversible* reactions. Attempting to apply our proof to their model would result in the first inequality $\Phi(S_{i+1}) - \Phi(S_i) \geq \tau$ being replaced by $\Phi(S_{i+1}) - \Phi(S_i) \geq \tau - 1$, which would result in a final lower bound of 0, instead of $t/2$, for $\Phi(S)$. Intuitively, the reversibility of reactions implies that attachment and detachment have symmetric effects on the free energy. But this also implies that their system requires driving the system forward through an unbiased random walk, taking n^2 steps on average to proceed by n net forward steps. Any attempt to speed up the reaction to make the forward rate of reaction faster than the reverse rate of reaction would introduce the imbalance in their respective effects on free energy that allows our proof to work. Therefore this tradeoff in speed versus reusability of tiles is fundamental.

4 Turing Machine Simulation

Throughout this section, fix some finite alphabet Σ. We first describe the class of functions that we will compute, which are intuitively those computable by a constant number of rows of assembly (although the number of columns is unbounded) in the standard aTAM. See [2], for example, for a formal definition of the standard aTAM model. Briefly, it is the same as the model defined in Section 2, but glue strengths are non-negative and are only positive between equal glues.

Definition 1. *Let T be a set of tile types, and let $e : T \to \Sigma$. We say that a row of tiles (a connected subassembly of some assembly of height 1) $t_1, t_2, \ldots, t_k$ e-*encodes *a string $x \in \Sigma^k$ if $e(t_1) = x[1], e(t_2) = x[2], \ldots, e(t_k) = x[k]$, where $x[i] \in \Sigma$ is the i^{th} symbol in x. A function $f : \Sigma^* \to \Sigma^*$ is* constant-row computable *if there exists a tile set T, a function $e : T \to \Sigma$, and a constant c such that, for each $x \in \Sigma^*$, there is a height-1 stable assembly $\sigma_x : \mathbb{Z}^2 \dashrightarrow T$ e-encoding x such that the tile assembly system $\mathcal{T} = (T, str, \sigma_x, 2)$ (with $str(g_1, g_2) > 0 \iff g_1 = g_2$) has the unique terminal assembly α, the height of α is c, the bottom row of α is σ_x, the top row of α e-encodes $f(x)$, and the leftmost column of any row of α is no further left than the bottom row.*

The widths of the rows representing the input and output may be different (i.e., possibly $|x| \neq |f(x)|$). In this case, we require only that the leftmost and rightmost tiles of each row have their glues specially marked to distinguish them from the tiles interior to the row.

Our construction shows how to design a tile set that will compute iterations of any constant-row computable function f, ensuring that no intermediate assembly grows larger than the size of the input or output processed by any *individual* invocation of f. Examples of such functions include the function f that, given a configuration of a single-tape Turing machine outputs the next configuration of this Turing machine, or that increments a counter represented in binary.

Figure 2 shows a high-level overview of the entire construction, in terms of a general constant-row computable function f. For concreteness, think of f as the function that, given a configuration of a t-time-bounded, s-space-bounded, single-tape Turing machine, outputs the next configuration of this Turing machine (extending the tape on the right side only). The construction proceeds as follows, each label corresponds to a picture in Figure 2.

(a) First, the scaffold tiles (green) connect to the x data assembly (white). The scaffold tiles initiate the computation of f (blue).
(b) The scaffold "detects" when the computation is finished, in the sense that the green row above $f(x)$ tiles cannot complete until all of $f(x)$ is present. Then the scaffold tiles grow back to the first scaffold tile to initiate the removal of $f(x)$ from the tiles surrounding $f(x)$.

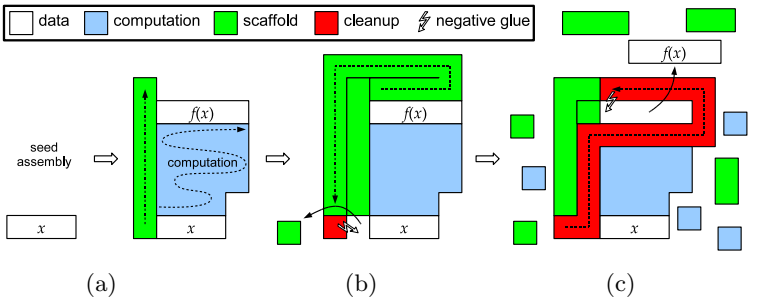

Fig. 2. High-level overview of assembly for computation of a constant-row computable function f

(c) The removal tiles (red) each use a negative glue strength against the tile "in front of" (on the path show by the arrows) it, and once this tile is removed, a new removal tile grows in its place to continue the removal. The path and bond placements and strengths are carefully chosen to ensure that no portion of $f(x)$ is removed, until the last step when $f(x)$ detaches whole from the rest of the tiles.

Note that since f is constant-row computable, the height of the scaffold and removal parts are bounded by a constant and therefore may be hard-coded into the tile set, whereas special glues mark the horizontal endpoints so that the length of x and $f(x)$ are not constrained.

The simulation of the Turing machine for t steps will then consist of executing this assembly process for t iterations, using the output assembly $f(x)$ as the input assembly x for the next execution. After each iteration, the width of the remaining "junk" assembly is a constant plus $O(1) + \max\{|x|, |f(x)|\}$, and the height is constant since f is constant-row computable, so the size of the intermediate assemblies is $O(\max\{|x|, |f(x)|\})$.

Figures 3, 4, and 5 give some more details for the three main steps of Figure 2, respectively (a), (b), and (c), using the specific example of f mapping a configuration of a single-tape Turing machine to its next configuration.

Figure 3 shows an example of tiles implementing step (a) of Figure 2, i.e., the computation of f. The example shows one transition of a single-tape Turing machine, with tape contents 01_0 (_ standing for blank), in state q, with tape head on the rightmost cell, transitioning to state p, moving the tape head right, changing the cell's symbol from 0 to 1, and encountering a blank on the new rightmost cell. In this case, a new rightmost cell is needed, illustrating how our construction handles dynamically changing space requirements, but if the tape head were further left in the row, it would simply fill in copy tiles to the right, just as to the left as shown above, and the row would stay the same width. At the start and end of a computation, the configuration is copied so that any strength > 1 bonds used in the computation are on the interior of the computation tiles, ensuring that only strength-1 bonds must later be broken to separate the data

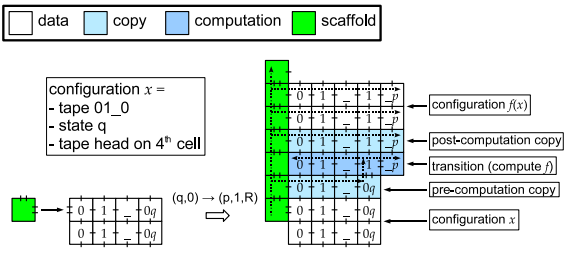

Fig. 3. Example of tiles implementing the computation step. Arrows within tiles show order of growth. In this case f is constant-row computable with constant $c = 1$. The first and last copy rows, shown in lighter shade than the center computation tiles, are always present no matter the function f, and their placement is initiated by the scaffold tiles. However, there is no interaction between the center computation and scaffold tiles. Note that the data tiles are two rows with strength 1 glues; this is to make them stable at temperature 2 but not producible (without additional scaffolding) as they would be if they were a single row connected with strength 2 glues.

tiles. Each data assembly on either end of the computation tiles is represented by a two-row assembly with only single-strength bonds on its interior, which ensures that when detached, the data assembly will be stable, but that it cannot form on its own without help from the scaffold tiles (which would happen if it were only a single row connected with strength-2 bonds). Each vertical position is hard-coded into the tile set; i.e., the scaffold tile set "knows" the required height to compute f. However, the absolute horizontal positions are not encoded into the tiles, only the leftmost and rightmost tiles of the configuration are specially marked, and all interior tiles representing the same data are identical.

Figure 4 shows the tiles implementing step (b) of Figure 2, positioning the tiles for cleanup. The top two rows must use cooperation to tell where the end of the

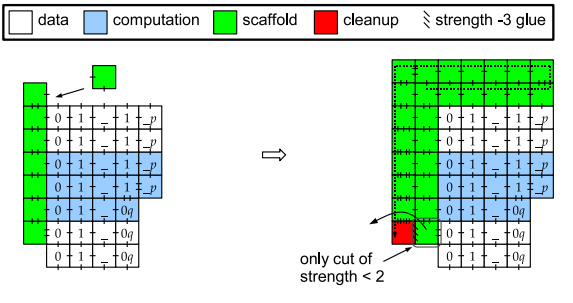

Fig. 4. Tiles that position the cleanup tiles. Here the "copy" tiles from Figure 3 are depicted in the same shade as the computation tiles; now that $f(x)$ has been computed our goal is to remove all of them from the subassembly representing $f(x)$. The order of growth of the scaffold tiles ensures that cleanup does not begin until all of $f(x)$ is present.

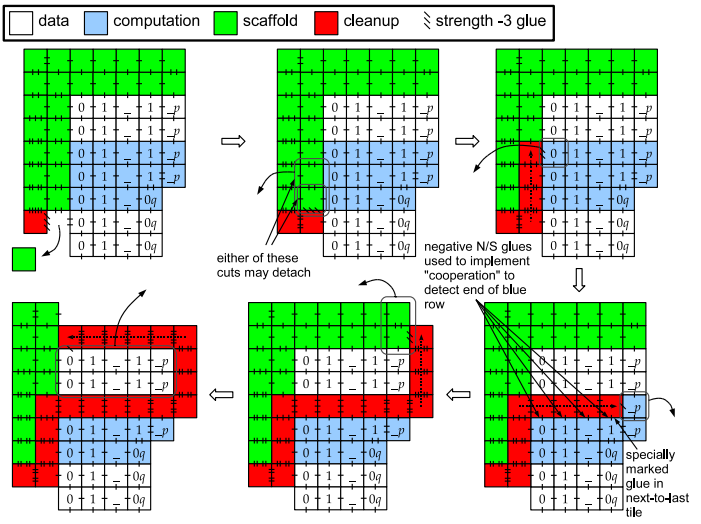

Fig. 5. Tiles that "clean up" the connections between the output data and the scaffold and computation tiles to separate them and allow the data tiles to be computed on again

row underneath is, since the width of the output row is unknown. The strengths of bonds on the leftmost downward-growing column must be sufficiently large to ensure that only the proper cut is made when the first negative-strength glue is applied.

Figure 5 shows the tiles implementing step (c) of Figure 2, "cleaning up" by removing the output $f(x)$ from the scaffold, computation, and x tiles. Though not shown, negative strength interactions are necessary between the second-to-top row of computation tiles and some of the right-growing cleanup tiles, to ensure that the right end of the row is properly detected. That is, there are two types of cleanup tiles growing right, one to detach the interior tiles, and one to detach the final rightmost computation tile. Since the east-west bonds between cleanup tiles are greater than 1, the negative north-south glue strengths between interior cleanup tiles and the second-to-rightmost blue tile – and between the rightmost cleanup tile and the interior computation tiles – must be strength -2 to ensure that the second-to-rightmost blue tile cannot stably attach except where intended.

5 Conclusion

We have shown two main results in the aTAM with negative glue strengths, under the standard assumption of *irreversible* attachment, meaning attachments that only occur with strength at least the temperature τ. The first result is that the amount of tile reuse afforded by the ability to detach tiles with negative glue strengths is fundamentally limited. After t steps of assembly, $\Omega(t)$ tiles are

permanently bound, unable to detach via negative glue strengths, and can only be detached by supplying external energy. The second result is a positive result that attempts to make do with this limitation: an $s(n)$-space-bounded Turing machine may be simulated for arbitrarily many steps, while ensuring that no intermediate assembly grows larger than $O(s(n))$.

Space-bounded "computation" as an end goal is not the only application of negative glue strengths, of course. Doty, Lutz, Patitz, Summers, and Woods [2] study the problem of generating uniform random distributions on the finite sets using the independent flips of a fair coin afforded by the random selection of competing tile types in the aTAM (a non-trivial problem when the cardinality of the set is not a power of the number of competing tile types), and find a tradeoff between the closeness to uniformity of the distribution obtained and the space required for sampling. They exhibit a construction imposing a perfectly uniformly distribution on the set $\{0, 1, \ldots, n-1\}$ that assembles a structure of width $\lfloor \log n \rfloor + 1$ and *expected* height at most 2, essentially implementing von Neumann's *rejection method* of flipping $\lfloor \log n \rfloor + 1$ fair coins repeatedly and stopping the first time that they encode a number smaller than n. It is very unlikely (probability at most 2^{-20}) to take more than 20 attempts. But using this method in a construction such as that of [3],[7] in which many (perhaps more than 2^{20}) copies of this experiment repeat throughout assembly, could increase the likelihood of growing too high. Even a single occurrence of a too-high subassembly will destroy the entire construction. Though we omit the details in the extended abstract, it is straightforward to augment the construction of [3] (which uses a variant of the random number selector of [2]) with negative glue strengths to implement perfectly uniform selection of random numbers, thus improving the fidelity of the simulation of [3], while providing an absolute guarantee on the space bound.

There are other uses of negative glues in the aTAM. For instance, we are able to improve the best known tile complexity (number of tile types) required to uniquely assemble a "thin" rectangle, i.e., an $n \times k$ rectangle with $k < \log n / (\log \log n - \log \log \log n)$. In the standard aTAM the tile complexity of this shape is known to be $\Omega(\frac{n^{1/k}}{k})$ [1]. With the model of negative glue strengths we are able to improve this to $O\left(\sqrt{\log n}\right)$, by first building a thick rectangle and using negative glues to "cut out" a thinner rectangle of the same length.

Other questions related to this work include the experimental aspects of such a model, for example, how repulsive forces can be realized on DNA tiles, and how to "clean" and "recycle" the junk introduced during the assembly.

[7] The main construction of [3] shows how a "universal" tile set can be constructed that can be "programmed" through appropriate selection of a seed assembly to simulate the growth of any tile assembly system in a wide class of systems termed "locally consistent" (see [3] for details). In this discussion, we are concerned only with the fact that the construction of [3] 1) requires random numbers to be generated in a bounded space at many points throughout assembly, and 2) would be improved if the distribution of these numbers were perfectly uniform instead of "close to uniform" as in [3].

Acknowledgments. The authors are grateful to Shinnosuke Seki for insightful discussions and anonymous referees for their suggestions.

References

[1] Aggarwal, G., Cheng, Q., Goldwasser, M.H., Kao, M.-Y., Moisset de Espanés, P., Schweller, R.T.: Complexities for generalized models of self-assembly. SIAM Journal on Computing 34, 1493–1515 (2005); Preliminary version appeared in SODA 2004

[2] Doty, D., Lutz, J.H., Patitz, M.J., Summers, S.M., Woods, D.: Random number selection in self-assembly. In: Calude, C.S., Costa, J.F., Dershowitz, N., Freire, E., Rozenberg, G. (eds.) UC 2009. LNCS, vol. 5715, pp. 143–157. Springer, Heidelberg (2009)

[3] Doty, D., Lutz, J.H., Patitz, M.J., Summers, S.M., Woods, D.: Intrinsic universality in self-assembly. In: STACS 2010: Proceedings of the 27th International Symposium on Theoretical Aspects of Computer Science. Leibniz International Proceedings in Informatics (LIPIcs), vol. 5, pp. 275–286 (2010)

[4] Reif, J.H., Sahu, S., Yin, P.: Complexity of graph self-assembly in accretive systems and self-destructible systems. In: Carbone, A., Pierce, N.A. (eds.) DNA 11. LNCS, vol. 3892, pp. 257–274. Springer, Heidelberg (2006)

[5] Rothemund, P.W.K., Winfree, E.: The program-size complexity of self-assembled squares (extended abstract). In: STOC 2000: Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, pp. 459–468 (2000)

[6] Schoen, R., Yau, S.-T.: On the positive mass conjecture in general relativity. Communications in Mathematical Physics 65(45), 45–76 (1979)

[7] Seeman, N.C.: Nucleic-acid junctions and lattices. Journal of Theoretical Biology 99, 237–247 (1982)

[8] Soloveichik, D., Cook, M., Winfree, E., Bruck, J.: Computation with finite stochastic chemical reaction networks. Natural Computing 7(4), 615–633 (2008)

[9] Wang, H.: Proving theorems by pattern recognition – II. The Bell System Technical Journal XL(1), 1–41 (1961)

[10] Winfree, E.: Algorithmic Self-Assembly of DNA. PhD thesis, California Institute of Technology (June 1998)