Progressive Alignment of Shapes

Ashwin Gopinath* David Kirkpatrick† Paul W. K. Rothemund‡ Chris Thachuk§

Abstract

We introduce a natural property of shape pairs, that, starting from any initial overlapping configuration, they can be brought into a unique configuration of maximum overlap by a continuous motion that monotonically increases their overlap. The identification of such shapes is motivated by applications of self-assembly, driven by molecular forces, in nanofabrication processes.

1 Introduction

Certain shapes, a disk is the simplest example, have the property that, starting from any initial placement with non-zero overlap with a target placement, there is a continuous rigid motion (a simple translation between centres, in the case of a disk) that takes the shape to its target placement with monotonically increasing overlap. We are interested in designing–and certifying–such self-aligning shapes. More generally, we are interested in shape-target pairs, under the additional restriction that the motion terminates in a unique placement that maximizes the overlap with the target. (Because of their rotational symmetry, disks clearly do *not* satisfy this restriction.)

We formalize the notion of shape alignment with some preliminary definitions. A (planar) *shape* is a connected bounded subset of $\Re^2$. A *placement* λ of a shape A is a proper rigid transformation of A (expressed as a translation and rotation of A, with respect to its centre of mass). We will frequently refer to both λ and the resulting shape, denoted A_λ, as a placement of A. Given a target shape T, the *T-overlap* of a placement A_λ is just area of intersection of T with A_λ. Placements with non-zero T-overlap are said to be *T-proximate*. A placement A_λ is locally (resp. globally) *optimal* with respect to target shape T if the T-overlap of A_λ is locally (resp. globally) maximum in the space of all placements of A. A *T-alignment* of shape A (or simply *alignment*, when T is understood), from placement λ_0 to placement λ_1, is a continuous function μ from $[0,1]$ to T-proximate placements of A, with $\mu(0) = \lambda_0$ and $\mu(1) = \lambda_1$. An alignment μ of A is *T-progressive* if the T-overlap of $A_{\mu(j)}$ exceeds that of $A_{\mu(i)}$, for all $0 \leq i < j \leq 1$. The pair (T, A) is an *optimally aligning pair* if there is a T-progressive alignment of A to a unique globally-optimal placement, from every T-proximate placement of A. (Note that in this case A has exactly one locally optimal placement.) If (A, A) is an optimally aligning pair then we say that A is a *self-aligning shape*.

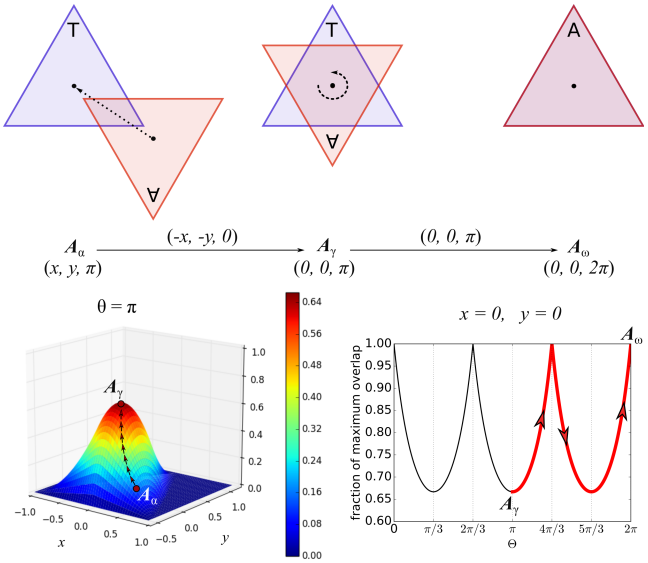

Figure 1: a two phase alignment pathway $\Pi = A_\alpha, A_\gamma, A_\omega$ for an equilateral triangle A and target placement T. In the first phase, shape A is translated (from its initial placement A_α to placement A_γ) along the vector that brings its centre coincident with that of T. In the second phase, A is rotated by π until it coincides with T (placement A_ω). Shown above are the relative placements of A and T. Shown below is the relevant projection of the alignment landscape that contains the pathway Π (the directed path highlighted in red).

We begin with a simple example to illustrate some of the concepts defined above (see Fig. 1). Consider an equilateral triangle A, with sides of length 1, centred at $(0, 0)$. An alignment of A, from some initial placement A_α (with rotation π) to some final placement A_ω (with rotation 2π), is a continuous sequence of placements, from A_α to A_ω.

The shape-pair (T, A) defines an *alignment landscape*

*Bioengineering, California Institute of Technology, ashwing@caltech.edu

†Department of Computer Science, University of British Columbia, kirk@cs.ubc.ca

‡Bioengineering and Computing & Mathematical Sciences, California Institute of Technology, pwkr@caltech.edu

§Computing & Mathematical Sciences, California Institute of Technology, thachuk@caltech.edu

in 4-dimensional space. An example pathway in that landscape is illustrated in Fig. 1. Due to the 3-fold rotational symmetry of A, and since A and T are congruent, the alignment landscape has three global maxima where A has been rotated by angle $\theta \in \{0, 2\pi/3, 4\pi/3\}$. After the centre of A has become coincident with that of T (placement A_γ) the alignment path rotates A through a first and then a second global maximum. Note that the example path is *not* T-progressive, but the prefix of the path ending at the placement with rotation $4\pi/3$ is progressive.

Observe that the absence of self-similarity alone is not sufficient to guarantee that a shape is self-aligning (see Fig. 2).

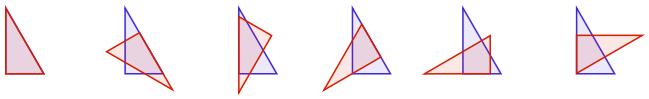

Figure 2: a right triangle is not self-aligning since many placements are local maxima.

This leaves us with the question: what shapes are self-aligning or, more generally, what shape-pairs are optimally aligning? One strategy is to separately model each specific alignment landscape with sufficient precision to effectively rule out the existence of undesirable local maxima. Our goal is to address this question using more precise geometric arguments that (i) bring more clarity to our understanding (in terms of parameter settings for certain families of shapes) of the shape characteristics that support self-alignment, and (ii) help to inspire the design of novel shapes with other desirable characteristics.

In the next section, we describe some of the motivation for the study of optimally-aligning shape-pairs that derives from applications in nanofabrication. We also review some related work and identify the key differences in the analytic approach used in this paper.

Section 3 sets out our results on optimally-aligning shape-pairs in a staged fashion, starting with shapes that progressively align using pure translations or pure rotations, and moving on to those whose progressive alignment draws on a combination of these motions.

Our presentation is more illustrative than comprehensive, hopefully raising more interesting questions than we have settled.

2 Background and motivation

The problem of designing/characterizing self-aligning shapes, or, more generally, optimally aligning shape pairs, is a natural geometric problem, worthy of study without further motivation. However, it turns out to have significant practical importance as well. Consider the task, arising in the context of nanofabrication, of producing a particular pattern on a flat surface. The pattern may be printed or etched (*e.g.* lithography), or individual components placed precisely until the pattern emerges. These are examples of pattern assembly directed by an external process. In contrast, autonomous *self-assembly* is the process by which a designed pattern emerges by assembling itself from constituent parts.

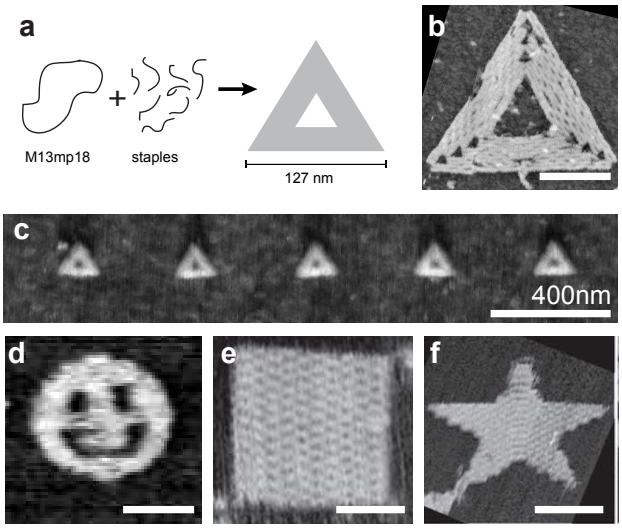

Figure 3: (a) DNA origami technology can fabricate 2D shapes on the nanoscale, such as a triangular annulus. (b) An atomic force microscope image of the triangular annulus. (c) Triangular annuli in a lattice arrangement on a surface. (d)–(f) Atomic force microscope images of other shapes realized with DNA origami. Scale bars in (b) and (d)–(f) are 50 nanometers.

A very successful example of self-assembly, driven by molecular forces, is DNA origami [6] which has become a foundational technology in nanoscience [2]. The process involves a long *scaffold* sequence of DNA — typically a circular plasmid — and a collection of short *staple* sequences that are complementary to two or more regions on the scaffold. When annealed in an appropriate buffer solution, each staple strand "pinches" regions of the scaffold together according to their designed complementarity. Typically, this process can result in the assembly of $\approx 10^{10}$ copies of a designed 2D or 3D shape, with feature resolution of 6 nanometers (nm) [6]. In contrast, the smallest resolution of any feature in current CMOS processes is 14 nm.

Efforts have sought to place these DNA origami on a surface. Kershner *et al.* [4] and Gopinath & Rothemund [3] demonstrated a combination of directed- and self-assembly by creating a surface with a regular lattice of triangular patches, via ebeam lithography, to which triangular shaped origami could bind (see Fig. 3c).

That work was the original motivation for our study of the progressive alignment of shapes since (i) origami

can initially land (from solution) onto the surface[1] in any orientation or translation relative to a patch, and (ii) the electrostatic force between patch and origami, coupled with stochastic perturbation, drives a process to improve the binding between the pair. To ensure an origami cannot become "stuck" in a degenerate placement, the energy landscape exhibited by the patch-origami pair should have the following property: from *all* initial placements there should exist a path, to one or more of the designated final placements, that monotonically increases the number of chemical bonds (in this case salt-bridges) between the origami and surface. We model this in the shape alignment problem by ensuring there is a progressive alignment from any initial placement to a designated final placement.

While equilateral triangles have been demonstrated experimentally to place well, each origami may be in one of three final states due to the 3-fold rotational symmetry of the patch-origami pair. It is natural to ask whether there exists a patch-origami pair that has a *unique* local maximum of overlap. Such a pair would enable placement of a single molecule on a surface, with 6 nm precision, since molecules can be attached to an origami at this same resolution. In this way, DNA origami acts as a molecular *breadboard*.

Böhringer and co-authors [5, 7, 1] encountered the problem of designing self-aligning shapes in the context of the study of surface-tension driven self-assembly of micro-components in microelectromechanical systems (MEMS). Their analysis is based on a first-order approximation of the energy model that reduces the problem to the self-alignment question posed above. They performed an experimental evaluation of a wide variety of shapes, and an exact analysis of a family of shapes formed by addition and subtraction of solid disks.

Their exact analysis led to a complete characterization of the shapes, called death-stars below, that are guaranteed to self-align. In a nutshell, this analysis relies on the simplicity of regions formed by intersecting disks to give a precise expression for the area of intersection of death-stars in a specified placement, as well as its derivative with respect to a specified direction of motion. This approach is simply not feasible for more complicated shapes that necessarily arise in our intended application. Furthermore, it is focused on optimizing certain shape parameters, ultimately through the numerical solution of certain systems of constraints, that, for the kind of simple alignment paths that arise in practice, are rather easy to approximate using more straightforward geometric arguments.

The analysis techniques used in this paper are based on a geometric description of motion paths, and can be used to show a larger class of shape-pairs have the requisite properties. This is important in the context of DNA origami where shapes are rasterized; those whose boundary has low curvature can be approximated well, those with high curvature cannot. Our approach is to view the area of intersection of a target shape T and a placement A_λ of a movable shape A as an integral of infinitesimal strips (referred to as *cuts*) parallel to a specified direction of motion. Then, to confirm that motion from A_λ in this specified direction leads to an increase in the T-overlap, it suffices to analyze and accumulate the change within each cut. (In effect we differ from the previous approach by simply interchanging the order of integration and differentiation.) The utility/versatility of this approach stems from the fact that the boundary of $T \cap A_\lambda$ can have a simple description in terms of portions of the boundary of T and A_λ. Thus, the change within each cut can be characterized by the boundary-type of the endpoints of the intervals of $T \cap A_\lambda$ that live within that cut. We make this approach more concrete in the next section.

3 Examples of certified progressive alignments

3.1 Progressive alignment by translation

Consider again the case of a simple disk shape. Fig. 4 (left) illustrates a typical T-proximate placement of a red disk A relative to a target placement T (blue). If A is translated horizontally, bringing its centre closer to that of T, the area of T-overlap clearly increases. One way of seeing, and quantifying, this is to imagine slicing the overlap lens into infinitesimal cuts, parallel to the direction of translation, and to analyse the (instantaneous) change of length of each cut. As illustrated, every cut is bounded on the left by the boundary of A and on the right by the boundary of T, and so each cut increases in length in proportion to the length of the translation of A. It follows that the rate of increase in overlap is proportional to the vertical extent (i.e. the length of the vertical projection) of that portion of the boundary of A that forms part of the left boundary of the intersection.

Note that if shape A is replaced by any shape whose smallest enclosing disk coincides with A, then we can, in a similar way, measure the rate of increase in overlap by the length of the vertical projection of that portion of the boundary of A that forms part of the left boundary of the intersection, minus the (necessarily smaller) length of the vertical projection of that portion of the boundary of A that forms part of the right boundary of the intersection.

To see a slightly more general situation, consider the case of an annulus (see Fig. 4 (right)). In this case the intersection in a typical T-proximate placement is a region that, when sliced into infinitesimal horizontal cuts as before, has some cuts bounded on the left by the

[1] A 2D (flat) origami can be designed to ensure that only one of its two faces can stick to the surface.

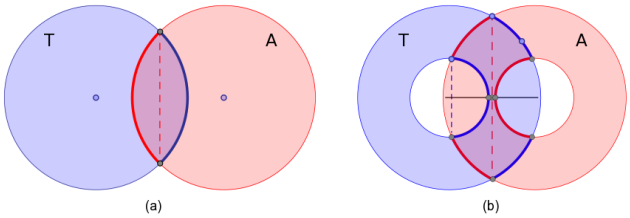

Figure 4: (a) disk and (b) annulus placements.

boundary of A and others by the boundary of T. Similarly for the right end of the cuts. As before, cuts that are bounded on the left by the boundary of A and on the right by the boundary of T increase in length as A is translated left. However, just the opposite is true for cuts with opposite bounding conditions. Furthermore, cuts bounded on both ends by the boundary of the same shape do not change in length under infinitesimal motion of A. Thus we can measure the total rate of change by accumulating the length of the vertical projection of the portion(s) of the boundary of A that coincide with the left boundary of the intersection, minus the length of the vertical projection of the portion(s) of the boundary of A that coincide with the right boundary of the intersection, plus the length of the vertical projection of the portion(s) of the boundary of T that coincide with the right boundary of the intersection, minus the length of the vertical projection of the portion(s) of the boundary of T that coincide with the left boundary of the intersection.

To ensure that a translation taking the centre of A onto the centre of T is progressive (i.e. the rate of change is always positive) for an annulus, it suffices to choose the radius of the inner circle of the annulus in such a way that for all T-proximate placements the height of the lens (red dashed segment) is at least twice the height of the inner and outer circle intersection (dashed blue segment). It is easy to confirm that, subject to this constraint, the rate of change is minimized when the annulus "holes" are disjoint. It follows that, with this same choice of inner radius, the same annulus and target, with arbitrary content inside the inner circles, will progressively align to a configuration in which the annulus and target centres coincide.

3.2 Progressive alignment by rotation

Of course, even if a shape is self-aligning, it is only coincidentally possible to achieve this by translation alone. Fortunately, our analysis of progressive alignments using pure translation has a direct counterpart for pure rotations. As before it is helpful to illustrate this with a simple example. Fig. 5 (left) illustrates a proximate placement of a yin-yang shape A (red) with respect to a congruent target T (blue). It is not surprising that translation to a configuration (Fig. 5 (right)) in which the centres coincide is not progressive. Nevertheless, once the centres coincide a rotation about the centre suffices to complete the alignment in a progressive fashion.

In the case of the T-proximate yin-yang configuration in Fig. 5 (right), observe that at every radius r no larger than the outer radius the intersection of the circle of radius r (the counterpart of an infinitesimal strip) with the current placement-target overlap is an arc bounded on the counter-clockwise end by a portion of the boundary of A and on the clockwise end by a portion of the boundary of T. Thus a counter-clockwise rotation of A will increase the length of all such arcs, ensuring that the intersection increases until it reaches it unique maximum.

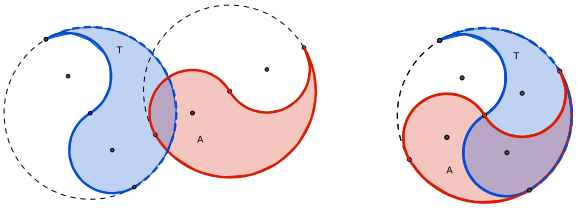

Figure 5: yin-yang shape placements.

Of course, many other shapes, including a half circle, enjoy this same progressive rotational alignment property. We have illustrated the yin-yang shape because its central, but non-reflective, symmetry turns out to be useful in certain applications.

3.3 Progressive alignment by translation then rotation

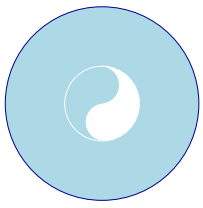

Figure 6: a self-aligning shape.

Combining the observations of the previous two sections, it follows directly that an annulus with a sufficiently small inner radius whose hole circumscribes a copy of the yin-yang shape (see Fig. 6) is self-aligning.

3.3.1 Self-alignment of disk with offset hole

Another way to realize a self-aligning shape is to break the symmetry of an annulus and offset the "hole". As

previously noted, such a family of shapes was the focus of a detailed analysis by Böhringer et al. [5, 7]. Here we observe that a straightforward application of the approach discussed in the previous two subsections allows us to draw conclusions about the specifications for such shapes, and generalizations, that ensure the self-aligning property.

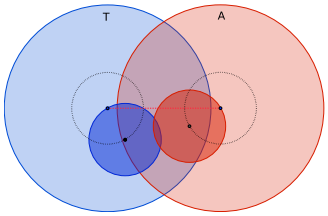

Figure 7: a generic death-star placement.

For obvious reasons we refer to the unit-radius disk with an offset hole of radius r as a "death-star" (although our intended application involves a realization at a considerably different scale). Fig. 7 illustrates a death-star A (red) in a proximate placement with respect to a target T (blue). Here holes are depicted in bold color to help in visualizing their interaction.

We will maintain the invariant that the star centres lie on the x-axis. Our progressive alignment in all cases consists of two phases. Phase 1 is a simple translation, taking the centre of star A onto the centre of star T. Phase 2 is a rotation of star A about its centre so as to reduce the separation of the hole centres to zero. It is easy to see that, provided the holes overlap the star centre, phase 2 monotonically increases the star overlap, since the hole intersection increases monotonically. Thus it remains to analyze the behaviour of phase 1.

We first observe that in the absence of intersections between holes and the opposite star, which is unavoidable when the distance between the star centres is at least $2r + 1$, a translation of A towards T increases the intersection of the stars in proportion to the height of the lens formed by the intersection of the disks A and T (just as in the case of intersecting disks discussed earlier). More generally, for smaller star separations, when the holes may intersect their opposite star, we can, as before, view the star intersection (which of course is a subset of the lens) as being made up of infinitesimal cuts, whose length may increase, decrease or remain unchanged as a result of such a translation.

As described in the analysis of annulus alignments, cuts are of four types depending on the boundary types of their left and right delimiters. Cuts delimited on the left by a portion of the A star boundary and on the right by a portion of the T star boundary, increase in length in proportion to the length of the translation. Cuts delimited on the left and right by portions of the boundary of the same star do not change length with such translations. Finally, cuts delimited on the left by a portion of the T star boundary and on the right by a portion of the A star boundary, decrease in length in proportion to the length of the translation. Since the right (respectively, left) boundary of the A disk (respectively, T disk) is disjoint from the lens, it follows that every shrinking cut connects a point on the right boundary of the T hole to a point on the left boundary of the A hole, and every unchanging cut involves a point on the right boundary of the T hole or a point on the left boundary of the A hole. Thus there is an expanding cut at every height that intersects the lens, but neither of the hole boundaries within the lens. Hence the star intersection increases at least in proportion to the height of the lens minus the sum of lengths of the vertical projections of the hole boundaries within the lens. (Here we have used the fact that cuts that intersect both hole boundaries are double counted in the vertical projections. Furthermore, we have not discounted the possible hole intersection, which can only reduce the shrinking cuts, or those portions of the right boundary of the B hole, or the left boundary of the B hole, that intersect the lens, which can only add to the expanding cuts.)

Thus, it remains to demonstrate that in all configurations the height of the lens exceeds the sum of lengths of the vertical projections of the hole boundaries within the lens. Note that this is obvious if the distance between A and T is less than $4r$ (a necessary condition for the holes to intersect) provided r is chosen so that the height of the lens is at least $4r$, in this case.

When the distance between A and T is at least $4r$, we observe that the length, as well as the vertical projection, of both hole boundaries within the lens is maximized when the distance between the centre of A and the centre of the T hole (respectively, the centre of T and the centre of the A hole) is minimized, i.e. the hole centres lie on the x-axis (see Fig.8). So, we hereafter study configurations in which both hole centres lie on the x-axis. Note: this does not suggest that our transformation rotates the stars to realize such a configuration; in fact such a rotation would in general reduce the start intersection.

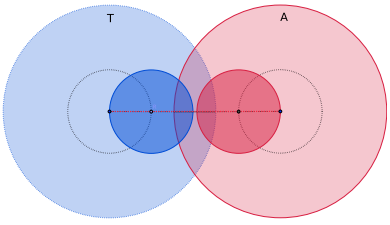

Figure 8: a placement when both hole centres lie on the x-axis.

It remains to analyse the rate of change of the star overlap as A is translated to make its centre coincident

with that of T. As in the case of the annulus, this can be achieved by studying the length of the vertical projections of the portions of the boundary of A and T that delimit the left and right boundaries of $A \cap T$. Since these projections change in a monotonic way with a decrease in the separation of A and T, it is straightforward to establish a largest value for r that will ensure that this translation phase is progressive.

As with the annulus we note that if the star A is modified in such a way that its hole is a proper subset of the A hole then the translation of this shape to T remains progressive. Of course, there is no guarantee that the second (rotation) phase would remain progressive with such a modification, but our argument for progressive rotational alignments allows us to make such an assertion in certain special cases, for example when the A hole is reduced to an inscribed square.

3.3.2 Progressive alignment of regular polygons

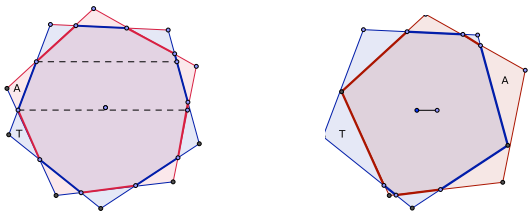

Figure 9: placements of regular 5-gons.

As a final example of the application of our approach to the certification of progressive alignments using translation and rotation, we consider the natural question of whether every k-regular polygon A progressively aligns to one of its maximally overlapped configurations with a congruent target T (see Fig. 9). It is straightforward to see that this is true if the intial and final configurations share a common centre. Thus it suffices to argue that, for any T-proximate placement A_λ, a translation taking A to a placement A_0 whose centre coincides with the centre of T, is progressive.

To confirm that this is the case we imagine any translation taking A from a centre-aligned placement A_0 in any direction, and argue that the overlap $A \cap T$ decreases with distance. This follows from several simple observations. First, we note that the boundary of the intersection $A_0 \cap T$ is semi-regular convex polygon consisting of $2k$ equal length sides drawn in alternating fashion from the sides of A_0 and T. (This is easily demonstrated by appealing to mirror symmetry across any line through the common centre and any point of intersection of the boundaries of A_0 and T.)

It follows from this that contracting the sides of the boundary of $A_0 \cap T$ that correspond to boundary edges of A_0 (respectively T) forms a regular k-gon. Thus, when we accumulate the vertical projections of portions of the boundaries of A_0 and T that determine the boundary of $A_0 \cap T$ we get a net change of zero for any instantaneous translation from the placement A_0.

The second observation is that, for any placement formed by an infinitesimal translation from placement A_0, the magnitude of the associated projections that count positively in the accumulated projection length decrease with distance, and those that contribute negatively increase. This demonstrates that for all placements in the immediate neighbourhood of a centre-aligned placement the derivative of the change of overlap, with respect to distance, is negative, confirming that any centre-aligned placement is a local maximum with respect to all translations.

To confirm that every centre-aligned placement is in fact a *global* maximum, with respect to all translations, we note that the second observation extends to translations that take A from A_0 to any placement whose symmetric difference with T contains three or more components. (To demonstrate this we cut the overlapped placements into strips, parallel to the direction of translation, bounded by lines through each of the polygon vertices, and accumulate the change of overlap in each strip separately.)

Finally, we note that if A has been translated to a T-proximate placement whose symmetric difference with T has two components, then an argument analogous to that used in the limiting case of disks (recall Fig. 4(a)) shows that such placements can be improved by any infinitesimal translation in the direction that takes the centres back into alignment.

4 Conclusion and Future Work

We have identified an interesting property of planar shapes that models a desirable attribute of self-assembling nanofabricated structures. We have identified a simple approach to the certification of this progressive alignment property and illustrated it with some fundamental examples, some of which are beyond the reach of earlier approaches, and all of which serve as potential building blocks for the design of shapes with other desirable properties in a nanofabrication context.

Our study of self-aligning shapes is still in a preliminary state. One indication of this is that at this point we still do not know if there exist convex shapes that are self-aligning. Furthermore, while we have reasonable tools for certifying progressive alignments that are composed of pure translations and pure rotations, these do not allow us to address arbitrary rigid motions, which presumably could be required for some progressive alignments.

Acknowledgments

The authors are supported by the Natural Sciences and Engineering Research Council of Canada (DK), a Banting Fellowship (CT), the Office of Naval Research (award N000141410702) (AG, PWKR), the US National Science Foundation (award 1317694) (AG, PWKR), and the Molecular Programming Project (http://molecular-programming.org) (AG, CT, PWKR). We thank W. Evans for helpful discussions and the anonymous reviewers for their feedback.

References

[1] K. F. Böhringer, U. Srinivasan, and R. T. Howe. Modeling of capillary forces and binding sites for fluidic self-assembly. In *Micro Electro Mechanical Systems, 2001. MEMS 2001. The 14th IEEE International Conference on*, pages 369–374. IEEE, 2001.

[2] Editorial. Returning to the fold. *Nature Materials*, 15(3):245–245, 2016.

[3] A. Gopinath and P. W. Rothemund. Optimized assembly and covalent coupling of single-molecule DNA origami nanoarrays. *ACS nano*, 8(12):12030–12040, 2014.

[4] R. J. Kershner, L. D. Bozano, C. M. Micheel, A. M. Hung, A. R. Fornof, J. N. Cha, C. T. Rettner, M. Bersani, J. Frommer, P. W. Rothemund, et al. Placement and orientation of individual DNA shapes on lithographically patterned surfaces. *Nature Nanotechnology*, 4(9):557–561, 2009.

[5] S.-H. Liang, X. Xiong, and K. F. Böhringer. Towards optimal designs for self-alignment in surface tension driven micro-assembly. In *Micro Electro Mechanical Systems, 2004. 17th IEEE International Conference on.(MEMS)*, pages 9–12. IEEE, 2004.

[6] P. W. Rothemund. Folding DNA to create nanoscale shapes and patterns. *Nature*, 440(7082):297–302, 2006.

[7] X. Xiong, S.-H. Liang, and K. F. Böhringer. Geometric binding site design for surface-tension driven self-assembly. In *Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on*, volume 2, pages 1141–1148. IEEE, 2004.