Journal of Computational Physics 228 (2009) 3656–3668

Contents lists available at ScienceDirect

Journal of Computational Physics

journal homepage: www.elsevier.com/locate/jcp

Effect of excluded volume on 2D discrete stochastic chemical kinetics

Sotiria Lampoudi [a,*], Dan T. Gillespie [b], Linda R. Petzold [a]

[a] Dept. of Computer Science, University of California, Santa Barbara, CA 93106, USA
[b] Dan T Gillespie Consulting, 30504 Cordoba Place, Castaic, CA 91384, USA

ARTICLE INFO

Article history:
Received 20 October 2008
Received in revised form 24 December 2008
Accepted 3 February 2009
Available online 12 February 2009

PACS:
02.50.Ga
02.70.Ns
05.10.Ln
05.70.Ln

ABSTRACT

The stochastic simulation algorithm (SSA) is widely used in the discrete stochastic simulation of chemical kinetics. The propensity functions which play a central role in this algorithm have been derived under the point-molecule assumption, i.e., that the total volume of the molecules is negligible compared to the volume of the container. It has been shown analytically that for a one-dimensional system and the A + A reaction, when the point-molecule assumption is relaxed, the propensity function need only be adjusted by replacing the *total* volume of the system with the *free* volume of the system. In this paper we investigate via numerical simulations the impact of relaxing the point-molecule assumption in two dimensions. We find that the distribution of times to the first collision is close to exponential in most cases, so that the formalism of the propensity function is still applicable. In addition, we find that the area excluded by the molecules in two dimensions is usually higher than their close-packed area, requiring a larger correction to the propensity function than just the replacement of the total volume by the free volume.

1. Introduction

The stochastic simulation algorithm (SSA) [1] is the workhorse algorithm for discrete stochastic simulation of networks of coupled chemical reactions. The physical system, in this case, is a collection of molecules of various chemical species that move around inside a fixed volume, and are subject to a set of chemical reactions in which the molecules may be reactants or products or both. The chemical reactions are all assumed to be "elementary" in the sense that they occur essentially instantaneously. Elementary reactions will invariably be either unimolecular or bimolecular; all other types of reactions (trimolecular, reversible, etc.) will consist of a series of two or more elementary reactions. If the system is well-stirred, we can define its state simply by giving the vector $\mathbf{x}$ of the molecular populations of the various chemical species. In that circumstance, it is usually possible to describe the dynamics of each reaction channel R_j by a "propensity function" $a_j(\mathbf{x})$, defined so that if the system is in state $\mathbf{x}$, then $a_j(\mathbf{x})dt$ gives the probability that the reaction will occur somewhere inside the system in the next infinitesimal time interval dt.The magnitude of $a_j(\mathbf{x})$ thus measures the "propensity" of reaction R_j to occur in the immediate future.

The propensity function is very close to, and sometimes numerically equal to, what in deterministic chemical kinetics is called the "reaction rate". But the propensity function does not make the assumption that reactions occur continuously and deterministically, and its product with dt is mathematically treated as a probability. The outcome of such a set of assumptions is the chemical master equation (CME) and the stochastic simulation algorithm (SSA), as discussed in numerous articles over the past three decades (for an overview of relevant contributions to the development of the chemical master equation and the stochastic simulation algorithm, see the review article [1] and references therein). In the thermodynamic limit

* Corresponding author. Tel.: +1 8052598536.
E-mail address: slampoud@cs.ucsb.edu (S. Lampoudi).

0021-9991/$ - see front matter © 2009 Elsevier Inc. All rights reserved.
doi:10.1016/j.jcp.2009.02.002

(infinite populations and infinite system volume with finite concentrations), the CME and SSA almost always reduce to the ordinary differential equations of deterministic chemical kinetics.

The SSA generates times τ between successive reactions as samples of an exponential distribution whose mean is equal to the inverse of the sum of the propensity functions. The most commonly used propensity functions are of a mass action form, according to which the rate of a reaction is proportional to the combinatorial product of the reactants' populations.

Mass action propensity functions for elementary reactions have been rigorously derived in a *well-stirred, dilute* hard sphere setting [2]. In this setting molecules are represented by hard spheres moving ballistically in a vacuum. We refer to them as *point molecules*, because, although they must have non-zero diameter l in order to collide, the volume of all the molecules combined is negligible compared to the volume of their container. If the point-molecule assumption is relaxed, to what extent does the volume occupied by the reactant molecules themselves affect the rates of the reactions in which they participate?

We will be studying the effect of reactant-excluded volume, in a simple but computationally tractable physical model. Specifically, we will attempt to answer the following two questions. First, is the time between successive reactions in a well-stirred, non-point molecule system exponentially distributed, as it must be for the stochastic process theory which underlies the SSA to hold? Second, if the reaction times are exponentially distributed, what is the mathematical form of the propensity functions in this setting?

Since bimolecular reactions are always initiated by a collision, the probability of a reaction between two molecules can be broken down into (a) the probability that the two molecules will collide, times (b) the probability that they will react given that they have collided. Throughout this work we make the simplifying assumption that (b) is unity, and thus use the terms *collision* and *reaction* probability (and inter-*collision* and inter-*reaction* time) interchangeably.

We have previously shown how, for a *one-dimensional* system, the mass action propensity functions need to be modified when the volume of the reactant molecules is comparable to the total system volume [3]. We analytically derived the following exact formula for the reaction probability, in the next infinitesimal time dt, of the reaction $A + A \rightarrow products$ in a one-dimensional system of N non-overlapping hard rods of length l moving ballistically in a volume of length L:

$$p_{\mathrm{col}}(dt) = \frac{N(N-1)s_{\mathrm{rel}}}{2(L-Nl)}\,dt. \tag{1}$$

(In the limit of $l \rightarrow 0$ this is equal to the usual dilute gas reaction rate law.) The propensity function for the reaction is, by definition, this probability divided by dt. In Eq. (1), s_{rel} is the mean relative speed of two randomly chosen rods. The correctness of this formula was then confirmed through an extensive series of exact hard rod molecular dynamics simulations.

An analogous treatment of the two-dimensional hard disk system has proved to be challenging. The difficulty arises when trying to find an analytical inter-molecular distance distribution function for non-overlapping, non-zero sized hard disks in a finite area. The one dimensional case, given in [3] is essentially a consequence of the Tonks result [4]. But, to the best of our knowledge, a two or three-dimensional exact version has not been reported in the literature, and we have not been able to derive it ourselves.

Thus, in this paper we use the hard spheres molecular dynamics simulation methodology to *computationally* investigate the effect of molecule size on the propensity for the $A + A \rightarrow products$ reaction in the two-dimensional version of the system. We consider a system of N hard disks, each of diameter l, initially distributed uniformly randomly with no overlap inside a circular container with hard reflective walls and diameter L. The choice of hard instead of periodic boundaries was made after careful consideration. We believe that hard boundaries bring our simple system closer to being "realistic". Periodic boundaries would introduce the unphysical "appearance" of molecules from nowhere, as they cross the boundary. Also, for molecules that have non-zero diameter, periodic boundaries make choices regarding initial random placement and inter-molecular collision detection awkward, if not arbitrary.

The molecules move ballistically, and their initial velocities are drawn from a Maxwell–Boltzmann distribution. These initial conditions represent a well-stirred system in thermal equilibrium. For this system, we collect statistics for the time τ from the initialization of the system until the first inter-molecular collision. We will *not* be concerned with the evolution of the system *beyond* the first collision, because our goal here is simply to study the form of the propensity functions when the well-stirred condition, which is assumed by the SSA, holds before each reaction. The question of under what conditions such a system will *return* to a well-stirred state is both interesting and important, but we do *not* address that question in this paper. We do, however, briefly consider the effect of container shape on our results.

We find that the distribution of inter-collision times τ in this system is approximately exponential, but with noticeable deviations in certain circumstances. We study how the τ distribution varies with the parameters l and L, which, for a fixed N, determine the area density of the system (defined as the ratio of the area of the molecule disks to the total area of the system).

For small numbers of molecules, it appears that three types of τ distribution are present: at intermediate values of the area density, the distribution is indistinguishable from an exponential; as the system tends to the point-molecule limit (low area density), *long* inter-molecular collision times are over-represented; as the area density of the system becomes high, *short* inter-molecular collision times are over-represented.

It is known that the choice of container shape affects the degree of ergodicity of the molecules' trajectories, with some container shapes encouraging trajectories that sample only small parts of the container's area. In the low population and small molecule size limit, we find that the small number of molecules, combined with a choice of non-ergodic container shape (e.g. circular, as opposed to "stadium"), gives rise to over-represented long times, while the τ distribution for short

times is just as predicted theoretically. Either increasing the number of molecules, or improving the shape of the boundary, ameliorates this effect.

As the area density of the system becomes high (i.e. for large l/L), and for low population, we find that the excluded area inferred from the simulation measurements is larger than what one might expect from taking into account the area of the molecule disks, or their close-packed area, or even several less dense packings [5]. However, as the number of molecules is increased, the excluded area approaches the close-packed area, as one might have expected.

These results suggest that excluded area in two and three-dimensional, finite, dense systems has a somewhat predictable impact for systems with a large number of molecules, but may have greater impact than one might have initially supposed for systems with low population.

In Section 2 we present the physical system under consideration and the computational algorithms used to simulate its kinetics. In Section 3 we give a brief derivation of the probability of collision from first principles. In Section 4 we present the results of our simulations. Section 5 summarizes and attempts to explain our findings.

2. Hard sphere molecular dynamics

2.1. The hard sphere molecular dynamics algorithm

The hard sphere molecular dynamics simulation algorithm is a simple billiard balls simulator. Given randomly uniform initial positions and random Maxwellian velocities for the molecules, it returns the time to the first inter-molecular collision.

The algorithm, as used for the two-dimensional problem, has the following steps:

1. Initialize N molecules, each a disk of diameter l, with:
   a. uniform random non-overlapping positions in a finite container of area A (see placement algorithms in Sections 2.2 and 2.3);
   b. velocities distributed according to the Maxwell–Boltzmann distribution.
2. For each molecule, compute the next putative collision time with the boundary of the system (ignoring all the other molecules in the system).
3. For each pair of molecules, compute the next putative inter-molecular collision time, if it exists.
4. Choose the smallest of all molecule-boundary and inter-molecular collision times, and advance the time by that amount.
5. If the collision was:
   a. inter-molecular, then report the time and exit;
   b. between a molecule and the boundary, then reflect the colliding molecule according to a specular reflection formula; advance the positions of the remaining molecules, and go back to step 2.

The velocities of the molecules are Maxwell–Boltzmann distributed, meaning that each Cartesian component of the velocity vector is a statistically independent normal random variable with mean 0 and temperature-determined variance σ^2.

Statistics for the distribution of times τ to the first collision are obtained by running an ensemble of 100,000 realizations of the hard spheres molecular dynamics algorithm (for a given N, l, and L), each with different random seeds.

2.2. Generating exact initial positions

For the *one-dimensional* problem it was possible to derive a rejection-free Monte Carlo algorithm for generating samples of the molecules' initial positions. An analogous procedure for two dimensions has proven elusive (this problem is equivalent to finding the analytical formula for the inter-molecular distance distribution function), so a rejection-based Monte Carlo algorithm was used to generate uniform random non-overlapping positions for the molecules in two dimensions.

The rejection-based initial placement algorithm has the following steps:

For molecule index $i = 1, \ldots, N$:

1. Generate putative coordinates (x_i, y_i) for the center of molecule i uniformly and randomly inside the container.
2. Check whether the proposed placement will overlap any of the $(i-1)$ previously placed molecules
   a. If an overlap is discovered, discard the entire set of placed molecules (indices $1, \ldots, i$). Reset the index to $i = 1$ and restart the procedure.
   b. If no overlap is discovered, accept the placement, increment the index and go to step 1.

Step (2a) is the only non-obvious step of the placement algorithm: we discard the *entire set* of already placed molecules, as opposed to only the very last one, which produced the overlap. The latter procedure will result in a biased distribution, as can be demonstrated both analytically and by numerical simulation in a one-dimensional setting.

The computational cost of this rejection-based method increases with the total number of molecules, and with the area density of the system. The simulations we were able to perform using this placement method reached up to area densities of 40% for small values of N (e.g. $N = 6$).

2.3. Generating approximate initial positions

One approximate alternative to generating the initial positions of the molecules by the rejection-based Monte Carlo method is to use a *pre-stirring* procedure. In this scheme we initialize the positions of the molecules on a regular grid and then allow the molecules to bounce around until their positions are randomized. The success of such an initialization depends on choosing a good stopping condition. This can be: (a) a time by which an ensemble of positions is guaranteed to have become well-stirred, (b) a target value for a function (e.g. the radial distribution function), which when reached would denote a well-stirred state, or (c) noticing that a function (again, possibly the radial distribution function) appears to have converged, without explicitly having a target value for it. We were unable to locate either (a) or (b) in the literature, and our experience with checking the radial distribution function for convergence suggested that it was too noisy for systems with small numbers of molecules.

There is a subtle aspect to choosing the stopping criterion when the system is simulated not by traditional Molecular Dynamics (i.e. integration of the laws of motion), but by our exact hard sphere algorithm, which steps from collision to collision. If we stop pre-stirring the system *immediately* after an inter-molecular collision has occurred, the two molecules which have just collided will be touching. Thus, choosing a stopping criterion that is based on the idea that an *exact* number of collisions must occur before stopping will yield ensembles of initial positions which systematically contain two molecules touching. One must remember to remove such a bias by evolving the system for some time past the last collision, or avoid introducing such a bias in the first place, by choosing a stopping condition expressed as a duration of time rather than number of collisions. We do the latter.

In our experiments exploring different amounts of pre-stirring we found that unexpectedly long amounts of pre-stirring time were necessary for other indicators, e.g. the mean of τ, which is the focus of this paper, to reach the values obtained by the rejection-based method. For a given system (i.e. given N, l, L) we quantify the duration of pre-stirring as $FN\tilde{\mu}_\tau$ units of time, where F is a scaling factor, N is the number of molecules, and $\tilde{\mu}_\tau$ is the expected time to the first collision for this system, under the assumption that $V_e = V_d = Na$ (see Section 3). Fig. 1 shows, for a system with $N = 9$, $L = 10$, $l = 2$, how the mean time to the first collision computed by simulations initialized with pre-stirring (circles) approaches that computed by a simulation initialized by rejection-based Monte Carlo (solid line with dashed confidence interval), with increasing F.

For a factor $F = 300$, and the system of Fig. 1 ($N = 9$), we see on average 4113 inter-molecular collisions, i.e. each molecule experiences on average 457 collisions with another molecule. We also see on average 2401 collisions with the boundary, i.e. on average 267 collisions with the boundary per-molecule. So F is a conservative estimate of the number of inter-molecular collisions each molecule experiences during the pre-stirring simulation.

While the pre-stirring method of obtaining initial positions makes it possible to work with dense, high population systems which we would otherwise not have been able to examine, it is nevertheless very time consuming. As can be seen from Fig. 1, $F = 300$ is suboptimal in terms of achieving good stirring (if the mean of the τ-distribution is used as a convergence criterion). However it would take 10 times as long to simulate with $F = 3000$, in order to achieve the dubious improvement of having the 95% confidence interval of the pre-stirring method overlap that of the rejection-based Monte Carlo. Simulating ensembles of size 100,000 with a stirring factor of $F = 300$ already takes days on a workstation, so the results we show in this paper are based on simulations with $F = 300$, rather than a higher value.

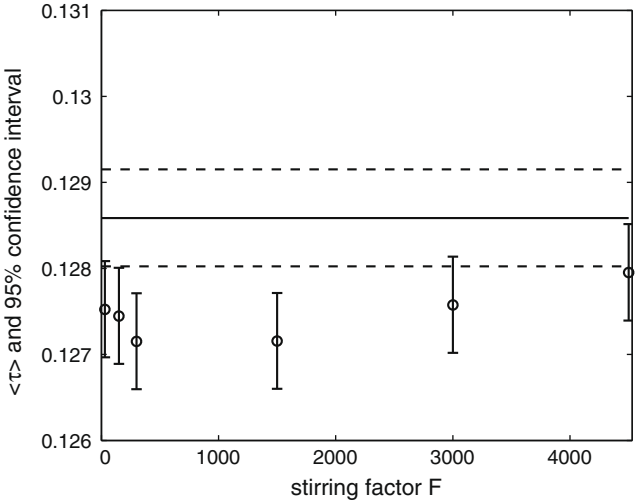

Fig. 1. The mean of the τ-distribution, as calculated from simulations initialized with rejection-based Monte Carlo, is given by the solid lines. The dashed lines give the 95% confidence interval in the estimate of the mean. The circles and error bars give the means and 95% confidence intervals of the means of the τ-distribution, as calculated from simulations initialized with pre-stirring with F given by the respective x-coordinates. The system is $N = 9$, $L = 10$, $l = 2$.

Due to the lack of a rigor of our pre-stirring method, we do not consider it to be on equal footing with the rejection-based Monte Carlo. Nearly all the results in this paper were obtained using simulations initialized by rejection-based Monte Carlo. The single exception is Fig. 7, which would have been impossible to obtain by rejection-based Monte Carlo, and which we therefore consider *speculative*.

3. Theory for the two-dimensional case

Here we derive an expression for the probability $p_{\mathrm{col}}(dt)$ of an inter-molecular collision occurring in the next infinitesimal time interval dt. This probability is given by the product of: (the number of ways one can choose a random pair of molecules) times (the probability that a randomly chosen pair of molecules will collide in the next dt). Because of the randomly uniform spatial distribution, the second factor can be further decomposed into the ratio: (the area one molecule will sweep out relative to the other molecule in the next dt) over (the total area inside the container accessible to the other molecule).

The initial velocities of the molecules are assumed to follow the equilibrium Maxwell–Boltzmann distribution, i.e. their Cartesian components are normal random variables with means 0 and variances $\sigma^2 \equiv k_B T/m$, where T is the absolute temperature of the system, m is the mass of the molecules and k_B is Boltzmann's constant. The mean relative speed of two randomly chosen molecules in such a distribution can be shown to be

$$s_{\mathrm{rel}} = \langle |v_{\mathrm{rel}}| \rangle = \sqrt{\pi}\sigma. \tag{2}$$

Suppose we randomly choose a pair of molecules. We can do this in $\frac{1}{2}N(N-1)$ ways. Now we change our frame of reference so that we are standing on one of the chosen molecules. Then the area swept out by the other molecule in the next infinitesimal time increment dt relative to the center of the one we are standing on is (see Fig. 2) $2ls_{\mathrm{rel}}\,dt$. To get the probability of our randomly chosen pair of molecules colliding in the next dt, we must divide this area by the total area available to the molecules. If we assume that the molecules have no extent ($l = 0$), the area A available to the molecules is the total area of the system. If the shape of the container is circular, and the diameter is L, the area is given by $A_c = \frac{1}{4}\pi L^2$. For a "stadium" container (formally known as a Bunimovich stadium [6]), in which semi-circles of diameter L are separated by a square of side L, the area is $A_s = \frac{1}{4}\pi L^2 + L^2$.

If we relax the point-molecule assumption (i.e. $l > 0$), the area available to the molecules will be less than the total area A of the system by *at least* the area excluded by the disks of the molecules themselves ($V_d = Na = N\frac{1}{4}\pi l^2$). The question of exactly how large is the excluded area has prompted this research. We will call the excluded area as estimated from the simulation data the *effective excluded area* V_e, and we will show in the results section that it is quite a bit larger than V_d.

Finally, combining all of the above, we find that the probability of an inter-molecular collision in the next t is given by

$$p_{\mathrm{col}}(dt) = \frac{N(N-1)}{2}\frac{2ls_{\mathrm{rel}}}{A - V_e}\,dt \tag{3}$$

with $V_e = 0$ at the limit $l = 0$.

If we assume that the same probability $p_{\mathrm{col}}(dt)$ holds for each successive infinitesimal time increment dt, from the initialization of the system until the first inter-molecular collision occurs, then the times τ will be samples of the exponential distribution whose mean is the inverse of the coefficient of dt in Eq. (3). If we further assume that every collision results in a reaction, that coefficient will be the propensity function of the stochastic simulation algorithm and the master equation for that reaction [1].

4. Simulation results

4.1. Collision time distribution methodology

To fully characterize the distributions of times to the first collision, τ, we asked the following questions: what is the empirical distribution of τ at different area densities, and what is the relationship between this empirical distribution and the analytical exponential in Eq. (3) with the most conservative estimate of excluded volume ($V_e = V_d = Na$)?

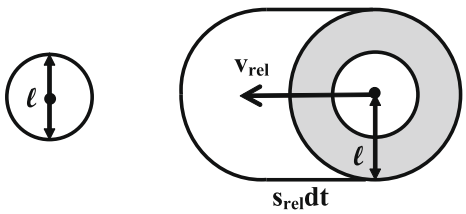

Fig. 2. To calculate the probability of collision of two randomly chosen molecules of diameter l, moving with relative velocity v_{rel}, in the next infinitesimal time dt we compute the area that one molecule sweeps relative to the other in that amount of time: $2ls_{\mathrm{rel}}\,dt$.

Upon casual visual inspection, all the distributions certainly appear exponential. To quantitatively *test* for exponentiality, we used a common two-sample comparison test, the Kolmogorov–Smirnoff test [7]. To test a τ distribution (100,000 values) for exponentiality, we generated the same number of random samples from an exponential with the same mean as the τ distribution. We then used Matlab's *kstest* routine at the default significance level $\alpha = 0.05$. The routine rejects the null hypothesis that the samples came from the same distribution (exponential) if the so-called p-value returned by the test is < 0.05; it fails to reject the null hypothesis that the samples came from the same distribution, if the p-value is > 0.05.

4.2. Collision time distribution for N = 6

We began by considering the low molecule count case $N = 6$, with $\sigma = 1$ in a circular container of diameter L. Fig. 3 shows the regions of (L, l) phase space where the distribution of τ is or is not statistically indistinguishable from an exponential according to the K–S test. At intermediate values of area density, for ensembles denoted with '+', the distribution is indistinguishable from the exponential. At low and high area density values, for ensembles denoted with 'o', the distribution deviates from the exponential.

However, the p-value of the K–S test does not reveal *how* the empirical τ distribution differs from an exponential distribution. To gain some insight into this issue, we note that the standard deviation of an exponential distribution is equal to its mean. We therefore computed the mean μ_τ and the standard deviation σ_τ of the empirical τ distribution, and then examined the quantity $I_\tau \equiv (\sigma_\tau/\mu_\tau) - 1$. If τ were exponentially distributed, we would have $I_\tau = 0$. If $I_\tau < 0$ the empirical τ distribution would be more peaked around its mean than an exponential distribution would be, and if $I_\tau > 0$ the empirical τ distribution would be more spread out than an exponential distribution.

The grayscale portion of Fig. 3 gives a plot of I_τ, again for $N = 6$ and $\sigma = 1$, in the (L, l)-plane. The plot shows where the distribution of τ is more spread out or less spread out than an exponential. One possible interpretation of this plot is the following: a large portion of the distribution of the τ sample is exponential, but superimposed on that is a component that ruins the overall exponentiality. In areas with $I_\tau < 0$ the values in the added component are close to the mean, i.e. small. In areas with $I_\tau > 0$ the values in the added component are larger than the mean. The latter component, in fact, contains some very large outliers.

This can be seen by comparing the hypothesized analytical exponential, whose mean is given by the inverse of the coefficient of dt in Eq. (3), to the histogram of the τ data. An instance of a system for which $I_\tau > 0$ is the "small" molecule case $\sigma = 1, N = 6, L = 35, l = 0.1$. The mean for the data is ~191 and the standard deviation is ~204. Fig. 4 breaks down the pdf of the τ distribution into three segments for clarity (note the different scales on the vertical axes). For shorter times (top plot), the empirical τ distribution (solid line) follows our analytical prediction (dashed line) very well. Approximately two standard deviations to the right of the mean (~600, middle plot), the data becomes heavier than that model prediction, and is better described by an exponential with the empirical mean (dotted line). Finally, at long times (bottom plot), the data is heavier than both exponentials. So it appears that, in the $I_\tau > 0$ regime, the data follows the analytical exponential at shorter τ values, but has heavier than exponential tails for long τ values.

To better understand this behavior, we looked at some of the realizations that contributed the outlier τ samples, i.e. very long times. The trajectories of the molecules in those systems were notable for their non-ergodicity; more specifically, there

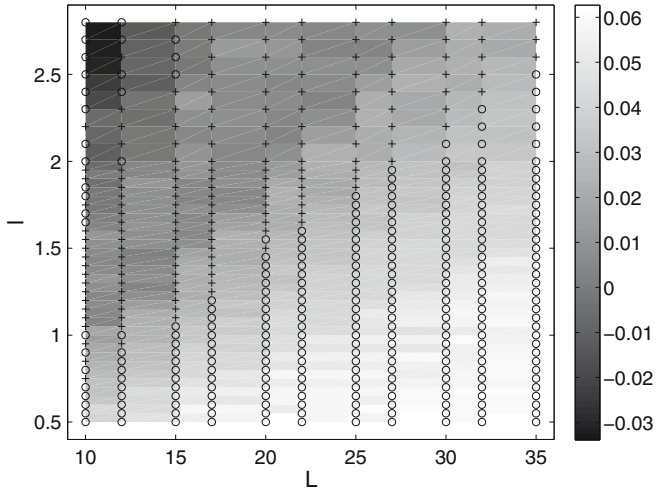

Fig. 3. Results of K–S test for exponentiality of the τ-distribution at significance level $\alpha = 0.05$, superimposed on a grayscale plot of the quantity $I_\tau \equiv (\sigma_\tau/\mu_\tau) - 1$. The population is $N = 6$ for all samples; the standard deviation of the Cartesian velocity components is $\sigma = 1$ for all samples; the vertical axis is l, the diameter of one molecule, and the horizontal axis is L, the diameter of the system. '+' denotes a τ-distribution which is indistinguishable from exponential (p-value > 0.05), while 'o' denotes a τ-distribution which deviates from the exponential (p-value < 0.05). Positive I_τ values imply a distribution with more spread than an exponential, while negative values imply a distribution with less spread than an exponential.

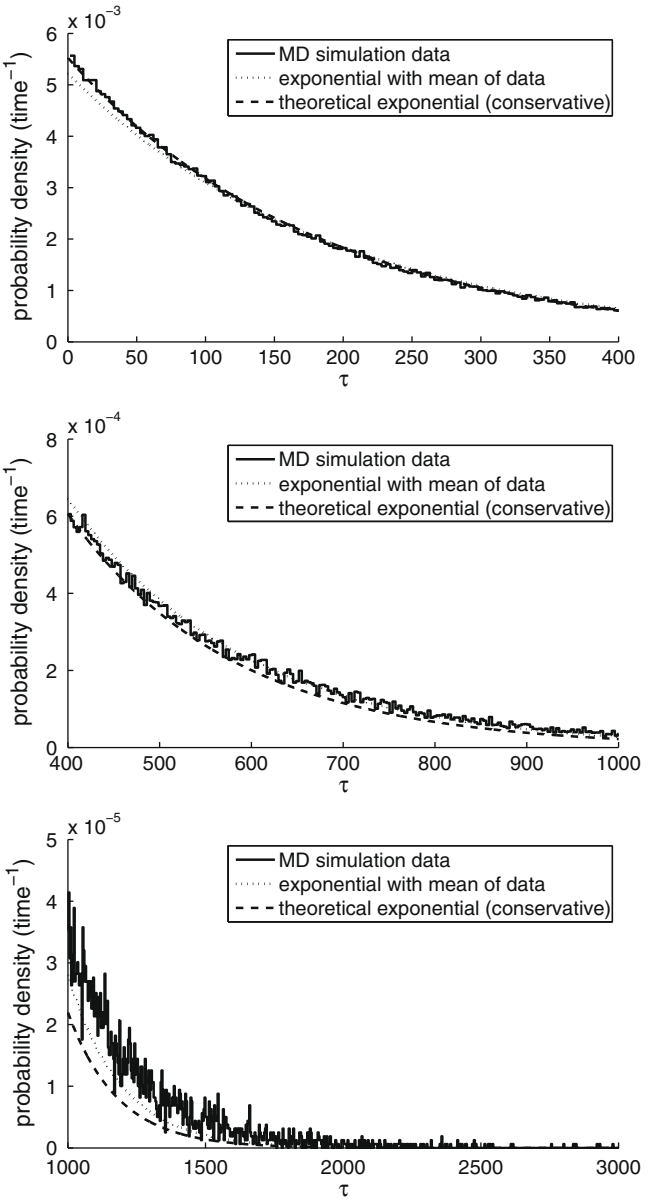

Fig. 4. Piecewise histogram of the τ distribution for $\sigma = 1, N = 6, L = 35, l = 0.1$ vs. analytical pdfs of two exponentials. The dotted exponential has the same mean as the τ distribution ($\sim$191). The dashed exponential is the most conservative theoretical prediction (with $V_e = V_d = Na$, where a is the area of one molecule).

would frequently be one or more molecules that moved along the circle boundary (a "whispering gallery" mode [8]), or that crossed very close to the center (and therefore bounced back and forth along the diameter). These trajectories are non-ergodic in the sense that the molecules are not sampling the entire container volume. This has the effect of reducing the number of potential pairings of molecules, which is small to begin with, thereby reducing the probability that *some* pair will collide. In this way, the *shape* of the container boundary combined with the small number of molecules contributes to the increased number of long-time outliers in the τ distribution.

One is also led to suspect the boundary's role from the observation that the number of boundary collisions preceding each inter-molecular collision increases dramatically as one approaches the point-molecule limit – something that does *not* happen in a one-dimensional system. To test the hypothesis that the boundary contributes to the over-represented long times, we tried several modifications to the hard circular boundary that we initially studied.

One class of modifications left the shape of the boundary intact, but changed how the molecules were reflected after they struck it. Since our hypothesis is that specular reflections from the circular boundary tend to orchestrate the molecules' trajectories in such a way that they do not always sample the entire space, we modified the reflection formulas in several ways,

in an attempt to increase the randomness in the trajectory of the molecules as they depart the boundary. First, we completely randomized the departing velocity of the molecule after a collision, while keeping the speed constant. Surprisingly, this had the effect of further lengthening the mean collision time. We next tried adding a small random angle to the departure angle after a reflection. This had the effect of very slightly shifting the mean collision time towards the model mean. Both of these were ways that intuitively seemed to us as though they would increase the ergodicity of the trajectories in the system; yet, they yielded very inconclusive results.

The other class of modification that we tried retained specular reflection but changed the *shape* of the boundary to one which is supposed to discourage non-ergodic trajectories. More specifically, we moved from a circular boundary to a "Bunimovich stadium" [6]. This is simply the interior of a boundary made by joining two semi-circles of diameter L with a square of side L in between. This boundary shape is supposed to create fewer non-ergodic trajectories than either a square or a circle boundary alone. Indeed, we observed that it affected the τ distribution by moving its mean slightly, but noticeably, closer to the analytically predicted mean, and also the standard deviation closer to the mean (as it should be for an exponential). For instance when comparing two 400,000 run ensembles with $\sigma = 1$, $N = 6$, $l = 0.1$, one in a circular container of diameter $L = 30$, $A = 706.9$ and the other in a stadium with $L = 19.9$, $A = 707.0$, we found a mean τ of 140.9 for the circle and 138.9 for the Bunimovich stadium, along with a standard deviation of 149.5 for the circle and 145.9 for the stadium. The analytical mean

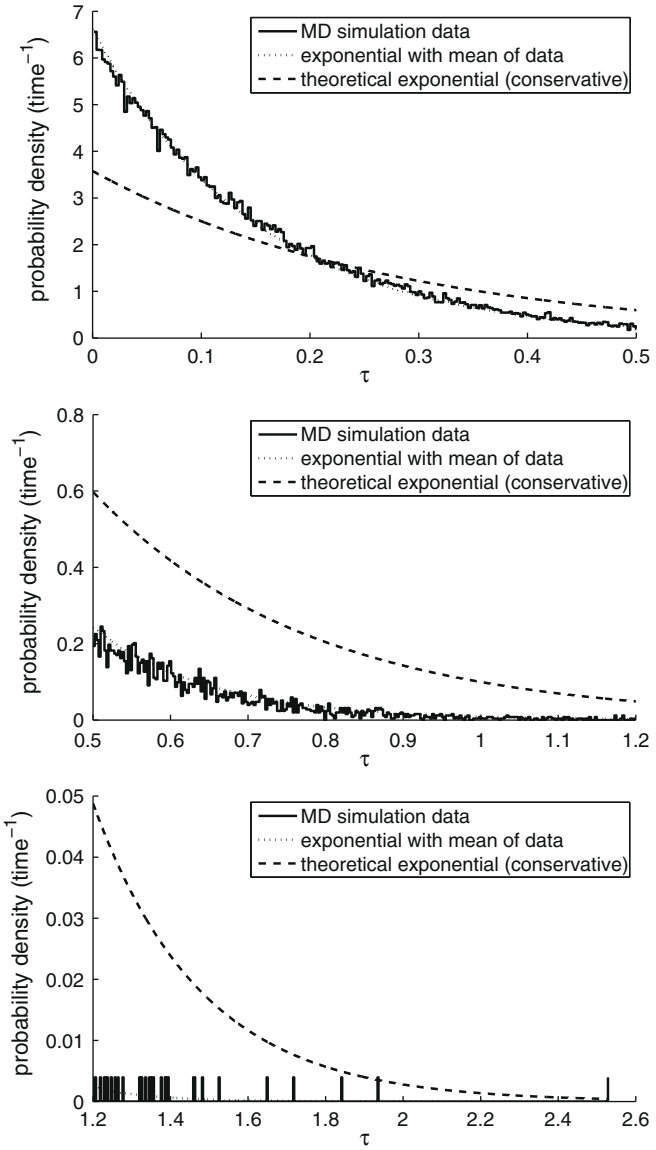

Fig. 5. Piecewise histogram of the τ distribution for $\sigma = 1, N = 6, L = 10, l = 2.8$ vs. analytical pdfs of two exponentials. The dotted exponential has the same mean as the τ distribution ($\sim$0.1511). The dashed exponential is the most conservative theoretical prediction (with $V_e = V_d = Na$, where a is the area of one molecule).

and standard deviation of τ for a volume of this area are 132.9. We confirmed that the effect is general by noticing that the p-values for the K–S exponentiality test imply that the test is much closer to accepting the distributions as exponential for Bunimovich stadium volumes than for circular volumes of the same area.

It seems to us that whether a boundary's shape and reflection characteristics contribute to non-ergodic trajectories is not a simple yes-or-no question, but rather one of degree. Some boundaries will cause molecules' trajectories to sample the container's area in a shorter amount of time than others. Our simulations would be most affected by the degree of ergodicity within a set finite amount of time (the expected time to a collision), not "eventually". Thus, we suspect that there is no such thing as an "ideal" boundary that would completely eliminate the effects of non-ergodic trajectories. If this is so, then it is not surprising that the Bunimovich stadium did not completely abolish the long-time tails; instead, it is satisfying that it produced a measurable change in the expected direction.

At the other extreme of the area density, we see a different picture. An example system for which $I_\tau < 0$ is given by the "large" molecule case $\sigma = 1, N = 6, L = 10, l = 2.8$. The mean for the data is $\sim$0.1511 and the standard deviation is $\sim$0.1460. Fig. 5 gives a breakdown of the pdf of the τ distribution for those parameters. The most noticeable feature of the pdf is the complete mismatch between the model exponential curve (dashed line) and the data (solid line). This is due to the fact that the model exponential (dashed line) is computed using the very conservative estimate that $V_e = Na$. But we see that the exponential curve with the empirical mean (dotted curve) follows the data (solid line) rather well. The major deviation in that regard is that at about two standard deviations to the right of the mean ($\sim$0.5), the data slightly undershoots the dotted exponential curve, which has the effect of biasing the mass of the distribution closer to the mean.

4.3. Collision time distribution for large population

We have shown that in the low molecule population and low area density case the small number of molecules conspires with the non-ergodic boundary to introduce long-time outliers in the τ-distribution. Fig. 6 shows that increasing the number of molecules while maintaining the low area density abolishes this effect, restoring exponentiality to the distribution.

But is it also the case that the low population but *high* area density non-exponentiality, which we described in the previous section, can be abolished by increasing the number of molecules? A definitive answer to this question could be given if our exact simulation methodology were tractable on dense, high N systems. However, the rejection-based Monte Carlo initialization of the positions of the molecules takes prohibitively long to complete for such systems. A *tentative* answer, which is hopefully a hint in the right direction, can be obtained using the *pre-stirring* based molecule position initialization routine. Fig. 7(a) shows that increasing the number of molecules while maintaining a high area density restores exponentiality to the τ-distribution. (We will discuss Fig. 7(b) and (c) in the next section.)

4.4. Excluded volume

Our initial goal in this effort was to investigate the effect of reactant-excluded volume on the kinetics of the $A + A \rightarrow$ *products* reaction in our simulation experiments. The most straightforward way to estimate the excluded volume felt by the molecules in these experiments is to proceed as follows: Assume the distribution of τ to be exponential; then estimate the propensity P as the inverse of the mean μ_τ of the empirical τ-distribution; finally, compute the *effective excluded volume* by solving the following equation for V_e:

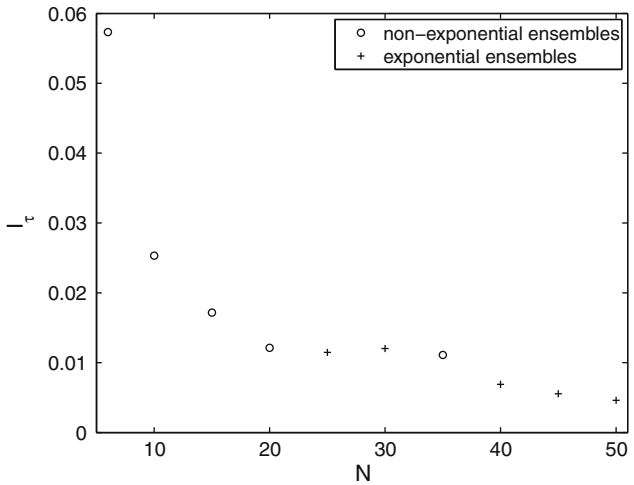

Fig. 6. Plot of the indicator $I_\tau \equiv (\sigma_\tau/\mu_\tau) - 1$ for fixed low area density 0.5%, $L = 10$, and varying N. '+' denote τ-distributions which are indistinguishable from exponential according to the K–S test at significance 0.05, while 'o' denote τ-distributions which are non-exponential.

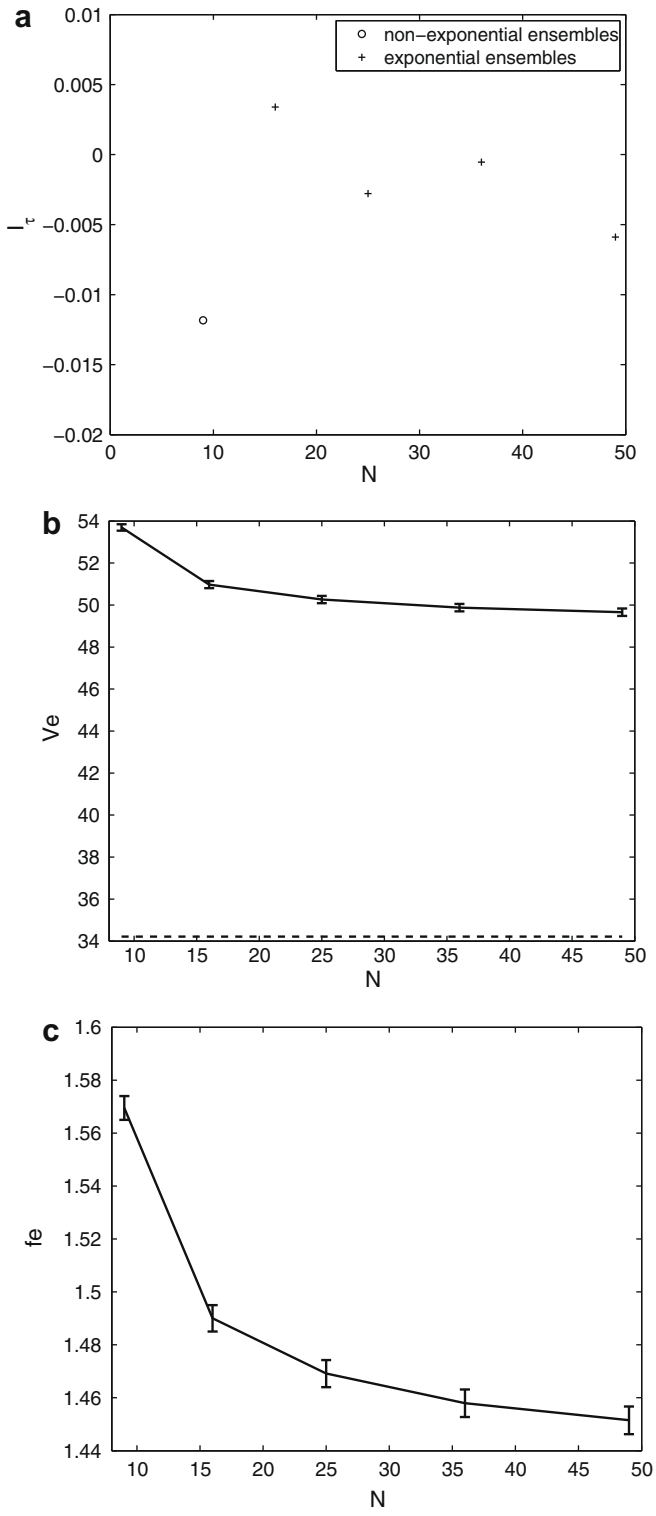

Fig. 7. (a) Plot of the exponentiality according to the K–S test, and the indicator $I_\tau \equiv (\sigma_\tau/\mu_\tau) - 1$ for fixed high area density ∼42%, $L = 10$, and varying N. '+' denote τ-distributions which are indistinguishable from exponential, according to the K–S test at significance 0.05, while 'o' denote τ-distributions which are not exponential by the K–S test. The y-coordinate of the '+' and 'o' gives the I_τ value for each distribution. Plot of the effective excluded volume V_e (b), and inverse packing fraction f_e (c), as a function of N, for the same fixed high area density ∼42%, and $L = 10$. The dashed line in (b) gives $V_d = Na$.

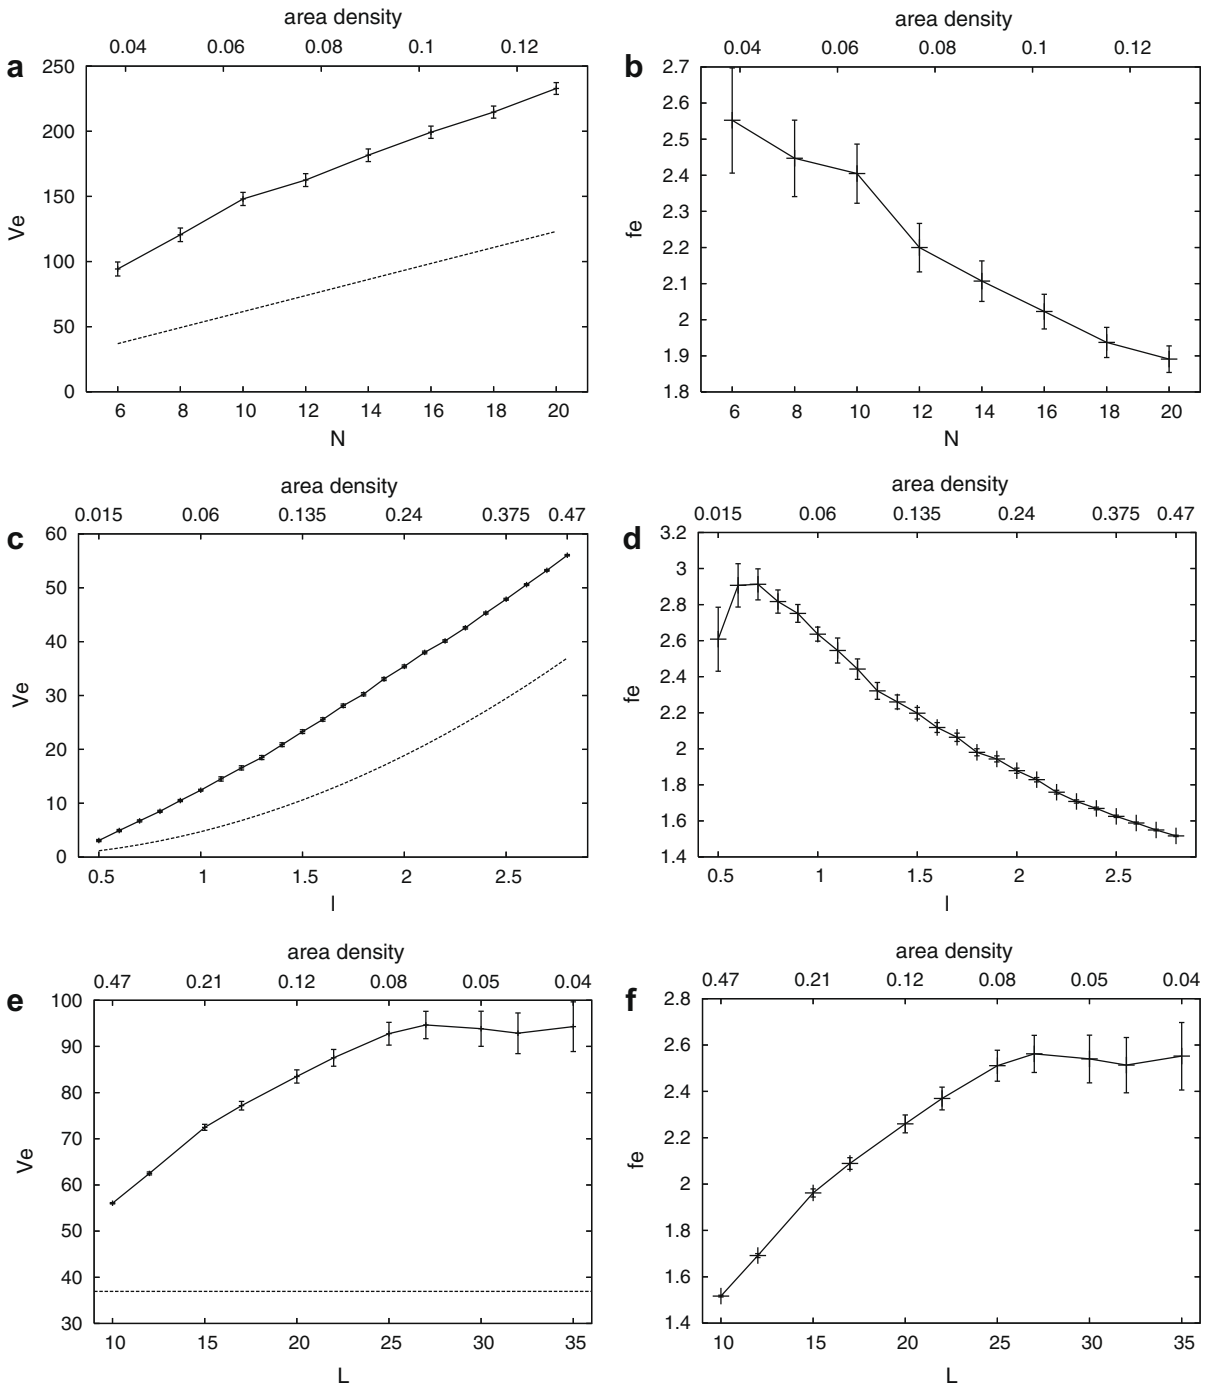

Fig. 8. Plots of the effective excluded volume V_e (left) and inverse packing fraction f_e (right), as a function of the three parameters L, l, and N, keeping two parameters fixed, and varying the other. For all plots we have $\sigma = 1$, ensembles of size 100,000, and 95% confidence intervals. The dashed lines give $V_d = Na$. In the top plots we vary N, the number of disks, while maintaining $L = 35, l = 2.8$; in the middle plots we vary l, the diameter of the molecule disks, while maintaining $N = 6$, $L = 10$; in the bottom plots we vary L, the diameter of the container circle, while maintaining $N = 6$, $l = 2.8$.

$$\frac{1}{\mu_\tau} = P = \frac{N(N-1)}{2}\frac{2ls_{\mathrm{rel}}}{A - V_e}. \tag{4}$$

The effective excluded volume computed this way can then be compared to several "theoretically plausible" excluded volume formulas. The simplest such formula takes into account only the area excluded by the molecule discs themselves: $V_d = Na$.

Another possibility would be to consider volumes of the form $V_i = Na/\phi_i$, where ϕ_i is a *packing fraction*. Some packing fractions that have been theoretically studied by others are: $\phi_{cp} = \pi/2\sqrt{3} \approx 0.90$, the "close packing" fraction; $\phi_f = 0.69$, the "freezing" packing fraction; and $\phi_c = 0.82$, the "random close" packing fraction [5]. It should be noted that in one dimension, all three of these packing fractions are equal to 1; therefore, if we were to find that in two dimensions *any one* of these fractions is the desired factor, the theory would limit nicely to the lower dimensional result.

Since all the proposed excluded volumes above are of the form $V = fNa$, with f the inverse of a packing fraction, a reasonable quantity to visualize is the *effective inverse packing fraction*, $f_e = V_e/Na$. Higher f_e is associated with looser packings, i.e. higher per-molecule excluded volume. The inverse "close packing" fraction is $f_{cp} \approx 1/0.90 \approx 1.1$; the inverse "freezing" packing fraction is $f_f \approx 1.45$; and the inverse "random close" packing fraction is $f_c \approx 1.22$.

The three parameters of the simulation are N, the number of molecules, L, the diameter of the circular container, and l, the diameter of the molecules. In each of the plots in Fig. 8, we keep two of these parameters fixed and vary the other. In the left plots we show V_e (solid line), as estimated from solving Eq. (4), along with $V_d = Na$ for comparison (dashed line). In the right plots we show f_e. Note that f_e is the ratio of the solid and dashed lines from the left plots.

The first two left plots (in which we vary N and l) show that the effective excluded volume is higher than just the disk volume, as expected. But does the effective excluded volume correspond to some packing fraction? The first two plots on the right address just that question: it seems that no constant packing fraction can account for the effective excluded volume we observe, *across the whole range* of area densities and populations. The effective inverse packing fraction, f_e, decreases with increasing area density and with increasing population (see Fig. 7(b) and (c), as well, for a speculative result based on the pre-stirring initialization of positions).

This implies that the excluded volume situation for finite, reflective boundary containers is not as simple in two dimensions as it is in one dimension. Since the effective excluded volume is a function of the mean μ_τ of the empirical τ-distribution (Eq. (4)), the reason why the packing appears more compact as the number of molecules increases must relate back to the τ-distribution. But we have already shown that in situations with low population, the τ-distribution contains artifacts introduced by the reflective boundary, which ruin its exponentiality, and which obviously impact its mean. So it is reasonable that in situations where the τ-distribution is not exponential, the packing estimated from the empirical is surprising, in this case for its looseness.

At high population and intermediate density, e.g. $N = 20$ and area density $\sim$12%, at the right end of Fig. 8(b), the packing reaches $f_e = 1.8$. At high population and high area density (*but* simulated with pre-stirring), e.g. $N = 50$ and area density $\sim$42%, at the right end of Fig. 7(c), the packing reaches $f_e = 1.45$, the freezing packing fraction. So, it appears that for densities at which the excluded area is a significant portion of the area of the system, at high population, the effective excluded area is reasonably close to close packed.

5. Summary and discussion

We have used computer simulations of hard disk dynamics to study the effect of reactant size on the rates of inter-molecular collisions in the ballistic setting. We have found that the distribution of collision times is close to exponential. It can be thought of as mostly exponential, except at low population, where we observe an additional mode which depends on the area density of the system.

At low population and low area density, non-ergodic trajectories contribute to longer than expected inter-molecular collision times. At low population and high area density shorter inter-molecular collision times are over-represented. We conclude, on the basis of exact simulations, that increasing the population abolishes the low area density effect; on the basis of approximate simulations, we speculate that the same is true at high area density.

At intermediate and high area densities, the volume excluded by the reactants ranges from about 2.5 times the area of the molecules' disks, at low population or low area density, to near the close-packed area, at high population or high area density.

In spite of differences in the details of the formula for the effective excluded volume in different population and area density regimes, it is clear that the probability of collision, and hence the reaction's propensity function, will be higher in a system with large molecules, compared to a system with the same number but smaller molecules.

Acknowledgments

Support for one of the authors (DTG) was provided by the University of California under Consulting Agreement 054281A20 with the Computer Science Department of its Santa Barbara campus; and by the California Institute of Technology, through Consulting Agreement 21E-1079702 with the Beckman Institute's Biological Network Modeling Center, and through Consulting Agreement 102-1080890 with the Control and Dynamical Systems Department pursuant to NIH Grant R01 GM078992 from the National Institute of General Medical Sciences, and through Contract 82-1083250 pursuant to NIH Grant R01 EB007511. Support for two of the authors (SL and LRP) was provided by the U.S. Department of Energy under DOE award No. DE-FG02-04ER25621; by the National Science Foundation under NSF awards CCF-0428912, CTS-0205584, CCF-326576; by the NIH under grants GM075297, GM078993 and EB007511; and by the Institute for Collaborative Biotechnologies through grant DAAD19-03-D-0004 from the U.S. Army Research Office.

References

[1] D.T. Gillespie, Stochastic simulation of chemical kinetics, Annu. Rev. Phys. Chem. 58 (2007) 35.
[2] D.T. Gillespie, A rigorous derivation of the chemical master equation, Physica A 188 (1992) 404.
[3] D.T. Gillespie, S. Lampoudi, L.R. Petzold, Effect of reactant size on discrete stochastic chemical kinetics, J. Chem. Phys. 126 (2007) 034302.
[4] L. Tonks, The complete equation of state of one, two and three-dimensional gases of hard elastic spheres, Phys. Rev. 50 (1936) 955–963.
[5] S. Torquato, Nearest-neighbor statistics for packings of hard spheres and disks, Phys. Rev. E 51 (1995) 3170–3182.
[6] L.A. Bunimovich, On the ergodic properties of nowhere dispersing billiards, Commun. Math. Phys. 65 (1979) 295–312.
[7] D.C. Boes, F.A. Graybill, A.M. Mood, Introduction to the Theory of Statistics, third ed., McGraw-Hill, New York, 1974.
[8] <http://www.chm.bris.ac.uk/~chjpr/Research/WGM.htm> (accessed 21.08.07).