JOURNAL OF COMPUTATIONAL PHYSICS 6, 449–472 (1970)

Gravitational N-Body Calculation in a Discrete Phase Space

RICHARD H. MILLER

Department of Astronomy, Committee on Information Sciences, and The Institute for Computer Research, University of Chicago, Chicago, Illinois 60637

Received March 23, 1970

An *n*-body calculation designed for astronomical studies in stellar dynamics is able to handle large number of particles by restricting the set of points over which force values are required. The resulting model can be treated exactly in a computer by using integer arithmetic; the "data" representing the physical system are modified to conform to the model. Reversibility and an exact Liouville theorem result, making the model close to the physics in the sense that exactly conserved quantities of the model correspond to essential features of the physics. The cost of this pleasing correspondence is a reduction in the accuracy of conserving some of the conventional integrals. While the formulation affords a useful viewpoint for considering the relationship between computer models and the physical systems they represent that is quite different from that underlying conventional calculations, numerical considerations differ only in that the numerical methods used are coarser than with conventional calculations. The numerical approximations and the features responsible for the exact properties are presented in detail. The calculation has been successfully used as a tool for numerical experiments on spiral structure and gravitational stability.

1. INTRODUCTION

A gravitational *n*-body program that permits calculations with large numbers of bodies (125,000 typically) has been running for some three years [1, 2, 3, 4]. Although others have run similar numbers of particles for gravitational [5, 6] and plasma calculations [7] our calculation has some novel features that we feel may have wider application. The unusual features principally arise in the point of view concerning the relations between the underlying physics and the computer model. The physics is expressed in certain invariance properties which are exactly incorporated into the model. An exactly calculable model results, a model that is pleasing to work with because of the feeling of precision that it affords. The data have been modified to fit the model.

The transition from the physics to the model is the principal subject of this paper. Not all of the physical invariance properties can be built into a calculation; a matter of taste is involved in the selection of those deemed most important. We

feel that the model is close to the physics because it incorporates the physically essential features exactly. The point of view might be expressed by stating that the physics does not live in the differential equations—they are at best one way of describing what is going on.

The motivation for seeking models with exact properties came from a demonstration of the difficulty of making satisfactory numerical investigations of gravitational n-body problems (Ref. [8]; the results are also briefly described in Ref. [9]). The numerical difficulties manifested themselves in the irreversibility of the calculation, which developed much more rapidly than had been expected. Some quantities can be computed adequately [10–14] but Lecar's [15] comparative studies show how difficult it is to draw detailed conclusions from gravitational n-body calculations, even apart from problems of interpretation like those expressed by Zwanzig and Ailawadi [16] and by Lebowitz et al. [17].

The invariance properties are incorporated through a simple finite-difference scheme that can be made to represent a discrete phase space by starting it with integer values and thereafter giving it only integer values to work with. The size of the interval between allowed discrete locations is equivalent to the roundoff in a conventional calculation; a matter of some concern is precisely what effects result from intentionally making roundoff and truncation error as large as seems reasonable to be in keeping with the spirit of the problem.

The problem in a discrete phase space may in turn be approached from either of two viewpoints. In the first, which is adopted here, the discrete representation may be regarded as a model approximating the dynamics of continuous systems. The model is capable of exact representation in the computer. All the approximations are made between the physics and the model. The advantage is that the approximations occur at a well-defined stage of the argument, and once made, leave an exactly calculable system. The investigations of numerical properties then describe the extent to which the model does or does not describe a physical situation. The other viewpoint is to argue, from the outset, that the mechanics appropriate to a discrete system is being developed, with less concern for its possible applicability in the real world. The mechanics of a discrete representation is an interesting problem which will not be developed here; Greenspan [18] seems to have made some progress along these lines.

Our calculation is reminiscent of some aspects of the "PIC" code in hydrodynamics [19]. The PIC code is quite different in spirit, as it does not rely on nearly so literal an interpretation of the notion of ascribing properties to cells or to discrete locations. Some properties are freed from the lattice, and all are regridded from time to time. None of these things happen in our calculation.

The model is introduced as a "game" in Section 2 (not in the sense of game theory); details of the model and of the relationship to physical systems appear in later sections.

2. A Game

A point is moved over a two-dimensional lattice according to these rules:

1. It moves between locations whose coordinates (x, u) are integers.

2. It may move from one allowed location to another, but must always alternate a move in which only x changes with one in which only u changes. Let the value of x immediately following the n-th step be $x^{(n)}$, the value of u following its next change after $x^{(n)}$ was reached be $u^{(n+1/2)}$.

3. The value of $x^{(n+1)}$ is given by

$$x^{(n+1)} = x^{(n)} + u^{(n+1/2)}. \tag{1a}$$

4. There is a table that gives the rule for changing u according to the present value of x. Let the value read out of this table following the n-th step be $f^{(n)}$. Then the value of $u^{(n+1/2)}$ is given by

$$u^{(n+1/2)} = u^{(n-1/2)} + f^{(n)}. \tag{1b}$$

The pair of moves is a complete step.

As an example (Fig. 1), the particle might start from location "a" with $x^{(0)} = +2$, $u^{(-1/2)} = 0$; the table of f's gives $f^{(0)} = -2$, so $u^{(+1/2)} = -2$, as shown at "b" in Fig. 1. The shift in x next takes the point to "c" of Fig. 1, and so on until the point has returned to the location from which it started. Thereafter, it would endlessly repeat the same "orbit."

Suppose there were many points in such a diagram, each moving about the lattice according to the stated rules, without regard for the other points. Then each point on a given row of $u = \text{const}$ would be moved the same distance to the right by the value of u at integer values of n as shown in Fig. 2a. The subsequent moves at half-odd integer values of n would be upward along a column of fixed x by the amount specified by f (Fig. 2b).

Whatever the table of f values (even if the table is changed at each n), (1) The contents of a location are transferred to another location as a unit. Thus, the contents of two locations cannot come to occupy the same location at a later time, nor can the contents of one location split to occupy two locations. (2) If the process were run backward after n steps (properly reversing the sequence of operations), all points would return to their original locations.

The game just described bears several suggestive resemblances to dynamics. The lattice is like a discrete phase space (distinctions about evaluating x's and u's at different times aside) for a one-dimensional problem if the identifications $x \leftrightarrow$ coordinate, $u \leftrightarrow$ velocity, $f \leftrightarrow$ force per unit mass, $n \leftrightarrow$ time are made. The

property that occupants of a lattice location stay together is the Liouville theorem (measure-preserving flow). The game describes the reversible flow of an incompressible fluid in the (discrete) phase space. Since the collision-free Boltzmann equation describes this kind of motion in phase space, the game approximates to solutions of that equation.

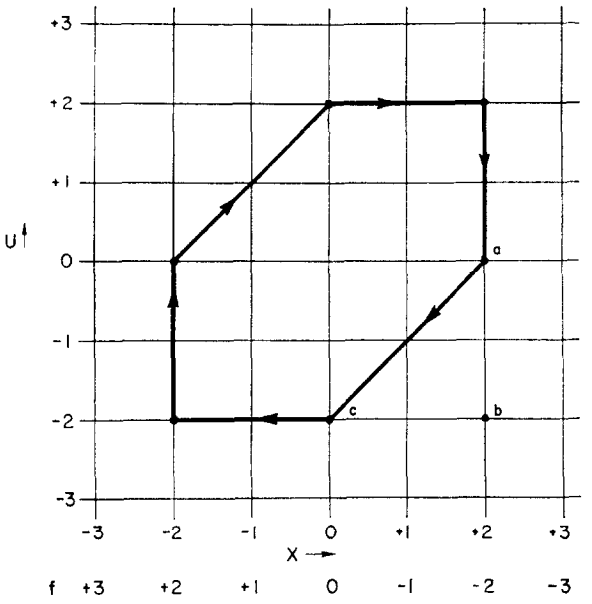

FIG. 1. Moves in an example of the game.

Equation (1) is a finite-difference scheme for integrating $\dot{x} = u$; $\dot{u} = f$, if the value of the time step be incorporated into the definitions of u and of f to make the entire system dimensionally compatible. The game is like any other technique for numerical integration in that the integration is truncated at a certain order. Here it is truncated at an unusually low order.

The game contains the physical features of the phase-space description exactly. The game carries out an approximate integration without doing arithmetic. It clearly can be generalized to more dimensions. We have used it both in a bit-shifting mode as suggested by Fig. 2, and in an arithmetic mode as suggested by Eq. (1). For stellar dynamics, the arithmetic mode utilizes storage space more effectively.

Because the Liouville theorem and reversibility summarize the essential physics of stellar dynamics (in the relaxation-free limit), this scheme is close to the basic physics of the problem. There is a large number of integrals of the motion; the

occupancy of any phase location, as the motion maps the phase space into itself, is an integral of the motion. The motion is also exactly reversible (apart from questions as to how the force values $f^{(n)}$ turned out to be integers; this is discussed later). A quantity that looks like the conventional energy integral can be formulated (cf. Ref. [18]), but it differs in important respects from the usual energy integral. To be more precise, it does not seem to be possible to derive the forces from a scalar potential in an exact way (Section 8).

Throughout this paper, the present calculation is contrasted with conventional gravitational n-body calculations. The term "conventional," as applied to gravitational n-body calculations, means a calculation which, in spirit, attempts to treat everything exactly to the precision allowed by the computer. The forces are usually obtained by explicit summation of contributions due to each pair of

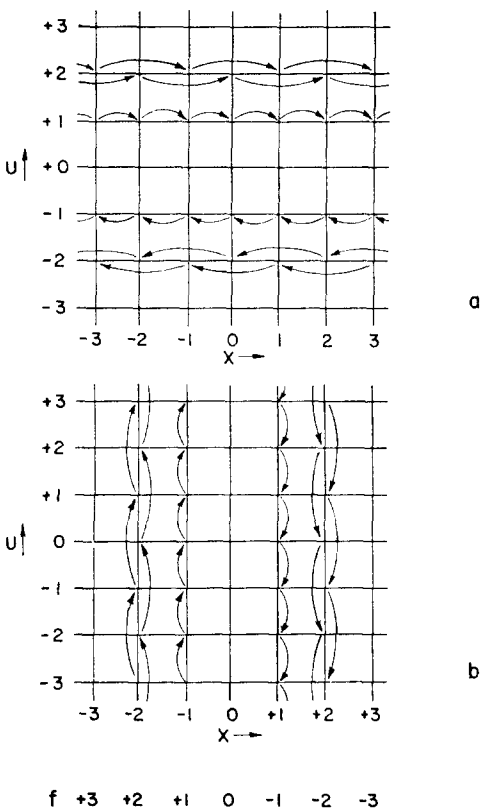

FIG. 2. Systematic moves of many particles in the game. (a) The moves in which x is changed according to the present value of u; (b) the moves in which u is changed by the amount of $f(x)$.

particles. The calculations reported in [10–15] are conventional in this sense. As a practical matter, it has not proved possible to handle more than about 250 particles in such calculations.

The features setting the game apart from conventional gravitational n-body calculations are quantitative and not qualitative. Other representations are discrete, although with much finer steps than we use. Other integration schemes are finite-difference, although usually the ignored terms are of higher order than in this calculation. The conventional calculations have at least one exact integral (shared by this calculation): The number of particles. Thus, the working out of numerical details for this calculation indicates how they might similarly be worked out for other calculations. The distinction is a matter of degree, not of kind. There are certainly some features that are handled much better in the conventional calculations than in this one. Angular momentum conservation is such a feature; in our calculation, angular momentum is expected to be conserved only in a statistical sense. But these disadvantages are more than offset by the ability to handle large numbers of particles.

Another representation of an n-body system is useful for certain fundamental properties. The Γ space ($6n$-dimensional for three-dimensional problems) might also be discrete. The $6n$-dimensional lattice can be mapped onto a one-dimensional array in which the representation of the system is now a vector with all elements zero except for one that describes the n-body system. After a time step, the system can again be represented by such a vector. Since the motion in the Γ space is completely determined (classical system), the transformation of one of these vectors into the other may be represented by a square matrix that contains just one nonzero element in each row and column. The matrix does not depend on the system configuration. The result of two time steps must be capable of representation by a transformation matrix having the same properties. Since the matrix is finite (for bounded motion), it is an element of a cyclic group. The order of the group is the number of time steps in a Poincaré recurrence time. This representation makes it intuitively clear that the game is stable as an initial-value problem.

3. The Finite-Difference Scheme

A. *A Norm for the State Vectors*

Discussions of the stability of the finite-difference scheme require a norm that does not change as the step size is altered. This seems to be most conveniently done in the following way. Define a (dimensionless) state vector

$$\xi^{(n)} = \begin{pmatrix} x^{(n)} \\ u^{(n-1/2)} \end{pmatrix}, \tag{2}$$

which, according to the difference-scheme (1) obeys

$$\begin{aligned}\xi^{(n)} &= A\xi^{(n-1)} + Bf^{(n-1)} = \begin{pmatrix}1 & 1\\ 0 & 1\end{pmatrix}\xi^{(n-1)} + \begin{pmatrix}1\\ 1\end{pmatrix}f^{(n-1)} \\ &= A^n\xi^{(0)} + \sum_{r=1}^{n} A^{r-1}Bf^{(n-r)},\end{aligned} \tag{3}$$

the usual iterated discrete-time relations. Although the system may easily be generalized to unequal time steps, that is not necessary for this calculation. The integer $x^{(n)}$ represents a physical length $X^{(n)} = x^{(n)}\Delta l$, where Δl is the (dimensioned) spacing of the allowed discrete locations. This, in turn, may be related to a dimensionless number representing displacements in units of some (dimensioned) normalization length L,

$$\mathbf{X}^{(n)} = \frac{X^{(n)}}{L} = x^{(n)}\frac{\Delta l}{L}. \tag{4}$$

The same thing can be done with the velocities and the forces by introducing the time step Δt and a normalization time T,

$$\begin{aligned}\mathbf{U}^{(n-1/2)} &= u^{(n-1/2)}\frac{\Delta l}{L}\frac{T}{\Delta t} \\ \mathbf{F}^{(n)} &= f^{(n)}\frac{\Delta l}{L}\left(\frac{T}{\Delta t}\right)^2.\end{aligned} \tag{5}$$

Then a new state vector can be defined,

$$\Xi^{(n)} = \begin{pmatrix}\mathbf{X}^{(n)}\\ \mathbf{U}^{(n-1/2)}\end{pmatrix} = N\xi^{(n)} = \frac{\Delta l}{L}\begin{pmatrix}1 & 0\\ 0 & T/\Delta t\end{pmatrix}\xi^{(n)}, \tag{6}$$

which can be used to define a norm,

$$\| \xi^{(n)} \|^2 = (\Xi^{(n)})^T(\Xi^{(n)}). \tag{7}$$

The superscript T denotes transpose. The independence of stepsize works as follows: Suppose Δl is decreased. Then x must increase to represent the same physical length (requiring both the old and the new x to be integers places some constraints on the way Δl is reduced; these only enter the discussion later). But the values of the elements of Ξ are not changed, so the norm is unchanged.

The separation of $f^{(n)}$ into an inhomogeneous term allows the particle interactions to be introduced in a way that permits the large matrix that advances the positions and velocities of all particles to be block-decomposed in the manner of Eq. (3).

B. *Stability of the Finite-Difference Scheme*

Stability is usually defined as the uniform boundedness of the difference between the computed solution (represented by $\xi^{(n)}$) and the exact solution to the system of differential equations, $v(n\,\Delta t)$, starting from the same initial condition ($\xi^{(0)} = v(0)$), and requiring that the two kinds of solution be compared at corresponding times. In this sense, a necessary and sufficient condition for stability (see, for example Wendorff [20]) is that the norm of $\| A \| \leqslant 1 + C\,\Delta t + O(\Delta t^2)$. In fact, for this case,

$$\| A \| = \sup \frac{\| A\xi \|}{\| \xi \|} = 1 + \frac{1}{2}\left(\frac{\Delta t}{T}\right) + O\left(\frac{\Delta t}{T}\right)^2. \tag{8}$$

An interesting feature of this calculation is that $(\Delta l/L)$ cancels out in this definition of $\| A \|$ through $\| \xi \|$. This is to be expected because uniformity requires that noninteger x's and u's be used, and there is nothing about the difference scheme, per se, that requires the x's and u's to be integers (or rational). The discrete phase space notion is tacked on to the difference scheme, and is not essential to it.

Intermediate degrees of reduction of time step can also be investigated. Let the normalization length L be the spatial lattice interval of the original calculation. Then let $m\,\Delta l = L$, so that increasing m refines the space lattice. Similarly, let $q\,\Delta t = T$, and let the calculation run for τ time steps of the initial calculation, so $n = q\tau$. After τ time steps of the original calculation, A^n is involved, giving

$$\| A^n \|^2 \to 1 + \frac{\tau^2}{2} + \frac{\tau}{2}\sqrt{4 + \tau^2} \leqslant 1 + \tau + \tau^2. \tag{8a}$$

The m and q dependence has dropped out, but the norm stays bounded for finite τ. This argument does not make use of any continuity properties of m, q, Δl, or Δt.

4. The Discrete Phase Space

The nice features of this calculation, such as reversibility and the special kind of Liouville theorem (particles initially together remain together), derive from the discrete phase space. That, in turn, comes from the property of the difference scheme (1) that the x's and u's will remain integers if the f's are all integers. It now remains to investigate how this modification affects the stability and accuracy of the calculation. It may be stressed again that this is a modification of the data to force it into the mold of a given model, such that the model can be treated exactly.

There is nothing unusual about selecting an initial condition such that all the

x's and u's happen to be integers. The effect of the calculation with integers appears through the inhomogeneous term of Eq. (3), the term involving $f^{(n)}$.

Let $\xi^{(n)}$ be the vector representing the system with integer forces $f^{(n)}$; let $\zeta^{(n)}$ be a corresponding vector representing a system with "exact" forces, $\psi^{(n)}$. Take the two initially equal, $\xi^{(0)} = \zeta^{(0)}$. Then the norm of the difference vector describes the additional effect of approximations in the force values,

$$D^2 = \| \xi^{(n)} - \zeta^{(n)} \|^2 = \sum_{r=1}^{n} \sum_{s=1}^{n} (NA^{r-1}B)^T (NA^{s-1}B)\, \varphi^{(n-s)}\varphi^{(n-r)}$$
$$= \left(\frac{\Delta l}{L}\right)^2 \sum_{r,s=1}^{n} \left[rs + \left(\frac{T}{\Delta t}\right)^2\right] \varphi^{(n-s)}\varphi^{(n-r)}, \tag{9}$$

where $\varphi^{(k)} = f^{(k)} - \psi^{(k)}$ is the difference between the approximate and the "exact" forces at step k.

The series (9) can be summed under a variety of circumstances. Only two extreme cases will be given, with indications of the important features in more realistic situations.

Case (a). Completely uncorrelated, $\varphi^{(j)}\varphi^{(k)} = \langle\varphi^2\rangle\, \delta_{jk}$. Then

$$D_A{}^2 = \left(\frac{\Delta l}{L}\right)^2 \langle\varphi^2\rangle \left[\frac{n(n+1)(2n+1)}{6} + n\left(\frac{T}{\Delta t}\right)^2\right]$$
$$\approx \left(\frac{\Delta l}{L}\right)^2 \langle\varphi^2\rangle \left[\frac{n^3}{3} + n\left(\frac{T}{\Delta t}\right)^2\right]. \tag{10a}$$

Case (b). Completely correlated, $\varphi^{(j)}\varphi^{(k)} = \langle\varphi^2\rangle$ independently of j, k. Then

$$D_B{}^2 = \left(\frac{\Delta l}{L}\right)^2 \langle\varphi^2\rangle \left[\left(\frac{n(n+1)}{2}\right)^2 + n^2\left(\frac{T}{\Delta t}\right)^2\right]$$
$$\approx \left(\frac{\Delta l}{L}\right)^2 \langle\varphi^2\rangle \left[\frac{n^4}{4} + n^2\left(\frac{T}{\Delta t}\right)^2\right]. \tag{10b}$$

This discussion is independent of the cause of the variation in the forces; in particular it describes the usual roundoff in computations as well as our more drastic roundoff to produce integer force values. Case (a) describes situations in which there is no memory in the roundoff from one evaluation of the forces to the next. Most roundoff situations are *assumed* to be of this kind. A more realistic description results from admitting that rounding in one direction on one integration step is likely to be followed by another rounding in the same direction on the next integration step. Although Case (b) is much more stringent than that situation, it

contains the essential features. The fourth power of n is typical of those cases with "memory" even if the "memory" lasts over a small fraction of the total duration of the calculation.

The approximate expressions in (10a) and (10b) can be used to investigate the effect of refining the spatial grid and the time step. The two cases yield

$$D_A^2 \approx \frac{q^3}{m^2}\left[\frac{\tau^3}{3} + \tau\right]\langle\varphi^2\rangle \tag{11a}$$

and

$$D_B^2 \approx \frac{q^4}{m^2}\left[\frac{\tau^4}{4} + \tau^2\right]\langle\varphi^2\rangle. \tag{11b}$$

With the correlation "memory" extending over shorter times, the dependence of D^2 is very similar to (11b). For a "correlation" that dies out linearly with time, going to zero at τ_c,

$$D_C^2 \approx \frac{q^4}{m^2}\tau_c\left[\frac{\tau^3}{3} + \tau\right]\langle\varphi^2\rangle. \tag{11c}$$

The q^4/m^2 dependence remains, as might be expected. In this expression, the correlation lasts for the same time interval independently of the number of integration steps required to cover that time interval. The partial memory is a reasonably good picture of the rounding of force values in an actual calculation. Particles that don't move much from one integration step to the next, in force fields that don't change much from one integration step to the next, are likely to round in the same direction in successive integration steps. But after a while, the particle should have moved into a region where the rounding is substantially independent of the earlier history. In a system containing many "particles" with quite different motions, no one "correlation" history can describe the situation, so some kind of "average particle" or even a "worst-case particle" is described by τ_c. This matter of rounding the force values to integers is essentially the only point at which our calculation differs in spirit from other n-body calculations. The correlation term arising in our calculation is matched by similar terms in more conventional calculations. The difference is one of degree.

An interesting aspect of D^2 in (11b) and (11c) is that it is natural to consider alterations in which the time step and spatial interval are changed keeping q^2/m constant; this keeps the forces constant and is the simplest change that does not affect the mass of the "particles." Under such a change, a build-up of D^2 with "uncorrelated" force roundoff terms would go to zero in the limit $q \to \infty$, but that contribution to D^2 due to the "correlated" force roundoff is not affected. The limiting process must be carried out in such a way that the spatial lattice is

refined faster than the square of the refinement of the time step to make D^2 go to zero.

The "correlated" contribution to D^2 can be avoided by using rounding rules that eliminate the correlation between successive values of the force. The force values can be rounded by a rule like ENTIER (F + RANDOM), where RANDOM produces a "pseudorandom number" in the interval (0, 1), rather than by the usual ENTIER (F + 1/2), for example. This makes the correlation go to zero at the cost of doubled variance. But it can also lose the reversibility of the calculation.

The uncertainty introduced through the roundoff can be quite large. Taking as an example $\langle \varphi^2 \rangle = 1/12$ from the usual rounding, with $q = m = 1$ and $\tau = 10$, (11a) gives $\sqrt{D^2} \approx 6$ spatial units, and (11b) gives ≈ 16. In conventional calculations where the lattice spacing is much finer, these terms are less harmful.

5. Reversibility

Much has been written on the reversibility of machine calculations (see, e.g., Buneman [21]). Some of the discussions may leave the impression that finding a suitable reversible algorithm is the main problem. Reversible algorithms may be generated easily by deriving them from variational principles, but the resulting algorithms may be implicit or otherwise ill-suited to machine calculations. Even so, a reversible algorithm may not run reversibly if a different numerical value is computed at some point of the calculation as the system passes in the forward and reverse directions. It is instructive to see how the performance of the difference scheme (1) matches expectations. The formalism that is developed in the course of doing this clarifies the essential features of reversibility of machine calculations.

Reversibility can be discussed in much the same way as the growth of roundoff errors. The same process that causes a calculation to be irreversible could cause two separate ("parallel") calculations to differ if they started from identical initial conditions and ran forward together. Such "parallel" calculations were used to study reversibility in n-body calculations earlier [8].

Two calculations, using the difference scheme (1), with integer values of x, u, and f, can differ only if the value of f at some part of the system has been rounded differently in the two cases. Let the two systems be described at the n-th step by vectors $\xi^{(n)}$, $\eta^{(n)}$, with $\xi^{(0)} = \eta^{(0)}$. The forces governing the system described by ξ are the same as those described by η, but each is to be computed separately. Most of the time, the forces in the two systems will round identically, $f_\xi^{(n)} = f_\eta^{(n)}$. But once in a while, presumably very rarely, the two will differ somewhere. As soon as this happens, the two systems will rapidly become quite different. The different rounding was the irreversible process. The notion of reversibility implies that it was somehow beyond the experimenter's control. The number of separate inte-

grations that were successfully done before this happened might be used as a measure of the "reversibility."

This definition of reversibility may seem too stringent; it could be argued that all that is required is that the reversed system return to a place "close to" the original starting point [11]. The definition used here does not make explicit use of the norm that would be required to give meaning to "close to." This definition also seems preferable because it is free of some of the complications of interpretation that go with explicit evaluations of $\| \xi^{(n)} - \xi^{(0)} \|$ (see, e.g., Ref. [8]).

Using the language of statistics (with substantial reservations about the language, but the comparison of two systems already implies an ensemble of sorts), we are led to inquire about the probability for the two force values to be different. Let F be the ("exact") force that would be computed by a perfect process, and let the computed value be $(F + \delta F)$. There is somewhere a roundoff boundary R. The probability for an error in the rounded force is the sum of the probabilities for making a rounding error in either direction. Near one rounding boundary this is $\mathscr{E}(|\delta F|)$ (the "mean absolute deviation"). If $\mathscr{E}(|\delta F|)$ is essentially independent of F, the total probability of making a rounding error (in either direction) summed over all possible rounded results is approximately $\epsilon = \mathscr{E}(|\delta F|)$. This can be obtained by considering that the probability for rounding incorrectly between R and $R + 1$ is $\mathscr{E}(|\delta F|) \times p(R)$, where $p(R)$ is the probability to get R as a rounded value. When a sum over all possible R's is evaluated, $\mathscr{E}(|\delta F|)$ factors out and $\sum p(R) = 1$ since $p(R)$ is a probability. This estimate of the chance of making an error in rounding is based on the notion of an underlying "reasonable" probability distribution governing the force values. This is an oversimplification, but it seems to provide the only reasonable way to estimate the chance of making a rounding error, and thus of determining how good (or how bad) the calculation looks from the standpoint of reversibility.

The probability of making a rounding error on one particle in one integration step is $\epsilon \approx \mathscr{E}(|\delta F|)$. In the entire system with many particles and many integration steps, one error somewhere, at some time, will throw the system off the track. This is a "go-no-go" process that is described by a binomial distribution characterized by ϵ. With w the total number of tries, the probability that no error has been made is $(1 - \epsilon)^w \approx 1 - w\epsilon$. When $w \approx 1/\epsilon$, there is a good chance that something irreversible has happened. Then w is the desired "reversibility measure."

In the large n-body calculation, force values are computed for each possible location in the space lattice. The number of force evaluations per integration step is twice the number of occupied lattice points (for two components of force in a two-dimensional calculation).

For convenience, assume that all locations are occupied. In our early calculations, there were $(128)^2 = 2^{14}$ locations. The force calculation provided nearly full single-register precision for the IBM 360 series with floating point numbers

scaled to about 16 for the largest magnitude, giving $\delta F \approx 2^{-19}$. This leads to $w = 2 \times$ (number of integration steps) $\times 2^{14} \approx 2^{19}$, or about 16 steps to an irreversible error. This is to be compared with our experiments [1] in which an exact fourfold symmetry was held for about 20 steps. The symmetry test can easily be seen to relate to the same criterion, since symmetrically disposed points were treated differently in the force calculation.

The reversibility of the calculation will be affected by changes in the time step or in the lattice spacing. According to Eq. (5), the values of the integer force scale according to m/q^2, and in a floating point calculation $\epsilon \approx \mathcal{E}(|\delta F|)$ will scale by the same amount. As m increases, the number of occupied lattice points will increase, but it cannot exceed the number of particles. Since the number of particles is a few times the number of allowed lattice points, (two–ten times the two-dimensional calculations; with three-dimensional calculations this ratio may be substantially higher) this won't make much difference. Then

$$\begin{aligned} w\epsilon &\to q\tau \frac{m}{q^2} \epsilon_0 \text{ [number of occupied lattice points]} \\ &\approx \tau \frac{m}{q} \epsilon_0 \text{ [min(number of allowed locations, number of particles)].} \end{aligned} \tag{12}$$

The number of integration steps (at the original time step) is τ and ϵ_0 is that "probability" of error in one step at one location. But if the problem is scaled to keep q^2/m constant, then the reversibility deteriorates with increasing q. A rule for scaling that simultaneously improves both D^2 and reversibility seems to be impossible.

Random rounding suggested in the previous section [illegible] means to reduce D^2

"random-number generators" are not strictly random; it is possible to construct one that runs reversibly in the sense that it can recover any previous random output. With a "reversible pseudorandom number generator" it would be possible to design a reversible calculation with "random-number" rounding. This violates the notion that the "random" element is outside the experimenter's control.

6. Liouville Theorem

A second advantage of the difference scheme is that occupancy of phase locations is conserved in the motion. Again, this results from the fact that integers are used as allowable position, velocity, and force values. A further feature assuring this conservation is that the forces are computed at each of the allowed configuration-

space locations. All particles that occupy a given configuration-space location experience the same force. With the exactness that integer arithmetic provides, this assures that all particles in a given phase location go to the same new location. Similarly, there is no way in which particles from two different phase-space locations can wind up together. This is true even if there are errors in computing the forces, because a table of rounded integer forces is constructed and used. This method of handling the forces is the feature of the calculation that permits large numbers of particles to be handled.

For calculations that are formulated differently from this one, the degree of phase conservation in the Liouville theorem can be discussed along lines of the reversibility argument of the previous section. The essential parameter is the chance that two particles, initially in the same phase-space location, will wind up at different locations. Given a difference scheme like Eq. (1), this is the chance that the two particles will be found to have different forces acting on them. Our calculation is set up so that this cannot happen because forces are rounded before being entered in the table; even rounding by adding a random number before entering in the table would not destroy the Liouville theorem property of our calculation.

Just as rounding with random numbers provides an experimental way for testing reversibility effects, a modification of the present calculation can be devised that permits testing for effects resulting from violations of the Liouville theorem. This would simply involve rounding the force values separately each time the force associated with a certain configuration space point is fetched. While the test for reversibility with random rounding can be carried out without affecting the behavior of the system with respect to the Liouville theorem, the reverse is not true. If the experiment is set up to test the effects of violating the Liouville theorem, reversibility will also be violated unless a "reversible pseudorandom number generator" is used.

For certain experiments on spiral structure, we have intentionally violated the Liouville theorem by rearranging the velocity distributions at a given configuration space location. This was done with a portion of the particles that were to approximate the dissipative dynamics of "gas" rather than the dynamics of stars, leading to a quite different behavior of the resulting system. The details are given elsewhere [2]; because the process was dissipative as well as being in violation of the Liouville theorem, it was not an adequate test of the effects of violating the Liouville theorem alone.

7. Integration Accuracy

The calculation approximates dynamics to the extent that the difference scheme (1) leaves out terms $O(\Delta t^3)$. In this sense it again differs from other, more

sophisticated, schemes only quantitatively. Some calculations have been run in which the ignored terms get dangerously large. Because the coefficient of the Δt^3 term includes $(\mathbf{u} \cdot \nabla)\,\mathbf{f}$ (interpreted in terms of the physical problem), the situation is worst for high velocity particles in regions where the force gradients are large. We try to attach physical interpretations to experimental results only where the integration criteria indicate that the integration is proceeding safely. The integration improves as the problem is refined by increasing m and q, if q^2/m is held constant.

Optimization of the calculation requires choosing the physical parameters in such a way that the truncation errors (errors involved in the finite difference scheme as approximating a derivative) in the worst case are about one lattice spacing after one integration step. This is larger than is customarily allowed, because the lattice spacing is large. The higher order ignored terms must be investigated; they may even be larger. We have partially avoided the problem by using a near-cutoff in the force law $((a^2 + x^2 + y^2)^{-1})$ which suppresses high derivatives of the forces. Even so, we have had to look into the problem.

8. Forces

The discussion so far might apply to any kind of nonrelaxing system. The relation to stellar dynamics, or to the gravitational n-body problem, arises through the forces used. These force values are computed just as if they were not to be rounded to integers, and then they are, in fact, so rounded. It is the fact that the forces need not be determined to high precision and can thus be obtained from the Poisson equation at a fairly small number of locations that lets the game handle very large numbers of particles. The time required for the force calculation is independent of the number of particles. It is important, however, that the force calculation take proper account of large numbers of particles at substantial distances. This is done by calculating in floating point, and only rounding to integer values at the very last step. The force calculation is the only part of the "integration" that requires arithmetic, when the calculation is done in the bit-manipulation mode described in Section 2.

The computational efficiency resulting from the use of discrete representations is obtained at the price of using a bounded space, which in turn means introducing models for the density distribution beyond the boundaries of the system. These models always seem somewhat artificial. For the square regions in a 2-dimensional problem with periodic boundary conditions that we have used, the density model consists of replicated regions everywhere over the plane. A zero of potential at the boundaries would imply image systems that are mirror reflections of the original system. (In plasma problems, a zero of potential corresponds to a conducting

boundary. There is no gravitational analog of a conductor.) An inverse-square force law between particle pairs in the original space corresponds to a sum over all the particles appearing in the replicated systems of some (a priori unknown) force law. This peculiarity may be avoided by using a force law that corresponds to r^{-2} between all pairs of particles, including those in the image systems. The resulting force law, of course, does not look like r^{-2} between particles of the original system, when the images are ignored. The principal difference is at large separations. The method we have used to compute the forces, properly summed over image systems, is described in this section. These boundary conditions are convenient with forces calculated by Fourier transform methods.

Isolated systems may be represented by using model density distributions that contain wide regions of zero density [4, 5]. Suppose that an infinite system is broken up into allowed and forbidden regions with boundaries such that, if a particle strays across the boundary into a forbidden region, that particle is deleted from the system. If there is no vector separation of points wholly within one allowed region that, suitably displaced, can connect points of two different allowed regions, then a force law can be constructed such that particles behave as if they interact only with other particles of the same allowed region. With squares or cubes of edge L, this means an empty space of width L between image squares (or cubes). Larger spacings, even infinite ones, are permissible. The Fourier transform, with an empty space of width L, provides the most efficient way to compute forces for isolated systems. The few particles that spill over the boundaries into the forbidden regions are easily handled.

A. *Method of Calculating the Forces*

Fourier transform techniques, with two-dimensional calculations, have the advantage that the singular density distribution can be handled analytically[1]:

$$\nabla^2\psi = 4\pi G\sigma(x, y)\,\delta(z). \tag{13}$$

If ψ and σ are expanded into two-dimensional Fourier transforms, the Poisson equation reduces to

$$\psi_{jk} = \frac{2GNL}{\sqrt{(j^2 + k^2)}} \exp\left(-\frac{2\pi}{NL}\sqrt{j^2 + k^2}\,|\,z\,|\right)\sigma_{jk}\,. \tag{14}$$

The notation is evident: j and k are dimensionless wavenumbers in the x and y directions (components of a vector in the reciprocal lattice), N is the number of allowed discrete locations in a periodic lattice cell length L (a dimensioned variable).

[1] The notation of this section is different from that of previous sections. Re-used symbols may have different meanings.

Both the potential ψ and the surface density σ carry dimensions. The x and y components of forces may be obtained by differentiating the potential (analytically).

This form of the potential is clearly an example of the Convolution Theorem in Fourier transform theory. The prescription for using that theorem is to put a particle at the origin and to calculate the force field around it. Then the Fourier transform of the forces generated by an arbitrary density distribution may be obtained by a term-by-term multiplication of the Fourier transform of the density ([illegible] by the Fourier transform of the force field around the particle at the origin

problem, in a form amenable to machine calculation. Aliasing will be taken up later.

The Fourier transform permits faster calculation of the force components. The fast Fourier transform leads to a number of operations proportional to $N^2 \log(N^2)$. This may be compared to N^4 calculations for direct convolution.

We have used a force derivable from a potential $(x^2 + y^2 + a^2)^{-1/2}$ summed over the direct periodic lattice. This is clearly equivalent to setting $z = a$ in the expression (14); it is the potential a units (of the small distance between discrete permitted locations) above the plane in which the particles move.

In handling the isolated system by Fourier transform methods (particularly in two dimensions), the forbidden regions may be considered to be the result of multiplying the actual (surface) density distribution by a function that is unity at all points of the allowed region, but which is zero at all points of the forbidden regions. When the Fourier transform is taken, the product becomes a convolution in the reciprocal space. Then, for the same reasons of computational economy that made the use of Fourier transform methods attractive in the first place (to avoid N^4 operations), it is best to handle this as a product in the direct space. This means that a machine code that will handle $N \times N$ configuration space points with periodic conditions can, with small modifications (to the crossing boundaries and generation of force convolution coefficient routines) handle the isolated problem on $(N/2) \times (N/2)$ configuration space points.

B. *Generation of Convolution Coefficients*

When the forces about an isolated particle can be directly generated with adequate precision, the convolution coefficients are best obtained by Fourier-transforming the set of force values. This happens, for example, with the isolated systems where $1/r^2$ forces are required. When the image particles are included, a different approach may be required.

The (periodic) lattice sums converge slowly, making it difficult to get good approximations to the force over the discrete space. An alternative formulation of the problem of summing over the periodic lattice is provided by the Ewald summation method, commonly used in crystal dynamics problems. In fact, the

full method is not used, but the underlying idea is useful and provides a formulation that can be used for three-dimensional problems as well as for two.

The essential trick of the Ewald method is to introduce an integral substitution

$$\frac{1}{x} = \frac{2}{\sqrt{\pi}} \int_0^\infty d\rho \, e^{-x^2\rho^2}. \tag{15}$$

Let a vector in the direct space be decomposed into the sum of a vector to the origin of a periodic cell, $\mathbf{x}(l)$, and a vector within the cell, $\mathbf{x}$. A sum over l represents a sum over all equivalent points of the periodic lattice. Then the potential (apart from $4\pi Gm$ factors) at the point $\mathbf{x}$ due to a particle at the origin of each periodic cell is

$$\phi(\mathbf{x}) = \sum_l \frac{1}{|\,\mathbf{x}(l) - \mathbf{x}\,|} = \int_0^\infty d\rho \left\{ \frac{2}{\sqrt{\pi}} \sum_l e^{-[\mathbf{x}(l)-\mathbf{x}]^2\rho^2} \right\}. \tag{16}$$

The function in the curly brackets is periodic with the periodicity of the lattice. It can therefore be expanded into a Fourier series,

$$\{\ \} = \sum_h g[\rho; \mathbf{y}(h)]\, e^{2\pi i \mathbf{y}(h)\cdot\mathbf{x}}, \tag{17}$$

where $\mathbf{y}(h)$ are the vectors of the reciprocal lattice. An essential part of this argument is that the scalar products of $\mathbf{y}(h) \cdot \mathbf{x}(l)$ for vectors of the direct and reciprocal lattices have integer values. These terms then drop out of the Fourier series.

For the two-dimensional problem, let ξ be a vector in the plane, and let ζ be a vector out of the plane. The problem may be considered in a four-dimensional space, so ξ and ζ are four-vectors spanning orthogonal subspaces. The "offset" a may be regarded as being along the fourth axis. Then ξ decomposes into periodic lattice parts $\xi(l) - \xi$, and the potential is

$$\phi(\mathbf{x}) = \sum_l \frac{1}{|\,\xi(l) - \xi + \zeta\,|}. \tag{18}$$

With the procedure used earlier, this is

$$\{\ \} = e^{-|\zeta|^2\rho^2} \sum_l \frac{2}{\sqrt{\pi}} e^{-[(l)-\xi]^2\rho^2}, \tag{19}$$

since the scalar product of $\xi \cdot \zeta = 0$. The Fourier transform should only be in the

ξ plane: let $\eta(h)$ be the reciprocal lattice vector in that plane. Then the Fourier coefficients are functions of the remaining spatial directions, and

$$g[\rho;\, \eta(h),\, z] = \frac{1}{A_{\text{cell}}} \frac{2}{\sqrt{\pi}} \int_{\zeta=\text{const}} d\xi\, e^{-[-\xi+\zeta]^2\rho^2 + 2\pi i \eta(h)\cdot[-\xi+\zeta]}$$

$$= \frac{2\sqrt{\pi}}{A_{\text{cell}}} e^{-|\zeta|^2\rho^2} \frac{e^{-\frac{\pi^2}{\rho^2}|\eta(h)|^2}}{\rho^2}. \tag{20}$$

In arriving at this form there is some shifting of integration limits from integrals over a cell with a subsequent sum over cells to an integral over the entire plane. The details may conveniently be found in Born and Huang [22]. An alternative way of arriving at this result is to note that the function g is a special case of the elliptic function ϑ_3, and that the integral result is just Jacobi's imaginary transformation (Ref. [23], Section 21.51; Ref. [24]). For that reason, this is called the "theta-function transformation." The decomposition with an orthogonal subspace that does not get Fourier-transformed does not appear in the literature cited.

In the usual Ewald summation process, the equality of two integrands is used to break the integral over ρ into two parts, each of which is rapidly convergent. For our calculation, we want the Fourier-transformed part, and the a term makes the convergence so rapid (by suppressing high-order Fourier components) that it is most economical to generate the convolution coefficients directly. Expressions like that of Eq. (20) can therefore be integrated over infinite limits, with the following results (return to the j, k notation as components of η made dimensionless): For two dimensions

$$g(j, k;\, z = 0,\, a) = \text{const} \frac{e^{-2\pi\frac{a}{N}\sqrt{j^2+k^2}}}{\sqrt{(j^2+k^2)}}, \tag{21}$$

and for three (obtained by letting ξ and η span a three-dimensional subspace with ζ spanning a one-dimensional subspace orthogonal to ξ)

$$g(j, k, l;\, a) = \text{const} \frac{2a}{N} \frac{K_1\left(2\pi\frac{a}{N}\sqrt{j^2+k^2+l^2}\right)}{\sqrt{(j^2+k^2+l^2)}}, \tag{22}$$

where $K_1(x)$ is a modified Bessel function (see, e.g., Ref. [25], Section 9.6) and l is the third wavenumber. The constants of Eq. (21) are the same as in Eq. (14) when all the loose ends are gathered up.

These expressions give Fourier coefficients for computing the potential over all of the three-dimensional space. Thus, the potential represents point particles;

in the two-dimensional case they are constrained to move on a plane. They are *not* rods.

There is an infinite potential to subtract. This may be done by setting the $j = k = 0$ (three dimensions, $j = k = l = 0$) term to zero. This is equivalent to setting the total mass in a cell to zero; there is a background of negative mass uniformly distributed over the cell (and thus over all of space).

The forces may be derived from the potential; the array of convolution coefficients for generating the force is

$$\mathscr{F}_x(j, k) = \text{const } j \frac{e^{-2\pi \frac{a}{N}\sqrt{j^2+k^2}}}{\sqrt{(j^2 + k^2)}} \tag{23}$$

in two dimensions, with an obvious extension to three. The array of convolution coefficients (23) can be generated once and stored; only one array is required, as the other arrays can be obtained from it by permuting indices. The value of the constant in Eq. (23) is to be separately determined for each problem as described in the paper setting forth the method [1].

C. *The "Aliasing" Problem*

In a discrete periodic lattice, $N \times N$ Fourier coefficients are sufficient to describe any (reasonable) function defined on each of the discrete points. Beyond N values in one direction (of the reciprocal lattice), the values of $e^{2\pi i \mathbf{n}\cdot\mathbf{x}}$ repeat at each of the permitted locations. Only N values of j and N values of k are needed, but there is no rule to specify in which part of the j, k plane the region used is to be located. (This region is a Brillouin zone). Analytic expressions for ψ or F imply that all (integer) j, k are valid and will be used, but then we only use a limited set. This problem appears whenever Fourier series are used to represent phenomena otherwise regarded as continuous; in power spectral work, it is called "aliasing" [26]. In effect, a means is needed for summing over the high-order harmonics so that the calculation need extend only over the first Brillouin zone.

The Fourier transform of the density is defined for all j, k as well. It depends on the exact shape assumed for each of the "particles." Again, only N values of j and N of k are used, so the density is represented adequately by a small portion of the j, k plane. Let $\sigma(j, k)$ represent the Fourier transform of the density, and $\phi(j, k)$ represent the Fourier transform of the potential about a point particle at the origin. Then the convolution theorem gives

$$\psi(j, k) = \phi(j, k)\, \sigma(j, k), \tag{24}$$

where ψ is the Fourier transform of the actual potential, and j, k assume all integer values. Let j_R, k_R represent the values of j and k reduced modulo N:

$j = j_R + JN$ for integer J. We seek a function $\psi_R(j_R, k_R)$ that will correctly represent the potential at each allowed discrete location. The easiest way to obtain a good approximation for $\psi_R(j_R, k_R)$ is to assume that $\sigma(j_R, k_R) = \sigma(j, k)$, i.e., σ is periodic in j and k with period N. Then

$$\psi_R(j_R, k_R) = \left[\sum_{J=-\infty}^{\infty} \sum_{K=-\infty}^{\infty} \phi(j_R + JN, k_R + KN)\right] \sigma(j_R, k_R) \tag{25}$$

represents a proper transformation of the entire j, k plane by using only some zone of N values of j and of k. The model for a particle is then a Dirac delta function repeated with the periodicity of the lattice (not the distance between allowed discrete points: a one-dimensional form of this is the "infinite Dirac comb" of Blackman and Tukey, Ref. [26], Section A2).

This is not the only way to handle the aliasing problem: ϕ might be modified to be a sum of infinite Dirac combs, with some shape ascribed to σ. Functions might be used for which the sums can be done explicitly. However it is handled, the matter of high-order terms cannot be ignored.

The calculation now uses force convolution coefficients obtained from a sum like that of Eq. (25) with individual terms like those of Eq. (23). Because the a term in Eq. (23) causes very rapid convergence, the sums can be evaluated in a computer. In the two-dimensional problem, four terms in each J and K are sufficient to

[illegible] any pair of points given by $x[x^2 + y^2 + a^2]^{-3/2}$. Forces are generated directly, not by taking differences of potentials.

The same kind of problem of disposing of high-order terms will appear in any geometry. Truncation of the series implies certain assumptions about the character of the particles, and serious trouble can result if this is not properly taken into account.

D. *Potentials*

The forces might be derived from a potential by differencing successive values of potentials obtained by using the array of Eq. (21) as convolution coefficients rather than the differentiated form of Eq. (23). This has the advantage of assuring irrotational forces (except for rounding effects). The nature of the approximation involved is conveniently studied by the following argument. A component of the force at some location is the gradient of ψ or may be the difference of successive values of some other function χ,

$$F_x = \frac{\partial \psi}{\partial x} = \frac{1}{2\Delta l}[\chi(x + \Delta l) - \chi(x - \Delta l)]. \tag{26}$$

Then the relation between the two functions is

$$\psi = \frac{1}{2\Delta l} \int_{x-\Delta l}^{x+\Delta l} dt\, \chi(t). \tag{27}$$

This function is not unique: a single χ may not exist such that both F_x and F_y can be found from it by differencing. The approximations in F by differencing methods follow directly from Eq. (27). With grids like 64×64 or 256×256, the approximations are not small.

Note that it is important, in differentiating a potential to get a force, that the Fourier representation including the entire j, k plane be used; not the collapsed representation of Eq. (25). Setting the 0–0 term of the aliased convolution coefficients to zero is equivalent to defining the additive constant in the potential in such a way that the potential sums to zero over all permitted discrete points in a periodic cell. The potential no longer averages to zero when integrated over a periodic cell. Potentials may not be directly calculated at points other than those permitted in the discrete lattice, using these Fourier techniques. Interpolation is permissible, however. In attempting to evaluate a potential energy (per periodic cell), the self-energy of particles must be eliminated, and the potential of overlapping particles must be properly evaluated.

A total potential for the (periodically replicated) system can be defined. This requires subtracting a uniform density from the system so the total mass inside a periodic cell is zero. Then the total potential energy per unit cell can be defined in the usual way. The effects of interactions between periodically replicated cells can be eliminated just as they are in crystal physics. It should even be possible to recover a kind of virial theorem for a periodic system such as is in fact calculated when the Fourier techniques are used. We have not done any of this so far.

9. Conclusions

The particularly simple difference scheme used in the large n-body calculation makes it easy to investigate the numerical and computational properties of the system, even though restricting the "particles" to a discrete phase space as coarse as the one used here forces consideration of some aspects of the difference schemes that are usually ignored. The appearance of spatial lattices as well as time steps makes this look a little bit like the problem that arises in discussing partial differential equations. Indeed, from one point of view, we are solving a partial differential equation by integrating along its characteristics—the individual particle orbits.

The finite difference scheme is stable as an initial-value calculation. Modifications due to the discrete phase space and discrete force values can be introduced into

the inhomogeneous term in the difference scheme and explicitly evaluated. The features of the system that give rise to reversibility and to the special form of the Liouville theorem appropriate to this problem can be isolated and tested experimentally. We have not yet conducted such experiments and do not know how important these features are in the representation of real stellar systems. However, they are the features that make the approximation appealing because it is close to the physics of the problem. The scheme is also appealing because it leads to a model that can be computed exactly. All the approximations are between the physical system and the model. The estimates of this note are attempts to see what has happened to the physics in setting up the model.

The idea underlying this calculation, that of modifying the data to fit the calculation, rather than the other way around, may have wider applicability in physical calculations than just to the n-body calculation described here. The description has been given in detail to show that the approximations can be fairly thoroughly analyzed.

The notions of discrete spaces can be directly applied to several kinds of calculations. Prendergast [27] has devised a similar calculation for hydrodynamics; there, the system is made to relax to an assumed form for a distribution function after each integration step so that what were "particles" in the n-body calculation become "streams." Prendergast reports very effective handling of shock-like discontinuities. Again, the discrete phase space notion seems to have direct application to a kind of semiclassical Fermi gas approximation to crystal or molecular dynamics—the exclusion principle can be rigorously built in. Indeed, the method is directly applicable to any system that can be described in terms of an equation like the Boltzmann equation. But the idea of exactly calculable models may have substantially wider applicability.

We have used the programs described here to study persistent spiral structure in self-gravitating systems, and for certain studies in the stability of plane self-gravitating systems. The results are being published elsewhere [2].

ACKNOWLEDGMENTS

It is a pleasure to thank Prof. K. H. Prendergast of Columbia University, who has been deeply involved in all phases of this work from its inception, and Mr. William J. Quirk, also of Columbia, who joined the project shortly after it started. The bulk of the computations have been carried out at the Goddard Institute for Space Studies in New York, through the courtesy of Dr. Robert Jastrow, Director. The work has been assisted by the U. S. Atomic Energy Commission under contract AT(11-1)-614. Part of it was done with the aid of a National Research Council Senior Postdoctoral Resident Research Associateship at the Goddard Institute for Space Studies, supported by the National Aeronautics and Space Administration.

REFERENCES

1. R. H. MILLER AND K. H. PRENDERGAST, *Astrophys. J.* **151** (1968), 699.
2. R. H. MILLER, K. H. PRENDERGAST, AND W. J. QUIRK, *Astrophys. J.*, **161** (1970), 903.
3. R. H. MILLER AND N. ALTON, "Three Dimensional *N*-Body Calculation," Quarterly Report, Institute for Computer Research, University of Chicago, No. 18, Aug. 1, 1968, unpublished.
4. R. H. MILLER, "Forces for the *N*-Body Calculation," Quarterly Report, Institute for Computer Research, University of Chicago, No. 17, May 1, 1968, unpublished.
5. F. HOHL AND R. W. HOCKNEY, *J. Comput. Phys.* **4** (1969), 306.
6. R. W. HOCKNEY AND F. HOHL, *Astron. J.* **74** (1969), 1102.
7. R. L. MORSE AND C. W. NIELSON, *Phys. Rev. Lett.* **23** (1969), 1087.
8. R. H. MILLER, *Astrophys. J.* **140** (1964), 250.
9. R. H. MILLER, *J. Comput. Phys.* **2** (1967), 1.
10. S. J. AARSETH, *Mon. Notic. Roy. Astron. Soc.* **126** (1963), 223.
11. S. J. AARSETH, *Mon. Notic. Roy. Astron. Soc.* **132** (1966), 35.
12. E. M. STANDISH, Thesis, Yale University, 1968, unpublished.
13. R. WIELEN, Ver. Astron. Rechen-Instituts Heidelberg, No. 19 (1967).
14. R. WIELEN, *Bull. Astron. (3)* **3**, (1968), 127.
15. M. LECAR, *Bull. Astron. (3)* **3**, (1968), 91.
16. R. ZWANZIG AND N. K. AILAWADI, *Phys. Rev.* **182** (1969), 280.
17. J. L. LEBOWITZ, J. K. PERCUS, AND L. VERLET, *Phys. Rev.* **153** (1967), 250.
18. D. GREENSPAN, *in* "Studies in Applied Mathematics," Mathematical Association of America, to appear.
19. F. H. HARLOW, *Methods Comput. Phys.* **3** (1964), 319.
20. B. WENDORFF, "Theoretical Numerical Analysis," Academic Press, New York, 1966.
21. O. BUNEMAN, *J. Comput. Phys.* **1** (1967), 517.
22. M. BORN AND K. HUANG, "Dynamics of Crystal Lattices," Oxford University Press, New York/London, 1954.
23. E. T. WHITTAKER AND G. N. WATSON, "A Course in Modern Analysis," Cambridge University Press, New York/London, 1952.
24. M. BORN AND M. BRADBURN, *Proc. Cambridge Phil. Soc.* **39** (1943), 104.
25. M. ABRAMOWITZ AND I. S. STEGUN, "Handbook of Mathematical Functions," U. S. Government Printing Office, Washington, D. C., 1964.
26. R. B. BLACKMAN AND J. W. TUKEY, "The Measurement of Power Spectra," Dover Press, New York, 1959.
27. K. H. PRENDERGAST, private communication.