Data Reduction Schemes in Davidson Subspace Diagonalization for MR-CI

Wilfried N. Gansterer[1], Wolfgang Kreuzer[1,⋆], and Hans Lischka[2,⋆]

[1] Institute of Distributed and Multimedia Systems, University of Vienna
[2] Institute for Theoretical Chemistry, University of Vienna

Abstract. In this paper, we investigate several data reduction schemes to improve the computational efficiency in the multi reference configuration interaction (MR-CI) method, one of the main quantum chemical approaches for solving the electronic Schrödinger equation. The basic idea is to take advantage of the often relatively low accuracy requirements on the solution of the resulting large eigenvalue problem, whose dimension may reach several hundred millions or even more. We will discuss some approaches to reduce the amount of data to be accessed and to be transferred within the Davidson subspace diagonalization method. We also show experimental results achieved with the COLUMBUS code.

1 Introduction

Quantum chemical methods provide very important procedures for the computation of molecular properties and for the computer simulations of chemical reactions which cannot be solved exactly. They are based on the electronic Schrödinger equation; a complicated many-particle differential equation, which cannot be solved analytically. Therefore, numerical approximations have to be used. These are usually very involved, extremely time consuming and require very large amounts of data flow.

In this paper, we focus on the *multireference configuration interaction* (*MR-CI*) approach [9]. It allows for accurate calculations of molecular systems based on the original many-particle Schrödinger equation. The MR-CI method is especially important in "difficult" cases, e. g., for the calculation of dissociation processes and electronically excited states.

A basis expansion of the many-particle wave function leads to an eigenvalue problem of enormous size. Although the Hamiltonian matrix H of this eigenproblem is very sparse, its dimension n can easily reach several hundred millions or even billions. H tends to be diagonally dominant and usually very few eigenpairs (or only a single one) need to be computed.

Given this setup, the Davidson method is a suitable approach for solving problems of this type numerically. However, in contrast to the matrix H, the subspace vectors and the Ritz vectors arising in this process are in general *dense*, and not sparse. Due to their enormous dimension, simply storing and handling

⋆ Supported by the project P14817-N03 of the Austrian Science Fund FWF.

I. Lirkov, S. Margenov, and J. Waśniewski (Eds.): LSSC 2005, LNCS 3743, pp. 564–571, 2006.

those vectors becomes a major performance bottleneck. This is especially severe when solving the eigenproblem in parallel (which is required for large problems of interest), because intensive memory accesses and transfer of huge amounts of data leads to an often prohibitive communication overhead. The data reduction approach developed in this paper will help to overcome these problems.

Related Work. In earlier work, Dachsel and Lischka [2] have developed first ideas for a data compression approach in the Davidson subspace diagonalization. They propose bitwise data reduction in the subspace vectors based on an estimator for the resulting error. Their approach assumes diagonal dominance of the matrix H. Alternative approaches based on fixed-point truncation schemes have been developed by Harrison and Handy [4], Knowles [5], and Olson [8].

In this paper, we analyze and extend the approach of Dachsel and Lischka, we compare it with various newly developed error estimators, and also discuss aspects related to the implementation of such schemes. Moreover, based on numerical experiments with the COLUMBUS code [6, 7], we also give some insight into how much data reduction can be achieved in practical situations.

2 Multi Resolution Configuration Interaction

The stationary, nonrelativistic, clamped-nuclei, electronic Schrödinger equation is given as

$$\mathcal{H}\Psi = E\Psi$$

with the total Hamiltonian $\mathcal{H} = \sum_i h_i + \sum_{i<j} g_{ij} + V_{kk}$, where h_i is the one-electron operator for electron i containing the kinetic energy and Coulomb attraction, g_{ij} is the electron-electron repulsion term, V_{kk} is the nuclear repulsion and Ψ is the many-electron wave function.

Expanding Ψ into a many-electron basis (configurations state functions $\{\Phi_i\}$) and applying the Ritz variational principle leads to the matrix eigenproblem

$$Hc = Ec \tag{1}$$

with symmetric $H \in \mathbb{R}^{n\times n}$, where n tends to be extremely large (several hundred millions up to a billion). H is sparse, but it is prohibitively costly (in terms of computation as well as in terms of storage requirements) to construct it explicitly ($H_{st} = \langle\Phi_s|\mathcal{H}|\Phi_t\rangle$). Usually, the lowest eigenpair (or a few of the lowest eigenpairs) of H need to be computed.

2.1 Davidson Subspace Diagonalization

Given the properties of the eigenproblem (1), *Davidson's method* [3, 1] is a suitable approach for solving it. The basic structure of this method is shown in Algorithm 1. Starting with an initial vector v_0 or subspace V_0, a basis for a (small) subspace V is constructed, in which approximations $(\bar{E}, u)$ of the desired eigenpairs (E, c) can be computed cheaply (r denotes the associated residual).

Davidson [3] suggested to expand this basis in each iteration by the solution of the correction equation shown in Algorithm 1. This procedure is very successful for strongly diagonally dominant H.

Algorithm 1. Basic structure of Davidson subspace diagonalization

$V = V_0, W = HV_0, v = v_0$
repeat
  $w = Hv$, $V = [V, v], W = [W, w]$
  $\bar{H} = V^\top HV = V^\top W$ ... projected matrix
  $[\bar{E}, \alpha] = \mathrm{eig}(\bar{H}, V^\top V)$ ... Ritz value and Ritz vector
  $u = V\alpha$, $r = Hu - \bar{E}u$
  solve $\left[\mathrm{diag}(H) - \bar{E}\right] v = r$ for v ... correction equation
  reduce V, W if necessary
until $|r| < tol$
return $\bar{E} \approx E, u \approx c$

2.2 Limitations

Davidson's method requires the storage and retrieval of the n-dimensional vectors spanning the subspaces V and W. Due to their size, these vectors are either stored on external disk devices or—in particular for parallel calculations—distributed over the memory of the individual nodes. Storage and retrieval of such large amounts of data constitutes a serious communication bandwidth bottleneck. Moreover, the subspace vectors and their memory requirements are the origin of serious limitations in parallel calculations. Although I/O requirements have been reduced significantly [10], the problem still remains. Therefore, reduction of the memory requirements for storage of the subspace vectors by means of data reduction schemes as described in this paper will lead to substantial alleviation of the bandwidth problem and, in parallel computations, will free memory needed for the local computation of the matrix-vector product Hv.

3 Data Reduction

The basic idea investigated in this paper is the following: Based on an accuracy tolerance τ determined by the user, reduce the amount of data to be stored for the v and w vectors in Davidson's method. The larger τ the more significant the data reduction. This reduction of data to be stored obviously also corresponds to a reduction of data accesses and transfers, which is especially important in parallel calculations.

In order to guarantee satisfactory accuracy of the computed spectral information of H, we develop error estimators and bounds for controlling the errors in E and c due to suggested perturbations in v and w. The main focus in this paper is on controlling the error in the energy E, analysis of the corresponding error in c is the subject of ongoing work.

3.1 Error Estimators and Bounds

One approach for deriving estimators of the effect of the data reduction on the computed energy is to examine the difference $\Delta\bar{E} := \bar{E} - E$. By definition,

$$\Delta\bar{E} = \frac{u^\top H u}{u^\top u} - \frac{c^\top H c}{c^\top c}.$$

Neglecting third and higher order terms in the Taylor series expansion with respect to u around c ($u = c + \Delta u$) leads to (first order terms are zero!)

$$\Delta\bar{E} \approx \Delta u^\top (H - EI)\Delta u\ . \tag{2}$$

From this relation, several estimators can be derived. In the following, we use the notation $e := (1, \ldots, 1)$ and for a matrix A we define $|A|$ by taking absolute values elementwise, that is, $|A|_{ij} := |A_{ij}|$.

Dachsel and Lischka [2] have proceeded by replacing the unknown eigenvalue E by the Ritz value $\bar{E}$ and the matrix H by its diagonal D in (2) which yields

$$\Delta\bar{E} \approx \mathrm{Est}_{DaLi} := \Delta u^\top (D - \bar{E}I)\Delta u = \sum_i \sum_j (D_i - \bar{E}I)\Delta u_i \Delta u_j$$

With an upper bound on the componentwise perturbation in the vector u, $|\Delta u_i| \le \beta$, this leads to the bound

$$|\Delta\bar{E}| \approx |\mathrm{Est}_{DaLi}| \le \beta^2 \sum_i \left|D_i - \bar{E}I\right|$$

Asking for $\Delta\bar{E} < \tau$ consequently corresponds to requiring

$$\beta < \sqrt{\frac{\tau}{e^\top \left|D - \bar{E}I\right| e}}\ . \tag{3}$$

Alternatives. The estimator constructed in [2] relies on two approximations: E is replaced by the Ritz value $\bar{E}$ in (2) (which may not always be appropriate because this is precisely the error to be estimated), and H is replaced by its diagonal D (which requires diagonally dominant H). Pursuing a more general approach, we can rewrite (2) as

$$\Delta\bar{E} \approx \Delta u^\top \left(H - \left(\bar{E} - \Delta\bar{E}\right) I\right) \Delta u,$$

leading to the estimator Est_H

$$\Delta\bar{E} \approx \mathrm{Est}_H := \frac{\Delta u^\top (H - \bar{E}I)\Delta u}{1 - \Delta u^\top \Delta u}\ . \tag{4}$$

In the following we assume that $1 - \Delta u^\top \Delta u > 0$ which implies $|\Delta u_i| < \sqrt{1/n}$. With this additional constraint we derive analogously to before an upper bound for the allowable componentwise perturbation in u:

$$\beta < \sqrt{\frac{\tau}{e^\top \left|H - \bar{E}I\right| e + n\tau}}\ . \tag{5}$$

If the matrix H happens to be diagonally dominant (as it tends to be the case in the MR-CI context), we can derive an estimator Est_D by replacing H with its diagonal D in (4), leading to the bound

$$\beta < \sqrt{\frac{\tau}{e^\top \left|D - \bar{E}I\right| e + n\tau}} \; . \tag{6}$$

Finally, we describe a fourth approach. Using the fact that due to the construction of the Davidson eigenvectors from a generalized eigenvalue problem $||u|| = 1$, we have

$$|\Delta\bar{E}| = |c^\top H c - u^\top H u| = |(u + \Delta u)^\top H (u + \Delta u) - u^\top H u| \; . \tag{7}$$

Based on (7) we require

$$\left|\Delta\bar{E}\right| = |\mathrm{Est}_q| := \left|2\Delta u^\top H u + \Delta u^\top H \Delta u\right| < \tau \; .$$

Analogously to before, this corresponds to requiring

$$\beta < -\frac{e^\top |Hu|}{e^\top |H| e} + \sqrt{\left(\frac{e^\top |Hu|}{e^\top |H| e}\right)^2 + \frac{\tau}{e^\top |H| e}}. \tag{8}$$

3.2 Implementation

Several aspects related to the efficient implementation of the concepts developed so far deserve some more detailed discussion.

Evaluation of Bounds. For performance reasons, the Hamiltonian matrix H is not constructed explicitly in the context of MR-CI calculations. For the same reasons, we introduce some additional simplifications in order to reduce the computational effort for evaluating the estimators and bounds derived before.

Our formulation of the bounds (3), (5), (6), and (8) is based on applying the triangle inequality, which leads to expressions with sums of absolute values to be evaluated in every iteration (because $\bar{E}$ changes). To further improve performance, we approximate the sums of absolute values by the absolute value of the sums in those bounds. Based on this simplification, we have to evaluate these sums only *once* during the entire process, and information from the huge matrix H only needs to be accessed once. This approximation tends to work well in our context since we need to approximate the *smallest* eigenvalue(s) E of H. Due to the fact that almost all diagonal entries D_i of H tend to have the same sign and due to the diagonal dominance of H, $\sum_i |D_i - \bar{E}| \approx |\sum_i D_i - \bar{E}|$.

Implementation of Data Reduction. The information about allowable perturbations in the u vectors contained in the bounds (3), (5), (6), and (8) can be translated into *bitwise* reductions in the mantissas of the entries in the v and w vectors as described in [2]. In this paper, we compare this approach with an

elementwise data reduction strategy, where vector entries with absolute value below a certain threshold (determined by the bounds) are eliminated, whereas all entries above this threshold are stored with full precision.

The advantage of elementwise reduction is that it can be implemented more easily and more efficiently, the disadvantage is that the data reduction achieved is always lower than or at best equal to that of bitwise reduction. Experimental evidence presented in Section 4 shows that the memory requirements with elementwise data reduction can actually be significantly higher than with bitwise data reduction.

4 Experiments

In this section we present some experimental evidence that the amount of data reduction achieved in practice is significant.

Test Case. We evaluated the effect of data reduction with the different estimators and bounds (3), (5), (6), (8) for a "small" but representative test problem of dimension $n = 14558$ computing the MR-CI energy and wave function for the Ethylene molecule C_2H_4.

The calculations were performed with the COLUMBUS program system, a collection of programs for high-level ab initio molecular electronic structure calculations [6, 7]. Double precision accuracy, i. e., a mantissa length of 52 bits, was used, the threshold for the stopping criterium for the Davidson subspace method was set to tol $= 10^{-4}$, and the tolerance τ for the data reduction was set to $\tau = 10^{-8}$.

Convergence History. Fig. 1 illustrates that the artificial perturbations introduced in the vectors v and w of Davidson's method hardly influence its convergence behavior. In our test case, only one additional iteration step is required until convergence for all versions of data reduction. Moreover, it can be observed that for the test case considered the various estimators and bounds hardly differ.

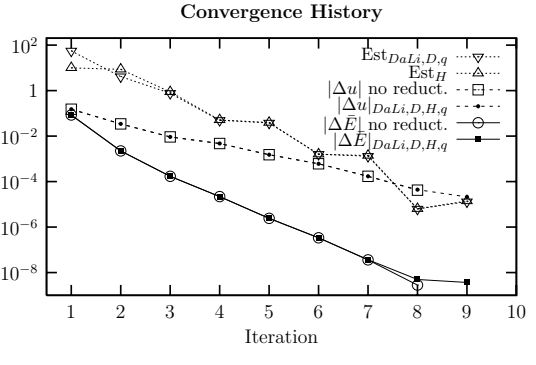

Fig. 1. Convergence history of Davidson's method with data reduction (based on various estimators and bounds) as well as without data reduction

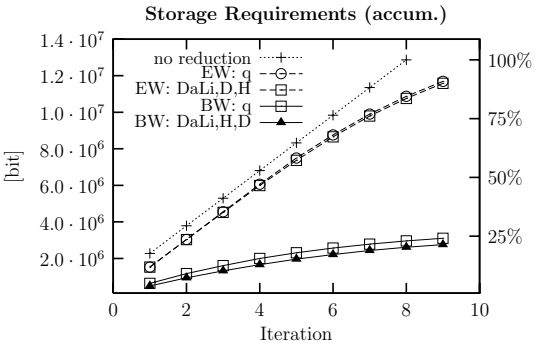

Fig. 2. Accumulative representation of the number of bits to be stored for different variants of data reduction (elementwise "EW", bitwise "BW", various estimators and bounds) as well as without data reduction

Storage Savings. In this example, we focus on achieving maximum storage efficiency, which means that per vector only one additional compression factor, representing the maximum difference between a vector entry and its value after data reduction, is stored. The actual length of the mantisse of each vector entry can then be calculated using this factor and the exponent of the entry. The starting vector for the Davidson procedure is created by diagonalizing a reference subspace. Most of the entries in this vector are zero. Therefore, this vector is submitted to the data reduction process at the beginning of the calculation as well.

Fig. 2 illustrates very high savings in storage requirements from bitwise data reduction. In contrast, for this test case the benefits of elementwise data reduction are much smaller.

5 Conclusions

A data reduction strategy for reducing the data transfer during the solution of the high dimensional eigenvalue problem arising in the MR-CI method for solving the electronic Schrödinger equation has been investigated. It is based on the estimation of the error in the eigenvalues resulting from componentwise perturbations in the subspace vectors. We have evaluated an earlier data compression scheme introduced by Dachsel and Lischka [2] and compared it to new alternatives based on new error estimators. Using the COLUMBUS program package, the following observations could be made for a realistic test case: (*i*) the data reduction achieved hardly depends on the specific error estimator used, (*ii*) with bitwise data reduction, the amount of data to be handled can be reduced to about one fourth, and (*iii*) for the test case considered, bitwise data reduction is much more attractive than elementwise reduction.

Current and future research investigates how this data reduction is best translated into performance gains in terms of actual runtime reductions. First experiences indicate that even for relatively fast communication, runtime reductions

may be achieved if data reduction is beyond (roughly) 50% (even earlier for slower communication). We need to point out that the potential benefits from our technique presented here do not only come from reducing runtimes, but, sometimes even more important, from making the solution of problems feasible which could not be handled before due to storage restrictions. Nevertheless, we will also investigate alternative (and more efficient) implementation strategies beyond the elementwise and bitwise reduction discussed here.

References

1. Crouzeix, M., Philippe, B., Sadkane, M.: The Davidson method. SIAM J. Sci. Comput. **15** (1994) 62–76
2. Dachsel, H., Lischka, H.: An efficient data compression method for the Davidson subspace diagonalization scheme. Theor. Chim. Acta **92** (1995) 339–349
3. Davidson, E.R.: The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real symmetric matrices. J. Comput. Phys. **17** (1975) 87–94
4. Harrison, R.J., Handy, N.C.: Full CI calculations on BH, H_2O, NH_3 and HF. Chem. Phys. Lett. (1983) 386–391
5. Knowles, P.J.: Very large configuration interaction calculations. Chem. Phys. Lett. (1989) 513–517
6. Lischka, H., Shepard, R., Pitzer, R.M., Shavitt, I., Dallos, M., Müller, T., Szalay, P.G., Seth, M., Kedziora, G.S., Yabushita, S., Zhang, Z.: High-level multireference methods in the quantum-chemistry program system COLUMBUS: Analytic MR-CISD and MR-AQCC gradients and MR-AQCC-LRT for excited states, GUGA spin-orbit CI and parallel CI density. Phys. Chem. Chem. Phys. **3** (2001) 664–673
7. Lischka, H., Shepard, R., Shavitt, I., Pitzer, R.M., Dallos, M., Müller, T., Szalay, P.G., Brown, F.B., Ahlrichs, R., Böhm, H.J., Chang, A., Comeau, D.C., Gdanitz, R., Dachsel, H., Ehrhardt, C., Ernzerhof, M., Höchtl, P., Irle, S., Kedziora, G., Kovar, T., Parasuk, V., Pepper, M.J.M., Scharf, P., Schiffer, H., Schindler, M., Schüler, M., Seth, M., Stahlberg, E.A., Zhao, J.G., Yabushita, S., Zhang, Z.: COLUMBUS, an ab initio electronic structure program, release 5.9 (2004)
8. Olson, J., Jørgensen, P., Simons, J.: Passing the one-billion limit in full configuration-interaction (FCI) calculations. Chem. Phys. Lett. (1990) 463–472
9. Shavitt, I.: The method of configuration interaction. In: Methods of Electronic Structure Theory. H. F. Schaefer, III, ed. Plenum Press, New York (1977) 189–275
10. Shepard, R., Shavitt, I., Lischka, H.: Reducing I/O costs for the eigenvalue procedure in large-scale configuration interaction calculations. J. Comp. Chem. **23** (2002) 1121–1125