# On the Use of Incomplete LU Decomposition as a Preconditioning Technique for Density Fitting in Electronic Structure Computations

Rui Yang<sup>1</sup>, Alistair P. Rendell<sup>1</sup>, and Michael J. Frisch<sup>2</sup>

<sup>1</sup> Department of Computer ScienceCollege of Engineering and Computer Science The Australian National University, Canberra, ACT 0200, Australia <sup>2</sup> Gaussian Inc.340 Quinnipiac St Bldg 40 Wallingford, CT 06492 USA {Rui.Yang, Alistair.Rendell}@anu.edu.au, frisch@gaussian.com

Abstract. Incomplete factorization preconditioners combined with Krylov subspace accelerators are currently among the most effective methods for iteratively solving large systems of linear equations. In this paper we consider the use of a dual threshold incomplete LU factorization (ILUT) preconditioner for the iterative solution of the linear equation systems encountered when performing electronic structure calculations that involve density fitting. Two questions are addressed, how the overall performance of the ILUT method varies as a function of the accuracy of the preconditioning matrix, and whether it is possible to make approximations to the original matrix on which the LU decomposition is based and still obtain a good preconditioner. With respect to overall performance both computational and memory storage requirements are considered, while in terms of approximations both those based on numerical and physical arguments are considered. The results indicate that under the right circumstances the ILUT method is superior to fully direct approaches such as singular value decomposition.

**Keywords:** ILUT preconditioning, Krylov subspace method, electronic structure calculation, density fitting.

# **1** Introduction

In computational science we are frequently required to solve systems of linear equations of the form:

$$Ax = b \tag{1}$$

where A and b are respectively a matrix and vector of known values, while x is a vector the values of which we wish to determine. Although there are a variety of approaches for solving such problems, if the dimension of the problem is large, and particularly if matrix A is sparse, then it is common to use iterative approaches such as the Krylov subspace method [1]. In these methods the algorithm proceeds by

essentially guessing an initial form for x, and then refining it through a series of iterative updates. To improve the efficiency and robustness of this procedure a number of preconditioning techniques have been proposed [2]. One such technique, that will be considered here, is the dual-dropping incomplete LU factorization technique (ILUT) [3].

The particular systems of linear equations that are of interest to us are those that arise when using an auxiliary basis set to fit the electronic density in electronic structure calculations. Specifically, in many implementations of Kohn-Sham density functional theory (KS-DFT) the electronic density ( $\rho(r)$ ) is expressed in terms of a product of one-particle atom-centered basis functions ( $\mu(r)$  and  $\nu(r)$ ):

$$\rho(r) = \sum_{\mu\nu}^{N} D_{\mu\nu} \mu(r) \nu(r)$$
<sup>(2)</sup>

where  $D_{\mu\nu}$  is an element of the density matrix and there are a total of N functions in the orbital basis set. Within this representation the total Coulomb energy  $(E_J)$  is given by:

$$E_{J} = \frac{1}{2} \int dr_{1} \int dr_{2} \frac{\rho(r_{1})\rho(r_{2})}{r_{12}} = \frac{1}{2} \sum_{\mu\nu\lambda\sigma}^{N} D_{\mu\nu} D_{\lambda\sigma} \left(\mu\nu|\lambda\sigma\right)$$
(3)

where  $(\mu \eta \lambda \sigma)$  are the two-electron repulsion integrals (ERI). Formally evaluation of Eqn. (3) scales as the fourth power of the number of basis functions  $(O(N^4))$ , however, if the density is expanded in terms of auxiliary basis set:

$$\tilde{\rho}(r) = \sum_{\alpha} c_{\alpha} \alpha(r) \tag{4}$$

this drops to  $O(N^2)$ , albeit  $O(N^2)$  where N is now the number of functions in the auxiliary basis set. It is in the evaluation of these fitting coefficients ( $c_{\alpha}$ ) that it is necessary to solve a set of linear equations.

With respect to Eqn. (1), the elements of matrix A represent Coulomb integrals between two auxiliary fitting basis functions, x the expansion coefficients  $(c_{\alpha})$ , and the elements of b correspond to the Coulomb potential in the auxiliary basis set generated by the electron density as expanded by the density matrix. Solving this system of linear equations is problematic in that the dimension of A can become quite large - in the order of ten thousand - making it both hard to store in memory and computationally expensive to solve using direct techniques such as singular value decomposition (SVD). (SVD is used since A is often ill-conditioned, reflecting near linear dependencies in the fitting basis set).

As the name suggests ILUT performs an approximate LU factorization of matrix A. The accuracy of this factorization is controlled by two parameters,  $\tau$  and p and is denoted as ILUT( $\tau$ , p). Parameter  $\tau$  serves as a threshold for the magnitude of entries retained in the LU factorization, while parameter p limits the maximum number of

non-zero entries retained in any given row of the factored matrix. Thus, while  $\tau$  provides no control over the memory required to store the LU factorization, p can be used to limit memory usage. In the limit that  $\tau \rightarrow 0.0$  and  $p \rightarrow n$  (where n is the dimension of matrix A) the LU factorization is exact, and the preconditioning step will solve the real problem. Conversely as  $\tau$  and p move away from these extremes preconditioning becomes ever more approximate resulting in larger number of subspace iterations. In this work we use the ILUT preconditioner approach of Saad [3] combined with the iterative Generalized Minimal Residual subspace method (GMRES) [4].

This paper seeks to explore two inter-related issues:

- 1. Given an exact representation of matrix *A*, can ILUT preconditioning be used to substantially speed-up the time taken to solve the linear equation system required when using density fitting?
- 2. As the elements of *A* represent Coulomb interactions that decay with distance, is it possible to use either a numerical threshold or chemical knowledge to construct a sparse approximation to matrix *A* from which it is possible to derive a good preconditioning matrix?

Finally, we note that while we have introduced density fitting in the context of auxiliary basis sets for performing KS-DFT calculations, density fitting also offers significant advantages for multi-configurational SCF (MC-SCF) [5-10], second order Møller-Plesset perturbation theory (MP2) [11-15], coupled cluster methods [16-19] and more recently, explicitly correlated MP2-R12 [20-22] calculations. Thus the work undertaken here has widespread applicability.

In the following sections we first describe the ILUT preconditioning technique and density fitting problem in general, before exploring the use of ILUT preconditioning to solve the density fitting problem for a variety of test cases. Conclusions and general discussion are given in section 5.

# 2 ILUT Preconditioning

Incomplete factorization preconditioners combined with Krylov subspace accelerators are currently among the most effective iterative techniques for solving large, sparse irregularly structured linear systems of equations [23]. The incomplete factorization technique involves a decomposition of the form A=LU-R=M-R where L and U obey the specific non-zero pattern P, and R is the residual of the preconditioning matrix M. If P has the same non-zero pattern as A, the LU decomposition is referred to as ILU(0). That is, the L and U matrices have the same non-zero structure as the lower and upper parts of A, respectively, with drop-offs in the LU decomposition depending only on the structure of A without considering the numerical values in LU decomposition are dropped based on their values rather than their locations [3].

As mentioned above the dual-dropping ILUT approach has two parameters: a threshold drop tolerance  $(\tau)$ , and a fill number (p) that specifies what fraction of the factorization is kept. Ideally, these two parameters should be chosen to balance the ILUT construction time with the iterative processing time. The basic ILUT algorithm is shown below:

```
Algorithm 2.1. ILUT(\tau, p):
              For a N×N dimension matrix A,
              Do i=1,.., N
              Step 1: Read in the i th row elements of A into
{w};
                        Do j=1, i-1
                              w_i = w_i / a_{ii}
                              Step 2: Applying a dropping rule to
Wi
                              Step 3: If w_i \neq 0 Then
                                          Do k=j+1, N
                                           W_{k} = W_{k} - W_{j} \cdot u_{jk}
                                          End Do
                                        End If
                        End Do
              Step 4: Applying a dropping rule to {w}
              Step 5: l_{j,j} = w_j for j = 1, ..., i - 1
                       u<sub>iii</sub>=w<sub>i</sub> for j=i,...,N
              Reset {w}=0
              End Do
```

At Steps 2 and 5, all entries with a magnitude less than  $\tau$  multiplied by the norm of the current row are dropped. Furthermore, at Step 5, only the largest *p* entries in each row of the L and U factorization are retained (in additional to the diagonal elements). Thus *p* is a parameter that helps control memory usage, while  $\tau$  also helps to reduce the computational cost. In the work presented here we have set *p* equal to the dimension of the problem, so that the accuracy of the preconditioner is determined solely by the parameter  $\tau$ . A small value of  $\tau$  implies a more accurate preconditioner and fewer Krylov iterations, but the preconditioner will be more expensive to construct. While a large value of  $\tau$  has the opposite effect.

At each preconditioning iterative step, a linear system of the form Me = r is solved, where M is the preconditioner that approximates A, r is the residual of the current iteration and e is the correction vector. The preconditioning can be applied to the left or right of the original linear equation system or in split forms, although the general consensus is that such variations make relatively little difference [24]. In this work initial tests using right and left side preconditioning supported this view. For the results presented here right side preconditioning is used with GMRES and ILUT routines that are derived from the SLATEC [25] and SPARSKIT [26] libraries respectively. For further details of the ILUT preconditioning process the reader is referred to Ref. [3, 24].

## **3** Density Fitting

The error in the fitted density (Eqn. (4)) for a two electron projection operator  $\omega_{12}$  is defined as:

$$\Delta \omega = \left(\rho - \tilde{\rho} |\omega_{12}|\rho - \tilde{\rho}\right) = \left(\rho |\omega_{12}|\rho\right) - 2c_{\alpha}\left(\rho |\omega_{12}|\alpha\right) + c_{\alpha}\left(\alpha |\omega_{12}|\beta\right)c_{\beta}$$
(5)

where  $c_{\alpha}$  are the fitting coefficients. Differentiating with respect to  $c_{\alpha}$  and minimizing gives rise to:

$$\frac{\partial \Delta \omega}{\partial c_{\alpha}} = -2(\rho |\omega_{12}|\alpha) + 2(\alpha |\omega_{12}|\beta)c_{\beta} = 0$$
<sup>(6)</sup>

which can be written as a set of linear equations of the form:

$$\sum_{\beta} A_{\alpha\beta} x_{\beta} = b_{\alpha} \tag{7}$$

with  $A_{\alpha\beta} = (\alpha | \omega_{12} | \beta), x_{\beta} = c_{\beta}$  and  $b_{\alpha} = D_{\nu,\mu} (\nu \mu | \omega_{12} | \alpha).$ 

Although there are a number of possibilities for the two electron projection operator  $\omega_{12}$ , it is widely acknowledged that use of the Coulomb operator gives the best results for energy evaluations [27,28]. Using this operator the linear system given in Eqn. (7) involves the following three-center and two-center repulsion integrals:

$$\sum_{\nu,\mu} D_{\nu,\mu} (\nu \mu | \alpha) = \sum_{\beta} (\alpha | \beta) c_{\beta}$$
(8)

where 
$$(\nu\mu|\alpha) \equiv \int dr_1 \int dr_2 \frac{\nu(1)\mu(1)\alpha^*(2)}{r_{12}}$$
 and  $(\alpha|\beta) \equiv \int dr_1 \int dr_2 \frac{\alpha(1)\beta^*(2)}{r_{12}}$ .

In practice solution of the linear equation system is a little more complex than suggested above, since the expansion coefficients must be constrained so that the total charge is constant, i.e.

$$\sum_{\alpha} c_{\alpha} S_{\alpha} = n \tag{9}$$

where n is the total charge and  $S_{\alpha} = \int \alpha(r) d^3 r$ . Practical implementation of this

constraint requires an extra orthogonalization step in each GMRES step. Furthermore, as matrix A is not positive definite (but semi-definite), the diagonal elements are scaled to modify its condition number.

In principle, matrix A is dense owing to the long tail of the Coulomb interactions involved in computing each element of this matrix. In practice, however, if two functions are well separated we would expect the value of the corresponding element of A to be relatively small, and therefore a numerical threshold could be used to determine whether it should be kept. Alternatively, since the fitting functions are normally chosen to be atom centered, it may be possible to construct a sparse representation of A based on knowledge of the atoms in the system. One obvious approach is to consider a sparse representation of A where the only non-zero elements involve those interactions between fitting functions that are located on the same atomic centre. This has the effect of producing a block diagonal representation of A. Less dramatic approximations might be based on including all interactions between fitting functions that are within the same functional group.

Finally, it should be noted that the KS-DFT method is in itself iterative, involving an initial guess of the density matrix  $(D_{\mu\nu})$  that is refined during each iteration until a "self-consistent field" (SCF) is reached. Within this process the density fitting equations must be solved at each iteration of the SCF procedure, but because the location of the fitting functions does not change from iteration to iteration the ILUT representation of matrix *A* remains the same for all SCF iterations. What does change, however, is the value of the *b* vector which must be re-evaluated at each SCF iteration. (This vector changes as it involves a contraction of the current guess for the density matrix with the relevant 3-center integrals.) The implication of this is that an ILUT factorization of *A* can be done just once before the start of the SCF procedure, and then used during every SCF iterations to improve performance when solving to fit the current density.

## 4 Numerical Results and Discussion

To explore the performance of the ILUT method for density fitting calculations four different computations were considered:

System 1: is a zeolite fragment ( $Si_8O_7H_{18}$ ) containing 33 atoms, and utilizing a 6-31g\* basis set. The fitting basis contains 1489 functions and is obtained using the scheme implemented in Gaussian 03 to automatically generate fitting basis sets [29, 30]. 14 iterations are required to converge the SCF.

System 2: this is identical to system 1, but employs a larger cc-pVDZ basis set that gives rise to 2048 automatically generated fitting basis functions. 13 iterations are required to converge the SCF.

System 3: is a Valinomycin molecule  $(C_{54}H_{90}N_6O_{18})$  containing 168 atoms, and utilizing a 3-21g basis set. This gives rise 3018 fitting functions. 11 iterations are required to converge the SCF.

System 4: is identical to system 3, but uses a 6-31g basis set. This gives rise to 4182 fitting functions. 11 iterations are required to converge the SCF.

All calculations were performed on the 900 MHz SPARC v9 processor running Sun Solaris 10 with code compiled using Sun Studio 11.

#### 4.1 Approximating the LU Decomposition of A

In this section we consider the case when matrix A is exact, but the accuracy of the LU decomposition is varied by changing parameter  $\tau$ . We allowed a maximum of 2000 Kyrlov iterations, and assume convergence to be satisfied when the 2-norm residual is reduced by a factor of 109. There appears to be no generally applicable guidelines for choosing a value for parameter  $\tau$ . We choose several sample values for  $\tau$  of 10<sup>-2</sup>, 10<sup>-3</sup>, 10<sup>-4</sup>, 10<sup>-5</sup> and 10<sup>-10</sup>, and consider the value of 10<sup>-10</sup> as corresponding to a complete LU decomposition. The sparsities of the preconditioner for the four different test systems and the five different values for  $\tau$  are given in Table 1. These results show that even with a value for  $\tau$  of 10<sup>-2</sup> the LU decomposition contains roughly 50% non-zero elements for all the systems considered. And that if the system size is held constant while the fitting basis is expanded, the sparsity decreases even further (i.e. in going from system 1 to system 2, or system 3 to system 4). These results might be expected since the fill-in that occurs during the ILUT process is controlled only by the numerical value of the fill-in, not by whether there is a non-zero element in the same location in the original matrix A.

τ	System 1	System 2	System 3	System 4
10-10	1.000	1.000	1.000	1.000
$10^{-5}$	0.868	0.886	0.904	0.905
$10^{-4}$	0.761	0.781	0.782	0.789
$10^{-3}$	0.608	0.647	0.608	0.627
$10^{-2}$	0.476	0.627	0.471	0.572

Table 1. Sparsity of the ILUT preconditioner with different  $\tau$  values for all studied systems

We now consider the overall performance of the ILUT method, and in particular the influence of  $\tau$  on performance. As was discussed in section 3, there are two aspects to using ILUT with density fitting. The first involves the incomplete factorization of the *A* matrix and occurs once at the start of the SCF process. The second involves use of the ILUT factorized *A* matrix to solve the density fitting equations during every SCF iteration (where the only difference in the density fitting equations between SCF iterations is in the form of the right hand side). For the purpose of this paper we will refer to the first aspect as the "preconditioning time", while the second aspect is referred to as the "GMRES time". The combined time is referred to as the "density fitting total time". As sparsity of the preconditioner plays a key role in determining the performance and storage requirements for density fitting, we plot the preconditioning time, GMRES time and density fitting total time for System 1 as a function of the sparsity of the preconditioner in Figure 1. Also shown are the corresponding  $\tau$  values, and the total time taken if the density fitting problem is solved using the SVD direct approach.



Fig. 1. The dependence of the preconditioning time, GMRES time and the density fitting total time on the sparsity of the ILUT preconditioner for System 1



**Fig. 2.** The dependence of the preconditioning time, GMRES time and the density fitting total time on the sparsity of the ILUT preconditioner for System 2

The results in Figure 1 show that the preconditioning time increases monotonically as the sparsity of the preconditioner decreases, and for  $\tau=10^{-10}$  (i.e. when the LU decomposition is complete) the total time is similar to that required when using SVD. For the GMRES component decreasing the quality of the preconditioner to the level of 60% sparsity has minimal effect on the overall GMRES time, but moving beyond this level dramatically increases the GMRES time. This behavior reflects the fact that the number of GMRES iterations changes only slightly from 14 GMRES steps for  $\tau=10^{-10}$  to 87 steps when  $\tau=10^{-3}$ , however, for  $\tau=10^{-2}$  this number explodes to 994 GMRES steps; at this point the preconditioning is so poor that the GMRES algorithm has problems converging. Clearly, the goal is to pick the value of  $\tau$  that minimizes the overall time, and for this benchmark it appears to be a value of around  $10^{-3}$ , at which point the ILUT approach is about twice as fast as using SVD.

Equivalent performance results for the other 3 systems are shown in Figures 2-4. These all show similar behavior with steadily decreasing computational time that reaches a minimum before increasing dramatically if the value of  $\tau$  becomes too large. Interestingly, the value of  $\tau$  that works best appears to be roughly the same at  $10^3$  for all 4 systems. The ratio of the density fitting time using ILUT preconditioner with  $\tau$ =10<sup>3</sup>, to the SVD time for system 1, 2, 3 and 4 are 0.47, 0.78, 0.46 and 0.55 respectively. This shows that between system 1 and system 2, or system 3 and system 4, the maximum relative advantage of using ILUT over SVD is smaller the larger the fitting basis set. This is to be expected since larger fitting sets exhibit greater linear dependency giving rise to a more ill-conditioned *A* matrix. (We note that for system 2 the SVD shows 8 eigenvalues below  $10^{-5}$ , while it is full rank for the other systems.)



**Fig. 3.** The dependence of the preconditioning time, GMRES time and the density fitting total time on the sparsity of the ILUT preconditioner for System 3



**Fig. 4.** The dependence of the preconditioning time, GMRES time and the density fitting total time on the sparsity of the ILUT preconditioner for System 4



**Fig. 5.** The variation of the scaling preconditioning time of all three tests with the sparsity of the ILUT preconditioner

It is of interest to compare the preconditioning time as a function of the sparsity of the ILUT factorization across the different test systems. To do this it is necessary to scale the preconditioning time obtained for a given test by the preconditioning time obtained when the sparsity was equal to 1.0. These results are shown for all four systems in Figure 5. This shows that the preconditioning time exhibits a uniform decrease as the sparsity of the preconditioner increases, and that this rate of decrease is very similar for all test systems. Clearly the ILUT preconditioning time depends solely on the sparsity of the ILUT preconditioner. If we assume that the cost of the GMRES

iterations varies little with  $\tau$  until we reach the "tipping point", then it would be relatively easy to develop a performance model that can predict what level of sparsity in the ILUT preconditioner is required in order to achieve a given level of performance. Such a model might then be used to weigh up the potential gain associated with using ILUT to solve the density fitting problem over a direct approach like SVD.

In summary, it can be concluded that use of the ILUT method can enhance the performance of the density fitting process over use of a direct approach like SVD, although the performance gain depends greatly on the threshold used for  $\tau$ . It also appears unlikely that by using a simple numerical cutoff we will be able to exploit greater than approximately 50% sparsity in the representation of the LU factorization (as beyond this threshold the GMRES iterations tend to increase dramatically).

#### 4.2 Pre-screening of the A Matrix

In the above it was shown that the ILUT preconditioner failed to reach convergence if the value of  $\tau$  was smaller than about 10<sup>-3</sup>. At this point the sparsity of the preconditioner was still quite large, with around 50% of the elements being non-zero. While this is a useful memory reduction it is hardly dramatic, so it is pertinent to examine whether further memory reductions are possible by removing elements from *A* prior to performing the ILUT factorization. Two options are considered, i) the use of a pure numerical threshold ( $\Delta_{PRE}$ ) to set elements of *A* to zero, and ii) the removal of elements of *A* based the underlying physical problem [31]. Specifically with respect to (ii) we consider use of a sparse block diagonal preconditioner where the only elements of *A* to be considered are those that occur between functions located on the same atomic center. In what follows we use A' to denote the *A* matrix after certain elements have been set to zero.

As has been mentioned before the A matrix in the density fitting problem is essentially a Coulomb integral matrix in which the elements come from the Coulomb interaction between two fitting basis functions. Since the fitting functions are (usually) Gaussian functions located on different atomic centers the value of this integral will depend on both the distance between the two functions and the values of the exponents of the two Gaussian functions involved. At a coarse level we can, however, ignore the exponent values and assume only distance between two fitting functions will determine the value of the corresponding element in A.

In Table 2 we show the sparsity of A' obtained using a variety of different drop-off thresholds ( $\Delta_{PRE}$ ) and also using block diagonal sparsity. These tables also contrast the sparsity of A' with the sparsity of the preconditioner assuming an LU decomposition with unrestricted fill-in. These results show that even if A' is sparse the LU decomposition is significantly less sparse, e.g. for system 1 Table 2 shows that with  $\Delta_{PRE} = 10^{-2} A'$  has a sparsity of 0.105, but the preconditioner has over 45% of its elements non-zero. By contrast when using the block diagonal algorithm to derive A', fill-in is considerably less since it cannot exceed the block structure of A'. Thus in Table 2 we find that both the A' matrix and the preconditioner have very high sparsity with just 4% and 6.6% of their elements non-zero respectively when using

**Table 2.** Sparsity of the A' matrix and the resulting preconditioning matrix (*M*) obtained when employing numerical screening with criteria  $\Delta_{PRE}$  and the block diagonal scheme for all studied systems

	4 -	The Sparsity of the $A'(M)$ matrix				
	$\Delta_{PRE}$	System 1	System 2	System 3	System 4	
Numeric Screening	$10^{-10}$	0.97(1.00)	0.99(1.00)	1.00(1.00)	1.00(1.00)	
	10-5	0.61(0.95)	0.63(0.97)	0.72(1.00)	0.70(1.00)	
	$10^{-4}$	0.45(0.86)	0.43(0.89)	0.49(0.97)	0.47(0.98)	
	10-3	0.27(0.71)	0.23(0.75)	0.25(0.84)	0.24(0.86)	
	10-2	0.11(0.46)	0.08(0.44)	0.08(0.56)	0.07(0.57)	
Block Diagonal		0.04(0.07)	0.02(0.04)	0.01(0.01)	0.01(0.01)	



Fig. 6. Comparison of density fitting total time for computing System 1 using numerical screening and block diagonal screening on A' matrix



Fig. 7. Comparison of density fitting total time for computing System 2 using numerical screening and block diagonal screening on A' matrix

block diagonal screening; values that are over one order of the magnitude less than those for the original *A* matrix.

In Figures 6-9 we plot the total density fitting time for the 4 test systems that are obtained if prescreening of *A* is combined with use of the various  $\tau$  values to control fill-in during the ILUT process. The times are plotted as a function of sparsity in the preconditioner. Also shown are the SVD times. The results show that in terms of overall computation time the minimum is achieved at a sparsity level of around 50% non-zero elements. Removing elements from *A* by using a threshold of  $10^4$  for  $\Delta_{PRE}$  has relatively little effect on performance, except in the case of system 2 where a slightly tighter threshold is required. Going to the extreme of only keeping the block diagonal elements of *A* has a dramatic effect on the sparsity of the preconditioner, but in general it leads to an overall increase in the density fitting time in comparison to SVD. The exception is for system 3, where a block diagonal preconditioner is slightly faster than SVD or ILUT with a tight threshold.



Fig. 8. Comparison of density fitting total time for computing System 3 using numerical screening and block diagonal screening on A' matrix



Fig. 9. Comparison of density fitting total time for computing System 4 using numerical screening and block diagonal screening on A' matrix

#### **5** Discussion and Conclusions

We have investigated the use of ILUT preconditioning combined with the GMRES subspace method for iteratively solving the sort of linear equations systems that are encountered when using density fitting techniques in electronic structure calculations. Our results show that under the right circumstances it is possible to obtain a performance advantage from using the ILUT approach compared with a direct method like SVD, however, this requires careful choice for  $\tau$  (the numerical threshold parameter in the ILUT algorithm). Moreover, as  $\tau$  increases we can very quickly transition from having a beneficial preconditioning matrix to having one that is rather poor – causing a huge increase in the number of GMRES iterations required. Somewhat disappointingly it also appears that for the preconditioning matrix to be beneficial it requires over 50% of the matrix elements to be non-zero.

Using an alternative approach that approximates both *A* and the LU decomposition of *A* we found some encouraging results were obtained when using physical insight to zero out all elements in *A* except for those corresponding to interactions between basis functions on the same centre. This block diagonal approach dramatically decreases the number of elements in the LU decomposition, and this may be advantageous if memory usage is a bottleneck. In comparison to SVD, for small systems the block diagonal ILUT method was found to be slower, but for larger systems and moderate fitting sets it was found to be slightly faster. This raises the question whether an even better block diagonal preconditioner can be found, perhaps by expanding the size of the diagonal blocks to correspond to functional groups or small fragments of the total system. Work along these lines is currently in progress.

Acknowledgments. This work is funded by Australian Research Council Linkage Grants LP0347178 and LP0774896, and is in association with Gaussian Inc. and Sun Microsystems. Provision of computer time from the Australian Partnership in Advanced Computing is gratefully acknowledged, as is the donation of computing resources from Alexander Technology.

## References

- 1. Simoncini, V., Szyld, D.B.: Recent computational developments in Krylov subspace methods for linear systems. Numer. Linear Algebra Appl. 14, 1–59 (2007)
- Benzi, M.: Preconditioning Techniques for Large Linear Systems: A survey. J. Compu. Phys. 182, 418–477 (2002)
- Saad, Y.: ILUT: a dual threshold incomplete LU preconditioner. Numer. Linear Algebra Appl. 1, 387–402 (1994)
- Saad, Y., Schultz, M.H.: GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems. SIAM J. Sci. Statist. Comput. 7, 856–869 (1986)
- 5. Ten-no, S., Iwata, S.: Three-center expansion of electron repulsion integrals with linear combination of atomic electron distributions. Chem. Phys. Lett. 240, 578–584 (1995)

- Kendall, R.A., Früchtl, H.A.: The impact of the resolution of the identity approximate integral method on modern ab initio algorithm development. Theoret. Chem. Acc. 97, 158–163 (1997)
- 7. Früchtl, H.A., Kendall, R.A., Harrison, R.J., Dyall, K.G.: An implementation of RI-SCF on parallel computers. Int. J. Quantum Chem. 64, 63–69 (1997)
- 8. Weigend, F.: A fully direct RI-HF algorithm: Implementation, optimized auxiliary basis sets, demonstration of accuracy and efficiency. Phys. Chem. Chem. Phys. 4, 4285–4291 (2002)
- Polly, R., Werner, H.J., Manby, F.R., Knowles, P.J.: Fast Hartree-Fock theory using local density fitting approximations. Mol. Phys. 104, 2311–2321 (2004)
- 10. Ten-no, S., Iwata, S.: Multiconfiguration self-consistent field procedure employing linear combination of atomic-electron distributions. J. Chem. Phys. 105, 3604–3611 (1996)
- Feyereisen, M.W., Fitzgerald, G., Komornicki, A.: Use of approximate integrals in ab initio theory. An application in MP2 energy calculations. Chem. Phys. Lett. 208, 359–363 (1993)
- 12. Bernholdt, D.E., Harrison, R.J.: Large-scale correlated electronic structure calculations: the RI-MP2 method on parallel computers. Chem. Phys. Lett. 250, 477–484 (1996)
- 13. Weigend, F., Häser, M., Patzelt, H., Ahlrichs, R.: RI-MP2: optimized auxiliary basis sets and demonstration of efficiency. Chem. Phys. Lett. 294, 143–152 (1998)
- 14. Weigend, F., Köhn, A., Hättig, C.: Efficient use of the correlation consistent basis sets in resolution of the identity MP2 calculations. J. Chem. Phys. 116, 3175–3183 (2002)
- Werner, H.J., Manby, F.R., Knowles, P.J.: Fast linear scaling second order Møller-Plesset perturbation theory (MP2) using local and density fitting approximations. J. Chem. Phys. 118, 8149–8160 (2003)
- 16. Hättig, C., Weigend, F.: CC2 excitation energy calculations on large molecules using the resolution of the identity approximation. J. Chem. Phys. 113, 5154–5161 (2000)
- 17. Hättig, C.: Optimization of auxiliary basis sets for RI-MP2 and RI-CC2 calculations: Core-valence and quintuple- basis sets for H to Ar and QZVPP basis sets for Li to Kr. Phys. Chem. Chem. Phys. 7, 59–66 (2005)
- Rendell, A.P., Lee, T.J.: Coupled-cluster theory employing approximate integrals: An approach to avoid the input/output and storage bottlenecks. J. Chem. Phys. 101, 400–408 (1994)
- Schütz, M., Manby, F.R.: Linear scaling local coupled cluster theory with density fitting. Part I: 4-external integrals. Phys. Chem. Chem. Phys. 5, 3349–3358 (2003)
- Manby, F.R.: Density fitting in second-order linear-R12 Møller-Plesset perturbation theory. J. Chem. Phys. 119, 4607–4613 (2003)
- 21. Ten-no, S., Manby, F.R.: Density fitting for the decomposition of three-electron integrals in explicitly correlated electronic structure theory. J. Chem. Phys. 119, 5358–5363 (2003)
- Klopper, W.: A hybrid scheme for the resolution-of-the-identity approximation in second-order Møller-Plesset linear-r12 perturbation theory. J. Chem. Phys. 120, 10890–10895 (2004)
- Golub, G.H., van der Vorst, H.A.: Closer to the solution: iterative linear solvers. In: Duff, I.S., Watson, G.A. (eds.) The State of the Art in Numerical Analysis, pp. 63–92. Clarendon Press, Oxford (1997)
- 24. Saad, Y.: Iterative Methods for Sparse Linear Systems. PWS Publishing, New York, NY (1996)
- 25. SLATEC Common Mathematical Library, Version 4.1 (1993), http://www.netlib.org/slatec/
- 26. SPARSKIT, A basic tool-kit for sparse matrix computations (Version 2), http:// www- users. cs.umn.edu/ saad/software/SPARSKIT/sparskit.html

- Whitten, J.L.: Coulombic potential energy integrals and approximations. J. Chem. Phys. 58, 4496–4501 (1973)
- Dunlap, B.I., Connoly, J.W.D., Sabin, J.R.: On first-row diatomic molecules and local density models. J. Chem. Phys. 71, 4993–4999 (1979)
- Frisch, M.J., Trucks, G.W., Schlegel, H.B., Scuseria, G.E., Robb, M.A., Cheeseman, J.R., Montgomery, Jr A.J., Vreven, T., Kudin, K.N., Burant, J.C., Millam, J.M., Iyengar, S.S., Tomasi, J., Barone, V., Mennucci, B., Cossi, M., Scalmani, G., Rega, N., Petersson, G.A., Nakatsuji, H., Hada, M., Ehara, M., Toyota, K., Fukuda, R., Hasegawa, J., Ishida, M., Nakajima, T., Honda, Y., Kitao, O., Nakai, H., Klene, M., Li, X., Knox, J.E., Hratchian, H.P., Cross, J.B., Adamo, C., Jaramillo, J., Gomperts, R., Stratmann, R.E., Yazyev, O., Austin, A.J., Cammi, R., Pomelli, C., Ochterski, J.W., Ayala, P.Y., Morokuma, K., Voth, G.A., Salvador, P., Dannenberg, J., Zakrzewski, V.G., Dapprich, S., Daniels, A.D., Strain, M.C., Farkas, O., Malick, D.K., Rabuck, A.D., Raghavachari, K., Foresman, J.B., Ortiz, J.V., Cui, Q., Baboul, A.G., Clifford, S., Cioslowski, J., Stefanov, B.B., Liu, G., Liashenko, A., Piskorz, P., Komaromi, I., Martin, R.L., Fox, D.J., Keith, T., Al-Laham, M.A., Peng, C.Y., Nanayakkara, A., Challacombe, M., Gill, P.M.W., Johnson, B., Chen, W., Wong, M.W., Gonzalez, C., Pople, J.A.: Gaussian 03, Revision C.02, Gaussian, Inc., Wallingford CT (2004)
- Yang, R., Rendell, A. P., Frisch, M. J.: Automatically Generated Coulomb-Fitting Basis Sets: Design and Accuracy for Systems Containing H to Ne. J. Chem. Phys. (2007) (In Printing)
- Saad, Y., Zhang, J.: BILUM: Block Versions of Multielimination and Multilevel ILU Preconditioner for General Sparse Linear Systems. Society for Industrial and Applied Mathematics 20, 2103–2121 (1999)