

The pursuit of the whole NChilada: Virtual petaflops using multi-adaptive algorithms for gravitational systems

G. Lake
T. Quinn
D. C. Richardson
J. Stadel

“The orbit of any one planet depends on the combined motion of all the other planets, not to mention the actions of all these on each other. To consider simultaneously all these causes of motion and to define these motions by exact laws allowing of convenient calculation exceeds, unless I am mistaken, the forces of the entire human intellect.” –Isaac Newton, 1687

We describe the keys to meeting the challenges of N-body simulation: adaptive potential solvers, adaptive integration, and volume renormalization. With these techniques and a dedicated teraflop facility, simulation can keep pace with observations of the universe. We also describe some problems in simulating the formation and stability of planetary systems.

Introduction

Astrophysical N-body simulation

Simulations play multiple roles in science. They may elucidate fundamental phenomena such as why proteins fold quickly or how astrophysical structures from the solar system through large-scale structure form and evolve. Simulations can also be key for data assimilation—from using isolated weather observations to build a full model of the state of the earth’s ocean and atmosphere to using observations of the universe to aid in our understanding of its nature, including the roles of dark matter and energy.

On the cosmological side, we have epochal surveys that have thrown down the gauntlet for cosmological simulation. We want to use those surveys to understand a variety of questions, the most important one at the moment being the equation of state of dark energy. For the solar system, we want to understand how trillions of small rocky bodies form so few planets and how the inner

solar system is cleansed of small bodies so well that the last major assault on the earth was when the dinosaurs became extinct.

We have reached a key point in astrophysical N-body simulation. N-body simulations were useful for showing qualitative features in the classic “hypothesis-driven” mode of science. A wonderful example was the classic study of bridges, tails, and mergers by Toomre and Toomre [1]. However, it was exceedingly dangerous to use simulations in a “discovery science” mode [2], as most “discoveries” were dominated by numerical artifacts, leading for example to the belief that it was impossible to form clusters of galaxies, systems that were bound and stable with substructures that lasted for many dynamical times. Simulations had always shown that all substructure detonated spontaneously, forming a single overmerged object. This matched galaxies and made clusters mysterious. With higher-quality simulations, we have discovered that substructure is preserved at all levels

©Copyright 2004 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

0018-8646/04/\$5.00 © 2004 IBM




Figure 1

Left panel shows a spatial slice from a simulation of the universe. Clusters and filaments are clearly evident. Right panel shows the same slice as seen by an observer at the bottom left corner. The observer sees a position on the sky and measures a “redshift” that bundles the expansion of the universe together with peculiar velocities. As a result, all of the clusters show an elongation that points at the observer (a phenomenon known as “fingers of God”). With a simulation, we can easily produce either view. The observer sees only the view at the right.

in the Cold Dark Matter (CDM) model, making the observed overmerging in galaxies the mystery instead [3]. This result has been reproduced by a number of groups and is now recognized as one of the lingering problems in a theory that has worked extremely well on larger scales [4]. We have also discovered that galaxies do influence one another in clusters. When small disk galaxies experience the tidal shocks of fast fly-by collisions by large galaxies, they are shaken to elevated star formation for a few Gyr (billion years) and then beaten into featureless spheroidal galaxies. We called this *galaxy harassment* [5]. At the time of the discovery of galaxy harassment, the consensus was that it had to be wrong because it was such a big effect that it should not have been missed before. However, it was easy to miss an effect that transformed galaxies when the numerical artifacts were completely destroying them!

In this paper, we describe the techniques that we use to address these challenges and place them in the context of cosmology as well as the formation and stability of planetary systems. With planetary systems, the challenge is to perform accurate integrations that retain Hamiltonian properties for 10^{13} timesteps.

Cosmological N-body simulations

Simulations are required to calculate the nonlinear final states of theories of structure formation as well as to

design and analyze observational programs. Galaxies have six coordinates of velocity and position, but observations determine just two coordinates of position and the line-of-sight velocity that bundles the expansion of the universe (the distance via Hubble’s law) together with random velocities created by the mass concentrations (see **Figure 1**). To determine the underlying structure and masses, and separate the physical structures from those created in velocity space, we must use simulations. If we want to determine the structure of a cluster of galaxies, how large must the survey volume be? Without using simulations to define observing programs, the scarce resource of observing time on \$2 billion space observatories may be misspent. Finally, to test theories for the formation of structure, we must simulate the nonlinear evolution to the present epoch.

This relationship to observational surveys defines our goal for the next decade. The Sloan Digital Sky Survey (SDSS) [6] will produce fluxes and sky positions for 5×10^7 galaxies with redshifts for the brightest 10^6 . Our ambitious observational colleagues have cut steel and ground glass to survey a “fair volume” that we must simulate, but we would need $N = 10^{12}$ particles to do this in a simplistic brute-force way. Direct summation of the gravitational forces using fixed timesteps would take 10^{10} teraflop-years.

We explain why this is a unique time to survey the universe, as well as describing the technical breakthroughs required to create a better survey of the cosmos. We then present the three keys to a realistic floating-point operation (flop) count:

1. Spatially adaptive potential solvers.
2. Temporally adaptive integrators.
3. Volume renormalizations.

Another goal of this paper is to define “high-quality simulations” and the niche science that can be done with $N \sim 10^8$ particles.

The progress of simulations

Over the last 20 years, the N of our simulations has increased as $\log_{10} N = 0.3(\text{year} - 1973)$. **Figure 2** shows the relative contributions of hardware and algorithms. The power of computers has doubled every 18 months (open circles, log scale to the right), with algorithmic advances having a slightly more rapid pace (filled circles, log scale to the left). Together, the doubling time in power is eight months, accumulating to a trillionfold increase in less than three decades. We cannot wait to simulate 10^{12} particles; we would have to invent algorithms that are a thousand times faster!

There are two constraints on our choice of N . The cost of computing a full cosmological simulation is $\sim 10^{5.7} N^{4/3}$ flops. Here, we have scaled N within a fixed physical volume. The scaling with $N^{4/3}$ arises from the increased time resolution needed as interparticle separation decreases. This extra 1/3 power in the scaling with resolution elements is an inherent part of nearly all physical simulation. For example, in a hydro grid code, one must follow sound waves across an element of the grid which decreases in size as the number of grid points to the 1/3 power if the overall physical size is fixed. Clearly, the scaling would be different if N represented the number of residues in a simulation of protein dynamics, so that the physical scales were changing. A simulation with a particularly large N was the “Hubble Volume Simulation” done with 10^9 particles by the Virgo consortium [7]. Since this simulation used an extremely large volume (roughly 100 times the size of the Sloan survey volume), the resolution was very low, and the timesteps could be large. The size of the dataset made it particularly difficult to analyze at the time that it was done, and there were no scientific results that made use of the larger volume. Clearly, both the resolution and the sample size are important to the science goals of N -body simulations.

This means that the “one byte per flop” rule of system design that applied to gigaflop machines is radically different for petaflop machines, where one likely needs

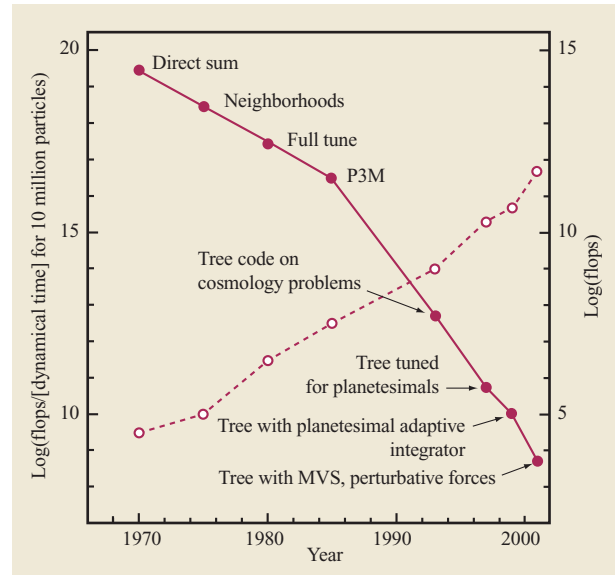


Figure 2

Gains in hardware and algorithms are compared for the N -body simulations. Algorithms are depicted by filled circles with the scale on the left, while hardware is depicted by open circles with the scale on the right.

only ~ 30 – 100 terabytes of memory for most physical problems.

The memory needed to run a simulation is $\sim 10^2 N$ bytes. If we fix N by filling memory, the time needed to run a simulation is $25 \text{ days} \times (\text{bytes/flop rate}) \times (N/100 \text{ million})^{1/3}$. Current machines are well-balanced for our simulations. With gigaflops and gigabytes, we have run simulations with $N \sim 10^{7.5}$. With teraflops and terabytes, we can simulate the behavior of 10^{10} particles or that of fewer particles in more depth. Simulations with $N \sim 10^{12}$ require ~ 100 TB and ~ 2 months on a petaflop machine. If we limit ourselves to runs of about a month, then ~ 50 TB suffices on a petaflop machine.

There are a variety of problems in which $N \sim 10^6$ represents a minimum ante. For example, clusters of galaxies are extremely important for determining cosmological parameters such as the density of the universe. Within a cluster, the galaxies are 1–10% of the mass, and there are roughly 10^3 of them. If the galaxies have fewer than 10^2 – 10^3 particles [8], they dissolve before the present epoch owing to two-body relaxation in the tidal field of the cluster. To prevent this, we need $N > 10^6$ particles per cluster. With the advances in machines and algorithms, these simulations are becoming possible on inexpensive group servers or small clusters. However, if we want to get a good sample of clusters within a Sloan volume, this would require $N \sim 10^{12}$ particles.

There are $\sim 10^{20}$ solar masses within the SDSS volume, so even 10^{12} particles is a paltry number, since each particle would represent 10^8 solar masses. We need ten times more to represent the internal structure of galaxies. N will always be far smaller than the true number of particles in the universe and will compromise the physics of the system at some level. We can only make sure that

1. The physics being examined has not been compromised by discreteness effects owing to N -deprivation.
2. Gravitational softening, discrete timesteps, force accuracy, and simulation volume do not make matters even worse.

N is not always the figure of merit in simulations, but it should be! The N -body Constitution [9] provides a set of necessary but not sufficient guidelines for N -body simulation.

The main physical effect of discreteness is the energy exchange that results from two-body collisions. Gravity has a negative specific heat owing to the total negative energy (sum of gravitational binding and kinetic energy) of a bound ensemble, like a star cluster. As a star cluster evolves, stars are scattered out by close gravitational encounters, leaving with positive energy. The remaining stars have greater negative energies, the cluster shrinks, the gravitational binding energy increases, and the stars move faster. In galaxies and clusters of galaxies, the time scale for this to occur is 10^3 to 10^6 times the age of the universe. In many simulations, the combination of discreteness in mass, time, and force evaluation can make the time scale much shorter, leading to grossly unphysical results. Thus, we must use sufficiently large N to ensure that physical heating mechanisms dominate over numerical artifacts, or that the numerical heating time scale is much longer than the time we simulate. We have inventoried all of the physical heating mechanisms experienced by galaxies in clusters and discovered a unique new phenomenon we call galaxy harassment [5].

Parallel spatially/temporally adaptive N -body solvers with “volume renormalization”

Performance gains of the recent past and near future rely on parallel computers that reduce CPU-years to wallclock-days. The challenge lies in dividing work among the processors while minimizing the latency of communication.

The dynamic range in densities demands that spatially and temporally adaptive methods be used. Our group has concentrated on tree codes [10] that can be made fully spatially and temporally adaptive. These codes use multipole expansions to approximate the gravitational acceleration on each particle. A tree is built with each node storing its multipole moments. Each node is recursively divided into smaller subvolumes until the final

leaf nodes are reached. Starting from the root node and moving level by level toward the leaves of the tree, we obtain a progressively more detailed representation of the underlying mass distribution. In calculating the force on a particle, we can tolerate a cruder representation of the more distant particles, leading to an $O(N \log N)$ method. We use a rigorous error criterion to ensure the accuracy of forces.

As the number of particles in a cosmological simulation grows, so do the density contrasts and the range of dynamical times ($\propto 1/\sqrt{\text{density}}$). If we take the final state of a simulation and weight the work done on particles inversely with their natural timesteps, we find a potential gain of ~ 50 . The leapfrog time evolution operator, $D(\tau/2)K(\tau)D(\tau/2)$, is the one most often used:

$$\begin{aligned} \text{Drift, } D(\tau/2); \quad r_{n+1/2} &= r_n + \frac{1}{2} \tau v_n, \\ \text{Kick, } K(\tau); \quad v_{n+1} &= v_n + \tau a(r_{n+1/2}), \\ \text{Drift, } D(\tau/2); \quad r_{n+1} &= r_{n+1/2} + \frac{1}{2} \tau v_{n+1}, \end{aligned}$$

where r is the position vector, v is the velocity, a is the acceleration, and τ is the timestep. This operator evolves the system under the Hamiltonian

$$H_N = H_D + H_K + H_{\text{err}} = \frac{1}{2} v^2 + V(r) + H_{\text{err}},$$

where H_{err} is of order τ^2 [11]. The existence of this surrogate Hamiltonian ensures that the leapfrog is symplectic (that is, volume-conserving in phase space, so that total energy is conserved); it is the exact solution of an approximate Hamiltonian. H_{err} explores the ensemble of systems close to the initial system rather than an ensemble of non-Hamiltonian time evolution operators near the desired one. Leapfrog is a second-order symplectic integrator requiring only one costly force evaluation per timestep and only one copy of the physical state of the system. These properties are so desirable that we have concentrated on making an adaptive leapfrog. Unfortunately, simply choosing a new timestep for each leapfrog step evolves (r, v, τ) in a manner that may not be Hamiltonian; hence, it is neither symplectic nor time-reversible. The results can be unmanageable [12]. Time reversibility can be restored [13] if the timestep is determined implicitly from the state of the system at both the beginning and the end of the step. This requires backing up timesteps, throwing away expensive force calculations, and using auxiliary storage. However, we can define an operator A that “adjusts” the timestep yet retains time reversibility and calculates a force only if it is used to complete the timestep [14]. This is done by choosing A such that it commutes with K , so that the sequence $DAKD$ is equivalent to $DKAD$. Since K changes

only the velocities, an A operator that depends entirely on positions satisfies the commutation requirement. The “natural definition” of a timestep, $1/\sqrt{\text{density}}$, is ideal but difficult to define when a few particles are close together. Synchronization is maintained by choosing timesteps that are a power-of-2 subdivision of the largest timestep, τ_s . That is, $\tau_i = (\tau_s/2^i)$, where τ_i is the timestep at a level of the timestep hierarchy. We are currently experimenting with this approach and encourage others to look at variants.

“Volume renormalization” uses a large-scale simulation with modest resolution to identify regions of particular interest: sites of galaxy/QSO formation, large clusters of galaxies, etc. Next, initial conditions are reconstructed using the same low-frequency waves but adding higher spatial frequencies. We have achieved 10^3 -parsec resolution within a cosmological volume of size 10^8 parsecs to study the origin of quasars [15].

There are methods of calculating potentials that are advertised as being $O(N)$ rather than $O(N \log N)$, such as the fast multipole method (FMM) [16]. We use a number of great tricks from FMM (our tree codes now use multipoles and combine and translate the multipole expansions). Nonetheless, there are several reasons why we do not use FMM. The first is that there is no method that is truly $O(N)$. This can be seen with one-dimensional gravity that satisfies Gauss’ theorem. In this case, the gravitational force is proportional to the number of particles to one’s left, minus the number of particles to one’s right. As a result, gravity is equivalent to a sort which cannot be reduced to $O(N)$ for arbitrary real numbers. The one-dimensional gravity is a particularly useful example, as the scaling is $O(N \log N)$ to calculate the relative positions, but then one performs only one gravitational interaction per particle [so it is not just $O(N)$, but truly “ N ”]. Similarly, the adaptive FMM must first build a tree which involves a sort [$O(N \log N)$], but then calculates $O(N)$ gravitational interactions. The interaction calculation is expensive enough that studies have generally seen the algorithm scale as $O(N)$. These arguments are a bit simplistic. We have assumed that the sort necessary to build a tree is $O(N \log N)$. There are $O(N)$ sorting routines that work as long as the distribution of points is not pathological, as in a Cantor set. However, the late stages of cosmological simulations become progressively more like a Cantor set! Requiring that the simulation be stably integrated means that particles cannot move too much, so tree reconstruction after each step should certainly be faster than a full build and should scale more like $O(N)$.

Finally, we use a tree code to employ adaptive timesteps. There is a tendency to focus on reducing the simple calculation of forces from the “sum over all

particles” N^2 scheme to one that is $O(N \log N)$ or $O(N)$. However, we need to solve the full dynamical problem that is naively $N^{7/3}$, where the additional $N^{1/3}$ comes from the need to take finer steps to implement higher resolution. With the tree code, one divides work into the calculation of forces on a particular particle. In this way, one can employ adaptive timestepping, with each particle having an individual timestep. In the FMM, work is effectively divided by scale, and the scales cannot be decoupled in any way to implement temporal adaptivity. The best sign of speed and relative efficiencies of the methods is that the tree build is a large fraction of the cost of a force evaluation in the tree code with adaptive timesteps, whereas its scaling can be seen in the FMM implementations.

One last class of methods is based on fast Fourier transforms (FFTs) combined with corrected particle–particle sums on small scales. The adaptive versions of this code by Couchman [17] with further refinements by Couchman, Thomas, and Pearce [18, 19] have powered the large Virgo consortium runs [7]. These codes use nested meshes with tailored Green functions with a final particle–particle summation on the smallest scales. For a single potential evaluation, this method is somewhat faster than the tree code, with the tree code winning only when the distributions are heavily clustered owing to its more efficient use of adaptive timestepping. There is also a hybrid composed of a particle-mesh code and a tree code [20]. One pathology of the tree code is that it requires the use of many more terms when the distribution of particles is nearly uniform, such as the early stages of a cosmological simulation. This is just the condition for which mesh codes are blindingly fast and accurate. This is one of the motivations for the more flexible simulation package discussed in the last section.

The distinctions between adaptive mesh, FMM, and tree code are becoming increasingly blurred. As one last example of this, our current method of collecting the interactions for each particle, or “tree-walking,” has been improved by applying a technique inspired by FMM. Instead of collecting all interactions for a single particle (or a small nearby set of particles), we consider which interactions are common to all subcells of a given cell at each level of the tree. This is done in a top-down manner, maintaining a “checklist” structure of cells that remain to be checked *either* for inclusion on an interaction list or to be opened and have their subcells placed on the checklist [21]. This results in many fewer comparisons in deciding which interactions to perform, thereby reducing one of the big integer-operation-limited steps of tree codes. With this tree-walk and the use of multipoles that are combined and translated, one could argue that our current code is as close to being a variant of FMM [16] as it is to the original tree codes [10, 22].

Current cosmological challenges

We now have multiple mysteries in cosmology as we continue on the path set by Copernicus, who showed that the earth was not the center of the solar system. We now know that our solar system resides on the outskirts of a galaxy which is at the edge of a cluster of galaxies, and that the entire system moves toward a greater attractor [23]. Not only that, but the universe appears to consist more of dark matter than of identifiable matter. Still more of the universe is in dark energy, which has a behavior entirely different from that of matter [24, 25]. Dark energy has a negative pressure or tension that makes the universe expand faster and faster, whereas the gravitational tug of normal matter slows it down. Strangely, we seem to live in a special epoch in which dark energy has just become dominant over normal and even dark matter. We still do not know what either the dark matter or the dark energy might be. A first step will be to determine its equation of state: Does it behave just like Einstein's cosmological constant, or differently?

Matching clusters, groups, galaxies, and cosmic flows has been a sensitive test of cosmological models. The key data for the next decade will be the Sloan Digital Sky Survey. Our proposed program to simulate the Sloan volume has been as follows:

1. Simulate the entire volume $(800 \text{ Mpc})^3$ with $N = 10^8$, where Mpc = one million parsecs.
2. "Renormalize" dozens of groups, clusters, etc. and simulate with 10^8 particles.
3. Run selected subvolumes that are 10^{-3} of the total volume or $(800 \text{ Mpc})^3$ with 10^8 particles to examine the evolution of groups and follow the early evolution of "rare objects."

The first simulation took 10^{16} floats or ~ 10 petaflop-seconds and has been run for a few cosmological models. Like many other groups, we found that matching structures in the universe required that the universe be dominated by a dark energy or the cosmological constant Λ that was introduced by Einstein. The dark energy was also needed to fit observations of distant supernovae [26] and has now been determined to very high precision by the Wilkinson Microwave Anisotropy Probe (WMAP) observations of the cosmic microwave background [27]. This satellite mission determined to high precision the age of the universe, its expansion rate, and the relative proportions of normal matter, dark matter, and dark energy. However, these experiments have not been sensitive to the dark energy equation of state, which will stimulate a new generation of dynamical probes coupled with simulations. A recent study [28] finds that there is a strong sensitivity of the cosmic Mach number [29] (the ratio of the flow velocity to the local velocity dispersion)

to the equation of state of the dark matter. The next step for using this sensitivity will be to obtain very accurate Mach numbers for a variety of flows.

The fate of the solar system

Advances in hardware and numerical methods have finally enabled us to integrate the solar system over its lifetime. Such an integration represents a thousandfold advance over the best, longest, most accurate integration ever performed [30] and can address numerous questions such as the following: Is the solar system stable? Do all of the planets remain approximately in their current orbits over the lifetime of the solar system, or are there drastic changes, or perhaps even an ejection of a planet?

What is the stability of other planetary systems?

Integration of the solar system will allow us to address other questions, such as the following.

What is the effect of orbital changes on the planetary climates? According to the Milankovich hypothesis, climatic variations on the earth are caused by insolation changes arising from slow oscillations in the earth's orbital elements and the direction of the earth's spin [29]. Remarkably, the geophysical data (primarily the volume of water locked up in ice as determined by the $^{18}\text{O}/^{16}\text{O}$ ratio in sea-bed cores) covers a longer time than any accurate ephemeris.

How does weak chaos alter the evolution of the solar system? Why does the solar system appear stable if its Lyapunov time, that is, the time scale for two nearly identical trajectories to diverge significantly, is so short?

How are the giant planets related to terrestrial planets in the "habitable zone" between boiling and freezing of water by the central star? Without a cleansing of planetesimals from the solar system by giant planets [30], the bombardment of the earth by asteroids would be steady and frequent throughout the main sequence lifetime of the sun [31]. The chaos produced by Jupiter and Saturn may have played a role in ensuring that planetesimals collided to form the terrestrial planets,¹ but too much chaos will eject planets in the habitable zone. While a search for giant planets is the only technically feasible one today, it may be the ideal way to screen systems before searching for terrestrial planets.

Integrating nine planets for 10^{11} dynamical times

When Laplace expanded the mutual perturbations of the planets to first order in their masses, inclinations, and eccentricities, he found that the orbits could be expressed as a sum of periodic terms—implying stability. Poincaré [32] showed that these expansions do not converge because of the occurrence of resonances. Using the

¹ In Ancient Greek, chaos was "the great abyss out of which Gaia flowed."

Table 1 Summary of direct solar system integrations.

| <i>Year</i> | <i>Team</i> | <i>Length (Myr)</i> | <i>No. of planets</i> | <i>General relativity?</i> | <i>Moon?</i> | <i>Special hardware?</i> |
|-------------|----------------------------|-------------------------|---------------------------|--------------------------------|--------------|------------------------------|
| 1951 | Eckert et al. [35] | 0.00035 | 5 | no | no | no |
| 1965 | Cohen and Hubbard [36] | 0.12 | 5 | no | no | no |
| 1973 | Cohen et al. [37] | 1. | 5 | no | no | no |
| 1984 | Kinoshita and Nakai [38] | 5. | 5 | no | no | no |
| 1986 | Applegate et al. [39] | 217. | 5 | no | no | yes |
| 1986 | Applegate et al. [39] | 3. | 8 | no | no | yes |
| 1987 | Carpino et al. [40] | 100. | 5 | yes | no | no |
| 1988 | Sussman and Wisdom [41] | 845. | 5 | no | no | yes |
| 1989 | Richardson and Walker [42] | 2. | 9 | no | no | no |
| 1991 | Quinn et al. [28] | 6. | 9 | yes | yes | no |
| 1992 | Sussman and Wisdom [43] | 100. | 9 | partly | yes | partly |
| 2002 | Ito and Tanikawa [44] | 10,000. | 5 | no | no | no |
| 2002 | Ito and Tanikawa [44] | 5,000. | 9 | no | no | no |
| 2003 | Varadi et al. [45] | 5,000. | 9 | no | no | no |
| 2003 | Varadi et al. [45] | 90. | 9 | yes | yes | no |
| 2003 | Armstrong et al. [46] | 1,000. | 9 | yes | yes | no |

Kolmogorov–Arnold–Moser (KAM) theorem, Arnold [33] derived constraints on planetary masses, eccentricities, and inclinations sufficient to ensure stability. The solar system does not meet his stringent conditions, but this does not imply that it is unstable. Laskar [34] tested the quasi-periodic hypothesis by numerically integrating the perturbations calculated to second order in mass and fifth order in eccentricities and inclinations, $\sim 150,000$ polynomial terms. Fourier analysis of his 200-million-year integration reveals that the solution is not a sum of periodic terms and implies an instability that is surprisingly short, just 5 Myr (million years). The second method for attacking the stability problem is to explicitly integrate the orbits of the planets (**Table 1**).

As early as 1965, Pluto’s behavior was deemed suspicious. In the last ten years, it has become clear that the solar system is chaotic. However, the source of the chaos is unclear, because the system of resonances is complex and the Lyapunov exponents appear to be sensitive to fine details of initial conditions. Nonetheless, the solar system is almost certainly chaotic. Laskar [47], examining the fate of Mercury, estimates that there is a significant chance of ejection during the lifetime of the solar system. Our belief in the apparent regularity of the solar system may be due to our inability to know that before the last few ejections, there may have been 10, 11, or even 12 planets a few billion years ago. At the very least, the chaotic motion leads to a horizon of predictability for the detailed motions of the planets. With a divergence time scale of 4–5 Myr, an error as small as 10^{-10} in the initial conditions will lead to a 100% discrepancy in 100 Myr. Every time that NASA launches a rocket, it can turn winter to spring in a mere 10 Myr.

Are the integrations meaningful given this sensitivity to the initial conditions? We investigate Hamiltonian systems that are as close to the solar system as possible. KAM theory tells us that the qualitative behavior of nearby Hamiltonians should be similar. While the exact phasing of winter and spring is uncertain after millions of years, the severity of winter or spring owing to changes in the earth–sun distance and the obliquity are predictable.

We have started a 9-Gyr integration: 4.5 Gyr into the past, when the solar system was formed, and 4.5 Gyr into the future, when the sun becomes a red giant. One basic requirement is a computer with fast quadratic precision to overcome roundoff problems. The IBM Power Series* has been the machine of choice. Our first digital orrery was an IBM POWER2* workstation that evolved the solar system $\sim 10^9$ times faster than “real time”; this was two to three orders of magnitude faster than other available CPUs owing to IBM’s implementation of quadratic precision. We have now completed a 1-Gyr integration [46]. To understand any chaos, we will need to see it by an independent means and devise methods to determine its underlying source.

A parallel method does not seem promising, since there are only nine planets to distribute among processors. We employ a different form of parallelism—the “time-slice concurrency method” (TSCM) [48]. In this method, each processor takes a different time slice; the initial conditions for processor 2 are the final conditions for processor 1, and so on. The trick is to start processor 2 with a good prediction for what processor 1 will eventually output, and iterate to convergence. This is analogous to the waveform relaxation technique used to solve some partial differential equations [49]. However, Kepler ellipses are a good approximation to the orbits for a time scale that is

Table 2 Direct simulations of the formation of the inner planets.

| Year | Team | N | t (yr) | Δa (AU) | Giants? |
|------|-----------------------------|--------|-------------------|-----------------|---------|
| 1986 | Lecar and Aarseth [55] | 200 | 6×10^4 | 0.5–1.5 | no |
| 1990 | Beaugé and Aarseth [53] | 200 | 6×10^5 | 0.6–1.6 | no |
| 1992 | Ida and Makino [56, 57] | 800 | 5000 | 0.3 | no |
| 1993 | Ida and Makino [54] | 800 | 5000 | 0.3 | no |
| 1993 | Aarseth et al. [58] | 400 | 1.2×10^4 | 0.04 | no |
| 1996 | Kokubo and Ida [59] | 5000 | 2×10^4 | 0.4 | no |
| 1998 | Kokubo and Ida [60] | 5000 | 2×10^4 | 0.4 | no |
| 1998 | Chambers and Wetherill [61] | 100 | 10^8 | 0.5–2.0 | yes |
| 1998 | Richardson et al. [62] | 10^5 | 1000 | 1.2–3.6 | yes |
| 2000 | Richardson et al. [63] | 10^6 | 1000 | 0.8–3.8 | yes |

proportional to the ratio of the mass of the sun to that of Jupiter. Tests show that it is extremely efficient to iterate to convergence in double precision (typically 14 iterations, each costing 10–15% of a quad iteration), then perform just two iterations to get convergence in quad. In this way, the total overhead of the full 16 iterations can be less than a factor of 4.

There are still many algorithmic issues to be addressed. For long-term integrations, TSCM has been formulated in a way that preserves the Hamiltonian structure and exploits the nearness to an exactly soluble system; otherwise, errors grow quadratically with time. TSCM should enable us to integrate ~ 5 Gyr per day on a 512-node IBM POWER4*, a speedup over real time of 10^{12} . This should make it feasible to study the instability of other solar systems. Detailed development and implementation will be much more challenging than for previous methods, and our high-quality serial integration will be required for comparison and validation.

Finally, we use a new technique to gauge the origin of instabilities (the “tangent equation method”) [50]. In the past, it was common to integrate orbits from many slightly different initial conditions. While that works, it is more rigorous and also more economical to integrate the linearized or tangent equations—the equations for differences from nearby orbits. We integrate the tangent equations along with the main orbit equations.

Cosmology meets cosmogony: Planetary system formation

Theories of solar system formation are traditionally divided into four stages [51]: collapse of the local cloud into a protostellar core and a flattened rotating disk (nebular hypothesis); sedimentation of grains from the cooling nebular disk to form condensation sites for planetesimals; growth of planetesimals through binary collision and mutual gravitational interaction to form protoplanets (planetesimal hypothesis); and the final assembly to planets, with the remaining disk cleansed by

ejections from chaotic zones. Our cosmology code is ideal for the third stage of solar system formation, particularly in the inner regions where gas was not a primary component and gravitational interactions dominated the evolution. The first stage entails magnetohydrodynamics, the complicated small-particle physics and gas dynamics of the second stage is still not well understood, and the fourth stage is the purview of long-term stability codes.

All that is required for a detailed simulation of the third stage is a model of the collisional physics and a code capable of dealing with a large number of particles. However, previous direct simulations of the planetesimal stage (summarized in **Table 2**) fall far short of capturing the full dynamic range of the problem. Our cosmology code has the potential to treat as many as 10^7 particles simultaneously for 10^7 dynamical times, a ten-millionfold improvement that makes us enthusiastic! Only statistical methods [52] employing prescriptions for the outcomes of gravitational encounters have been used to take a look at this regime. We reach an important threshold at $N = 10^7$ in our ability to follow planetesimal evolution. At early times, the relative velocities between planetesimals are small, and inelastic physical collisions lead to “runaway” growth of planetary embryos [53]. Eventually, gravitational scattering increases the planetesimal eccentricities to such an extent that collisions result in fragmentation, not growth. The embryos continue to grow owing to their large mass, but at a slower rate as their “feeding zones” are depleted [54]. The total mass of our planetary system is 448 times that of the earth ($448M_{\text{earth}}$) or $3.6 \times 10^4 M_{\text{moon}}$, while the inner planetesimal disk amenable to simulation had a mass of $10^2 M_{\text{moon}}$. To capture both growth and fragmentation [52] requires a minimum particle mass of $10^{-5} M_{\text{moon}}$, leading to our $N = 10^7$.

A detailed direct simulation of planet formation can address a variety of important questions, including the following: Was there runaway growth of a few embryos, or a continuously evolving homogeneous mass distribution? How does the primordial surface density alter the

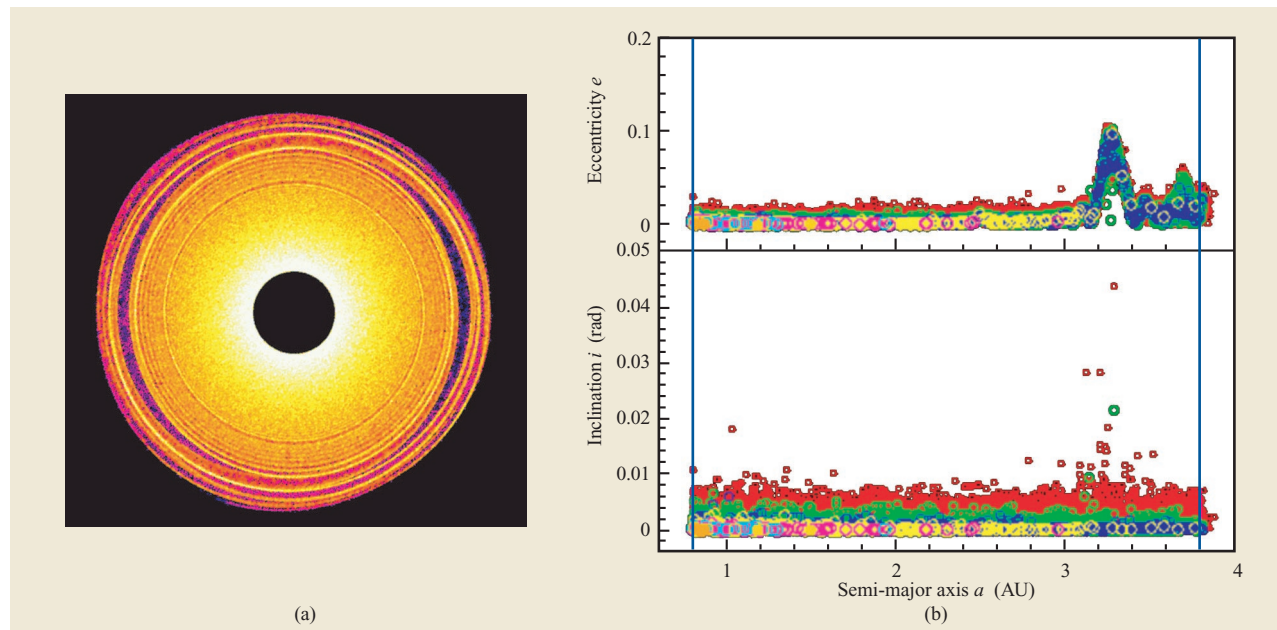


Figure 3

(a) Mass density of a 10^6 -particle simulation after 100 years. Bright colors represent regions of high density. The dot at the left is Jupiter. The disk extends from 0.8 AU, just inside the earth's present-day orbit, to 3.8 AU near the outer edge of the asteroid belt. The gaps and spiral structures in the disk are associated with Jupiter's mean-motion resonances. (b) Plot of semi-major axis a vs. eccentricity e and inclination i for a 10^6 -particle simulation, showing only every hundredth particle to prevent overcrowding. The peaks in e and i (strongest in e) at specific a values correspond to mean-motion resonances with Jupiter. Circle size is proportional to particle radius. The particles are also color-coded by size, with the smallest red circles indicating pristine material, larger green circles indicating bodies that have undergone one collision, blue circles two collisions, and so on (the color sequence repeats at larger masses). Reprinted from [63], © 2000, with permission from Elsevier.

evolution? What fixes the spin orientation and period of the planets—uniform spin-up from planetesimal accretion [64], or a stochastic process dominated by the very last giant collisions [65]? Is it feasible that the earth suffered a giant impact late in its growth that led to the formation of the moon [29]? How much radial mixing was there, and can it explain observed compositional gradients in the asteroid belt [66]? Finally, what is the dominant physical mechanism that drives the late stages of growth—are intrinsic gravitational instabilities between embryos sufficient, or are perturbations by the giant gas planets required?

This last point is of key importance to future searches for terrestrial planets. We strongly suspect that the end result of our research may be the assertion that one should concentrate searches for terrestrial planets in those systems that have giant planets. We have begun to address these issues with a modified version of the cosmology code. Collisions are detected (rather than “softened away”), and the outcomes are determined by the impact energy, the lowest energies generally leading to mergers and the highest energies leading to fragmentation. In our current code, merging and bouncing are implemented,

while fragmentation is in the testing stage. Integrations are carried out in the heliocentric frame and may include the giant planets as perturbers. Additional programs are used to generate appropriate initial conditions and to analyze the results of the simulation, but the main work is performed by the modified cosmology code.

Figure 3 shows the mass density (a) and semi-major axis a vs. both eccentricity e and inclination i (b), respectively, at the end of a 100-yr run that began with 10^6 identical cold planetesimals in a disk from 0.8 to 3.8 AU (astronomical units) with surface density proportional to $r^{-3/2}$. The present-day outer planets were included in the calculation. The simulation took 60 hours to finish on a Cray T3E** computer² with 128 dedicated processors using a fixed timestep of 0.01 yr. The effect of Jupiter on the disk, which extends well into the present-day asteroid belt, can be seen clearly in the density plot: There is a large density gap at the 2:1 resonance at 3.2 AU and a narrow groove at the 3:1 at 2.5 AU, along with spiral wave patterns and other telltale features. Corresponding features in Figure 3(b) show how Jupiter stirs up

² Cray, Inc., Seattle, WA.

planetesimals at the mean-motion resonances. Note that conservation of the Jacobi integral accounts for the slight bending of the e peaks toward smaller a . Meanwhile, planetesimal growth has proceeded unmolested in the inner region of the disk (under the assumption of perfect accretion). The largest planetesimal at the end of the run is eight times its starting size.

As far as we are aware, this is the largest simulation of a self-gravitating planetesimal disk that has ever been attempted. The figures show, however, that to reach the regime of runaway growth ($\sim 10^4$ – 10^5 yr), a new timestepping approach is needed. We are currently developing a technique to exploit the near-Keplerian motion of the planetesimals. For weakly interacting particles, we divide the Hamiltonian into a Kepler component, implemented using Gauss' f and g functions to step along elliptical rather than linear segments, and a perturbation component owing to the force contributions of all of the other particles. In this regime, timesteps can be of the order of the dynamical (i.e., orbital) time, resulting in computational speedups of 10–100. For strongly interacting particles (defined as particles with overlapping Hill spheres), the Hamiltonian is factored into the standard kinetic and potential energy components, with the central force of the sun as an external potential. In this regime, particles are advanced in small steps, which allows for the careful determination of collision circumstances. It also allows the detection of collisions in the correct sequence even if a single particle suffers more than one collision during the interval. This new technique will transfer well to the protein folding problem.

The code that was used to do this simulation has considerable room for improvement. Nonetheless, we can compare the time to solution with GRAPE (GRAVity PipE) computer³/HARP computer⁴ simulations [59, 60]. We found that our code had an order of magnitude faster time to solution in the host computer used by this group (a Compaq EV5 Alpha computer⁵) than they obtain by employing the special-purpose hardware back end. Since all hardware is a moving target, we are working on a new version of the code that will have a similar relationship to the current-generation GRAPE-6 (a 64-teraflop computer for stellar dynamics studies). The gains that we anticipate are summarized in **Table 3**.

The challenge is to predict when particles will change between the regimes of weak and strong interaction. One method we are considering is to construct a new binary tree ordered by perihelion and aphelion. Those particles with orbits separated by less than a Hill sphere are flagged for further testing. This screening has a cost of $N \log N$ and is performed only once per long Kepler step. Flagged

pairs of particles with phases that are certain to stay separated over the integration step are reset. The remaining particles are tested by solving Kepler's equation in an elliptical cylindrical coordinate system to determine the time of actual Hill sphere overlap. Switching between Hamiltonians is not strictly symplectic, but it occurs infrequently enough for any given particle that it is not a concern. Dissipating collisions are inherently non-symplectic anyway. (There is some skepticism from colleagues about switching the integration scheme. True conservatives are also worried that the normal tree codes are not momentum-conserving. There is a momentum-conserving variant [67], but it loses this property with adaptive timestepping.) Once particles separate beyond their Hill spheres (or merge), they are returned to the Kepler drift scheme. Although much work remains to be done, the reward will be the first self-consistent direct simulation of planetesimals evolving into planets in a realistic disk. The results can be used to study related problems, such as the formation of planetary satellites, orbital migration of giant planets in a sea of planetesimals, and ultimately the ubiquity and diversity of extra-solar planetary systems.

To resolve the collisions and their outcomes, we are testing a scheme in which the particles are replaced by "rubble piles" when they collide. Rubble piles have little or no tensile strength; they are assemblages of particles that are held together only by gravity [68]. This approach captures the dynamics of gravitational reaccumulation explicitly without relying on questionable scalings of strength-regime laboratory experiments to gravity-regime km-size planetesimals. After a fixed interval, particles that have accumulated into sufficiently large bodies are replaced with single large particles, while the remaining debris is treated as "dust" that slowly accumulates on larger particles and drags on their orbits over time. We believe that this mechanism may be responsible for the low orbital eccentricities of present-day terrestrial planets, a property that has so far eluded all numerical simulations of planet formation.

Combining special hardware and advanced algorithms

In all of the codes that we run, gravity plays an important role. The calculation of gravitational interactions is essential and is normally a large fraction of the simulation cost. The GRAPE consortium has built special hardware to accelerate the gravitational calculations. We have noted that our algorithms lead to faster times to solution on their host machine than would be achieved by using their special-purpose hardware and simpler algorithms. What about using *both* spatially and temporally adaptive algorithms and their special hardware?

³ RIKEN, Wako, Japan.

⁴ Sundance Corporation, Reno, NV.

⁵ HP Compaq Corporation, Palo Alto, CA.

Table 3 Software and hardware for integrating planetesimals.

| | <i>Gain factors</i> |
|--|---------------------|
| I. Improved software gain (factors do not all accumulate; adopt 20, with 5 as a reserve factor) | 20 |
| Gravity module (now coding) | 4 |
| Improve tree build and walk | |
| Use tensor multipoles for recursion | |
| Enable local expansions | |
| Collision module (testing) | 4 |
| New collision-detection scheme | |
| Timestepping (testing) | 5–30 |
| Use MVS with collision detection | |
| II. Total hardware gain | 240 |
| Use 512 CPUs (available) | 4 |
| Run 600 hours (available) | 10 |
| Run on current CPUs | 10 |
| (special functions used extensively; baseline via CRAY T3E (300 MHz), boosted via Compaq EV5 Alpha) | |
| III. Evolution | |
| Particle number decreases as system evolves | 10 |
| Increase initial cross section to accelerate number decrease | 3 |
| Total gains | 10,000 |
| Emergency reserve factors | 15 |
| Total speedup relative to HARP-2 for 10^6 particles | ~500,000,000 |
| Alpha EV5/HARP-2 ratio on test case = 10 | |
| Algorithmic scaling including HARP-2 efficiency gains | 2,000 |
| 200× speedup on 512-node machine | |
| 10× gain from next-generation CPUs | |
| 20× gain from software | |

Note: In this table, we have included some “optimistic” projections as reserve factors (shown in italics). We think that much of this reserve of 15× speedup is likely to be realized, but we expect to encounter some problems achieving every goal in the table, so we hold it in reserve to be reasonably conservative.

The effort involved in creating a typical gravity-based code comprises the following: 1) tree construction and traversal (including neighbor list determinations), 20%; 2) computation of pairwise gravitational interactions, 50%; and 3) computation of periodic boundary conditions, special functions for planetary integrations, SPH (smooth particle hydrodynamics) gas physics, etc., 30%. To date, special-purpose hardware has focused on making the pairwise gravitational interactions blindingly fast. In the limit that it is infinitely fast, this provides only a factor of 2 speedup. However, obtaining even half of this gain is a severe challenge to the I/O designs of the special hardware. For example, GRAPE was originally designed for the specific problem of globular cluster dynamics, where it has had a number of successes [69]. These systems have $\sim 10^6$ stars, so there is no reason to use a larger N . This sets a very modest limit to the amount of memory needed in the special box. With a “sum over all pairs,” it means that for every 24 bytes transferred for the particle positions (and returned for the particle velocities), one does $\sim 3 \times 10^7$ floating-point operations. Thus, the ratio of the I/O channel speed to the floating-point speed need be only roughly 2 MB/s for every teraflop, so a

peripheral component interconnect (PCI) bus is sufficient to connect to a teraflop of hardware. However, in our simulations, there are typically only a few hundred pairwise interactions calculated, so the ratio needed to spend more time on computation than on data transfer is $\sim 10\beta$ GB/s for every teraflop, where β is the number of times that the particle is loaded into the special hardware (it has to be there every time it appears on an interaction list, so with efficient traversals and caching, we might keep $\beta \sim 10$). Thus, it is not straightforward to use special hardware with aggressive algorithms.

Looking toward the future

The astute reader will notice that the phenomenal algorithmic progress shown in Figure 2 stopped a couple of years ago. There are both scientific and social reasons for this. It is extremely important for the fidelity of the scientific results to ensure that the integrators are symplectic. There are many subtle attributes of these aggressive algorithms that will alter the Hamiltonian character of the problem. We were surprised to discover that some of the problems were easy to address, while others were far more difficult. One example was that

all past cosmological simulations had performed the integrations in noncanonical coordinates [21], resulting in a strong transient energy change in the early stages of the integration when the expansion rate of the universe was high. Since it had been included in every simulation, it became a “feature.”

We are concerned about a number of other “features” and have been engaged in extensive tests to understand them. These derive from the fact that the forces in a tree code are normally not symmetrical, so that Newton’s law of equal and opposite forces (which ensures conservation of global momentum) is not strictly upheld. When a code explicitly conserves a quantity, it is a bug check. When it does not, as in this case, it becomes something to monitor as a reality check on error propagation. We are more concerned about putting in “switches” to the integrators that change timesteps in a fixed integration scheme or completely change the integration scheme, as in the collisions of planetesimals. This requires extensive testing.

The last section makes it clear that the codes have enough complexity that when one element is greatly speeded up, the bottlenecks move elsewhere. This is also true at a higher level. The simulation code stopped being the main bottleneck to scientific production. We had a several-year push to bring the machines and algorithms to the level at which we could do simulations fairly easily with $N \sim 10^7$, and it required only mild heroics to run $N \sim 10^8$. This also meant that $N \sim 10^6$ became a problem for workstations. This opened up a wealth of scientific possibilities, which was the main driving force in the first place. It also meant that we could generate enormous data sets that require analysis for scientific results.

What we need now is more flexible analysis code and a greater range of physics represented in the codes. We also want to create a venue to test and compare codes. There is demand for using Couchman’s adaptive mesh techniques for the early stages of cosmological simulation and then switching to tree methods. In some cases, it would certainly be useful to incorporate GRAPE or Blue Gene* hardware [70]. For such cases, we might want to couple our framework to Blue Matter [71], the molecular simulation software to be used on Blue Gene. Both of these are broader community needs and must incorporate work from multiple groups. To this end, we started the “Whole NChilada Project”⁶ to tackle the simulations, data handling, and analysis. This process mirrors the evolution that has occurred in a number of simulation projects and is reflected in the changes of support from the HPCC (High Performance Computing and Communications) projects in multiple agencies to the newer initiatives such as the NSF (National Science Foundation) ITR

(Information Technology Research), the NASA Earth System Modeling Framework and its broader Computing Technologies Project, and the interagency National Virtual Observatory. It also reflects the success of the first-generation University of Washington *N*-body shop; many of the former students and postdoctoral fellows are now group leaders at other universities. NChilada is a vehicle for realizing collective synergies from the efforts of allied but individual groups.

Summary: Virtual petaflops

Past planetesimal simulations used codes with an algorithmic complexity that would be similar to the point labeled “full tune” in Figure 2 and computers with speeds of ~ 10 Mflops. (Special-purpose “GRAPE” hardware of $\sim 10^6$ Mflops has been used, but such implementations involve sums over all interactions, and so they are closer to the direct-sum case in floating-point cost [60].) Our algorithms result in a collective speedup of $\sim 10^8$ for simulations with $N \sim 10^7$ (with rough accounting for the reduction in N over the course of the simulation in the solar system formation runs).

We have emphasized that both hardware and algorithm improvements are required to achieve the performance gains that we require. Figure 2 shows that the speedup factor from algorithms vastly exceeds that of the hardware. It is not sufficient to simply wait for computers to get better, nor does it seem to pay to build special hardware. McMillan et al. [72] asserted that if GRAPE-6 were built, it could use its petaflop speed to follow 10^6 planetesimals by 2003. They claimed that this would offer a seven-year advantage over general-purpose computers that would only be able to follow 10^4 particles by the year 2000. As of September 2003, a 64-teraflop GRAPE-6 system has been built, and smaller GRAPE-6 systems garnered IEEE Gordon Bell prizes in 2000 and 2001.

Our test simulations without the new integrator were already ten years ahead of their projections [63]. Our approach has outpaced special hardware efforts over the last several years, and we believe that it will continue to do so for the foreseeable future. Sir Isaac would love to see the enhancement of “the entire human intellect” by high-performance computing.

Acknowledgments

This project is supported by the NASA High Performance Computing and Communications/Earth and Space Sciences (HPCC/ESS) program, Innovative Research Program (IRP), Applied Information Systems Research (AISR), and the NIST Advanced Technology Program (ATP), the NSF Knowledge and Distributed Intelligence (KDI) and Information Technology Research (ITR) programs, and the Washington State University Band Professorship fund.

⁶ <http://hpcc.astro.washington.edu/nchilada/bin/view/Nchilada/WebHome/>.

*Trademark or registered trademark of International Business Machines Corporation.

**Trademark or registered trademark of Cray Inc.

References

1. A. Toomre and J. Toomre, "Galactic Bridges and Tails," *Astrophys. J.* **178**, 623 (1972).
2. T. Ideker, T. Galitsky, and L. Hood, "A New Approach to Decoding Life: Systems Biology," *Ann. Rev. Genomics & Human Genetics* **2**, 343 (2001).
3. B. Moore, S. Ghigna, F. Governato, G. Lake, T. Quinn, J. Stadel, and P. Tozzi, "Dark Matter Substructure Within Galactic Haloes," *Astrophys. J.* **524**, L19 (1999).
4. J. P. Ostriker and P. Steinhardt, "New Light on Dark Matter," *Science* **300**, 1909 (2003).
5. B. Moore, N. Katz, G. Lake, A. Dressler, and A. Oemler, "Galaxy Harassment and the Evolution of Clusters of Galaxies," *Nature* **379**, 613 (1996).
6. D. G. York, J. Adelman, J. E. Anderson, S. F. Anderson, J. Annis, N. A. Bahcall, J. A. Bakken, R. Barkhouser, S. Bastian, E. Berman, W. N. Boroski, S. Bracker, C. Briegel, J. W. Briggs, J. Brinkmann, R. Brunner, S. Burles, L. Carey, M. A. Carr, F. J. Castander, B. Chen, P. L. Colestock, A. J. Connolly, J. H. Crocker, I. Csabai, P. C. Czarapata, J. E. Davis, M. Doi, T. Dombeck, D. Eisenstein, N. Ellman, B. R. Elms, M. L. Evans, X. Fan, G. R. Federwitz, L. Fiscelli, S. Friedman, J. A. Friedman, M. Fukugita, B. Gillespie, J. E. Gunn, V. K. Gurbani, E. de Haas, M. Haldeman, F. H. Harris, J. Hayes, T. M. Heckman, G. S. Hennessy, R. B. Hindsley, S. Holm, D. J. Holmgren, C. Huang, C. Hull, D. Husby, S. Ichikawa, T. Ichikawa, Ž. Ivezić, S. Kent, R. S. J. Kim, E. Kinney, M. Klane, A. N. Kleinman, S. Kleinman, G. R. Knapp, J. Korienek, R. G. Kron, P. Z. Kunszt, D. Q. Lamb, B. Lee, R. F. Leger, S. Limmongkol, C. Lindenmeyer, D. C. Long, C. Loomis, J. Loveday, R. Lucinio, R. H. Lupton, B. MacKinnon, E. J. Mannery, P. M. Mantsch, B. Margon, P. McGehee, T. A. McKay, A. Meiksin, A. Merelli, D. G. Monet, J. A. Munn, V. K. Narayanan, T. Nash, E. Neilsen, R. Neswold, H. J. Newberg, R. C. Nichol, T. Nicinski, M. Nonino, N. Okada, S. Okamura, J. P. Ostriker, R. Owen, A. G. Pauls, J. Peoples, R. L. Peterson, D. Petravick, J. R. Pier, A. Pope, R. Pordes, A. Prosapio, R. Rechenmacher, T. R. Quinn, G. T. Richards, M. W. Richmond, C. H. Rivetta, C. M. Rockosi, K. Ruthmansdorfer, D. Sandford, D. J. Schlegel, D. P. Schneider, M. Sekiguchi, G. Sergey, K. Shimasaku, W. A. Siegmund, S. Smei, J. A. Smith, S. Snedden, R. Stone, C. Stoughton, M. A. Strauss, C. Stubbs, M. SubbaRao, A. S. Szalay, I. Szapudi, G. P. Szokoly, A. R. Thakar, C. Tremonti, D. L. Tucker, A. Uomoto, D. Vanden Berk, M. S. Vogeley, P. Waddell, S. Wang, M. Watanabe, D. H. Weinberg, B. Yanny, and N. Yasuda, "The Sloan Digital Sky Survey: Technical Summary," *Astronom. J.* **120**, 1579 (2000).
7. A. Jenkins, C. S. Frenk, S. D. M. White, J. M. Colberg, S. Cole, A. E. Evrard, H. M. P. Couchman, and N. Yoshida, "The Mass Function of Dark Matter Haloes," *Monthly Notices Roy. Astronom. Soc.* **321**, 372 (2001).
8. B. Moore, N. Katz, and G. Lake, "On the Destruction and Overmerging of Dark Halos in Dissipationless N-Body Simulations," *Astrophys. J.* **457**, 455 (1996).
9. G. Lake, N. Katz, T. Quinn, and J. Stadel, "Cosmological N-Body Simulation," *Proceedings of the 7th SIAM Conference on Parallel Processing*, 1995, p. 307.
10. J. Barnes and P. Hut, "A Hierarchical $O(N \log N)$ Force-Calculation Algorithm," *Nature* **324**, 446 (1986).
11. P. Saha and S. Tremaine, "Long-Term Planetary Integration with Individual Time Steps," *Astronom. J.* **108**, 1962 (1994).
12. M. P. Calvo and J. M. Sanz-Serna, "The Development of Variable-Step Symplectic Integrators, with Application to the Two-Body Problem," *SIAM J. Sci. Comput.* **14**, 936 (1993).
13. P. Hut, J. Makino, and S. McMillan, "Building a Better Leapfrog," *Astrophys. J.* **443**, L93 (1995).
14. T. R. Quinn, J. Stadel, and G. Lake, "Time Stepping N-Body Simulations," *Astrophys. J.*, submitted, 2003.
15. N. Katz, T. Quinn, E. Bertschinger, and J. M. Gelb, "Formation of Quasars at High Redshift," *Monthly Notices Roy. Astronom. Soc.* **270**, L71 (1994).
16. L. Greengard and J. Rockhlin, "A Fast Algorithm for Particle Simulations," *Comput. Phys.* **73**, 325 (1987).
17. H. M. P. Couchman, "Mesh-Refined P3M—A Fast Adaptive N-Body Algorithm," *Astrophys. J.* **368**, L23 (1991).
18. H. M. P. Couchman, P. A. Thomas, and F. R. Pearce, "Hydra: An Adaptive-Mesh Implementation of P 3M-SPH," *Astrophys. J.* **452**, 797 (1995).
19. F. R. Pearce and H. M. P. Couchman, "Hydra: A Parallel Adaptive Grid Code," *New Astron.* **2**, 411 (1997).
20. P. Bode, J. P. Ostriker, and G. Xu, "The Tree Particle-Mesh N-Body Gravity Solver," *Astrophys. J. Suppl.* **128**, 561 (2000).
21. J. G. Stadel, "Cosmological N-Body Simulations and Their Analysis," Ph.D. thesis, University of Washington, Pullman, 2001.
22. A. W. Appel, "An Efficient Program for Many-Body Simulations," *SIAM J. Sci. Stat. Comput.* **6**, 85 (1961).
23. A. Dressler, "The Supergalactic Plane Redshift Survey—A Candidate for the Great Attractor," *Astrophys. J.* **321**, 519 (1988).
24. A. G. Riess, P. E. Nugent, R. L. Gilliland, B. P. Schmidt, J. Tonry, M. Dickinson, R. I. Thompson, T. Budavári, S. Casertano, A. S. Evans, A. V. Filippenko, M. Livio, D. B. Sanders, A. E. Shapley, H. Spinrad, C. C. Steidel, D. Stern, J. Surace, and S. Veilleux, "The Farthest Known Supernova: Support for an Accelerating Universe and a Glimpse of the Epoch of Deceleration," *Astrophys. J.* **560**, 49 (2001).
25. C. L. Bennett, M. Halpern, G. Hinshaw, N. Jarosik, A. Kogut, M. Limon, S. S. Meyer, L. Page, D. N. Spergel, G. S. Tucker, E. Wollack, E. L. Wright, C. Barnes, M. R. Greason, R. S. Hill, E. Komatsu, M. R. Nolta, N. Odegard, H. V. Peiris, L. Verde, and J. L. Weiland, "First-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Preliminary Maps and Basic Results," *Astrophys. J. Suppl.* **148**, 1 (2003).
26. F. Governato, L. Mayer, J. Wadsley, J. P. Gardner, B. Willman, E. Hayashi, T. Quinn, J. Stadel, and G. Lake, "The Formation of a Realistic Disk Galaxy in Lambda Dominated Cosmologies," *Astrophys. J.*, submitted, 2003.
27. J. P. Ostriker and Y. Suto, "The Mach Number of the Cosmic Flow—A Critical Test for Current Theories," *Astrophys. J.* **348**, 378 (1990).
28. T. R. Quinn, S. Tremaine, and M. Duncan, "A Three Million Year Integration of the Earth's Orbit," *Astronom. J.* **101**, 2287 (1991).
29. A. Berger, J. Imbrie, G. Hayes, G. Kukla, and B. Saltzman, *Milankovich and Climate; Understanding the Response to Astronomical Forcing*, Reidel, Dordrecht, The Netherlands, 1984.
30. M. J. Duncan and T. Quinn, "The Long-Term Dynamical Evolution of the Solar System," *Ann. Rev. Astron. & Astrophys.* **31**, 265 (1993).
31. G. W. Wetherill, "Possible Consequences of Absence of Jupiters in Planetary Systems," *AP & SS* **212**, 23 (1994).
32. H. Poincaré, *Les Méthodes Nouvelles de la Mécanique Celeste*, Gauthiers-Villars, Paris, 1892.
33. V. I. Arnold, "Small Denominators and the Problem of the Stability of the Solar System," Report presented at

- the IVth All-Union Mathematical Congress, Leningrad, 1961.
34. J. Laskar, "A Numerical Experiment on the Chaotic Behaviour of the Solar System," *Nature* **338**, 237 (1989).
 35. W. J. Eckert, D. Brouwer, and G. M. Clemence, "Coordinates of the Five Outer Planets, 1653–2060," *Astronomical Papers Prepared for the Use of the American Ephemeris and Nautical Almanac*, Vol. 12, xxxi, 327 pp., U.S. Government Printing Office, Washington, DC, 1951.
 36. C. J. Cohen and E. C. Hubbard, "Libration of the Close Approaches of Pluto to Neptune," *Astronom. J.* **70**, 10 (1965).
 37. C. J. Cohen, E. C. Hubbard, and C. Oesterwinter, *Elements of the Outer Planets for One Million Years*, Nautical Almanac Office, U.S. Naval Observatory, Washington, DC; for sale by the Superintendent of Documents, U.S. Government Printing Office, 1973.
 38. H. Kinoshita and H. Nakai, "Motions of the Perihelions of Neptune and Pluto," *Celest. Mech.* **34**, 203 (1984).
 39. J. H. Applegate, M. R. Douglas, Y. Gursel, G. J. Sussman, and J. Wisdom, "The Outer Solar System for 200 Million Years," *Astronom. J.* **92**, 176 (1986).
 40. M. Carpino, A. Milani, and A. M. Nobili, "Long-Term Numerical Integrations and Synthetic Theories for the Motion of the Outer Planets," *Astronom. & Astrophys.* **181**, 182 (1987).
 41. G. J. Sussman and J. Wisdom, "Numerical Evidence that the Motion of Pluto Is Chaotic," *Science* **241**, 433 (1988).
 42. D. L. Richardson and C. F. Walker, "Numerical Simulation of the Nine-Body Planetary System Spanning Two Million Years," *J. Astronaut. Sci.* **37**, 159 (1989).
 43. G. J. Sussman and J. Wisdom, "Chaotic Evolution of the Solar System," *Science* **257**, 56 (1992).
 44. T. Ito and K. Tanikawa, "Long-Term Integrations and Stability of Planetary Orbits in our Solar System," *Monthly Notices Roy. Astronom. Soc.* **336**, 483 (2002).
 45. F. Varadi, B. Runnegar, and M. Ghil, "Successive Refinements in Long-Term Integrations of Planetary Orbits," *Astrophys. J.* **592**, 620 (2003).
 46. J. C. Armstrong, C. B. Leovy, and T. R. Quinn, "A 1 Gyr Climate Model for Mars: New Orbital Statistics and the Importance of Seasonally Resolved Polar Processes," in preparation, 2003.
 47. J. Laskar, "Large-Scale Chaos in the Solar System," *Astronom. & Astrophys.* **287**, L9 (1994).
 48. P. Saha, J. Stadel, and S. Tremaine, "A Parallel Integration Method for Solar System Dynamics," *Astronom. J.* **114**, 409 (1997).
 49. C. Gear, "Waveform Methods for Space and Time Parallelism," *J. Comput. & Appl. Math.* **38**, 137 (1951).
 50. P. Saha, "The Use of Tangent Equations to Detect Chaos in Solar System Dynamics," in preparation, 2004.
 51. J. J. Lissauer, "Planet Formation," *Ann. Rev. Astron. & Astrophys.* **31**, 129 (1993).
 52. G. W. Wetherill and G. R. Stewart, "Accumulation of a Swarm of Small Planetesimals," *Icarus* **77**, 330 (1989).
 53. C. Beaugé and S. J. Aarseth, "N-Body Simulations of Planetary Formation," *Monthly Notices Roy. Astronom. Soc.* **245**, 30 (1990).
 54. S. Ida and J. Makino, "Scattering of Planetesimals by a Protoplanet—Slowing Down of Runaway Growth," *Icarus* **106**, 210 (1993).
 55. M. Lecar and S. J. Aarseth, "A Numerical Simulation of the Formation of the Terrestrial Planets," *Astrophys. J.* **305**, 564 (1986).
 56. S. Ida and J. Makino, "N-Body Simulation of Gravitational Interaction Between Planetesimals and a Protoplanet; I—Velocity Distribution of Planetesimals," *Icarus* **96**, 107 (1992).
 57. S. Ida and J. Makino, "N-Body Simulation of Gravitational Interaction Between Planetesimals and a Protoplanet; II—Dynamical Friction," *Icarus* **98**, 28 (1992).
 58. S. J. Aarseth, D. N. C. Lin, and P. L. Palmer, "Evolution of Planetesimals. II. Numerical Simulations," *Astrophys. J.* **403**, 351 (1993).
 59. E. Kokubo and S. Ida, "On Runaway Growth of Planetesimals," *Icarus* **123**, 180 (1996).
 60. E. Kokubo and S. Ida, "Oligarchic Growth of Protoplanets," *Icarus* **131**, 171 (1998).
 61. J. E. Chambers and G. W. Wetherill, "Making the Terrestrial Planets: N-Body Integrations of Planetary Embryos in Three Dimensions," *Icarus* **136**, 304 (1998).
 62. D. C. Richardson, T. Quinn, J. Stadel, and G. Lake, "Direct Simulation of Planet Formation with a Million Planetesimals: First Results," *Bull. Amer. Astronom. Soc.* **30**, 1052 (1998).
 63. D. C. Richardson, T. Quinn, J. Stadel, and G. Lake, "Direct Large-Scale N-Body Simulations of Planetesimal Dynamics," *Icarus* **143**, 45 (2000).
 64. J. J. Lissauer and V. S. Safronov, "The Random Component of Planetary Rotation," *Icarus* **93**, 288 (1991).
 65. L. Dones and S. Tremaine, "On the Origin of Planetary Spins," *Icarus* **103**, 67 (1993).
 66. J. C. Gradie, C. R. Chapman, and E. F. Tedesco, "Distribution of Taxonomic Classes and the Compositional Structure of the Asteroid Belt," *Asteroids II*, R. P. Binzel, T. Gehrels, and M. S. Matthews, Eds., University of Arizona Press, Tucson, 1989, p. 316.
 67. W. Dehnen, "A Very Fast and Momentum-Conserving Tree Code," *Astrophys. J.* **536**, L39 (2000).
 68. D. C. Richardson, Z. M. Leinhardt, H. J. Melosh, W. F. Bottke, and E. Asphaug, "Gravitational Aggregates: Evidence and Evolution," *Asteroids III*, W. F. Bottke, Jr., A. Cellino, P. Paolicchi, and R. P. Binzel, Eds., University of Arizona Press, Tucson, 2002, p. 501.
 69. J. Makino, "N-Body Simulation on GRAPE-6," *ASP Conference Series 263: Stellar Collisions, Mergers and Their Consequences*, 2002, p. 389.
 70. Blue Gene Team, "Blue Gene: A Vision for Protein Science Using a Petaflop Supercomputer," *IBM Syst. J.* **40**, 310 (2001).
 71. B. G. Fitch, R. S. Germain, M. Mendell, J. Pitera, M. Pitman, A. Rayshubskiy, Y. Sham, F. Suits, W. Swope, T. J. C. Ward, Y. Zhestov, and R. Zhou, "Blue Matter, an Application Framework for Molecular Simulation on Blue Gene," *J. Parallel & Distr. Comput.* **63**, 759 (2003).
 72. S. McMillan, P. Hut, J. Makino, M. Norman, and F. J. Summers, "Design Studies on Petaflops Special Hardware: Hardware for Particle Simulation," presented at the Symposium on the Frontiers of Massively Parallel Computing, 1997.

Received October 16, 2003; accepted for publication December 17, 2003

George Lake *Department of Physics, Washington State University, P.O. Box 642814, Pullman, Washington 99164 (lake@wsu.edu).* In 2003 Dr. Lake became the William Band Professor of Theoretical Physics at Washington State University. He received a B.A. degree in physics and astronomy from Haverford College in 1975, and an M.A. degree in 1977 and a Ph.D. degree in 1980, both in physics, from Princeton University. Dr. Lake held a NATO fellowship at the University of Cambridge and remains a Fellow of Churchill College. He was a member of the technical staff at AT&T Bell Labs before moving to the University of Washington in 1985. In 2000, he became Professor and CIO at the Institute for Systems Biology in Seattle. Dr. Lake has served as the Project Scientist for the NASA High Performance Computing and Communications/Earth and Space Sciences Project, and he has been the Chief Scientist at the Arctic Regional Supercomputing Center. He was also appointed to an endowed chair in High Performance Computing at the University of Alaska. Dr. Lake's interests range from the simulations presented here to the origin of the Cambrian explosion and mammalian radiation as well as the properties of biological networks.

Thomas Quinn *Department of Astronomy, University of Washington, P.O. Box 351580, Seattle, Washington 98195 (trq@astro.washington.edu).* Dr. Quinn is a Professor of Astronomy in the Department of Astronomy and an Adjunct Professor in the Department of Physics at the University of Washington in Seattle. He received a B.Sc. degree in engineering physics in 1982 from Lehigh University and a Ph.D. degree in astrophysical sciences in 1986 from Princeton University. He held two postdoctoral positions, first at the Canadian Institute for Theoretical Astrophysics from 1986 to 1989 and then at the University of Oxford from 1989 to 1993. He joined the University of Washington as an Assistant Research Professor in 1993. Dr. Quinn's current work includes surveys of small bodies in the solar system, the dynamics of planet formation, the structure of galaxy halos, and large-scale structure formation.

Derek C. Richardson *Department of Astronomy, University of Maryland, College Park, Maryland 20742 (dcr@astro.umd.edu).* Dr. Richardson is an Assistant Professor of Astronomy in the Department of Astronomy at the University of Maryland at College Park. He received B.Sc. degrees in astronomy and physics in 1990 from the University of British Columbia and a Ph.D. degree in astrophysics in 1993 from the University of Cambridge. He held two postdoctoral positions, first at the Canadian Institute for Theoretical Astrophysics from 1993 to 1996, then at the University of Washington from 1996 to 1999. He was promoted briefly to Research Assistant Professor at the University of Washington before accepting his current position at the University of Maryland in 2000. Dr. Richardson's current work consists primarily of numerical studies of planetesimal dynamics and planet formation, with an emphasis on the collisional evolution of small self-gravitating bodies such as asteroids, comets, and the particles in planetary rings.

Joachim Stadel *University of Zurich, Institute for Theoretical Physics, Winterthurerstrasse 190, 8057 Zurich, Switzerland (stadel@physik.unizh.ch).* Dr. Stadel is a Research Associate in Theoretical Physics with the newly established Computational Cosmology Group at the University of Zurich.

He received a B.Sc. degree in physics and mathematics in 1991 from Brandon University and a Ph.D. degree in astronomy in 2001 from the University of Washington. From 2000 to 2002, he held a postdoctoral position at the University of Victoria, where he was a CITA (Canadian Institute for Theoretical Astrophysics) National Fellow and a PIMS (Pacific Institute for Mathematical Sciences) Fellow. From 1993 to 1995, Dr. Stadel worked as a research astronomer for the NASA (National Aeronautics and Space Administration) HPCC/ESS (High Performance Computing and Communications/Earth and Space Sciences) project in Seattle. He is currently studying large-scale structure using computer simulations, and also the chaotic dynamics of few-body systems. His interests also include high-performance computing and the development of new algorithms and computational techniques.