

# Speedup to Virtual Petaflops using Adaptive Potential Solvers and Integrators for Gravitational Systems \*

George Lake<sup>††</sup>    Thomas Quinn<sup>‡</sup>    Derek C. Richardson<sup>‡</sup>    Joachim Stadel<sup>‡</sup>

## Abstract

*The orbit of any one planet depends on the combined motion of all the 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

Epochal surveys are throwing down the gauntlet for cosmological simulation. We describe three keys to meeting the challenge of N-body simulation: adaptive potential solvers, adaptive integrators and volume renormalization. With these techniques and a dedicated Teraflop facility, simulation can stay even with observation of the Universe.

We also describe some problems in the formation and stability of planetary systems. Here, the challenge is to perform accurate integrations that retain Hamiltonian properties for  $10^{13}$  timesteps.

## 1 COSMOLOGICAL N-BODY SIMULATION

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, 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 \$2Billion space observatories may be mispent. 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) (Gunn and Knapp 1992) will produce fluxes and sky positions for  $5 \times 10^7$  galaxies with redshifts for the brightest  $10^6$ . Our ambitious observational

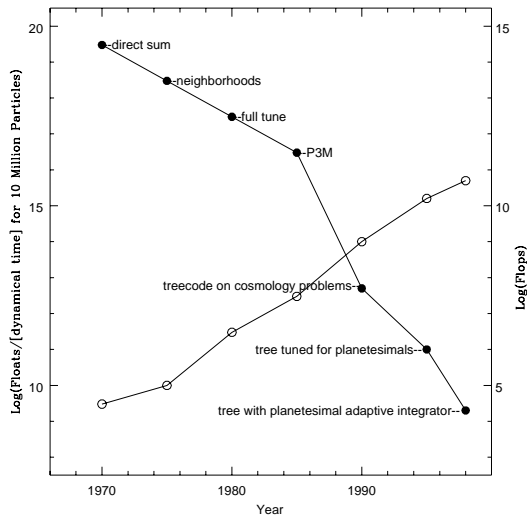


Figure 1: Gains in hardware and algorithms are compared for the N-body simulations. Algorithms are shown as filled points with the scale to the left, while hardware is open points with the scale on the right. The final algorithmic point should be considered a hopeful projection for 1999.

colleagues have cut steel and ground glass to survey a “fair volume” that *we must simulate*, but we need  $N = 10^{12}$  to do this. Direct summation of the gravitational forces using fixed timesteps would take  $10^{10}$  Teraflop-years.

We will 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 will then present the three keys to a realistic float count: 1) spatially adaptive potential solvers, 2) temporally adaptive integrators and 3) volume renormalizations. Another goal of this paper is to define “high quality simulations” and the niche science that can be done with  $N \sim 10^8$ .

## 2 THE PROGRESS OF SIMULATIONS

Over the last 20 years, the  $N$  of our simulations has increased as:  $\log_{10} N = 0.3 * (Year - 1973)$ . Figure 1 shows the relative contributions of hardware and algorithms. We can’t wait to simulate  $10^{12}$  particles,

\*Support by NASA HPCC/ESS, IRP and ATP

†NASA HPCC/ESS Project Scientist

‡Department of Astronomy, University of Washington

we have to invent the algorithms that are a thousand times faster! The power of computers has doubled every 8 months (open circles, log scale to the right) with algorithmic advances keeping the same pace (closed circles, log scale to the left). Together, the doubling time in power is 8 months, accumulating to a trillion-fold increase in less than 3 decades. We can't wait to simulate  $10^{12}$  particles, we have to invent the 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  $\sim 10^{5.7}N^{4/3}$  floats (the scaling with  $N^{4/3}$  arises from the increased time resolution needed as inter-particle separation decreases). The memory needed to run a simulation is  $\sim 10^2N$  bytes. If we fix  $N$  by filling memory, the time to run a simulation is 10 days  $\times$  (bytes/flop rate( $N/30\text{Million}$ ) $^{1/3}$ ). Current machines are well balanced for our Grand Challenge simulations. With Gigaflops and Gigabytes, we can perform simulations with  $N \sim 10^{7.5}$ . With Teraflops and Terabytes, we can simulate  $10^{10}$  particles. Simulations with  $N \sim 10^{12}$  lie in the nether world of Petaflops and Petabytes.

There are a variety of problems where  $N \sim 10^6$  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  $10^3$  particles, they dissolve before the present epoch owing to two-body relaxation in the tidal field of the cluster. To prevent this, we need  $N > 10^7$  per cluster. Scaling to the Sloan Volume yields  $N \sim 10^{12}$ .

There are  $\sim 10^{20}$  solar masses within the SDSS volume, so even  $10^{12}$  is a paltry number as each particle would represent  $10^8$  solar masses. We need a ten-fold 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 and 2) gravitational softening, discrete timesteps, force accuracy and simulation volume don't make matters even worse.  $N$  is not the figure of merit in most reported simulations—it should be! *The N-body Constitution* (Lake *et al.* 1995) 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 nega-

tive total 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 collisions leaving with positive energy. The remaining stars remaining have greater negative energies, the cluster shrinks, the gravitational binding energy increases and the stars move faster. In galaxies and clusters of galaxies, the timescale 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 timescale much shorter leading to grossly unphysical results. So, we must use  $N$  sufficient that physical heating mechanisms dominate over numerical or the numerical heating timescale is much longer than the time we simulate. We inventoried all the physical heating mechanisms experienced by galaxies in clusters and discovered a unique new phenomena we call “galaxy harassment” .

### 3 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 wall-clock-days. The challenge lies in dividing work amongst 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 forsaken adaptive mesh codes to concentrate on tree-codes (Barnes and Hut 1986) that can be made fully spatially and temporally adaptive. The latter 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 insure accurate 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 of  $\sim 50$ .

The leapfrog time evolution operator,  $D(\tau/2)$

$K(\tau)D(\tau/2)$ , is the one most often used:

$$\begin{aligned} \text{Drift, } D(\tau/2); \quad \mathbf{r}_{n+1/2} &= \mathbf{r}_n + \frac{1}{2}\tau\mathbf{v}_n, \\ \text{Kick, } K(\tau); \quad \mathbf{v}_{n+1} &= \mathbf{v}_n + \tau\mathbf{a}(\mathbf{r}_{n+1/2}), \\ \text{Drift, } D(\tau/2); \quad \mathbf{r}_{n+1} &= \mathbf{r}_{n+1/2} + \frac{1}{2}\tau\mathbf{v}_{n+1} \end{aligned}$$

where  $\mathbf{r}$  is the position vector,  $\mathbf{v}$  is the velocity,  $\mathbf{a}$  is the acceleration, and  $\tau$  is the timestep. This operator evolves the system under the Hamiltonian

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

where  $H_{err}$  is of order  $\tau^2$  (Saha and Tremaine 1994). The existence of this surrogate Hamiltonian ensures that the leapfrog is symplectic—it is the exact solution of an approximate Hamiltonian. Errors explore 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  $(\mathbf{r}, \mathbf{v}, \tau)$  in a manner that may not be Hamiltonian, hence it is neither symplectic nor time-reversible. The results can be awful (Calvo and Sanz-Serna 1993). Time reversibility can be restored (Hut, Makino and McMillan 1994) 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 that “adjusts” the timestep,  $A$ , yet retains time reversibility and only calculates a force if it is used to complete the timestep (Quinn *et al.* 1997). This is done by choosing  $A$  such that it commutes with  $K$ , so that  $DAKD$  is equivalent to  $DKAD$ . Since  $K$  only changes the velocities, an  $A$  operator that depends entirely on positions satisfies the commutation requirement. The “natural definition” of timestep,  $\propto 1/\sqrt{\text{density}}$ , is ideal but it is difficult to define when a few particles are close together. Synchronization is maintained by choosing timesteps that are a power-of-two subdivision of the largest timestep,  $\tau_s$ . That is,  $\tau_i = \frac{\tau_s}{2^{n_i}}$ , where  $\tau_i$  is the timestep of a given particle. 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 a resolution of we can achieve  $10^3$  parsec resolution within a cosmological volume of size  $10^8$  parsec to study the origin of quasars (Katz *et al.* 1994).

#### 4 SIMULATING THE SLOAN VOLUME

Our proposed program to simulate the Sloan Volume before the millenia is as follows:

- Simulate the entire volume  $(800 \text{ Mpc})^3$  with  $N = 10^{10}$ , each with a mass of  $10^{10.5}M_\odot$ .
- “Renormalize” dozens of groups, clusters, etc. and simulate with  $10^8$ – $10^9$  particles.

The total cost for the first simulation is roughly a Teraflop-year and requires a machine with a Terabyte of memory. The second sequence of simulations should be designed to have roughly equal computational cost, but will require less memory.

#### 5 THE FATE OF THE SOLAR SYSTEM

Advances in hardware and numerical methods finally enable us to integrate the solar system for its lifetime. Such an integration is a 1,000 fold advance on the best longest accurate integration ever performed (Laskar, Quinn and Tremaine 1992) and can address numerous questions:

**Is the Solar System stable?** Do all 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 affect of orbital changes on the planetary climates?** According to the Milankovich hypothesis, climate 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 (Berger *et al.* 1984). Remarkably, the geophysical data (primarily the volume of water locked up in ice as determined by the  $^{18}O/^{16}O$  ratio in seabed 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 is so short?

**What is the stability of other planetary system?** How are the giant planets related to terrestrial planets in the “inhabitable zone” between boiling and freezing of water by the central star? Without such a cleansing of planetesimals from the solar system by giant planets (Duncan and Quinn 1993), the bombardment of the Earth by asteroids would be steady and frequent throughout the main sequence lifetime of the Sun (Wetherill 1994). The chaos produced by the Jupiter and Saturn may have played a role in insuring that plan-

etesimals collided to form the terrestrial planets <sup>1</sup>, 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.

## 6 INTEGRATING NINE PLANETS FOR 10<sup>11</sup> 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 (1892) showed that these expansions don't converge owing to resonances. Using the KAM theorem, Arnold (1961) derived constraints on planet masses, eccentricities, and inclinations sufficient to insure stability. The solar system does not meet his stringent conditions, but this does not imply that it is unstable.

Laskar (1989) tested the quasi-periodic hypothesis by numerically integrating the perturbations calculated to second order in mass and fifth order in eccentricities and inclinations, ~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.

The second method for attacking the stability problem is to explicitly integrate the planets' orbits (Table 1). As early as 1965, Pluto's behaviour was suspicious. In the last ten years, it has become clear that the solar system is chaotic. However, the source of the chaos is unclear as the system of resonances is complex and the the Lyapunov exponent appear to be sensitive to fine details of initial conditions.

Nonetheless, the Solar System is almost certainly chaotic. Laskar (1994) looked at the fate of Mercury and estimates the chance of ejection in the next few billion years approaches 50%. Our belief in the apparent regularity of the solar system may owe to our inability to know that before the last few ejections, there were 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 timescale of 4-5 Myr time, an error as small as 10<sup>-10</sup> 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.<sup>2</sup>

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

<sup>2</sup>Don't let this go beyond this room, environmental impact statements are already tough enough! 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

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 quad precision to overcome roundoff problems. The IBM Power 2 series is the current machine of choice, evolving the solar system at ~ 10<sup>9</sup> times faster than "real time", this is 1-3 orders of magnitude faster than other available cpus. To understand any chaotic, we will need to see it by an independent means and devise methods to determine its underlying source.

**Table 1: Solar System Integration History**

Year Ref	Length (Myr)	# Planets	GR?	Earth's Moon?
1951 Eckert...	0.00035	5	no	no
1965 Cohen...	0.12	5	no	no
1973 Cohen...	1.	5	no	no
1986 Applegate...	217.	5	no	no
	3.	8	no	no
1986 Nobili	100.	5	yes	no
1988 Sussman...	845.	5	no	no
1989 Richardson...	2.	9	no	no
1991 Quinn...	6.	9	yes	yes
1992 Sussman...	100.	9	yes	yes
1999 us	10,000	9	yes	yes

**A Parallel Method** doesn't 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*) (Saha, Stadel and Tremaine 1997). In this method, each processor takes a different time-slice; processor 2's initial conditions are processor 1's final conditions 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 (Gear 1991). However, Kepler ellipses are a good guess to the orbits for a timescale that is proportional to the ratio of the Sun's mass to Jupiter's. 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

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.

and exploits the nearness to an exactly soluble system; otherwise errors grow quadratically with time. *TSCM* will enable us to integrate  $\sim 0.5$  Gyr per day on a 512 node SP2—a speed-up over real-time of  $10^{11}$ . This will make it feasible to study the stability 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 will use a new technique to gauge the origin of instabilities (the “tangent equation method”) (Saha 1997). 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 the linearized or tangent equations—the equations for *differences* from nearby orbits. We will integrate the tangent equations along with the main orbit equations.

## 7 COSMOLOGY MEETS COSMOGONY: PLANETARY SYSTEM FORMATION

Theories of Solar System formation are traditionally divided into four stages (Lissauer 1993): 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 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 1) 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 simulataneously for  $10^7$  dynamical times, a ten-million-fold improvement that makes us enthusiastic! Only statistical methods (Wetherill and Stewart 1989) employing prescriptions for the outcomes of gravitational encounters have been used to peek at this regime.

We reach an important threshold at  $N \sim 10^7$  in

ref	$N$	$t$ (yr)	$\Delta a$ (AU)	Col?	G?
Lecar....	200	$6 \times 10^4$	0.5–1.5	a	n
Beaugé...	200	$6 \times 10^5$	0.6–1.6	abcf	n
Ida...	800	5000	0.3	–	n
Aarseth...	400	$1.2 \times 10^4$	0.04	ab	n
Chambers...	100	$10^8$	0.5–2.0	a	y
Kokubo...	5000	$2 \times 10^4$	0.4	a	n
Richardson...	$10^5$	$10^3$	1.2–3.6	a	y
goal	$10^7$	$10^7$	0.5–2.0	abcf	y

Table 1: Highlights of advances in direct simulations of the formation of the inner planets.  $N$  is the maximum number of planetesimals used in the simulation,  $t$  is the longest integration time, and  $\Delta a$  is either the width of the simulation region at 1 AU or the actual range in orbital distance. If collisions are included in the simulations, details are noted by: a=agglomeration; b=bouncing; c=cratering; f=fragmentation. The final column shows whether perturbations from one or more giant planets are included.

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 (Beaugé and Aarseth 1990). Eventually gravitational scattering increases the planetesimal eccentricities to such an extent that collisions result in fragmentation, not growth. The embryos will continue to grow owing to their large mass, but at a slower rate as their “feeding zones” are depleted (Ida and Makino 1993). The total mass of our planetary system is  $448M_{\oplus}$  or  $3.6 \times 10^4 M_{lunar}$ , while the inner planetesimal disk amenable to simulation had a mass  $\sim 10^2 M_{lunar}$ . To capture both growth and fragmentation (Wetherill and Stewart 1989) requires a minimum particle mass of  $10^{-5} M_{lunar}$ , leading to our target  $N \sim 10^7$ .

A detailed direct simulation of planet formation can address a variety of important questions, including: Was there runaway growth of a few embryos, or a continuously evolving homogeneous mass distribution? How does the primordial surface density alter the evolution? What fixes the spin orientation and period of the planets—uniform spin-up from planetesimal accretion (Lissauer and Safronov 1991), or a stochastic process dominated by the very last giant collisions (Dones and Tremaine 1993)? Is it feasible that the Earth suffered a giant impact late in its growth that led to the formation of the Moon (Benz *et al.* 1986)? How much radial mixing was there and can it explain observed compositional gradients in the asteroid belt (Gradie, Chapman and Tedesco 1989)? Finally, what is the dominant physical mechanism that drives the late stages of growth—




Figure 2: Mass density of a  $10^6$ -particle simulation after 250 yr. Bright shades represent regions of high density. The dot at the top right is Jupiter. The disk extends from 0.8 AU, just inside 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 mean-motion resonances.

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 (presently merging and bouncing are implemented). Integrations are carried out in the heliocentric frame and may include the giant planets as perturbers. Auxilliary 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.

Figures 2 and 3 show the mass density and  $a$  vs.  $e$ , respectively, at the end of a 250-yr run that began with  $10^6$  identical cold planetesimals in a disk from 0.8 to 3.8 AU 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 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




Figure 3: Plot of semi-major axis  $a$  vs. eccentricity  $e$  for the  $10^6$ -particle simulation, showing only every 100<sup>th</sup> particle to prevent overcrowding. The peaks in  $e$  correspond to mean-motion resonances with Jupiter; there are similar features in plots of  $a$  vs. inclination  $i$  (not shown). The circles are scaled by mass for emphasis.

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. There are corresponding features in Fig. 3 which show how Jupiter stirs up 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 8 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 get to 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, and a perturbation component owing to the force contributions of all the other particles. In this regime, timesteps can be of order 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.

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 only performed 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. 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.

## 8 Summary: Virtual Petaflops

Past planetesimal simulations used codes with an algorithmic complexity that would be similar to the point labeled "full tune" in Fig. 1 and computers with speeds of  $\sim 10$  Mflops. (Special purpose "GRAPE" hardware of  $\sim 10^6$  Mflops has been used, but such implementations involve sums over all interactions so they are closer to the direct sum case in floating point cost [cf. Kokubo and Ida 1998].) Our algorithms result in a collective speed-up  $\gtrsim 10^8$  for simulations with  $N \sim 10^7$  (with rough accounting for the reduction in  $N$  over the course of the simulation). It must be emphasized that to attain the desired performance, both hardware and algorithm improvements are required. Figure 1 shows that the speed-up 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.* (1997) asserted that if Grape-6 is built, it could use its Petaflop speed to

follow  $10^6$  planetesimals by the year 2000. They claimed that this would be a seven-year advantage over General Purpose Computers that would only be able to follow  $10^4$  particles by the year 2000. Our test simulations *without the new integrator* are already 10 years ahead of their projections. We argue that our approach will beat efforts that rely on special purpose hardware with encoded algorithms for at least the next decade.

Sir Isaac would love to see the enhancement of "the entire human intellect" by high performance computing.

## REFERENCES

- Aarseth, S. J.; D. N. C. Lin; and P. L. Palmer. 1993. *Evolution of planetesimals. II. Numerical simulations*. Ap.J., 403, 351–76.
- Applegate, J. H.; M. R. Douglas; Y. Gursel; G. J. Sussman and J. Wisdom. 1986. *The outer solar system for 210 million years*. A.J., 92, 176–94.
- Arnold, V. I. 1961. *Small Denominators and the Problem of Stability in Classical and Celestial Mechanics*. In Report to the IV<sup>th</sup> All-Union Math. Congress, Leningrad.
- Barnes, J. 1986. *An efficient N-body algorithm for a fine-grain parallel computer*. In The Use of Supercomputers in Stellar Dynamics, P. Hut and S. McMillan, eds., Springer Verlag, New York, 175–80.
- Barnes J. and P. Hut. 1986. *A Hierarchical  $O(N \log N)$  Force-Calculation Algorithm*. Nature, 324, 446–50.
- Beaugé, C. and S. J. Aarseth. 1990. *N-body simulations of planetary formation*. M.N.R.A.S., 245, 30–9.
- Benz, W.; W. L. Slattery and A. G. W. Cameron. 1986. *The origin of the Moon and the single-impact hypothesis. I*. Icarus, 66, 515–35.
- Berger, A.; J. Imbrie; J. Hayes; G. Kukla and B. Saltzman. 1984. *Milankovitch and climate: understanding the response to astronomical forcing*. Reidel. Dordrecht.
- Calvo, M. P. and J. M. Sanz-Serna. 1993. *The development of variable-step symplectic integrators with application to the two-body problems*. SIAM J. Sci. Comput., 14, 936–52.
- J. E. Chambers and G. W. Wetherill, *N-body simulations of the formation of the inner planets*, B.A.A.S., 28, (1996), 1107.
- Cohen, C. J. and E. C. Hubbard, 1965. *Libration of the Close Approaches of Pluto to Neptune*. A.J., 70, 10–22.
- Cohen, C. J.; E. C. Hubbard and C. Oesterwinter. 1973. *Elements of the outer planets for a million years*. Astron. Pap. Amer. Ephem., 22, 1–5.
- Dones L. and S. Tremaine. 1993. *On the origin of planetary spins*. Icarus, 103, 67–92.
- Duncan M. and T. Quinn. 1993. *The Long-Term Dynamical Evolution of the Solar System*. Ann. Rev. Astr. Ap., 31, 265–89.

- Eckert, W. J.; D. Brouwer and G. Clemence. 1951. *Coordinates of the five outer planets 1653–2060*. Astron. Pap. Amer. Ephem., 12, 1–47.
- Gear, C. 1991. *Waveform methods for space and time parallelism*. J. Comp. and Appl. Math., 38, 137–47.
- Gradie, J. C.; C. R. Chapman and E. F. Tedesco. 1989. *Distribution of taxonomic classes and the compositional structure of the asteroid belt*. In Asteroids II, R. P. Binzel *et al.* eds., Univ. Arizona Press, Tucson, 316–35.
- Gunn J. E. and G. R. Knapp. 1992. *The Sloan Digital Sky Survey*. In A.S.P. conf., #43, 267–79.
- Hut, P.; J. Makino and S. McMillan. 1994. *Building a better leapfrog*. Ap.J.Lett., 443, L93–6.
- Ida, S. and J. Makino, 1992a. *N-body simulation of gravitational interaction between planetesimals and a protoplanet. I. Velocity distribution of planetesimals*, Icarus, 96, pp. 107–20.
- Ida, S. and J. Makino, 1992b. *N-body simulation of gravitational interaction between planetesimals and a protoplanet. II. Dynamical friction*, Icarus, 98, pp. 28–37.
- Ida S. and J. Makino. 1993. *Scattering of Planetesimals by a Protoplanet: Slowing Down of Runaway Growth*. Icarus, 106, 210–27.
- Katz, N.; T. Quinn; E. Bertschinger and J. M. Gelb. 1994. *Formation of quasars at high redshift*. M.N.R.A.S., 270, L71–4.
- Kokubo, E. and S. Ida, 1996, *On runaway growth of planetesimals*, Icarus, 123, pp. 180–91.
- Kokubo, E. and S. Ida, 1998, *Oligarchic growth of planetesimals*, Icarus, 131, pp. 171–78.
- Lake, G.; N. Katz; T. Quinn and J. Stadel. 1995. “*Cosmological N-body Simulation*”. Proc 7th SIAM Conf. on Parallel Processing for Sci. Comp., 307–12
- Laskar, J. 1989. *A numerical experiment on the chaotic behaviour of the Solar System*. Nature, 338, 237–8.
- Laskar, J. 1994 *Large-scale chaos in the solar system*. A.A., 287, L9–12
- Lecar, M. and S. J. Aarseth, 1986 *A simulation of the formation of the terrestrial planets*, Ap.J., 305, pp. 564–79.
- Lissauer, J. J. 1993 *Planet formation*. Ann. Rev. Astr. Ap., 31, 129–74.
- Lissauer, J. J. and V. S. Safronov. 1991. *The random component of planetary rotation*. Icarus, 93, 288–97.
- Laskar, J.; T. Quinn and S. Tremaine. 1992 *Confirmation of resonant structure in the solar system*. Icarus, 95, 148–52.
- McMillan, S., P. Hut, J. Makino, M. Norman, and F. J. Summers. 1997 *Design studies on petaflops special purpose Hardware for particle simulation*. in *Frontiers: 6<sup>th</sup> Symposium on the Frontiers of Massively Parallel Computing*.
- Moore, B.; N. Katz; G. Lake; A. Dressler and A. Oemler. 1996. *Galaxy harassment and the evolution of clusters of galaxies*. Nature, 379, 613–16.
- Nobili, A. M. 1986. *LONGSTOP and the masses of Uranus and Neptune*. In Solid Bodies of the Outer Solar System. ESTEC. Noordwijk, 9–13.
- Poincaré, H. 1892. *Les methodes nouvelles de la mecanique celeste*. Gauthiers-Villars. Paris.
- Quinn, T., N. Katz, J. Stadel and G. Lake. 1997. *Time stepping N-body simulations*. In preparation.
- Quinn, T.; S. Tremaine and M. J. Duncan. 1991. *A three million year integration of the Earth’s orbit*. A.J., 101, 2287–305.
- Richardson, D. L. and C. F. Walker. 1987. *Multivalued integration of the planetary equations over the last one-million years*. Astrodynamics 1987, Soldner *et al.*, eds. Univelt. San Diego. 1473–56.
- Richardson, D. L. and C. F. Walker. 1989. *Numerical simulation of the nine-body planetary system spanning two million years*. J. Astronautical Sci., 37, 159–82.
- Richardson, D. C., G. Lake, T. Quinn, and J. Stadel, *Direct simulation of planet formation with a million planetesimals: A progress report*, B.A.A.S., 30, (1998), 765.
- Saha P. and S. Tremaine. 1992. *Symplectic integrators for solar system dynamics*. A.J., 104, 1633–40.
- Saha, P. and S. Tremaine. 1994. *Long-term planetary integration with individual time steps*. A.J., 108, 1962–9.
- Saha, P., J. Stadel and S. Tremaine. 1997. *A Parallel Integration Method for Solar System Dynamics*. A.J., 114, 409–15.
- Saha, P. 1997. *The Use of Tangent Equations to Detect Chaos in Solar System Dynamics*. in preparation.
- Stadel, J. and T. Quinn. 1998. *A generalization of Ewald summation to arbitrary multipole order*. In preparation.
- Sussman, G. J. and J. Wisdom. 1988. *Numerical evidence that the motion of Pluto is chaotic*. Science, 241, 433–7.
- Sussman, G. J. and J. Wisdom. 1992. *Chaotic Evolution of the Solar System*. Science, 257, 56–62.
- Wetherill, G. W. 1994. *Possible consequences of absence of “Jupiters” in planetary systems*. Ap. Sp. Sci., 212, 23–32.
- Wetherill G. W. and G. R. Stewart. 1989. *Accumulation of a swarm of small planetesimals*. Icarus, 77, 330–57.