

The Complexity of N -body Simulation*

John H. Reif and Stephen R. Tate

Computer Science Department, Duke University, Durham, NC 27706

Abstract. The *n-body simulation problem* is stated as follows: Given initial positions and velocities of n particles that have pair-wise force interactions, simulate the movement of these particles so as to determine the positions of the particles at a future time.

In this paper, we give the first known n -body simulation algorithms with rigorous proofs of bounded error. The *reachability problem* is to determine if a specific particle will reach a certain region at some specified target time. In the case we require $poly(n)$ bits of accuracy and where the target time is $poly(n)$, our complexity bounds are surprisingly PSPACE.

We also have matching lower bounds for n -body simulation problem (in comparison all previous lower bound proofs required either artificial external forces or obstacles). We show that the reachability problem for a set of interacting particles in three dimensions is PSPACE-hard.

1 Introduction

The *n-body problem*, is the problem of simulating a set of n charged particles in three dimensions, where the particles interact under the induced electrostatic or gravitational potential field. Generally the simulations are done by time stepping. See [1, 6, 7, 8, 10] for details. These simulations are one of the heaviest users of super computer cycles (for example at the CRAY-YMP, at RTP, a study by MCNC recent showed that over 30 percent of all compute time was used for n -body simulation by molecular chemists), and are widely used by astronomers, chemists, and biochemists, and to a lesser degree physicists (note: certain physicists prefer other methods based on energy minimization).

The equations of motion for each body are in fact given by Newton's second law of motion applied to each body; this results in a system of n ordinary differential equations. These equations can be approximately solved from initial positions and velocities by stepping in time, using the equations of motion and numerical integration to determine approximations to incremental movements and velocity changes of the bodies due to the forces exerted by the other particles. The force vector associated with these potentials is calculated by taking the partial derivatives of the potential in each direction. This is the basis for most computer simulations of n -body systems. The main computational task

* This research was supported by DARPA/ISTO Contracts N00014-88-K-0458, DARPA N00014-91-J-1985, N00014-91-C-0114, NASA subcontract 550-63 of prime contract NAS5-30428, US-Israel Binational NSF Grant 88-00282/2, and NSF Grant NSF-IRI-91-00681.

Preprint of paper appearing in Proceedings of the 20th Annual Colloquium on Automata, Languages and Programming (ICALP'93), Lund, Sweden, July, 1993, p. 162-176.

is the calculation of the potential field due to all other bodies, at the current location of each body. The naive algorithm for this potential computation requires quadratic work; however, the potential can be approximated to p bits of accuracy in $O(np^2 \log p)$ time using the multipole method of Greengard and Rokhlin [4, 5], or by the recent modified multipole method of Reif and Tate [11] that has time complexity $O(np^2)$.

1.1 The n -body Reachability Problem

In this paper, we consider the complexity of simulating a set of n charged particles in three dimensions, where the particles interact under the induced electrostatic potential field.

Throughout this paper, n denotes the number of bodies. We require that the number of bits of the input description is polynomial in n . A k -bit rational is a ratio of two k -bit integers. Consider a set of n points satisfying a fixed electrostatic potential law. We assume that we are given an initial $poly(n)$ -bit rational position and velocity as well as a destination position, given by a ball, where the ball's position and radius is n -bit rational. The *n -body reachability problem* concerns the trajectories of these bodies; in particular we wish to determine if a given particle reaches a position within the given destination ball within a given time bound, where the ball's position and radius are $poly(n)$ -bit rationals.

We give the first known n -body simulation algorithms with rigorous proofs of bounded error. In the case we require $poly(n)$ bits of accuracy and where the target time is $poly(n)$, our complexity bounds are surprisingly PSPACE. Our algorithm requires the additional assumption that there is at least an exponentially small separation between all pairs of particles at all times during the simulation.

Molecular Computers, Molecular Castles, and Our Lower Bounds for the n -body Simulation Problem. We also give the first lower bounds for these simulations, and show that the *reachability problem* for n -body simulations is PSPACE-hard. We prove this lower bound for the most practical version of this problem: inverse-square law forces in three dimensions.

The hardness proof is via a reduction to machine simulation, and is novel due to the nature of the problem under study. In particular, non-trivial problems to overcome include the fact that machines work in discrete time steps and particle simulations are continuous, and the fact that realistic machines perform transitions based on local state whereas the particle simulations have the property that all particles induce a force on all other particles (so all effects are global). To our knowledge, this is the first hardness proof to overcome these problems.

The techniques involved in our lower bound proof include constructions in which the time-averaged potential of a small set of particles is almost identical to the potential due to solid, uniformly-charged plates which persists for exponential time. A side effect of this construction is a proof that given any set of polygons fixed at rational positions in 3D with f faces and sides (say a castle),

we can construct in 3D a stable configuration (for exponential time) of a set of $O(f)$ charged particles (which move according to Newton's and Coulomb's laws) which generate a time averaged very close approximation (with exponentially small error) to this set of polygons. The physical implications of such a construction is interesting in its own right.

Our PSPACE results for n -body reachability indicate that there is no polynomial time computable closed form representation of the equations for the motion or trajectory of n -body systems with above a certain constant number of particles n , unless $P = PSPACE$.

Theoretically, if the particles are placed with initial position and velocity of sufficient precision (polynomial bit precision will suffice), these PSPACE-hard n -body systems can be viewed as general computing machines executing in the real time at molecular sizes. However, we strongly caution that these n -body systems may not be practical, since the above assumptions of the classical laws of potential theory may not always hold in the physical world for small displacements. Our constructions can be scaled to as small a dimension as possible, and quantum mechanical effects minimized by using large mass, large charge particles at a small spatial scale — in this case, as the masses and charges of the particles increase, the motion of the particles approaches the motion of an ideal Newtonian system. While this is fine in theory, actual subatomic particles with fixed charge and mass do indeed exhibit quantum mechanical effects that deviate from the simple Newtonian force laws that we use in this paper.

Related Work in Hardness Results. Here we note that all previous hardness results for n -body simulation depend on sharp discontinuities in space (i.e., obstacles) or force in order for the proof to work. In contrast, our lower bound applies when the forces are realistic inverse-square law forces, and there are no obstacles present to produce discontinuities.

A related result has been obtained by Moore [9] who investigated unpredictability in dynamical systems. He showed that motion of a body with as few as three degrees of freedom in the presence of a fixed, immobile potential field can simulate universal computation. His construction requires an artificially defined potential field that does not satisfy the usual potential laws for far distances (in fact, potentials from even moderately distant interactions are assumed by Moore to be 0, contradicting the classical potential laws). In contrast, in our work on the complexity of the n -body reachability problem, we assume the classical potential laws are in effect and no other external potential fields are assumed; therefore our lower bound requires a sophisticated construction to overcome the errors which accrue from each pair of bodies potential force. This causes a difficult to overcome accumulation of error due to the fact that all n bodies exert non-zero potential on each other, depending only on their distance.

In addition, it follows from the “Billiard Ball Computer” construction of Fredkin and Toffoli [3] that the reachability problem for particles under elastic collisions (but no potential fields) is PSPACE-hard. However, this construction relies on the presence of fixed non-movable obstacles and thus is not applicable

to n -body simulations. A related lower bound result was obtained by Reif and Yoshida for the optical ray tracing problem [12]. Again, the use of fixed polygonal obstacles is vital to their proof.

2 Simulation Algorithm and Error Bounds

In this section, we present an algorithm for simulating a system of n charged particles under the induced electrostatic forces for some time T . The algorithm is the basic Euler method for evaluating a system of differential equations, and we prove error bounds that are specific to the problem of n -body simulation. In order for the error bounds to be reasonable, we require that no pair of particles is ever closer than unit distance apart (we explain below how we can handle situations where the particles can get exponentially close).

The input to the simulation consists of the following information:

Notation	Description
m_1, m_2, \dots, m_n	Particle masses
q_1, q_2, \dots, q_n	Particle charges
$\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n$	Initial positions
$\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_n$	Initial velocities
C	Constant for Coulombic forces
T	The amount of time to simulate
ϵ	The maximum allowable (position) error

The simulation time T must be bounded by n^c for some constant c , all masses must be at least n^{-c} , and the *lengths of the binary representations* of all the remaining variables must be bounded by n^c . The constant C is such that the force magnitude between two particles i and j is

$$C \frac{q_i q_j}{\|\mathbf{x}_i - \mathbf{x}_j\|^2}.$$

From this equation, it is clear that if the distance between all pairs of particles is lower bounded by $\Delta = 2^{-n^c}$ for some constant c , we can rescale space and particle charges by a factor of $1/\Delta$ so that there is at least unit distance between all pairs of particles and the motion of all the particles is exactly the same as in the original system (but on the larger spatial scale). Specifically, we can create a new system with $\mathbf{x}_{\text{new},i} = (1/\Delta)\mathbf{x}_i$, $\mathbf{v}_{\text{new},i} = (1/\Delta)\mathbf{v}_i$, $q_{\text{new},i} = (1/\Delta)q_i$, and $C_{\text{new}} = (1/\Delta)C$. The induced force is thus a factor of $(1/\Delta)$ of the force in the original system, which is exactly what we need for the increased spatial scale. This system induces the exactly same particle motion as the original system, and yet it meets all of the assumptions in our problem statement — the length of each scaled variable grows at most by an additive polynomial factor, and the minimum distance between pairs of particles is always at least 1.

The simulation takes place by taking discrete time steps of length τ , which defines approximate trajectories for each particle. The continuation of the discrete approximation position and velocity functions are $\tilde{\mathbf{x}}_k(t)$ and $\tilde{\mathbf{v}}_k(t)$, and we

denote the discrete “samples” as $\tilde{\mathbf{x}}_{k,j} = \tilde{\mathbf{x}}(j\tau)$ and $\tilde{\mathbf{v}}_{k,j} = \tilde{\mathbf{v}}(j\tau)$. The simulation can then be described by

$$\tilde{\mathbf{v}}_{k,j} = \begin{cases} \tilde{\mathbf{v}}_{k,j-1} + \frac{\tilde{\mathbf{f}}_{k,j-1}}{m_k} \tau & \text{if } j > 0 \\ \mathbf{v}_k & \text{if } j = 0 \end{cases}$$

and

$$\tilde{\mathbf{x}}_{k,j} = \begin{cases} \tilde{\mathbf{x}}_{k,j-1} + \tilde{\mathbf{v}}_{k,j-1} \tau + \frac{\tilde{\mathbf{f}}_{k,j-1}}{2m_k} \tau^2 & \text{if } j > 0 \\ \mathbf{x}_k & \text{if } j = 0 \end{cases} .$$

In order to bound the error of these approximations, it is useful to note that the velocity and position are easily upper bounded. It is a simple calculation to calculate the total energy of the system to be simulated. Since all pairs of particles are required to stay unit distance apart at all times, we can lower bound the potential energy at any time, and thus upper bound the kinetic energy of the system, since the total energy is conserved. In particular, let B_x denote the bound on the position norm, and B_v denote the bound on the velocity norm. It can easily be seen that B_x and B_v are both bounded by e^{n^c} , for some constant c . This ensures us that the velocity and position of all particles can always be represented by a polynomial number of bits, and since the only changing state in the simulation is the velocity and position of each particle, the described simulation is clearly in PSPACE.

We use $\epsilon_{x,k,j}$, $\epsilon_{v,k,j}$, and $\epsilon_{f,k,j}$ to denote the error in the approximation $\tilde{\mathbf{x}}_{k,j}$, $\tilde{\mathbf{v}}_{k,j}$, and $\tilde{\mathbf{f}}_{k,j}$, respectively. For example, $\epsilon_{x,k,j} = \mathbf{x}_k(j\tau) - \tilde{\mathbf{x}}_{k,j}$. The error values are vectors, and we will denote the norm of the vectors by replacing the ϵ with a δ . For example, $\delta_{x,k,j} = \|\epsilon_{x,k,j}\|$.

Our main result for this problem is the following:

Theorem 1. *If $T \leq n^b$ for some constant b , there exists a constant c such that if the above simulation is run with timestep $\tau = e^{-n^c}$, then the final approximate position of each particle after simulating T/τ time steps is within ϵ of its correct position. Furthermore, all of the values required in the simulation can be represented by n^c bits, so the simulation is in PSPACE.*

Proof. (Sketch) Now we sketch the proof of this theorem. More careful analysis of the constants involved will be included in the full paper. We first notice that at step j the exact velocity of particle k can be written as

$$\mathbf{v}_k(j\tau) = \mathbf{v}_k((j-1)\tau) + \frac{\mathbf{f}_k((j-1)\tau)}{m_k} \tau + \frac{1}{2} \mathbf{x}_k^{(3)}(\xi_1) \tau^2 ,$$

where $\mathbf{x}_k^{(3)}(t)$ is the third derivative of the position function, and $(j-1)\tau \leq \xi_1 \leq j\tau$. The updated approximation velocity is therefore

$$\begin{aligned} \tilde{\mathbf{v}}_{k,j} &= \mathbf{v}_k((j-1)\tau) + \epsilon_{v,k,j-1} + \frac{\mathbf{f}_k((j-1)\tau) + \epsilon_{f,k,j-1}}{m_k} \tau \\ &= \mathbf{v}_k((j-1)\tau) + \frac{\mathbf{f}_k((j-1)\tau)}{m_k} \tau + \epsilon_{v,k,j-1} + \frac{\epsilon_{f,k,j-1}}{m_k} \tau \\ &= \mathbf{v}_k(j\tau) - \frac{1}{2} \mathbf{x}_k^{(3)}(\xi_1) \tau^2 + \epsilon_{v,k,j-1} + \frac{\epsilon_{f,k,j-1}}{m_k} \tau . \end{aligned}$$

The error is therefore updated as

$$\epsilon_{v,k,j} = \epsilon_{v,k,j-1} + \frac{\epsilon_{f,k,j-1}}{m_k} \tau - \frac{1}{2} \mathbf{x}_k^{(3)}(\xi_1) \tau^2 \quad .$$

Similarly, the position error is updated as

$$\epsilon_{x,k,j} = \epsilon_{x,k,j-1} + \epsilon_{v,k,j-1} \tau + \frac{1}{2} \frac{\epsilon_{f,k,j-1}}{m_k} \tau^2 - \frac{1}{6} \mathbf{x}_k^{(3)}(\xi_2) \tau^3 \quad .$$

By writing

$$\mathbf{f}_k(t) = \sum_{\substack{i=1 \\ i \neq j}}^n \frac{\mathbf{x}_k - \mathbf{x}_i}{\|\mathbf{x}_k - \mathbf{x}_i\|^3} C q_i q_k \quad , \quad (1)$$

some algebraic manipulation yields

$$\delta_{f,k,j} \leq 4n \delta_{x,k,j} + \delta_c \quad ,$$

where δ_c is the error introduced in the computation of (1), say by multipole approximation. We can, of course, choose δ_c to be such that $\delta_{f,k,j} \leq 8n \delta_{x,k,j}$. Furthermore, by noticing that $\mathbf{x}_k''(t) = \mathbf{f}_k(t)/m_k$, we can bound

$$\begin{aligned} \|\mathbf{x}_k^{(3)}(t)\| &\leq \sum_{\substack{i=1 \\ i \neq j}}^n \|\mathbf{x}_k'(t)\| + \frac{3}{2} \|\mathbf{x}_k(t)\| \|\mathbf{x}_k'(t) - \mathbf{x}_i'(t)\| \\ &< n B_v + \frac{3}{2} n B_x B_v = n B_v \left(1 + \frac{3}{2} B_x\right) \quad . \end{aligned}$$

Set $B_3 = n B_v \left(1 + \frac{3}{2} B_x\right)$, so B_3 is a bound on the third derivative of the position. Putting all of this together we can bound the error propagation by the following formula.

$$\begin{pmatrix} \delta_{x,k,j} \\ \delta_{v,k,j} \end{pmatrix} \leq \begin{pmatrix} 1 + \frac{4n}{m_k} \tau^2 & \tau \\ \frac{8n}{m_k} \tau & 1 \end{pmatrix} \begin{pmatrix} \delta_{x,k,j-1} \\ \delta_{v,k,j-1} \end{pmatrix} + \begin{pmatrix} \frac{1}{6} B_3 \tau^3 \\ \frac{1}{2} B_3 \tau^2 \end{pmatrix} \quad . \quad (2)$$

For any linear recurrence $\mathbf{y}_j \leq A \mathbf{y}_{j-1} + \mathbf{b}$ with $\mathbf{y}_0 = 0$, we can bound the norm of vector \mathbf{y}_j by

$$\|\mathbf{y}_j\| \leq \frac{\lambda_{\max}^j - 1}{\lambda_{\max} - 1} \|\mathbf{b}\| \quad , \quad (3)$$

where λ_{\max} is the largest eigenvalue of A . By the Gerschgorin Circle Theorem (see, for example, [2, p. 489]), it is easy to bound

$$\lambda_{\max} \leq 1 + \left(\frac{4n}{m_k} \tau + \frac{8n}{m_k} + 1 \right) \tau$$

for matrix equation (2). Using (3), we can bound

$$\begin{aligned}
\delta_{x,k,T}/\tau &\leq \frac{m_k}{(4n\tau+8n+m_k)\tau} \left(1 + \frac{4n\tau+8n+m_k}{m_k}\tau\right)^{T/\tau} \frac{1}{2}B_3\tau^2 \\
&= \left(1 + \frac{4n\tau+8n+m_k}{m_k}\tau\right)^{\frac{m_k}{(4n\tau+8n+m_k)\tau}T} T^{\frac{4n\tau+8n+m_k}{m_k}} \frac{m_k B_3}{2(4n\tau+8n+m_k)\tau} \\
&< e^{T\frac{4n\tau+8n+m_k}{m_k}} \frac{m_k B_3}{2(4n\tau+8n+m_k)\tau} \\
&< \epsilon^{n^c} \tau,
\end{aligned}$$

for some constant c . Therefore, setting $\tau = \epsilon\epsilon^{-n^c}$ insures that $\delta_{x,k,T}/\tau \leq \epsilon$ for all k . \square

3 The N -body Reachability Problem is PSPACE-hard

In this section we examine the *n-body reachability problem*: given an n -body system (namely, a set of n bodies that interact according to a harmonic potential function, and with no external forces present), an initial position and velocity of each body, and some fixed ball B , does a given body eventually reach B ? We always assume that the initial position and velocity of each of the bodies are vectors of rationals, and that the destination ball B has rational coordinates. We prove a lower bound for the most practical version of this problem: inverse-square law (repulsive) forces in three dimensions.

Our PSPACE-hardness proof uses a number of log-space reductions [13] between various problems, of interest themselves. In the following, let a *rectangular obstacle environment* be a finite set of immobile rectangular surfaces in three dimensional space whose face-planes are described by linear equations with *poly(n)*-bit rational coefficients. The problems we consider are the following:

1. *n-body reachability problem*, as defined above,
2. The *fixed potential field 1-body reachability problem* is the 1-body reachability problem augmented with a fixed potential field generated by a rectangular obstacle environment, where each obstacle is a surface with uniform electrostatic charge.
3. The *bouncing particle reachability problem* is a 1-body reachability problem with no potential fields, but with a rectangular obstacle environment where we assume that the body bounces on the obstacle surfaces with perfect elastic collisions. We actually consider a restricted version of this problem, called the *$\sigma(n)$ -centered obstacle bouncing particle reachability problem*, where we guarantee that if the particle hits an obstacle, it hits it “near” the center of the obstacle. Specifically, if d_c is the distance from the point of impact to the center of the obstacle and d_e is the distance to the closest edge, then $\frac{d_e}{d_c} \geq \sigma(n)$.

The following class of functions will be useful in the following discussion.

Definition 2. We use the notation $\exp(f(n))$ to denote the function $2^{f(n)}$. By writing simply exp-poly, we mean the class of all functions bounded by an $\exp(n^c)$ function. In other words,

$$\text{exp-poly} = \bigcup_{c>0} O(2^{n^c}) .$$

The functions $\exp(n^c)$ have a very important property, and that is that for any two constants $1 < c_1 < c_2$ and sufficiently large n , it is true that $2^n \exp(n^{c_1}) \ll \exp(n^{c_2})$. In other words, $\exp(n^{c_2})$ is more than an *exponential* factor greater than $\exp(n^{c_1})$, so any $O(2^n)$ factor of $\exp(n^{c_1})$ is negligible when compared to $\exp(n^{c_2})$.

In the remaining part of this section, we give the reductions between the three previously mentioned problems. To keep the goal in mind, we quote our final result here.

Theorem 8. *The n -body reachability problem, as defined above, is PSPACE-hard.*

First, modifying the PSPACE-hardness result for the ray tracing problem of Reif, Tygar, and Yoshida [12] to the $\sigma(n)$ -centered bouncing particle reachability problem is straightforward.

Theorem 3. *For any $\sigma(n) \in \text{exp-poly}$, the $\sigma(n)$ -centered bouncing particle reachability problem is PSPACE-hard.*

Proof. The proof of Reif, Tygar, and Yoshida that the raytracing problem with only reflective surfaces is PSPACE-hard [12] can be directly extended to show that the bouncing particle reachability problem is PSPACE hard. To show that the $\sigma(n)$ -centered bouncing particle reachability problem uses a similar construction, but a base $\sigma(n)$ encoding is used to encode the tape contents. In particular, if the tape of the simulated reversible Turing machine has contents $(a_{n-1}, \dots, a_1, a_0)$, this is represented by the distance

$$\sum_{i=0}^{n-1} a_i (\sigma(n))^{n-i} .$$

In this way, all configurations with $a_0 = 0$ are within distance $\sigma(n)^{n-1}$ of each other, but are distance at least $\frac{\sigma(n)^n}{2}$ away from all configurations with $a_0 = 1$. In this way, the splitter obstacles (the only mirrors that violate the $\sigma(n)$ -centered constraint in [12]) can be centered in such a way that the collisions are all $\sigma(n)$ centered. \square

We next give a reduction from the $\sigma(n)$ -centered bouncing particle reachability problem to the fixed potential field 1-body reachability problem. Every obstacle of the bouncing particle problem will be replaced by a box containing a single charged plate. Ideally, a particle entering the box should be

repelled by the plate in such a way that, to an observer standing outside the box, the particle seems to have bounced off the obstacle in the same way as in the bouncing particle problem. For our traveling particle, we choose a particle with unit charge and unit mass — all other parts of the construction will be scaled to these units.

To see how a particle can be repelled in such a way, consider the following special case: the point of the particle's closest approach to the plate is directly over the center of the plate (see Fig. 3). When the traveling particle is far away from the plate, the force on the particle is very small and the particle travels in approximately a straight line. In fact, as the distance is taken to infinity, the trajectory of the particle is a straight line. This line is marked on Fig. 3 as the dashed asymptote. By a simple symmetry argument, it can be shown that fixed potential field problems like this one have the following property: if a particle has trajectory $x(t)$ from time 0 to time T , then if the particle is *started* at position $x(T)$ with velocity $-x'(T)$, then the trajectory will be exactly $x(T - t)$ — it exactly reverses the original trajectory.

Let p be the position of the particle's closest approach to the charge plate, and let v_p be the velocity vector at this point. By the above property, if we start a particle at point p with velocity $-v_p$, it should exactly reverse the original trajectory. Since the potential field is symmetric about the center line, this is exactly the reflected outgoing trajectory of the original trajectory, so the entire trajectory is symmetric about this line. The immediate consequence of this argument is that the angle at which the particle is repelled from the plate is exactly the same as the angle at which the angle approaches the plate. As simple as this seems, there are two non-trivial problems that arise in this construction.

First, it is impossible to make a charged plate act exactly as an obstacle in the bouncing particle problem. In particular, if the particle's closest approach to the plate is not exactly above the center of the plate, then the reflected angle will be different from the incoming angle. We avoid this by guaranteeing that the closest approach of the particle is above a point on the plate that is *near* the center of the plate while never reaching the plate itself (recall the definition of the $\sigma(n)$ -centered bouncing particle reachability problem from above). Thus we can bound the error induced in the potential field model.

Secondly, in the bouncing particle problem, when the particle is traveling between boxes it always follows a straight line, since there are no forces acting upon it. However, in the potential field problem, there is an electrostatic force at *all* points in space, including at points between the boxes.

Both of the above problems involve error induced by approximating the bouncing particle with repelling potential fields. If the path that the bouncing particle takes is defined by the function $x_b(t)$, and the path the potential field particle takes is defined by the function $x_p(t)$, then we define an error function $\text{err}(t)$ by

$$\text{err}(t) = x_b(t) - x_p(t) \ .$$

Let $f(t)$ denote the electrostatic force acting on the particle traveling in the potential field problem. We will decompose this force into two components,

intuitively the “good” force $f_g(t)$ and the “bad” force $f_b(t)$. The good force will be the force necessary to simulate the bouncing particle, and the bad force will be the part that induces errors. Clearly, whenever the particle is outside of a box there should be no force acting on it, so all of the force is bad force. Inside a box, the good force is exactly the force needed to turn the path of the particle so that it leaves the box exactly the same way as it would have in the bouncing particle problem. Any additional force is defined to be bad force. Defined in this way, *all* of the error is induced by bad force, so if we can bound $f_b(t)$, then we can bound the error of our bouncing particle simulation.

By the careful construction given below, we can ensure that at all times $\|f_b(t)\| \leq \exp(-n^{c+1})$ for any constant c , so for a time $T = \exp(n^c)$ simulation the accumulated error is at most

$$\|\text{err}(T)\| \leq \frac{1}{2} \exp(-n^{c+1}) (\exp(n^c))^2 \leq \exp(-n^c) ,$$

for n sufficiently large.

To bound the bad force, we use the following lemma.

Lemma 4. *Consider a point p with unit electrostatic charge, and a plate with uniform electrostatic charge Q . If d is the distance from p to the closest point on the plate, then the magnitude of the force induced on p is at most $\frac{Q}{4d}$.*

To build a box that simulates a particle bouncing off an obstacle, we will place a uniformly charged, square plate centered at the location of the original obstacle. For a box with sides of length s , we want to set the distance of the particle’s closest approach to the plate to be d . Noticing that the kinetic energy of the particle at infinity is $\frac{1}{2}$ plus the error (which is $\exp(-n^c)$), we would like for the potential energy at the point distance d above the center of the plate to be exactly $\frac{1}{2}$. We can achieve this by setting the total charge of the plate appropriately. We should note here that the charge required on the plate grows linearly in d .² By using obstacles that are only $\frac{s}{2}$ wide (so the distance from the plate to the edge of the box is at least $\frac{s}{4}$) and setting $d = s \exp(-n^c)$ for a c that we specify in the following proof, we can insure a small error.

Theorem 5. *The fixed potential field 1-body reachability problem is PSPACE-hard.*

Proof. The proof of this theorem is essentially an analysis of the error introduced in our simulation, which will be shown to be very small. In this construction,

² The exact equation for potential over the plate is quite messy, but the growth rate can be proved to be linear in d by considering the growth rate of Q for a uniformly charged disk. The potential due to a charged disk is a fairly simple formula that grows as $Q(\sqrt{d^2 + r^2} - d)$, where r is the radius of the disk. It is easy to see that Q grows linearly in x when solving for a fixed potential, and this must be the same as the growth rate of the charge required by a square plate (consider inscribed and circumscribed disks).

assume we want to achieve a simulation in which the bad force is always bounded by $\exp(-n^{c_1})$. We perform the above construction with $c_2 = c_1 + 1$.

To analyze the amount of “bad” force due to this construction, first consider the bad force at positions outside of any box. The bad force due to any particular box of side-length s can be bounded by $\frac{16Q}{s^2} = O\left(\frac{s \exp(-n^{c_2})}{s^2}\right) = O(\exp(-n^{c_2}))$. Therefore, the total bad force at points outside a box is $O(n \exp(-n^{c_2}))$, which is much smaller than $\exp(-n^{c_1})$ since $c_2 > c_1$.

We now bound the bad force within a box can be bounded by defining regions of the plate that induce the bad force. In particular, for a particle entering the box, we hit “near” the center of the box. Clearly, there is some subset of the plate that can be used to generate the good force: we take a subset such that the closest approach of the traveling particle is above the center of this plate. Now the trajectory due to the force induced by the subplate looks exactly like a reflection, by the previous argument. Since we hit near the center of the plate, the parts of the plate not included in our “good” region are small and at the boarder of the plate. We can easily bound the distance to the closest point on the “bad” region by $\frac{s}{8}$, and the total charge of the bad region is significantly less than Q ; a very loose bound on the bad force induced is therefore $\frac{8Q}{s^2} = O\left(\frac{s \exp(-n^{c_2})}{s^2}\right)$. Once again, this is much smaller than $\exp(-n^{c_1})$. \square

3.1 Molecular Castles

Finally, we reduce the fixed potential field 1-body reachability problem to the n -body reachability problem by simulating surfaces by rapidly moving points. For example, consider the simple problem of simulating the effect of a potential field generated by a line segment with uniform charge distribution. To simulate this, consider a point moving very quickly back and forth between two massive particles of similar charge (see Fig. 1).

In the following discussion, we make the simplifying assumption that the outside masses are stationary; we will show how to remove the assumption later. In this idealized case, the lighter particle would bounce back and forth between the outside particles forever, repeating the same exact trajectory over and over. Let Δt denote the amount of time it takes the bouncing particle to trace out its trajectory once. In other words, if $x(t)$ and $x'(t)$ denote the position and velocity of the particle at time t , then $x(t) = x(t + i\Delta t)$ and $x'(t) = x'(t + i\Delta t)$ for all $i = 1, 2, \dots$. Now it can be seen that if the potential at a point x is averaged over one time span (0 to Δt , for example) we can see that the average potential is a good approximation of the potential of the line segment. In fact, if we restrict our attention to the potential at points that are in an area of width $\exp(-n^c)$ times the distance between the endpoints, then we see that the velocity of the bouncing particle is constant with an error term of $\exp(-n^c)$. Furthermore, the charge of the bouncing particle is chosen so that the contribution to the average potential from times when the particle is outside this range is another exponentially small, $\exp(-n^c)$. This means that at all points in our range of interest, the average potential of the bouncing particle approximates the potential of the bar with an error term of $\exp(-n^c)$.

Of course, there is additional error involved because the above discussion made the simplifying assumption that the outside masses were stationary. Of course, in the general n -body problem, no particles have fixed position, so the outside masses will move slightly. However, by making the outside masses exponentially more massive than the inside masses, they will move too little to cause any problems in our limited-time simulation. In particular, if the mass of the bouncing particle is $\exp(n^{c_1})$, then by making the outside masses have mass $\exp(n^{c_2})$ for some $c_2 > c_1$, the outside masses will move only an exponentially small amount over the polynomially bounded time of the simulation. By increasing c_2 above, we can make the amount of movement in the outside masses arbitrarily small. The only additional error introduced into the system is related to the ratio of the time-step and the inverse of the velocity of the approaching particle. This error is analyzed in the proof of the following lemma.

Lemma 6. *Let c_1 and c_2 be constants such that $1 < c_1 < c_2$. In the above construction, if the bouncing particle bounces between its endpoints with frequency $\exp(n^{c_2})$, then the trajectory of a unit velocity particle approaching the line of the bouncing particle will differ from the trajectory of the same particle approaching a uniformly charged bar by at most $\exp(-n^{c_1})$.*

Proof. In the text preceding the lemma, we showed that the average potential field of the bouncing point can be made arbitrarily close to the potential field of a uniformly charged bar, where the average is taken over a time interval of length Δt . To bound the error on the trajectory of the moving particle, we can consider the following equivalent formulation: the particle is moving toward a uniformly charged bar, but the force on the particle over time in the range $(i\Delta t, (i+1)\Delta t]$ is taken to constant with the value of the force at the beginning of the interval. Since we have assumed that the traveling particle has unit mass, the acceleration is identical to the force at all times. Next we bound the error introduced by using this approximation for a single time step.

Let $a_{err} = \max_{0 \leq t \leq \Delta t} (a(t) - a(0))$. Then the position error (denoted x_{err}) induced between time 0 and time Δt can be bounded using elementary equations of motion:

$$x_{err} = \int_0^{\Delta t} \int_0^{\tau} (a(t) - a(0)) dt d\tau \leq \int_0^{\Delta t} \int_0^{\tau} a_{err} dt d\tau \frac{1}{2} a_{err} (\Delta t)^2.$$

If we do this repeatedly over time T , then the number of steps required is $\frac{T}{\Delta t}$, so the accumulated error is $a_{max} \Delta t T$, where a_{max} is the maximum acceleration induced on the particle by the bar over the simulation time T . This is obviously bounded since there is a finite minimum distance between the particle and the bar. By taking $\Delta t = \exp(-n^{c_2})$ we make this error negligible.

Therefore, the total error introduced by our simulation is bounded by the error of the average potential field plus the error from the discrete time stepping. Both quantities can be made $\exp(-n^{c_2})$ for arbitrarily large c_2 , so by making $c_2 > c_1$ the simulation meets the desired error bound of $\exp(-n^{c_1})$. \square

To construct the more complex structure of a square sheet with uniform potential, we can use the above construction to make the four repelling line segments that bound the desired square, and then set another particle sweeping out the area enclosed by the line segments (see Fig. 2). If viewed from a distance, the average potential (taken over the time of an entire sweep of the square) induced by the bouncing particle is very close to that of a uniformly charged plate. If the horizontal distance between successive vertical sweeps is $\exp(-n^{c_2})$, then any particle at distance $\exp(-n^{c_1})$ will experience an average potential field that differs from the potential field of a square plate by only an exponentially small amount, say $\exp(-n^{c_3})$ for $c_3 > c_2$. This is formalized in the following lemma.

Lemma 7. *Let c_1 , c_2 , and c_3 be constants such that $1 < c_1 < c_2 < c_3$. Assume that the particle sweeping out the area of the square makes $\exp(n^{c_3})$ sweeps of the entire square (each one involving $\exp(n^{c_2})$ vertical sweeps) per time unit. Then the difference between the trajectory of a unit-velocity particle approaching this construction differs from the trajectory of the same particle approaching a uniformly charged plate by at most $\exp(-n^{c_1})$.*

Proof. First, we need to prove that the average potential of the sweeping particle is a good approximation of a charged plate. If we could sweep a solid bar over the area, then the proof of this fact would be identical to the proof in the preceding lemma. The vertical sweeps of the sweeping particle look (on the average) like a set of vertical bars, with small error as described in the previous lemma. Intuitively, one would expect that if the vertical bars are packed densely enough, then the average potential field of all these bars would be only slightly different from the potential field of the plane, and in fact, this is true. The potential due to a charged plate can be calculated by integrating the potential due to a bar over the width of the square. By using discretely placed bars, we are essentially using numerical integration to estimate the potential of the charged plate, and by making the integration step small enough ($\exp(-n^{c_2})$ of the square width), we make the numerical integration error very small ($\exp(-n^c)$).

We have shown that the average potential over a small timestep ($\exp(-n^{c_3})$) has an exponentially small error from that of a charged plate, so we can do a timestepping error analysis exactly like the previous lemma to complete the proof here. \square

By replacing each uniformly charged plate in the 1-body fixed potential field problem with the above construction, we have reduced the 1-body problem to the n -body reachability problem. We can make the error of the new simulation $\exp(-n^c)$ for arbitrarily large c , completing the proof of our main lower bound result.

Theorem 8. *The n -body reachability problem, as defined above, is PSPACE-hard.*

References

1. D. Beveridge and W.L. Jorgensen, Eds. *Computer Simulation of Chemical and Biochemical Systems*. Ann NY Acad. Sci. Vol. 482, 1986. Proceedings of a 1986 conference.
2. R. L. Burden, and J. D. Faires. *Numerical Analysis, Fourth Edition*, PWS-KENT Publishing Company, Boston, MA, 1989.
3. E. Fredkin and T. Toffoli. Conservative Logic. In *Int. J. of Theo. Phys.*, Vol. 21, pp. 219–253, 1982.
4. L. Greengard and V. Rokhlin. Rapid Evaluation of Potential Fields in Three Dimensions. Yale University Research Report YALEU/DCS/RR-515, 1987.
5. L. Greengard and V. Rokhlin. On the Efficient Implementation of the Fast Multipole Algorithm. Technical Report RR-602, Yale University Department of Computer Science, 1988.
6. W.F. van Gunsteren and P.K. Weiner, Eds. *Computer Simulations of Biomolecular Systems*. ESCOM, Leiden, 1989.
7. R.W. Hockney and J.W. Eastwood. *Computer Simulation Using Particles*. McGraw-Hill, New York, 1981.
8. M. Karpus and G.A. Petsko. Molecular Dynamics Simulations in Biology. In *Nature* 347, pp. 631-639, 1990.
9. C. Moore. Unpredictability and Undecidability in Dynamical Systems. In *Phy. Rev. Lett.* **64**, pp. 2354-2357, 1990.
10. K. Nabors and J. White. Fastcap: A Multipole Accelerated 3-D Capacitance Extraction Program. Technical Report, MIT Department of Electrical Engineering and Computer Science, 1991.
11. J. Reif and S. Tate. Fast Approximation Algorithms for Trummer's Problem and n -body Potential Field Evaluation. Technical Report, Duke University, July 1992.
12. J. Reif, D. Tygar, and A. Yoshida. The Computability and Complexity of Optical Beam Tracing. In *Proc. 31st Annual Symposium on Foundations of Computer Science I*, pp. 106-114, 1990.
13. W.J. Savitch. Relations Between Nondeterministic and Deterministic Tape Complexities. In *J. Comput. Systems Sci.* **4**, pp. 177-192, 1970.

Fig. 1. Simulating the potential field of a line segment using three points.

Fig. 2. Simulating the potential field of a planar region.

Fig. 3. A particle being repelled by a uniformly charged plate with the closest point of approach being over the center of the plate.