

Persistent voids: a new structural metric for membrane fusion

Peter M. Kasson^{1,2,*}, Afra Zomorodian³, Sanghyun Park², Nina Singhal⁴,
Leonidas J. Guibas⁴ and Vijay S. Pande^{2,*}

¹Medical Scientist Training Program, ²Department of Chemistry, Stanford University, Stanford CA 94305,

³Department of Computer Science, Dartmouth College, Hanover NH 03755 and ⁴Department of Computer Science, Stanford University, Stanford CA 94305, USA

Received on February 10, 2007; revised on April 13, 2007; accepted on May 2, 2007

Advance Access publication May 11, 2007

Associate Editor: Anna Tramontano

ABSTRACT

Motivation: Membrane fusion constitutes a key stage in cellular processes such as synaptic neurotransmission and infection by enveloped viruses. Current experimental assays for fusion have thus far been unable to resolve early fusion events in fine structural detail. We have previously used molecular dynamics simulations to develop mechanistic models of fusion by small lipid vesicles. Here, we introduce a novel structural measurement of vesicle topology and fusion geometry: *persistent voids*.

Results: Persistent voids calculations enable systematic measurement of structural changes in vesicle fusion by assessing fusion stalk widths. They also constitute a generally applicable technique for assessing lipid topological change. We use persistent voids to compute dynamic relationships between hemifusion neck widening and formation of a full fusion pore in our simulation data. We predict that a tightly coordinated process of hemifusion neck expansion and pore formation is responsible for the rapid vesicle fusion mechanism, while isolated enlargement of the hemifusion diaphragm leads to the formation of a metastable hemifused intermediate. These findings suggest that rapid fusion between small vesicles proceeds via a small hemifusion diaphragm rather than a fully expanded one.

Availability: Software available upon request pending public release.

Contact: kasson@cmgm.stanford-edu or pande@stanford.edu

Supplementary information: Supplementary data are available at *Bioinformatics* online.

1 INTRODUCTION

Membrane fusion is a process of fundamental biological importance, yet its mechanisms remain incompletely understood. One of the great problems in the study of fusion has been defining the mechanisms by which two relatively stable start and end states (unfused and fused) are bridged through an energetically unfavorable intermediate space. This is both a problem central to the basic understanding of membrane function and potentially a window for therapeutic intervention into disease processes, such as viral infection, neurodegeneration and endocrine disorders such as diabetes. A hemifused

structure, in which the outer leaflets of the opposing membranes have fused but the inner leaflets remain unfused, has been suggested as a common intermediate in the fusion process (Chernomordik *et al.*, 1995; Chernomordik and Kozlov, 2005). Theoretical studies of membrane fusion have postulated progression from an early stalk-like stage (Kozlovsky and Kozlov, 2002; Markin *et al.*, 1984; Siegel, 1993) (with a point-like hemifusion diaphragm, similar to Fig. 1b) to a later stable hemifused stage in which the outer leaflets of the two vesicles have formed an expanded hemifusion diaphragm (similar to Fig. 1c) and lipids mix freely between the outer but not between the inner leaflets (Chernomordik *et al.*, 1995; Chernomordik and Kozlov, 2005; Cohen and Melikyan, 2004; Katsov *et al.*, 2004). Stable hemifused structures have been observed experimentally (Giraud *et al.*, 2005; Kemble *et al.*, 1994; Xu *et al.*, 2005), but the kinetic role of these structures in the fusion reaction remains unclear. However, the ability to generate terminal hemifused states using mutants of biological fusion proteins suggests that energetic modulation of these structures may provide a way to manipulate the fusion process. Experimental studies have shown biphasic kinetics for the formation of fused vesicles (Evans and Lentz, 2002), demonstrating the presence of multiple kinetically distinct fusion intermediates. To further elucidate these intermediates, we have recently employed molecular dynamics simulations of fusion to systematically develop a kinetic model for fusion (Fig. 1; Kasson *et al.*, 2006). Our computational predictions suggest a branched reaction pathway with a metastable intermediate, a pathway consistent with the experimentally observed kinetic data.

Despite the great importance of fusion intermediates for defining mechanisms of membrane fusion, characterization of fusion intermediates has been hampered by a lack of both high-resolution structural data and systematic structural measurement approaches to assess hemifusion structures. Using approaches from computational topology, we have developed a novel structural metric for the geometry of fusion intermediates, which we call persistent voids. High-resolution structural data are difficult to obtain from experimental sources because of the dynamic processes and large multi-molecular protein–lipid assemblies involved and because of the importance of transient intermediates. Molecular dynamics

*To whom correspondence should be addressed.

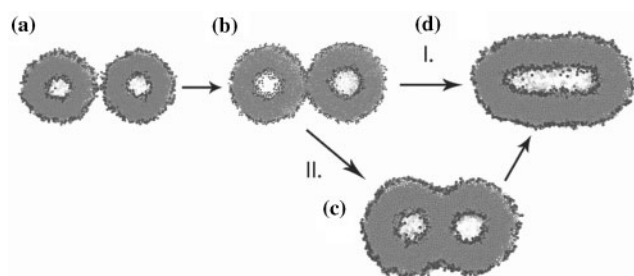


Fig. 1. Branching reaction pathway for vesicle fusion. Schematized here is the reaction pathway for the fusion of 15 nm POPE vesicles as computed from mixing-metric analysis of 10 000 molecular dynamics trajectories. Renderings show cross-sections through a representative vesicle structure for each reaction intermediate. Pathway (II) represents the canonical reaction pathway from unfused vesicles (a) through a stalk-like state (b) and a late-hemifusion intermediate (c) which we find to be metastable, to form fully fused vesicles. (d) Pathway (I) represents the additional rapid-fusion pathway that we observe in $\sim 20\%$ of reaction trajectories, in which vesicles in the stalk-like state proceed in a rapid manner to full fusion without residence in a late-hemifusion intermediate. Colour version of this figure is available as Supplementary material online.

simulation provides a means to generate high-resolution structural and dynamic models for fusion in a physically based manner. We have therefore applied this metric to molecular dynamics simulations of vesicle fusion to predict the role and nature of hemifusion intermediates in membrane fusion.

1.1 Using molecular dynamics simulations to build long-timescale models of fusion

Molecular dynamics simulation can aid in understanding membrane fusion because it provides physically based predictions at a level of structural and kinetic detail not currently available from experimental measurements. These predictions can also be combined with experimental data to generate more complete explanatory models for the fusion process. However, membrane fusion also poses substantial challenges for simulation. The fastest experimental measurement for membrane fusion is $\sim 200 \mu\text{s}$ (Llinas *et al.*, 1981), a timescale that, while fast by experimental standards, is substantially longer than the nanosecond timescale generally considered feasible for atomic-resolution molecular dynamics studies. In addition, the molecular systems involved in fusion comprise millions of atoms, making even nanosecond-level simulations a major undertaking.

To address this challenge, coarse-grained lipid molecular dynamics methods have been developed and applied to the problem of membrane fusion (Marrink *et al.*, 2004; Marrink and Mark, 2003a, b). These methods preserve the physical rules of traditional atomic-scale molecular dynamics but represent molecules at a coarser level of detail—one coarse-grained particle per 2–5 heteroatoms. Such an approach greatly decreases the computational requirements of fusion simulations while providing a quantitative approximation of atomic behavior, making 250 ns simulations

of fusion by small vesicles feasible on a timescale of 40 compute days.

Coarse-grained simulation methodology provides an important step in making molecular dynamics simulation of membrane fusion feasible. Initial validation of this coarse-grained approach to fusion has been addressed previously; observations of fusion in this context (Marrink and Mark, 2003a; Muller *et al.*, 2003; Stevens *et al.*, 2003) have suggested important hypotheses for fusion but do not provide sufficient sample sizes to make accurate quantitative predictions regarding the fusion process. In addition, although the formation and expansion of fusion stalks play an important part in many of these hypotheses, techniques for systematic measurement of fusion stalk widths from structural or molecular dynamics data have not yet been established.

We have demonstrated the combined power of distributed computing to generate molecular dynamics trajectories with extensive sampling of the relevant free-energy landscape and Markovian state models to use those trajectories to predict long-timescale kinetic behavior with high statistical accuracy (Singhal *et al.*, 2004; Snow *et al.*, 2005). Simulating thousands of molecular dynamics trajectories and constructing Markovian state models allow us to accurately model reactions that have characteristic timescales much longer than the length of any single simulation. Applying this methodology to the study of fusion, we have generated a detailed model of the branching reaction pathway for vesicle fusion (Fig. 1), including predictions of reaction rates (Kasson *et al.*, 2006). Our initial analysis utilized lipid mixing between the inner and outer leaflets, respectively, of the two vesicles as a structural metric to identify the branching reaction pathway. The results of these simulations are in good qualitative agreement with existing experimental data and previous theoretical analyses of fusion energetics. This analysis did not, however, yield an explanation of why some vesicles fuse rapidly from the stalk-like state while others form stably hemifused intermediates. We have therefore developed a novel structural metric to explain the divergence of the rapid and slow fusion pathways as well as to provide a systematic means of measuring fusion stalk widths.

1.2 Persistent voids: a novel method for fusion event detection

Previous work on vesicle fusion relied upon visual analysis or indirect readouts such as lipid order and lipid mixing. In this work, we introduce a robust computational method for measuring structural changes during fusion. Intuitively, we may view each leaflet of a vesicle as a monolayer enclosing the interior contents. For each leaflet, a fusion trajectory begins with two monolayers enclosing the interior of the two vesicles that fuse to form one enclosure. As such, fusion is a topological event as it changes the connectivity of the underlying space. We could describe this event in two different ways. The first viewpoint regards fusion as the merging of two components. Computing connected components is a well-studied problem in computer science, and efficient data structures exist for maintaining collections of disjoint dynamic sets (Cormen *et al.*, 2001). In practice, however, the spatial sampling error for

components is of comparable magnitude to the distance between two vesicles that are closely apposed but not yet fused, preventing this approach from reliably detecting fusion. Instead, we interpret the fusion event as the merging of two enclosing membranes, or voids. Generally speaking, a void is an empty space completely enclosed by a surface. Liang *et al.* (Liang *et al.*, 1998a, b) have developed a related approach to detect protein cavities by defining a discrete flow based on alpha shapes. Our use of voids requires advanced techniques from algebraic topology and allows more robust detection of relevant structures in membrane fusion. In particular, the formalism of homology unifies both the connected component and void viewpoints by extending the notion of a cycle to all dimensions: components are zero-dimensional and voids are 2D cycles. Using voids, we may detect the fusion event more reliably (see Supplementary Fig. S1).

Our approach to measuring the fusion event consists of two techniques: *alpha shapes* from computational geometry (Edelsbrunner and Mücke, 1994) and *persistent homology* from the emerging area of computational topology (Zomorodian, 2005; Zomorodian and Carlsson, 2005). Alpha shapes are geometric constructs for describing the structure of a set of sampled points. We employ α -shapes to represent leaflet geometry, generating the sampled points by extracting phosphate coordinates at a given trajectory snapshot. The method approximates the space underlying the sampled data by placing a ball of radius α centered at each sample point. As we increase α , the set of balls change connectivity, and the α -shapes method captures this connectivity at each scale in a family of combinatorial structures.

Persistent homology tracks topological attributes across scale. For each structural snapshot, we consider the components and voids formed over a range of α values. With increasing α , larger connected components are formed, and enclosed voids are created, subdivided and are ultimately filled in. Persistent homology measures this process in terms of the lifetimes of topological attributes—the values of α at which each attribute is created and destroyed. Geometrically significant attributes persist longer and therefore have longer lifetimes. By applying this technique to the voids formed by each vesicle leaflet, we are able to detect whether one or two significant voids exist at any particular simulation time point and pinpoint the fusion event.

Persistent void measurements constitute a novel topologically based means of characterizing vesicle fusion intermediates. These measurements scale with the fusion neck cross-sectional radius and allow us to track the formation and widening of hemifusion diaphragms and fusion pores in a dynamic fashion over the course of a fusion reaction. We use this measurement on 10 000 molecular dynamics trajectories of fusion by small (15 nm) lipid vesicles to determine the dynamic relationship between hemifusion neck widening and formation of a full fusion pore. We find that rapid fusion of vesicles proceeds in a coordinated rather than a stepwise fashion, while an isolated enlargement of the hemifusion diaphragm is strongly associated with the formation of a metastable hemifused intermediate. This suggests that, while hemifusion likely plays an important role in the modulation of biological fusion, a fully expanded hemifusion diaphragm

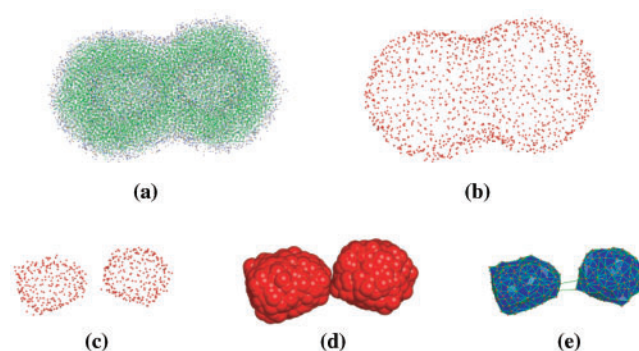


Fig. 2. Visualization of balls and α -shapes from molecular structures. For each vesicle conformation (a), we generate two point sets consisting of the outer (b) and inner (c) phosphate group coordinates. To understand the topology of a membrane leaflet, we place balls of radius α around its points, shown here for the inner leaflet (d). Alpha shapes is an efficient method that captures the connectivity of the ball set (d) in a combinatorial structure called an α -complex (e).

does not represent an obligate intermediate in the fusion pathway.

2 METHODS

2.1 Molecular dynamics simulation of vesicle fusion

Molecular dynamics were performed as previously described (Kasson *et al.*, 2006) using 15 nm palmitoylcholinephosphoethanolamine (POPE) vesicles represented using the coarse-grained forcefield reported by Marrink and Mark (Marrink *et al.*, 2004; Marrink and Mark, 2003a). Vesicles were placed in explicit solvent and linked by a 1 nm crosslinker molecule to generate a lipidic starting conformation roughly analogous to the postulated pre-fusion complex formed in viral or neuronal vesicle fusion. Molecular dynamics simulations were performed using GROMACS (Van der Spoel *et al.*, 2005) and the Folding@Home distributed computing architecture (Shirts and Pande, 2000) to generate 10 000 independent simulation trajectories.

2.2 Measuring fusion geometry using persistent voids

Persistent void measurements provide a more direct means of tracking fusion and hemifusion pore formation and development. For each structural snapshot from a fusion simulation (Fig. 2a), we consider the inner and outer leaflets separately and represent each leaflet using its phosphate group coordinates (Figs 2b and c). The alpha shapes method captures the connectivity of the ball set derived from each leaflet (Fig. 2d) in a combinatorial structure called an α -complex (Fig. 2e). As α increases, the complexes become larger: components are created and connected, and voids are created and are eventually filled in. In Figure 3, we show four α -complexes for a conformation with unfused inner leaflets. Persistent homology captures the lifetimes of topological attributes in a multiset of intervals. Since the distance between the samples is approximately the same as the distance between the vesicles, it is not possible to detect the fusion event through persistent components. Instead, we utilize the lifetimes of voids as a function of α (shown above the α -axis in Fig. 3).

Most voids have insignificant lives, but two voids stand out, corresponding to the two unfused inner leaflets. These persistent voids have statistically identical lifetimes—they are created and destroyed at the same α values—indicating that they describe two

separate and identical unfused inner leaflets. The particular snapshot we examine here describes a *hemifused* structure—the inner leaflets are unfused but the outer leaflets have fused with a wide hemifusion diaphragm. Like its unfused counterpart, a fused outer leaflet (Fig. 4) also exhibits two significant voids in its α -complexes, but the voids here are created at different α -radii and have unequal lifetimes. The void with the longest life embodies the volume of the single fused membrane. It is created when the surface of the outer leaflet is completely tiled over and is destroyed when the volume is completely filled in (Fig. 4). The second void is created when a membrane partitions the first void at the

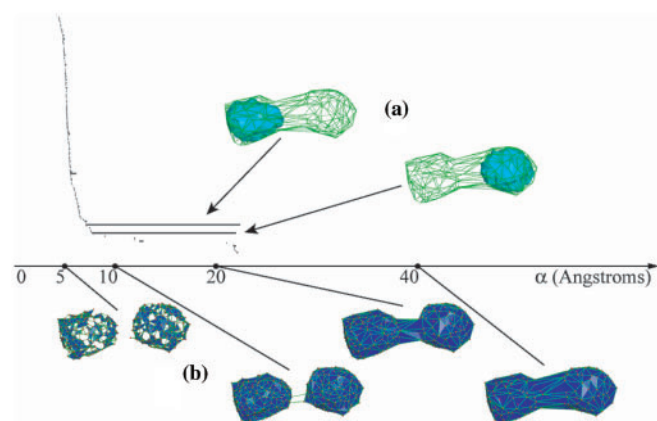


Fig. 3. Visualization of inner leaflet voids. Inner-leaflet voids (a), corresponding to the structural snapshot shown in Figure 2, are rendered above the axis; the generated α -complexes (b) are rendered below. We may vary the radius α of the balls in Figure 2d to get a family of growing α -complexes. Here, we show four α -complexes for a conformation with an unfused inner leaflet. Persistent homology captures the lifetimes of voids in a multiset of 146 intervals, drawn above the α -axis. Most voids have very short lives, but there are two significant voids. Statistically, these voids share lifetimes: they are created and destroyed at the same α value. We visualize them using descriptions, computed with persistent homology, nested inside a framework of outside edges of the rightmost α -complex shown.

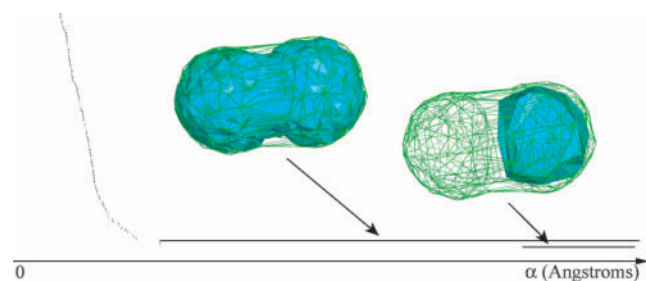


Fig. 4. Visualization of outer leaflet voids. We show the 111 void intervals for a conformation with a fused outer leaflet. Unlike the voids for the unfused inner leaflets in Figure 3, the two significant voids here have different interval lengths and are thus structurally distinct. The void with the longest interval represents the volume of the fused membrane. The second void is a subset of the first void created when the α -complex partitions the first void at the fusion neck. Consequently, we use the second void's creation radius as a geometric measure of the neck. The topological creation of a second void corresponds to the geometric partitioning of one enclosed volume into two.

fusion neck. As such, the second void's creation, α value scales approximately with the radius of the fusion neck and provides a measurement of the geometric and topological changes in membrane fusion.

To compute the persistent voids, we derived two point sets from each conformation, corresponding to the phosphate group of the inner and outer leaflet, respectively. The inner sets each contained 516 points and the outer sets each contained 1238 points. For each point set, we computed the family of α -complexes using the 3D alpha shapes software, due to Ernst Mücke and others (Edelsbrunner and Mücke, 1994). We next computed the persistence intervals, using the two longest intervals as a compact description for each leaflet. This process required about half a second for each point set. We repeated this process for the 173 526 3D point sets from the 10 000 trajectories.

2.3 Construction of predictive models using Markovian state models

Markovian state models allow systematic analysis and long timescale characterization of complex reactions by representing the reaction as a series of 'macrostates' with Markovian transitions between them (Swope *et al.*, 2004a, b). These macrostates are derived via clustering of the underlying molecular dynamics data in some reaction coordinate space. Molecular dynamics trajectories were broken into structural 'snapshots' taken at 20 ns intervals. These snapshots were then mapped to reaction coordinate space. Lipid-mixing metrics, as previously described, (Kasson *et al.*, 2006) were used as the reaction coordinate; persistent void measurements alone yield a poor reaction coordinate using current methods (Supplementary Material). Macrostates were identified using k-means clustering on the normalized reaction space. Each macrostate can now be considered a node in a graph-theoretic representation of the reaction and each molecular dynamics trajectory a series of edges between the nodes. We use this clustering approach both to classify structural snapshots of vesicle fusion and to construct long-timescale dynamic models of fusion (see Supplementary Materials for details).

3 RESULTS

Theoretical work on hemifusion intermediates has postulated both an early hemifused, or stalk-like, state and a late hemifused state marked by the presence of an enlarged hemifusion diaphragm (Chernomordik *et al.*, 1995; Chernomordik and Kozlov, 2005; Cohen and Melikyan, 2004; Kozlovsky and Kozlov, 2002; Markin *et al.*, 1984; Siegel, 1993). Our molecular dynamics simulation results yield states corresponding to each of these structures (Fig. 1) and allow us to characterize the kinetic relationships between them. Previous analyses using measurements of lipid vesicle mixing have allowed us to identify a branching reaction pathway in the simulation data; our new metrics for vesicle geometry allow more precise analyses of the structural differences between these pathways.

We can track the process of fusion via the second void creation radius for the outer and inner leaflets, respectively (Fig. 5). This radius gives the value of α required to topologically separate the two vesicles into separate enclosed volumes, thus quantifying the extent of fusion between them. Our persistent voids method for the first time provides a measurement of the actual geometric and topological changes in membrane fusion. We have used this means of assessing fusion intermediate geometry to track the development of

fusion pores in 10 000 molecular dynamics trajectories of the fusion of two small lipid vesicles (Kasson *et al.*, 2006). In our previous analyses of vesicle fusion simulations, we identified two reaction pathways: one that proceeds rapidly to fusion

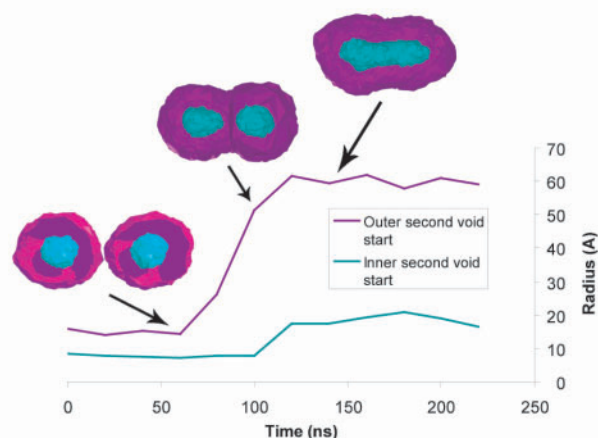


Fig. 5. Progress of the fusion reaction assessed by persistent voids. A given fusion trajectory may be followed using the fusion neck width (approximated by the second void creation radius) as a reaction coordinate. The second void creation radius is plotted here for both the outer and inner vesicle leaflets, and cut away renderings of the outer and inner voids are shown at the unfused, hemifused and fully fused stages. Outer leaflet voids are rendered in magenta, and inner leaflet voids are rendered in cyan.

(Fig. 1, pathway I) and one that remains metastably in a late-hemifused state before fusing on a timescale of micro-seconds (Fig. 1, pathway II). Assuming a fusion mechanism that proceeds from a stalk-like intermediate, one could hypothesize two possible fusion mechanisms: either that the widening of the hemifusion diaphragm and the formation of a fusion pore are tightly correlated or that they are independent events. In the tightly correlated case, one would expect to observe hemifusion diaphragm widening and pore formation to occur together or separated by some small and relatively constant δt . If pore formation is an independent event from hemifusion diaphragm widening, one would expect the δt of separation to be a random variable.

The characteristic patterns of fusion neck width evolution for these two mechanisms are shown in Figure 6. We observe both these events, coordinated and stepwise, in our simulation data, but with vastly different frequencies. Correlation coefficients between second void persistence lengths of the inner and outer leaflets were used to assess coordinated fusion neck widening. To allow for a time lag δt , we also use phase-shifted correlation coefficients for each trajectory (Equation 1).

$$r(\delta) = \frac{\sum (x_i - \bar{x})(y_{i+\delta} - \bar{y})}{\sqrt{\sum (x_i - \bar{x})^2 \sum (y_{i+\delta} - \bar{y})^2}} \quad (1)$$

Correlation coefficients were significantly higher in trajectories that fused compared to trajectories that did not ($P < 1 \times 10^{-10}$ via Wilcoxon rank sum test).

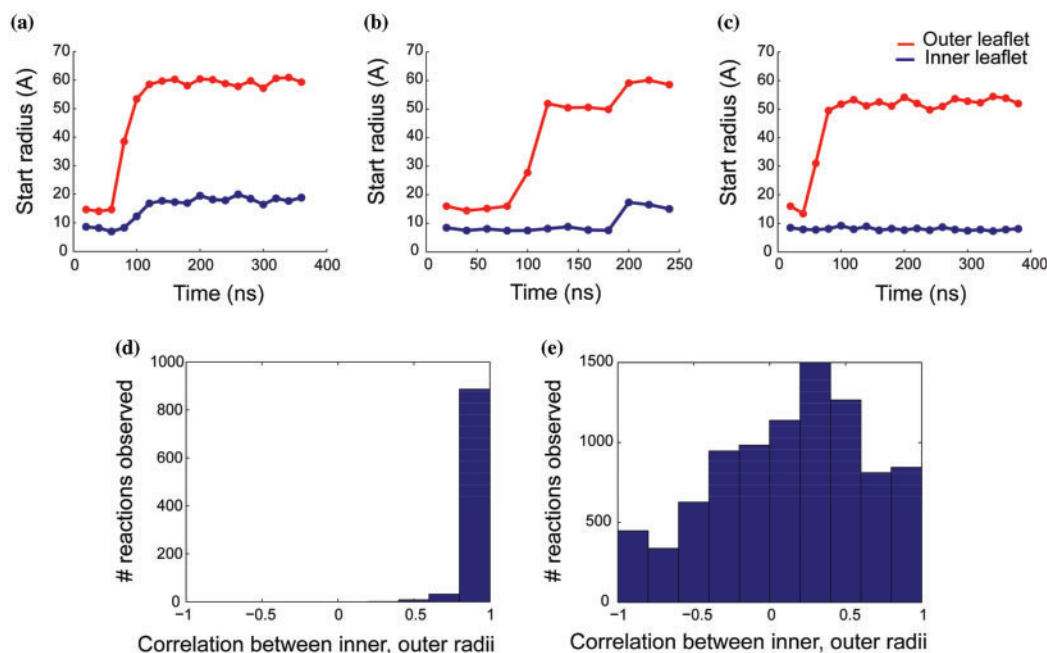


Fig. 6. Rapid vesicle fusion involves coordinated fusion-neck expansion. Fusion trajectories assessed by fusion neck width fall into one of three general patterns. For rapid fusion, the most common pattern is that of coordinated widening of the outer and inner fusion neck widths (a). Less commonly, there is a variable time lag between the outer and inner leaflet widening events (b). The third pattern is that of metastable hemifusion (c), in which the outer leaflet fusion neck widens hundreds of nano seconds to micro seconds before the inner leaflet. Panels (d) and (e) provide quantitative confirmation that most rapid fusion proceeds in a coordinated manner by comparing the outer–inner correlation functions for rapid fusion (d) and hemifusion (e).

We have also used our Markovian State Model construct to categorize trajectories into those proceeding directly to fusion (Fig. 1, pathway I) and those fusing via a late-hemifusion intermediate (Fig. 1, pathway II). Correlations between the inner and outer fusion neck measurements are shown in Figure 6d and e. For trajectories proceeding through a late-hemifusion intermediate, there is no statistically significant correlation between the inner and outer neck widths over time, while for trajectories fusing via the rapid pathway these two measurements are tightly correlated ($P < 1 \times 10^{-10}$ via Kolmogorov–Smirnov test using correlation coefficients from hemifusion trajectories as a reference sample). When a phase shift δt is added, 75% of trajectories in the rapid fusion pathway have an optimal $\delta t < 20$ ns, while only 22% of the trajectories proceeding through the late-hemifused pathway have an optimal $\delta t < 20$ ns (Fig. 7). This correlation and the distribution of phase shifts suggest that the rapid fusion pathway proceeds via coupled rather than independent widening of the hemifusion diaphragm and formation of a fusion pore.

4 CONCLUSIONS

In this work, we describe the development of a new structural metric for vesicle fusion, persistent voids, that measures the neck width of fusion pores and hemifusion diaphragms. This novel measurement approach allows systematic assessment of fusion intermediate geometries and is also generally applicable to the analysis of curved membrane structures. We apply this metric to data from molecular dynamics simulations, which offer a means of generating detailed models for vesicle fusion in

a physically based manner. We have performed massively distributed simulations of vesicle fusion and developed the subsequent systematic analysis of reaction pathways with good statistical sampling and quantitative prediction of long time-scale kinetics. In previous work, we have shown that these molecular dynamics simulations suggest a branched reaction pathway for the fusion of small lipid vesicles, in good accordance with experimental data. In this branched pathway, an early stalk-like intermediate can either react rapidly to form a fusion pore or can proceed more slowly through a metastable hemifused intermediate.

We have used persistent voids to characterize more precisely, the transition from a stalk-like intermediate to the formation of a fusion pore in our molecular dynamics simulations. In the rapid fusion pathway, we find widening of the hemifusion diaphragm to proceed in a manner tightly coupled to the formation of a fusion pore. This coordination between hemifusion diaphragm and fusion pore widening is the hallmark that differentiates the rapid fusion pathway from the stepwise diaphragm and pore widening that predominates in the slow fusion pathway.

These findings and our previous free energy calculations (Kasson *et al.*, 2006) suggest the following model for fusion. Isolated expansion of the hemifusion diaphragm energetically stabilizes the outer leaflet but traps the system in a metastable state. Coordinated expansion of both the hemifusion diaphragm and the fusion pore harnesses the outer-leaflet energetic strain created in the stalk-like state to drive full fusion. Modulation of these two pathways is likely dependent on the lipid properties in the region of the stalk. Our results are most directly applicable to the small, highly curved vesicles that we simulate. We hypothesize that the underlying mechanistic pathways for fusion are likely conserved across vesicle sizes; however, the relative energetics and frequency of different pathways will vary.

Fast fusion events have been observed in a number of experimental systems; however, kinetically resolving the early fusion events has been difficult. Our model suggests that, in these fast fusion events, inner and outer leaflet fusion proceeds in a coordinated rather than in a stepwise-independent manner. In this case, the radial width of the hemifusion diaphragm prior to inner-leaflet fusion would be small; this diaphragm undergoes further expansion in conjunction with inner-leaflet fusion. Our model is derived from physics-based simulation; final confirmation of our hypothesis will await experimental methods to observe these fast structural changes directly.

ACKNOWLEDGEMENTS

The authors thank O. Troyanskaya and T. Fenn for many helpful discussions. This work was supported in part by the Simbios NIH Center for Biomedical Computing and by a fellowship from the Berry Foundation to P.K. Research by A.Z. and L.G. was partially supported under NSF grants ITR-0205671, CARGO-0310661, NIH grant GM072970, and DARPA grant 32905 at Stanford and under grant HR0011-06-1-0038 at Dartmouth.

Conflict of Interest: none declared.

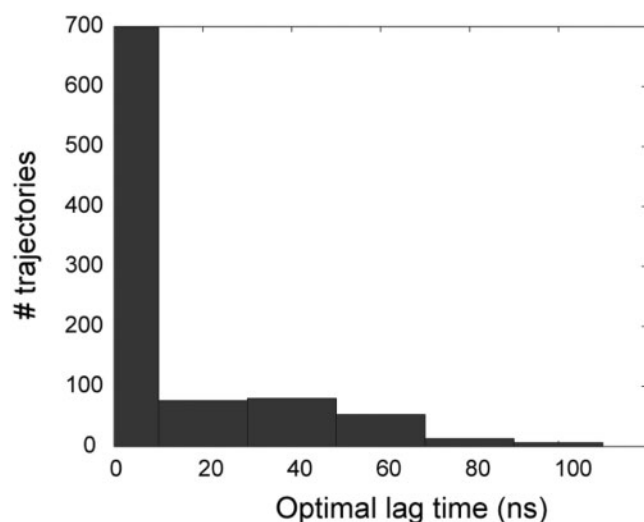


Fig. 7. Coordinated expansion of outer and inner leaflets shown via optimal lag times. Correlation coefficients are computed between the outer and inner leaflet voids for each trajectory that undergoes rapid fusion using a variable phase shift from 0 to 100 ns. For each trajectory, the phase shift was selected that maximizes the correlation coefficient, and a histogram of these optimal time lags is plotted. Greater than 75% of trajectories have an optimal time lag < 20 ns, indicating coordinated expansion of the inner and outer leaflets during rapid fusion. Colour version of this figure is available as Supplementary material online.

REFERENCES

- Chernomordik, L. *et al.* (1995) The hemifusion intermediate and its conversion to complete fusion – Regulation by membrane-composition. *Biophys. J.*, **69**, 922–929.
- Chernomordik, L.V. and Kozlov, M.M. (2005) Membrane hemifusion: crossing a chasm in two leaps. *Cell*, **123**, 375–382.
- Cohen, F.S. and Melikyan, G.B. (2004) The energetics of membrane fusion from binding, through hemifusion, pore formation, and pore enlargement. *J. Membr. Biol.*, **199**, 1–14.
- Cormen, T.H. *et al.* (2001) *Introduction to Algorithms*. MIT Press, Cambridge, MA.
- Edelsbrunner, H. and Mücke, E.P. (1994) 3-Dimensional alpha-shapes. *AcM Trans. Graph.*, **13**, 43–72.
- Evans, K.O. and Lentz, B.R. (2002) Kinetics of lipid rearrangements during poly(ethylene glycol)-mediated fusion of highly curved unilamellar vesicles. *Biochemistry*, **41**, 1241–1249.
- Giraud, C.G. *et al.* (2005) SNAREs can promote complete fusion and hemifusion as alternative outcomes. *J. Cell Biol.*, **170**, 249–260.
- Kasson, P.M. *et al.* (2006) Ensemble molecular dynamics yields sub-millisecond kinetics and intermediates of membrane fusion. *Proc. Natl Acad. Sci. USA*, **103**, 11916–11921.
- Katsov, K. *et al.* (2004) Field theoretic study of bilayer membrane fusion. I. Hemifusion mechanism. *Biophys. J.*, **87**, 3277–3290.
- Kemble, G.W. *et al.* (1994) Lipid-anchored influenza hemagglutinin promotes hemifusion, not complete fusion. *Cell*, **76**, 383–391.
- Kozlovsky, Y. and Kozlov, M.M. (2002) Stalk model of membrane fusion: solution of energy crisis. *Biophys. J.*, **82**, 882–895.
- Liang, J. *et al.* (1998a) Analytical shape computation of macromolecules: I. Molecular area and volume through alpha shape. *Proteins*, **33**, 1–17.
- Liang, J. *et al.* (1998b) Analytical shape computation of macromolecules: II. Inaccessible cavities in proteins. *Proteins*, **33**, 18–29.
- Llinas, R. *et al.* (1981) Relationship between presynaptic calcium current and postsynaptic potential in squid giant synapse. *Biophys. J.*, **33**, 323–351.
- Markin, V.S. *et al.* (1984) On the theory of membrane fusion. The stalk mechanism. *Gen. Physiol. Biophys.*, **3**, 361–377.
- Marrink, S.J. *et al.* (2004) Coarse grained model for semiquantitative lipid simulations. *J. Phys. Chem. B*, **108**, 750–760.
- Marrink, S.J. and Mark, A.E. (2003a) The mechanism of vesicle fusion as revealed by molecular dynamics simulations. *J. Am. Chem. Soc.*, **125**, 11144–11145.
- Marrink, S.J. and Mark, A.E. (2003b) Molecular dynamics simulation of the formation, structure, and dynamics of small phospholipid vesicles. *J. Am. Chem. Soc.*, **125**, 15233–15242.
- Muller, M. *et al.* (2003) A new mechanism of model membrane fusion determined from Monte Carlo simulation. *Biophys. J.*, **85**, 1611–1623.
- Shirts, M. and Pande, V.S. (2000) Computing – screen savers of the world unite! *Science*, **290**, 1903–1904.
- Siegel, D.P. (1993) Energetics of intermediates in membrane-fusion – comparison of stalk and inverted micellar intermediate mechanisms. *Biophys. J.*, **65**, 2124–2140.
- Singhal, N. *et al.* (2004) Using path sampling to build better Markovian state models: predicting the folding rate and mechanism of a tryptophan zipper beta hairpin. *J. Chem. Phys.*, **121**, 415–425.
- Snow, C.D. *et al.* (2005) How well can simulation predict protein folding kinetics and thermodynamics? *Annu. Rev. Biophys. Biomol. Struct.*, **34**, 43–69.
- Stevens, M.J. *et al.* (2003) Insights into the molecular mechanism of membrane fusion from simulation: Evidence for the association of splayed tails. *Phys. Rev. Lett.*, **91**(18): 188102.
- Swope, W.C. *et al.* (2004a) Describing protein folding kinetics by molecular dynamics simulations. 1. Theory. *J. Phys. Chem. B*, **108**, 6571–6581.
- Swope, W.C. *et al.* (2004b) Describing protein folding kinetics by molecular dynamics simulations. 2. Example applications to alanine dipeptide and beta-hairpin peptide. *J. Phys. Chem. B*, **108**, 6582–6594.
- Van der Spoel, D. *et al.* (2005) GROMACS: fast, flexible, and free. *J. Comput. Chem.*, **26**, 1701–1718.
- Xu, Y. *et al.* (2005) Hemifusion in SNARE-mediated membrane fusion. *Nat. Struct. Mol. Biol.*, **12**, 417–422.
- Zomorodian, A.J. (2005) *Topology for Computing*. Cambridge University Press, Cambridge, UK, New York.
- Zomorodian, A. and Carlsson, G. (2005) Computing persistent homology. *Discrete Computat. Geom.*, **33**, 249–274.