



Soft Matter

Thermodynamic Stability versus Kinetic Accessibility: Pareto Fronts for Programmable Self-Assembly

Journal:	<i>Soft Matter</i>
Manuscript ID	SM-ART-05-2021-000681.R1
Article Type:	Paper
Date Submitted by the Author:	16-Jun-2021
Complete List of Authors:	Trubiano, Anthony; New York University Courant Institute of Mathematical Sciences, Mathematics Holmes-Cerfon, Miranda; New York University Courant Institute of Mathematical Sciences, Mathematics

SCHOLARONE™
Manuscripts

Cite this: DOI: 00.0000/xxxxxxxxxx

Thermodynamic stability versus Kinetic Accessibility: Pareto Fronts for Programmable Self-Assembly[†]

Anthony Trubiano and Miranda Holmes-Cerfon

Received Date

Accepted Date

DOI: 00.0000/xxxxxxxxxx

A challenge in designing self-assembling building blocks is to ensure the target state is both thermodynamically stable and kinetically accessible. These two objectives are known to be typically in competition, but it is not known how to simultaneously optimize them. We consider this problem through the lens of multi-objective optimization theory: we develop a genetic algorithm to compute the Pareto fronts characterizing the tradeoff between equilibrium probability and folding rate, for a model system of small polymers of colloids with tunable short-ranged interaction energies. We use a coarse-grained model for the particles' dynamics that allows us to efficiently search over parameters, for systems small enough to be enumerated. For most target states there is a tradeoff when the number of types of particles is small, with medium-weak bonds favouring fast folding, and strong bonds favouring high equilibrium probability. The tradeoff disappears when the number of particle types reaches a value m_* , that is usually much less than the total number of particles. This general approach of computing Pareto fronts allows one to identify the minimum number of design parameters to avoid a thermodynamic-kinetic tradeoff. However, we argue, by contrasting our coarse-grained model's predictions with those of Brownian dynamics simulations, that particles with short-ranged isotropic interactions should generically have a tradeoff, and avoiding it in larger systems will require orientation-dependent interactions.

1 Introduction

Designing novel materials or structures out of mesoscopic building blocks, such as colloids, is a widespread challenge, with applications in drug delivery, optical metamaterials, and environment sensing micro-robots, among others^{1–6}. A fundamental issue is to design the building blocks so the desired target state is both *thermodynamically stable*, i.e. it remains near the target state for long times at finite temperature, and *kinetically accessible*, i.e. the system self-assembles into the target state over a reasonable timescale. In general, guaranteeing a target state is thermodynamically stable does not guarantee it is kinetically accessible, and vice versa, with these two objectives often working against each other^{7–12}. Yet, we know from the plethora of successfully self-assembling material and biological systems, that it is possible to achieve both objectives simultaneously given the right choice of parameters. Consider proteins: while randomly-designed proteins take astronomically long times to fold into their

native states^{13–15}, the proteins found in nature fold quickly, a great many orders of magnitude faster than it would take them to sample their thermodynamic equilibrium distribution, a phenomenon known as Levinthal's paradox^{16,17}. The reason why is said to be that naturally occurring proteins have evolved so their free energy landscape is “funnel-shaped”, with the folded structure at the bottom of the funnel and most dynamic pathways leading downhill, encountering only small free energy barriers along the way^{13,18–20}.

Inspired by proteins, can we design material systems to similarly fold quickly to a deep metastable state? That is, given a collection of building blocks with various design parameters (such as interaction energies, concentrations, shapes, etc), can we choose parameter values so the system achieves both objectives of thermodynamic stability and kinetic accessibility simultaneously? This question lies at the heart of most theoretical studies of programmable self-assembly, but it is a challenge to optimize two objectives simultaneously. To simplify, earlier studies of self-assembly considered evaluating measures of thermodynamic stability and kinetic accessibility as design parameters are varied along an axis; such studies found there is a narrow range of parameter values in the observation regions where assembly may occur, and furthermore that the finite time yield of the target state is non-monotonic in the binding energy^{7,9,11}. More recent

Courant Institute of Mathematical Sciences, New York University, New York, New York 10012, USA. Email: trubiano@cims.nyu.edu, holmes@cims.nyu.edu

[†] Electronic Supplementary Information (ESI) available: See DOI: 10.1039/cXsm00000x/

The authors contributed equally to this work.

studies have considered self-assembly in an optimization framework, by building methods to optimize a single objective measuring progress toward the target state. One such objective is the partition function, measuring thermodynamic stability; much progress has been made in understanding how to choose pair interaction energies between distinct particles to make a target structure low free energy^{21–26}, motivated partly by the spectacular self-assembly of DNA-brick based structures²⁷. Adjusting the concentrations of the components of a target structure can further lower its free energy²⁸. Kinetic effects, while harder to study, have been controlled via nucleation barriers²⁹, geometrical measures of the ruggedness of the free energy landscape³⁰, or by maximizing the yield of a target structure at a fixed time^{31,32}, an approach which couples thermodynamic and kinetic considerations into a single objective. Other measures of kinetics may be tuned by building automatic differentiation into molecular dynamics simulations, which promises to be a powerful tool for optimizing design parameters of a self-assembling system³³.

In this paper we take an alternative approach: rather than optimizing for a single measure of the success of programmable self-assembly, we consider the two objectives of thermodynamic stability and kinetic stability separately, and seek to understand and ultimately to minimize the tradeoff between these two objectives as the design parameters are varied. We focus our examples on a model system, namely polymers of colloids interacting with short-ranged isotropic attractive forces, where the design parameters are the energies of the attractive interactions. Such particles can be synthesized experimentally by gluing single-stranded DNA onto the surfaces of the colloids, and are a promising set of building blocks for designing new materials^{34–38}. We consider small collections of colloids, which have been thoroughly studied both experimentally^{39,40} and theoretically^{21,22,41}. We coarse-grain their dynamics to efficiently search over parameters, which allows us to exhaustively characterize their thermodynamic-kinetic tradeoffs for this coarse-grained system, a pedagogical characterization that has not been illustrated before.

We characterize the thermodynamic-kinetic tradeoff for a target structure by computing its *Pareto front*, namely the set of values attained by the objectives such that it is not possible to further increase both objectives⁴². The Pareto front is part of the boundary of everything that is achievable by the objective functions as the design parameters are varied. This curve gives qualitative insight into what sets the tradeoff between thermodynamics and kinetics for different target structures, and it is useful as a practical design tool: one should always choose parameter values that give objective values on this curve, since otherwise one could increase both objectives using different parameters. Furthermore, the Pareto front allows a user to *control* the tradeoff between the objectives, by deciding how many design parameters to vary. As the number of design parameters increases, the Pareto front necessarily becomes steeper, implying the tradeoff between the objectives becomes less significant. Strikingly, in our examples, there is a sharp transition where the tradeoff suddenly becomes minimal. We argue that this is a signature of a “funnel-shaped” landscape.

The transition to a steep Pareto front gives the minimal number of design parameters to avoid a tradeoff, within the context of

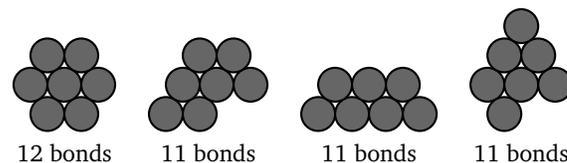


Fig. 1 A system of 7 identical disks that illustrates the tradeoff between thermodynamic stability and kinetic accessibility. The leftmost, “flower” cluster, has the highest equilibrium probability when the interactions are strong, because it has the most bonds. However, the average time to form the flower increases exponentially with the energy of the interaction, because most trajectories hit another clusters first^{43,44}.

our coarse-grained dynamical model. Sometimes the same design parameters also lead to fast folding in Brownian dynamics simulations, and sometimes they don’t. When they don’t, we show it is because of “chiral traps”, configurations that cannot reach the target state without breaking a strong bond, but that have the same bonds as a cluster on the correct folding pathway. We argue that because of these inevitable chiral traps, particles with short-ranged isotropic interactions should generically have a tradeoff between thermodynamics and kinetics; avoiding this tradeoff in larger systems will require particles with orientation-dependent interactions.

2 There is a tradeoff between thermodynamic stability and kinetic accessibility

A simple example that illustrates the tradeoff between thermodynamic stability and kinetic accessibility for particles with short-ranged attractive interactions is shown in Figure 1⁴³. This figure shows clusters of 7 disks in the plane, that are energetic local minima when the particles interact with a short-ranged, isotropic pairwise attractive potential. For identical disks, when the bond energy E is large, then the leftmost, “flower” cluster has the lowest free energy, because it has 12 bonds whereas the other clusters have only 11 bonds. Therefore, as E increases, the equilibrium probability of the flower (when sufficiently confined) approaches 1, making it thermodynamically stable. However, as E increases, the average time it takes to form the flower increases too, because most dynamical pathways hit one of 11-bond clusters first^{43,44}, and breaking a bond to reach the flower happens on a timescale proportional to $e^{E/kT}$. Therefore, the flower becomes less kinetically accessible as it becomes more thermodynamically stable.

This tradeoff occurs even when the particles are non-identical. Consider now a system of 6 spheres which are constrained to stay in a chain, as a polymer. We add this backbone constraint, partly to mimic proteins, and partly to obtain a system that does not evaporate without introducing undesired features like walls. We assume here and throughout the paper that particles interact with a pairwise Morse potential, $U(r) = E \left(e^{-2\rho(r-d)} - 2e^{\rho(r-d)} \right)$, with diameter $d=1$, range parameter $\rho=40$, and an energy E of the interaction that varies depending on the particle pair. There are two clusters that are energetic local minima when E is the same for all pairs, an octahedron and a polytetrahedron (Figure 2(a)). Both have 12 bonds hence the same potential energy, but their differing rotational and vibrational entropies imply that when $E=12kT$,

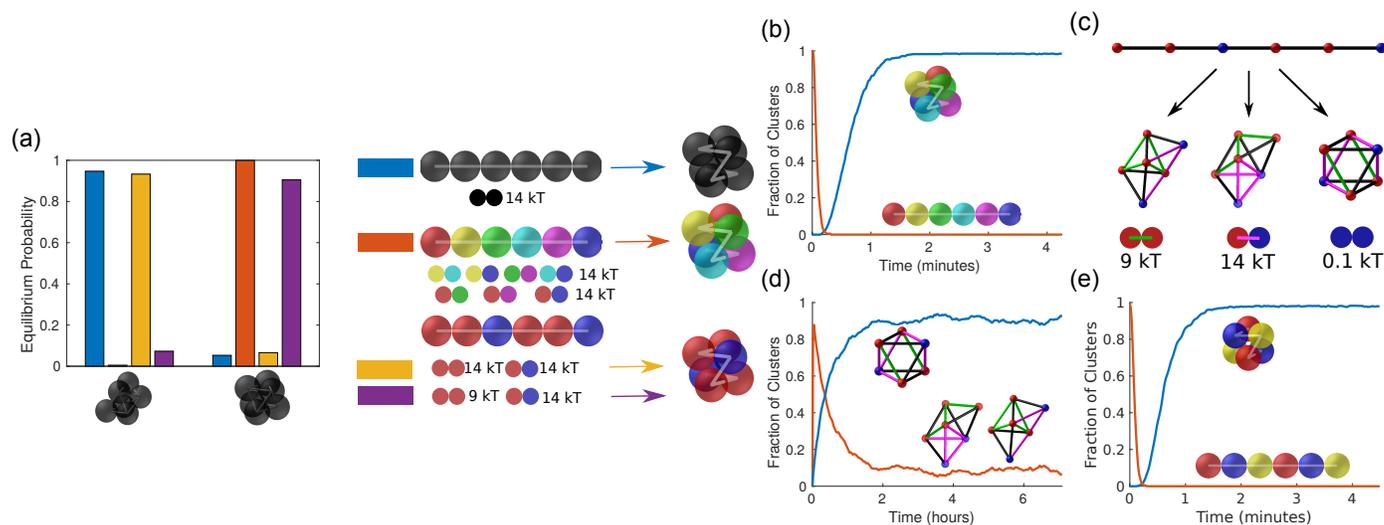


Fig. 2 There is a thermodynamic-kinetic tradeoff to form an octahedron of 6 spheres, but this can be avoided with the right choice of particle types and parameters. (a) Equilibrium probabilities for low energy clusters of polymers of 6 spheres, the polytetrahedron (left) and octahedron (right), with different design parameters (bond energies). The bond energies are listed in the adjacent legend, for: identical particles, distinct particles, two types of particles arranged in an AABAAB configuration (A=red, B=blue). Any unlabeled interactions are $0.1kT$. (b) The octahedron forms within *minutes* in a Brownian dynamics simulation, when formed from distinct particles with interactions given above. The timescale is estimated for $1.3 \mu\text{m}$ colloids⁴⁰. (c) Competing ground states for the AABAAB chain, showing AA bonds (green), AB bonds (pink), and backbone bonds (black). (d) The octahedron forms within *hours* when made from two types of particles (interactions given in purple legend), because of the presence of kinetic traps. (e) With three types of particles and optimal interactions $E_{AB}=E_{AC}=E_{BC}=14kT$ (all other pairs $0.1kT$), the octahedron forms within minutes, as for distinct particles. There is no more thermodynamic-kinetic tradeoff.

the octahedron has an equilibrium probability of only 0.0535 – compare this to the polytetrahedron’s equilibrium probability of 0.9456 (SI Appendix[†], Section 1). The remaining probability is associated with floppy clusters.

Let’s try to make the octahedron have high probability, by letting the particles be distinct. Following the principles of^{21,23,41}, which claim that optimal assembly occurs when bonds in the target state are identically strong and other bonds are identically weak, choose $E=14kT$ for interactions in the target octahedron, while $E=0.1kT$ for all other interactions. The weak bonds have non-zero energy, to avoid computational difficulties with exact zeros, and because it is hard to make DNA-coated colloids that have strictly zero attraction⁴⁵. With these interactions, the octahedron has equilibrium probability 0.99995 (Figure 2(a)). It furthermore assembles quickly, with few kinetic traps: Figure 2(b) shows that in Brownian dynamics simulations starting from a linear chain, the octahedron reaches a nearly steady probability of 0.985 after about 1 minute, where the timescale is estimated for micron-scale colloids using the experimental parameters in⁴⁰ (SI Appendix[†], Section 2). The yield is not perfect because the chain sometimes reaches a floppy state in which existing bonds geometrically block additional correct bonds from forming. We return to this issue in Section 5, but for now use these results for six particle types as a benchmark.

Making all particles distinct with controlled interactions is currently not feasible for DNA-coated colloids, and furthermore, biology tells us we should not require a number of particle types proportional to system size – proteins form tens of thousands of distinct structures from strings of hundreds to thousands of only 20 types of amino acids. To see if we can make the octahedron from

fewer types of particles, we consider a chain with only two particle types, A and B, ordered along the chain as AABAAB. The target configuration has two AA bonds and five AB bonds. Naively applying the same principle as before, we set $E_{AA}=E_{AB}=14kT$ and $E_{BB}=0.1kT$. The octahedron now has an equilibrium probability of 0.08. This is still small, because there are two polytetrahedra with the same total number of AA and AB bonds, hence the same potential energy, but which are favoured entropically (Figure 2(c)).

As might have been expected, making target bonds strong and non-target bonds weak is not sufficient to form a target state when the number of particle types is restricted. However, a bond-counting argument suggests the octahedron can form if we additionally allow bonds with *medium* strength. The octahedron has five AB bonds while the competing states only have three or four, and it has two AA bonds while the competing states have four or three (Figure 2(c)). By making AB bonds strong and AA bonds have medium strength, the potential energy of the octahedron should be lowered compared to the polytetrahedra. With $E_{AB}=14kT$, $E_{AA}=9kT$, the equilibrium probability for the octahedron is now 0.905 (Figure 2(a)). As $E_{AA}, E_{AB} \rightarrow \infty$ with $E_{AB} - E_{AA} \rightarrow \infty$, the equilibrium probability approaches 1.

Having identified bond energies that give the octahedron high equilibrium probability, we test whether it forms in a dynamical simulation. Figure 2(d) shows the octahedron reaches a nearly steady probability of about 0.9, but only after about 2 hours. Compared to the timescale of 1 minute for distinct particles, this is significantly slower, more than two orders of magnitude slower.

The octahedron forms slowly because the polytetrahedra are kinetic traps: as the octahedron’s free energy decreases (by in-

creasing E_{AA}, E_{AB}), so too does the polytetrahedron's free energy, and simultaneously the energy barrier to transition from a polytetrahedron to the octahedron increases. Many trajectories hit a polytetrahedron before the octahedron, and must break AA bonds to form the octahedron, which happens slowly. With two types of particles arranged as AABAAB, it appears there is a tradeoff between the thermodynamic stability and the kinetic accessibility of the octahedron.

3 Quantifying the tradeoff between thermodynamic stability and kinetic accessibility

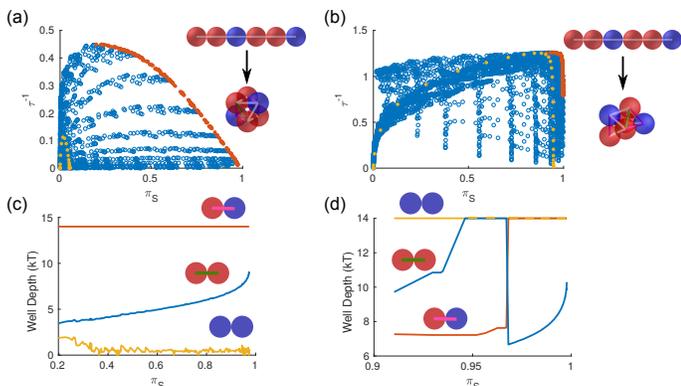


Fig. 3 Quantifying the tradeoff between equilibrium probability π_S and folding rate τ^{-1} for a chain of AABAAB spheres. (a,b) Scatter plots (π_S, τ^{-1}) -values (outcomes) to each of the two rigid clusters, (a) octahedron and (b) polytetrahedron. The octahedron forms the permutation shown plus one other, while the polytetrahedron forms 12 energetically equivalent permutations which are lumped together. Blue points are outcomes evaluated with bond energies E_{AA}, E_{AB}, E_{BB} on a grid, each taking 20 equally spaced values in $[0.1, 14]$. Yellow points are outcomes with $E_{AA}=E_{AB}=E_{BB}$, hence correspond to identical particles. Orange points, on the boundary of the set of outcomes, form the Pareto front, computed using our genetic algorithm. (c,d) Bond energies along the Pareto fronts for each cluster, as a function of the equilibrium probability π_S .

Can we choose bond energies such that the octahedron forms quickly yet has high equilibrium probability? And can we further understand the tradeoff, for this and other systems? We answer these questions by framing them as a multi-objective optimization problem: we wish to simultaneously maximize two functions, one measuring thermodynamic stability and one measuring kinetic accessibility. Usually it is not possible to maximize both functions at once, in which case we wish to quantify the tradeoff between the two.

To proceed, let us measure thermodynamic stability by π_S , the equilibrium probability of all configurations S that correspond to a target state, such as the octahedron, and let us measure kinetic accessibility by τ^{-1} , a function we call the *rate*, where τ is the mean first passage time (mfpt) from the linear chain to the target state. The functions π_S and τ^{-1} are our *objective functions*, or *objectives*, the functions we wish to maximize. We have a collection of design parameters, such as the interaction energies between particles of different types. The values of the objectives change as the design parameters vary.

A technical but important issue is evaluating the objectives for

different design parameters. Evaluating these by brute-force simulation is extremely time-consuming, so we instead construct a coarse-grained model that lets us efficiently evaluate the objectives for different design parameters. Our model starts by lumping together configurations that share the same adjacency matrix, and then approximates the dynamics as a Markov chain on the set of feasible adjacency matrices, including those corresponding to floppy configurations. The rates of forming a bond are estimated from simulations and are assumed to be independent of bond energies, and the rates of breaking a bond are set by detailed balance, where the equilibrium probabilities for each node of the Markov chain (each adjacency matrix) are estimated by a Monte Carlo sampling procedure with fixed bond energies⁴⁶ (SI Appendix[†], Section 3). Obtaining the initial rates and equilibrium probabilities with one set of bond energies is time-consuming, but once these are calculated we may reweight them for other energies. Then, calculating π_S and τ^{-1} simply requires solving a linear algebra problem for each different set of bond energies. This approach is only feasible for systems small enough that all the adjacency matrices (including permutations) are possible to enumerate.

Let's explore our model for different types of particles. With one type of particle we may vary the energy E of the common interaction. Figure 3(a) (yellow dots) shows the rate τ^{-1} is maximized for medium-weak energies, $E \approx 5.2$. At this energy π_S is small because the floppy configurations are also probable. As E increases beyond this value, τ^{-1} decreases and π_S increases, approaching their limiting values of $(\pi_S, \tau^{-1}) = (0.054, 0)$ as $E \nearrow \infty$. The set of achievable (π_S, τ^{-1}) -values, called the set of *outcomes*, describes the tradeoff between stability and accessibility for identical particles.

Now consider a chain of two types of particles arranged along a chain as AABAAB. We vary the three energies $\vec{E} = (E_{AA}, E_{AB}, E_{BB})$ on a grid. Figure 3(a) (blue dots) shows a scatter plot of the outcomes (π_S, τ^{-1}) . These outcomes appear to fill a two-dimensional region of (π_S, τ^{-1}) -space. This region has expanded well beyond the outcomes for identical particles, with a maximum rate about four times higher and a maximum equilibrium probability approaching 1. Yet, there are still no outcomes that maximize π_S and τ^{-1} simultaneously. As for identical particles, the parameters that give the maximum rate, $\vec{E} = (3.5, 14, 1.9)$ are not all large – fast folding appears to require some weak bonds. Also as for identical particles, there are parameters such that $\pi_S \rightarrow 1$, but in this limit $\tau^{-1} \rightarrow 0$.

A key observation is that the set of outcomes appears to be bounded – in between the outcomes with maximum rate, and maximum equilibrium probability, the outcomes appear to lie on one side of a curve. This curve, which we have approximated using an algorithm to be described later (orange points), will play a key role in our analysis. We argue that if one wishes to design a self-assembling system, one should choose parameters that give outcomes on this curve. To see why, suppose one chooses parameters that give an outcome away from the curve, in the interior of the set of outcomes. Then there exist outcomes, to the right and upward, that have higher equilibrium probability (without decreasing the rate), higher rate (without decreasing the equilib-

rium probability), or both. One can continue to vary parameters to improve either objective until one reaches the boundary of the set of outcomes, where one can no longer increase both objectives simultaneously.

In the language of multi-objective optimization, an outcome such that no objective value can be increased without decreasing the value of another is called *Pareto optimal*⁴². The collection of all such outcomes is called the *Pareto front*. The curve that bounds the set of outcomes in Figure 3(a) is a Pareto front.

The Pareto front is a powerful design tool. With this curve in hand, a designer can choose an outcome that weights the objectives according to their interests, and then read off the parameter values (bond energies) that give this outcome. For example, a designer who wishes to obtain at least 50% equilibrium yield, might choose parameters that give the maximum τ subject to $\pi_S > 0.5$ so the system also assembles as quickly as possible; a designer with a limited time budget might choose parameters that give a value of τ comparable to their desired timescale, ensuring the system is likely to reach the target state but also has the best probability of remaining at it. This approach of providing a set of optimal values, which is common in design problems that consider multiple objectives (e.g. building a product that has high quality but is not too costly, creating a cancer radiation treatment plan that is effective but safe, designing airfoils that are energy efficient while minimizing sound pressure)⁴², is agnostic to which objective function is better than the other; it allows different designers to place different levels of importance on the objectives.

4 Pareto fronts for a collection of examples

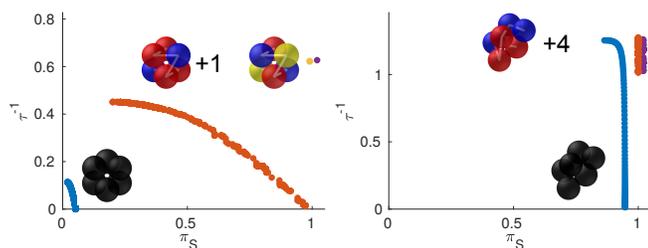


Fig. 4 Pareto fronts for polymers of 6 spheres with m particle types: $m=1$ (blue), $m=2$ (orange), $m=3$ (yellow), $m=6$ (purple), for the octahedron (left) and polytetrahedron (right). The octahedron's Pareto front with $m=3$ types is a single point, implying there is no thermodynamic-kinetic tradeoff. The polytetrahedron's Pareto front with $m=2$ types is essentially vertical, implying the same. The optimal particle labels are AABAAB and ABCABC for the octahedron, and ABAABB for the polytetrahedron.

Since the Pareto front describes the fundamental tradeoff between thermodynamic stability and kinetic accessibility as the bond energies are varied, we turn our attention to computing this curve for systems described by our coarse-grained model. We compute the Pareto front using a genetic algorithm.

The genetic algorithm works as follows. We initialize a population of P members with bond energies sampled randomly from a uniform distribution on $[0.1, E_M]$, where $E_M = 12-14$ depending on the system. At each step of the algorithm, we evaluate the objectives π_S, τ^{-1} for each member and then use these outcomes

to compare members. We say member X is *dominated* by member Y if member Y has larger values in each objective than member X . We sort the members by the number of dominating members, from least to greatest, and use this list to construct the next generation. All non-dominated members, those dominated by no other members, are carried to the next generation. The rest of the population dies, and is replaced by offspring. The offspring are created by choosing two random parents, from the top $p\%$ of the current generation, and choosing the child's parameters by randomly assigning them from each parent. There is also an $r\%$ chance the child undergoes mutation for any given parameter, meaning that parameter is re-sampled from the original distribution. We used $r = 10$ and p in the range $[30, 50]$. We repeat this procedure for each generation until one of two end conditions is met: either the iteration limit is reached, or all members are non-dominated (Appendix, Section 4).

This algorithm also allows us to consider different particle labelings, e.g. AABBA, ABBBB, BAAAA, etc, by allowing particle type to be an inheritable parameter in our genetic algorithm, assuming a fixed number m of types. During mating, the type of particle i in the child configuration is chosen randomly from the two parents. If it mutates, the types is sampled uniformly from the collection of m types. This adds a negligible amount of computation per iteration, at the cost of increasing the number of iterations until convergence.

4.1 Octahedron and polytetrahedron

The Pareto front for the AABAAB octahedron is shown in Figure 3(a) (orange points). This curve lies on the boundary of the set of outcomes, verifying the algorithm is working correctly. From the members on the Pareto front one obtains a parameterization of the bond energies as a function of either objective. Figure 3(c) shows the bond energies as a function of π_S . The strong bond E_{AB} is constant at the maximum value, the weak bond E_{BB} fluctuates near the minimum value, and E_{AA} increases smoothly along the Pareto front, from a medium-small value where the rate τ^{-1} is highest, to a larger value where the equilibrium probability π_S is highest.

We repeated this calculation for the AABAAB polytetrahedron (Figure 3(b)). The Pareto front is a union of two curves, one nearly horizontal and one nearly vertical. Each curve favors a different permutation of the polytetrahedron with different numbers of AA, AB bonds. The Pareto front does not extend down to $\tau^{-1} \rightarrow 0$ (when E_{BB} is bounded) because as E_{AA} increases, the equilibrium probability begins to decrease as the particles become identical. A notable difference from the octahedron is that the polytetrahedron's Pareto front is much steeper. This means the tradeoff between thermodynamic stability and kinetic accessibility is smaller, since a steeper slope means the system can still achieve a high rate without sacrificing much equilibrium probability.

How can we design a system so its Pareto front is steep? So far we have fixed the particle labels, but allowing them to vary gives greater design flexibility. Figure 4 shows the Pareto fronts for the octahedron and polytetrahedron when particle label is included

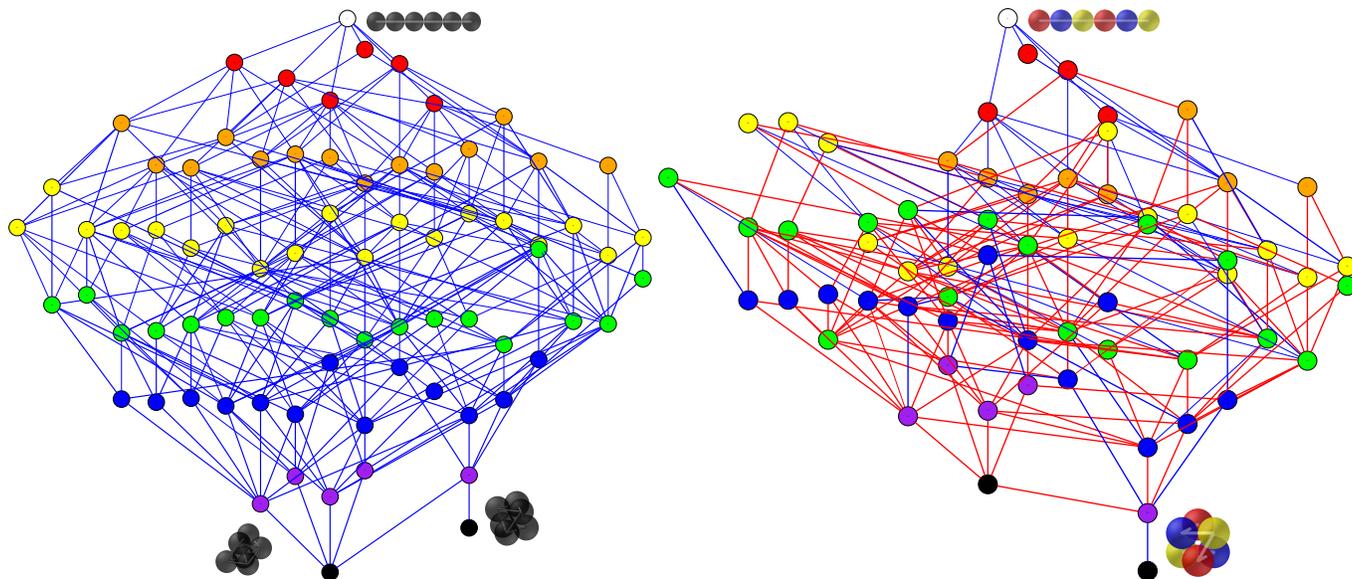


Fig. 5 Free energy diagrams for the self-assembly of 6 spheres with identical particles with interaction energy $8kT$ (left) and three particle types with Pareto optimal interactions for an octahedron (right). The optimal interactions lead to a diagram that appears qualitatively to be “funnelled.” The vertical positions of the nodes are proportional to the state’s free energy. Node coloring denotes the number of non-trivial bonds; the chain in white, one additional bond for each color of the rainbow, and black for ground states. Edge width is proportional to the absolute net transition rate, $|Q_{ij} - Q_{ji}|$, where Q is the transition matrix. Edge color denotes the preferred transition direction: blue for forming a bond, and red for breaking a bond.

as a design parameter, for various numbers of types m . With 2 types, the octahedron’s optimal labeling is the AABAAB that we have been considering and its Pareto front has a shallow slope. With 3 types, however, the Pareto front is a *single point* – equilibrium probability and rate are maximized simultaneously. There is no thermodynamic-kinetic tradeoff – we have found a system with a funnelled energy landscape. To make this analogy with funnelled landscapes more explicit, Figure 5 shows a free energy diagram over the coarse-grained states. We see that for identical particles, the landscape is entirely downhill, with most probability flowing into the global free-energy minimum, the polytetrahedron. For three particle types with the Pareto optimal parameters, all probability flows into the octahedron, which now lies at the bottom of the funnel, while probability only flows out of the polytetrahedron; it is no longer a kinetic trap.

To check if this optimization works outside the assumptions of our coarse-grained model, we performed Brownian dynamics simulations of chains of 6 spheres using the optimal ordering ABCABC and optimal bond energies found by the algorithm. Figure 2(e) shows the octahedron forms with high probability in just over 1 minute. This result is comparable to our earlier Brownian dynamics simulation with distinct particles, showing that we can achieve the same assembly efficiency using three types of particles as with six types.

4.2 Polymers of disks

Figure 6 shows the Pareto fronts for all rigid clusters formed from chains of 6 & 7 discs, optimized over particle labels for a variety of types m . There are many things to observe about these curves.

They all have a critical value $m=m_$ for which the Pareto front becomes essentially vertical.* This is perhaps our most striking observation. In Figure 6, $m_*=2$ for (b,c,d,e) and $m_*=3$ for (a,f,g). Although the Pareto front does not always collapse to a single point, as it did for the octahedron, when it is very steep the trade-off between thermodynamic stability and kinetic accessibility has been eliminated: a target state may have equilibrium probability close to 1, and still fold rapidly. We suggest that a steep Pareto front is a signature of a funnelled energy landscape.

*Some clusters can fold faster when the number of types is increased beyond m_** For example, it appears that the chevron can fold slightly faster with $m=3$, even though it has $m_*=2$. We confirmed this observation by Brownian dynamics simulation, where we observe a 20% reduction in the mean first passage time by using $m=3$ instead of $m=2$ (See Figure 3(a) in the SI Appendix[†]).

Pareto fronts can be discontinuous. For example, see Figure 6(g), $m=2, 3$. Each continuous fragment of the Pareto front is associated with different optimal particle labels.

Shallow-sloped Pareto fronts have medium-weak bonds where folding is fastest, and strong bonds where free energy is lowest. For $m=2$ types, shallow Pareto fronts occur in Figure 6(a,f,g) and for the octahedron, Figure 4(a). For all of these, as π_S increases, the strongest bond remains constant while a weaker bond increases in strength, from a small (but not the smallest) value, to close to the maximum value. This case bears comparing with proteins, which are known to be marginally stable, i.e. with energies just low enough to be stable but not as low as could be possible⁴⁷. The reasons for proteins’ marginal stability are not fully known; one speculation is that it aids proteins’ functionality, by giving them some flexibility and also allowing them to be easily taken apart; another is that it is a result of neutral evolution, which

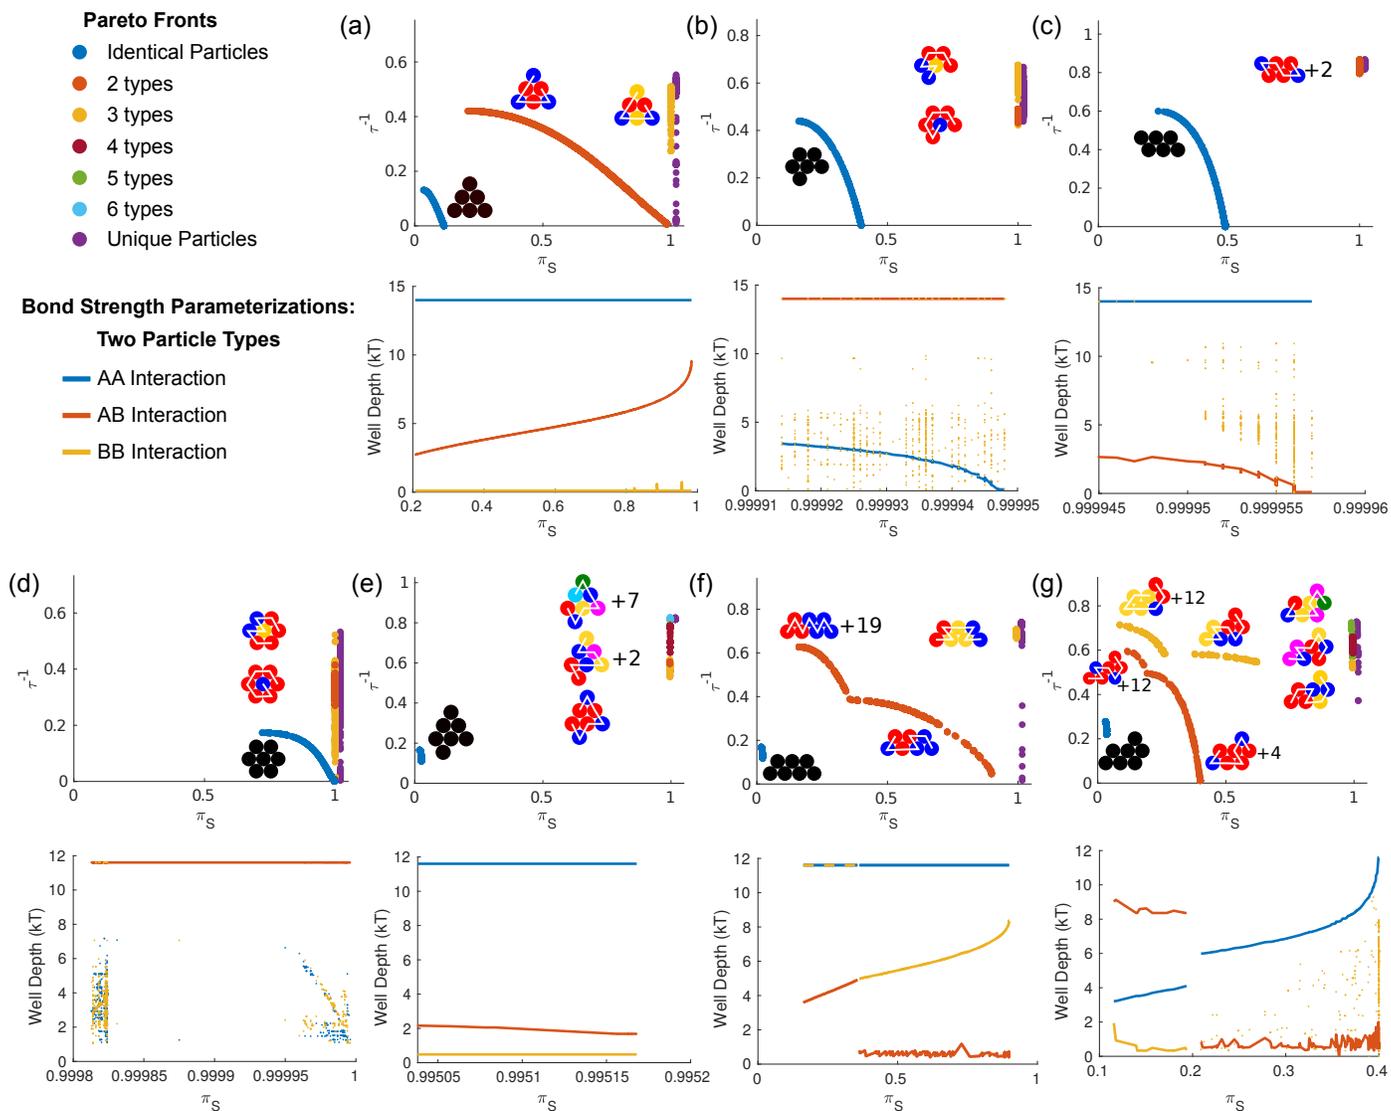


Fig. 6 Pareto fronts for all rigid clusters of 6 and 7 disks, with varying number of particle types m . The target states include all permutations of a given cluster. For each target states we calculate the Pareto front for distinct particles (purple) to determine an upper bound on τ^{-1} . Then, we increase m until the assembly is as efficient as having distinct particles. Pareto fronts for distinct particles are shifted rightwards so they do not overlap other Pareto fronts. Cluster plots show which permutation generate the Pareto front. In some cases, several permutations with approximately equal probability generate the front, indicated by '+c', where $c+1$ is the total number of permutations. Below each Pareto front there is a parameterization of the 2-type Pareto front in terms of the three bond energies. Dotted data indicates the Pareto front is insensitive to the given interaction, for example a BB interaction in a system with only one B -type particle.

favours native structures that can be folded by the most amino acid sequences⁴⁷. Our results suggest that weaker bonds could also help proteins fold more quickly, by minimizing the effect of kinetic traps.

Steeply-sloped Pareto fronts favor (slightly) weaker bonds where free energy is lowest. For $m=2$ types, steep Pareto fronts occur in Figure 6(b,c,d,e) and for the polytetrahedron, Figure 4(b). For most of these examples, the strong bond is still constant but the weaker bond decreases slightly as π_S increases. We suggest three possible explanations. The first is an effect of permutation lumping. When computing the mfpt to a target, all permutations are considered, even if some have vanishingly small equilibrium probability due to the presence of a weak bond. Since the mfpt cal-

culatation only considers hitting a target state, and not staying in it, including these unlikely permutations can slightly enhance the rate. We tested this hypothesis using the chevron (Figure 6(b)) as an example. If we restrict the target state to the 2-type permutation shown in the figure, the Pareto front reduces to a single point, supporting the hypothesis. However, the same is not true for the triangle (Figure 6(a)); we still find a near-vertical front when restricting the target to the three particle type permutation shown in the figure. A second possibility is that adding auxiliary weak bonds can increase the rate of forming a target, even if that bond is not present in the target state. We could not find numerical evidence for this possibility (SI Appendix[†], Section 6). A third is that this observation is a numerical artifact, since we compute

only an approximation to the true Pareto front.

4.3 Lattice polymer

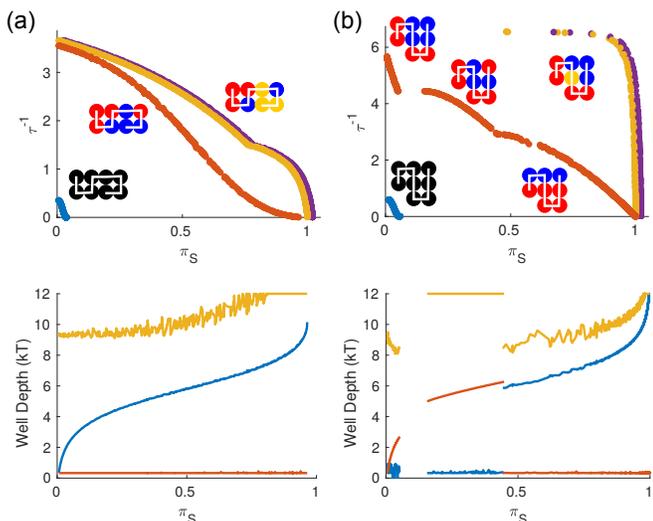


Fig. 7 Pareto fronts for two maximally bonded configurations of a lattice polymer. The information plotted is the same as in Figure 6.

As a final example we compute Pareto fronts for lattice polymers formed from a chain of 8 particles (SI Appendix[†], Section 7). Figure 7 shows Pareto fronts for two different target states, each with a fixed permutation. For both target structures, the best folding is achieved at $m=3$, where the Pareto front is indistinguishable from the Pareto front for $m=8$. Many of the qualitative features of these fronts are the same as for disks, but a notable exception is the Pareto front for Figure 7(a) does not become nearly vertical, even for distinct particles (although the slope near $\pi_S=1$ approaches ∞ .) We believe the absence of a nearly-vertical Pareto front is because our lattice model allows two bonds to form at once, so it captures some of the kinetic traps that get averaged out by our coarse-grained continuum model, an issue we discuss later.

4.4 Pareto Front Sensitivity

How sensitive are the Pareto fronts to the underlying assumptions of the model, such as the initial condition, estimation of free energies, model for transition rates, choice of pair potential and bond cutoff, neglect of hydrodynamic interactions, etc? We computed Pareto fronts for a small collection of examples under two types of perturbations: one which alters the transition rates, and another which considers different initial conditions (SI Appendix, Section 5.) Figure 8 shows that while the numerical values of the rates along the Pareto fronts do change, their qualitative shape does not, and, remarkably, the bond parameterizations along the Pareto fronts appear insensitive to any of these perturbations. These findings suggest that the Pareto optimal interactions are mainly a function of the system's connectivity network, i.e. its topology, rather than its detailed kinetics.

Such a finding implies that strategies that attempt to minimize the number of accessible pathways to mis-folded states, such as

was done for a minimal, 3-level toy model of self-assembly⁹, may successfully avoid kinetic traps, without detailed knowledge of the folding kinetics.

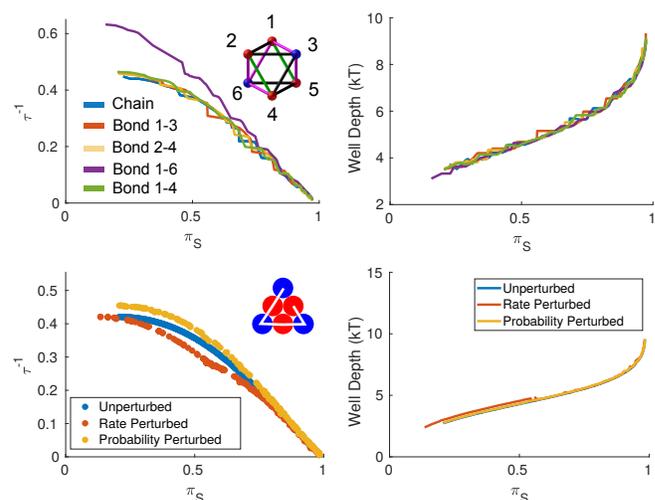


Fig. 8 Top Row: Pareto fronts and non-trivial bond energy parameterizations (AA interaction) for the octahedron state with two particle types, for various 1-bond starting states. With the exception of the linear chain, all starting states are sampled from the equilibrium measure in that state. Starting with a mis-folded bond has nearly no effect, as the bond can quickly break, while starting in a loop increases the maximum rate, as this bond is the slowest in the assembly process. Bottom Row: Pareto fronts and parameterizations (AB interaction) for the two particle type triangle after perturbing model parameters. Exit rates out of coarse grained states are perturbed by adding Gaussian noise with standard deviation 20% of each state's rate ($\sigma = 0.2/\tau_i$.) Probabilities are perturbed by making each transition out of a state equally likely.

5 Chiral traps: the problem with short-ranged isotropic interactions

We encountered a problem when we tested whether the design parameters found by our genetic algorithm, lead to fast folding in Brownian dynamics simulations. Sometimes the Brownian dynamics simulations agreed with our predictions of fast folding, but sometimes the simulations got stuck in kinetic traps. For example, Figure 9 shows Brownian dynamics simulations of a triangle with 3 particle types and interactions found by the algorithm to have no tradeoff. The yield stagnates at about 0.5; it will eventually reach 1 but only over a much longer timescale. The other 50% of the trajectories are stuck in a floppy state, which requires breaking a strong bond to leave. This floppy state has the same adjacency matrix as a floppy state which is on the folding pathway to the triangle. Our coarse-grained model does not distinguish clusters with the same adjacency matrix: it lumps the kinetic trap and the on-pathway cluster into the same state, effectively allowing the system to tunnel between the two configurations, and thereby removing the energy barrier for leaving the trap. This problem occurs even for distinct particles: the kinetic trap always exists for the original, continuous system, even if it is removed in our coarse-grained model.

This problem also occurs for the octahedron, although more

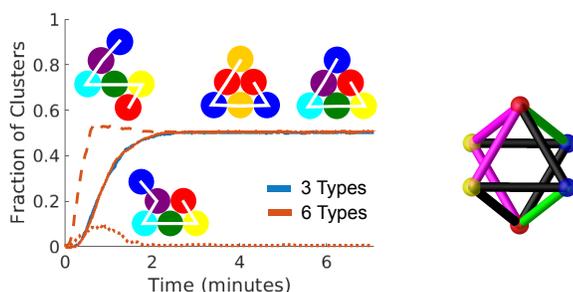


Fig. 9 Chiral traps prevent efficient assembly. Left: Brownian dynamics simulations of polymers of 6 disks, with 3 types (ABCBC) and 6 types. The yield stagnates at about 0.5 in both cases, because half the trajectories get stuck in a kinetic trap (dashed line), which requires breaking a strong bond to form the triangle. The kinetic trap has the same adjacency matrix as an on-pathway cluster (dotted line), so it is impossible to remove this trap using short-ranged isotropic interactions. All yields were calculated from 400 trajectories initialized as a linear chain. The three-type chain had energies $E_{AA}=E_{AC}=12kT$, while the 6-type chain set all target energies to $12kT$; all other interactions were $0.1kT$. Right: The floppy kinetic trap for the octahedron, with the strong AB and AC bonds in green and pink respectively, and backbone bonds in black. This cluster has the same adjacency matrix as a cluster where the frontmost yellow and blue particles are interchanged, which can then form strong blue-yellow bonds with their neighbours.

mildly; about 1-2% of the trajectories got stuck in a floppy trap with the same adjacency matrix as an on-pathway cluster, shown in Figure 9. This same floppy trap has been observed experimentally in DNA coated colloids assembling from a gas, with interaction energies similar to the ones we identified as optimal⁴⁸.

We call such kinetic traps *chiral traps*, because there is no way to distinguish them from on-pathway clusters given only the matrix of particle contacts. Chiral traps pose a fundamental problem for self-assembly or folding using short-ranged isotropic interactions. There is no way to remove these traps using interactions that do not depend on particle orientation, even if particles are all distinct. Indeed, such traps have been observed in simulations of dozens of distinct particles²².

Yet, all is not lost – the success of our coarse-grained model in finding vertical Pareto fronts, suggests that were the chiral traps to be removed, or distinguished from their on-pathway counterparts, one could achieve fast, reliable assembly. One way to do this is with particles with patchy interactions^{49–52}. It could be that patchy particles have qualitatively similar Pareto fronts as with our coarse-grained model. This idea is supported by computational studies that observe high fidelity yields of icosahedra using patchy particles⁸; yield reductions were due primarily to aggregation effects, which are not present in our system, instead of mis-folding effects. We also remark that the coarse-grained model is successful within the context of its own assumptions. Figure 9 shows that simulating a chain of distinct particles, gives the same yield curve as a chain with 3 particle types. Therefore, the model can identify when assembly is as efficient with m particle types as it is with distinct particles.

6 Extension to Larger Systems

So far we have computed Pareto fronts for systems whose state space can be fully enumerated, and we furthermore lumped states sharing the same adjacency matrices so that states can be easily identified⁵³. This allowed us to efficiently evaluate measures of thermodynamic stability and kinetic accessibility, but it will not be feasible for larger systems whose state spaces cannot be fully enumerated.

A straightforward idea for handling non-enumerable systems would be to simulate them, and estimate the objectives from an ensemble of trajectories. We do not expect this to be feasible – first, because computing a Pareto front requires evaluating the objectives thousands of times, and second because many of the parameter values of interest lead to extremely slow rates of folding, which also imply very long simulation times. However, there is flexibility in choosing the objectives; other functions might be just as (or more) informative, yet possible to evaluate cheaply and accurately by simulation.

We have tentatively explored a selection of other objectives. One idea that appears promising, is to replace π_S with p_S , an average energy barrier for leaving the state, and τ^{-1} with k_A , a measure of the energy of mis-folded bonds when a trajectory first becomes trapped (see SI Appendix[†], Section 8 for details). These objectives may be estimated reasonably efficiently from an ensemble of trajectories. Figure 10(a) shows that for a lattice polymer of size $N = 8$, the Pareto front computed for these new objectives, gives bond energies that lie mostly on the Pareto front for the original objectives. This shows the parameterization of the Pareto front is relatively unchanged by this change of objectives, a promising result. We then ran our genetic algorithm for a lattice polymer of size $N = 16$, a system big enough that enumeration is impossible, and obtained converged Pareto optimal parameters (SI Appendix[†], Section 8.) However, for both lattice polymers, the Pareto optimal parameters do not give a good yield of the target structure: Figure 10(b) shows that the finite time target structure yield is significantly eroded by the presence of chiral traps. Since chiral traps contain no mis-folded bonds, they do not contribute to k_A , and thus are undetected by our sampling. By inspection, we see these trap states could also be eliminated by introducing directional bonds.

7 Discussion

We computed Pareto fronts describing the tradeoff between equilibrium probability and rate of folding, for small polymers of spheres and disks with programmable interaction energies. We showed that a genetic algorithm can compute these Pareto fronts, and can handle the discrete optimization problem of determining particle labels. This approach makes no assumptions on the intrinsic timescales of the assembly (in contrast to approaches which optimize yield at a fixed final time), allowing the assembly timescales to emerge naturally through our optimization algorithm, and, it can identify collections of optimal parameters in a high-dimensional parameter space.

A key approximation that let us efficiently evaluate the equilibrium probability and rate of folding, was to first construct

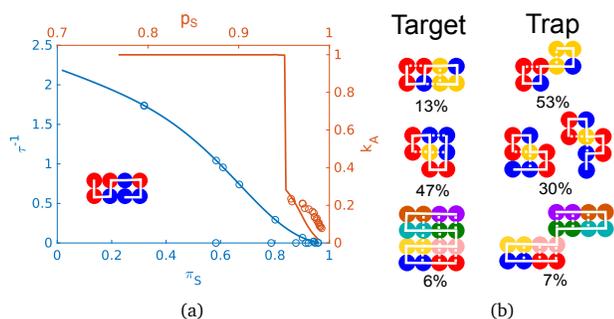


Fig. 10 (a) Mapping of the Pareto front for the $N = 8$ particle rectangle state on a lattice in the original objectives (π_S, τ^{-1}) (solid blue curve) to the new objectives (p_S, k_A) (solid red curve), computed with the coarse-grained model. The genetic algorithm was then applied to the new objectives by sampling, using $P = 100$, $p = 0.25$, $T = 500$, $t_{\text{trap}} = 300$ (see SI Appendix[†], Section 8 for details), 200 generations, and 500 samples to estimate averages. The non-dominated portion of the population is shown as unfilled red circles. The parameters for the red circles were extracted, and plugged into the coarse-grained model to evaluate the original objectives (blue circles), which can be seen to still lay mostly on the original Pareto front. (b) Yields of the target state and most common kinetic trap, estimating using 5000 samples, for the two $N = 8$ ground states at $T = 5000$, as well as a square of $N = 16$ particles at $T = 50000$, using the Pareto optimal parameters from sampling. Solid white lines denote the backbone, and dashed white lines denote bonds with $E = 12kT$. Note that the contact matrix for the trapped states are consistent with the target state, showing the kinetic traps are chiral traps.

a coarse-grained model which approximates the dynamics as a Markov chain on the set of adjacency matrices. The main limitation of this approximation is that it removes some kinetic traps for isotropic particles that sometimes prevent a real system from assembling efficiently. We believe other possible limitations are negligible – indeed, we found the design parameters on the Pareto front are remarkably insensitive even to large perturbations.

We exhaustively calculated the Pareto fronts for the ground states of small polymers with every possible number of types of particles, for this coarse-grained dynamical model. Apart from being useful as a design tools, these Pareto fronts give qualitative insight into what determines the thermodynamic-kinetic trade-off. A striking observation was that for nearly all target states, there was a critical number m_* of particle types, less than the total number of particles, that led to a vertical or point-like Pareto front. When this happens there is no tradeoff; one may choose interactions such that the polymer folds rapidly to a low-energy target state.

Our philosophical approach to calculating Pareto fronts is not limited to polymers of spherical particles, but could be adapted to handle particles assembling from a gas, particles with orientation-dependent interactions, particles with different shapes, etc; additional objectives could also be considered, resulting in Pareto fronts that are hypersurfaces with dimension one less than the number of objectives. Our technical approach to coarse-graining the configuration space, however, is limited to systems small enough that one can enumerate all the possible ways to put particles in contact. The main challenge in extending our results will be to efficiently evaluate measures of thermodynamic stability

and kinetic accessibility, in systems where this enumeration is no longer feasible. We tinkered with other objectives, which seemed promising when evaluated for small lattice polymers, however we were discouraged by the increasing number of chiral traps that occur for larger lattice polymers, which prevent the system from ever folding efficiently. Any further work in this direction must consider non-isotropic interactions.

We were motivated by the puzzling success of protein folding, and suggest that a steep Pareto front is a signature of a “funnelled” energy landscape. This funnel metaphor has proven invaluable in studying how proteins navigate their complex energy landscape to efficiently fold into their native states, but it has not yet found use as a practical design tool in materials science. Framing the design problem in the language of multi-objective optimization, as we have done, and building upon existing tools to address the tradeoffs between objectives, may provide an alternative method to design complex self-assembling systems given realistic experimental constraints. A fascinating question would be to understand how the critical number m_* of particle types grows with system size or target structure complexity. Information theory^{54,55} combined with statistical mechanics considerations may give additional insight into this question. Will this number always be less than 20, as for proteins?

8 Conflicts of Interest

There are no conflicts of interest to declare.

9 Acknowledgements

The authors are grateful to Eric Vanden-Eijnden for introducing them to this problem. This work was supported by the Research Training Group in Modeling and Simulation funded by the National Science Foundation via grant RTG/DMS – 1646339, and by the United States Department of Energy under Award no. DE-SC0012296. M. H.-C. acknowledges support from the Alfred P. Sloan Foundation.

10 Data Availability Statement

The raw data and scripts to make figures, as well as all the code to generate our data is available on Github⁵⁶.

Notes and references

- 1 G. A. Garg T, Rath G, *Critical Reviews in Therapeutic Drug Carrier Systems*, 2015, **32(2)**, 89–147.
- 2 M. Beija, R. Salvayre, N. Lauth-de Viguerie and J.-D. Marty, *Trends in Biotechnology*, 2012, **30**, 485–496.
- 3 J.-H. Huh, K. Kim, E. Im, J. Lee, Y. Cho and S. Lee, *Advanced Materials*, 2020, **32**, 2001806.
- 4 J. A. Fan, Y. He, K. Bao, C. Wu, J. Bao, N. B. Schade, V. N. Manoharan, G. Shvets, P. Nordlander, D. R. Liu and F. Capasso, *Nano Letters*, 2011, **11**, 4859–4864.
- 5 S. Ebbens, *Current Opinion in Colloid & Interface Science*, 2016, **21**, 14–23.
- 6 S. A. Mallory, C. Valeriani and A. Cacciuto, *Annual Review of Physical Chemistry*, 2018, **69**, 59–79.

- 7 M. F. Hagan and D. Chandler, *Biophysical Journal*, 2006, **1**, 42–54.
- 8 A. W. Wilber, J. P. K. Doye, A. A. Louis, E. G. Noya, M. A. Miller and P. Wong, *The Journal of Chemical Physics*, 2007, **127**, 085106.
- 9 J. Grant, R. L. Jack and S. Whitelam, *The Journal of Chemical Physics*, 2011, **135**, 214505.
- 10 C.-A. Palma, M. Cecchini and P. Samorì, *Chemical Society Reviews*, 2012, **41**, 3713–3730.
- 11 S. Whitelam and R. L. Jack, *Annual Review of Physical Chemistry*, 2015, **66**, 143–163.
- 12 G. Bisker and J. L. England, *Proceedings of the National Academy of Sciences*, 2018, **115**, E10531–E10538.
- 13 J. D. Bryngelson, J. N. Onuchic, N. D. Socci and P. G. Wolynes, *Proteins-Structure Function and Genetics*, 1995, **21**, 167–195.
- 14 H. S. Chan and K. A. Dill, *The Journal of Chemical Physics*, 1994.
- 15 H. Li, C. Tang and N. S. Wingreen, *Proceedings of the National Academy of Sciences*, 1998, **95**, 4987–4990.
- 16 C. Levinthal, *Mössbauer spectroscopy in biological systems*, 1969, **67**, 22–24.
- 17 R. Zwanzig, A. Szabo and B. Bagchi, *Proceedings of the National Academy of Sciences*, 1992, **89**, 20–22.
- 18 J. N. Onuchic, *Proceedings of the National Academy of Sciences*, 1997, **94**, 7129–7131.
- 19 K. Dill and H. Chan, *Nature Structural Biology*, 1997, **4**, 10–9.
- 20 G. R. Bowman and V. S. Pande, *Proceedings of the National Academy of Science*, 2010.
- 21 S. Hormoz and M. P. Brenner, *Proceedings of the National Academy of Sciences of the United States of America*, 2011, **108**, 5193–8.
- 22 Z. Zeravcic, V. N. Manoharan and M. P. Brenner, *Proceedings of the National Academy of Sciences*, 2014, **111**, 15918–15923.
- 23 L. O. Hedges, R. V. Mannige and S. Whitelam, *Soft Matter*, 2014, **10**, 6404–6416.
- 24 M. Z. Miskin, G. Khaira, J. J. de Pablo and H. M. Jaeger, *Proceedings of the National Academy of Sciences*, 2016, **113**, 34–39.
- 25 J. Madge and M. A. Miller, 2017, 1–13.
- 26 A. Das and D. T. Limmer, *The Journal of Chemical Physics*, 2021, 1–13.
- 27 Y. Ke, L. L. Ong, W. M. Shih and P. Yin, *Science*, 2012, **338**, 1177–1183.
- 28 A. Murugan, J. Zou and M. Brenner, *Nature Communications*, 2015, **6**, 6203.
- 29 W. M. Jacobs, A. Reinhardt and D. Frenkel, *Proceedings of the National Academy of Sciences*, 2015, **112**, 6313–6318.
- 30 Z. Yan and J. Wang, *Proceedings of the National Academy of Sciences*, 2020, **117**, 27218–27223.
- 31 A. W. Long and A. L. Ferguson, *The Journal of Physical Chemistry B*, 2014, **118**, 4228–4244.
- 32 S. Whitelam and I. Tamblin, *Phys. Rev. E*, 2020, **101**, 052604.
- 33 C. P. Goodrich, E. M. King, S. S. Schoenholz, E. D. Cubuk and M. P. Brenner, *Proceedings of the National Academy of Sciences*, 2021, **118**,.
- 34 R. J. Macfarlane, B. Lee, M. R. Jones, N. Harris, G. C. Schatz and C. A. Mirkin, *Science*, 2011, **334**, 204–208.
- 35 W. Rogers, W. Shih and V. Manoharan, *Nature Reviews Materials*, 2016.
- 36 Y. Wang, Y. Wang and X. e. a. Zheng, *Nature Communications*, 2015.
- 37 Y. Wang, I. C. Jenkins, J. T. McGinley, T. Sinno and J. C. Crocker, *Nature Communications*, 2017, **8**, 1–8.
- 38 A. McMullen, M. Holmes-Cerfon, F. Sciortino, A. Y. Grosberg and J. Brujic, *Physical Review Letters*, 2018, **121**, 138002.
- 39 G. Meng, N. Arkus, M. P. Brenner and V. N. Manoharan, *Science*, 2010, **327**, 560–563.
- 40 R. W. Perry, M. C. Holmes-Cerfon, M. P. Brenner and V. N. Manoharan, *Phys. Rev. Lett.*, 2015, **114**, 228301.
- 41 Z. Zeravcic, V. N. Manoharan and M. P. Brenner, *Rev. Mod. Phys.*, 2017, **89**, 031001.
- 42 T. Stewart, O. Bandte, H. Braun, N. Chakraborti, M. Ehr Gott, M. Göbelt, Y. Jin, H. Nakayama, S. Poles and D. Di Stefano, in *Real-World Applications of Multiobjective Optimization*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2008, pp. 285–327.
- 43 E. D. Klein, *PhD thesis*, Harvard, 2019.
- 44 S. Barkley, A. Neophytou, A. Trubiano, E. D. Klein, D. Chakraborti, M. Holmes-Cerfon and V. N. Manoharan, *In preparation*, 2021.
- 45 M. H. Huntley, A. Murugan and M. P. Brenner, *Proceedings of the National Academy of Sciences*, 2016, **113**, 5841–5846.
- 46 M. Holmes-Cerfon, *Journal of Chemical Physics*, 2020, **153**, 164112.
- 47 D. M. Taverna and R. A. Goldstein, *Proteins: Structure, Function, and Genetics*, 2002, **46**, 105–109.
- 48 J. W. Collins, *PhD thesis*, 2014.
- 49 Z. Zhang and S. C. Glotzer, *Nano letters*, 2004, **4**, 1407–1413.
- 50 F. Romano, E. Sanz and F. Sciortino, *The Journal of Chemical Physics*, 2010, **132**, 184501.
- 51 Q. Chen, S. C. Bae and S. Granick, *Nature*, 2011, **469**, 381–384.
- 52 Y. Wang, Y. Wang and D. e. a. Breed, *Nature*, 2012.
- 53 E. Zappa and M. Holmes-Cerfon, *Journal of Nonlinear Science*, 2019, **29**, 2021–2053.
- 54 D. Soloveichik and E. Winfree, *SIAM Journal on Computing*, 2007, **36**, 1544–1569.
- 55 E. D. Demaine, S. Eisenstat, M. Ishaque and A. Winslow, *International Workshop on DNA-Based Computers*, 2011, pp. 100–114.
- 56 A. Trubiano, *CPfold*, <https://github.com/onehalfatsquared/CPfold>, 2021.