Using Markov state models to study self-assembly

Matthew R. Perkett and Michael F. Hagan^{a)}

Martin Fisher School of Physics, Brandeis University, Waltham, Massachusetts 02474, USA

(Received 7 February 2014; accepted 30 April 2014; published online 2 June 2014)

Markov state models (MSMs) have been demonstrated to be a powerful method for computationally studying intramolecular processes such as protein folding and macromolecular conformational changes. In this article, we present a new approach to construct MSMs that is applicable to modeling a broad class of multi-molecular assembly reactions. Distinct structures formed during assembly are distinguished by their undirected graphs, which are defined by strong subunit interactions. Spatial inhomogeneities of free subunits are accounted for using a recently developed Gaussian-based signature. Simplifications to this state identification are also investigated. The feasibility of this approach is demonstrated on two different coarse-grained models for virus self-assembly. We find good agreement between the dynamics predicted by the MSMs and long, unbiased simulations, and that the MSMs can reduce overall simulation time by orders of magnitude. © 2014 AIP Publishing LLC. [<http://dx.doi.org/10.1063/1.4878494>]

I. INTRODUCTION

The assembly of basic units into structures with increased size and complexity is central to biology, where examples of assembled structures include viruses (e.g., Refs. 1–4), cell membranes, cytoskeletal filaments,⁵ and ordered layers of proteins on bacterial surfaces.⁶ Assembly is also increasingly important to nanoscience, where interactions between colloidal particles are being engineered to drive assembly into sophisticated, functional materials (e.g., Refs. 7–15) and DNA origami promises the ability to build structures of nearly limitless complexity (e.g., Refs. 16–18). An important focus of current research is understanding how the interactions between individual components determine assembly pathways, timescales, and fidelity for a target structure. Computational modeling can play a key role in determining assembly pathways and mechanisms, since most intermediates are transient and thus not readily characterized in experiments. However, simulating assembly is challenging because target structures can be orders of magnitude larger than their constituent components and assembly pathways typically surmount large free energy barriers, leading to timescales which greatly exceed computational limitations.

This paper is concerned with using Markov State Models (MSMs) to overcome the gap between assembly times and computationally accessible timescales. Many powerful enhanced sampling techniques have been developed to efficiently harvest computational trajectories that include barrier crossings or other rare events (e.g., Refs. 19–34). However, many of these methods are limited by requiring *a priori* knowledge of reaction coordinates, one or few pathways to completion, or only several metastable minima. In contrast, MSMs can be used to study assembly reactions characterized by multiple free energy barriers, a diverse ensemble of pathways, and pleomorphic products. Furthermore, MSMs are one

of only a few methods³⁵ that can describe non-stationary, out-of-equilibrium dynamical processes.

Previous works on assembly of capsids or other structures have postulated Markov state models in which the state space and transition rates are pre-assumed based on physical considerations.^{36–45} In the approach described here, the state space and transition rates emerge from particle-based dynamics simulations, and the validity of the Markov assumption is explicitly tested against the microscopic dynamics.

While such MSMs have been extensively developed in the context of protein folding,^{46–60} existing approaches cannot describe the assembly of disconnected, permutable subunits. Here, we present a method to construct MSMs that is applicable to a wide variety of such assembly reactions. We test our approach on two models, which respectively describe the assembly of viral proteins around rigid nanoparticles and flexible polymers. While straightforward dynamics simulations with similar models have led to important insights about capsid assembly (e.g., Refs. 61–76, reviewed in Refs. 4 and 77), these investigations have been limited to parameters for which nucleation barriers are small. We show that the MSM approach enables simulation over a much wider range of experimentally relevant parameter values.

II. METHODS

A. Existing implementations of MSMs

In this section, we review how relatively short, unbiased simulations can be used to build an MSM and how this technique has been applied to study protein folding or conformational transitions. The procedure begins by partitioning configurations from the short simulations into states such that conformations which interconvert rapidly are collected into the same state. The separation of timescales resulting from this partitioning ensures that the model is Markovian on timescales longer than a “lag time” τ , meaning that the probability of transitioning to a new state only depends upon the

^{a)}Electronic mail: hagan@brandeis.edu

current state. Taking $\vec{P}(0)$ to be the vector of probabilities of being in each of the possible states of the system at time $t = 0$, the state probability at time t_f is given by $\vec{P}(t_f) = \mathbf{T}(\tau)^n \vec{P}(0)$ where $n = t_f/\tau$ and $\mathbf{T}(\tau)$ is the stochastic matrix of interstate transition probabilities estimated from the simulations at lag time τ .

Determining a state decomposition that achieves the separation of timescales described above is a crucial aspect of building an MSM. If the states do not sufficiently distinguish values of all of the slow degrees of freedom in the system, then the lag time τ at which the system becomes Markovian will be comparable to its longest relaxation timescale. Since simulations must be greater than τ in length, the “short” simulations will approach the length of long, unbiased trajectories and thus the method will offer no computational savings.

Several approaches to determining state decompositions have been developed in the context of all-atom protein simulations.^{46–60,78–93} In cases where a set of collective coordinates describing all of the slow degrees of freedom is known (or guessed) *a priori*, biased sampling can be used to determine the free energy landscape as a function of these coordinates. States can then be defined based on local free energy minima (e.g., Refs. 49 and 78–81). Since it is rare to have *a priori* knowledge of good collective coordinates, alternative approaches have been developed in which configurations are clustered based on geometric criteria, such as structural similarity (e.g., Refs. 82–85). Chodera *et al.*⁵⁹ developed an algorithm to refine an initial geometric partitioning of “microstates” into “macrostates” based on kinetics. Open source software packages such as MSMBUILDER^{86,87} and EMMA⁹⁴ provide a suite of tools for building MSMs in this fashion and analyzing them. This algorithm and similar approaches have been shown to be extremely powerful for the study of proteins, allowing for prediction of folding pathways and rates on even the supra-millisecond timescale (e.g., Refs. 54, and 87–90) as well as identifying hidden allosteric sites.⁹¹ Recently, systematic approaches to find optimal coordinates for constructing MSMs have been developed.^{92,93}

B. Building MSMs for assembly systems

In contrast to protein folding, where each residue in the protein has a unique index, assembly subunits are permutable and thus cannot be indexed. Therefore, existing algorithms for determining state decompositions cannot be directly applied to self-assembly. Here, we describe several approaches to creating a state decomposition based on the network of subunit interactions and their positions relative to heterogeneous nucleation sites.

We consider systems in which subunits can assemble either through homogeneous nucleation to form a single component structure (e.g., an empty virus capsid shell) or through heterogeneous nucleation around a scaffold (e.g., a polymer or a nanoparticle) to form a multicomponent structure. To simplify the presentation, we consider one type of subunit and focus on only the largest assembled structure in the system at any given time, but the approach can be generalized to multiple subunit species and assemblies. Our approach can be

applied to systems which assemble via reversible or irreversible interactions; in both cases, we will describe a strong interaction as a “bond.”

To generate a state decomposition, we categorize subunits into three classes: class I: bonded subunits in the assemblage, class II: subunits bonded to the scaffold, but not the assemblage, and class III: unbonded, free subunits. In the case of homogeneous nucleation, there are no class II subunits (since there is no scaffold), and class III subunits can be ignored if subunit association to the scaffold is sufficiently reaction-rate limited that the density of free subunits is spatially uniform.

For systems that assemble into well-defined structures, fluctuations in bond distances and angles are fast compared to bond formation and breakage. By averaging over these short timescale fluctuations, the unique structure of a growing cluster can be defined by the class I subunit bonding network. More precisely, each cluster is converted into an undirected graph with nodes corresponding to subunits and edges corresponding to bonds between the subunits (see Fig. 11 in Appendix B). This ensures a consistent state definition that is unaffected by exchanges of subunits with different indices, short-timescale conformational fluctuations, or rigid body motions of the assemblage. Class I subunits can be further sub-partitioned by including the distance from the scaffold to each subunit, but this was unnecessary for the systems that we considered. Class II subunits can be handled in a similar manner by considering the subunit-scaffold bonding network, but we found that it was only necessary to track the total number of class II subunits.

Class III subunits must be included in the state decomposition when their association to the cluster or scaffold approaches the diffusion limit, which results in density inhomogeneities. Free subunit positions fluctuate rapidly since they are not involved in strong interactions, and thus we use their density distribution rather than their positions to decompose states. We follow the approach developed by Gu *et al.*⁹⁵ to include solvent degrees of freedom in protein folding MSMs. We define a vector $\mathbf{w}$, in which each index corresponds to a subunit in the cluster or a residue of the scaffold and each component corresponds to a distance-weighted density of free subunits around the indexed subunit,

$$w_i \equiv \sum_{j \in \text{free subunits}} e^{-|\vec{r}_j - \vec{r}_i|^2 / \sigma_d} \quad (1)$$

with $\vec{r}_i$ and $\vec{r}_j$ as the positions of cluster/scaffold subunit i and the free subunit j . σ_d is an adjustable decay length that sets the scale for relevant interactions between free subunits and the scaffold or growing cluster. This definition weighs nearby subunits more heavily since they are more likely to associate with the cluster or scaffold. Whether cluster subunits, scaffold residues, or both need to be considered in this definition depends on which association reactions approach the diffusion limit. For example, in the case considered below where subunit adsorption onto the nanoparticle approaches the diffusion limit but subunit-subunit associations do not, it is only necessary to include the nanoparticle in $\mathbf{w}$.

Finally, if scaffold internal degrees of freedom (e.g., conformations of a polymer) evolve slowly, these should also be

included in the state definition. Because the scaffold residues (i.e., polymer segments) are indexable, these degrees of freedom can be treated via the existing RMSD-based approach (e.g., in MSMBuilder^{86,87} and EMMA⁹⁴).

1. Reducing the number of states

As the size of the target structure increases, the number of distinct assembly intermediates and hence the number of unique graphs grows rapidly. The number of states could become intractable if class II or class III subunits were included in the state definition. There are two routes to reduce the number of states. First, kinetic data can be used to group microstates that interconvert rapidly into macrostates by following the approach based on Perron cluster analysis in MSMBuilder^{60,86,87} and EMMA.⁹⁴ Second, *a priori* knowledge of the system can be used to reduce the number of unique states. Since assembly into a target structure with high fidelity generally requires weak subunit-subunit interactions,^{4,96,97} subunits in a cluster with only one bond rapidly dissociate. Thus, the edges corresponding to these interactions can generally be neglected when building graphs. For the models considered here, we found that fluctuations about the most compact, highly bonded structures are rapid enough that it was sufficient to consider only the number of subunits in a cluster. Several alternative simplified descriptions are discussed in Appendix B.

C. Generating the transition matrix

The transition probability matrix $\mathbf{T}(\tau)$ is calculated by column-normalizing the count matrix $\mathbf{C}(\tau)$, in which each element C_{ji} gives the total number of transitions from state i to state j measured at a lag time τ . The count matrix can be calculated from many, relatively short, unbiased trajectories run in parallel. Because of the Markov property, the initial conditions for these trajectories can be chosen to efficiently generate good statistics for all of the relevant transition elements.

When no information is available *a priori* about which transition elements are most significant, one can use a ratcheting procedure (Appendix C). Many simulations are run in parallel for a time t_s , which must be longer than the lag time τ but can be much shorter than the longest relaxation timescale. Microstates are then determined from coordinates saved during these trajectories, and a new ensemble of trajectories is started with initial conditions preferentially chosen from the microstates with the poorest sampling. This procedure is repeated until $\mathbf{T}$ has satisfactorily converged. Once sufficient statistics have been gathered to crudely estimate $\mathbf{T}$, more systematic adaptive sampling⁴⁷ can be used to choose initial conditions that will reduce the statistical uncertainty of the MSM. However, the initial ratcheting procedure already allows for tremendous speed up in comparison to long, brute force (without enhanced sampling) simulations as it enables the system to cross free energy barriers in linear rather than exponential simulation times. Note that because the protocol does not generate initial conditions according to the equilibrium distribution, the count matrix $\mathbf{C}$ should not be symmetrized when

simulating assembly dynamics. In fact, even if $\mathbf{C}$ is estimated from long, unbiased trajectories that achieve formation of the target structure it should not be symmetrized when calculating dynamics, since assembly to the target structure is an out-of-equilibrium process.

D. Analysis of MSMs

In this subsection, we briefly review analysis of constructed MSMs and discuss an application which is useful for analyzing assembly reactions. Upon spectral decomposition of the transition matrix, the time-dependent state probabilities can be written as

$$\vec{P}(t; \tau) = \sum_{i=1}^N |i\rangle \langle i| \vec{P}(0) e^{-\lambda_i t}, \quad (2)$$

$$\lambda_i = -\log(\omega_i)/\tau,$$

where ω_i is the i th eigenvalue of $\mathbf{T}(\tau)$ and $\langle i|$ and $|i\rangle$ are the corresponding left/right eigenvectors, which are assumed to be normalized. Since $\mathbf{T}(\tau)$ is generally not Hermitian, the left and right eigenvectors are not equivalent. Because the rate matrix is stochastic, there is only one unit eigenvalue, whose associated right eigenvector corresponds to the equilibrium distribution, while all other eigenvalues are positive and real.⁵⁰ The implied timescale, λ_i^{-1} corresponds to the relaxation timescale for eigenmode i . For lag times on which the system satisfies the Markov assumption, the calculated implied timescales are nearly independent of τ .⁸⁸ Checking the convergence of the implied timescales is useful for selecting an appropriate τ , but does not guarantee that the model is Markovian, which also requires converged eigenvectors.⁸⁸

1. Self-assembly reactions

It is often useful to calculate the completion fraction $f_c(t)$, which is defined as the fraction of structures in the target state as a function of time. This quantity can be compared to light scattering or size exclusion chromatography experiments.^{43,98–105} With $\vec{P}$ ordered such that index 1 corresponds to the initial, unassembled state and the largest index N corresponds to the target state, f_c is given by $P_N(t)$. Inserting $P_i(0) = \delta_{i,1}$, with δ the Kronecker delta, into Eq. (2) gives

$$f_c(t) = \sum_{i=1}^N |i\rangle_N \langle i|_1 e^{-\lambda_i t}, \quad (3)$$

where $\langle i|_n$ indicates the n th index of $\langle i|$. The mean completion time $\bar{t}_c = \int_0^\infty (\frac{d}{dt} f_c(t)) t dt$ then follows as

$$\bar{t}_c = \sum_{i=2}^N |i\rangle_N \langle i|_1 \frac{1}{\lambda_i} e^{-\lambda_i t}. \quad (4)$$

E. Analysis using transition path theory

Insight into assembly mechanisms can be obtained from MSMs using Transition Path Theory (TPT),^{106,107} which has been developed in the context of MSMs in Refs. 54, 90, 108,

and 109. We state two of these results which are particularly useful for analyzing assembly reactions here. The microstates that correspond to the transition state ensemble can be identified by calculating the committor probability for each state, which is the probability that a dynamical trajectory initiated from a given state will subsequently visit the target state.^{110,111} We define A as the set of mostly unassembled states that rapidly interconvert (the reactant states), B as the target structure (the product state), and I as all other states (the intermediate states). The forward committor probability, q_i^+ , is the probability that a trajectory started in state i will visit B before A and is given by solving⁵⁴

$$q_i^+ - \sum_{j \in I} T_{ij} q_j^+ = \sum_{j \in B} T_{ij}. \quad (5)$$

Similarly, the backward committor probability q_i^- is the probability that the system was more recently in state A than B . For an equilibrium system, the committor probabilities are related by⁵⁴ $q_i^- = 1 - q_i^+$, but for the models considered here, the target structure acts as an absorbing state (see Appendix C), which gives $q_i^- = 1$ for $i \notin B$.

The relative probabilities of different assembly pathways can be calculated from the flux between states, which is given by⁵⁴ $f_{ij} = \pi_i q_i^- T_{ij} q_j^+$, with π_i as the stationary probability of being found in state i . Since f_{ij} contains non-productive loops that are not on the pathway to completion, the forward flux f_{ij}^+ is defined by subtracting out these contributions:⁵⁴

$$f_{ij}^+ = \text{Max}(0, f_{ij} - f_{ji}). \quad (6)$$

III. MODELS

To test and benchmark our MSM framework, we consider two previously studied models for viral assembly, which differ in their level of detail and more importantly in the type of cargo being packaged. Both models represent capsid protein subunits as rigid bodies with excluded volume geometries and orientation-dependent interactions (following, e.g., Refs. 61–76, and 112–114), designed such that the lowest energy structure is an icosahedron with 20 subunits. Each subunit can be thought of as describing a trimer of proteins that form a $T = 1$ capsid.⁴

A. Patchy sphere model

The first model is motivated by experiments in which capsid proteins from brome mosaic virus (BMV) or Hepatitis B virus (HBV) capsid proteins assemble around nanoparticles functionalized with negative charge.^{115–126} Following the approach developed by Schwartz *et al.*,⁶¹ the subunit excluded volume is spherically symmetric and three attractive patches (bond vectors) are rigidly fixed to the subunit, with each pair of bond vectors forming an angle of 108° (see Fig. 1 and Eq. (A1)). There is a favorable interaction between subunits when (1) the ends of bond vectors nearly overlap, (2) the bond vectors are nearly anti-parallel, and (3) the secondary bond vectors are nearly coplanar. Twenty subunits realizing these conditions results in the minimum energy target structure (a complete capsid) shown in Fig. 1. The interaction strength

FIG. 1. Patchy sphere model geometry. (a) Geometry of two attracting subunits with the bond vectors depicted as arrows and attractors colored teal. The angle between each of the subunit bond vectors is 108° , and the interactions are described in Appendix A (Eq. (A1)). The dihedral angle ϕ_{ij}^b is not shown. (b) A cutaway view of the complete capsid showing the nanoparticle. (c) The complete capsid, which contains 20 subunits arranged with icosahedral symmetry.

is tuned by the parameter ε_B . The nanoparticle has a spherical excluded volume and short-range attractive interactions with capsid subunits, which are tuned by the parameter ε_S and qualitatively represent screened electrostatic attractions. More details about this model can be found in Appendix A 1.

B. Triangles model

This model represents a capsid protein subunit using multiple, spherical “excluders” that enforce excluded volume and spherical “attractors” with short-range, pairwise attractions tuned by the parameter ε_{cc} . Excluders and attractors are arranged so that the minimum energy capsid is an icosahedron (Fig. 2). The cargo considered here is a self-avoiding bead-spring polymer with a persistence length comparable to that of single-stranded RNA (ssRNA) with no base pairing. Polymer beads experience short range attractive interactions with polymer-attractors located on the bottom of protein subunits (see Fig. 2); the interaction strength is tuned by the parameter ε_{cp} . This model has previously been used to study assembly around a polymer⁷² and empty capsid assembly.⁷³ Similar models were applied to empty capsids by Rapaport^{69–71} and Nguyen *et al.*⁶³ More details are given in Appendix A 2.

C. Simulations and units

Subunit positions and orientations are propagated using overdamped Brownian dynamics according to a second order predictor-corrector algorithm.¹²⁷ To represent an experiment with excess capsid protein, each simulation includes a single scaffold (nanoparticle or polymer) and is coupled to a bulk solution by performing grand canonical Monte Carlo moves, in which subunits at the periphery of the simulation

FIG. 2. Triangle model geometry. (a) Subunit geometry with grey excluders, green polymer attractors, and teal subunit attractors. Subunits experience an attractive interaction when the subunit attractors nearly overlap. (b) A cutaway view of the complete capsid with the encapsulated polymer shown in blue. (c) The complete capsid, which contains 20 subunits.

box are exchanged with a reservoir at fixed chemical potential with a frequency consistent with the diffusion limited rate.^{72,122} To obtain dimensionless units, we rescale energies by $k_B T$ and times by a characteristic diffusion timescale (see Appendix A).

IV. RESULTS

We performed simulations over a wide range of parameters to test the ability of the MSMs to accurately reproduce assembly dynamics and to determine the extent of computational speed up in comparison to brute force calculations. To evaluate accuracy, we compare the MSM (Eq. (3)) and brute force dynamics predictions for the cumulative distribution of assembly times, $f_c(t)$ in Figs. 3 and 4. We find that this comparison provides a stringent test of the MSM, as it requires an accurate estimate of all statistically relevant elements of the transition matrix. In particular, capturing the assembly lag phase (the time before the first target structures appear, see Fig. 3) requires accurate estimates for as many as 20 implied timescales and their associated eigenvectors.

The assembly time distributions are accurately predicted from the MSMs for a wide range of parameter values that represent different assembly mechanisms (Figs. 3 and 4). In the patchy sphere model, assembly for the weakest subunit-nanoparticle interaction ($\epsilon_S = 6.6$) is heavily nucleation-dominated, while for the strongest value ($\epsilon_S = 9$) the nucleation timescale is comparable to the elongation timescale¹²⁸ (i.e., the time required for a critical nucleus to grow to completion¹²⁹). These two scenarios are distinguished by the spectrum of implied timescales (Fig. 5). MSMs corresponding to nucleation-dominated parameter sets are characterized by a wide separation between the two largest implied timescales, whereas the MSM corresponding to $\epsilon_S = 9$ yields a dense spectrum of implied timescales. As discussed in Sec. IV C,

FIG. 3. Patchy sphere model: Fraction complete (f_c) as a function of time from the brute-force calculations (symbols) and the MSM calculations (lines). The subunit-subunit binding energy parameter is $\epsilon_B = 10$ for all simulations and the values of the subunit-nanoparticle interaction parameter ϵ_S are indicated on the plot. MSM states are defined by the number of subunits adsorbed to the nanoparticle and the largest cluster size, except for $\epsilon_S = 9$, which also includes Eq. (1) for density variations around the nanoparticle with $\sigma_d = 4\sigma$. Notice that a semi-log scale is used to accommodate a wide range of assembly times. The brute force estimates of $f_c(t)$ are calculated from the completion times for 500–1000 unbiased simulations.

FIG. 4. Triangles model: Fraction complete as a function of time for the brute-force calculations (symbols) and the MSM calculations (lines) for indicated values of the subunit-subunit (ϵ_{cc}) and subunit-polymer (ϵ_{cp}) interaction parameters. It was not possible to simulate assembly using brute-force calculations for $\epsilon_{cp} = 3.0$. MSM states are defined by the number of subunits adsorbed to the polymer and the largest cluster size.

the different parameter values give rise to very different assembly pathways as well.

For most parameter sets, we found that the minimal state definition capable of reproducing assembly dynamics included the number of subunits in the largest cluster and the number of subunits adsorbed to the nanoparticle or polymer. However, it is worth noting that the ratcheting procedure used to estimate the transition matrix considered not only the cluster size, but also the number of intra-cluster bonds (Appendix C). Ratcheting was less efficient when only the cluster size was considered, which suggests that the bonding network within a cluster is important. For the parameter set with the strongest subunit-nanoparticle interaction strength ($\epsilon_S = 9$ in Fig. 3), subunits rapidly adsorbed onto the nanoparticle and it was also necessary to include the free subunit density distribution (Eq. (1)) when building the MSM.

Although a simple state definition yields accurate results for these models, a more detailed order parameter that describes the cluster structure will be required in other situations. To show that MSM construction is feasible even with our most general definition, we also generated MSMs using the graph coordinate (described in the Appendix B). As shown in Fig. 6 the predicted assembly time distributions are identical to those predicted from the simpler coordinates.¹³⁰

FIG. 5. Implied timescales λ_i^{-1} for the patchy sphere model. The five largest timescales are shown for (a) slow, nucleation dominated assembly ($\epsilon_B = 10$, $\epsilon_S = 6.6$) and (b) rapid assembly ($\epsilon_B = 10$, $\epsilon_S = 9$). For (a), the largest implied timescale (excluding the unit eigenvalue) corresponds to the nucleation timescale.

FIG. 6. Feasibility of the graph coordinate. The fraction complete for the patchy sphere model calculated from MSMs defined using the number of sub-units adsorbed to the nanoparticle and either the largest cluster size (dashed line) or the graph coordinate (solid line) for indicated parameter values. Data from long, unbiased simulations is shown as points.

A. Testing convergence

To test that a system is Markovian on timescales corresponding to the lag time τ used to build the transition matrix,

FIG. 7. MSM convergence data for patchy spheres ($\epsilon_B = 10$, $\epsilon_S = 6.6$). (a) The scaled error (Eq. (7)) of the mean completion time $\bar{t}_c$ as a function of total simulation time t_T for the MSM calculation and straightforward dynamics. (b) Convergence of the fraction complete f_c as more trajectories are used to build the MSM. The mean completion time is about 100 times the length of the short simulations used for this parameter set. The MSM has converged after about 1000 short simulations, which corresponds to the total CPU time of 10 brute force simulations. Without the MSM, hundreds of brute force simulations are required to achieve comparable precision.

one typically checks that the implied timescales are nearly lag-time-independent for lag times equal to or exceeding τ (Fig. 5). For assembly systems, convergence can be more stringently tested by determining if the predicted assembly

FIG. 8. Patchy sphere committor probabilities (Eq. (5)). The top row compares committor probabilities calculated using (a) the MSM and (b) brute force dynamics for the same parameter set, while (c) and (d) show committor probabilities calculated using the MSM for two parameter sets for which assembly proceeds by different pathways. Parameter values are indicated on each plot.

time distribution f_c , which depends on all of the implied timescales and associated eigenvectors, becomes independent of lag time (Fig. 7).

B. Simulation time

To assess the computational speed up afforded by MSMs for our systems, we calculated a scaled error Θ for the estimated mean assembly time $\bar{t}_c$:

$$\Theta(t_T) = \left(\frac{\bar{t}_c(t_T) - \bar{t}_c(t_f)}{\bar{t}_c(t_f)} \right)^2. \quad (7)$$

Here t_T is the total simulation time accrued during the short trajectories used to estimate the transition probability matrix; t_T was varied by changing the number of short trajectories. We neglected computational overhead associated with initializing short trajectories and spectral decomposition of the transition matrix, as these factors were negligible. We also calculated Θ as a function of simulation time for straightforward dynamics, with t_T varied by changing the number of trajectories used to estimate $\bar{t}_c(t_f)$. As shown in Fig. 7, the MSM calculation converges with an order of magnitude less simulation time than the estimate based on straightforward dynamics. As expected, the magnitude of speedup depends on the separation of timescales, with greater speed up for large nucleation barriers and only limited speed up in growth-dominated regimes with dense timescale spectra. Importantly, the method is not hindered by multiple, large nucleation barriers such as tend to

occur for low values of ε_B in these capsid assembly systems.⁶² The MSM calculation shown for the lowest polymer-subunit interaction in Fig. 4 ($\varepsilon_{cp} = 3.0$) demonstrates that efficient convergence can be achieved for parameters which are inaccessible to unbiased simulations; a typical unbiased simulation at these parameters would require 2×10^4 cpu hours (over 2 cpu years).

C. Analysis using transition path theory

As a further test of the MSM approach and its ability to provide mechanistic insight, we applied Transition Path Theory (TPT)⁵⁴ (see Sec. II E) to MSMs for several patchy sphere parameter sets. First, we calculated committor probabilities as functions of the cluster size n and the number of adsorbed subunits n_s . Committors calculated using the MSM (Fig. 8(a)) closely agree with those calculated from straightforward dynamical trajectories (Fig. 8(b)). Provided that the chosen collective variables are suitable reaction coordinates,¹¹⁰ states with committor probabilities near $q_i = 1/2$ correspond to critical nuclei, or coordinates from which complete assembly or complete disassembly are equally probable. Figure 8 reveals that the ensemble of critical nuclei depends sensitively on parameter values. For instance, with moderate strengths of subunit-subunit and subunit-nanoparticle interactions (Fig. 8(a)) critical nuclei occur for a narrow range of adsorbed subunits $n_s \approx 15$ but include a broad range of $n = 10 - 15$ subunits in the cluster, indicating that subunit

FIG. 9. (a) and (b): The total forward flux (Eq. (6)) through each state as calculated from the MSM for patchy spheres, with indicated parameter values. Note that values of the forward flux are colored according to a log scale. (c) and (d): Snapshots depicting the typical assembly pathways for the parameter sets in (a) and (b), respectively. The Roman numeral labels indicate points on the flux diagrams to which the snapshots correspond. Bonded subunits are colored red, unbonded subunits adsorbed to the nanoparticle are colored blue, and the nanoparticle is rendered translucent.

adsorption is the controlling degree of freedom, while the assembly state of the adsorbed subunits is fluctuating rapidly during nucleation. In contrast, for weak subunit-subunit and strong subunit-nanoparticle interactions (Fig. 8(d)), for which subunit adsorption is rapid but assembly requires a rare fluctuation to a large cluster size, the critical nuclei include a narrow range of n .

To determine how parameter values influence assembly mechanisms, we used Eq. (6) to calculate the total forward flux through each Markov state (Fig. 9). The sequence of states with the largest flux connecting the initial and target configurations corresponds to the most probable assembly pathway. We see that for nucleation dominated parameters ($\varepsilon_B = 10$, $\varepsilon_S=6.6$, Figs. 9(a) and 9(c)) on average 15 of the 20 subunits required to form a capsid have adsorbed onto the nanoparticle in a disordered manner before any significant assembly occurs, indicating that assembly proceeds according to the *en masse* mechanism defined in Refs. 72, 122, and 131. For weaker subunit-nanoparticle interactions and stronger subunit-subunit interactions ($\varepsilon_S=4$, $\varepsilon_B = 13.5$, Figs. 9(b) and 9(d)), the predominant assembly pathway changes to a nucleation and growth scenario in which the number of adsorbed subunits and the cluster size proceed in lockstep, because adsorption of unassembled subunits is only transient. The calculated assembly pathways closely agree with those determined from brute force dynamics.

V. DISCUSSION AND CONCLUSIONS

In this paper, we propose a general approach to construct MSMs for assembly systems, in which undirected graphs categorize the interaction network of growing clusters. We also discuss several simplified descriptors of the assembly state. A Gaussian-based signature⁹⁵ can be used to describe the density distribution of free subunits; however, this parameter can be neglected when subunit association is reaction-rate limited. We find that MSMs constructed from trajectories that are orders of magnitude shorter than the mean assembly time accurately describe dynamics on all assembly timescales for two different model systems. Analysis of the constructed MSMs using transition path theory⁵⁴ reveals predominant assembly pathways and the transition state ensembles as functions of control parameter values. This information can be used to infer mechanisms and to identify key intermediates that could be targeted to effect changes in assembly yields or structures. Furthermore, the method scales trivially to many processors, which should enable simulation of larger assembling structures than is currently possible.

A. Limitations and possible extensions of the method

A key requirement of the MSM approach is determining a state decomposition for which the system is Markovian on a timescale short in comparison to the assembly time, which requires that the state decomposition distinguishes values of all slow degrees of freedom. Based on the success of our MSM calculation for two different models, we expect our approach to state decomposition to apply to a wide variety of assembly reactions. Our structure-based coordinates can be

applied to any reaction in which subunit association and dissociation timescales are slow in comparison to cluster internal vibrations (i.e., fluctuations around preferred interaction distances and angles). The coordinates described in this paper do not explicitly describe the polymer degrees of freedom and thus require that these are fast in comparison to assemblage properties. However, since polymer segments can be indexed, polymer configurations can be categorized via the existing RMSD-based approaches.^{86,87,94} Other degrees of freedom which might potentially be included in MSM definitions include nucleic acid secondary and tertiary structure and capsid protein conformational states (e.g., Ref. 124).

B. Expected scaling of the method

The minimum time required for simulating assembly of a target structure with N subunits without enhanced sampling, in the absence of nucleation barriers, scales as $\sim N^2$. This assumes that the elongation time (defined in Sec. IV) scales $\sim N$ (see Ref. 129) and that the average number of particles in each simulation (and thus the CPU time/step) scales $\sim N$. As shown in Appendix D, the approach described in this article achieves similar scaling even in the presence of nucleation barriers. Use of more sophisticated adaptive sampling can be expected to further reduce the required CPU time.⁴⁷ In addition to overcoming slow relaxation timescales, a significant advantage of the MSM approach over unbiased dynamics is that the short trajectories used to estimate the transition rates are trivially parallelizable, and thus can be extended to many more processors than unbiased dynamics.

Similar estimates can be applied to scaling with other parameters. For example, while the CPU time for unbiased dynamics increases with decreasing total subunit concentration c_0 according to $\sim c_0^{n^*}$ with n^* the critical nucleus size, the MSM approach will scale as the minimum assembly time, $\sim c_0^{-1}$ (Ref. 129). This scaling could be further improved upon using coordinates that consider positions of free subunits (e.g., w described in Sec. II B) or methods that implicitly account for diffusive dynamics (e.g., Refs. 45 and 132).

C. Application to non-templated assembly

The state definitions described here can be readily generalized to account for multiple clusters to describe multiple scaffolds or empty capsid assembly. However, the size of the state space will increase exponentially with the number of clusters. Thus the method is most practical when only a few clusters assemble at any given time in a spatially correlated region, so that tracking only the largest few clusters is sufficient for state decomposition. Also see Ref. 141 for further discussion on simulations of small assembly systems.¹³³ This situation occurs in empty capsid assembly when nucleation barriers are sufficiently large that nucleation timescales are slow in comparison to growth times.^{128,129} While this condition might seem limiting, it is precisely the regime in which enhanced sampling methods can be expected to provide significant speedup, since there is a significant separation of timescales. Furthermore, most experiments are performed under these conditions, with concentrations of

intermediates often so low as to be undetectable, because assembly is most robust against kinetic traps when nucleation is rate limiting.^{4,99,128,129} The MSM approach described here enables simulating assembly at these experimentally relevant parameter values.

ACKNOWLEDGMENTS

We thank Vijay Pande for enlightening discussions and for suggesting to develop a coordinate based on subunit densities. This work was supported by Award Number R01GM108021 from the National Institute of General Medical Sciences. We also acknowledge support by NSF-MRSEC-0820492. Computational resources were provided by the National Science Foundation through XSEDE computing resources (Open Science Grid and Trestles) and the Brandeis HPCC.

APPENDIX A: MODEL DETAILS

1. Patchy sphere model

In this model, adapted from Ref. 122, the minimum energy structure is a complete capsid of 20 subunits encapsulating the nanoparticle. Subunits have a spherical excluded volume with three attractive patches, or bond vectors, that are separated by 108° and rotate rigidly with the subunit. The attractive interaction between two complementary bond vectors on respective subunits i and j is maximized when (1) the distance between the attractors r_{ij}^b is minimized, (2) the angle θ_{ij}^b between bond vectors is minimized, and (3) the dihedral angle ϕ_{ij}^b calculated from two secondary bond vectors, which are not involved in the primary interaction, is minimized. The schematic of subunit interactions is shown in Fig. 1. Minimizing ϕ_{ij}^b creates an interaction that resists torsion and enforces angular specificity commensurate with a complete capsid. The potentials are given by Eq. (A1)¹²²

$$\begin{aligned}
 U &= U_{\text{rep}}(R_{ij}) + \sum_b U_{\text{att}}(r_{ij}^b) S(\theta_{ij}^b, \phi_{ij}^b), \\
 U_{\text{rep}}(R_{ij}) &= \mathcal{L}_{12}(R_{ij}, 2^{\frac{1}{6}}\sigma, \sigma), \\
 U_{\text{att}}(r_{ij}) &= \varepsilon_B \mathcal{L}_{12}((r_{ij} + 2^{\frac{1}{6}}\sigma), 2^{\frac{1}{6}}\sigma, r_c), \\
 S(\theta, \phi) &= \frac{1}{4} \Theta(\theta - \theta_c) \Theta(\phi - \phi_c) \\
 &\quad (\cos(\pi\theta/\theta_c) + 1) (\cos(\pi\phi/\phi_c) + 1),
 \end{aligned} \tag{A1}$$

with $\mathcal{L}_p$ a generalized truncated and shifted Lennard-Jones function:

$$\begin{aligned}
 \mathcal{L}_p(x, x_c, \sigma) &\equiv 4 \left(\left(\frac{x}{\sigma} \right)^{-p} - \left(\frac{x}{\sigma} \right)^{-p/2} \right. \\
 &\quad \left. - \left(\frac{x_c}{\sigma} \right)^{-p} + \left(\frac{x_c}{\sigma} \right)^{-p/2} \right) \Theta(x - x_c). \tag{A2}
 \end{aligned}$$

In Eq. (A1), the index b sums over pairs of complementary bond vectors, $\Theta(x)$ is the Heaviside step function and R_{ij} is the subunit center-to-center distance.

The nanoparticle is modeled as a spherical excluded volume with a Lennard-Jones potential whose argument is shifted so that the minimum occurs when subunits are on the

surface:

$$\begin{aligned}
 U_S(r) &= \varepsilon_S \mathcal{L}_{12}(r_{\text{eff}}, r_c, \sigma), \\
 r_{\text{eff}} &\equiv r - R_S + 0.5\sigma,
 \end{aligned} \tag{A3}$$

with r as the nanoparticle-subunit center-to-center distance, R_S as the radius of the nanoparticle, and ε_S as a tuneable subunit-nanoparticle interaction strength. This potential qualitatively represents subunit electrostatic interactions with functionalized gold nanoparticles.^{115–121}

a. Units and parameter values

Lengths have units of σ , the subunit diameter, energies have units of $k_B T$, and times have units of $t_0 = \sigma^2/D$, where D is the subunit diffusion constant. The box side length is 18σ , the nanoparticle radius is $R_S = 0.9\sigma$, grand canonical particle exchanges occur at least 9σ from the nanoparticle surface, and the reservoir subunit concentration is $c_0 = 0.005\sigma^{-3}$ (defined as $c_0 = N\sigma^3/L^3$ with N the number of subunits). To qualitatively estimate the subunit concentration, following the approach described in Ref. 72, we set the diameter of a subunit to $\sigma = 7$ nm based on the satellite tobacco mosaic virus (STMV) structure. The reservoir concentration is then $25 \mu\text{M}$. The attractor cutoff values are $r_c = 2.5\sigma$, $\theta_c = 1$ and $\phi_c = \pi$.

2. Triangles model

In this model taken from Ref. 72, the minimum energy structure is a complete capsid of 20 subunits assembled around a bead-spring polymer. The truncated pyramidal subunits are composed of two layers of excluders, which enforce excluded volume interactions, and two layers of capsomer attractors on the edges, which mimic hydrophobic and electrostatic attractions (see Fig. 10). The attractive interaction between subunits is maximized when the capsomer attractors are perfectly overlapping. In this situation, excluders on either side of the interface are separated by the cut off distance of their potential.

The polymeric cargo is a freely jointed chain of spherical monomers, which have an excluded volume that represents steric interactions and screened electrostatic repulsions. The model represents the polymer in good solvent, which

FIG. 10. Triangles model subunit geometry with distances marked. Excluders are gray, subunit attractors are teal, and the black circles depict the overlapping arrangement of excluders (polymer attractors not shown). Note that the position of the k th attractor on subunit i is given by $\mathbf{r}_i^k \equiv \mathbf{R}_i + \mathbf{a}_i^k$, and the position of the l th excluder on subunit j is given by $\mathbf{r}_j^l \equiv \mathbf{R}_j + \mathbf{b}_j^l$ (as used in Eqs. (A4) and (A5)).

behaves as a self-avoiding random walk with radius of gyration $R_g = 0.21N_p^{3/5}\sigma_b$; N_p is the number of monomers and σ_b is the monomer diameter. The polymer has short-range interactions with the polymer attractors on the bottom of each subunit, which represent screened electrostatic interactions. All interactions in this model can be decomposed into pairwise potentials.

a. Capsomer-capsomer interactions

Capsomer subunit interactions decompose into pairwise interactions between their constituent building blocks—the excluders and attractors. Excluders in a subunit i experience a truncated Lennard-Jones-like potential when within a cutoff radius r_c of excluders in subunit $j \neq i$ (i.e., subunits excluders repel each other when overlapping). Attractors in subunit i experience an attractive Lennard-Jones-like potential with *commensurate attractors* from another subunit $j \neq i$. Commensurate attractors are defined to be the attractors which would overlap in a complete capsid:⁷²

$$U_{cc}(i, j) = \sum_{k, l \in \text{excluders}} \mathcal{L}_8(|\mathbf{r}_i^k - \mathbf{r}_j^l|, 2^{1/4}\sigma_b, \sigma_b) + \sum_{k, l \in \text{attractors}} \varepsilon_{cc} \mathcal{L}_4(|\mathbf{r}_i^k - \mathbf{r}_j^l| - 2^{1/2}\sigma_a, 4\sigma_a, \sigma_a), \quad (\text{A4})$$

where the first sum is over all pairs of excluders between subunits i and j and the second sum is over all commensurate attractor pairs between subunits i and j , with ε_{cc} as an adjustable parameter that tunes the attraction strength. The coordinates of the k th excluder/attractor are given by $\mathbf{r}_i^k$ and $\mathcal{L}_4$ is given by Eq. (A2).

b. Capsomer-polymer interactions

The capsomer-polymer attraction is similar to the capsomer-capsomer interaction, with an attractive potential that is minimized when polymer beads overlap subunit attractors and a repulsive potential to account for excluded volume. For a capsid subunit i and polymer subunit j , the potential is

$$U_{cp}(i, j) = \sum_{k \in \text{excluders}} \mathcal{L}_8(|\mathbf{r}_i^k - \mathbf{r}_j|, 2^{1/4}\sigma_{bp}, \sigma_{bp}) + \sum_{k \in \text{poly attractors}} \xi_k \varepsilon_{cp} \mathcal{L}_8(|\mathbf{r}_i^k - \mathbf{r}_j| + 2^{1/4}\sigma_p, 4\sigma_p, \sigma_p), \quad (\text{A5})$$

$$\sigma_{bp} \equiv \frac{1}{2}(\sigma_b + \sigma_p),$$

with the first sum over all subunit excluders and the second sum over all polymer attractors on the subunit. ε_{cp} parameterizes the attraction strength for each attractor. σ_p is the diameter of a polymer bead and is set to $0.4\sigma_b$. ξ_k is a factor that decreases the interaction strength for the outmost three polymer attractors on the subunit ($\xi_k = \frac{1}{2}$) to compensate for these sites overlapping in a complete capsid. ξ_k is set to one for all other attractors.

c. Polymer-polymer interactions

The polymer potential includes a “bonded” interaction between monomers that are nearest neighbors along the chain and a nonbonded, excluded volume interaction with all other polymer beads. For bead coordinates $\mathbf{r}_i$ and $\mathbf{r}_j$ ($i \neq j$)

$$U_{pp}(r_{ij} \equiv |\mathbf{r}_i - \mathbf{r}_j|) = \begin{cases} \mathcal{L}_8(r_{ij}, 2^{1/4}\sigma_p, \sigma_p) & : r_{ij} < 2^{1/4}\sigma_p \\ \mathcal{L}_8(2^{5/4}\sigma_p - r_{ij}, 2^{5/4}\sigma_p, \sigma_p) & : r_{ij} > 2^{1/4}\sigma_p \\ & \& \{i, j\} \text{ bonded} \\ 0 & : \text{otherwise.} \end{cases} \quad (\text{A6})$$

d. Units and parameter values

The capsomer subunits have anisotropic translational and rotational diffusion constants calculated using Hydrosb7.C.^{72,134} Lengths are scaled by σ_b and times are scaled by t_0 , which is the Brownian time for a sphere with diameter σ_b . The box side length is $40\sigma_b$, grand canonical moves are performed at least $10\sigma_b$ from the center of the polymer, the reservoir concentration is $0.00047\sigma_b^{-3}$, and the polymer length is 200. To qualitatively estimate the subunit concentration, we set the pseudoatom diameter to $\sigma_b = 2.36$ nm based on the STMV structure (see Ref. 72), which results in a reservoir concentration of $59 \mu\text{M}$. The attractor and polymer bead diameters are respectively $\sigma_a = 0.2\sigma_b$ and $\sigma_b = 0.4\sigma_b$.

APPENDIX B: COLLECTIVE COORDINATES

1. Graph isomorphism

Here we describe a general approach to state decomposition for self-assembly reactions, which represents the bonding network of a growing cluster as an undirected graph (Fig. 11), with subunits as the nodes and strong interactions,

FIG. 11. (a) and (c) show clusters of patchy spheres growing on the nanoparticle during assembly while (b) and (d) show their respective graph representations. Attractors are not shown, but bonds (strong attractions) between subunits are depicted as teal cylinders. In this work, only the largest cluster was considered when constructing the MSM, which would exclude the lower dimer in (c) and (d) from consideration.

which we denote as “bonds,” as the edges (see Fig. 11). Structures which differ only in the indices that label their subunits or fluctuations of bond geometries have the same graphs; to account for subunit index changes. We only applied this order parameter to class I subunits (i.e., the largest growing cluster) as described in Sec. II B, but it can be generalized to class II subunits by considering subunit interactions with the scaffold. Isomorphic graphs were identified using the algorithm in the Boost C++ libraries.^{135,136} While identification was fast for all cases we considered, faster algorithms are available (e.g., Refs. 137 and 138), with further improvements possible in the case where all graphs are planar. The search process can be trivially parallelized by initially separating structures based on quantities such as number of subunits, bonds, and cycles.

2. Alternative structure-based order parameters

Through *a priori* knowledge of the system, it is possible to devise state decompositions which result in fewer states. For example, subunits attached to a cluster by only a single interaction are transient under typical assembly conditions.⁶² Graphs can therefore be simplified by pruning nodes which are connected to the rest of the graph by only a single edge. Further simplifications can be made by only recording features of the graph. For systems in which single bonds are relatively unstable, clusters tend to grow by sequential completion of polygons.⁷³ One can then record the set of all cycles (which correspond to complete polygons) in the graph, or one can account only for the number of subunits n and number of bonds n_b within a growing cluster. These two quantities can be combined into a single coordinate as

$$\gamma = a(n_b - n) + n_b \quad (\text{B1})$$

with a parameter a that prevents degeneracy among typical structures by separating structures with different numbers of complete polygons. We used $a = 5$ or $a = 10$.

APPENDIX C: SIMULATIONS AND PROTOCOL FOR ESTIMATING THE TRANSITION MATRIX

This section describes the ratcheting procedure that we used to estimate the elements of the transition matrix. We began by performing n_s simulations (with $n_s = 50 - 100$), each started from initial configurations in which subunits had random positions and orientations (except that subunits were not allowed to overlap). Each simulation was run for a time t_s , and snapshots were saved regularly to a database. Saved configurations were then classified into states. This classification could be based on any of the state decomposition approaches described in this article; we used the coordinate defined in Eq. (B1). A second iteration was then started, in which n_s new simulations of length t_s were performed with snapshots regularly saved to the database. To efficiently estimate important elements of the transition matrix, initial configurations were preferentially chosen from states from which fewer simulations had already been initialized. Iterations were then repeated until the MSM converged (Fig. 7). This procedure is useful when no *a priori* knowledge is available about the

system, since it both focuses sampling on poorly sampled states and identifies new states that arise through the natural dynamics of the system. However, once enough states have been gathered to construct an MSM, more sophisticated adaptive sampling is possible.⁴⁷

Optimal values of the parameters n_s and t_s depend on both the system being simulated and the available computational resources and thus need to be chosen through trial and error. The simulation time t_s must be longer than the lag time τ , but should not be too long in order to efficiently ratchet the system over free energy barriers. We found that a t_s of around $20\tau - 100\tau$ usually worked well. The total number of simulations required to generate a converged MSM varied depending on the parameters, but we found that about 1000 simulations were sufficient for most parameter sets. We used a modified version of MSMBuild⁸⁶ to construct the transition matrix and to calculate its eigenvalues and eigenvectors. For most parameter sets, it was sufficient to calculate only the top 20 eigenvectors.

For the models that we consider here, dissociation of a subunit from a complete capsid occurs only on timescales long compared to the assembly time for most parameter sets. Thus, well-formed target states are effectively absorbing states on simulated timescales (see Refs. 4, 129, and 139 for discussion); however, on rare occasions, capsids would form in strained configurations and eventually dissociate. When building the transition matrix, it was important to acquire sufficient statistics from short trajectories in the target state to balance the rare dissociations from strained target states. This need could be avoided by either more stringently defining the target state or by manually defining the target state as an absorbing state during transition matrix construction. Note that if one were interested in capturing the much longer time scale associated with dissociation of a well-formed target structure, additional coordinates describing fluctuations of the subunit-subunit bonds would be required to efficiently sample this transition.

APPENDIX D: EXPECTED SCALING OF THE METHOD WITH TARGET STRUCTURE SIZE

In this appendix, we consider scaling of the specific approaches described in this article with target structure size N . However, we note that the size of the state space and the number of simulations used to estimate transition rates can be considerably reduced by lumping and adaptive sampling.⁴⁷

The total CPU time required depends on the number of transition rates to be estimated and the number of interacting particles in the short simulations used for estimation. We begin with the most general state decomposition, based on the graphs, which leads to the largest state space. While the number of possible graphs for a given model is exponential in the number of subunits N in the target structure, productive assembly reactions only visit a relatively small subset of possible structures.^{73,140} Since the MSM state space includes only those states which are visited during estimation of the transition rate matrix, we expect only polynomial scaling. Although additional degrees of freedom, such as the number of subunits interacting with the scaffold (n_s), increases the size

FIG. 12. (a) The number of states as a function of completed capsid size N . (We have included the trivial point $N = 1$ for which there are two structures.) (b) The number of unique transitions between states as a function of N . Both plots are from MSMs built using either (n_s, n) or $(n_s, \text{graphIso})$. n_s is the number of subunits adsorbed to the nanoparticle, n is the largest cluster size, and graphIso is the graph isomorphism. The model for a size 20 capsid is described in Sec. III; sizes 30 and 60 were developed in Refs. 62 and 122 (respectively). Size 12 is the same as the size 20 model, except with modified bond vectors for pentamer subunits to make a dodecahedron.

of the state space, only a subset of n_s are encountered for most values of the structural coordinate. Finally, since most transitions involve association or dissociation of one or a few subunits, the transition matrix is sparse, with the number of nonzero elements roughly proportional to its linear size.

To approximately quantify these scaling estimates, we built MSMs for patchy sphere models with sizes $N = 12, 20, 30, 60$. This analysis was complicated by the fact that the number of structures visited depends sensitively on the parameters (more aggressive assembly reactions visit more structures) and the details of how ratcheting is performed (longer or poorly designed ratcheting protocols find more statistically irrelevant structures). To minimize these variations as much as possible, we used parameter sets that led to similar assembly pathways at each size, and we built MSMs from long unbiased simulations. The numbers of graph structures and the total number of non-zero transition matrix elements for each model size N are shown in Fig. 12. For comparison, and to illustrate the effectiveness of other approaches, the data for MSMs built using n_s and n is also plotted. The number of states scales between N and N^2 , whereas the number of unique transitions scales roughly linearly with N . It is worth noting that by eliminating transition rate elements with 5 or fewer recorded transitions, the total number of transition elements in the graph estimate was reduced by a factor of about 20 without affecting the accuracy of the MSM. A comparable reduction was achieved for the simplified state decomposition.

To estimate the total CPU time required to build an MSM, we note that the average number of particles in each simulation will scale $\sim N$. Thus, based on Fig. 12, the total CPU time required will scale as $\sim N^2$.

¹D. L. D. Caspar and A. Klug, *Cold Spring Harbor Symp. Quant. Biol.* **27**, 1 (1962).

²A. Zlotnick and S. Mukhopadhyay, *Trends Microbiol.* **19**, 14 (2011).

³J. A. Speir and J. E. Johnson, *Curr. Opin. Struct. Biol.* **22**, 65 (2012).

⁴M. Hagan, *Adv. Chem. Phys.* **155**, 1 (2014).

⁵Y. Yang, R. Meyer, and M. F. Hagan, *Phys. Rev. Lett.* **104**, 258102 (2010).

⁶S. Whitlam, *Phys. Rev. Lett.* **105**, 088102 (2010).

⁷S. Sacanna, W. T. M. Irvine, P. M. Chaikin, and D. J. Pine, *Nature (London)* **464**, 575 (2010).

⁸E. Duguet, A. Desert, A. Perro, and S. Ravaine, *Chem. Soc. Rev.* **40**, 941 (2011).

⁹Y. Yang, L. Gao, G. P. Lopez, and B. B. Yellen, *ACS Nano* **7**, 2705 (2013).

¹⁰S. Whitlam, I. Tamblin, T. K. Haxton, M. B. Wieland, N. R. Champness, J. P. Garrahan, and P. H. Beton, *Phys. Rev. X* **4**, 011044 (2014).

¹¹Y. Wang, Y. Wang, D. R. Breed, V. N. Manoharan, L. Feng, A. D. Hollingsworth, M. Weck, and D. J. Pine, *Nature (London)* **491**, 51 (2012).

¹²Y. Wang, A. D. Hollingsworth, S. K. Yang, S. Patel, D. J. Pine, and M. Weck, *J. Am. Chem. Soc.* **135**, 14064 (2013).

¹³A. Walther and A. H. E. Mueller, *Chem. Rev.* **113**, 5194 (2013).

¹⁴G. van Anders, N. K. Ahmed, R. Smith, M. Engel, and S. C. Glotzer, *ACS Nano* **8**, 931 (2014).

¹⁵A. H. Groeschel, A. Walther, T. I. Loebling, F. H. Schacher, H. Schmalz, and A. H. E. Mueller, *Nature (London)* **503**, 247 (2013).

¹⁶P. Rothmund, *Nature (London)* **440**, 297 (2006).

¹⁷B. Sacca and C. M. Niemeyer, *Angew. Chem., Int. Ed.* **51**, 58 (2012).

¹⁸T. Torring, N. V. Voigt, J. Nangreave, H. Yan, and K. V. Gothelf, *Chem. Soc. Rev.* **40**, 5636 (2011).

¹⁹A. C. Pan, D. Sezer, and B. Roux, *J. Phys. Chem. B* **112**, 3432 (2008).

²⁰V. Ovchinnikov, M. Karplus, and E. Vanden-Eijnden, *J. Chem. Phys.* **134**, 085103 (2011).

²¹P. G. Bolhuis, D. Chandler, C. Dellago, and P. L. Geissler, *Annu. Rev. Phys. Chem.* **53**, 291 (2002).

²²S. Fischer, K. W. Olsen, K. Nam, and M. Karplus, *Proc. Natl. Acad. Sci. U.S.A.* **108**, 5608 (2011).

²³R. Elber, *Biophys. J.* **92**, L85 (2007).

²⁴F. Pietrucci, F. Marinelli, P. Carloni, and A. Laio, *J. Am. Chem. Soc.* **131**, 11811 (2009).

²⁵M. Lei, M. I. Zavodszky, L. A. Kuhn, and M. F. Thorpe, *J. Comput. Chem.* **25**, 1133 (2004).

²⁶D. Moroni, P. G. Bolhuis, and T. S. van Erp, *J. Chem. Phys.* **120**, 4055 (2004).

²⁷A. Dickson and A. R. Dinner, *Annu. Rev. Phys. Chem.* **61**, 441 (2010).

²⁸R. J. Allen, P. B. Warren, and P. R. ten Wolde, *Phys. Rev. Lett.* **94**, 018104 (2005).

²⁹J. Pfandtner, D. Branduardi, M. Parrinello, T. D. Pollard, and G. A. Voth, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 12723 (2009).

³⁰A. Barducci, M. Bonomi, and M. Parrinello, *Biophys. J.* **98**, L44 (2010).

³¹B. W. Zhang, D. Jasnow, and D. M. Zuckerman, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 18043 (2007).

³²G. A. Huber and S. Kim, *Biophys. J.* **70**, 97 (1996).

³³A. L. Ferguson, A. Z. Panagiotopoulos, I. G. Kevrekidis, and P. G. Debenedetti, *Chem. Phys. Lett.* **509**, 1 (2011).

³⁴T. S. van Erp, "Dynamical rare event simulation techniques for equilibrium and nonequilibrium systems," in *Kinetics and Thermodynamics of Multistep Nucleation and Self-Assembly in Nanoscale Materials*, Advances in Chemical Physics Vol. 151, edited by G. Nicolis and D. Maes (Wiley & Sons, 2012), pp. 27–60.

³⁵N. B. Becker, R. J. Allen, and P. R. ten Wolde, *J. Chem. Phys.* **136**, 174118 (2012).

- ³⁶F. Jamalyaria, R. Rohlf, and R. Schwartz, *J. Comput. Phys.* **204**, 100 (2005).
- ³⁷T. Keef, C. Micheletti, and R. Twarock, *J. Theor. Biol.* **242**, 713 (2006).
- ³⁸M. Hemberg, S. Yaliraki, and M. Barahona, *Biophys. J.* **90**, 3029 (2006).
- ³⁹E. C. Dykeman, P. G. Stockley, and R. Twarock, *Phys. Rev. E* **87**, 022717 (2013).
- ⁴⁰B. Sweeney, T. Zhang, and R. Schwartz, *Biophys. J.* **94**, 772 (2008).
- ⁴¹T. Q. Zhang and R. Schwartz, *Biophys. J.* **90**, 57 (2006).
- ⁴²N. Misra, D. Lees, T. Q. Zhang, and R. Schwartz, *Comput. Math. Method Med.* **9**, 277 (2008).
- ⁴³M. S. Kumar and R. Schwartz, *Phys. Biol.* **7**, 045005 (2010).
- ⁴⁴L. Xie, G. Smith, X. Feng, and R. Schwartz, *Biophys. J.* **103**, 1545 (2012).
- ⁴⁵G. R. Smith, L. Xie, B. Lee, and R. Schwartz, *Biophys. J.* **106**, 310 (2014).
- ⁴⁶G. R. Bowman, V. A. Voelz, and V. S. Pande, *J. Am. Chem. Soc.* **133**, 664 (2011).
- ⁴⁷G. R. Bowman, D. L. Ensign, and V. S. Pande, *J. Chem. Theory Comput.* **6**, 787 (2010).
- ⁴⁸G. R. Bowman, K. A. Beauchamp, G. Boxer, and V. S. Pande, *J. Chem. Phys.* **131**, 124101 (2009).
- ⁴⁹W. C. Swope, J. W. Pitera, F. Suits, M. Pitman, M. Eleftheriou, B. G. Fitch, R. S. Germain, A. Rayshubski, T. J. C. Ward, Y. Zhestkov, and R. Zhou, *J. Phys. Chem. B* **108**, 6582 (2004).
- ⁵⁰W. C. Swope, J. W. Pitera, and F. Suits, *J. Phys. Chem. B* **108**, 6571 (2004).
- ⁵¹J.-H. Prinz, J. D. Chodera, V. S. Pande, W. C. Swope, J. C. Smith, and F. Noe, *J. Chem. Phys.* **134**, 244108 (2011).
- ⁵²S. Park, D. L. Ensign, and V. S. Pande, *Phys. Rev. E* **74**, 066703 (2006).
- ⁵³V. S. Pande, K. Beauchamp, and G. R. Bowman, *Methods* **52**, 99 (2010).
- ⁵⁴F. Noe, C. Schutte, E. Vanden-Eijnden, L. Reich, and T. R. Weikl, *Proc. Natl. Acad. Sci. U.S.A.* **106**, 19011 (2009).
- ⁵⁵T. J. Lane, G. R. Bowman, K. Beauchamp, V. A. Voelz, and V. S. Pande, *J. Am. Chem. Soc.* **133**, 18413 (2011).
- ⁵⁶G. Jayachandran, V. Vishal, and V. S. Pande, *J. Chem. Phys.* **124**, 164902 (2006).
- ⁵⁷J. D. Chodera, W. C. Swope, F. Noe, J.-H. Prinz, M. R. Shirts, and V. S. Pande, *J. Chem. Phys.* **134**, 244107 (2011).
- ⁵⁸N. S. Hinrichs and V. S. Pande, *J. Chem. Phys.* **126**, 244101 (2007).
- ⁵⁹J. D. Chodera, N. Singhal, V. S. Pande, K. A. Dill, and W. C. Swope, *J. Chem. Phys.* **126**, 155101 (2007).
- ⁶⁰P. Deuffhard and M. Weber, *Linear Algebra Its Appl.* **398**, 161 (2005).
- ⁶¹R. Schwartz, P. W. Shor, P. E. Prevelige, and B. Berger, *Biophys. J.* **75**, 2626 (1998).
- ⁶²M. F. Hagan and D. Chandler, *Biophys. J.* **91**, 42 (2006).
- ⁶³H. D. Nguyen, V. S. Reddy, and C. L. Brooks, *Nano Lett.* **7**, 338 (2007).
- ⁶⁴A. W. Wilber, J. P. K. Doye, A. A. Louis, E. G. Noya, M. A. Miller, and P. Wong, *J. Chem. Phys.* **127**, 085106 (2007).
- ⁶⁵H. Nguyen and C. Brooks, *Nano Lett.* **8**, 4574 (2008).
- ⁶⁶H. D. Nguyen, V. S. Reddy, and C. L. Brooks, *J. Am. Chem. Soc.* **131**, 2606 (2009).
- ⁶⁷I. G. Johnston, A. A. Louis, and J. P. K. Doye, *J. Phys.: Condens. Matter* **22**, 104101 (2010).
- ⁶⁸A. W. Wilber, J. P. K. Doye, A. A. Louis, and A. C. F. Lewis, *J. Chem. Phys.* **131**, 175102 (2009).
- ⁶⁹D. C. Rapaport, J. E. Johnson, and J. Skolnick, *Comput. Phys. Commun.* **121–122**, 231 (1999).
- ⁷⁰D. Rapaport, *Phys. Rev. E* **70**, 051905 (2004).
- ⁷¹D. Rapaport, *Phys. Rev. Lett.* **101**, 186101 (2008).
- ⁷²O. Elrad and M. F. Hagan, *Phys. Biol.* **7**, 045003 (2010).
- ⁷³M. F. Hagan, O. M. Elrad, and R. L. Jack, *J. Chem. Phys.* **135**, 104115 (2011).
- ⁷⁴J. P. Mahalik and M. Muthukumar, *J. Chem. Phys.* **136**, 135101 (2012).
- ⁷⁵J. D. Perlmutter, C. Qiao, and M. F. Hagan, *eLife* **2**, e00632 (2013).
- ⁷⁶D. C. Rapaport, *Phys. Rev. E* **86**, 051917 (2012).
- ⁷⁷S. Hicks, "Statistical mechanical models of virus capsid assembly," Ph.D. thesis (Cornell University, 2010).
- ⁷⁸S. Sriraman, I. G. Kevrekidis, and G. Hummer, *J. Phys. Chem. B* **109**, 6479 (2005).
- ⁷⁹J. D. Chodera, W. C. Swope, J. W. Pitera, and K. A. Dill, *Multiscale Model. Simul.* **5**, 1214 (2006).
- ⁸⁰E. J. Sorin and V. S. Pande, *Biophys. J.* **88**, 2472 (2005).
- ⁸¹S. P. Elmer, S. Park, and V. S. Pande, *J. Chem. Phys.* **123**, 114903 (2005).
- ⁸²M. Andrec, A. K. Felts, E. Gallicchio, and R. M. Levy, *Proc. Natl. Acad. Sci. U.S.A.* **102**, 6801 (2005).
- ⁸³B. L. de Groot, X. Daura, A. E. Mark, and H. Grubmüller, *J. Mol. Biol.* **309**, 299 (2001).
- ⁸⁴M. E. Karpen, D. J. Tobias, and C. L. Brooks, *Biochemistry* **32**, 412 (1993).
- ⁸⁵N. Singhal, C. D. Snow, and V. S. Pande, *J. Chem. Phys.* **121**, 415 (2004).
- ⁸⁶G. R. Bowman, X. Huang, and V. S. Pande, *Methods* **49**, 197 (2009).
- ⁸⁷K. A. Beauchamp, G. R. Bowman, T. J. Lane, L. Maibaum, I. S. Haque, and V. S. Pande, *J. Chem. Theory Comput.* **7**, 3412 (2011).
- ⁸⁸J.-H. Prinz, H. Wu, M. Sarich, B. Keller, M. Senne, M. Held, J. D. Chodera, C. Schutte, and F. Noe, *J. Chem. Phys.* **134**, 174105 (2011).
- ⁸⁹K. A. Beauchamp, R. McGibbon, Y.-S. Lin, and V. S. Pande, *Proc. Natl. Acad. Sci. U.S.A.* **109**, 17807 (2012).
- ⁹⁰V. A. Voelz, G. R. Bowman, K. Beauchamp, and V. S. Pande, *J. Am. Chem. Soc.* **132**, 1526 (2010).
- ⁹¹G. R. Bowman and P. L. Geissler, *Proc. Natl. Acad. Sci. U.S.A.* **109**, 11681 (2012).
- ⁹²G. Perez-Hernandez, F. Paul, T. Giorgino, G. De Fabritiis, and F. Noé, *J. Chem. Phys.* **139**, 015102 (2013).
- ⁹³C. R. Schwantes and V. S. Pande, *J. Chem. Theory Comput.* **9**, 2000 (2013).
- ⁹⁴M. Senne, B. Trendelkamp-Schroer, A. S. Mey, C. Schutte, and F. No, *J. Chem. Theory Comput.* **8**, 2223 (2012).
- ⁹⁵C. Gu, C. Huang-Wei, L. Maibaum, V. S. Pande, G. E. Carlsson, and L. J. Guibas, *BMC Bioinf.* **14**, 1 (2013).
- ⁹⁶J. Grant, R. L. Jack, and S. Whitelam, *J. Chem. Phys.* **135**, 214505 (2011).
- ⁹⁷A. Zlotnick, *Virology* **315**, 269 (2003).
- ⁹⁸P. E. Prevelige, D. Thomas, and J. King, *Biophys. J.* **64**, 824 (1993).
- ⁹⁹A. Zlotnick, J. M. Johnson, P. W. Wingfield, S. J. Stahl, and D. Endres, *Biochemistry* **38**, 14644 (1999).
- ¹⁰⁰A. Zlotnick, R. Aldrich, J. M. Johnson, P. Ceres, and M. J. Young, *Virology* **277**, 450 (2000).
- ¹⁰¹G. L. Casini, D. Graham, D. Heine, R. L. Garcea, and D. T. Wu, *Virology* **325**, 320 (2004).
- ¹⁰²C. Chen, C. C. Kao, and B. Dragnea, *J. Phys. Chem. A* **112**, 9405 (2008).
- ¹⁰³C. Berthet-Colominas, M. Cuillel, M. H. J. Koch, P. Vachette, and B. Jacrot, *Eur. Biophys. J.* **15**, 159 (1987).
- ¹⁰⁴S. Kler, R. Asor, C. Li, A. Ginsburg, D. Harries, A. Oppenheim, A. Zlotnick, and U. Raviv, *J. Am. Chem. Soc.* **134**, 8823 (2012).
- ¹⁰⁵M. Tsiang, A. Niedziela-Majka, M. Hung, D. Jin, E. Hu, S. Yant, D. Samuel, X. Liu, and R. Sakowicz, *Biochemistry* **51**, 4416 (2012).
- ¹⁰⁶W. E and E. Vanden-Eijnden, *Annu. Rev. Phys. Chem.* **61**, 391 (2010).
- ¹⁰⁷P. Metzner, C. Schütte, and E. Vanden-Eijnden, *J. Chem. Phys.* **125**, 084110 (2006).
- ¹⁰⁸F. Noé and S. Fischer, *Curr. Opin. Struct. Biol.* **18**, 154 (2008).
- ¹⁰⁹P. Metzner, C. Schutte, and E. Vanden-Eijnden, *Multiscale Model. Simul.* **7**, 1192 (2009).
- ¹¹⁰C. Dellago, P. G. Bolhuis, and P. L. Geissler, *Adv. Chem. Phys.* **123**, 1 (2002).
- ¹¹¹J.-H. Prinz, M. Held, J. C. Smith, and F. No, *Multiscale Model. Simul.* **9**, 545 (2011).
- ¹¹²A. Levandovsky and R. Zandi, *Phys. Rev. Lett.* **102**, 198102 (2009).
- ¹¹³S. D. Hicks and C. L. Henley, *Phys. Rev. E* **74**, 031912 (2006).
- ¹¹⁴D. J. Wales, *Philos. Trans. R. Soc. A* **363**, 357 (2005).
- ¹¹⁵I. Tsvetkova, C. Chen, S. Rana, C. C. Kao, V. M. Rotello, and B. Dragnea, *Soft Matter* **8**, 4571 (2012).
- ¹¹⁶J. Sun, C. DuFort, M. C. Daniel, A. Murali, C. Chen, K. Gopinath, B. Stein, M. De, V. M. Rotello, A. Holzenburg, C. C. Kao, and B. Dragnea, *Proc. Natl. Acad. Sci. U.S.A.* **104**, 1354 (2007).
- ¹¹⁷X. Huang, L. M. Bronstein, J. Retrum, C. Dufort, I. Tsvetkova, S. Anigayei, B. Stein, G. Stucky, B. McKenna, N. Remmes, D. Baxter, C. C. Kao, and B. Dragnea, *Nano Lett.* **7**, 2407 (2007).
- ¹¹⁸S. K. Dixit, N. L. Goicochea, M. C. Daniel, A. Murali, L. Bronstein, M. De, B. Stein, V. M. Rotello, C. C. Kao, and B. Dragnea, *Nano Lett.* **6**, 1993 (2006).
- ¹¹⁹M.-C. Daniel, I. B. Tsvetkova, Z. T. Quinkert, A. Murali, M. De, V. M. Rotello, C. C. Kao, and B. Dragnea, *ACS Nano* **4**, 3853 (2010).
- ¹²⁰C. Chen, E. S. Kwak, B. Stein, C. C. Kao, and B. Dragnea, *J. Nanosci. Nanotechnol.* **5**, 2029 (2005).
- ¹²¹C. Chen, M. C. Daniel, Z. T. Quinkert, M. De, B. Stein, V. D. Bowman, P. R. Chipman, V. M. Rotello, C. C. Kao, and B. Dragnea, *Nano Lett.* **6**, 611 (2006).
- ¹²²M. F. Hagan, *Phys. Rev. E* **77**, 051904 (2008).
- ¹²³M. F. Hagan, *J. Chem. Phys.* **130**, 114902 (2009).
- ¹²⁴O. M. Elrad and M. F. Hagan, *Nano Lett.* **8**, 3850 (2008).

- ¹²⁵A. Siber, R. Zandi, and R. Podgornik, *Phys. Rev. E* **81**, 051919 (2010).
- ¹²⁶L. He, Z. Porterfield, P. van der Schoot, A. Zlotnick, and B. Dragnea, *ACS Nano* **7**, 8447 (2013).
- ¹²⁷A. Branka and D. Heyes, *Phys. Rev. E* **60**, 2381 (1999); D. Heyes and A. Branka, *Mol. Phys.* **98**, 1949 (2000).
- ¹²⁸D. Endres and A. Zlotnick, *Biophys. J.* **83**, 1217 (2002).
- ¹²⁹M. F. Hagan and O. Elrad, *Biophys. J.* **98**, 1065 (2010).
- ¹³⁰Convergence required a 50% increase in total simulation time because statistics were reduced by the increased number of states (18000 states with the graph compared to 220 with the simple state definition). However, this increase could be at least partially eliminated through lumping of microstates.
- ¹³¹R. F. Garmann, M. Comas-Garcia, A. Gopal, C. M. Knobler, and W. M. Gelbart, *J. Mol. Biol.* **426**, 1050 (2013).
- ¹³²J. van Zon and P. ten Wolde, *J. Chem. Phys.* **123**, 234910 (2005).
- ¹³³Note that characterization of the full time course of assembly kinetics for a system in which subunits are the limiting reagent will require multiple simulations at different subunit concentrations.
- ¹³⁴J. G. de la Torre and B. Carrasco, *Biopolymers* **63**, 163 (2002).
- ¹³⁵S. Fortin, "The graph isomorphism problem," Technical Report No. 96-20 (Department of Computer Science, University of Alberta, 1996).
- ¹³⁶N. Reingold and Deo, *Combinatorial Algorithms: Theory and Practice* (Prentice Hall, 1977).
- ¹³⁷B. D. McKay, *Practical Graph Isomorphism* (Congressus Numerantium, 1981).
- ¹³⁸E. M. Luks, *J. Comput. Syst. Sci.* **25**, 42 (1982).
- ¹³⁹A. Zlotnick, *J. Mol. Biol.* **366**, 14 (2007).
- ¹⁴⁰P. Moisan, H. Neeman, and A. Zlotnick, *Biophys. J.* **99**, 1350 (2010).
- ¹⁴¹T. E. Ouldridge, A. A. Louis, and J. P. K. Doye, *J. Phys.: Condens. Matter* **22**, 104102 (2010).