

Pathway complexity of model virus capsid assembly systems

Navodit Misra^a, Daniel Lees^b, Tiequan Zhang^c and Russell Schwartz^{c*}

^aDepartment of Physics, Carnegie Mellon University, Pittsburgh, PA, USA; ^bDepartment of Computer Science, Amherst University, Amherst, MA, USA; ^cDepartment of Biological Sciences, Carnegie Mellon University, Pittsburgh, PA, USA

(Received 21 January 2008; final version received 29 April 2008)

As computational and mathematical studies become increasingly central to studies of complicated reaction systems, it will become ever more important to identify the assumptions our models must make and determine when those assumptions are valid. Here, we examine that question with respect to viral capsid assembly by studying the ‘pathway complexity’ of model capsid assembly systems, which we informally define as the number of reaction pathways and intermediates one must consider to accurately describe a given system. We use two model types for this study: ordinary differential equation models, which allow us to precisely and deterministically compare the accuracy of capsid models under different degrees of simplification, and stochastic discrete event simulations, which allow us to sample use of reaction intermediates across a wide parameter space allowing for an extremely large number of possible reaction pathways. The models provide complementary information in support of a common conclusion that the ability of simple pathway models to adequately explain capsid assembly kinetics varies considerably across the space of biologically meaningful assembly parameters. These studies provide grounds for caution regarding our ability to reliably represent real systems with simple models and to extrapolate results from one set of assembly conditions to another. In addition, the analysis tools developed for this study are likely to have broader use in the analysis and efficient simulation of large reaction systems.

Keywords: capsid assembly; simulation; stochastic; differential equation; pathway

1. Introduction

The study of viral capsids has attracted considerable attention from the fields of mathematical and computational modelling, in large part because of capsid assembly’s value as a general model of the self-assembly of large macromolecular complexes. Capsids are remarkable for the size, complexity and diversity of the structures they build, as well as the high efficiency and fidelity of their assembly processes. They therefore serve as an important test of our ability to understand and predict self-assembly dynamics in both biological and artificial systems. By learning the principles by which viral capsids assemble so effectively, we hope to learn not only about viral biology, but also about how similar principles can apply to the many simpler examples of self-assembly in biology. Likewise, by understanding these principles, we may learn how to apply them to achieve similarly remarkable behaviours in self-assembling nanotechnology [15,29,31]. There is currently no feasible experimental method to directly observe a rapid, nanometer-scale assembly process such as that involved in building a capsid. Capsid assembly has thus been an important model for numerous simulation studies, where they

*Corresponding author. Email: russells@andrew.cmu.edu

have helped us elucidate general principles that may guide self-assembly [1,4,28,35], map out parameter spaces of possible assembly products and pathways [6,8,10,11,14,21,26], and perform model-based interpretation of indirect experimental measures of assembly progress in real virus systems [3,9,38,39].

One intriguing observation to emerge from many of these computational studies is the fact that a single assembly system can exhibit very different assembly pathways under only moderately different assembly conditions. Several simulation studies of capsid assembly have suggested that changing parameter values (*e.g.*, binding free energies, concentrations or configurational tolerances of binding) so as to promote more rapid assembly can abruptly shift a system from a productive nucleation-limited assembly pathway to an unproductive pathway dominated by kinetically trapped incomplete structures [4,14,17,18,20,21,32,34]. Other studies have shown that similarly small changes in assembly conditions can shift pathways so as to alter the morphology of final assembled structures [1,14,17,18,21]. Experimental support from diverse real viral assembly systems supports these simulation predictions, showing that modest changes in assembly conditions can indeed shift capsid assembly from productive to kinetically trapped assembly [4,13,16,24,36], or induce malformations or other altered morphologies [5,12,19,22,37]. Recent modelling work by Sweeney *et al.* [26] has further suggested that even a single end state may be reached by very different pathways depending on small changes in local binding parameters. They showed that a simple, two-parameter $T = 1$ icosahedral model could exhibit three prominent pathways depending on parameter choices: one proceeding by a fifth-order pentamer nucleation event followed by elongation through successive monomer additions (monomer pathway); a second proceeding by a rapid accumulation of dimers, a third-order nucleation through a trimer of dimers and elongation by dimer additions (dimer pathway); and a third proceeding by rapid accumulation of pentamers, nucleation through a trimer of pentamers, and elongation by pentamer additions (pentamer pathway). Modest shifts in rates or concentrations of subunits, within the ranges accessible to *in vitro* experiments, were sufficient to shift the system from one pathway to another or to a kinetically trapped region in which all three pathways converge.

All of these results point to a common conclusion: assembly pathways need not be an inherent property of a viral assembly system. Rather, the specific pathways used by a capsid system can shift quite dramatically in response to relatively small changes in assembly conditions. This conclusion has important implications for the study of viral assembly and for its use as a model system for complex self-assembly in general. First, it suggests that basic assembly mechanisms determined from *in vitro* capsid assembly systems could be quite different from those used by the same systems *in vivo*, a potentially serious problem given that we currently have a very limited ability to monitor capsid assembly *in vivo*. Accurately determining the assembly pathways of a given system is important not only as a basic research question, but also potentially for attempts to manipulate or interfere with those pathways, as in current attempts to develop capsid-targeted antiviral drugs [23,25,27,36]. Second, the sensitivity of pathway choice has implications for how we can model assembly mathematically and computationally. Capsid assembly is a complicated and computationally expensive process to model and there are thus tremendous advantages to working with highly simplified models of the process, provided such models are accurate. Endres and Zlotnick [6] have shown that just a few pathways are often adequate to describe almost perfectly the behaviour of a simple capsid model with hundreds of pathways in principle accessible to it. Likewise, mathematical models of capsid assembly based on nucleation theory [30] provide a powerful tool for understanding and predicting capsid behaviour from first principles, provided assembly is under conditions allowing for a simple nucleation/elongation pathway. Studies from stochastic models have suggested that even for complicated capsid assembly models, a large fraction of the accessible

parameter space will fall into one of a small number of distinct parameter domains each explicable by a simple pathway model [6,11,14,26].

These studies have, however, left open the possibility that pathway complexity might rise significantly outside of, or between, these discrete domains. It may be the case that as we vary parameters to move between distinct simple pathway domains, we experience rapid ‘phase transitions’, with simple models providing accurate explanations for almost all parameter values. Or it may be the case that the parameter space contains broad border regions between the domains, in which simple pathway models break down and with them some of our most powerful simulation and analysis tools. In the present work, we attempt to distinguish between these two possibilities by examining how pathway complexity varies as we move across the assembly parameter space. We apply two complementary modelling approaches. First, we use the ordinary differential equations (ODE) framework developed by Zlotnick and Endres [6,35] to examine the trade-off between model complexity and accuracy in a fully deterministic, but relatively simplified setting. This model allows us to follow Endres and Zlotnick in asking rigorously how effectively one can prune the pathway set from a model without substantially affecting overall reaction progress. We also conduct a parallel study using stochastic capsid models [33], which allow us to drop some assumptions of the ODE models and look at a substantially larger model system, but at the cost of limiting our analysis tools. By examining intermediate distributions in this model, we distinguish regions of parameter space that appear well described by simple nucleation–elongation models from those that appear unlikely to be describable without an extremely large number of distinct pathways and reaction intermediates. Both lines of inquiry support the conclusion that the ability of a system to be described by a compact model can vary quite dramatically given relatively small parameter changes, suggesting that the notion of reaction pathways is not as cleanly defined for complicated assembly systems as we might hope. Our results provide guidance for issues of model selection and model-based interpretation of experimental results for capsid assembly and other complicated self-assembly systems. In addition, the analysis tools developed for this study may have broader applicability in efficiently modelling complex reaction systems in general.

2. Models and methods

We applied two complementary modelling methods to examine the variability of pathway complexity across assembly parameter spaces. ODE models [35] allow us to exhaustively determine fractional pathway usage, but at the cost of requiring highly simplified capsid models and limiting allowed assembly pathways. Stochastic simulation algorithm (SSA) models [7] allow us to perform a random sample of pathways accessible to a far more complicated model system, but their high computational cost and the stochastic noise in their results limit our ability to perform the kind of exhaustive analysis possible in the ODE case. This section describes each model type in turn and explains how it was used to examine pathway complexity across the parameter space.

2.1 ODE models

We began our analysis, using an ODE model based on the approach of Endres and Zlotnick [6]. As in that work, we used a simplified dodecamer capsid model, which can be considered a model of $T = 1$ assembly from pentameric capsomer subunits. We applied two variants of this model. In the first model, the dodecahedral shell is assumed to assemble only through the addition of single subunits per reaction, each representing one of the 12 faces of the dodecamer. This assumption of assembly by successive monomer additions is commonly used to counter the

rapid proliferation in possible oligomer/oligomer binding pathways with structure size, which would otherwise make the model computationally infeasible. A heuristic justification is provided by the argument that in typical self-assembly systems the equilibrium concentration of monomers is much larger than that of the intermediates. However, as has been shown recently using stochastic methods, there are regions of parameter domain where the intermediate concentrations can far exceed their equilibrium values during the assembly process. Therefore, we used a second model, which allowed dimer/oligomer binding also in addition to monomer/oligomer binding. This allows us to assess the effects of adding more pathways on the self-assembly kinetics. The distinction is illustrated in Figure 1. Figure 1(a) shows an example of a monomer addition reaction, where a monomer and an octomer bind to form a 9-mer. This reaction would be allowed in either model. Figure 1(b) shows a dimer reaction, with a monomer and a septamer binding to form a 9-mer. This reaction is allowed in the dimer/oligomer pathway set, but not in the monomer/oligomer pathway set. Figure 1(c) shows binding of a trimer to a hexamer to build a 9-mer; this reaction would not be allowed in either of the ODE pathway sets we consider. With either pathway set, we assume that the forward rate constants are identical for all reactions, in effect ignoring the relative variation in diffusion coefficient with size of the oligomer.

2.1.1 Defining the ODE pathway set

Each species in the assembly tree is identified by two indices (size, type), where size is the number of subunits in the species and type is an arbitrary assignment used to distinguish two sister species. Let us represent the molar concentrations of each species (\mathbf{j}, \mathbf{k}) as $[j, k]$ and the degeneracy of forward reaction (monomer/dimer addition) between (\mathbf{j}, \mathbf{k}) and (\mathbf{m}, \mathbf{n}) as $a_{j,k}^{m,n}$. Similarly, the backward reaction degeneracy is represented by $b_{m,n}^{j,k}$. $a_{j,k}^{m,n}$ and $b_{m,n}^{j,k}$ capture the reaction degeneracies produced by symmetries of the structures, a topic explained in greater detail when we describe our protocol for testing graph isomorphism below. The relative stability of each species is determined by the Arrhenius factor $\exp(-G/RT)$, where G is the molar free

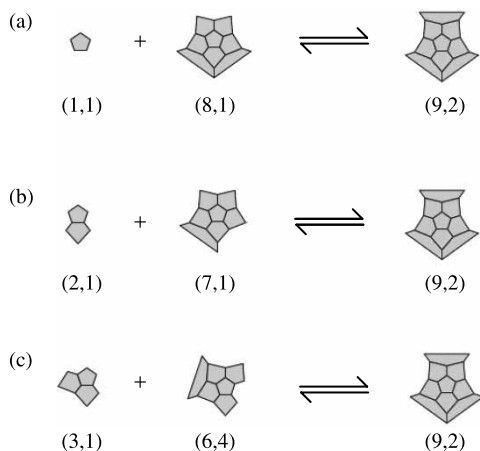


Figure 1. An example of three out of several possible ways to form the species (9,2) in the ODE dodecamer model. (a) Monomer binding to the 8-mer (8,1). (b) Dimer binding to the 7-mer (7,1). (c) Trimer (3,1) binding to the 6-mer (6,4). The two ODE models implemented in this paper allow either only pathway [(a); monomer model] or both [(a) and (b); monomer/dimer model]. Pathway (c) is disallowed in both the models. An analogous dodecamer model conducted with the SSA simulator described in Section 2.2 would have allowed all three pathways, as well as many other possible oligomer/oligomer reactions.

energy of the structure, R is the gas constant and T is the absolute temperature. Our model assumes that each bond contributes the same amount of molar free energy, denoted by ΔG . If species (\mathbf{j}, \mathbf{k}) has $c_{j,k}$ contacts or bonds, the relative stability of (\mathbf{j}, \mathbf{k}) with respect to (\mathbf{m}, \mathbf{n}) is

$$s_{m,n}^{j,k} = \exp(-\Delta G^*(c_{m,n} - c_{j,k})/RT).$$

The forward reaction rate also depends on the symmetry of the binding unit (monomer or dimer), denoted by $O(n)$, $\{n = 1, 2\}$. The constant forward reaction rate k_{on} is then irrelevant to the kinetics and only sets the overall time scale. In terms of these quantities the set of kinetic equations can be written as:

$$\begin{aligned} \frac{d[j, k]}{dt} = & k_{on} \sum_{m,n} \left(b_{m,n}^{j,k} s_{m,n}^{j,k} [m, n] - a_{j,k}^{m,n} O(m-j)[j, k][m-j] \right) \\ & - k_{on} \sum_{p,q} \left(b_{j,k}^{p,q} s_{j,k}^{p,q} [j, k] - a_{p,q}^{j,k} O(j-p)[j-p][p, q] \right), \end{aligned} \quad (1)$$

where $(\mathbf{m} - \mathbf{j})$ represents either the monomer or the dimer.

The primary task in constructing the assembly tree is identifying all *distinct* intermediate oligomers and computing the forward and backward degeneracies for each reaction pair. Since the dodecahedral structure can be projected onto a unit sphere as 12 symmetrically placed points, each of the 12 sites are identified by their spherical-polar coordinates $\{\theta_i, \phi_i\}$ and nearest neighbour locations can be stored as an incidence matrix. Our procedure represents graphs for any species G by a density distribution: $\rho_G(\theta, \phi) = \sum_{i \in G} \delta(\theta - \theta_i) \delta(\phi - \phi_i)$ where $\{\theta_i, \phi_i\}$ is the representation of the i th site. In graph theoretical terms, each intermediate can be represented by a planar graph. Our representation of binding sites suggests a natural definition for distinguishing species, in analogy with rigid body transformations: if two species can be translated over the unit sphere and rotated about their centre of mass in such a manner that their configurations overlap then they are non-distinct. We iteratively construct the state space by successively adding a subunit to each available binding site of each species identified and testing the resulting oligomers for isomorphism. Each time a species is repeated, the forward degeneracy for the corresponding reaction is incremented by one. We are now in a position to give a precise definition of which graphs are identical:

Definition 1. Two graphs G_1 and G_2 are *non-distinct* iff $\exists \tau \in SO(3)$ such that the corresponding density functions satisfy: $\rho_{G_1}(\mathbf{r}) = \rho_{G_2}(\tau \mathbf{r})$.

Here, τ stands for both a member of $SO(3)$ and its 3D representation in terms of the Euler matrices.

We have devised an algorithm to test planar graphs for isomorphism under this restricted set of transformations. The algorithm works in worst case time, $O(N^2 * \min\{N, N_{\text{capsid}}^{1/2}\})$, but takes $O(N * \min\{N, N_{\text{capsid}}^{1/2}\})$ steps for most pairs. We now give an outline of the algorithm used:

- (1) Compute the centroid of this density distribution:

$$\begin{aligned}\langle \theta \rangle_G &\equiv \arccos \langle \cos \theta \rangle_G = \arccos \frac{\sum_{i \in G} \cos \theta_i}{N} \\ \langle \phi \rangle_G &\equiv \arctan \frac{\langle \sin \theta \sin \phi \rangle_G}{\langle \sin \theta \cos \phi \rangle_G} = \arctan \frac{\sum_{i \in G} \sin \theta_i \sin \phi_i}{\sum_{i \in G} \sin \theta_i \cos \phi_i}.\end{aligned}$$

- (2) Given the orientation of the centroid $\mathbf{n}_{G_i} = \{\langle \theta \rangle_{G_i}, \langle \phi \rangle_{G_i}\}$ and their radial magnitudes $\langle r \rangle_{G_i}; i = \{1, 2\}$, check a necessary condition for isomorphism: $\langle r \rangle_{G_1} = \langle r \rangle_{G_2}$.
- (3) If $\langle r \rangle_{G_1} = \langle r \rangle_{G_2} = 0$, check the graphs separately (using a routine to be described later). Assuming these are non-vanishing, perform *Euler* rotations $\tau_i, i = \{1, 2\}$ on both graphs, such that the centroid coincides with the north pole or $\theta = 0, \phi = 0$ point on the sphere. If G_1 and G_2 are *non-distinct* as in Definition 1, the transformed density distributions ρ_{G_i} can only differ under the group of transformations S such that $\forall \tau \in S$

$$\langle \tau \mathbf{n}_{G_1} \rangle = \langle \mathbf{n}_{G_1} \rangle = \langle \mathbf{n}_{G_2} \rangle. \quad (2)$$

This is just the $SO(2)$ subgroup of $SO(3)$. In terms of polar coordinates, this implies that the graphs differ only in rotations around the Z -axis.

- (4) Since any $\tau \in SO(2)$ leaves the $\{\theta_i\}, i \in G$ unchanged, the lists $\{\theta_{i_1}\}$ and $\{\theta_{i_2}\}$ ($i_1 \in G_1$ and $i_2 \in G_2$) are identical after sorting. The corresponding ϕ_i , however, need not be, since there might exist sites $i, j \in G$ such that $\theta_i = \theta_j$ and $\phi_i \neq \phi_j$. This degeneracy d_θ is maximal for $\theta_i = \pi/2$, in which case $d_\theta \lesssim \sqrt{N_{\text{capsid}}}$, since

$$\frac{d_\theta}{N_{\text{capsid}}} \lesssim \frac{\text{circumference}}{\text{surface area}} \sim \frac{1}{\text{radius}} \sim \frac{1}{\sqrt{N_{\text{capsid}}}}.$$

- (5) If we choose any $\theta = \theta_i, i \in G_1$, then assuming it has a degeneracy d_θ , it has d_θ possible candidates for its image in G_2 . Sequentially, test each of these cases. Let the set of possible images of i be $\mathbb{l}_i := \{j_1, \dots, j_{d_\theta}\} \subset G_2$. Perform a rotation $\mathbb{R}_{i,j}$ on G_2 , by an angle $\delta\phi_j := \phi_i - \phi_j$, for each $j \in \mathbb{l}_i$. If j is the image of i then the transformed Graph \tilde{G}_2 fulfils the identity $\tilde{G}_2 = \mathbb{R}_{i,j}G_2 = G_1$. If this condition is not met for any $j \in \mathbb{l}_i$ then the graphs cannot be related by an element of $SO(3)$ and hence by Definition 1 are distinct.
- (6) We now address the exceptional case of graphs where the centroid is located at the origin. Here, we make use of the fact that if two graphs of size N ($G_{i,N}$ with $i = \{1, 2\}$) are isomorphic then their parent graphs $G_{1,N-1}$ and $G_{2,N-1}$ must differ in only two sites (call them b_1 and b_2):

$$G_{i,N} = G_{i,N-1} \cup \{b_i\}, \quad (3)$$

for $i = \{1, 2\}$. If $G_{1,N-1}$ and $G_{2,N-1}$ are *non-distinct*, then $G_{1,N}$ and $G_{2,N}$ are trivially so. Consider the case where the parent graphs are distinct. Referring back to our Definition 1, the graphs are isomorphic *iff* $\exists \tau \in SO(3)$ such that:

$$G_{1,N} = \mathbb{R}_\tau G_{2,N}, \quad (4)$$

which implies that $\exists \nu_1 \in G_{1,N-1}$ and $\nu_2 \in G_{2,N-1}$ such that

$$\mathbb{R}_\tau^{-1} \nu_1 = b_2, \quad (5)$$

$$\mathbb{R}_\tau \nu_2 = b_1. \quad (6)$$

Search for this transformation τ by searching for the image of b_1 in $G_{2,N}$ and of b_2 in $G_{1,N}$. To this end, sequentially choose $\nu \in G_{1,N} - \{b_1\}$ and construct a new graph $G_\nu = G_{1,N} - \{\nu\}$. Then, $G_{i,N}$ are isomorphic iff G_ν and $G_{2,N-1}$ are isomorphic for some $\nu \in G_{1,N}$. Since the centroid of either of these does not coincide with the origin, we can check them for isomorphism by the usual procedure.

Once all the distinct species of a given size are identified, the backward reaction degeneracy can be computed using the principle of *detailed balance*. For example, monomer additions follow

$$\frac{b_{j,k}^{j-1,n}}{O(j,k)} = \frac{a_{j-1,n}^{j,k}}{O(j-1,n)}, \quad (7)$$

where $O(j,k)$ is the order of the symmetry group of species (\mathbf{j}, \mathbf{k}) . A simple way to compute $O(j,k)$ is to sum over all n in the previous equation. The right-hand side of the equation is already known in terms of the symmetry groups and forward degeneracies of the parent, while the left-hand side is $1/O(j,k)$ times the total number of ways (\mathbf{j}, \mathbf{k}) can decay, $\sum_n b_{j,k}^{j-1,n}$, which equals the size \mathbf{j} minus the number of articulation points of (\mathbf{j}, \mathbf{k}) . We used depth-first search to compute the articulation points, which eventually allows us to compute each individual degeneracy $b_{j,k}^{j-1,n}$ using Equation (7). An analogue of the previous equation for dimer–oligomer reactions allows us to compute $b_{j,k}^{j-2,n}$ from $a_{j-2,n}^{j,k}$.

A subset of the experiments were run using a pruned version of the ODE tree. We used a landscape approach similar to that of Endres and Zlotnick [6] to prune the assembly tree. The controlling parameter for the identification of intermediates to prune is the following probability-like quantity:

$$P(j,k) = \frac{\sum_l a_{j-1,l}^{j,k} P(j-1,l) \mu_{j-1,l}^{j,k}}{\sum_{k,l} a_{j-1,l}^{j,k} P(j-1,l) \mu_{j-1,l}^{j,k}}, \quad (8)$$

where $\mu_{j-1,l}^{j,k}$ is a control parameter used to assign relative weights to different reactions. One choice that, though *ad hoc*, has proven useful is to choose $\mu = 0.1$ for reactions that proceed by a single subunit addition and 1 otherwise. We have tried to use a similar choice, but for a broader concentration regime, specifically those where the probability distribution among the intermediate states is expected to be more uniform.

2.1.2 ODE experimental design

We first carried out a series of experiments to determine how the efficiency of landscape pruning varies with the initial monomer concentration for a simple pathway model. These experiments were conducted allowing for monomer accretion reactions, such as that in Figure 1(a), but excluding dimer accretion reaction such as that in Figure 1(b), as well as interactions of higher-order oligomers as in Figure 1(c). Simulations were conducted for eight values of the pruning parameter: 1 (unpruned), 0.999, 0.995, 0.987, 0.975, 0.95, 0.90 and 0.75. These values yield, respectively, 73, 65, 55, 47, 39, 34, 29 and 21 intermediate states, out of the 73 present in the full model. These values were repeated for three concentrations – 10 μM , 100 μM and 1 mM – intended to approximately span the range from the lowest concentrations likely to produce assembly *in vitro* to the largest concentrations likely to be found *in vivo*. All simulations were conducted assuming a free energy of binding of $\Delta G = -3.5$ kcal/mole. Simulations were

carried out by numerically integrating the ODE models using an adaptive step size embedded Runge–Kutta 4–5 method with relative error = 10^{-9} .

We next conducted simulations to estimate the contribution of oligomer/oligomer interactions to pathway kinetics. For each of the following five concentration levels – 10 μM , 100 μM , 200 μM , 500 μM and 1 mM – we conducted an additional simulation permitting reactions of dimers with arbitrary oligomers. That is, the simulations permit the reactions of both Figure 1(a) and (b), although, they still exclude higher-order reactions such as those of Figure 1(c). These simulations were likewise conducted with a free energy of binding of $\Delta G = -3.5$ kcal/mole and were carried out by numerically integrating the ODE models using an adaptive step size embedded Runge–Kutta 4–5 method with relative error = 10^{-9} .

For each of the 24 comparisons of pruned to unpruned models and five comparisons of monomer to monomer/dimer models, we examined the deviations in capsid concentrations vs. time between each pruned model and the full model. To perform this comparison, we define the *quality factor* (Q_r), a measure of error of the restricted model r with respect to the complete model c , as follows:

$$Q_r = \left(\frac{1}{T} \int_0^T \left(\frac{[\text{capsid}]_r - [\text{capsid}]_c}{[\text{capsid}]_c} \right)^2 dt \right)^{1/2}. \quad (9)$$

The quality factor is meant to capture the average fractional deviation between the two models across the time to equilibration. We use a relative measure of deviation normalized by capsid count to account for the fact that deviations between the models tend to have a large proportional difference, but a small absolute difference, early in their execution and to converge to the same equilibrium late in their execution. Q_r provides a sensitive measure of these large proportional changes early in simulations that would otherwise be difficult to detect. In the case of landscape pruning, the eventual equilibrium concentration depends on the model used, although, the difference is negligible under the experimental conditions examined here. In such cases, we defined the cut-off time T as the point at which the concentrations for the complete model had equilibrated to within 2%, *i.e.*, $([\text{capsid}]_{\text{eqb}} - [\text{capsid}]_c)/[\text{capsid}]_{\text{eqb}} \leq 0.02$, where $[\text{capsid}]_{\text{eqb}}$ is the equilibrium concentration of the fully formed capsid. This cut-off definition allows us to compare Q_r values for different pruning schemes.

2.2 Stochastic models

Stochastic simulations were conducted using the simulator of Zhang *et al.* [33], which implements a SSA model [7] of capsid assembly. The SSA model allows us to test the scaling of pathway control to larger structures and larger pathway spaces than we can explore with the ODE model. In an SSA model, the assembly system proceeds by the activation of successive reaction events. For the models constructed here, each reaction event consists of either the binding of two structures to form a larger structure or the breakage of a structure into two smaller substructures. Binding is controlled by a local rule model [1], which specifies allowed reactions through a set of ‘local rules’ that define how each subunit can bind to its neighbours in any capsid or substructure. The model allows for any binding reaction of either monomers or oligomers provided the product of the reaction is a subset of a correctly formed capsid. So, for example, an analogous SSA model of a dodecamer would allow all of the reactions depicted in Figure 1(a)–(c). The model does not, however, allow for the production of malformed structures. Reverse reactions are allowed by breaking any single bond that results in a separation of a current assembly into exactly two sub-assemblies. The simulator disallows reverse reactions that would require the simultaneous breaking of two or more binding interactions.

For the present simulations, we used a 60-subunit icosahedral model originally described in Ref. [34], representing a full $T = 1$ capsid assembly system at the monomer level. This model constructs icosahedrally symmetric 60-mers from an assembly subunit with three binding sites, two that bind to copies of one another to form pentamers of coat protein (capsomers) and one that binds copies of itself to link together capsomers to form the complete icosahedron. We refer to the former as intra-capsomer interactions and the latter as inter-capsomer interactions. We chose this model in part because it allows us to explore a model size beyond what we can feasibly simulate with the ODE approach and in part because the asymmetry of the assembly subunits makes it possible to independently vary two binding free energies to produce distinct productive assembly pathways, a capability that was previously exploited in Sweeney *et al.* [26].

All icosahedral assembly simulations were run with a fixed population of 6000 coat monomers. Binding kinetics in the model are parameterized by four rate constants: a forward binding rate for binding interactions within pentamers (the intra-capsomer association rate, k_{a+}), a reverse binding rate for interactions within pentamers (the intra-capsomer dissociation rate, k_{a-}), a forward binding rate for binding interactions between pentamers (the inter-capsomer association rate, k_{r+}), and a reverse binding rate for interactions between pentamers (the inter-capsomer dissociation rate, k_{r-}). k_{a-} and k_{r-} were fixed at 10^{-3} for all simulations. The two forward rates were varied independently to produce a total of 378 simulations, spanning a range of 18 inter-capsomer binding rates (k_{r+} from $10^{-2.8}$ to $10^{1.2}$ in factors of $10^{0.2}$) and 21 intra-capsomer binding rates (k_{a+} from $10^{-2.6}$ to $10^{0.8}$ in factors of $10^{0.2}$). These values were chosen based on an analysis of data from Sweeney *et al.* [26] indicating that this region would include examples of all three simple pathways (monomer, dimer and pentamer) and extend into the kinetically trapped region in which the three pathways merge with one another. Each simulation was run for sufficient time to reach a pseudo-equilibrium, as verified by manual examination. Numbers of monomers, dimers, pentamers and 60-mers were recorded vs. simulator time across each simulation.

For each of the 378 simulations, we assessed fractional usage of the three pathway types based on the total mass fraction of the major assembly intermediate of each pathway (monomers for the monomer pathway, dimers for the dimer pathway and pentamers for the pentamer pathway), integrated over the course of the simulation. While all three pathways would be expected to exhibit some transient appearance of each of the three key intermediates, each should have a dominant presence of a single intermediate until the pool of small oligomers is exhausted into capsids or large kinetically trapped forms. Integrated mass fractions of monomers (C_1), dimers (C_2) and pentamers (C_5) were computed by the following formulas:

$$C_1 = \sum \frac{(h_{1i} + h_{1(i+1)})(T_{i+1} - T_i)}{T_{50}}, \quad C_2 = \sum \frac{2(h_{2i} + h_{2(i+1)})(T_{i+1} - T_i)}{T_{50}},$$

$$C_5 = \sum \frac{5(h_{5i} + h_{5(i+1)})(T_{i+1} - T_i)}{T_{50}},$$

where h_{ji} is the count of assemblies of size j at time step i and T_i is the simulator time at time step i . T_{50} is the amount of simulator time required to reach 50% of the final yield for a given simulation, which we use as a normalizing constant to correct for varying timescales of different assembly reactions. We assessed fractional usage of the three pathways (P_1 for monomer fraction, P_2 for dimer fraction and P_5 for pentamer fraction) as follows:

$$P_1 = \frac{C_1}{C_1 + C_2 + C_5}, \quad P_2 = \frac{C_2}{C_1 + C_2 + C_5}, \quad P_5 = \frac{C_5}{C_1 + C_2 + C_5}.$$

3. Results

3.1 ODE simulation results

We first examined the complexity of pathway space in the ODE model allowing assembly only by accretion of successive monomers. Figure 2 shows the relationship between quality factor Q_r and degree of pruning for three concentration values. At all concentrations, we observe increasing Q_r and thus, decreasing quality of fit, as we prune increasingly more common intermediates. The effect is more pronounced the higher the concentration, suggesting that in high concentration domains a larger number of intermediates and pathways contribute significantly to the overall assembly kinetics. Note that the absolute magnitude of the deviation is small at all points, consistent with the observations of Endres and Zlotnick [6]. The highest total deviation observed, for a 21-intermediate model at 1 mM concentration, is approximately 4.5%. Furthermore, under all parameter conditions examined here, the final equilibrium concentration was not noticeably different between pruned and full models. Rather, deviations appear to predominantly take the form of large relative shifts in concentrations early in the assembly process. The difference in Q_r values between high and low concentration conditions appears to be predominantly a result of a slower rate of convergence of the pruned models with respect to the unpruned model when concentration is higher, as opposed to a greater amount of separation prior to convergence. Figure 3 illustrates these early deviations, showing the relative deviation as a function of time for a high- Q_r condition (1 mM, Figure 3(b)), with a low- Q_r condition (10 μ M, Figure 3(a)) provided for comparison. Collectively, the data show that highly pruned models do provide very good fits over a broad parameter range, although they often produce large proportional deviations early in assembly. Furthermore, the deviations they produce vary significantly with assembly conditions, increasing approximately 20-fold between

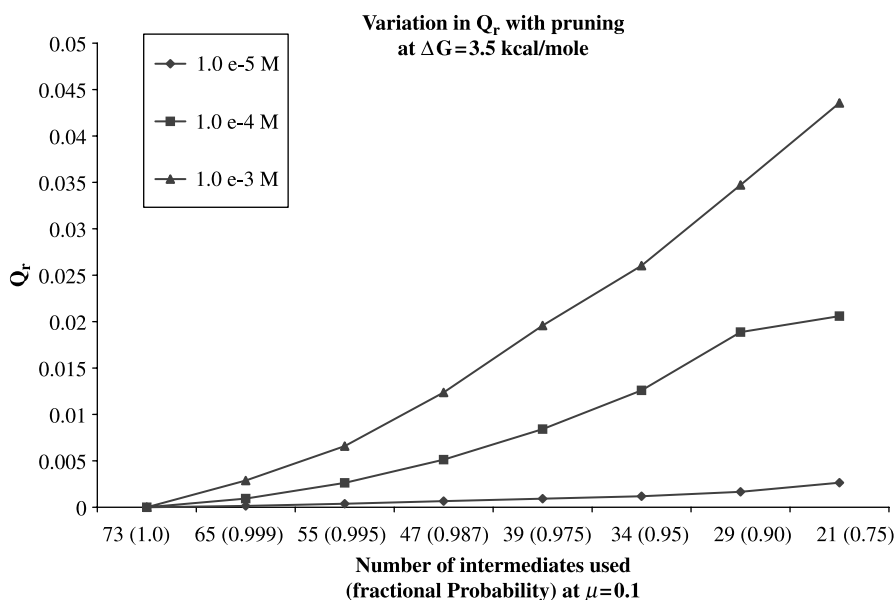


Figure 2. Variation in quality factor Q_r as a function of degree of pruning for the ODE model. The figure shows three curves, representing three simulation concentrations (10 μ M, 100 μ M and 1 mM), with each curve plotting Q_r across a range of pruning parameters. The X-axis is labelled with pruning parameter used, as well as the number of intermediates (from 73 possible) retained with that pruning parameter value.




Figure 3. Examples of fractional error $F_r = ([\text{capsid}]_r - [\text{capsid}]_c)/[\text{capsid}]_c$ as a function of time during the early stages of assembly for full vs. pruned models. Both results were obtained at pruning parameter $\mu = 0.1$ and $\Delta G = 3.5$ kcal/mole (a) Error vs. time for a model exhibiting low- Q_r (high accuracy) defined by concentration 10 μM . (b) Error vs. time for a model exhibiting high- Q_r (low accuracy) defined by concentration 1.0 mM.

a typical *in vitro* assembly concentration (10 μM) and a value not far above likely *in vivo* concentrations (1 mM).

The minimal effect of pruning away the less stable intermediates on the kinetics at late stages of assembly is to be expected as the ODE model can be reduced to a linear ODE model whose

slowest relaxation rate is bounded from above by the decay rate of the most stable intermediate. Introducing unstable intermediates with large decay rates would not affect the slowest relaxation rate appreciably. In that sense, kinetics in the late stages is primarily controlled by the equilibrium concentration. If a pruned model does not change the equilibrium distribution appreciably, it will display similar relaxation kinetics to the complete model. The value of the fully nonlinear ODE model is, then, for studying the kinetics at the early stages. As we examine early time scales, the unstable, fast decaying intermediates contribute more and more to the kinetics and any pruning strategy is bound to be inaccurate. This observation is clearly seen in Figure 3. The effects of pruning at extremely early times (~ 10 units of time as opposed to the time scale over which the system equilibrates, $\sim 10^5$ units of time) are almost identical for the $10\ \mu\text{M}$ and $1\ \text{mM}$ concentrations. The difference lies only in the rate at which the fractional errors decay in time. Note that the actual units of the time scale are arbitrary, since the shape of the curves depends on the ratio of the forward and reverse rates and the time axis can be scaled arbitrarily without altering that ratio.

While, the ODE approach does not allow us to feasibly examine the full ensemble of oligomer/oligomer pathways, we can estimate the effects of pruning out oligomer/oligomer reactions by adding in the subset of dimer/oligomer reactions. Figure 4 shows deviations between simulation progress by monomer reactions only vs. those allowing monomer and dimer reactions as a function of concentration, as assessed by the Q_r measure. The figure shows a rapid rise in Q_r from 10 to $100\ \mu\text{M}$, followed by a more gradual increase across the rest of the concentration range. We attribute the value at $10\ \mu\text{M}$ to the fact that that simulation is below the critical concentration of the system and therefore exhibits minimal growth under either model. The other data points show that pruning of oligomer/oligomer pathways does produce a noticeable error in results, although this error like those above, is small in absolute magnitude and large predominantly at the earliest stages of assembly.

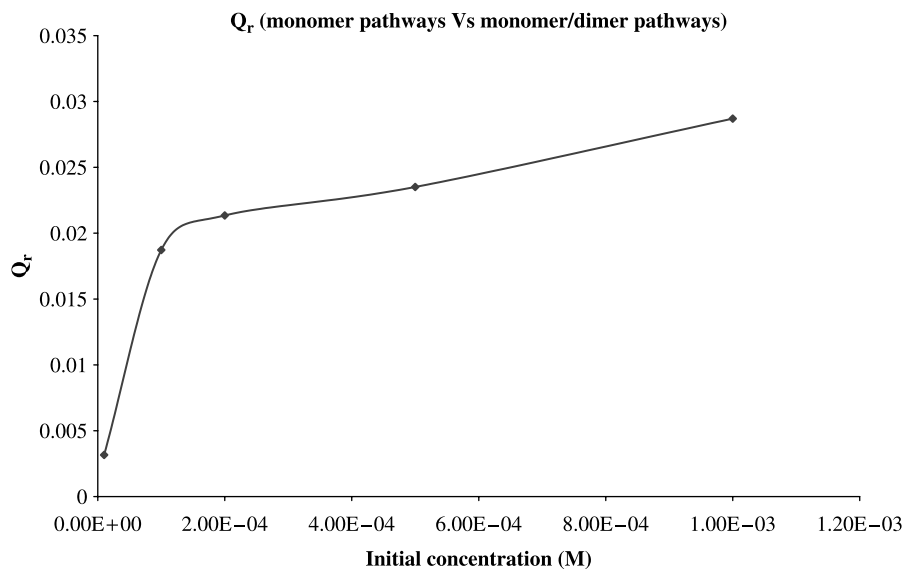


Figure 4. Estimated error induced by eliminating dimer–oligomer reactions as a function of concentration. The figure shows Q_r values plotted vs. concentration for comparisons of an ODE model allowing all 73 intermediates but only monomer-accumulation and loss reactions vs. a model allowing reactions involving accumulation or loss of either monomers or dimers.

3.2 Stochastic simulation results

We next examined how these results would apply to a more complicated but realistic system, defined by a 60-subunit $T = 1$ stochastic model allowing for independent variation of two distinct binding rates. We assessed pathway usage in this system by monitoring concentrations of key intermediates over time across the simulation. Figure 5 shows mass fractions of the three intermediates – monomer, dimer and pentamer – for the $T = 1$ capsid model as we vary intra-capsomer and inter-capsomer binding rates over a parameter range of 3.4 orders of magnitude in inter-capsomer binding rate and 4 orders of magnitude in intra-capsomer binding rate. Figure 5(a) shows the mass fraction of monomers for each point in parameter space, Figure 5(b) the mass fraction of dimers and Figure 5(c) the mass fraction of pentamers.

Collectively, the images confirm the existence of the three distinct pathway regions identified in Sweeney *et al.* [26]. In the bottom left of each image, corresponding to low rates of binding for both interaction types, we see simulations dominated by monomers with minimal appearance of dimers or pentamers, as we would expect for the monomer pathway. In the upper left, corresponding to a high rate of inter-capsomer binding and a low rate of intra-capsomer binding, we see simulations dominated by dimers, as we would expect for the dimer pathway. In the lower right, corresponding to high intra-capsomer binding rate and low inter-capsomer binding rate, we see simulations dominated by the appearance of large numbers of pentamers, as we would expect for the pentamer pathway.

The majority of Figure 5 does, however, appear to show significant hybrid use of multiple pathways. The upper left half of each figure shows significant usage of both monomers and dimers, while the lower right half shows significant usage of both monomers and pentamers. The figure therefore, supports the conclusion that there are sizable border regions between discrete pathway domains in which neither pathway alone provides a valid description of the overall pathway space. We observe a steady, continuous shift in mass fractions of the intermediates as we interpolate between different pathway domains, as opposed to a sudden phase transition. It is therefore difficult to define exact boundaries of the border regions. We can somewhat arbitrarily define a border region to be a region of parameter space in which the dominant intermediate type has less than an 80% mass fraction of all three. Using that definition, we find that the border region between monomer and dimer pathways has a width of approximately 1.5 orders of magnitude change in either binding rate and the border between monomer and pentamer pathways has a width of approximately 1.0 orders of magnitude. The more rapid shift from monomer to pentamer vs. monomer to dimer may reflect the comparatively high effective reaction order of pentamer formation relative to dimer formation. The two borders also have very different slopes, approximately 0.4 for the monomer/dimer border and 1.7 for the monomer/pentamer border. We attribute the 0.4 slope to the competition between second-order dimerization needed by the dimer pathway and fifth order pentamer formation needed for the monomer pathway. The 1.7 slope may be explained by the balance between the fifth-order effective rate of pentamer formation, needed by both monomer and pentamer pathways, vs. the third order formation of trimer-of-pentamer nuclei needed by the pentamer pathway. The nature of the model type makes it computationally infeasible to rigorously determine how many intermediates are actually necessary to the model, as we did in the ODE case. Nonetheless, we can reasonably expect that regions in which two small oligomers coexist at high concentrations during most of the assembly will be characterized by an exponential explosion in the number of distinct reactions being used, due to the numerous ways we can get from nucleus to complete capsid by combinations of additions of two sub-assemblies. The figure thus supports the conclusion that border regions are both sizable and are characterized by high pathway complexity.

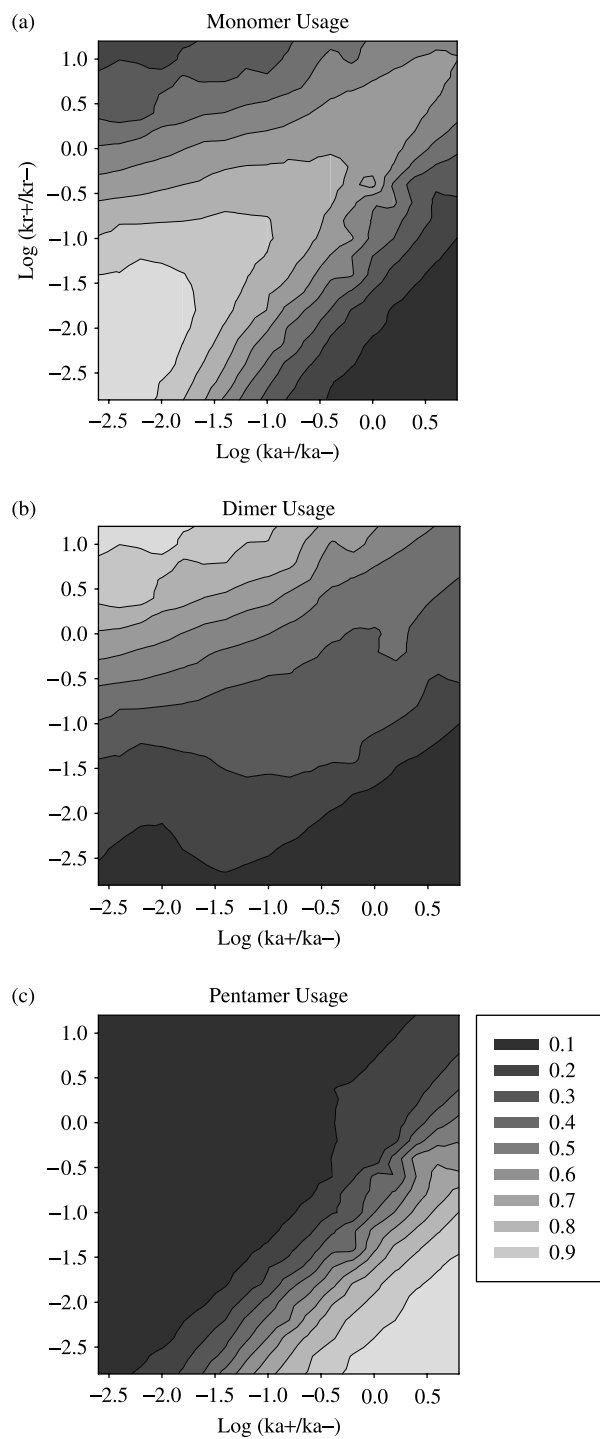


Figure 5. Fractional pathway usage as a function of two binding rates for a 60-mer icosahedral system. (a) P_1 , showing the fractional monomer pathway usage. (b) P_2 , showing the fractional dimer pathway usage. (c) P_5 , showing the fractional pentamer pathway usage.

It is possible that the wide border regions in the figure in part reflect the fact that the simulations were conducted close to the region in which kinetic trapping becomes significant, approximately corresponding to the upper limits of concentrations examined in the ODE simulations. Based on the analysis of Sweeney *et al.*, we would expect this region to correspond approximately to assembly at likely *in vivo* concentrations of subunits (roughly 500 μM). We can observe kinetic trapping becoming prominent in the upper right portions of the images, where all three intermediate types begin to appear with high frequency. We do not observe any appreciable change in the width of the border regions over the parameter range examined here, suggesting that a similarly wide border may extend much further into the low-rate domain than we can feasibly examine. Nonetheless, we cannot be certain that the border region does not become thinner at lower rates more typical of *in vitro* concentrations. Computational resources required for a simulation increase rapidly with reduced binding rate beyond the region shown here because nucleation rates fall with the third order or fifth order of binding rate, depending on the pathway used. A high-precision examination of the parameter space beyond that shown here to more likely *in vitro* values would therefore not be possible with current computational tools.

4. Discussion

We have used a combination of ODE and stochastic models to explore how pathway complexity varies across a space of binding rate parameters for simple models of virus capsid assembly. Simulations reveal that pathway complexity can vary substantially with relatively small changes in model parameters. ODE simulations using a simplified dodecamer model show that the number of intermediates, and thus pathways, we must consider to achieve high model accuracy substantially increases in high vs. low concentration domains. Likewise, the contribution of dimer/oligomer reactions to the assembly kinetics goes up by a significant fraction as concentration varies from levels typically used *in vitro* to levels likely to be feasible at sites of virus assembly *in vivo*. Stochastic models allow us to test how pathway complexity behaves in a more complicated model defined by a 60-mer icosahedron with two independently varying binding rates. In this system, we can map out a pathway space that exhibits not only several distinct simple pathways, but also broad border regions where hybrids of the simple pathways appear to operate. It therefore appears that in a large fraction of the parameter space, no single pathway operates in isolation. Such border regions can be expected to require a large number of distinct reactions in order to adequately describe their kinetics.

This work may also have relevance beyond capsid assembly. Our observations of the ease of shifting pathways between *in vitro* and *in vivo* models may help in interpreting *in vitro* assembly studies of many important self-assembling systems in biology. The same questions are also of relevance in understanding how we might prototype, model and optimize assembly reactions for self-assembling nanotechnology. The methods developed here may also be more broadly applicable to modelling complex reaction systems. We have developed an efficient algorithm for testing planar graph isomorphism for a general class of chemical species that works in time $O(N^2 * \min\{N, N_{\text{capsid}}^{1/2}\})$ and shown its utility for efficiently enumerating the pathway space of a complicated assembly model. This method can be extended to include structures with multiple types of bonds, which is one of the directions for further exploration using the ODE model. A similar method may also prove useful for hybrid discrete/stochastic models, a possible avenue for overcoming some of the limits of each of the two model types examined here.

Acknowledgements

We thank Blake Sweeney for assistance in executing stochastic simulations. This work was supported in part by US NSF award #0346981. D. Lees was supported by the Carnegie Mellon Biological Sciences Summer REU program.

References

- [1] B. Berger, P.W. Shor, L. Tucker-Kellogg, and J. King, *Local rule-based theory of virus shell assembly*, Proc. Natl Acad. Sci. USA 91 (1994), pp. 7732–7736.
- [2] B. Berger, J. King, R. Schwartz, and P. Shor, *Local rule mechanism for selecting icosahedral shell geometry*, Discrete Appl. Math. 104 (2000), pp. 97–111.
- [3] G.L. Casini, D. Graham, D. Heine, R.L. Garcea, and D.T. Wu, *In vitro papillomavirus capsid assembly analyzed by light scattering*, Virology 325 (2004), pp. 320–327.
- [4] P. Ceres and A. Zlotnick, *Weak protein–protein interactions are sufficient to drive assembly of hepatitis B virus capsids*, Biochemistry 41 (2002), pp. 11525–11531.
- [5] W. Earnshaw and J. King, *Structure of phage P22 coat protein aggregates formed in the absence of the scaffolding protein*, J. Mol. Biol. 126 (1978), pp. 721–747.
- [6] D. Endres, M. Miyahara, P. Moisant, and A. Zlotnick, *A reaction landscape identifies the intermediates critical for self-assembly of virus capsids and other polyhedral structures*, Protein Sci. 14 (2005), pp. 1518–1525.
- [7] D.T. Gillespie, *Exact stochastic simulation of coupled chemical reactions*, J. Phys. Chem. 81 (1977), pp. 2340–2361.
- [8] M.F. Hagan and D. Chandler, *Dynamic pathways for viral capsid assembly*, Biophys. J. 91 (2006), pp. 42–54.
- [9] D.E. Kainov, S.J. Butcher, D.H. Bamford, and R. Tuma, *Conserved intermediates on the assembly pathway of double-stranded RNA bacteriophages*, J. Mol. Biol. 328 (2003), pp. 791–804.
- [10] T. Keef, A. Taormina, and R. Twarock, *Assembly models for Papovaviridae based on tiling theory*, Phys. Biol. 2 (2005), pp. 175–188.
- [11] T. Keef, C. Micheletti, and R. Twarock, *Master equation approach to the assembly of viral capsids*, J. Theor. Biol. 242 (2006), pp. 713–721.
- [12] L. Montross, S. Watkins, R.B. Moreland, H. Mamon, D.L. Caspar, and R.L. Garcea, *Nuclear assembly of polyomavirus capsids in insect cells expressing the major capsid protein VP1*, J. Virol. 65 (1991), pp. 4991–4998.
- [13] S.D. Moore and P.E. Prevelige, Jr., *A P22 scaffold protein mutation increases the robustness of head assembly in the presence of excess portal protein*, J. Virol. 76 (2002), pp. 10245–10255.
- [14] H.D. Nguyen, V.S. Reddy, and C.L. Brooks III, *Deciphering the kinetic mechanism of spontaneous self-assembly of icosahedral capsids*, Nano Lett. 7 (2007), pp. 338–344.
- [15] A.J. Olson, Y.H. Hu, and E. Keinan, *Chemical mimicry of viral capsid self-assembly*, Proc. Natl Acad. Sci. USA 104 (2007), pp. 20731–20736.
- [16] K.N. Parent, A. Zlotnick, and C.M. Teschke, *Quantitative analysis of multi-component spherical virus assembly: Scaffolding protein contributes to the global stability of phage P22 procapsids*, J. Mol. Biol. 359 (2006), pp. 1097–1106.
- [17] D.C. Rapaport, *Self-assembly of polyhedral shells: A molecular dynamics study*, Phys. Rev. E 70 (2004), p. 051905.
- [18] D. Rapaport, J. Johnson, and J. Skolnick, *Supramolecular self-assembly: Molecular dynamics modeling of polyhedral shell formation*, Comput. Phys. Comm. 122 (1999), pp. 231–235.
- [19] D.M. Salunke, D.L. Caspar, and R.L. Garcea, *Polymorphism in the assembly of polyomavirus capsid protein VP1*, Biophys. J. 56 (1989), pp. 887–900.
- [20] R. Schwartz, P.E. Prevelige, Jr., and B. Berger, *Local rules modeling of nucleation-limited virus capsid assembly (MIT–LCS–TM-584)*, MIT Laboratory for Computer Science, 1998.
- [21] R. Schwartz, P.W. Shor, P.E. Prevelige, Jr., and B. Berger, *Local rules simulation of the kinetics of virus capsid self-assembly*, Biophys. J. 75 (1998), pp. 2626–2636.
- [22] R. Schwartz, R.L. Garcea, and B. Berger, *'Local rules' theory applied to polyomavirus polymorphic capsid assemblies*, Virology 268 (2000), pp. 461–470.
- [23] J. Sticht, M. Humbert, S. Findlow, J. Bodem, B. Muller, U. Dietrich, J. Werner, and H.G. Krausslich, *A peptide inhibitor of HIV-1 assembly in vitro*, Nat. Struct. Mol. Biol. 12 (2005), pp. 671–677.
- [24] S.J. Stray, P. Ceres, and A. Zlotnick, *Zinc ions trigger conformational change and oligomerization of hepatitis B virus capsid protein*, Biochemistry 43 (2004), pp. 9989–9998.

- [25] S.J. Stray, C.R. Bourne, S. Punna, W.G. Lewis, M.G. Finn, and A. Zlotnick, *A heteroaryldihydroypyrimidine activates and can misdirect hepatitis B virus capsid assembly*, Proc. Natl Acad. Sci. USA 102 (2005), pp. 8138–8143.
- [26] B. Sweeney, T. Zhang, and R. Schwartz, *Exploring the parameter space of complex self-assembly through virus capsid models*, Biophys. J. 94 (2008), pp. 772–783.
- [27] C.M. Teschke, J. King, and P.E. Prevelige, Jr., *Inhibition of viral capsid assembly by 1,1'-bi(4-anilinothalene-5-sulfonic acid)*, Biochemistry 32 (1993), pp. 10658–10665.
- [28] R. Twarock, *A tiling approach to virus capsid assembly explaining a structural puzzle in virology*, J. Theor. Biol. 226 (2004), pp. 477–482.
- [29] G.M. Whitesides and B. Grzybowski, *Self-assembly at all scales*, Science 295 (2002), pp. 2418–2421.
- [30] R. Zandi, P. van der Schoot, D. Reguera, W. Kegel, and H. Reiss, *Classical nucleation theory of virus capsids*, Biophys. J. 90 (2006), pp. 1939–1948.
- [31] S. Zhang, *Fabrication of novel biomaterials through molecular self-assembly*, Nat. Biotechnol. 21 (2003), pp. 1171–1178.
- [32] T. Zhang and R. Schwartz, *Simulation study of the contribution of oligomer/oligomer binding to capsid assembly kinetics*, Biophys. J. 90 (2006), pp. 57–64.
- [33] T. Zhang, R. Rohlf, and R. Schwartz, *Implementation of a discrete event simulator for biological self-assembly systems*, Proceedings of the 2005 Winter Simulation Conference (2005), pp. 2223–2231.
- [34] T. Zhang, W. Kim, and R. Schwartz, *Investigating scaling effects on virus capsid-like self-assembly using discrete event simulations*, IEEE Trans. Nanobioscience 6(3) (2007), pp. 235–241.
- [35] A. Zlotnick, *To build a virus capsid: An equilibrium model of the self assembly of polyhedral protein complexes*, J. Mol. Biol. 241 (1994), pp. 59–67.
- [36] A. Zlotnick and S.J. Stray, *How does your virus grow? Understanding and interfering with virus assembly*, Trends Biotechnol. 21 (2003), pp. 536–542.
- [37] A. Zlotnick, N. Cheng, J.F. Conway, F.P. Booy, A.C. Steven, S.J. Stahl, and P.T. Wingfield, *Dimorphism of hepatitis B virus capsids is strongly influenced by the C-terminus of the capsid protein*, Biochemistry 35 (1996), pp. 7412–7421.
- [38] A. Zlotnick, J.M. Johnson, P.W. Wingfield, S.J. Stahl, and D. Endres, *A theoretical model successfully identifies features of hepatitis B virus capsid assembly*, Biochemistry 38 (1999), pp. 14644–14652.
- [39] A. Zlotnick, R. Aldrich, J.M. Johnson, P. Ceres, and M.J. Young, *Mechanism of capsid assembly for an icosahedral plant virus*, Virology 277 (2000), pp. 450–456.