TOWARDS THE MILLION BODY PROBLEM ON THE COMPUTER – NO NEWS SINCE THE THREE-BODY-PROBLEM?

R. SPURZEM
Astronomisches Rechen-Institut, Mönchhofstraße 12-14, D-69120 Heidelberg, Germany
E-mail: spurzem@ari.uni-heidelberg.de

A. KUGEL
Dept. for Computer Science V, Univ. of Mannheim, B6-26, D-68131 Mannheim, Germany
E-mail: kugel@ti.uni-mannheim.de

Astrophysical Challenges which demand the solution of the one million (or more) gravitating body problem are briefly discussed for the fields of cosmology, galactic nuclei and globular star clusters. Results from the classical three-body problem are to be combined with methods of statistical mechanics or thermodynamics in order to provide a reasonable model for astrophysically relevant particle numbers of the order $10^4 - 10^8$. For high-precision modelling of star clusters and galactic nuclei over the required hundreds and thousands of crossing times specially tailored "brute-force" N-body integration methods are used together with special purpose hardware. A recent implementation of such code on a general purpose parallel computer is presented. After a comparison of the relevance of different N-body algorithms a new concept for a more flexible customized special purpose computer based on a combination of GRAPE and FPGA is proposed. It is an ideal machine for all kinds of N-body simulations using neighbour schemes, as the Ahmad-Cohen direct N-body codes and smoothed particle hydrodynamics (SPH) for systems including interstellar gas dynamics.

1 Introduction

Gravitation is the only one of the four fundamental forces in physics which cannot be shielded by particles of opposite charge and whose force at long range interactions does not have any (e.g. exponential) cutoff; the gravitational force is described by Newton's inverse square force law to the largest possible ranges in which classical mechanics applies. So it does not have a preferred scale and as a consequence gravitational forces play an important role for the dynamical evolution of astrophysical systems on practically all scales. Among the most challenging problems which have been treated by purely gravitational N-body simulations are structure formation in the universe, evolution of galactic nuclei and globular star clusters. Even if non-gravitative effects become important, like hydrodynamics or magnetohydrodynamic forces (e.g. for star formation or dissipative galaxy formation) some of the methods to solve the relevant dynamical equations represent the gas by particles interacting by a combination of gravitational and non-gravitative forces (smoothed particle hydrodynamics, SPH, [31,18]). Thus, simulations using particles to follow the dynamical evolution of astrophysical systems are one of the most important tools in computational astrophysics and have become a third independent experimental field of astrophysical research besides theory and observations.

This paper is organized as follows. First we present some key questions which are being addressed by present or future direct N-body simulations in the fields of globular star clusters, galactic nuclei, and cosmological models of structure formation (Sect. 2). Second, the main algorithms for astrophysical N-body simulations are briefly introduced and discussed in comparison with each other (Sect. 3). In Sect. 4 their implementation on special purpose hardware is discussed and a new concept for faster and more flexible hardware tailored to various kinds of direct N-body simulations including gas dynamics is presented.

2 Some Astrophysical Key Questions

2.1 Globular Star Clusters

There is an excellent review of the internal dynamics of globular star clusters [41]. Here we only want to stress some selected topics which are relevant to the subject of N-body simulations. Globular star clusters are nicely spherical (sometimes slightly flattened) star clusters consisting of 10^5 – 10^6 stars. Since their escape velocity is small compared to the typical velocities of stellar winds and explosions they are practically gas-free. They are ideal stellar dynamical laboratories because their relevant thermal (two–body relaxation) and dynamical timescales are smaller than their lifetime. Globular cluster systems exist around many other galaxies as well [12].

Due to their relative isolation, lack of observable interstellar gas and due to their symmetry globular clusters are well approximated by simplified theoretical models. Since the relaxation timescale is long compared to the dynamical time they develop through a sequence of dynamical (virial) equilibria. The fundamental kinetic equation in such case is the Fokker-Planck equation. The use of this equation for stellar dynamics was inspired by plasma physics [9]. Recent models of that type include the effects of anisotropy (differences between radial and tangential velocity dispersion, which can be present even in spherical systems) [68,69]. Another improvement includes for the first time the effect of rotation for those globular clusters which are slightly flattened [11]. Also anisotropic gaseous models based on a moment evaluation of the Fokker–Planck equation were successfully used [30,62].

In the presence of self-gravitation many concepts of thermodynamics, however, have to be used with care. So, for example, there is no global thermodynamic equilibrium, because a system of gravitating point masses can always achieve infinite amounts of binding energy just by moving two or more or all of the particles closer and closer together. At some limiting central velocity dispersion general relativity takes over and the cluster collapses by a dynamical instability towards a black hole [61]. Before reaching that limit most realistic astrophysical systems, however, will reach the limit of physical collisions and merging of their stars. Another alternative is that strong two–body correlations (close binaries) form which subsequently stop the global gravitational collapse by superelastic scatterings with field stars [7]. If there are too many binaries, however, the fundamental assumption of using the one-particle distribution function and the Fokker-Planck equation breaks down.

The second obstacle of thermodynamic methods to treat astrophysical ensembles of particles is simply that particle numbers are not large enough. Therefore stochastic fluctuations, deviations from thermodynamic expectation values in individual representations of e.g. star clusters are much more significant than in any laboratory gas; the amplitude of such fluctuations can be of a size comparable to that of the observed quantity. Hence it is by no means guaranteed, that a given individual globular cluster consisting of some one million or less particles strictly evolves according to models derived from statistical mechanics. In a seminal series of papers [14,15,16,17] the results of statistical models based on the Fokker-Planck approximation were compared with ensemble averages of a number of statistically independent direct N-body simulations [14,15,16,17]. From this and similar work [64,33] one can conclude that in spherical isolated clusters statistical models in spherical symmetry with standard two–body relaxation work fairly well, but already in the case of a galactic tidal field with an enhanced mass loss by stellar escapers, severe problems occur in understanding the results of the direct N-body models and their relation to the results of the Fokker-Planck results [4]. Note that such problems occur for one of the still most simplified globular cluster models; no rotation and no mass loss by stellar evolution was included, stars were considered as point masses and no effects of binary stellar evolution taken into account, no primordial binaries present, no time-dependent three-dimensional tidal field, and so on. Consequences from that are two–fold: first great care should be taken advancing Fokker-Planck and gaseous models to more complicated situations, second direct N-body models should be seen either as a theoretical tool to check and gauge the statistical models (eventually after a process of ensemble averaging) or they should employ a realistic particle number to directly model an individual, real star cluster. Besides the questions of gravothermal oscillations [7,33] in very large N-body systems and the scalability of cluster models in a galactic tidal field, we would like to stress here the importance to acquire information on the pre- and post-collapse evolution of N-body models of rotating globular clusters, for which very little is known yet, except a first series of Fokker-Planck type models [11].

Hence only an exact, direct N-body integrator should be used for problems of globular cluster and galactic nuclei dynamics, which treats the two-body relaxation by small angle gravitative encounters of all impact parameters with a maximal accuracy and simultaneously, accurately and efficiently follows the formation and evolution of very close binaries, whose timescales differ by many orders of magnitude from the dynamical timescale of the whole cluster. Such requirements are fulfilled by NBODY5 [2] and its successors (see Sect. 3). As a final remark for this subsection, we stress that in very young stellar clusters, like newly forming globular clusters seen around merging galaxies [59], the timescales to deplete the cluster from remaining protostellar gas by stellar winds, for star formation and evolution of massive stars, are comparable to the dynamical time of the cluster. Since mass segregation by two–body relaxation can be faster than the standard relaxation by a factor of M/m, where M are the most massive species, and m is the average particle mass in the cluster, even two-body effects are not completely negligible at cluster formation. Modelling such situation in a context of cooling and fragmentation of a gas cloud [48] including stellar dynamical effects would require an highly accurate N-body integrator in dynamical coupling with a gaseous component. It is interesting to note here, that the problem of star formation in molecular cloud complexes, which is another challenging problem of theoretical astrophysics, has been studied by a numerical model based on particles (SPH, smoothed particle hydrodynamics) and using special purpose computers (GRAPE, gravity pipe), on which we will elaborate below [26,25].

2.2 Galactic Nuclei

Another long–standing problem of collisional stellar dynamics is the question of the equilibrium system and dynamical evolution of a cusp of stars surrounding a massive central black hole. Such massive black holes are very likely to reside in the centres of galaxies as a fossile of earlier acticity [27]. Their formation as a result of collisionless dynamical general relativistic collapse and dissipative processes during galaxy formation is very likely but not yet fully understood [56]. In an earlier paper [13] the interplay between mass and energy transport by two–body relaxation and loss–cone accretion of stars on orbits with low angular momentum by the black hole was studied by semi-analytical scaling arguments; these results were confirmed by Monte-Carlo numerical models later [39] followed by multi–mass direct numerical solutions of the 1D Fokker–Planck equation for isotropic stellar cusps [47]. Only recently the first self-consistent N-body models of massive black holes including a sufficient number of stars in their surrounding cusps were done by use of hybrid N-body algorithms [54,55] or a high-speed special purpose computer for a direct summation algorithm [36,34]. However, the latter work was occupied mainly with the question of dynamical friction of black hole binaries in a galactic nucleus after a merger event. Still the standard picture of [13] has not yet carefully been checked by using a direct full N-body simulation. It is not certain, whether the assumption that two–body relaxation dominates the evolution is correct; it has been suggested that large angle close encounters of stars with each other and with the black hole compete with it, and that there may be non–standard relaxation processes at work [57]. These are interesting open question to tackle with high accuracy pure particle N-body simulations. Even more important as in the case of globular clusters, however, are the possible effects of gas produced by stellar collisions, which can accumulate in the centre due to the much deeper central potential, and form new stars [56,58]. We would like to conclude this subsection with the final remark that this is again a physical situation where highly accurate direct N-body models, eventually dynamically coupled with the dynamics of a gas component are very important for future understanding of such objects.

2.3 Cosmology and Structure Formation

In the standard paradigm of cosmological structure formation primordial quantum fluctuations grow gravitationally in a universe dominated by non-dissipative dark matter. In the non-linear regime the distribution of masses can be estimated by simple theory [53], later extended to N-body models [70,49]. On small scales gas physics, which (e.g. in the case of star formation) is only known approximately has to be included into the models [65]. Recently it has been shown, that softening of the gravitational potential, which was adopted in most of the models, causes spurious two–body relaxation effects [66]. Consistently [46] find that the structure of cold-dark-matter (CDM) haloes significantly changes if models with much higher resolution in particle number are used. Again, we want to conclude here that high resolution, high-accuracy N-body simulations, gravitationally coupled with a gas component are useful to study such questions.

3 Numerical Methods of N-body integration

Table 1. Algorithms for N-body Simulations

Number	Name	Scaling
NO PARTICLE-PARTICLE RELAXATION:		
1	PM – particle mesh	$\mathcal{O}(N) + \mathcal{O}(n^3)$
2.	Fast Multipole	$\mathcal{O}(N) + \mathcal{O}(nlm)$
3.	Self Consistent Field	$\mathcal{O}(N) + \mathcal{O}(nlm)$
"EXACT":		
4.	NBODY1 - 4	$\mathcal{O}(N^2)$
5.	NBODY5 - 6	$\mathcal{O}(NN_n) + \mathcal{O}(N^2)$
6.	KIRA	
"MIXED":		
7.	TREE	$\mathcal{O}(N \log N)$
8.	P^{3}M	$\mathcal{O}(N_n^2) + \mathcal{O}(N)\mathcal{O}(nlm)$

In Table 1 an overview over the most commonly used present algorithms for direct N-body simulations is given. The symbols used in the "Scaling" column denote the particle number N, a neighbour number N_n (compare Sect. 4), a grid resolution n or the number of terms nlm in a series evaluation of the gravitational potential. We want to comment only very briefly on each of the methods to give an overview for the reader. The first group has been labelled "no particle–particle relaxation" because it does not use direct gravitational forces between particles. The gravitational potential is computed from the particle configuration via an intermediate step, either through a mesh in coordinate space or an orthogonal function series. Reviews on classical particle mesh (PM) techniques can be found in [60]. "SUPERBOX" is a multi–grid method in a classical PM scheme suitable for high resolution problems and relaxing the inflexibility of conventional PM methods somehow [32]. Fast multipole methods used as presented by Greengard [19,20] can only efficiently be used for codes using the same timestep for particles, which makes them unfeasible for astrophysical problems with gravitating particles developing into highly structured and/or inhomogeneous states. Codes using an orthogonal series expansion (so-called "self consistent field" or SCF codes) have been introduced to the astrophysical community mainly by [22], although there are earlier similar approaches [8].

In all cases where a highly accurate computation of the gravitational potential with all its graininess due to individual particles, responsible for various relaxation effects, is necessary, there is no way to avoid a direct brute–force summation algorithm, where individual pairwise forces are computed. Such approach goes back to the early 60's [1,24]. Close encounters and the formation of binaries, whose binding energies are large compared to the thermal energy of the system have led to a special algorithm to treat hierarchical subsystems in such codes (named NBODY3 - 6 [3]) by a transformation to regularized variables [45,43,44] The more advanced code versions employ an individual or hierarchical block time step scheme and a high order time integrator [2,35], completed by an Ahmad-Cohen neighbour scheme [5] to reduce the number of full force computations (code versions NBODY5 - 6 [2,3]). The algorithm is well parallelizable and has been implemented on general purpose parallel computers [63]. Last but not least the KIRA code is mentined in Table 1, which is a fresh approach for a high accuracy direct N-body simulation code and has been used in studies of the dissolution of star clusters with stellar evolution in a tidal field [51,52]. At the end of Table 1 there are the "mixed" codes; one is the TREE-Code [6], using particle–particle forces in principle; it groups, however, subsets of particles in some distance from a test particle together, taking only their centres of masses into account (and if required also some multipole moments of the mass distribution). It is highly efficient for lumpy particle configurations, where the configuration has a small overall filling factor, and has been used very successfully for large-scale cosmological simulations and models of merging galaxies, partly even including a gas component treated by SPH [10,42]. Most TREE-Code implementations do not require very high accuracies, for example an energy error of up to a few percent is generally tolerated. Enforcing in a TREE–code very high accuracy as it is required for globular cluster models (10^{-3} %, a typical value achieved for direct Aarseth NBODY-integrators) leads to a very significant reduction of its efficiency [40]. Finally, another "mixed" code also used especially for cosmological simulations with an SPH gas component is the P^3M–code, for which we refer to a recent paper of [50].

4 Hardware

4.1 Introduction

The construction of special-purpose hardware to compute gravitational forces in direct N-body simulations was inspired by the fact that the total CPU time required for one time step of all particles scales as $T = \alpha N + \beta N^2$, with some numerical time constants α and β. For large particle number N the second part (the pairwise force calculation) consumes most of the time [67]. So our Japanese colleagues built the GRAPE hardware, whose development culminated in the presentation of the GRAPE-4 Teraflop computer [38,37]. The latter could be so fast, that according to Amdahl's law the other parts of the code (e.g. to advance the particles or compute neighbour related quantities) become the bottleneck of the simulation. This becomes even more serious if particle based methods are applied to solve the equations of hydrodynamics. The widely used SPH (smoothed particle hydrodynamics) represents the fluid by an ensemble of particles each carrying mass and momentum (as in the N-body problem). Thermodynamic observables are determined in a local averaging process over a given set of neighbouring particles. Hence the scaling properties of such codes, as well as the Ahmad-Cohen pure N-body codes using a neighbour scheme, can be modelled by a timing formula for the CPU time required per timestep as $T = \alpha N + \beta N^2 + \delta N \cdot N_n$, where δ is another time constant and N_n a typical neighbour number, for which neighbour forces (in case of NBODY6) or gas dynamical forces (in case of SPH) are to be calculated on a test particle.

However, the special purpose machines of the GRAPE series reach their highest efficiency only for problems, which can be tackled with pure and clean N-body algorithms such as NBODY4 or KIRA. For SPH or standard N-body simulations using an Ahmad-Cohen neighbour scheme or a very large number of close (so-called primordial) binaries, or even worse for molecular dynamics simulations with potentials other than the Coulomb potential (e.g. van der Waals) they are not the optimal choice. One commonly used solution is to use general purpose massively parallel machines as the CRAY T3E, for which a competitive implementation NBODY6++ exists using MPI and SHMEM [63]. While its performance compares well with one of the single GRAPE-4 boards, as can be seen in Fig. 1, a larger scale GRAPE machine or the coming GRAPE-6 are still much more efficient for the pure N-body case. There is still work in progress, however, to improve the implementation on the general purpose parallel computers.

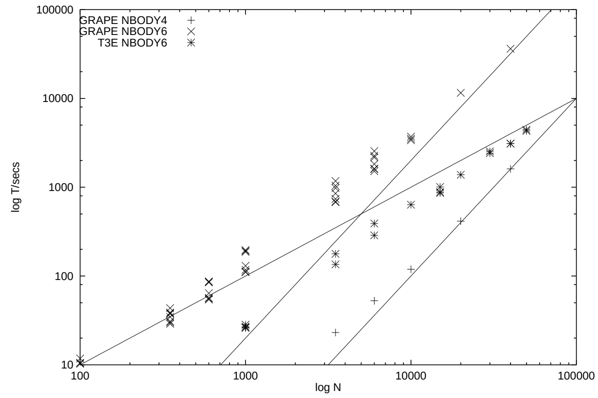

Figure 1. CPU time needed for one N-body time unit as a function of particle number N using the parallel version NBODY6++ – on special purpose computers GRAPE and CRAY T3E, and NBODY4 on GRAPE (keys for symbols given in the top left corner of the plot). For each N several data points are given for measurements with varying average neighbour number and processor/pipeline number, which are not individually discriminated in the figure. Power-Laws scaling with N and N^2 are given for comparison.

The new solution presented here is to build a hybrid machine, which uses for the intermediate range forces a reconfigurable custom computing machine – an FPGA processor. This new system will profit from both the extremely high performance of the GRAPE's for the $\mathcal{O}(N^2)$ gravitational force computation and the high degree of flexibility of the FPGA processor which lets it adapt to the needs of the various hydrodynamic (SPH) oriented computations in the $\mathcal{O}(NN_n)$ regime.

4.2 FPGA processors

The family of FPGA devices was introduced in 1984 by Xilinx. FPGA's feature a large number of relatively simple elements with configurable interconnects and an indefinite number of reconfiguration cycles with short configuration times. All configuration information is stored in SRAM cells. The basic processing element (PE) of all current mainstream FPGA's is a 4-input/1-output look-up-table (LUT) with an optional output register. The functionality of the FPGA is thus determined by the contents of the look-up-tables within the PEs and the wiring between these elements. Over the last few years FPGA performance has increased tremendously as it profits from both: Increased density by a factor of 24 from 1993 through 1998 (Xilinx XC4000: 400 to 18400 elements) increased speed by a factor of 3 from 1994 through 1998 (Xilinx XC4000: 133 to 400MHz internal toggle rate).

8 years of experience at the University of Mannheim with FPGA based computing machines shows that this new class of computers is an ideal concept for constructing special-purpose processors combining both the speed of a hardware and the flexibility of a software solution. The so–called FPGA processors consist of a matrix of FPGA's and memory forming the computational core. In addition there are a (programmable) I/O unit and an internal (configurable) bus system. As processing unit, I/O unit and bus system are implemented in separate modules, this kind of system provides scalability in computing power as well as I/O bandwidth.

FPGA processors have shown to provide superior performance in a broad range of fields, like encryption, DNA sequencing, image processing, rapid prototyping. The hybrid microprocessor/FPGA systems developed at the University of Mannheim [23] are in particular suitable for:

* acceleration of computing intensive pattern recognition tasks in High Energy Physics (HEP) and Heavy Ion Physics,
* subsystems for high-speed and high-frequency I/O in HEP,
* 2-dimensional industrial image processing,
* 3-dimensional medical image visualization and
* acceleration of multi-particle interaction calculations in astronomy

A commonly used procedure to adjust a hybrid system to different applications is modularity. ATLANTIS implements modularity on different levels. First of all there are the main entities host CPU and FPGA processor which allow to partition an application into modules tailored for either target. Next the architecture of the FPGA processor uses one board-type to implement mainly computing tasks and another board-type to implement mainly I/O oriented tasks. A backplane based interconnect system provides scalability and supports an arbitrary mix of the two board-types. Finally modularity is used on the sub-board level by allowing different memory types or different I/O interfaces per board type; it is an optimal machine to be applied to the astrophysical problems [29].

4.3 Architecture

Using FPGA's to accelerate complex computations using floating-point algorithms has not been considered a promising enterprise in the past few years. The reason is that general floating-point as well as particular N-body implementation have shown only poor performance on FPGA's. Usually N-body calculations and particle based fluid simulations need a computing performance in at least Teraflop range and are accelerated with the help of ASIC-based co-processors. Nonetheless we have recently investigated the performance of a certain sub-task of the SPH algorithm on the Enable++ system [28]. The results indicate that FPGA's can indeed provide even in this area a significant performance increase. Computation of the kernel function and the density determination of the SPH algorithm were implemented on 15 out of 16 core FPGA's of the Enable++ system making heavy use of the configurable interconnect structure. For the implementation a 28bit floating-point format was used: 1 sign-bit, 7 bits exponent, 20 bits mantissa. The maximum pipeline depth is 6 stages and a result is produced at every clock cycle. The total performance for the code in the loop is therefore $16 \cdot 13$ MHz = 208 Mflops with the XC4013-5 chips and $16 \cdot 32$ MHz = 512 Mflops with the XC4028-2 chip respectively. If the XC4036-3 implementation will allow – as we expect – that 2 instances can run in parallel, the performance will increase to 1.024 Mflops. Parallel I/O is also done with 52 or 128MB/s on the input side plus a few MB/s on the output side. ATLANTIS will support two instances of this code to run in parallel on one computing board.

4.4 AHA-GRAPE

For astrophysical particle simulations including self-gravity, the determination of the gravitational potential at each particles position is usually the most expensive step in terms of computational time required. This step shall be done by the special hardware GRAPE for force computation in N-body simulations, which proved highly efficient in the case of a pure point-mass simple algorithm (e.g. NBODY1) case. For many more realistic applications however, some parts of the code become important bottlenecks if the gravitational force calculation is done very fast. They are usually of order $\mathcal{O}(NN_n)$ – where N_n is a neighbour particle number — and comprise

1. Computation of the neighbour force when using SPH or a more complex, but more efficient N-body algorithm (about 20 flops per pairwise force, of which order N_n per particle per time-step have to be computed).

2. Computation of the kernel function in SPH, its derivatives, and terms related to gas dynamical quantities (about 100 flops per pairwise particle interaction, of which again $\mathcal{O}(N_n)$ per particle per time-step have to be computed).

3. Integration of binary motion in regularized coordinates as a function of near perturbers (order N_n); this is a very sophisticated algorithm and cannot easily be estimated now in its complexity [43,45].

4. Integration of SCF force (self-consistent field [22] to compute approximately the gravitational potential of distant particles) for hybrid N-body models .

The floating point operations related to 1) and 2) are in principle straightforward to map onto an FPGA processor, however critical for the performance is the word length which is sufficient for each component of the sum of pairwise forces. Test calculations have to be performed to clarify this. The two subsystems – GRAPE cluster and FPGA processor – will be connected to the host workstation by the PCI bus, either directly or via an interface. Within the subsystems the respective local buses will be used to broadcast sample data and intermediate results. In a second step and by close cooperation with the Tokyo group a hierarchical coupling of the FPGA device with GRAPE, including memory and control of the GRAPE, could be envisaged. This will further improve performance by parallelization of force computations and data communication. Fig. 2 displays the performance estimates for various systems with and without FPGA processor. The system with FPGA is called AHA-GRAPE (**A**daptive **H**ydrodyn**A**mics **GRAPE**), since it is ideally suited for high accuracy gravitational N-body simulations and hydrodynamical models using the SPH algorithm.

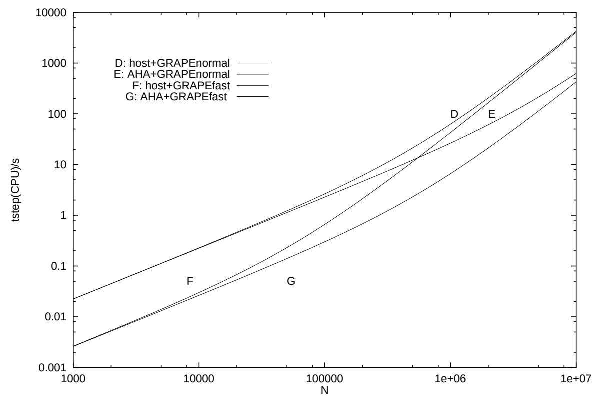

Figure 2. CPU time per step required for a simulation with direct gravitational force computation and neighbour scheme (SPH gas dynamics or Ahmad-Cohen N-body code) as a function of particle number for the proposed AHA-GRAPE machine (E, G) and a standard GRAPE-host combination (D, F), for a "normal" (D, E) and "fast" (F, G) GRAPE. Details see main text.

4.5 Status and plans

At present (February 1999) a test implementation of the SPH-loop/step1 on Enable++ is carried out to verify the estimated performance. By mid 1999 the new ATLANTIS system will be available where the full SPH-code has to be implemented. A communication library for LINUX must be developed, supporting simultaneous transfers between host/GRAPE and host/FPGA respectively. We expect the first prototype AHA-GRAPE system to be available in mid 2000. The key figures for this prototype are 50Mflops for the host workstation, 5Gflops for the FPGA processor and 500Gflops for the GRAPE subsystem, which have been used for the timing estimates depicted in Fig. 2. The presence of the FPGA processor will lead to an increase in performance by a factor of 10 and will allow us to handle up to $\approx 10^6$ particles in collision dominated N-body simulations and a few 10^7 particles in SPH.

Acknowledgments

It is a great pleasure to acknowledge in the name of the GRAPE user community in Germany the help and support received from D. Sugimoto, now continued by J. Makino and the members of his team, which originated from the time Professor Sugimoto visited as a Gauss Professor Göttingen observatory in 1983. Support by the DFG (German Science Foundation) under grants Sp 345/9-1,2 and Sonderforschungsbereich 439 (Galaxies in the Young Universe) is gratefully acknowledged. Computing time at NIC Jülich (former HLRZ), Höchstleistungsrechenzentrum Stuttgart (HLRS) and SSC Karlsruhe is gratefully acknowledged.

References

1. S.J. Aarseth, *Mon. Not. Royal Astron. Soc.* **126**, 223 (1963).
2. S.J. Aarseth in *Multiple time scales*, eds. J.U. Brackbill and B.I. Cohen (Acad. Press, Orlando, USA, 1985).
3. S.J. Aarseth in *Impact of Modern Dynamics in Astronomy*, IAU Coll. No. 172 (Kluwer Acad. Publ., Dordrecht, The Netherlands, 1999).
4. S.J. Aarseth and D.C. Heggie, *Mon. Not. Royal Astron. Soc.* **297**, 794 (1998).
5. A. Ahmad and L. Cohen, *Journ. Comp. Phys.* **12**, 389 (1973).
6. J. Barnes and P. Hut, *Nature* **324**, 446 (1986).
7. E. Bettwieser and D. Sugimoto, *Mon. Not. Royal Astron. Soc.* **208**, 493 (1984).
8. M. Clutton-Brock, *Astrophys. Space Science* **16**, 101 (1972).
9. H.N. Cohn, *Astroph. Journal* **242**, 765 (1980).
10. R. Davé, J. Dubinski and L. Hernquist, *New Astron.* **2**, 277 (1997).
11. C. Einsel and R. Spurzem, *Mon. Not. Royal Astron. Soc.* **302**, 81 (1999).
12. D.A. Forbes, J.P. Brodie and C.J. Grillmair, *Astron. Journal* **113**, 1652 (1997).
13. J. Frank and M.J. Rees, *Mon. Not. Royal Astron. Soc.* **176**, 633 (1976).
14. M. Giersz and D.C. Heggie, *Mon. Not. Royal Astron. Soc.* **268**, 257 (1994).
15. M. Giersz and D.C. Heggie, *Mon. Not. Royal Astron. Soc.* **270**, 298 (1994).
16. M. Giersz and D.C. Heggie, *Mon. Not. Royal Astron. Soc.* **279**, 1037 (1996).
17. M. Giersz and D.C. Heggie, *Mon. Not. Royal Astron. Soc.* **286**, 709 (1997).
18. R.A. Gingold and J.J. Monaghan, *Mon. Not. Royal Astron. Soc.* **181**, 375 (1977).
19. L. Greengard, *Computers in Phys.* **4**, 142 (1990).
20. L. Greengard and V. Rokhlin, *Journ. Comp. Phys.* **73**, 325 (1987).
21. D.C. Heggie *et al* in *Highlights of Astronomy Vol. 11*, ed. J. Andersen (Kluwer Acad. Publ., Dordrecht, The Netherlands, 1999). Highlights of Astronomy Vol. 11, Kluwer Acad. Publishers, in press
22. L. Hernquist and J.P. Ostriker, *Astroph. Journal* **386**, 375 (1992).
23. H. Högl *et al* in *Proc. 3rd IEEE Symp. on FPGA's for Custom Computing Machines* (Napa, CA, USA, 1995).
24. S. von Hoerner, *Z. f. Astroph.* **50**, 184 (1960).
25. R. Klessen, *Mon. Not. Royal Astron. Soc.* **292**, 11 (1997).
26. R. Klessen, A. Burkert and M. Bate, *Astroph. Journal Lett.* **501**, 205 (1998).

27. J. Kormendy and D. Richstone, *Ann. Rev. Astron. Astroph.* **33**, 581 (1995).
28. T. Kuberka, *Leistungsanalyse von Gleitkommaarithmetik auf Prozessoren an Hand der Implementierung des SPH-Algorithmus (Schritt 1)*, Diploma Thesis Univ. of Mannheim, November 1998 (in German).
29. T. Kuberka *et al*, subm. to *International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA)* (1999).
30. P.D. Louis and R. Spurzem, *Mon. Not. Royal Astron. Soc.* **244**, 478 (1991).
31. L. Lucy, *Astron. Journal* **82**, 1013 (1977).
32. R. Madejsky and R. Bien, *Astron. Astroph.* **280**, 383 (1993).
33. J. Makino, *Astroph. Journal* **471**, 796 (1996).
34. J. Makino, *Astroph. Journal* **478**, 58 (1997).
35. J. Makino and S.J. Aarseth, *Proc. Astron. Soc. Japan* **44**, 141 (1992).
36. J. Makino and T. Ebisuzaki, *Astroph. Journal* **465**, 527 (1996).
37. J. Makino and M. Taiji, *Scientific Simulations with Special-Purpose Computers* (J. Wiley & Sons, Chichester, New York, Weinheim, 1998).
38. J. Makino *et al*, *Astroph. Journal* **480**, 432 (1997).
39. A.B. Marchant and S.L. Shapiro, *Astroph. Journal* **239**, 685 (1980).
40. S.L.W. McMillan and S.J. Aarseth, *Astroph. Journal* **414**, 200 (1993).
41. G. Meylan and D.C. Heggie, *Astron. Astroph. Rev.* **8**, 1 (1997).
42. J.C. Mihos, J. Dubinski and L. Hernquist, *Astroph. Journal* **494**, 183 (1998).
43. S. Mikkola, *Celest. Mech. Dynam. Astron.* **68**, 87 (1997).
44. S. Mikkola and S.J. Aarseth, *Celest. Mech. Dynam. Astron.* **64**, 197 (1996)
45. S. Mikkola and S.J. Aarseth, *New Astron.* **3**, 309 (1998).
46. B. Moore *et al*, *Astroph. Journal Lett.* **499**, 5 (1998).
47. B.W. Murphy, H.N. Cohn and R.H. Durisen, *Astroph. Journal* **370**, 60 (1991).
48. S.D. Murray and D.N.C. Lin, *Astroph. Journal* **467**, 728 (1996).
49. J.F. Navarro, C.S. Frenk and S.D.M. White, *Astroph. Journal* **490**, 493 (1997).
50. F.R. Pearce and H.M.P. Couchman, *New Astron.* **2**, 411 (1997).
51. S.F. Portegies Zwart *et al*, *Astron. Astroph.* **337**, 363 (1998).
52. S.F. Portegies Zwart and K. Takahashi, *Astroph. Journal Lett.* **503**, 49 (1998).
53. W.H. Press and P.L. Schechter, *Astroph. Journal* **187**, 425 (1974).
54. G.D. Quinlan, *New Astron.* **1**, 35 (1996).
55. G.D. Quinlan and L. Hernquist, *New Astron.* **2**, 533 (1997).
56. G.D. Quinlan and S. Shapiro, *Astroph. Journal* **356**, 483 (1990).
57. K.P. Rauch and S. Tremaine *New Astron.* **1**, 149 (1996).
58. M.J. Rees, *Rev. Mod. Astroph.* **10**, 179 (1997).
59. F. Schweizer *et al*, *Astron. Journal* **112**, 1839 (1996).
60. J.A. Sellwood, *Ann. Rev. Astron. Astroph.* **25**, 151 (1987).
61. S. Shapiro and S.A. Teukolsky, *Astroph. Journal Lett.* **292**, L41 (1985).
62. R. Spurzem in *Ergodic Concepts in Stellar Dynamics*, eds. D. Pfenniger and V.G. Gurzadyan (Springer, Berlin, Heidelberg, 1994).
63. R. Spurzem and H. Baumgardt, subm. to *Mon. Not. Royal Astron. Soc.*, 1999.
64. R. Spurzem and S.J. Aarseth, *Mon. Not. Royal Astron. Soc.* **282**, 19 (1996).
65. M. Steinmetz *Mon. Not. Royal Astron. Soc.* **278**, 1005 (1996).
66. M. Steinmetz and S.D.M. White, *Mon. Not. Royal Astron. Soc.* **288**, 545

(1997).
67. D. Sugimoto *et al*, *Nature* **345**, 33 (1990).
68. K. Takahashi, PASJ **48**, 691 (1996).
69. K. Takahashi, PASJ **49**, 547 (1997).
70. S.D.M. White *et al*, *Astroph. Journal* **313**, 505 (1987).