N-body simulation on general- and special-purpose supercomputers.

Jun Makino
RIKEN AICS/TiTech ELSI/...

Dec 10, 2015, MODEST-15s, Kobe, Japan

Overview

- **History of direct N-body simulations**
- **Problem of individual timestep/neighbor scheme**
- **"Solution"**
- **Summary**

History of direct N-body simulations

- **Before 1986: Aarseth and scalar computers**
- **1986: blockstep scheme**
- **1992: 4th-order Hermite scheme**
- **1995, 2001: GRAPE-4 and GRAPE-6: without the neighbor scheme**
- **1999(?): NBODY6++ parallel neighbor scheme**
- **2006: Ninja scheme**
- **2011: $\mathrm{P^3T}$ scheme**
- **2012, 2015: NBODY6+GPU, NBODY6++ + GPU**

Before 1986: Aarseth and scalar computers

Individual timestep:

- **Each star has its own time t_i and timestep Δt_i.**
- **One with minimum $t_i + \Delta t_i$ is selected and updated, using the predicted position of other stars.**
- **Use 4-step, 4th-order predictor-corrector scheme with variable timestep.**

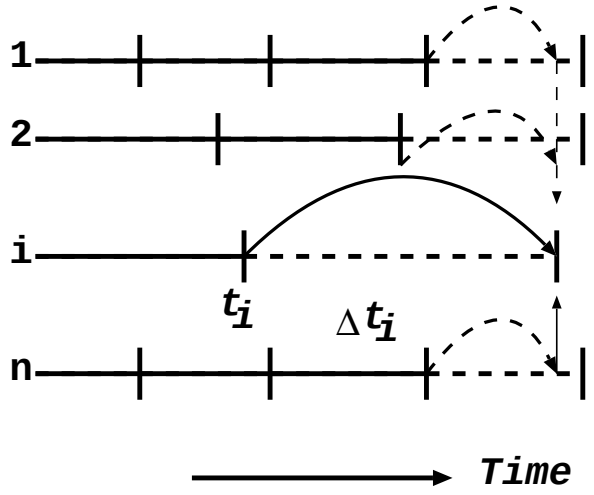

Neighbor scheme

- Separate total force on a star to two parts: neighbor and "regular" (the rest).
- Apply different timesteps to two parts.
- Change in the neighbor list need to be corrected at each regular step.

1986: blockstep scheme

THE VECTORIZATION OF SMALL-N INTEGRATORS

Stephen L. W. McMillan
Northwestern University
Department of Physics and Astronomy
Evanston, Illinois 60201

While it is very likely that supercomputers will greatly expand our understanding of the behavior of large self-gravitating systems, it is also probable that they will enable us to study exhaustively the dynamics of much smaller star clusters and associations. Through the simulation of large numbers of such systems (those containing less than, say, five hundred stars), one might hope eventually to put the small N-body problem on the same statistical footing as has already been achieved in the three- and four-body cases (Hills 1975, Hut and Bahcall 1983, Hut 1984, Mikkola 1984). Among the many issues that can be addressed are the timescales and modes of cluster dissolution, the formation and energy distribution of binary systems, and the detailed microphysics of individual stars' orbits and interactions.

In adapting existing codes for use on supercomputers, new, machine-dependent considerations inevitably arise. The current fastest machines derive much of their speed from their ability to operate in so-called "vector" mode, where the same operation is applied to successive elements of an array, at a rate much greater than would be attained if the calculation were done one element at a time (that is, in "scalar" mode). The details of the implementation of vector calculations vary widely from one machine to another, but, in general, the speed of an arithmetic operation can

1986: blockstep scheme

- "Quantize" the stepsize to powers of two. Update the stars with the same time in parallel.
- Much better parallel efficiency
- Reduction of the calculation cost of the predictor of other stars.

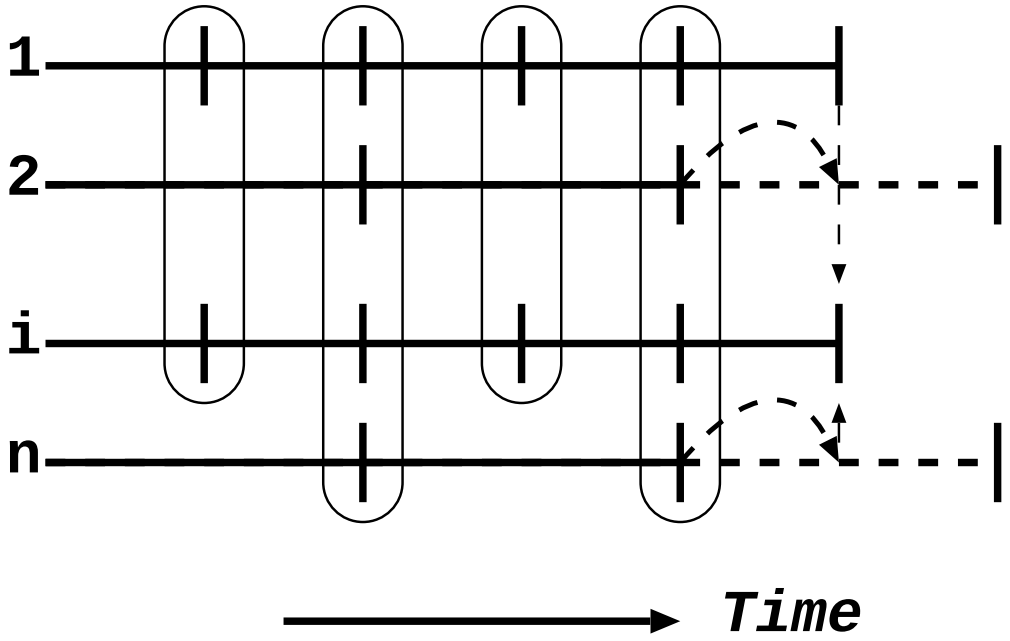

1992: 4th-order Hermite scheme

(JM and Aarseth 1992)

- Calculate "jerk" (first time derivative of acceleration) in addition to acceleration.
- Use two-point Hermite interpolation to construct 4th-order integrator. ("single step" scheme)
- Works with neighbor scheme as well.
- Several advantages
  - Easier to write
  - Longer timestep: Better efficiency on parallel/vector/special-purpose machines

1995, 2001: GRAPE-4 and GRAPE-6

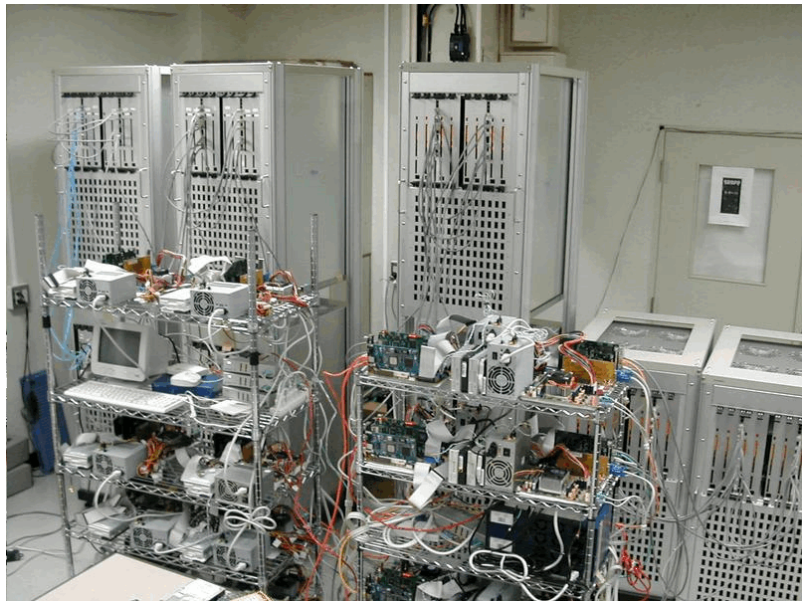

- Both rely on blockstep + Hermite scheme
- We did not implement the neighbor scheme: to minimize the work of the host CPU

1999: NBODY6++ parallel neighbor scheme

(Spurzem 1999)

- Each process keeps the complete copy of the system
- After the list of stars to be integrated in the current step is determined, each process determine its share to update.
- Updated results are exchanged between processes.
- Relies on blockstep, Hermite and neighbor scheme.
- Works great for moderate number of processes.

2006: Ninja scheme

(Nitadori+ 2006)

- Parallelize in two dimensions: Force calculation on one particle is parallelized as well.
- Communication is minimized.
- Scales to thousands of cores for $N \sim 10^5$ or less. (Kominami's poster)
- No one tried to implement neighbor scheme yet...

2011: P^3T scheme

(Oshino+ 2011, Iwasawa+ 2015)

- Latest (as far as I know) approach to use treecode to integrate collisional systems.
- Previous efforts: Jernigan and Porter 1989, McMillan and Aarseth 1993, (Richardson 1993).
- Previous works tried to combine tree and individual timestep
- We gave up, and split gravitational interaction into two terms using distance-dependent switching function (same as Mercury integrator for planetary dynamics)
- Apply treecode + leapfrog to long-range term, and Hermite to short-range term.
- Seems to be quite promising.

2012, 2015: NBODY6+GPU, NBODY6++ + GPU

(Nitadori and Aarseth 2012, Wang 2015)

- **Use GPU to the "regular" force.**
- **At present the fastest code available.**
- **The million-body problem has finally become feasible.**

Problem of individual timestep/neighbor scheme

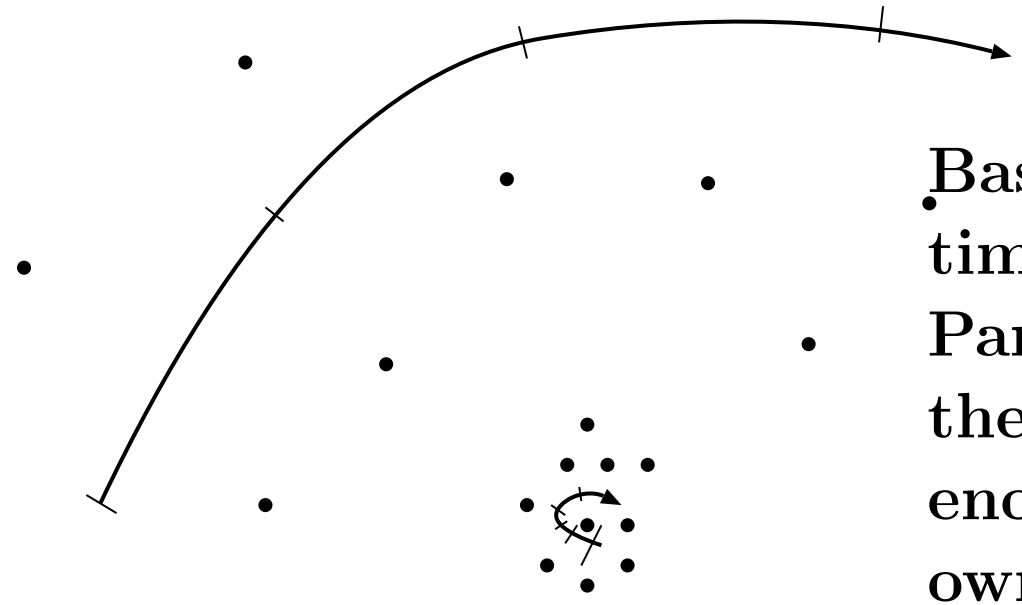

Basic idea of individual timestep:
Particles should have the timestep just enough to resolve their own orbits.

What happens to the forces from short-timescale particles to long-timescale particles?

Slide recycled from the Capri meeting

What's happening

They are integrated in a completely wrong way!

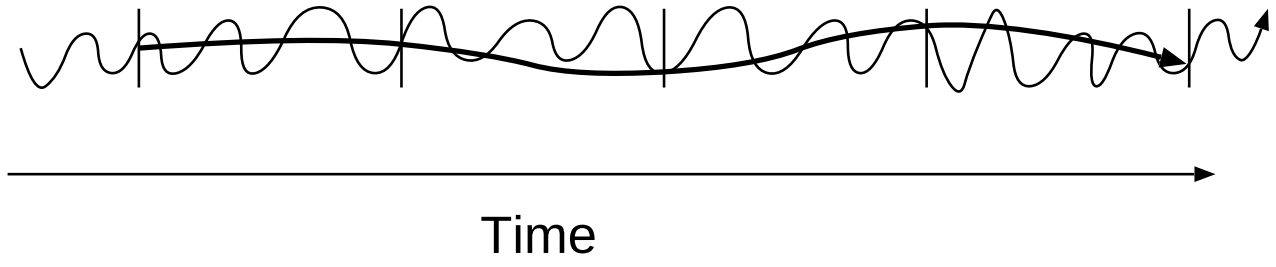

- **Forces do have rapidly changing components**
- **If the timestep is large, forces are sampled "randomly" (if the orbit is not periodic)**
- **Much more problematic with the Hermite scheme**
- **Even more problematic with the neighbor scheme**

Slide recycled from the Capri meeting

When does this happen?

- **When the orbital timescale of particles in the core becomes less than the timestep of typical particles in the cluster.**
- **Roughly speaking: If $r_c \ll r_h N^{-1/3}$**
- **Just before bounce: $r_c \sim r_h/N \ll r_h N^{-1/3}$**

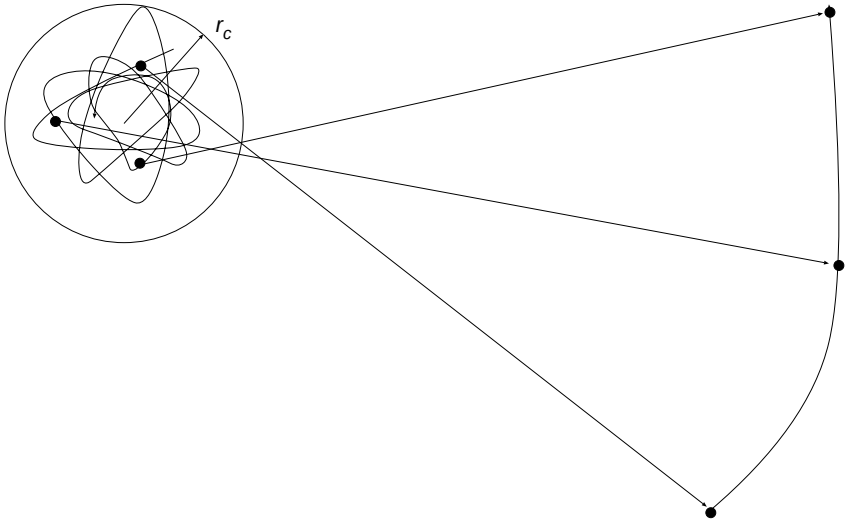

Slide recycled from the Capri meeting

Does this really matter?

- This error is actually visible: The reason why the energy error increases toward the core collapse.
- Reduction of timestep helps, but only as $\Delta t^{1.5}$
- The only way to suppress this error completely is to reduce the timesteps of all particles to less than the core crossing time
- Can we just let the error grow? No. With poor energy conservation we cannot say for sure that the calculation is "correct".

Impact on the calculation cost

- Hopefully not so severe for normal star clusters
  - the fraction of time for which the core size is small is small
  - Mass spectrum makes the core size larger
- Any system with central massive BH might be problematic.

Slide recycled from the Capri meeting

"Solution"

- With P^3T scheme, it is not unpractical to integrate all stars with the timestep smaller than the core crossing time
- Almost all calculation cost is spent by the simple shared-timestep treecode
- Parallelization and the use of accelerators are pretty efficient (can be done with our FDPS package)
- Currently we are working on this...

Some good news

P^{3}T scheme can actually conserve energy better than Hermite (Iwasawa+ 2015)

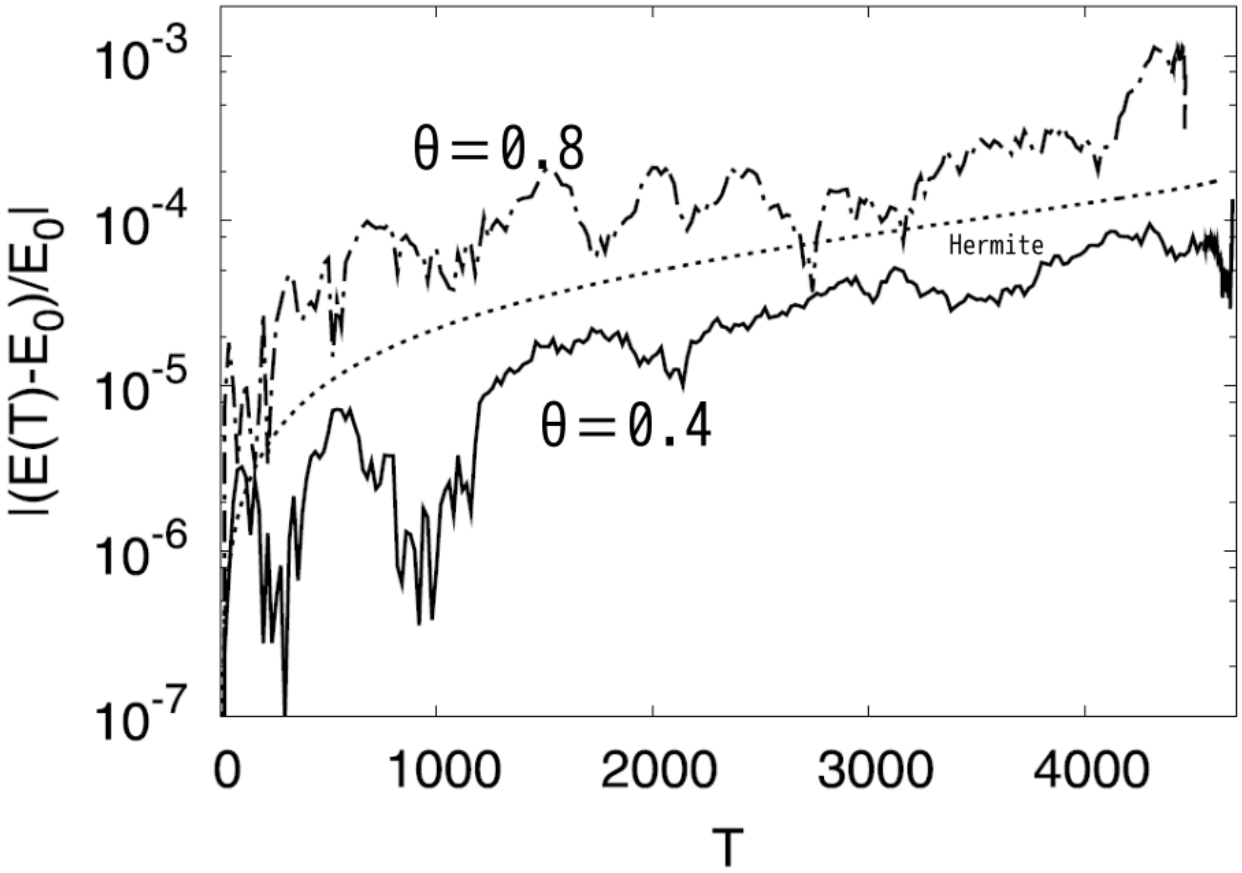

$N = 16$k, down to core collapse.

FDPS

Iwasawa+2015 (in prep...)

- **Please visit: https://github.com/FDPS/FDPS**
- **A Framework for Developing parallel Particle Simulation code**
- **FDPS offers library functions for domain decomposition, particle exchange, interaction calculation using tree.**
- **Can be used to implement pure Nbody, SPH, or any particle simulations with two-body interactions.**
- **Use essentially the same algorithm as used in our treecode implementation on K computer (GreeM, Ishiyama, Nitadori and JM 2012).**
- **Runs efficiently on K, Xeon clusters or GPU clusters (GPU version release two weeks from now).**

Summary

- **In the last three decades we have seen quite significant improvement on what we can do with direct N-body simulations.**
- **Steve's blockstep algorithm has been one of the critical ingredients.**
- **Currently, NBODY6(++) + GPU works great.**
- **To use even larger number of particle P^3T scheme with parallel treecode (on accelerators) might be necessary, and we are working on this.**