


***Ab initio* Condensed Matter Calculations on the QCD Teraflops Computer**

T.A. Arias, B.E. Larson, M. Galván and J.D. Joannopoulos

RLE Technical Report No. 576

February 1993



The RESEARCH LABORATORY *of* ELECTRONICS

MASSACHUSETTS INSTITUTE of TECHNOLOGY
CAMBRIDGE, MASSACHUSETTS 02139-4307

***Ab initio* Condensed Matter Calculations on the QCD Teraflops Computer**

T.A. Arias, B.E. Larson, M. Galván and J.D. Joannopoulos

RLE Technical Report No. 576

February 1993

**Research Laboratory of Electronics
Massachusetts Institute of Technology
Cambridge, Massachusetts 02139-4307**

Ab initio Condensed Matter Calculations
on the
QCD Teraflops Computer

T.A. Arias⁽¹⁾, B.E. Larson⁽²⁾,

M. Galván⁽¹⁾ and J.D. Joannopoulos⁽¹⁾

⁽¹⁾*Research Laboratory of Electronics
Massachusetts Institute of Technology
Cambridge, MA*

⁽²⁾*Thinking Machines Corporation
Cambridge, MA*

Recently proposed enhancements to the CM-5 architecture for QCD calculations will also allow Teraflop performance in *ab initio* material science and chemical calculations with current, well established techniques. Such performance brings fundamental computational studies to a *qualitatively* new echelon and unleashes the full predictive power of first principles' techniques on problems in several critical technologies. Additionally, the new CM-5 architecture is well poised to take advantage of two coming major developments in condensed matter calculation.

I. INTRODUCTION

Ab initio calculation of condensed matter systems has advanced greatly in recent years. This approach has yielded tremendous understanding of the thermodynamic properties of bulk materials[1], crystal surface structures[2] and dynamics[3], the nature of point defects[4], and the diffusion and interaction of impurities in bulk materials[5]. However, the computational demands of the approach and the attendant bounds on the size of systems which may be studied (roughly one hundred atoms) have limited the direct impact of the approach on materials and chemical engineering. Several *ab initio* applications which will benefit technology tremendously remain out of reach, requiring an order of magnitude increase in the size of addressable systems. Problems requiring the simultaneous study of thousands of atoms include defects in glasses (fiber optics communications), complexes of extended crystalline defects (materials' strength and processing), and large molecules (drug design). The computational demands of the *ab initio* technique scales as the third power of system size. Studies of thousands of atoms require a *three order of magnitude* increase in computational speed, and Teraflop performance is needed to perform such studies.

Negele[6] has recently proposed enhancing the performance of the computation nodes of the commercially available CM-5 to attain Teraflop performance in QCD calculations by bringing the CM-5 into computation-communication load balance for QCD. Improving the node design impacts tremendously on the machine's performance for this application because the QCD Hamiltonian is a *local* operator and, therefore, places a much higher premium on computational than communications performance. The most successful *ab initio* condensed matter Hamiltonians are also local operators[7], and we shall demonstrate in this report that condensed matter calculations will also attain Teraflop performance on the proposed machine. The new machine will be the first machine capable of addressing the aforementioned problems of pressing

technological importance from first principles.

In the next section, we describe the particular *ab initio* electronic structure technique which best suits the parallel computing environment and outline the major steps in a calculation performed with this technique. Section III then describes the basic layout of the calculation among the processing nodes of the enhanced CM-5 and defines the quantities needed to describe the rigors of the calculation. Section IV then presents a summary of the computational and communications requirements of *ab initio* electronic structure calculation on the enhanced machine. This section also gives a detailed account of the demands of each step in the calculation sufficiently specific to serve as a guide for programming. The report concludes in section V with a discussion of the results presented in section IV.

II. OUTLINE OF THE LCGTO-LSD METHOD

There are a variety of well established *ab initio* electronic structure methods[7]. The technique which best exploits the locality of the Schrodinger equation is the the linear combination of Gaussian-type orbitals-local spin density (LCGTO-LSD) method[8–11]. This method expands the electronic orbital wavefunctions in terms of sums of Gaussian functions and then solves the local Kohn-Sham Hamiltonian[12] self-consistently. The localized spatial nature of the Gaussian basis function ensures that the Hamiltonian matrix elements decay extremely rapidly with distance and renders application of the Kohn-Sham Hamiltonian *local* in the same sense as is the application of the QCD Hamiltonian described in [6].

The calculation for condensed matter, however, is somewhat more complex than in the QCD case. The self-consistent nature of the Kohn-Sham Hamiltonian necessitates an iterative process where the Hamiltonian must be continually re-computed. Further, the electronic structure problem requires the calculation of many filled Fermionic

electron states, rather than merely the ground state and a few excited states, as in the QCD case. Furthermore, orthonormality must be maintained among all of the electronic states. Consequently, the LCGTO-LSD calculation proceeds by the iterative repetition of three phases: 1) calculation of the self-consistent Kohn-Sham Hamiltonian, 2) application of that Hamiltonian to the electronic wavefunctions (and the addition of a fraction of the result to the wavefunctions), and 3) re-orthonormalization of the wavefunctions.

The self-consistent Kohn-Sham Hamiltonian in a Gaussian basis set $\{\mu\}$ representation is

$$H_{\mu\nu}^{(sc)} = H_{\mu\nu}^{(1)} + V_{\mu\nu}^{(dir)} + V_{\mu\nu}^{(xc)}, \quad (2.1)$$

where $H_{\mu\nu}^{(1)}$ is the sum of the kinetic energy and bare nuclear potential terms, $V_{\mu\nu}^{(xc)}$ is the exchange-correlation potential, and $V_{\mu\nu}^{(dir)}$ is the Hartree, or direct, potential. The exchange-correlation potential is computed in terms of $\rho(r_i)$, the electronic charge density evaluated on an integration grid for the exchange-correlation energy. The values of $\rho(r_i)$ are most efficiently computed from the density *matrix* $P_{\mu\nu} \equiv \sum_n C_{n\mu}^* C_{n\nu}$, where $C_{n\mu}$ is the coefficient of basis element μ for electronic wavefunction n . Evaluation of the Hartree potential is more complex because of the long-range nature of the Coulomb interaction. To fully exploit the parallel computing environment, it is best to replace the action-at-a-distance Coulomb interaction with local interactions mitigated by the electrostatic potential field. The method expands this field in terms of an augmented LCGTO basis which, to assure a proper representation of the field, includes simple Gaussian functions in the interatomic regions in addition to the basis used to expand the wavefunctions near the atoms. In the representation of the *extended* basis $\{i\}$, the familiar electromagnetic variational principle for the electrostatic potential becomes minimization of the quantity

$$E_{\Phi} \equiv \sum_{ij} \Phi_i L_{ij} \Phi_j + e \sum_{\mu\nu i} P_{\mu\nu} [\mu\nu i] \Phi_i, \quad (2.2)$$

where e is the electronic charge, the L_{ij} are the matrix elements of $\frac{\nabla^2}{8\pi}$, and $[\mu\nu i]$ is the overlap of the three basis functions μ , ν and i . The minimum of the sum of this quantity and the usual electronic energy functional then simultaneously gives the correct solution for both the electrostatic potential arising from the electronic charge density and the electronic eigenstates of (2.1) in the corresponding $V^{(dir)}$.

The second phase of the calculation requires only one intermediate quantity, the value of the Hamiltonian acting on the electronic wavefunctions, which may be computed directly once the Hamiltonian is known.

The subsequent orthonormalization of the wavefunctions, however, is more complex and must proceed in terms of a number of intermediate quantities. The LCGTO basis is not orthonormal; consequently, the first step in computing the the overlap between two states, m and n , is to multiply the coefficients of n by the *basis* overlap matrix to form C_n^μ . The overlaps between wavefunctions, $U_{nm} \equiv \langle m | n \rangle$, then become direct inner products between the $C_{m\mu}$ and the C_n^μ . Once U_{nm} is known, the transformation matrix required to bring the wavefunctions into orthonormality, A_{nm} , may be readily computed. (A common *approximation* is to take $A \equiv I_d - \frac{1}{2}(U - I_d)$ where I_d is the identity matrix[7].) The last step in orthonormalization is then to apply A to the wavefunctions.

III. ORGANIZATION OF THE CALCULATION

Assigning each vector unit (VU) to the processing of the basis functions and grid points in a confined region of space best exploits the locality of the calculation. In the LCGTO-LSD method, the basis functions are all centered on the coordinates of nuclei and the integration grids are similarly designed centered on the positions of the atoms. It is most natural to group the processing in terms of the atoms and to assign

some number of atoms to each processing unit. As we shall soon see, the machine is balanced when each of the $N_{VU} = 16,384$ vector units (VU's) is assigned a single atom. To allow a more general discussion we shall assign a atoms to each VU, where a is envisioned to be unity but may vary either greater or less than one.

Typically there are $\gamma = 20(a)$ and $g = 150(a)$ electronic basis functions and exchange-correlation grid points assigned to each VU, respectively. Experience with plane wave calculations shows that a 12 *Ry* cutoff is sufficient to represent the electrostatic potential in *interatomic* regions. This level of accuracy corresponds to an additional 25–30 basis functions in the extended basis set for the electrostatic potential beyond the basis functions representing the electrons. This represents a total of approximately $p = 50(a)$ *electrostatic potential* basis functions per VU.

To process its basis functions, each VU requires, in principle, knowledge of the coefficients of all $B = 20(a)N_{VU}$ elements in the basis. The local nature of the LSD Hamiltonian and the LCGTO basis ensure, however, that this is not necessary in practice. In practice, basis functions with centers more than 6Å apart have matrix elements sufficiently small that their interactions may be ignored[13]. At typical solid state densities this corresponds to a local neighborhood of approximately 50 atoms. Thus, each VU only requires knowledge of the roughly $\Gamma = 20 \times 50 = 1000$ basis functions in its neighborhood and must access only the corresponding $\gamma\Gamma$ elements of the self-consistent Hamiltonian. Note that the value of Γ is fixed by the *physical* range of 6Å and is independent of a . In a similar vein, each VU will periodically require information regarding the $G = 150 \times 50 = 7500$ grid points and $P = 50 \times 50 = 250$ electrostatic basis functions in its “within range” neighborhood. Finally, the calculation will require the calculation of approximately $S = 4(a)N_{VU}$ filled, orthogonal electronic states.

IV. COMPUTATIONAL REQUIREMENTS

Table I lists the per node computation, communication and memory requirements during an iteration for each of the major objects defined in section II above. The table gives the results both in terms of the quantities defined in section III and in terms of time estimates based on a mean VU processing rate of 40 MFLOPS and a mean VU-VU communications rate of $5MW/sec$. Three additional quantities enter Table I beyond those defined in section III: 1) M_3 , the floating point operations count to compute a three basis function overlap; 2) M_f , the operation count to evaluate a basis function at a given point in space; and 3) Σ , the time required to collect a *global* sum of one real number on each vector unit. The first object is on the order of 900 floating point operations, while M_f requires about 50. Commercially available CM-5's perform properly pipelined global sum operations at rate of $10\mu s/value$, a figure which is reducible to $2\mu s/value$ in the enhanced machine. We shall now discuss in detail the derivation of the results for each object in Table I.

A. Evaluation of $P_{\mu\nu}$

Each VU must calculate

$$P_{\mu\nu} \equiv \sum_n C_{n\mu}^* C_{n\nu} \tag{4.1}$$

for μ and ν among the Γ basis functions in range. This will require, first, the communication of the ΓS electronic coefficients from within the 6\AA neighborhood of the VU onto the unit and then $2\Gamma^2 S$ floating point operations. Also, the $P_{\mu\nu}$ will require Γ^2 words of storage.

B. Evaluation of $\rho(r_i)$

With the density matrix computed, the $\rho(r_i)$ are readily computed through

$$\rho(r_i) = \sum_{\mu\nu} P_{\mu\nu} b_\mu(r_i) b_\nu(r_i), \quad (4.2)$$

with μ and ν each taking the Γ values within range and r_i varying over the g points for which the VU is responsible. The values $b_\mu(r_i)$ may be stored in $g\Gamma$ words and are readily computed. These values may then be looked up and (4.2) computed with $3g\Gamma^2$ floating point operations. With all of the needed information compressed into the $P_{\mu\nu}$ and stored on the VU's, no additional communications are necessary to compute this object.

C. Evaluation of $V^{(xc)}$

With the $\rho(r_i)$ known, the exchange-correlation potential is then

$$V_{\mu\nu}^{(xc)} = \sum_i w_i v_{(xc)}(\rho(r_i)) b_\mu(r_i) b_\nu(r_i), \quad (4.3)$$

where $v_{xc}(\rho(r))$ is the standard approximation to the exchange-correlation potential functional[12]. This evaluation requires the communication of G values of v_{xc} in range of each processor as well as $(2M_f + 3)\gamma\Gamma G$ floating point operations, where M_f is the operations count to evaluate the value of a basis function at a particular point in space. Proper ordering of the evaluation may be able to avoid continual re-computation of the $b_\mu(r_i)$ and reduce the computational load of this part of the calculation. The values of $V_{\mu\nu}^{(xc)}$ require $\gamma\Gamma$ words of storage.

D. Evaluation of $\tilde{\rho}_i$

$V^{(dir)}$ is determined by the Φ_i , which come from minimizing E_{Φ} as given in (2.2) by a steepest descents' method which may occur simultaneously with the electronic relaxations. For each VU the first step in determining the gradient of (2.2) with respect to the Φ_i is to compute the overlaps of the charge density with the electrostatic basis,

$$\tilde{\rho}_i = \sum_{\mu\nu} P_{\mu\nu}[\mu\nu i]. \quad (4.4)$$

This must be done for each of the p electrostatic potential grid points for which the VU is responsible. This requires $(M_3 + 2)p\Gamma^2$ floating point operations, where M_3 is the operations count to determine a single three basis function overlap. No communications are required in this phase and the $\tilde{\rho}_i$ require a negligible amount of storage.

E. Evaluation of $L\{\Phi\}$

The second component of the gradient with respect to the Φ_i comes from the application of the Laplace operator L_{ij} defined in (2.2). The application of this operator to Φ requires knowledge of the P values of Φ_i within range of each VU and a total of $2pP$ operations. The needed pP values of L are assumed to be computed once and stored in the local memory, but may be recomputed continually to save memory at very little cost in computation.

F. Evaluation of $V^{(dir)}$

Once the Φ_i have been updated, the matrix elements of the direct potential may be computed as

$$V_{\mu\nu}^{(dir)} = \sum_{\sigma} \Phi_i[\mu\nu i]. \quad (4.5)$$

Calculation of this for the $\gamma \times \Gamma$ matrix elements needed on each node requires transmission of the P values of the updated Φ_i within range and $(M_3 + 2)\gamma\Gamma P$ floating point operations. No communications is required as the needed values of Φ_i were transmitted previously in the communications for calculating $L\{\Phi\}$. The results for $V^{(dir)}$ may be accumulated directly onto the $V_{\mu\nu}^{(xc)}$ above so that no additional storage is necessary.

G. Evaluation of $H\{\psi_n\}$

Once computed, the updated Hamiltonian must be applied to each wavefunction with some fraction of the result added back to the original wavefunction, according either to a conjugate gradient, steepest descents or second order equation of motion approach. Regardless of the technique employed, once $H\{\psi_n\}$ has been evaluated, combination of this object with the wavefunctions is a trivial task requiring on the order of γS operations in each VU. Application of the Hamiltonian to each wavefunction,

$$(\hat{H}C)_{n\mu} = \sum_{\nu} H_{\mu\nu}^{(sc)} C_{n\nu}; \quad (4.6)$$

however, requires $2\gamma\Gamma S$ operations and γS words of storage to keep the result. All needed elements of the Hamiltonian are computed directly in each VU, and the communications requirement for this phase is that each node accesses ΓS coefficients in its neighborhood.

H. Evaluation of C_n^μ

During the calculation, it is imperative to maintain orthonormality among the wavefunctions to keep the Fermionic wavefunctions from collapsing to the ground state. The electronic coefficients $C_{n\mu}$ are vectors of length B which must satisfy the following S^2 (≈ 4 billion) constraints,

$$\begin{aligned} U_{nm} &\equiv \sum_{\mu\nu} C_{m\mu}^* [\mu\nu] C_{n\nu} \\ &= \delta_{nm}, \end{aligned} \quad (4.7)$$

where $[\mu\nu]$ is the overlap between the basis functions μ and ν . Note the presence of the *double* basis sum (over μ and ν) necessitated by the non-orthonormality of the LCGTO basis. The first step in a practical calculation of (4.7) is to compute

$$C_n^\mu \equiv \sum_\nu [\mu\nu] C_{n\nu} \quad (4.8)$$

which requires knowledge of all of the band coefficients in the range, a communication of $S\Gamma$ floating point values to each processor. The computational requirement for this calculation is $2\gamma\Gamma S$ floating point operations. The storage requirements are those of another set of wavefunctions, (γS words), but the space used in accumulating $H\{\psi_n\}$ may be reused for the C_n^μ so that no additional storage is required.

I. Evaluation of the overlap matrix U_{nm}

This is the most demanding part of the entire calculation. Computing

$$U_{nm} = \sum_\mu C_{m\mu}^* C_n^\mu, \quad (4.9)$$

even exploiting the Hermitian symmetry, represents $\frac{1}{2}S^2$ global sums over the entire basis set. The partial sums within the VU's require γS^2 operations concurrent on each node, but then these results must be summed over all VU's. This represents a communications load of $\frac{1}{2}S^2$ global sums of individual values held in each VU. With Σ as the average time to compute such a sum when properly pipelined, the time for forming the global sums will be $\frac{1}{2}S^2\Sigma$. The $S \times S$ matrix U may be stored spread across all VU's requiring $\frac{S^2}{N_{VU}}$ words per VU. The transformation matrix as described in section II ($A \equiv I_d - \frac{1}{2}(U - I_d)$) is then readily computed with a minimal number of concurrent operations.

J. Evaluation of the orthonormalized wavefunctions, $A\{\psi_n\}$

Once A_{nm} has been determined, its S^2 elements are broadcast to all nodes, which then each perform a set of $2\gamma S^2$ floating point operations. (It is most efficient to broadcast batches of data corresponding to entire columns of A .) As it is built up, the result $A\{\psi_n\}$ requires γS words on each processor, but the same space employed

to store $H\{\psi_n\}$ may be reused for this part of the calculation so that no additional storage is needed for the application of A to the $C_{n\mu}$.

V. EVALUATION AND CONCLUSIONS

From Table I we see that the VU memory of $4MW$ is sufficient for calculations of up to 16,384 atoms. Furthermore, because an initial relaxation from random wavefunctions generally requires approximately 50 iterations, the new machine will perform calculations on such extremely large systems within times on the order 200–300 hours, making such calculations quite feasible.

We also find that, even with the enhanced nodes, the calculation is still computationally bound with computations requiring twice the time spent on communications. With only about 40% the machine's time spent communicating, calculations on these large systems will *sustain an average processing rate of about 0.4 Teraflops*. Because the calculation is still computationally bound, the factor of ten improvement in node compute performance is critical in making the calculations feasible; without it, several *months of machine time* would be required for each ionic configuration.

The new machine is not only well balanced for current *ab initio* condensed matter techniques, it will also serve well as the field develops over the next several years because its architecture is well suited to coming developments in the field. Galli *et al.*[14] recently have been developing techniques which reduce the demands of precisely those objects which dominate the computation (P_{uv} , U_{nm} and $A\{\psi_n\}$) and communications (U_{nm} and $A\{\psi_n\}$) of the LCGTO-LSD calculation to the level of the other, less demanding components. This development will maintain the good computation-communication balance and allow iterations on these large systems to be performed in *minutes*. Dynamics calculations could then be carried out over picosecond time scales, sufficiently long to compute *free energies*. Furthermore, to exploit the natural separa-

tion of space into regions of rapidly and slowly varying wavefunctions, practitioners of the *ab initio* condensed matter approach are experimenting with and moving toward wavelet bases[15]. These bases not only yield Hamiltonian matrix elements with a local structure very similar to that in LCGTO-LSD, they also have several great advantages over the LCGTO basis including the ability to systematically complete the basis in mathematically rigorous and well defined manner[16]. Approaches based on these new bases will place very demands on computer architectures similar to those of the LCGTO-LSD method. The new machine will be ready to exploit this cutting edge development to its fullest.

ACKNOWLEDGMENTS

We would like to thank Dennis Salahub for helpful discussions.

REFERENCES

- [1] F. Buda, R. Car and M. Parrinello, *Phys. Rev. B* **41**, 1680 (1990).
- [2] M. Needels, M.C. Payne, J.D. Joannopoulos, *Phys. Rev.* **B38**, 5543 (1988).
- [3] S. Ihara, S.L. Ho, T. Uda and M. Hirao, *Phys. Rev. Lett.* **65**, 1909 (1990).
- [4] M. Needels, J.D. Joannopoulos, Y. Bar-Yam, and S.T. Pantelides, *Phys. Rev.* **B43**, 4208 (1991).
- [5] C.G. Van de Walle, Y. Bar-Yam, S.T. Pantelides, *Phys. Rev. Lett.* **60**, 2761, (1988).
- [6] John W. Negele, "QCD Teraflops Computer," unpublished report, Center for Theoretical Physics, MIT, Cambridge MA.
- [7] See M.C. Payne, M.P. Teter, D.C. Allan, T.A. Arias and J.D. Joannopoulos, *Rev. Mod. Phys.* **64** 1045 (1992) and references therein.
- [8] H. Sambe and R.H. Felton, *J. Chem. Phys.* **62**, 1122 (1975).
- [9] B.I. Dunlap, J.W.D. Connolly, and J.R. Sabin, *J. Chem. Phys.* **71**, 3396 (1979).
- [10] D.R. Salahub in *Applied Quantum Chemistry Symposium*, edited by V. Smith, H.F. Schaefer III, K. Morokuma, Jr. (Reidel, Dordrecht, 1986).
- [11] G. Chiarello, J. Andzelm, R. Fournier, N. Russo, and D.R. Salahub, *Surf. Sci.* **202**, L621 (1988).
- [12] W. Kohn and L. Sham, *Phys. Rev. A* **140**, 1133 (1965).
- [13] D.R. Salahub, private communication.

- [14] Giulia Galli and Michele Parrinello, *Phys. Rev. Lett.* **69**, 3547 (1993).
- [15] P. Lam, K. Cho, T.A. Arias and J.D. Joannopoulos, submitted to *Physical Review Letters* (1993).
- [16] Ingrid Daubechies, *Ten Lectures on Wavelets*, Capital City Press, Montpelier, VT (1992).

TABLES

Table I. Summary of computational, communications and memory requirements of the LCGTO-LSD approach as proposed for execution on the enhanced CM-5. The objects required for the calculation are discussed and defined in section II. The results presented are for a calculation employing all 16,384 vector units in the machine each processing a atoms. (The remaining scaling variables are defined in sections III and IV.) An overall vector unit processing rate of 40 MFLOPS and an average VU-VU communications rate of 5 MW/sec is assumed.

| Object | Computation (per node) | | Communication | | Memory |
|-------------------------|---|--------------|------------------------|----------------------|-----------------------------------|
| | Operations | Time (min) | Words | Time (min) | |
| A) P_{uv} | $2\Gamma^2 S$ | $55(a)$ | ΓS | $0.22(a)$ | $\Gamma^2 = 1MW$ use[G] |
| B) $\rho(r_i)$ | $3g\Gamma^2$ | $0.19(a)$ | data in unit | | $g\Gamma = 150KW(a)$ use[I] |
| C) $V_{\mu\nu}^{(xc)}$ | $(2M_f + 4)\gamma\Gamma G$ | $6.5(a)$ | G | 2.5×10^{-5} | $\gamma\Gamma = 20KW(a)$ |
| D) $\tilde{\rho}_i$ | $(M_3 + 2)p\Gamma^2$ | $19(a)$ | data in unit | | $p < 1KW(a)$ |
| E) $L\{\Phi\}$ | $2pP$ | $10^{-4}(a)$ | P | 8.3×10^{-6} | $pP = 125KW(a)$ |
| F) $V_{\mu\nu}^{(dir)}$ | $(M_3 + 2)\gamma\Gamma P$ | $19(a)$ | P | 8.3×10^{-6} | $\gamma\Gamma$ use[C] |
| G) $H\{\psi_n\}$ | $2\gamma\Gamma S$ | $1.1(a^2)$ | ΓS | $0.22(a^2)$ | $\gamma S = 1.3MW(a)$ |
| H) C_n^μ | $2\gamma\Gamma S$ | $1.1(a^2)$ | ΓS | $0.22(a)$ | γS use[G] |
| I) U_{nm} | γS^2 | $36(a^3)$ | $\frac{1}{2}S^2\Sigma$ | $72(a^2)$ | $\frac{S^2}{N_{VV}} = 260KW(a^2)$ |
| J) $A\{\psi_n\}$ | $2\gamma S^2$ | $72(a^3)$ | S^2 | $14(a^2)$ | γS use[G] |
| $H_{\mu\nu}^{(1)}$ | Computed once and <i>saved</i> in unit | | | | $\gamma\Gamma = 20KW(a)$ |
| $\{\psi_n\}$ | Requires <i>uncorrupted</i> storage during each iteration | | | | $\gamma S = 1.3MW(a^2)$ |
| Total | $210\min(\frac{a}{2} + \frac{a^3}{2})$ | | $87\min(a^2)$ | | $3.2MW(a^2)$ |