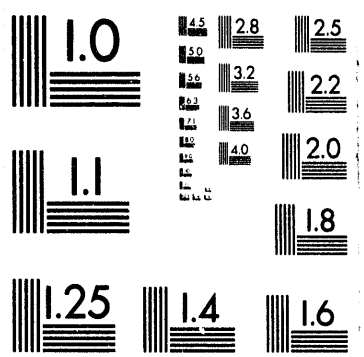
1.0
1.1
1.25
1.4
1.6
1.8
2.0
2.2
2.5
2.8
3.2
3.6
4.0
4.5
5.0
5.6
6.3
7.1
8.0
9.0
10
11
12
13
14
15
16
17
18
19
20
21
22

1 of 1

Conf-9309176--5

Many-Body Mean-Field Equations: Parallel Implementation

Michel Vallières[1], Sait Umar[2], Michael Strayer[3], and Clayton Chinn[2]

[1]Physics & Atmospheric Science, Drexel University, Philadelphia, PA 19104
[2]Dept. of Physics & Astronomy, Vanderbilt University, Nashville, TN 37235
[3]Physics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831

to be published

Proceedings

2nd International Conference on Computational Physics
Beijing, China
September 13-17, 1993

MANY-BODY MEAN-FIELD EQUATIONS:

PARALLEL IMPLEMENTATION

M. VALLIÈRES[1], A.S. UMAR[2], M.R. STRAYER[3] and C. CHINN[2]
[1] *Physics & Atmospheric Science, Drexel University, Philadelphia, PA 19104*
[2] *Dept. of Physics & Astronomy, Vanderbilt University, Nashville, TN 37325*
[3] *Physics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831*

ABSTRACT

We describe the implementation of Hartree-Fock Many-Body Mean-Field Equations on a Parallel Intel iPSC/860 hypercube. We first discuss the Nuclear Mean-Field approach in physical terms. Then we describe our parallel implementation of this approach on the Intel iPSC/860 hypercube. We discuss and compare the advantages and disadvantages of the domain partition versus the Hilbert space partition for this problem. We conclude by discussing some timing experiments on various computing platforms.

1. Introduction

Accurately modeling the global properties of nuclei presents a challenge to Nuclear Physicists. The masses, deformation, separation energies and the location of the proton and neutron drip line are very important properties that may depend very sensitively on the nucleon-nucleon potential parameters and unfortunately where predictions may also depend sensitively on the model used. This makes the prediction of these global properties very challenging.

Nuclear physics is always in quests of test cases where the adequacy of the effective interaction can be ascertained; the systematic and accurate prediction of the global properties of nuclei is certainly foremost in this aspect. There is a wealth of data of this type for nuclei of all masses against which to test the modeling. Furthermore, these data allow for testing the accuracy of the nuclear model used in the prediction, both from the physics and computational points of view.

The real challenge for nuclear physics is in extrapolating the model predictions to the proton- and neutron-rich nuclei. These nuclei are very important in astrophysics since they are the key elements in the nucleosynthesis chain. Nuclear formation is thought to occur via proton (X-ray bursts in binary stars[1]) and neutron (in supernova[1]) capture within nuclei far from stability until the nuclei become so unstable against beta decay that the particle capture chain breaks. This produces favored paths in the production of elements towards the production of ever larger elements. Explanations of nucleosynthesis therefore require a very precise knowledge of the nuclear masses and separation energies and beta decay rates.

The Radioactive Ion Beam (RIB) facility at Oak Ridge National Laboratory will allow the systematic exploration of properties of proton-rich nuclei. This facility

is based on a two-stage acceleration process; normal nuclei will be accelerated, followed by the acceleration of the radioactive ions produced when these normal nuclei hit a target. The facility is due on line within one year. It will allow for the first time for measuring global properties of proton-rich nuclei and therefore test the theoretical prediction of these properties.

A feasible approach to model the global nuclear properties is via the mean-field approach. In this model the nucleons are assumed in the average field of all nucleons in a nucleus. The Hartree-Fock method implements this idea.

Following a brief physical discussion of the nuclear mean-field approach, this talk will describe an accurate numerical solution approach to these equations based on the basis-spline collocation method. We will then discuss the implementation of this approach on the Intel iPSC/860 hypercube supercomputer, followed by a discussion of some timing experiments on various computing platforms. The emphasis throughout will be on the numerical implementation of the approach.

2. Hartree-Fock Approach

The nuclear many-body system to solve is characterized by a two- and three-body interaction

$$\hat{H} = \sum_{i}^{N} t_i + \sum_{i<j}^{N} V_{i,j} + \sum_{i<j<k}^{N} V_{i,j,k}. \tag{1}$$

The potential we use is the Skyrme interaction.[2] This phenomenological interaction compares well to the effective interaction obtained via many-body techniques. It depends on very few parameters which ought to be treated as constants for all nuclei once fitted. It has the advantage of rendering the Hartree-Fock equations local, thus greatly simplifying the solution to the problem. The mean-field equations (details can be found in Ref. [3]) follow from a variational principle applied to the action of the system,

$$S = \int_{t_1}^{t_2} \left\langle \Phi(t) | \hat{H} - i\hbar\partial_t | \Phi(t) \right\rangle \tag{2}$$

with respect to the single-particle wavefunctions. Choosing a Slater determinant for the many-body wavefunction

$$\Phi_0(\{r\}, t) = \frac{1}{\sqrt{N!}} det|\phi_\lambda(\vec{r}, t)| \tag{3}$$

yields Schrödinger-like equations to solve for the single-particle wavefunctions. Assuming a static solution yields the time-independent Hartree-Fock equations

$$\left\{ -\frac{\hbar^2}{2m}\nabla^2 + V(\rho, \tau, j, J) + V_{Coulomb} \right\} \begin{pmatrix} \chi_\alpha^+(\vec{r}) \\ \chi_\alpha^-(\vec{r}) \end{pmatrix} = e \begin{pmatrix} \chi_\alpha^+(\vec{r}) \\ \chi_\alpha^-(\vec{r}) \end{pmatrix}. \tag{4}$$

The wavefunctions are spin $\frac{1}{2}$ spinors. The potential occurring in the equations depends on four density functionals:

$$\rho(\vec{r}) \quad = \quad \sum_{\alpha} \omega_\alpha \left\{ |\chi_\alpha^-(\vec{r})|^2 + |\chi_\alpha^+(\vec{r})|^2 \right\} \tag{5}$$

$$\tau(\vec{r}) = \sum_{\alpha} \omega_{\alpha} \left\{ |\nabla \chi_{\alpha}^{-}(\vec{r})|^2 + |\nabla \chi_{\alpha}^{+}(\vec{r})|^2 \right\} \quad (6)$$

$$\vec{j}(\vec{r}) = \sum_{\alpha} \omega_{\alpha} \left\{ \chi_{\alpha}^{-*}(\vec{r}) \vec{\nabla} \chi_{\alpha}^{-}(\vec{r}) + \chi_{\alpha}^{+*}(\vec{r}) \vec{\nabla} \chi_{\alpha}^{+}(\vec{r}) \right\} \quad (7)$$

$$\vec{J}(\vec{r}) = -i \sum_{\alpha} \omega_{\alpha} \chi_{\alpha}^{\mu *}(\vec{r}) (\nabla \times \sigma) \chi_{\alpha}^{\mu'}(\vec{r}). \quad (8)$$

The Coulomb potential satisfies Poisson equation

$$\nabla^2 V_{Coulomb}(\vec{r}) = -4\pi\epsilon^2 \rho(\vec{r}). \quad (9)$$

The factors ω_{α} in the density equations are the occupation of each orbital. They result from modeling the short range part of the nucleon-nucleon interaction. A pairing piece is often added to the Skyrme interaction to simulate these short range correlations. In our approach we treat the pairing interaction via the usual BCS approximation or the Lipkin-Nogami [4] approximation. They lead to a set of non-linear algebraic equations which must be solved after each update of the single-particle energies, adding somewhat to the complexity of the calculations.

The time-independent problem now reduces to finding the solutions to the non-linear set of differential equations given above for all the single-particle wave-functions necessary to describe a given nucleus. In practice, the solution is obtained via an iterative scheme according to the following steps

1. "Guess" a set of orthogonal spinors (e.g., deformed harmonic oscillators)
2. Solve the pairing equations
3. Compute the densities
4. Compute the Hartree-Fock potential
5. Solve the Poisson equation for $V_{Coulomb}$
6. Solve the Schrödinger equation for the spinors
7. Orthogonalize the spinors
8. Repeat from step #2 until convergence (stability of observables)

These steps can be understood in terms of an imaginary time evolution by which the system is said to evolve toward the ground state.

In some instances constraints may be applied to this set of equations to force the system to exhibit some particular features. For instance an easy constraint to understand physically would be to impose a given "shape" for the solution, i.e., for the nucleus to be of quadrupole shape. Such constraints add extra pieces to the self-consistent potential appearing in the Schrödinger equation. We are currently looking into appropriate constraints to impose for studying the response function of the system corresponding to various decay modes.

The time-dependent set of equations is solved much in the same way as described above, except for two important differences: the iterations to achieve

^{16}O	Radial N=500 D=0.025	3D, M=3 N=22 D=0.9	3D, M=5 N=22 D=0.9	3D, M=7 N=22 D=0.9	3D, M=9 N=22 D=0.9	3D, M=7 N=33 D=0.6
EHF (MeV)	-127.73	-122.20	-127.18	-127.69	-127.73	-127.73
rms (fm)	2.6822	2.7179	2.6872	2.6826	2.6821	2.6822
Es1/2 (n) (MeV)	-33.307	-33.038	-33.305	-33.311	-33.308	-33.307
Ep3/2 (n) (MeV)	-19.882	-19.249	-19.814	-19.873	-19.880	-19.882
Ep1/2 (n) (MeV)	-13.551	-13.452	-13.515	-13.540	-13.545	-13.551
Es1/2 (p) (MeV)	-29.739	-29.523	-29.741	-29.743	-29.740	-29.739
Ep3/2 (p) (MeV)	-16.477	-15.908	-16.415	-16.468	-16.474	-16.475
Ep1/2 (p) (MeV)	-10.270	-10.226	-10.242	-10.260	-10.265	-10.269

	EHF	rms (fm)	Q20		
Radial	-341.11	3.4016	0	^{40}Ca	($M = 7$)
3D	-341.00	3.4023	0.0002		($N = 24$)

	EHF		
Radial	-1636.4	^{208}Pb	($M = 9$)
3D	-1636.3		($N = 22$)

Table 1. Convergence of the basis-spline approach for various spline order (M) and lattice size (N) for various spherical nuclei against a spherical calculation.

"self-consistency" now correspond to a real time evolution of the system. Further, the orthogonalization step is no longer necessary.

3. Basis-Splines

The equations are solved in discretized Cartesian configuration space to ensure that no bias is applied toward some specific geometry for the solution. The differential equations are solved via the basis spline expansion.[5] This basis is optimal for the current problem. The spline functions are piecewise-continuous polynomials joined at "knots"; the wavefunctions are expanded in the spline functions and *exactly* represented at collocation points. The whole problem of solving the differential equations transforms itself into that of solving a matrix equation for the wavefunctions themselves at the collocation points. The boundary conditions are taken into account within these matrix equations. The operators acquire a matrix representation which can be computed ahead of time once and for all. The details of this approach are to be found in Ref. [5].

This method is very accurate. As an example we quote in Table 1 the convergence of binding energy and single-particle energies for ^{16}O, ^{140}Ca and ^{208}Pb; these nuclei are spherical and therefore can be solved numerically "exactly" via a spherical code.

4. Parallel Implementation

The solution of the mean field equations as explained in the previous

HF steps	Domain Decomposition	Hilbert Space Decomposition
- densities	Global Broadcasts	Global Sums
- potentials	none	none
- Diff Eq Solutions	Nearest Neighbor (heavy)	none
- Orthogonalization	Global Sums	Global Broadcasts (heavy)
- BCS	Global Broadcasts (light)	Global Broadcasts (light)

Table 2. Communication overhead in the two parallelization scheme.

sections is very compute intensive. The calculation requires of the order of 1.3MB/wavefunction (150 to 200 wavefunctions needed for large nuclei) and approximately 15 Cray 2 CPU hours for *Pb*. This suggests that an implementation on parallel computers is needed. We will now describe such an implementation on the Intel iPSC/860 hypercube supercomputer.

The set of equations can be implemented in a parallel fashion either via *domain decomposition* or via *Hilbert space decomposition.* The former way is the most common (and often times the only possible way) approach to the parallelization of a physical problem; in this case the configuration space is divided into small sub-domains, each of which is treated on one of the nodes of the parallel computer. In problems where "few" functionals are sought, i.e., fluid flow, atmospheric problems, ..., this is the only feasible parallelization path. For our problem, on the other hand, a second parallelization approach exists, that of the *Hilbert space decomposition.* This stems from the fact that we must solve for many (hundreds for large nuclei) wavefunctions. Therefore it becomes possible to parcel out the wavefunctions among the nodes, each wavefunction residing entirely within a specific node. This is the way we chose to parallelize our approach.

The efficiency of a parallel implementation depends on our minimizing the communication overhead. Table 3 lists the communication burden in each of the two parallelization schemes. It shows that the heavy communication lies in two different parts of the calculation for the two parallelization schemes. The Hilbert space decomposition has the added advantage of being very simple to implement since most of the coding is identical to that on a scalar machine, in particular the differential equation solver, and that only the orthogonalization needs careful recoding. It allows a machine-independent implementation of the approach.

The algorithm for orthogonalization needs careful setup in Hilbert space decomposition since the burden of communication lies in this part of the calculation. The algorithm is shown in Fig. 1 and summarized below

1. Orthogonalize local wavefunctions on each node
2. Let node zero broadcast its local wavefunctions to nodes of higher node numbers
3. Let the nodes receive the broadcasted wavefunctions; if the broadcast came from a node of lower number, orthogonalize the received wavefunctions to the local ones
4. Repeat for all subsequent nodes

Size	Mass	# nodes	Total CPU	Orth CPU	Remaining CPU
12	24	4	357	8	349
	48	8	492	204	288
	96	16	831	531	300
16	24	4	1145	275	870
	48	8	1407	535	872
	96	16	2258	1365	893
20	24	4	2252	324	1928
	48	8	2878	935	1943
	96	16	4411	2450	1961

Table 3. CPU time per node for mass 24, 48 and 98 nuclei and size 12, 16 and 20 lattices.

5. Timing Experiments

Fig. 2 shows a comparison of Cray and iPSC/860 hypercube times as a function of lattice size. The Cray becomes more efficient for a larger lattice due to a better vectorization. Fig. 3 shows that the CPU time follows roughly an n^4 dependence. This easily traces back to the HF formula. Fig. 4 shows the same information concerning the orthogonalization times.

A criteria for the success of the parallelization of the mean-field approach, as described here, is whether the total CPU time per node for computing various nuclei remains constant. Namely, the ideal situation would be if computing ^{16}O and ^{208}Pb would require the same CPU time per node, and therefore the same wall clock time, upon proper scaling the size of the hypercube used. Table 3 shows our results for different mass nuclei where we have made sure to scale the hypercube size according to the number of wavefunctions needed. The code scales according to expectation except for the orthogonalization piece which still falls short of our goal.

6. Conclusions

We have described a parallel implementation of nuclear mean-field approach. We chose to implement a *Hilbert space decomposition* approach which was rendered possible by the fact we need to solve for hundreds of wavefunctions in the calculation. This allows a scaling of the calculation whereby the CPU time per node is almost constant for calculations of nuclei of various sizes, except for the orthogonalization piece of the calculation, upon scaling the hypercube size according to the number of wavefunctions. Note, however, that this scaling is fully achieved in the time-dependent Hartree-Fock approach which does not require orthogonalization.

Acknowledgement

This research has been supported in part by the U.S. Department of En-

ergy (DOE) Office of Scientific Computing under the High Performance Computing and Communications Program (HPCC), as a Grand Challenge titled the Quantum Structure of Matter, and in part by DOE under contract No. DE-AC05-84OR21400 managed by Martin Marietta Energy Systems, Inc., and under contract No. DE-FG05-87ER40376 with Vanderbilt University, and by the National Science Foundation. Some of the numerical calculations were carried out on the Intel iPSC/860 hypercube multicomputer at the Oak Ridge National Laboratory and on the Cray C90 at the National Energy Research Supercomputing Center, Livermore.

Bibliography

1. *Astrophysics and Nuclei far from Stability,* F.K. Thielemann, K.L. Kratz, B. Pfeiffer, T. Rauscher, L. van Wormer, and M.C. Wiescher, to appear in Nuc. Phys. **A**.
2. T.K.R. Skyrme, Nuc. Phys. **9**, 1959, 615.
3. A.S. Umar and M.R. Strayer, Comp. Phys. Comm. **63** , 1991, 179 and references therein.
4. H.J. Lipkin, Ann. Phys. **31**, 1960, 525; Y. Nogami, Phys. Lett. **15**, 1965, 335.
5. C. de Boor, *Practical Guide to Splines,* (Springer, New York, 1978); A.S. Umar, J. Wu, M.R. Strayer and C. Bottcher, J. Comp. Phys., **93**, 1991, 426.

Figure Captions

Fig. 1 Schematic representation of the orthogonalization algorithm.

Fig. 2 Hypercube and Cray CPU times for 50 iterations for a mass 48 nucleus.

Fig. 3 Total hypercube CPU times as a function of lattice size for a fixed number of nodes.

Fig. 4 Same as Fig. 3, except for the orthogonalization only.

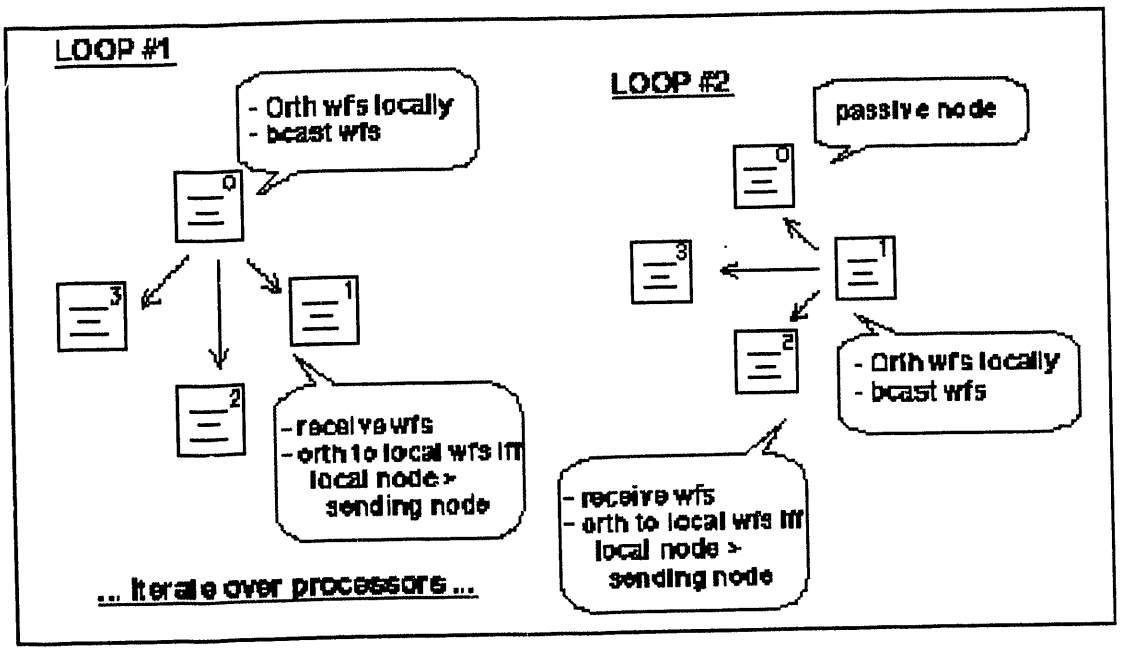

Figure 1

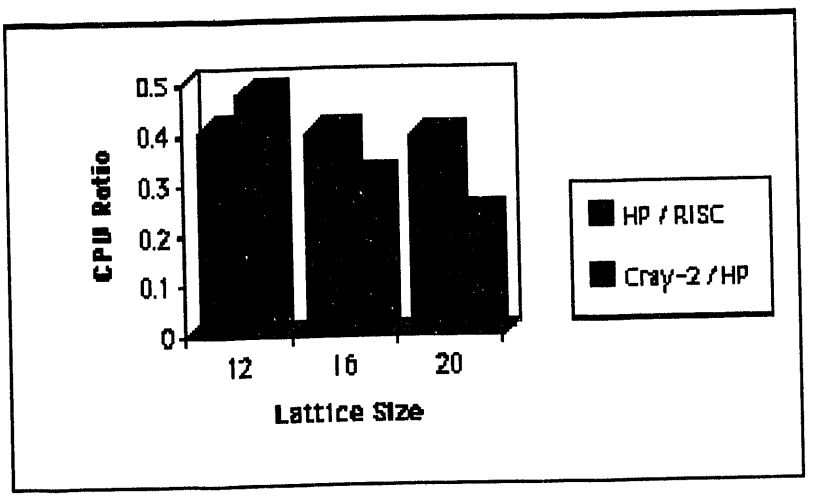

Figure 2

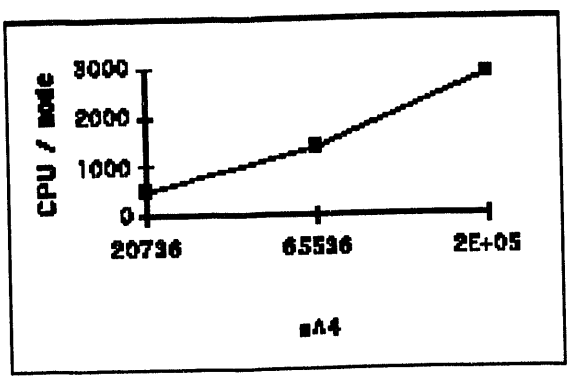

Figure 3

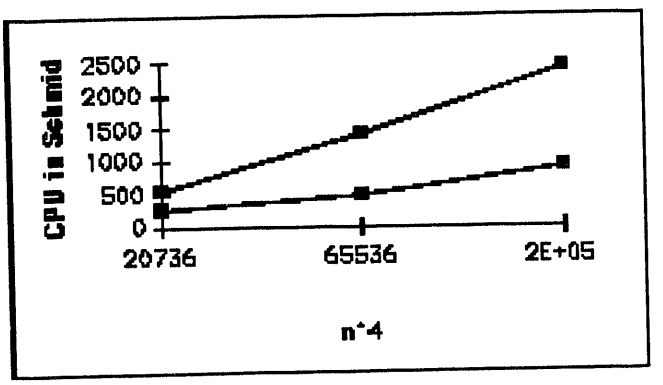

Figure 4

DATE
FILMED
5/10/94
END