SIMULATIONS OF GALAXY INTERACTIONS ON A MASSIVELY PARALLEL COMPUTER

A.H.NELSON,
Physics Department,
University of Wales College of Cardiff,
Cardiff, CF1 3TH, UK.

ABSTRACT. Computer simulations of the encounter of a satellite galaxy with a host galaxy are described. These are being carried out using tree-code gravity on an NCUBE parallel processor with 128 processors. The particles of the host - representing both stars and gas - are distributed over the processors, and the potential of the whole ensemble is calculated using a fast parallel communications algorithm to pass information between the processors, with the result that the amount of processing and communications scales as log N where N is the number of particles.

1 Introduction

Simulations of galaxies employing collisionless particles for stars and interacting particles for gas have typically used only a few tens of thousands of particles to-date. In order to minimise the effects of numerical noise, and increase the spatial resolution many more particles need to be used. The newly available parallel computers are one means of increasing the number of particles significantly at an affordable cost. The Physics department in Cardiff has recently acquired an NCUBE parallel computer with 128 nodes, and I describe here the implementation on it of the satellite encounter calculation.

2. Numerical Model

The 128 processor nodes of the NCUBE are interconnected as a 7-dimensional hypercube (order 7), and have a DMA connection to a SUN which acts as a graphical front end. A hypercube of any order is constructed from two hypercubes of the next order down (the daughter sub-cubes) with corresponding nodes connected (eg. an ordinary 3-d cube is formed by connecting two squares). This connectivity allows many communication topologies (grids etc.) to be mapped onto the nodes with neighbouring nodes in the topology being neighbouring nodes on the hypercube. In the present calculation we have used the simplest possible numerical model with the nodes being used as a processor farm controlled by a master process running on the SUN. Each node handles a subset of the star/gas particles which are scattered throughout the physical domain. The gravitational potential is calculated on a grid which is stored in its entirety on all the nodes. Each node calculates the contribution of its subset of particles to the potential, using a version of tree-code gravity due to Lars Hernquist (Barnes and Hut 1986, Hernquist 1987). Then the contribution of each node is communicated to the SUN which accumulates the total potential, and broadcasts it back to the nodes. The nodes then accelerate the

F. Combes and F. Casoli (eds.), Dynamics of Galaxies and Their Molecular Cloud Distributions, 353–355.
© 1991 *IAU. Printed in the Netherlands.*

https://doi.org/10.1017/S0074180900244501 Published online by Cambridge University Press

particles using the gradient of the potential interpolated to the position of each particle, update their positions, and the cycle begins again. For gas particles the Smooth Particle Hydrodynamics scheme is used (Lucy 1977, Gingold and Monaghan 1977), and the SPH forces are similarly extrapolated onto a grid, and broadcast via the SUN to interpolate the total SPH force at each gas particle.

The method by which the potential is communicated to the SUN is crucial. Simply sending the potential from each node directly to the SUN means that the communication process is serial in that the SUN can read from only one node at a time. However the parallel nature of the hypercube can be exploited by dividing the complete set of nodes in the calculation into its two daughter sub-cubes, and arranging for one of these subcubes to communicate its partial potential to the other, with each pair of nodes performing the communication step in parallel. The subcube with the partially accumulated potentials then divides into its two daughter subcubes, and performs the same process, so that by repeated subdivisions we end up with the total potential on one node which can then be communicated directly to the SUN. The computational expense of this procedure varies as the log n, n being the number of processors, rather than as n. It therefore varies as log N for N particles with a fixed number per node. The inverse procedure can be employed to re-broadcast the potential to the full set of nodes.

3) Results

Fig 1. shows some early results. The host has a spherical halo represented by a Toomre potential, with the mass of the halo being 80% of the host mass, and the star and gas masses being 10% each (the stars are initially distributed as an exponential disc, while the gas is distributed uniformly). The satellite mass has a mass of 10% of the host and a hyperbolic orbit with closest approach 1.5 times the gas disc radius, and with its orbital plane in the plane of the host so that the problem is essentially 2-d and we confine the particles to the plane. The results shown are for a 192,000 particle simulation (only 30,000 are shown) taking 17 secs per time step on a 50x50 mesh, i.e. 0.00009 secs per time step per particle, this can be compared with the 0.005 secs per time step per particle obtained on a CRAY XMP by Hernquist (1987) using full tree-code on a single processor.

I thank Phillip Fayers of the Physics Dept., UWCC, for invaluable assistance in programming the graphics and parallel communications code.

4) References

Barnes,J., & Hut,P. (1986), Nature, 324, 446.
Gingold, R.A., & Monaghan, J.J. (1977) MNRAS, 181, 375.
Hernquist,L. (1987) Ap.J.Suppl, 64, 715.
Lucy,L. (1977) Astron. J., 82, 1013.

https://doi.org/10.1017/S0074180900244501 Published online by Cambridge University Press

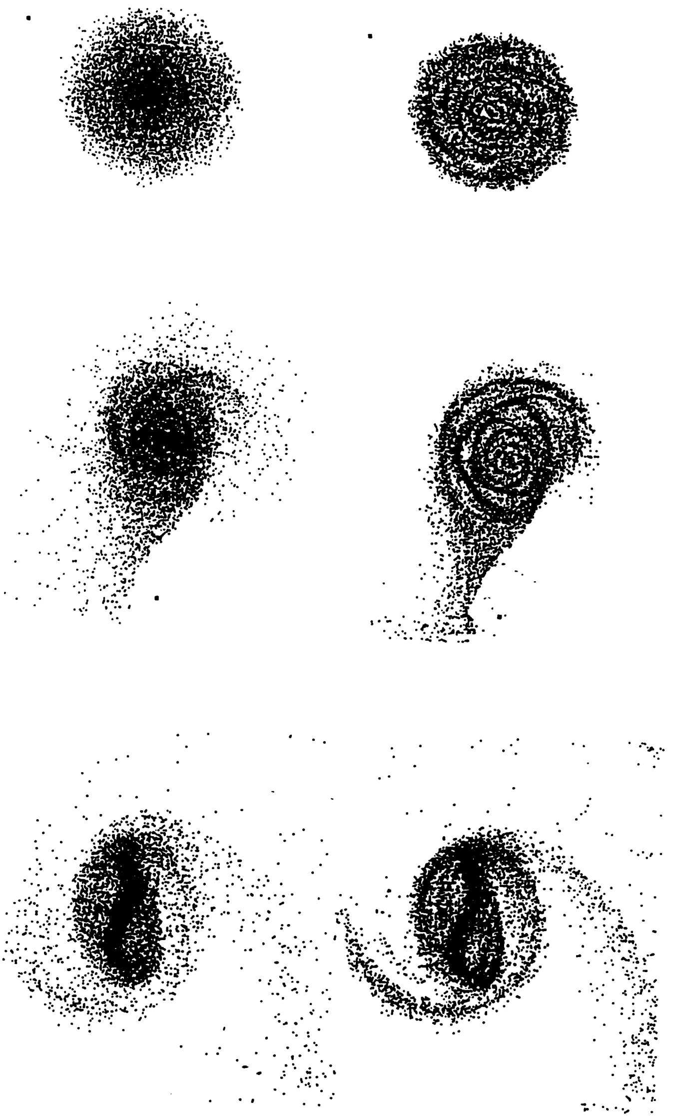

Fig.1 a) just before perigalacticon (stars on the right, gas on the left)
b) just after perigalacticon
c) approximately half a rotation period after perigalacticon.

https://doi.org/10.1017/S0074180900244501 Published online by Cambridge University Press