First-Passage Monte Carlo Algorithm: Diffusion without All the Hops

Tomas Opplestrup,[1,2] Vasily V. Bulatov,[1] George H. Gilmer,[1] Malvin H. Kalos,[1] and Babak Sadigh[1]

[1]*Lawrence Livermore National Laboratory, University of California, Livermore, California 94551, USA*

[2]*Royal Institute of Technology (KTH), Stockholm S-10044, Sweden*

(Received 16 September 2006; published 4 December 2006)

We present a novel Monte Carlo algorithm for N diffusing finite particles that react on collisions. Using the theory of first-passage processes and time dependent Green's functions, we break the difficult N-body problem into independent single- and two-body propagations circumventing numerous diffusion hops used in standard Monte Carlo simulations. The new algorithm is exact, extremely efficient, and applicable to many important physical situations in arbitrary integer dimensions.

DOI: 10.1103/PhysRevLett.97.230602 PACS numbers: 05.10.Ln, 02.60.Cb, 02.70.Tt, 02.70.Uu

A kinetic Monte Carlo (KMC) method is often applied to model systems that evolve through collisions among multiple random walkers. The range of applications is exceptionally wide, including diffusion-controlled chemical and biochemical reactions, recombination of carriers in semiconductors, surface growth, diffusion-limited aggregation, Ostwald ripening, and diverse studies like epidemiology and population dynamics. Its transparency, direct connection to the underlying physical mechanisms, and simplicity of implementation, have made KMC calculations a method of choice in computational science and its appeal has increased steadily over time [1]. Regrettably, the method's practical applicability is severely limited when the density of particles, or walkers, is so low that it takes many diffusion hops for any two walkers to find each other. In such important situations, much work is wasted while walkers explore space whereas the interesting events occur very infrequently. Several methods have been proposed to propagate particles more effectively using the exactly known statistics of collisions between two diffusing particles [2–5]. Unfortunately, the existing methods rely on approximations whose accuracy needs to be carefully assessed.

We present a new general KMC method for diffusion simulations that is both computationally efficient and approximation free. The method relies on time dependent Green's functions to propagate the walkers over long distances in a single KMC step. The need for approximations is removed by partitioning the space into nonoverlapping "protective domains" (PDs), each containing one walker, and then using appropriate first-passage Green's functions [6] to propagate walkers to the boundaries of their respective PDs, one walker at a time. Thus, numerous diffusive hops are replaced by large "superhops" guaranteeing that the statistics of random walks remains unaltered. This new method carries an additional computational overhead that is more than offset by a dramatic increase in propagation efficiency. The new method extends the applicability of KMC simulations well beyond their perceived computational limits, while maintaining exact statistical equivalence to existing methods.

The theory of first-passage processes treats diffusive propagation of particles to a boundary [6]. In a continuum-space continuum-time random walk starting at time $t = 0$ from position $\mathbf{r}_0 = 0$ inside a simply connected spatial domain Ω with boundary $\partial\Omega$, the probability that the walker has not reached $\partial\Omega$ by time t is

$$P(t) = \int_\Omega c(\mathbf{r}, t)d\mathbf{r},$$

where $c(\mathbf{r}, t)$ is the solution (Green's function) of the diffusion equation in Ω, $\partial_t c = D\Delta c$ with the initial condition $c(\mathbf{r}, 0) = \delta(\mathbf{r} - \mathbf{r}_0)$ and the absorbing boundary condition $c(\partial\Omega, t) = 0$. The latter condition means that, when the walker arrives at $\partial\Omega$, it disappears from Ω. The conditional probability density of finding the walker at point $\mathbf{r}$ at time t provided it is still inside Ω is simply $\frac{c(\mathbf{r},t)}{P(t)}$, while the probability that when the walker arrives to boundary $\partial\Omega$ at time τ it will exit Ω at point $\mathbf{r}^*$ is

$$\rho(\mathbf{r}^*, \tau) = \frac{\nabla c(\mathbf{r}^*, \tau) \cdot \mathbf{n}}{\int_{\partial\Omega} \nabla c(\mathbf{r}, \tau) \cdot \mathbf{n} dS}, \qquad \mathbf{r}^* \in \partial\Omega,$$

where the numerator is the density of probability current in the direction of the outward normal $\mathbf{n}$ to the boundary surface $\partial\Omega$, and the denominator is the total (integrated) outward probability current at time τ.

For domains Ω of simple geometries, e.g., spheres or cubes, simple solutions for the first-passage statistics are available to propagate walkers through space efficiently. The famous walk-on-spheres algorithm for simulations of diffusion-limited aggregation (DLA) is an example [7]. So far, the use of first-passage theory in Monte Carlo simulations has been limited to single walker propagation viewed as a time independent process [8,9]. The only exceptions known to us are the use of time dependent Green's functions for Monte Carlo calculations of the ground state of bosons [10,11]. In their algorithm [10], Kalos, Levesque, and Verlet (KLV) "protected" each hard-sphere walker by a larger sphere concentric with the walker position and then, for each walker independently, sampled the time of first arrival to the surfaces of their protective spheres from

appropriate distributions. Given that the protective spheres were constructed to be nonoverlapping, the statistics of each walker was guaranteed to be independent of any other walker during the time it takes for the first walker to reach its protective boundary (see Fig. 1). Thus, the KLV algorithm entails: (1) drawing nonoverlapping protective spheres, (2) sampling times of first arrival for each walker $\tau_1, \tau_2, \ldots, \tau_N$, (3) finding the shortest time of first arrival $\tau_{\min}$, (4) moving the respective walker to a randomly selected point on its protective sphere, (5) advancing the clock by $\tau_{\min}$, and (6) updating positions of all other $N-1$ walkers by sampling from the conditional probability distribution $c(\bar{r}, t)/P(t)$. Then, the cycle continues from step (1) to step (6). In KLV, first-passage solutions are used in steps (2) and (6).

Here we propose a seemingly minor but crucial modification of the KLV algorithm. First, not just the shortest time of first arrival $\tau_{\min}$ is selected, but a specially ordered list (heap) of all N first arrival times is maintained. As in KLV, on every cycle we sample a move associated with the PD having the shortest arrival time, but do not update positions of the remaining $N-1$ walkers. We redraw only the protective sphere for the walker just propagated so as not to overlap with any other of $N-1$ spheres, randomly sample a new first arrival time for this walker, and insert it into the existing list. The statistics of N random walkers remain independent of each other. By reducing the computational complexity from $O(N)$ in steps (1), (2), and (6) of the KLV algorithm to $O(1)$, this modification makes a dramatic difference in the complexity required to advance the clock. The new algorithm treats only one walker on each step rather than all N. This results in a spectacular speedup of KMC simulations while retaining the exact statistics.

Figure 2 demonstrates that the method is indeed accurate and efficient. This is a simulation of the diffusion-controlled reaction of irreversible coalescence $A + A \Rightarrow A$ on a discrete $1d$ lattice. On every MC step every particle hops to the right or to the left with equal probabilities and whenever two particles are found to occupy the same lattice site one of them is eliminated. We performed this simulation in two different ways: first using the standard KMC procedure [1] and then using the new algorithm in which on every MC step only one particle is advanced to the boundary of its PD (a segment in $1d$). Figure 2(a) shows the kinetics of this reaction averaged over 70 independent realizations. The two kinetics coincide within (small) statistical uncertainties. At the same time, it takes on average 5×10^9 particle hops to reach the end of this reaction using the standard KMC algorithm whereas the new algorithm accomplishes the same in just 4×10^4 steps. Similarly impressive agreement and speedup is achieved in simulations of a similar annihilation reaction $A + A \Rightarrow 0$, but this time performed in continuous space and continuous time in three dimensions [Fig. 2(b)].

Several additional issues had to be addressed to make the new algorithm practically efficient. One potentially troublesome situation arises when one of the walkers propagates itself into a small region of space "squeezed" between other PDs. In such cases, the size of the new PD that can be drawn around the walker after its propagation can be very small, limited by the surrounding PDs. The new first arrival time for the squeezed walker is then likely to be the shortest so that the same walker will be selected again and again for propagation. To improve efficiency, when such a situation is detected, the algorithm updates positions of one or several neighboring walkers whose PDs bound the squeezed walker, as in step 6 in the KLV algorithm described earlier: in the statistical distribution used for such an update, time is counted from the instant when the walker selected for an update was last propagated or updated. Such a partial repartitioning typically frees up more space for the squeezed walker. On the other hand, when two walkers are detected to be considerably closer to each other than to any other walker, it is advantageous to protect the two close walkers as a pair and to use an appropriate Green's function to propagate such a pair either to collision or to separation. A more detailed de-

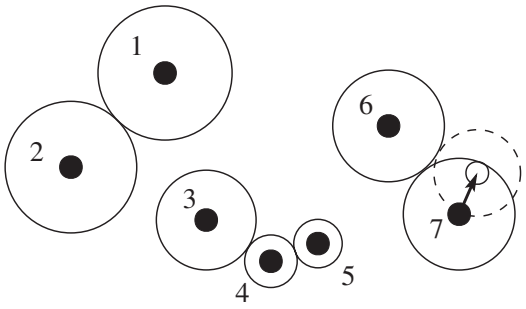

FIG. 1. Diffusing particles (small closed circles) numbered 1 to 7 are protected by nonoverlapping PDs (larger open circles). Particle 7 is propagated to the surface of its PD because its arrival time (τ_7) is shortest among all particles. It is moved to a random position on its protective circle, a new nonoverlapping protective circle is drawn around its new position and the time is advanced to τ_7.

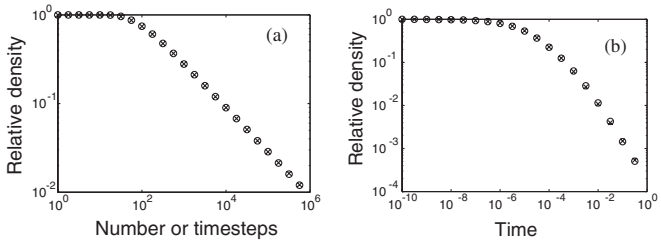

FIG. 2. Comparison between the standard KMC calculations (crosses) and the first-passage KMC (circles) results. (a) Discrete-time discrete-space simulations of $A + A \Rightarrow A$ coalescence reaction in 1D. (b) Continuous-time continuum-space simulations of $A + A \Rightarrow 0$ annihilation reaction in 3D. Note that each circle is filled with a cross, indicating agreement.

scription of collision handling and other algorithmic elements will be presented in a future publication.

We have thus far discussed our new algorithm when spherical walkers are protected by concentric spheres. However, the choice for the PDs shape is arbitrary. A particularly simple implementation arises when cubical walkers are protected by cubical domains. Then the evolution of a system with N walkers in d dimensions is equivalent, in most aspects of algorithmic implementation, to a simpler problem of Nd walkers in 1D. In particular, products of simple one-dimensional Green's functions can be used for propagation of the walker pairs to collisions. Green's functions for sampling binary collisions between spherical walkers are somewhat more complicated. It is possible to use combinations of PDs of different shapes. More generally, it is possible to exercise the freedom to partition the available space so as to optimize the overall performance of the algorithm by balancing the expense of making the random propagation steps and that of redrawing the PDs.

The efficiency of the new algorithm allows very large and very long KMC simulations for many diffusion-limited processes. A typical time to complete a simulation of $A + A \Rightarrow 0$ reactions among 10^8 particles in 1D is a few hours on a single CPU workstation. The previous largest simulation of the same reaction involved 16×10^6 particles, entailed considerably larger computational effort and used an approximate method [5,12]. In our simulations, the computational efficiency remains very high, taking on average 10.4 propagations between collision events. In 1D, the simulation efficiency is independent of the particle density but in higher dimensions the efficiency decreases slightly with decreasing density. A simulation of the same reaction with 125×10^6 particles in 3D took 36 hours on a single CPU workstation. Here the average number of single-particle propagations between subsequent collisions was 33.6. As shown in Fig. 3(a), the running average of this number increases inversely proportional to the cube root of the density of remaining particles. This is compensated, in the late stages of reaction, by the reduced calculations when the number of particles becomes small. Consequently, the overall efficiency remains high in 2D, 3D, and higher dimensions. It would have taken tens of years to complete the same 125×10^6 particle simulation in 3D using standard KMC calculations. [Figure 3(a) compares the computational efficiency of the new method and standard KMC calculations].

To further enhance the efficiency of our new algorithm in simulations of annihilation, coalescence, and similar reactions, we combine it with volume replication [13]. In a 1D simulation, as soon as the number of particles in the periodic simulation volume (line segment in 1D) decreases by a factor of 2, we double the volume by adding to the surviving particles their images in a neighboring periodic replica of the previous volume. Thus, the number of particles returns to its initial value after the replication. In dimension d, the same recipe requires waiting until the

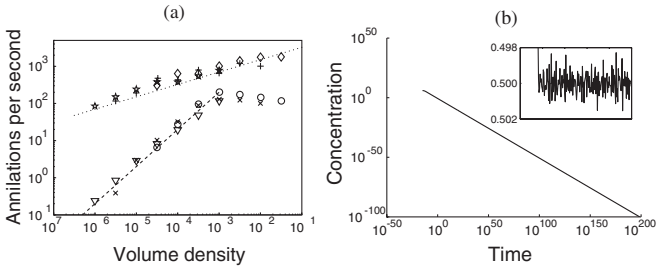

FIG. 3. (a) Computational performance as a function of particle density measured in six independent realizations of $A + A \Rightarrow 0$ annihilation reaction in 3D performed on a single CPU workstation using the standard KMC calculations with a finite hop distance (upper symbols) and the new algorithm (lower symbols). The dashed and dotted lines are theoretical estimates with slopes of 1 for the standard KMC calculations and $1/3$ for the new algorithm. (b) Kinetics of $A + A \Rightarrow 0$ annihilation reaction in 1D, using replication. The inset is the logarithmic derivative of the kinetic curve taken over the same time interval reproducing the theoretical exponent $\alpha = 0.5$ for this reaction.

number of particles decreases by a factor 2^d. The new replicated particles are treated as independent walkers after the replication. For such replications to preserve correct statistics, one must note that the correlation length in the system may grow with time and that the size of the simulation volume must remain larger than the correlation length. Given that correlations propagate by diffusion, an upper bound on the rate of growth of the correlation length is $\sim\sqrt{Dt}$ where D is the diffusion coefficient. Conservatively, the volume should be replicated whenever $\sqrt{Dt}$ grows to become an appreciable fraction of the linear box size L. Hence, the physical time t elapsed between replications should at least quadruple with each replication. In a particular case of $A + A \Rightarrow 0$ (or $A + A \Rightarrow A$) reaction in 1D, the asymptotic kinetics for particle concentration is $c(t) \sim t^{-1/2}$. The number of particles can be allowed to decrease by half between subsequent replications and does not have to grow at all in order to avoid propagation of correlations through the expanding volume. Remarkably, such a simulation can be continued over an arbitrary long physical time [see Fig. 3(b)].

In higher dimensions and for different reactions, the situation may be less favorable. In dimension d, assuming that the asymptotic kinetics is $t^{-\alpha}$, the number of particles will have to increase with each replication by a factor $2^{d-2\alpha}$ unless, for some reason, the correlation length grows slower than $\sqrt{Dt}$. Thus, in higher dimensions and/or for smaller decay exponent α, the number of particles should grow over continued replications. In such cases, it may be better to start from a small number of particles and let the system grow by replications to a maximum size afforded by the computer memory. This opposes the common practice to start from a maximum size that fits into memory and then simulate the reaction to the end. As has been noted before [12], in their late stages such simulations do not reproduce the kinetics of interest due to the propagation of

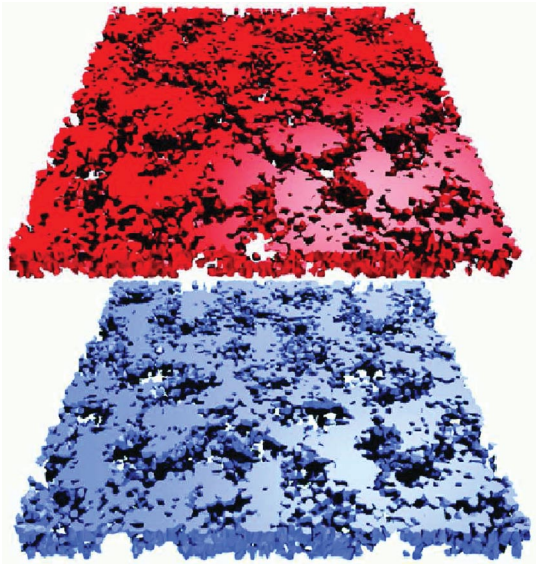

FIG. 4 (color). A thin slice through the domain structure formed in a simulation of $A + B \Rightarrow 0$ reaction in 3D. The boundaries of the A-rich (red) and B-rich (blue) domains are identified with the sides of Voronoi polyhedra shared by unlike particles. The two domains are complementary and fill the space when brought together.

the correlation length through the (fixed) simulation volume. The useful time horizon of even the largest simulations performed has heretofore been limited to 8–10 time decades. The replication trick makes it possible to significantly extend the time horizon of such simulations (in some cases to infinite time) at no additional cost.

The new method is quite general and trivially extendable to diffusion-limited processes in any number of dimensions, in complex confined geometries with absorbing or reflecting boundaries, anisotropic diffusion, simultaneous diffusion of slow and fast and/or large and small particles, diffusion with a constant bias, and various other extensions. As an example, consider the diffusion-limited two-species annihilation reaction $A + B \Rightarrow 0$ in 3D. As was noted in [14], the annihilation results in a gradual emergence of alternating A-rich and B-rich domains. Remarkably, unlike most other diffusion-limited processes in 3D, this reaction does not obey the mean-field kinetics when the number of A and B species is close to the stoichiometry: asymptotically, the concentration of species decays as $\sim t^{-3/4}$. One of important physical realizations of this reaction is the recombination of vacancies and interstitial defects produced in crystals by ion or electron irradiation [3]. It has not been possible to investigate in detail the long-term kinetics and geometry of domains emerging in reactions of this kind, precisely because the standard KMC calculations slow down dramatically when the density of remaining particles becomes low [14]. By contrast, the new algorithm presented here handles the same difficult cases with ease and in any number of dimensions. One such simulation of two-species annihilation reactions in 3D is shown in Fig. 4.

In summary, we present an extension of the diffusion Monte Carlo approach that overcomes severe computational limitations of standard KMC calculations when the density of diffusing particles is low. The new method is based on exact solutions for first-passage processes and, as a result, is exactly statistically equivalent to and, simultaneously, dramatically more efficient than the standard KMC calculations allowing for simulations of very large systems evolving over very long elapsed physical times on modest computers. The efficiency and accuracy of the new method is demonstrated here by simulations of well-studied diffusion-controlled reactions in a number of spatial dimensions. At the same time, the range of potential applications of the new first-passage algorithm is exceptionally wide given its ability to handle simulations of scale, length, and complexity that were impossible heretofore.

This work was performed under the auspices of the U. S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48. This work was supported by the Office of Laboratory Directed Research at LLNL and the Office of Basic Energy Sciences U. S. Department of Energy. The authors wish to express their gratitude to G. Martin, S. Redner, W. Cai, C. Mailhiot, D. Yi, and A. Donev for fruitful discussions.

[1] A. B. Bortz, M. H. Kalos, and J. L. Lebowitz, J. Comput. Phys. **17**, 10 (1975).
[2] J. Dalla Torre, J.-L. Bocquet, N. V. Doan, E. Adam, and A. Barbu, Philos. Mag. **85**, 549 (2005).
[3] C.-C. Fu *et al.*, Nat. Mater. **4**, 68 (2005).
[4] J. S. van Zon and P. R. ten Wolde, J. Chem. Phys. **123**, 234910 (2005).
[5] D. ben-Avraham, J. Chem. Phys. **88**, 941 (1988).
[6] S. Redner, *A Guide to First-Passage Processes* (Cambridge University Press, Cambridge, England, 2001).
[7] M. E. Muller, Ann. Math. Stat. **27**, 569 (1956).
[8] J. A. Given, J. B. Hubbard, and J. F. Douglas, J. Chem. Phys. **106**, 3761 (1997).
[9] M. Mascagni and N. A. Simonov, SIAM J. Sci. Comput. **26**, 339 (2004).
[10] M. H. Kalos, D. Levesque, and L. Verlet, Phys. Rev. A **9**, 2178 (1974).
[11] D. M. Ceperley and M. H. Kalos, in *Monte Carlo Methods in Statistical Physics*, edited by K. Binder (Springer-Verlag, Berlin, 1979).
[12] D. Zhong, R. Dawkins, and D. ben-Avraham, Phys. Rev. E **67**, 040101(R) (2003).
[13] M. Smith and T. Matsoukas, Chem. Eng. Sci. **53**, 1777 (1998).
[14] F. Leyvraz and S. Redner, Phys. Rev. Lett. **66**, 2168 (1991).