



Optimizing Tile Set Size While Preserving Proofreading with a DNA Self-assembly Compiler

Constantine G. Evans^{1,2(✉)} and Erik Winfree²

¹ Evans Foundation, Pasadena, CA, USA
cgevans@evansfmm.org

² California Institute of Technology, Pasadena, CA, USA

Abstract. Algorithmic DNA tile systems have the potential to allow the construction by self-assembly of large structures with complex nanometer-scale details out of relatively few monomer types, but are constrained by errors in growth and the limited sequence space of orthogonal DNA sticky ends that program tile interactions. We present a tile set optimization technique that, through analysis of algorithmic growth equivalence, potentially sensitive error pathways, and potential lattice defects, can significantly reduce the size of tile systems while preserving proofreading behavior that is essential for obtaining low error rates. Applied to systems implementing multiple algorithms that are far beyond the size of currently feasible implementations, the optimization technique results in systems that are comparable in size to already-implemented experimental systems.

1 Introduction

Self-assembling DNA tile systems provide a mechanism for implementing complex self-assembly behaviors at a molecular scale [22,30]. Both simple periodic structures with a range of attachment and lattice configurations, and “uniquely-addressed” structures of single copies of thousands of different monomers, have been demonstrated experimentally [21,26,32,33]. In between, algorithmic tile systems can employ a small number of monomer types that perform potentially arbitrary computation during growth to construct large structures with complex, small-scale details [7,34]. Additionally, using the choice of initial seed or presence of a particular monomer type as an input to a computation, the same algorithmic tile system can grow substantially different assemblies.

The number of monomer types in a tile system affects its cost and is a frequently-used measure of complexity [23,31]. The number of glues, however, is a potentially more limiting constraint. Implemented as short, single-stranded “sticky ends” of DNA, glues are limited by sequence space and spurious binding between subsequences of non-complementary sticky ends [10], particularly for the 5 to 6 nucleotide sticky ends in the widely-used double-crossover (DX) tile motifs [12]. Uniquely-addressed systems can assemble largely-correct structures

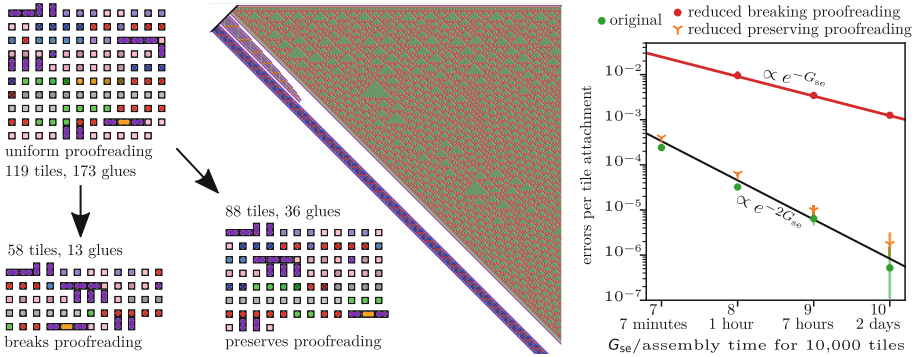


Fig. 1. An example of an algorithmic tile system designed for DNA tiles, simultaneously implementing 3 separate algorithms, reduced in size with and without preservation of proofreading behavior, and corresponding simulation results.

even if some tiles attach incorrectly: they may thus be less affected by non-complementary interactions, and sticky ends of inhomogeneous strengths may actually assist growth [15]. For algorithmic systems, however, a single incorrect tile attachment could completely change further growth, and thus spurious interactions can strongly limit the number of glues, limiting the complexity of implementable algorithms. A method to reduce the number of glues in a tile system could thus significantly increase the complexity of experimentally-implementable algorithms.

Ma and Lombardi defined the Pattern Assembling Tile-Set Synthesis (PATS) problem to consider the minimal algorithmic tile system required to assemble a unique terminal assembly of a given pattern, with research resulting in algorithms for finding such systems [6, 23, 24], but also establishing the problem as NP-hard [16, 17, 19, 20, 23]. These methods consider abstract tile systems in a model not allowing errors that assemble to unique final assemblies from single seeds, with all single-strength glues, in a single growth direction. While Göös et al. developed a measure of reliability for such systems in a more physically-relevant model that allows errors [14], as discussed in Sect. 4, model limitations prevented the measure from being preserved during optimization and from handling proofreading error reduction behavior [4, 5, 27, 31, 35]. Without proofreading behavior through either accident or design, systems will have error rates that decrease only with the square root of assembly speed. Thus, systems without proofreading have far higher error rates, whereas with basic proofreading, error rates decrease linearly with assembly speed. For simulations of our XOR example in Sect. 6, a reduced-size version preserving proofreading could grow a 1,000 tile assembly in 200 min with 99.7% probability, while one not preserving proofreading succeeded in the same conditions in only 26% of the trials. For slower assembly speeds, the difference becomes even more pronounced: thus, proofreading has become a practical necessity for complex tile systems used in experiments, along with other design principles such as nucleation control [9, 28, 29, 35].

one or more sites in a regular lattice and will have a set number of edges facing adjacent lattice sites. In physical systems, tiles can rotate and attach in multiple orientations, but in our model, tile orientation is fixed. We include the rotations of tiles as separate tiles with fixed orientations: we define rotation functions R_i^σ , which map a tile T of structure σ to a rotation of T that may have a different structure σ' .

A seed is some structure that presents a certain number of glues on certain edges of sites on a lattice: every producible assembly will grow from a seed. For the purposes of this paper, a tile system S is a set of tiles and potential seeds, $\{T_i\} \cup \{\Sigma_j\}$; a *rotatable* tile system is the closure of a tile system under all rotations.

As in the PATS problem, we will attempt to reduce the size of a tile system by making different glues, or tiles, equivalent.

Definition 1. *The glue merge transformation, for non-complementary glues a and b of the same glue structure and bond strength, is defined as $Mg_{b,a}(X) = X'$, where every instance of a or b in X is considered to be an identical glue in X' , and every instance of a^* or b^* is considered to be an identical glue. X may be any object containing glues, such as a glue, a tile, a seed, or a tile system (Fig. 3).*




Fig. 3. Glue merge $Mg_{b,a}$ with allowed and disallowed attachment sites shown.

Some tiles or seeds in a tile system S may be mapped to identical tiles or seeds in $Mg_{b,a}(S)$, if they are already identical except for glues a and b or a^* and b^* . Thus tile merges can be defined as compositions of glue merges (here $\prod_i M_i = M_1 \circ M_2 \circ \dots$ denotes composition):

Definition 2. *The tile merge transformation, for tiles T and U of the same tile structure and color, is defined as $Mt_{U,T} = \prod_i Mg_{g_i(U),g_i(T)}$, where $g_i(T)$ is the glue on the i^{th} edge of T (Fig. 4).*

As not all pairs of glues can be merged, not all pairs of tiles have a defined tile merge transformation.

For abstract, idealized growth, we will first consider whether a series of merges will continue to result in equivalent growth in the abstract Tile Assembly Model (aTAM). In this model, starting from an initial seed, tiles attach to empty lattice sites adjacent to an assembly if they can bind by matching (complementary) glues with bond strengths that sum to at least some threshold τ , sometimes called the temperature in other papers. Once attached, tiles never detach [34]. After establishing equivalence criteria at this abstract level, more physically-accurate models can be considered.




Fig. 4. Tile merge $Mt_{U,T}$.

3 aTAM Equivalence

Since different tile systems will not have the same tiles, we first define equivalence between assemblies in different tile systems. We define two assemblies A and A' as *color equivalent* if every lattice site is either empty in both or filled with a tile of the same structure and color in both (irrespective of glues and bonds). We similarly define two tile systems S and S' as color equivalent if, for every assembly A that is producible by one system, there exists a corresponding assembly A' producible by the other that is color equivalent with A . We will define the set of all assemblies that S can produce through correct growth as $\text{PA}(S)$. Then, after some composition of merges M , we can state

Lemma 1. *If $M(\text{PA}(S)) = \text{PA}(M(S))$, i.e., if M applied to the set of producible assemblies of S is equal to the set of every producible assembly of $M(S)$, then tile systems S and $M(S)$ are **color equivalent**.*

Proof. M preserves tile color and structure, so every assembly A is color equivalent to $M(A)$. For every assembly A in $\text{PA}(S)$, $M(\text{PA}(S)) = \text{PA}(M(S))$ means that $M(A)$ will be in $\text{PA}(M(S))$, and for every assembly A' in $\text{PA}(M(S))$, it means that there will be an assembly A in $\text{PA}(S)$ such that $M(A) = A'$. \square

Intuitively, all bonds between tiles possible in S will remain possible in $M(S)$, and thus all of the same assemblies will remain producible. However, with merged glues, there may be growth possible in S' that would not be possible in S . There may also be different growth pathways that construct color equivalent assemblies. To limit equivalence to tile systems with equivalent growth pathways, we define a more restrictive goal, which is a form of bisimulation [18, 25]:

Definition 3. *Tile systems S and $M(S)$ are **growth equivalent** if for every assembly $A \in \text{PA}(S)$, every attachment site that has $Q \subseteq S$ as the set of possible correct tile attachments has $M(Q)$ as the set of possible attachments in $M(A)$.*

Lemma 2. *If tile systems S and $M(S)$ are growth equivalent, then they are color equivalent.*

Proof. For growth equivalent S and $M(S)$, at every attachment site p in every assembly A in $\text{PA}(S)$, if tile T in S can attach to form an assembly $A +_p T$, then $M(T)$ can attach to the corresponding site in $M(A)$, resulting in an assembly $M(A) +_p M(T)$. Similarly, for every tile T' that can attach to an attachment site p in $M(A)$, resulting in $M(A) +_p T'$, growth equivalence requires that there must be a tile T in S that can attach to the corresponding site in A where $M(T) = T'$. Thus if A is in $\text{PA}(S)$, every attachment step $A +_p T$ will be in $\text{PA}(S)$, and if $M(A)$ is in $\text{PA}(M(S))$, every possible attachment step will result in an assembly of the form $M(A) +_p M(T) = M(A +_p T)$. Every seed Σ in S is in $\text{PA}(S)$ and corresponds to a seed $M(\Sigma)$ in $\text{PA}(M(S))$, there are no other seeds in $\text{PA}(M(S))$, and all assemblies start from seeds; therefore by induction, growth equivalence requires that $M(\text{PA}(S)) = \text{PA}(M(S))$. \square




Fig. 5. (a) shows an assembly and potential attachment sites, which can be viewed as local neighborhoods (b) allowing certain tiles to attach. For a subset of tile systems, use annotations allow (c) the generation of input neighborhoods for each used tile, though for systems with double tiles, assemblies of two tiles (d) must also be considered.

Growth equivalence implies something stronger than just color equivalence: the set of assemblies producible from a specific assembly A will always be equivalent to those producible from $M(A)$. As a special case, A is a terminal assembly (i.e., no further tiles can attach) if and only if $M(A)$ is, which is not necessarily the case for color equivalence.

Since the aTAM assumes that tiles attach on a regular lattice and glues are additive and have non-negative strengths, whether a tile can attach to an attachment site can be determined by just the *local neighborhood* of edges adjacent to the attachment site (Fig. 5(b)). For systems of single and double tiles, a local neighborhood will be a subset of the edges adjacent to either one empty lattice site (potentially allowing attachment of a single tile) or two empty lattice sites (potentially allowing attachment of a double tile of one orientation), each labeled with a glue as if from a pre-existing tile in an assembly. An attachment site adjacent to an assembly may have multiple local neighborhoods, each with different subsets of edges. We define $PN(S)$ to be the set of local neighborhoods present in every producible assembly of S , and $PT_S(L)$ for a local neighborhood L to be the set (possibly empty) of tiles that can correctly attach. Then,

Lemma 3. *For tile system S and merge M , if for all $L \in PN(S)$, $M(PT_S(L)) = PT_{M(S)}(M(L))$, then S and $M(S)$ are growth equivalent.*

Proof. Every attachment site in an assembly A in $PA(S)$ will have some local neighborhood L , which will be in $PN(S)$, and so at that attachment site, $Q = PT_S(L)$ will be the set of tiles that can attach. The corresponding site in $M(A)$ will have local neighborhood $M(L)$, where $PT_{M(S)}(M(L))$ can attach. Thus, if $PT_{M(S)}(M(L)) = M(Q) = M(PT_S(L))$, growth equivalence will be satisfied. \square

Whether a local neighborhood is present in any producible assembly of a tile system is undecidable: considering a tile system that implements an arbitrary program and produces a particular local neighborhood only if the program halts, the question can be reduced to the halting problem. Thus, we only consider tile systems satisfying additional constraints. First, we require that the system include (correct) *use annotations* for each tile actually used in producible assemblies, designating edges on tiles as being used as *inputs* (edges with which the tile attaches to assemblies) or *outputs* (edges where other tiles attach): while tiles in

a system may have multiple use annotations, we do not consider systems where tiles attach with intentionally mismatched glues, leaving certain edges with glues *unused*. Second, we require that for every local neighborhood $L \in \text{PN}(S)$ with glues of total bond strength of at least τ , there exists a tile $T \in S$ that can attach to L with no mismatched glues, and with attachments of only input edges on the tile to output edges on the local neighborhood.

With these assumptions, for systems containing only single tiles, every $L \in \text{PN}(S)$ will consist of output edges that bind to complementary glues on corresponding input edges of some tile T , so we can enumerate all possible local neighborhoods of S , at the cost of possibly including some that are not actually producible, by examining each tile in S with input-annotated edges. To do so, we will define $\text{IN}(T)$ to be the local neighborhood (or neighborhoods, if the tile bound with strength greater than τ) corresponding to the input edges of T (Fig. 5(c)). Since $\text{IN}(S) \supseteq \text{PN}(S)$, it is clear from Lemma 3 that

Theorem 1. *For tile system S (that uses only single tiles) and merge M , if for all $L \in \text{IN}(S)$, $M(\text{PT}_S(L)) = \text{PT}_{M(S)}(M(L))$, then S and $M(S)$ are growth equivalent.*

All parts of the requirement in this theorem—local neighborhoods from input annotations on tiles, and whether other tiles can correctly attach to them—are computable, so it can be used to verify that M and $M(S)$ are growth equivalent.

In the case of a system including double tiles, since double tiles may attach to local neighborhoods with edges adjacent to two lattice sites (Fig. 5(b)), $\text{IN}(S)$ must be extended to include input neighborhoods constructed from two single tiles with matching glues and use annotations (Fig. 5(d)).

Theorem 1 provides a way of determining whether a tile system S , after a series of merges M , will continue to have only equivalent correct attachments, and thus equivalent growth in the aTAM. In more physically-relevant models, however, tiles can attach incorrectly, resulting in errors. To ensure that error behavior remains similar after a series of merges, further criteria, in a more physically-relevant model, will be needed.

4 Sensitivity Profiles and kTAM Equivalence

In the kinetic Tile Assembly Model (kTAM), any tile can attach to any potential site, and will do so at a rate $r_f = \hat{k}_f e^{-G_{\text{mc}}}$ for a constant \hat{k}_f , where G_{mc} is an analogue of the tile concentration $[c] \propto e^{-G_{\text{mc}}}$. Instead of determining whether an attachment is possible, the total bond strength of matching glues b will instead determine the rate at which a tile detaches, $r_{r,b} = \hat{k}_f e^{-bG_{\text{se}}}$, where G_{se} is the (sign-reversed and unitless) free energy of a single bond. Thus, tiles will attach at the same rate, but will fall off faster if b is smaller. To approximate $\tau = 2$ aTAM, G_{mc} is typically set to $2G_{\text{se}} - \epsilon$ for some small ϵ , such that tiles bound by $b < 2$ will fall off faster than they attach, and tiles bound by $b = 2$ will attach slightly faster than they detach. In this regime, the growth rate of an assembly will be dependent upon the bond free energy G_{se} , scaling approximately as $e^{-2G_{\text{se}}}$.

The design of tile systems that robustly exhibit the same growth in the kTAM as in the aTAM is itself an area of continuing research. In the limit of infinitely slow growth, growth in the kTAM and aTAM is equivalent, as tiles attaching by bond strength 0 or 1 will fall off far faster than they attach [34]. Moving away from this limit, however, incorrect attachments may provide pathways for the growth of assemblies that could not be produced by the system in the aTAM, where the attachments remain as errors, and further undesired growth can continue via correct attachments [11].

One approach to approximating aTAM growth in the kTAM is to minimize the rate at which errors occur in some error rate model so that assemblies of a desired size assemble perfectly with high probability. The *kinetic trapping model* (KTM) provides a model for one type of error, a *growth error*, where a tile attaches incorrectly to a site where another tile could attach correctly, and allows further growth that effectively locks the error in place [11, 34]. Another type of error, a *facet nucleation error*, can occur when a tile incorrectly attaches to a location where not all adjacent output edges are present yet [4], but these errors are beyond the scope of our analysis.

For a tile system at $\tau = 2$ containing only single tiles and where all glues have bond strengths of one, the KTM considers kTAM transition rates between empty and filled states for a single local neighborhood (Fig. 6(a)). Starting from an empty state (E), the correct tile can attach, resulting in the “correct” state C, or a tile attaching by one correct bond and one mismatch can attach, resulting in the “almost-correct” state A. C or A can revert to E by the tile falling off at the kTAM detachment rate, or, at some rate r^* related to the growth rate, can be trapped in place by further tiles correctly attaching to the tile and surrounding assembly, resulting in the trapped correct (TC) and trapped almost-correct (TA) states. When there are m possible almost-correct attachments at a site and only one possible correct attachment, the KTM predicts a growth error rate (the probability of reaching TA from E) of $P_{\text{error}} = me^{-G_{\text{se}} + \epsilon}$, with $G_{\text{mc}} = 2G_{\text{se}} - \epsilon$; as the growth rate of a system usually scales as $e^{-2G_{\text{se}}}$, this means the error rate scales as the square root of the growth rate.

The validity of the KTM error rate estimate depends on two critical assumptions, which may not always be true and may depend on the tile system or the attachment site: first, that every attachment site during the growth of an assembly has exactly m almost-correct tiles that can attach instead of the correct tile attachment, and second, that both a correct and an almost-correct tile will become kinetically trapped at the same rate by subsequent attachments. We will consider the effect of these two assumptions in turn.

As almost-correct tile attachments need one matching glue, the number of almost-correct attachments will be determined by the tile system at an abstract level, and may also depend on the local neighborhood. To analyze these attachments, we will define *first-order sensitivity*, the first of a series of *sensitivity profiles*, that will enumerate pairs of tiles (T, U) where U allows the E to A pathway in the KTM to take place in the attachment site where T would attach correctly (i.e., in the input neighborhood of T). These sensitivity profiles were




Fig. 6. (a) shows pathways in the KTM, while (b) shows potential first-order sensitive pairs for a tile, (c) shows configurations resulting in second-order pairs (black boxes represent unused edge annotations), (d) shows 2×2 uniform proofreading, and (e) shows configurations resulting in 2×2 sensitive pairs.

originally derived from glue sensitivity classes developed to analyze the effect of non-complementary glue interactions [10].

Definition 4. A pair of tiles (T, U) in S are **first-order sensitive** if and only if some subset of the input edges of T contain glues that are the same as glues on corresponding edges (regardless of use annotation) of U , with total bond strength 1 or more.

If an almost-correct attachment of a tile U in the KTM is possible in any possible local neighborhood where a tile T can attach correctly, then (T, U) is a first-order sensitive tile pair (Fig. 6(b)). For any producible local neighborhood where a tile T can attach by bond strength $b = \tau$, an almost-correct attachment in the KTM requires a tile that can attach to that local neighborhood by $b = 1$. By definition, any tile that can do so, in a system with complementary glues, will be first-order sensitive with T .

First-order sensitivity is similar to the tile system *reliability* of Göös et al. defined within the context of the PATS problem—defined, for a system constructing only a single terminal assembly, as the probability of perfectly growing that assembly in the kTAM [14]. This was calculated by combining the KTM probability of correct growth at each site in the correct assembly, considering only the C and A states and the number of potential almost-correct attachments at each site, which resulted in a reliability that decreases as the number of possible almost-correct attachments increases. Starting from uniquely-addressed systems where no almost-correct attachments were possible, their reduction techniques tended to first decrease, and then increase, their reliability measure. Similarly, one option for our methods would be reduction that attempts to only apply merges that do not add new first-order sensitive pairs, or that seeks to minimize the number of pairs.

There are two limitations to approaches such as these, related to the two assumptions underpinning the KTM error rate estimate. One, which Göös et al. address and account for, is that the number of potential almost-correct tile attachments, m , can vary from site to site. Thus, tile set reductions that decrease

m will result in a lower KTM error rate estimate. However, for algorithmic self-assembly to take advantage of the computational power at $\tau \geq 2$, there must be some correct attachment determined by two glues rather than one, and thus there must be some potential almost-correct attachments, and $m \geq 1$ for some such sites. Otherwise, the tile system will be equivalent to a $\tau = 1$ system, with its accompanying computational limitations [8]. Consequently, the lowest error rate estimate provided by this use of the KTM will still be proportional to $e^{-G_{se}}$, while the growth rate is proportional to $e^{-2G_{se}}$, resulting in an error rate that scales no better than \sqrt{r} as the growth rate r is decreased.

The second limitation concerns the assumption that all almost-correct tile attachments have the potential to become trapped at the same rate as correct tile attachments. Göös et al. do not address this issue, which accounts for the main difference in our results. In particular, when an almost-correct attachment occurs in some tile systems, there may be no tile that can attach by at least τ bond strength in the resulting local neighborhoods to trap the error in place, even for tile systems where every local neighborhood in correct growth will allow tile attachment. In this case, another growth error would be required for growth to continue, making it more likely that the initial error will detach instead—and making the KTM error rate estimate invalid. Proofreading behavior [4, 5, 27, 31, 35], where almost-correct attachments cause exactly such impediments to further correct growth, can in principle reduce error rates, for a desired $k > 1$, to scale as $e^{-kG_{se}}$, or $r^{k/2}$ for a growth rate r . In practice, this is necessary for experimental systems to have low error rates, and therefore proofreading needs to be preserved in tile system reduction.

To this end, as with the E to A transition in the KTM, we can define a *second-order sensitivity* to enumerate pairs where the A to TA transition is possible, by considering shared glues on additional tiles that can attach to first-order sensitive pairs (Fig. 6(c)):

Definition 5. For a tile system containing only single tiles, a pair of tiles (T, U) are **second-order sensitive** if they are first-order sensitive and, for some output edge on T with glue b , some tile V that can attach by an input edge with a glue b^* , and some tile W that can attach to the corresponding edge on U with a glue g^* , taking every glue g_i that is on both an input edge of V and the corresponding edge of W , both g^* and at least one g_i are at least strength 1.

Theorem 2. For a tile system of only single tiles, only second-order sensitive pairs as defined above will have a potentially valid pathway in the KTM between almost-correct attachment and a trapped almost-correct attachment.

Proof. Consider an almost-correct state in the KTM of a tile U attaching where T would have attached correctly. As shown previously, the almost-correct tile U must be first-order sensitive with T . In order to reach the trapped-almost-correct state, an additional tile must be able to attach to U and the surrounding assembly. If T had attached instead, as a correct attachment, it would have done so by input edges, and any available glues on edges adjacent to empty lattice sites would be output edges. For any of those output edges, by our requirement that

correct growth in sites with glues on adjacent output edges always be possible, there must be a tile V that attaches to that output edge by an input edge and, unless the glue on the output edge has strength $b = \tau$, at least one other adjacent output edge in the local neighborhood by some input edge on V . Thus, if U were to attach instead of T , then for a tile W to lock U in place by filling the site where V would have attached, W must share a sufficient total strength of *matching* glues on edges of V labeled as inputs (as the local neighborhood where V could attach must have corresponding outputs available). These constraints, considered for each output edge of T , are the same as the criteria for second-order sensitivity. \square

In short, the KTM error rate estimate only applies for tiles that have second-order sensitivity; the true error rate in cases that have only first-order sensitivity will be insignificant in comparison. If a tile set reduction technique were to ensure that there are no second-order sensitive tile pairs, excellent proofreading error rates could be achieved. This is the aim of our methods, although we will somewhat soften this goal below.

The second-order definition and Theorem 2 above are valid only for single tiles; double tiles add additional complexity in that tile edges further away from the initial error may be involved in allowing the trapping second attachment to occur. Defining second-order sensitivity that accounts for double tiles is possible, but would require consideration of a large number of potential local configurations. Our current second-order sensitivity implementation treats double tiles as two single tiles with edges that can be inputs or outputs depending upon which results in second-order sensitivity; by doing so, it does not account for certain error pathways, but for many systems with few double tiles that do not have all six glues, it is sufficient to find most pathways.

In the ideal case of a system with any number of first-order pairs, but no second-order pairs, applying the KTM only where it is valid would predict *no* growth errors: no almost-correct attachment could be trapped in place. Proofreading transformations, however, usually satisfy the weaker constraint of requiring that an error prevent correct growth at some later point (rather than immediately), such that breathing of the growth front is still likely to remove the initial error. For 2×2 uniform proofreading [35], which is the simplest to implement experimentally, an initial almost-correct attachment of a tile U can, at worst, allow a further attachment on one edge of U , but will not allow any correct attachment on a second edge. To attempt to preserve such behavior, we can construct a further sensitivity profile to find pairs that could violate it:

Definition 6. *A pair of tiles (T, U) are 2×2 sensitive if they are second-order sensitive and the second-order criteria can be simultaneously satisfied on two different output edges of T , or the second-order criteria can be satisfied on one output edge of T , and U is a double tile that extends out along a second output edge of T .*

Intuitively, for tile systems with only strength-1 glues, only tile pairs that are 2×2 sensitive will allow growth to continue with no further impediment after an

almost-correct attachment. While a perfect 2×2 uniform proofreading design will have second-order sensitive pairs, it will have no 2×2 sensitive pairs, and thus the sensitivity profile is also useful for checking proofreading implementation. However, the sensitivity profile has no rigorous significance with respect to the KTM.

5 Considerations Beyond the kTAM: Lattice Defects and Spurious Hierarchical Assembly

The aTAM and kTAM both assume that tiles assemble into perfect regular lattices. Physical DNA tile lattices, however, have some degree of flexibility, and tiles can form bonds with other tiles that fall outside of a perfect lattice, creating *lattice defects*, as shown in Fig. 7. In general, lattice defect formation depends on numerous physical factors, and would be difficult to model rigorously. However, it would be beneficial to have a method for avoiding their formation, and ensuring that in reducing tile system size, their likelihood is not increased.

To do so, in an approach similar to sensitivity profiles, we search for small assemblies of tiles that could create neighborhoods where a tile could attach by two correct bonds and form a lattice defect. For tile systems of DAO-E tiles, we speculate that the simplest, smallest lattice defects, in the orientations likeliest to allow bonds to form between non-adjacent tiles, will be the likeliest, as shown in Fig. 7(b) and (c). If in every possible combination of tiles in the pattern of one of these defects, no tile can attach by two correct bonds, then lattice defects of that size should not be possible without previous errors or other growth directions.

Another potential concern beyond the kTAM is that, for growth in solution, assemblies may bind to other assemblies, a process that has been utilized in other “hierarchical” models of self-assembly [3], but is generally undesirable in the systems designed to grow by single tile attachments. Such *spurious hierarchical attachments* have been seen in some experiments [13], but the importance of design criteria to avoid them is unclear. Many systems designed by hand, for simplicity, use each glue consistently on only input or output edges of tiles, thus avoiding assembly-assembly interactions because no glue on the edges of




Fig. 7. (a) shows an AFM image of a DAO-E tile system lattice defect, while (b) shows a layout of a simple DAO-E lattice defect. (c) illustrates a simple algorithm for enumerating potential small lattice defects, in two orientations: circles represent arbitrary bonds of non-zero strength. (d) shows a spurious hierarchical attachment.

assemblies, which will be outputs, will be complementary to glues on edges of other assemblies. Such a distinction can be broken when merging glues, potentially allowing assemblies, while at lower concentrations than tiles, to attach by numerous bonds. The distinction can be preserved, however, by requiring that merges do not result in any glue being on both input and output edges of tiles.

6 Algorithm and Results

With the combination of criteria for aTAM equivalence, sensitivity profiles, lattice defect pathways, and spurious hierarchical attachment, a simple algorithm to attempt merges and check the resulting tile systems can be implemented. While more optimized search algorithms could improve performance [14, 24], this simple algorithm suffices to demonstrate our reduction methods. In principle, as tile merges are simply the compositions of glue merges, a search through all potential glue merges could minimize a tile system, as tile merges would result automatically. In practice, for our non-exhaustive searches, we have found that first attempting possible tile merges and then glue merges is most effective at reducing the numbers of both tiles and glues.

For both searches, our algorithm is the same. For every pair of tiles, or pair of glues (filtered, if desired, by spurious hierarchical attachment criteria), we attempt to merge the pair with a transformation M . We then perform checks for aTAM equivalence per Theorem 1, and for a desired set of sensitivity profiles, check whether every pair of sensitive tiles $(M(A), M(B))$ in the merged system has a corresponding pair (A, B) of the same profile in the initial system. If either criteria fails because of a pair of tiles, the algorithm attempts, recursively, to merge those two tiles until either the criteria are satisfied or the necessary merge would be impossible. If satisfied, the resulting tile system is checked for new potential lattice defects. If all these criteria are satisfied, subsequent merge attempts are applied to the resulting system. However, in all merges, use annotations and input neighborhoods are used from the original tile system for sensitivity profile and aTAM equivalence: if they are used to generate input neighborhoods from an equivalent merged system, they will result in more local neighborhoods that will not be present in any producible assembly, thus overly constraining reduction.

This reduction algorithm is currently implemented in Alhambra, a software package for tile system design and compilation [1]. To examine the effect of the reductions on physically-implementable tile systems, we used three tile systems we had previously designed in Alhambra: XOR, Crosshatch, and Rule 110, all shown in Fig. 8. Each of the three implements 2×2 uniform proofreading, and grows from an origami seed. Additionally, we combined the three reduced systems into a “Combined” tile system implementing all three simultaneously, and reduced this system again with the same parameters.

As the algorithm can preserve different sets of sensitivity profiles, size results for several choices are shown in Table 1. Ignoring sensitivity entirely results in aTAM-equivalent systems that are similar to what might be found by PATS




Fig. 8. Structures from example tile systems. The XOR system implements a ribbon of constant, seed-defined width, with the center of the ribbon implementing the XOR function to create a Sierpinski triangle pattern, and boundaries that reflect bits. The Crosshatch (XH) system implements a ribbon with “signals” that bounce back and forth diagonally: when a signal reaches the “north” boundary, it causes the ribbon to shrink, resulting in assemblies of a finite, input-specified size. The Rule 110 system (P110) implements the cellular automata system Rule 110 using expanding boundaries, one of which uses a zig-zag growth order.

methods applied to systems that are not uniquely-addressed [14,24]: they do not preserve proofreading error rates, and, as seen in simulation results in Fig. 9 (Fig. 1 for the Combined system), have error rates that scale as $e^{-G_{se}}$. Regardless of whether spurious hierarchical attachments are avoided, preserving either 2×2 sensitivity or both second-order and 2×2 almost always results in systems that are significantly reduced in size but preserve proofreading error rate scaling of $e^{-2G_{se}}$, with error rates orders of magnitude below the sets that ignore sensitivity. The exceptions to this scaling, P110 with second-order and 2×2 preserved, both have results that fail to preserve proofreading by chance because of a pathway for errors through double tile attachments, illustrating the need for a theory and implementation of second-order sensitivity that addresses double tiles in order to most effectively reduce the size of systems such as P110 which make more extensive use of them. Fortunately, although again by chance, reductions of P110 with less stringent criteria avoided the problems with double tiles—and resulted in smaller tile systems as well.

7 Discussion

The results from our reduction methods suggest that it may be possible to implement tile systems of considerable algorithmic complexity in experimentally-viable ways using numbers of tiles and glues that are comparable with existing experimental systems. In considering experimental implementations, the unreduced Combined system, with 173 glues, uses more glues than we speculate is currently experimentally feasible for algorithmic DX tile systems. The number has been approached in uniquely-addressed DX tile systems using 152 glues [32], but uniquely-addressed systems are much less sensitive to glue quality, and so the numbers are not directly comparable. The 2×2 -preserving combined system, however, preserves proofreading behavior while having fewer glues (36) than the

Table 1. Sizes of tile systems before and after reduction preserving different sets of sensitivity profiles: SHA refers to preventing spurious hierarchical attachments by restricting glue usage. Each reduction used 350 tile reduction trials, then 35 glue reduction trials on each of the 10 results with the least glues. “Combined” combines XOR, XH, and P110. Tile counts are of tiles needed for implementation, while glue are counted as sequence/complement pairs

	XOR		XH		P110		Combined	
	Tiles	Glues	Tiles	Glues	Tiles	Glues	Tiles	Glues
Before reduction	32	44	41	56	46	73	119	173
Preserve 2nd, 2×2 , SHA	30	25	38	31	38	30	103	62
Preserve 2nd, 2×2	29	18	36	23	36	23	99	46
Preserve 2×2 , SHA	27	22	33	28	31	24	91	50
Preserve 2×2	25	15	30	17	32	17	88	36
Ignore sensitivity	13	6	15	6	24	9	58	13

unreduced XOR system (44), which is well within the range of experimental feasibility of DX tile systems. Previously published algorithmic DX tile systems have used 35 [9], 34 [29], and 23 [2] glues, and sequence space searches have found sets around 80 glue sequences of comparable quality [10]. For systems compiled into other physical implementations, such as single-stranded tiles (SST), the larger sequence space could make even larger systems experimentally viable; uniquely-addressed SST assemblies have been demonstrated using over 10,000 glues [26]. Yet even with DX tiles, our reductions could allow experimental implementation of tile systems, using current experimental methods, far beyond the complexity of those we would otherwise be able to implement.

There are a number of other potential directions for optimizations in a tile system compiler. Preserving larger uniform proofreading transformations (e.g. $k \times k$ for $k \geq 3$) would require expanded sensitivity profiles. Other tile system properties, such as avoidance of facet nucleation errors [4] and barriers to spurious nucleation [28], would also be desirable to preserve, and it may be possible to develop similar criteria. Additionally, while our methods are intended for optimizing tile systems that are already algorithmic, it may be interesting to consider the combination of sensitivity criteria with uniquely-addressed systems and more sophisticated PATS search methods. Uniquely-addressed tile systems have no error pathways in the KTM and no sensitive pairs of any order; it is only in merging tiles that the KTM becomes applicable. Thus, preserving sensitivity profiles while reducing a uniquely-addressed assembly could result in systems that exhibit strong proofreading behavior without any need for a proofreading transformation, and avoid $e^{-G_{se}}$ error rate scaling.

Another approach might be to go beyond preserving desirable properties in tile systems that have already incorporated proofreading, and search for changes that add them to tile systems designed without proofreading. Rather than trying to reduce the size of a system, an optimizing compiler could try to split tiles and

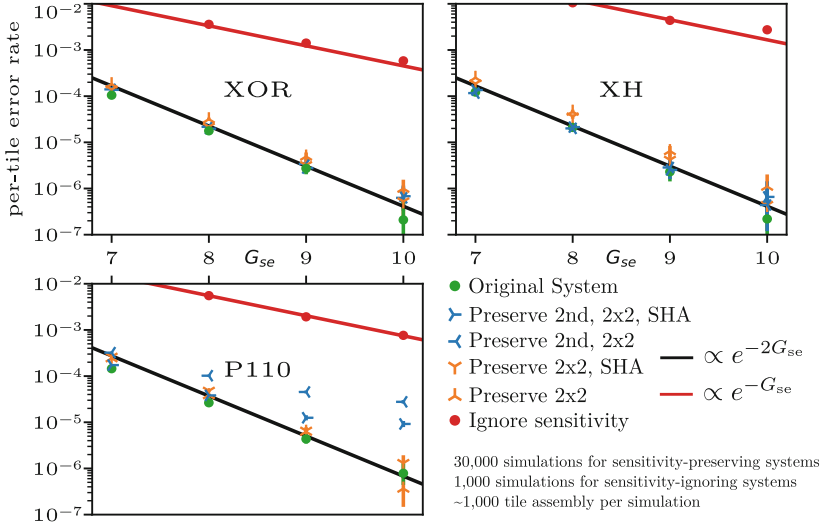


Fig. 9. Per-tile error rates in kTAM simulations, in Xgrow via Alhambra, varying G_{se} , with $G_{mc} = 2G_{se} - \log 2$. Assembly time in this regime scales approximately as $e^{2G_{se}}$; at $G_{se} = 7$, a 1,000 tile assembly will grow in about 4 min, and at $G_{se} = 9$, in about 4 h. Per-tile error rates were determined from the percentage of perfect assemblies.

glues to preserve aTAM behavior while improving or introducing proofreading, barriers to spurious nucleation, facet nucleation error rates, and lattice defect formation. Such a tile-system-specific approach could find systems with behaviors similar to those provided by general tile system transformations and design principles while being smaller, easier to design, and possibly more effective.

Acknowledgments. We thank Chigozie Nri, Philip Petersen, Lulu Qian, and Grigory Tikhomirov for discussions and collaboration on physical implementations and the Alhambra compiler, and Robert Johnson and William Poole for discussions on aTAM equivalence. This work was partially supported by the Evans Foundation and National Science Foundation award CCF-1317694.

References

1. Alhambra. <https://github.com/DNA-and-Natural-Algorithms-Group/alhambra>
2. Barish, R.D., Schulman, R., Rothmund, P.W.K., Winfree, E.: An information-bearing seed for nucleating algorithmic self-assembly. *PNAS* **106**(15), 6054–6059 (2009). <https://doi.org/10.1073/pnas.0808736106>
3. Cannon, S., et al.: Two hands are better than one (up to constant factors): self-assembly in the 2HAM vs. aTAM. In: Portier, N., Wilke, T. (eds.) *STACS 2013. LIPIcs*, vol. 20, pp. 172–184. Dagstuhl (2013). <https://doi.org/10.4230/LIPIcs.STACS.2013.172>

4. Chen, H.-L., Goel, A.: Error free self-assembly using error prone tiles. In: Ferretti, C., Mauri, G., Zandron, C. (eds.) DNA 2004. LNCS, vol. 3384, pp. 62–75. Springer, Heidelberg (2005). https://doi.org/10.1007/11493785_6
5. Chen, H.L., Schulman, R., Goel, A., Winfree, E.: Reducing facet nucleation during algorithmic self-assembly. *Nano Lett.* **7**, 2913–2919 (2007). <https://doi.org/10.1021/nl070793o>
6. Czeizler, E., Popa, A.: Synthesizing minimal tile sets for complex patterns in the framework of patterned DNA self-assembly. *Theor. Comput. Sci.* **499**, 23–37 (2018). <https://doi.org/10.1016/j.tcs.2013.05.009>
7. Doty, D.: Theory of algorithmic self-assembly. *Commun. ACM* **55**(12), 78–88 (2012). <https://doi.org/10.1145/2380656.2380675>
8. Doty, D., Patitz, M.J., Summers, S.M.: Limitations of self-assembly at temperature 1. *Theor. Comput. Sci.* **412**(1–2), 145–158 (2011). <https://doi.org/10.1016/j.tcs.2010.08.023>
9. Evans, C.G.: Crystals that count! Physical principles and experimental investigations of DNA tile self-assembly. Ph.D. thesis, California Institute of Technology (2014). <http://resolver.caltech.edu/CaltechTHESIS:05132014-142306756>
10. Evans, C.G., Winfree, E.: DNA sticky end design and assignment for robust algorithmic self-assembly. In: Soloveichik, D., Yurke, B. (eds.) DNA 2013. LNCS, vol. 8141, pp. 61–75. Springer, Cham (2013). https://doi.org/10.1007/978-3-319-01928-4_5
11. Evans, C.G., Winfree, E.: Physical principles for DNA tile self-assembly. *Chem. Soc. Rev.* **46**(12), 3808–3829 (2017). <https://doi.org/10.1039/C6CS00745G>
12. Fu, T.J., Seeman, N.C.: DNA double-crossover molecules. *Biochemistry* **32**, 3211–3220 (1993). <https://doi.org/10.1021/bi00064a003>
13. Fujibayashi, K., Hariadi, R., Park, S.H., Winfree, E., Murata, S.: Toward reliable algorithmic self-assembly of DNA tiles: a fixed-width cellular automaton pattern. *Nano Lett.* **8**(7), 1791–1797 (2008). <https://doi.org/10.1021/nl0722830>
14. Göös, M., Lempäinen, T., Czeizler, E., Orponen, P.: Search methods for tile sets in patterned DNA self-assembly. *J. Comput. Syst. Sci.* **80**(1), 297–319 (2014). <https://doi.org/10.1016/j.jcss.2013.08.003>
15. Jacobs, W.M., Reinhardt, A., Frenkel, D.: Rational design of self-assembly pathways for complex multicomponent structures. *PNAS* **112**(20), 6313–6318 (2015). <https://doi.org/10.1073/pnas.1502210112>
16. Johnsen, A., Kao, M.Y., Seki, S.: A manually-checkable proof for the NP-hardness of 11-color pattern self-assembly tileset synthesis. *J. Comb. Optim.* **33**(2), 496–529 (2017). <https://doi.org/10.1007/s10878-015-9975-6>
17. Johnsen, A.C., Kao, M.-Y., Seki, S.: Computing minimum tile sets to self-assemble color patterns. In: Cai, L., Cheng, S.-W., Lam, T.-W. (eds.) ISAAC 2013. LNCS, vol. 8283, pp. 699–710. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-45030-3_65
18. Johnson, R., Dong, Q., Winfree, E.: Verifying chemical reaction network implementations: a bisimulation approach. *Theor. Comput. Sci.* (2018). <https://doi.org/10.1016/j.tcs.2018.01.002>
19. Kari, L., Kopecki, S., Meunier, P.É., Patitz, M.J., Seki, S.: Binary pattern tile set synthesis is NP-hard. *Algorithmica* **78**(1), 1–46 (2017). <https://doi.org/10.1007/s00453-016-0154-7>
20. Kari, L., Kopecki, S., Seki, S.: 3-color bounded patterned self-assembly. *Nat. Comput.* **14**(2), 279–292 (2015). <https://doi.org/10.1007/s11047-014-9434-9>

21. Ke, Y., Ong, L.L., Shih, W.M., Yin, P.: Three-dimensional structures self-assembled from DNA bricks. *Science* **338**(6111), 1177–1183 (2012). <https://doi.org/10.1126/science.1227268>
22. Lin, C., Liu, Y., Rinker, S., Yan, H.: DNA tile based self-assembly: building complex nanoarchitectures. *ChemPhysChem* **7**(8), 1641–1647 (2006). <https://doi.org/10.1002/cphc.200600260>
23. Ma, X., Lombardi, F.: Combinatorial optimization problem in designing DNA self-assembly tile sets. In: 2008 IEEE International Workshop on Design and Test of Nano Devices, Circuits and Systems, pp. 73–76 (2008). <https://doi.org/10.1109/NDCS.2008.7>
24. Ma, X., Lombardi, F.: Synthesis of tile sets for DNA self-assembly. *IEEE Trans. Comput.-Aided Des. Integr. Circ. Syst.* **27**(5), 963–967 (2008). <https://doi.org/10.1109/TCAD.2008.917973>
25. Milner, R.: *Communication and Concurrency*. Prentice Hall, Upper Saddle River (1989)
26. Ong, L.L., et al.: Programmable self-assembly of three-dimensional nanostructures from 10,000 unique components. *Nature* **552**(7683), 72–77 (2017). <https://doi.org/10.1038/nature24648>
27. Reif, J.H., Sahu, S., Yin, P.: Compact error-resilient computational DNA tiling assemblies. In: Ferretti, C., Mauri, G., Zandron, C. (eds.) *DNA 2004*. LNCS, vol. 3384, pp. 293–307. Springer, Heidelberg (2005). https://doi.org/10.1007/11493785_26
28. Schulman, R., Winfree, E.: Programmable control of nucleation for algorithmic self-assembly. *SIAM J. Comput.* **39**(4), 1581–1616 (2010). <https://doi.org/10.1137/070680266>
29. Schulman, R., Yurke, B., Winfree, E.: Robust self-replication of combinatorial information via crystal growth and scission. *PNAS* **109**(17), 6405–6410 (2012). <https://doi.org/10.1073/pnas.1117813109>
30. Seeman, N.C., Sleiman, H.F.: DNA nanotechnology. *Nat. Rev. Mater.* **3**, 17068 (2017). <https://doi.org/10.1038/natrevmats.2017.68>
31. Soloveichik, D., Winfree, E.: Complexity of compact proofreading for self-assembled patterns. In: Carbone, A., Pierce, N.A. (eds.) *DNA 2005*. LNCS, vol. 3892, pp. 305–324. Springer, Heidelberg (2006). https://doi.org/10.1007/11753681_24
32. Wang, W., Lin, T., Zhang, S., Bai, T., Mi, Y., Wei, B.: Self-assembly of fully addressable DNA nanostructures from double crossover tiles. *Nucleic Acids Res.* **44**(16), 7989–7996 (2016). <https://doi.org/10.1093/nar/gkw670>
33. Wei, B., Dai, M., Yin, P.: Complex shapes self-assembled from single-stranded DNA tiles. *Nature* **485**(7400), 623–626 (2012). <https://doi.org/10.1038/nature11075>
34. Winfree, E.: Simulations of computing by self-assembly. Technical report, CaltechCSTR:1998.22, Pasadena, CA (1998). <https://doi.org/10.7907/Z9TB14X7>
35. Winfree, E., Bekbolatov, R.: Proofreading tile sets: error correction for algorithmic self-assembly. In: Chen, J., Reif, J. (eds.) *DNA 2003*. LNCS, vol. 2943, pp. 126–144. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24628-2_13