

Letter

Linear Scaling Solution of the Time-Dependent Self-Consistent-Field Equations

Matt Challacombe

Theoretical Division, Los Alamos National Laboratory, Los Alamos, NM 87545, USA;

E-Mail: mchalla@lanl.gov; Tel.: +1-505-667-6117

Received: 19 November 2013; in revised form: 3 March 2014 / Accepted: 4 March 2014 /

Published: 14 March 2014

Abstract: A new approach to solving the Time-Dependent Self-Consistent-Field equations is developed based on the double quotient formulation of Tsiper 2001 (*J. Phys. B*). Dual channel, quasi-independent non-linear optimization of these quotients is found to yield convergence rates approaching those of the best case (single channel) Tamm-Dancoff approximation. This formulation is variational with respect to matrix truncation, admitting linear scaling solution of the matrix-eigenvalue problem, which is demonstrated for bulk excitons in the polyphenylene vinylene oligomer and the (4,3) carbon nanotube segment.

Keywords: quasi-independent optimization; rayleigh quotient iteration; *J*-symmetry; random phase approximation; time-dependent density functional theory; inexact linear algebra

1. Introduction

The Time-Dependent Self-Consistent-Field equations together with models that include some portion of the Hartree-Fock (HF) exchange admit control over the range of self-interaction in the optical response [1–4], and are related to new models of electron correlation based on the Random Phase Approximation (RPA) [5–8]. Solving the TD-SCF equations is challenging due to an unconventional *J*-symmetric structure of the naive molecular orbital (MO) representation,

$$\begin{pmatrix} \mathbb{A} & \mathbb{B} \\ -\mathbb{B} & -\mathbb{A} \end{pmatrix} \begin{pmatrix} \vec{X} \\ \vec{Y} \end{pmatrix} = \omega \begin{pmatrix} \vec{X} \\ \vec{Y} \end{pmatrix} \quad (1)$$

where \mathbb{A} and \mathbb{B} are Hermitian blocks corresponding to 4th order tensors spanning transitions between occupied and virtual sub-spaces, ω is the real excitation energy and $\vec{v} = \begin{pmatrix} \vec{X} \\ \vec{Y} \end{pmatrix}$ is the corresponding

transition density. By construction, the MO representation allows strict separation between the dyadic particle-hole (ph) and hole-particle (hp) solutions, \vec{X} and \vec{Y} , for which specialized algorithms exist. Nevertheless, convergence of the naive J -symmetric problem is typically much slower than the corresponding Hermitian Tamm-Dancoff approximation (TDA), $\mathbb{A}\vec{X} = \omega\vec{X}$, which is of reduced dimensionality in the MO representation.

Several TD-SCF eigensolvers are based on the oscillator picture $\begin{pmatrix} 0 & \mathbb{K} \\ \mathbb{T} & 0 \end{pmatrix} \begin{pmatrix} \vec{p} \\ \vec{q} \end{pmatrix} = \omega \begin{pmatrix} \vec{p} \\ \vec{q} \end{pmatrix}$, with $\mathbb{K} = \mathbb{A} + \mathbb{B}$ and $\mathbb{T} = \mathbb{A} - \mathbb{B}$ the Hermitian potential and kinetic matrices, and the dual $\{\vec{p}, \vec{q}\} = \{\vec{X} - \vec{Y}, \vec{X} + \vec{Y}\}$ corresponding to position and momentum. This picture avoids the imbalance $\|\vec{X}\| \gg \|\vec{Y}\|$ whilst admitting conventional solutions based on the Hermitian matrix $\mathbb{G} = \mathbb{K} \cdot \mathbb{T}$, as shown by Tamara and Udagawa [9] and extended by Narita and Shibuya with second order optimization of the quotient $\omega^2[\vec{p}, \vec{q}] = \vec{q} \cdot \mathbb{G} \cdot \vec{p} / |\vec{p} \cdot \vec{q}|$ [10]. More recently, Tsiper [11] considered the quotients

$$\omega[\vec{p}, \vec{q}] = \frac{\vec{p} \cdot \mathbb{K} \cdot \vec{p}}{2 |\vec{p} \cdot \vec{q}|} + \frac{q \cdot \mathbb{T} \cdot \vec{q}}{2 |\vec{p} \cdot \vec{q}|} \quad (2)$$

and developed a corresponding dual channel Lanczos solver. Subspace solvers in this dual representation have recently been surveyed by Tretiak, Isborne, Niklasson and Challacombe (TINC) [12], with comparative results for semi-empirical models.

Another challenge is dimensionality and scaling. Writing Equation (1) in the general form $\mathbb{L} \cdot \vec{v} = \omega \vec{v}$, admitting arbitrary representation, the superoperator matrix \mathbb{L} is a $\sim N^2 \times N^2$ tetradic, with N the number of basis functions, assumed proportional to system size. In practice the action of \mathbb{L} onto \vec{v} is carried out implicitly as $\mathbf{L}[\mathbf{v}] = [\mathbf{F}, \mathbf{v}] + [\mathbf{G}[\mathbf{v}], \mathbf{P}]$, using an existing framework for construction of the effective Hamiltonian (Fockian) \mathbf{F} , where \mathbf{P} is the one-particle reduced density matrix, \mathbf{G} is a screening operator involving Coulomb, exchange and/or exchange-correlation terms and the correspondence between superoperator and functional notation is given by a tensorial mapping between diadic and matrix, $\vec{v}_{N^2 \times 1} \Leftrightarrow \mathbf{v}_{N \times N}$.

Recent efforts have focused on addressing the problem of dimensionality by employing linear scaling methods that reduce the cost of $\mathbf{L}[\cdot]$ within Density Functional Theory (DFT) to $\mathcal{O}(N)$. However, this remains an open problem for the Hartree-Fock (HF) exchange, an ingredient in models that account for charge transfer in the dynamic and static response, including the Random Phase Approximation (RPA) at the pure HF level of theory. Likewise, scaling of the TD-SCF eigenproblem remains formidable due to associated costs of linear algebra, even when using powerful Krylov subspace methods. Underscoring this challenge, one of the most successful approaches to linear scaling TD-DFT avoids the matrix eigenproblem entirely through explicit time-evolution [13,14].

Linear scaling matrix methods exploit quantum locality, manifest in approximate exponential decay of matrix elements expressed in a well posed, local basis; with the dropping of small elements below a threshold, τ_{mtx} , this decay leads to sparse matrices and $\mathcal{O}(N)$ complexity at the forfeit of full precision [15–17]. Likewise, linear scaling methods for computing the HF exchange employ an advanced form of direct SCF, exploiting this decay in the rigorous screening of small exchange interactions below the two-electron integral threshold τ_{2e} [18]. The consequence of these linear scaling approximations is an inexact linear algebra that challenges Krylov solvers due to nested error

accumulation, a subject of recent formal interest [19,20]. Consistent with this view, TINC found that matrix perturbation (a truncation proxy) disrupts convergence of Krylov solvers with slow convergence, *i.e.*, Lanczos and Arnoldi for the RPA, but has less impact on solvers with rapid convergence, *i.e.*, generically for the TDA or Davidson for the RPA. Relative to semi-empirical Hamiltonians, the impact of incompleteness on subspace iteration may be amplified with first principles models and large basis sets (ill-conditioning).

An alternative is Rayleigh Quotient Iteration (RQI), which poses the eigenproblem as non-linear optimization and is variational with respect to matrix perturbation. Narita and Shibuya [10] considered optimization of the quotient $\omega^2 [\vec{p}, \vec{q}]$ with second order methods, but these are beyond the capabilities of current linear scaling technologies and also, convergence may be slower by a power of $1/2$. For semi-empirical Hamiltonians, TINC found that optimization of the Thouless functional $\omega[\vec{v}] = \vec{v} \cdot \mathbb{L} \cdot \vec{v} / |\vec{v} \cdot \vec{v}|$, corresponding to the solution of Equation (1), was significantly slower for the RPA relative to the TDA, and also compared to subspace solvers. For first principles models and non-trivial basis sets, this naive RQI can become pathologically slow as shown in Figure 1. On the other hand, the Tsiper formulation exposes the underlying pseudo-Hermitian structure of the TD-SCF equations. Here, this structure is exploited with QUasi-Independent Rayleigh Quotient Iteration (QUIRQI), involving dual channel optimization of the Tsiper quotients coupled only weakly through line search. Although this work was first placed in the arXiv some time ago [21], it is offered here after review and revision, with changes primarily in the concluding remarks.

2. Theoretical Development

Our development begins with a brief review of the representation independent formulation developed by TINC, which avoids the $\mathcal{O}(N^3)$ cost of rotating into an explicit p - h , h - p symmetry. Instead, this symmetry is maintained implicitly via annihilation, $\mathbf{x} \leftarrow f_a(\mathbf{x}) = \mathbf{P} \cdot \mathbf{x} \cdot \mathbf{Q} + \mathbf{Q} \cdot \mathbf{x} \cdot \mathbf{P}$, with \mathbf{P} the first order reduced density matrix and $\mathbf{Q} = \mathbf{I} - \mathbf{P}$ its compliment. Likewise, the indefinite metric associated with the J -symmetry of Equation (1) is carried through the generalized norm $\langle \mathbf{x}, \mathbf{y} \rangle = \text{tr}\{\mathbf{x}^T \cdot [\mathbf{y}, \mathbf{P}]\}$. Introducing the operator equivalents, $\mathbf{L}[\mathbf{p}] \Leftrightarrow \mathbb{K} \cdot \vec{p}$ and $\mathbf{L}[\mathbf{q}] \Leftrightarrow \mathbb{T} \cdot \vec{q}$, the Tsiper functional becomes $\omega[\mathbf{p}, \mathbf{q}] = \frac{\langle \mathbf{p}, \mathbf{L}[\mathbf{p}] \rangle}{2|\langle \mathbf{p}, \mathbf{q} \rangle|} + \frac{\langle \mathbf{q}, \mathbf{L}[\mathbf{q}] \rangle}{2|\langle \mathbf{p}, \mathbf{q} \rangle|}$. Transformations between the transition density and the dual space involves simple manipulations and minimal cost, allowing Fock builds with the transition density and optimization in the dual space. The splitting operation is given by $\mathbf{p} = f_+(\mathbf{v}) = \mathbf{P} \cdot \mathbf{v} \cdot \mathbf{Q} + [\mathbf{Q} \cdot \mathbf{v} \cdot \mathbf{P}]^T$ and $\mathbf{q} = f_-(\mathbf{v}) = \mathbf{P} \cdot \mathbf{v} \cdot \mathbf{Q} - [\mathbf{Q} \cdot \mathbf{v} \cdot \mathbf{P}]^T$. Likewise, $\mathbf{L}[\mathbf{p}] = f_-(\mathbf{L}[\mathbf{v}])$ and $\mathbf{L}[\mathbf{q}] = f_+(\mathbf{L}[\mathbf{v}])$. The back transformation (merge) from dual to density is $\mathbf{v} = F(\mathbf{p}, \mathbf{q}) = (\mathbf{p} + \mathbf{q} + [\mathbf{p} - \mathbf{q}]^T) / 2$. This framework provides the freedom to work in any orthogonal representation, and to switch between transition density and oscillator duals with minimal cost.

QUIRQI is given in Algorithm 1. It begins with a guess for the transition density, which is then split into its dual (lines 2–3). The choice of initial guess is discussed later. Lines 4–24 consist of the non-linear conjugate gradient optimization of the nearly independent channels: In each step, the flow of information proceeds from optimization of the duals to builds involving the density and back to the duals in a merge-annihilate-truncate-build-split-truncate (MATBST) sequence. For the variables \mathbf{v} , \mathbf{p} and \mathbf{q} this sequence is comprised by lines 22–23 and 5–7, and lines 15–19 for the corresponding conjugate

gradients \mathbf{h}_v , \mathbf{h}_p and \mathbf{h}_q . Truncation is carried out with the *filter* operation as described in Reference [17] and also below, with cost and error determined by the matrix threshold τ_{mtx} .

Algorithm 1 QUIRQI

```

1: procedure QUIRQI( $\omega, \mathbf{v}$ )
2:   guess  $\mathbf{v}$ 
3:    $\mathbf{p} = f_+(\mathbf{v}), \mathbf{q} = f_-(\mathbf{v})$ 
4:   while  $e_{\text{rel}} > \epsilon$  and  $g_{\text{max}} > \gamma$  and  $\omega < \omega^{\text{old}}$  do
5:      $\mathbf{L}[\mathbf{v}] = [\mathbf{F}, \mathbf{v}] + [\mathbf{G}[\mathbf{v}], \mathbf{P}]$ 
6:      $\mathbf{L}[\mathbf{p}] = f_-(\mathbf{L}[\mathbf{v}]), \mathbf{L}[\mathbf{q}] = f_+(\mathbf{L}[\mathbf{v}])$ 
7:     filter( $\mathbf{L}[\mathbf{p}], \mathbf{L}[\mathbf{q}], \tau_{\text{mtx}}$ )
8:      $\omega_p = \frac{\langle \mathbf{p}, \mathbf{L}[\mathbf{p}] \rangle}{2|\langle \mathbf{p}, \mathbf{q} \rangle|}, \omega_q = \frac{\langle \mathbf{q}, \mathbf{L}[\mathbf{q}] \rangle}{2|\langle \mathbf{p}, \mathbf{q} \rangle|}, \omega = \omega_p + \omega_q$ 
9:      $\mathbf{g}_p = \mathbf{q} \omega_q - \mathbf{L}[\mathbf{p}], \mathbf{g}_q = \mathbf{p} \omega_p - \mathbf{L}[\mathbf{q}]$ 
10:     $e_{\text{rel}} = (\omega^{\text{old}} - \omega) / \omega$ 
11:     $g_{\text{max}} = \max_{i,j} \{ [\mathbf{g}_p]_{ij}, [\mathbf{g}_q]_{ij} \}$ 
12:     $\beta_p = \frac{\langle \mathbf{g}_p, \mathbf{g}_p - \mathbf{g}_p^{\text{old}} \rangle}{\langle \mathbf{g}_p^{\text{old}}, \mathbf{g}_p^{\text{old}} \rangle}, \beta_q = \frac{\langle \mathbf{g}_q, \mathbf{g}_q - \mathbf{g}_q^{\text{old}} \rangle}{\langle \mathbf{g}_q^{\text{old}}, \mathbf{g}_q^{\text{old}} \rangle}$ 
13:     $\omega^{\text{old}} \leftarrow \omega, \mathbf{g}_p^{\text{old}} \leftarrow \mathbf{g}_p, \mathbf{g}_q^{\text{old}} \leftarrow \mathbf{g}_q$ 
14:     $\mathbf{h}_p \leftarrow \mathbf{g}_p + \beta_p \mathbf{h}_p, \mathbf{h}_q \leftarrow \mathbf{g}_q + \beta_q \mathbf{h}_q$ 
15:     $\mathbf{h}_v = F(\mathbf{h}_p, \mathbf{h}_q), \mathbf{h}_v \leftarrow f_a(\mathbf{h}_v)$ 
16:    filter( $\mathbf{h}_p, \mathbf{h}_q, \mathbf{h}_v, \tau_{\text{mtx}}$ )
17:     $\mathbf{L}[\mathbf{h}_v] = [\mathbf{F}, \mathbf{h}_v] + [\mathbf{G}[\mathbf{h}_v], \mathbf{P}]$ 
18:     $\mathbf{L}[\mathbf{h}_p] = f_-(\mathbf{L}[\mathbf{h}_v]), \mathbf{L}[\mathbf{h}_q] = f_+(\mathbf{L}[\mathbf{h}_v])$ 
19:    filter( $\mathbf{L}[\mathbf{h}_p], \mathbf{L}[\mathbf{h}_q], \tau_{\text{mtx}}$ )
20:     $\{ \lambda_p, \lambda_q \} = \underset{\{ \lambda_p, \lambda_q \}}{\text{argmin}} \omega [\mathbf{p} + \lambda_p \mathbf{h}_p, \mathbf{q} + \lambda_q \mathbf{h}_q]$ 
21:     $\mathbf{p} \leftarrow \mathbf{p} + \lambda_p \mathbf{h}_p, \mathbf{q} \leftarrow \mathbf{q} + \lambda_q \mathbf{h}_q$ 
22:     $\mathbf{v} \leftarrow F(\mathbf{p}, \mathbf{q}), \mathbf{v} \leftarrow f_a(\mathbf{v})$ 
23:    filter( $\mathbf{p}, \mathbf{q}, \mathbf{v}, \tau_{\text{mtx}}$ )
24:   end while
25: end procedure

```

The Tsiper functional is the sum of dual quotients ω_p and ω_q , determined at line 8, followed by the gradients \mathbf{g}_p and \mathbf{g}_q computed at line 9. After the first cycle, the corresponding relative error e_{rel} (10) and maximum matrix element of the gradient g_{max} (11) are computed and used as an exit criterion at line 4, along with non-variational behavior $\omega > \omega^{\text{old}}$.

Next, the Polak-Ribiere variant of non-linear conjugate gradients yields the search direction in each channel, \mathbf{h}_p and \mathbf{h}_q (12–14). The action of $\mathbf{L}[\cdot]$ on to \mathbf{h}_p and \mathbf{h}_q is then computed, again with a MATBST sequence (15–19), followed by a self-consistent dual channel line search at line 20, as described below. With steps λ_p and λ_q in hand, minimizing updates are taken along each conjugate direction (22), and the cycle repeats with the MATBST sequence spanning lines 21–23 and 5–7.

Optimization of the Tsipper functional $\omega[\lambda_p, \lambda_q] \equiv \omega[\mathbf{p} + \lambda_p \mathbf{h}_p, \mathbf{q} + \lambda_q \mathbf{h}_q]$ involves a two dimensional line-search (line 20) corresponding to minimization of

$$\omega[\lambda_p, \lambda_q] = \frac{A_p + \lambda_p B_p + \lambda_p^2 C_p + A_q + \lambda_q B_q + \lambda_q^2 C_q}{R_{pq} + \lambda_p S_{pq} + \lambda_q T_{pq} + \lambda_p \lambda_q U_{pq}} \quad (3)$$

with coupling entering through terms in the denominator such as $U_{pq} = \langle \mathbf{h}_p, \mathbf{h}_q \rangle$. A minimum in Equation (3) can be found quickly to high precision by alternately substituting one-dimensional solutions one into the other until self-consistency is reached. This semi-analytic approach starts with a rough guess at the pair $\{\lambda_p, \lambda_q\}$ (eg. found by a coarse scan) followed by iterative substitution, where for example the p -channel update is

$$\begin{aligned} \lambda_p \leftarrow & \left\{ \left[(2C_p R_{pq} + 2C_p \lambda_q S_{pq})^2 - 4(C_p T_{pq} + C_p \lambda_q U_{pq}) \right. \right. \\ & \left. \left[B_p R_{pq} + B_p \lambda_q S_{pq} - (A_q + A_p + B_q \lambda_q + C_q \lambda_q^2) T_{pq} \right. \right. \\ & \left. \left. - (A_q \lambda_q + A_q \lambda_q + B_q \lambda_q^2 + C_q \lambda_q^3) U_{pq} \right] \right]^{1/2} \\ & \left. - 2C_p R_{pq} - 2C_p \lambda_q S_{pq} \right\} / [2C_p (T_{pq} + \lambda_q U_{pq})] \end{aligned} \quad (4)$$

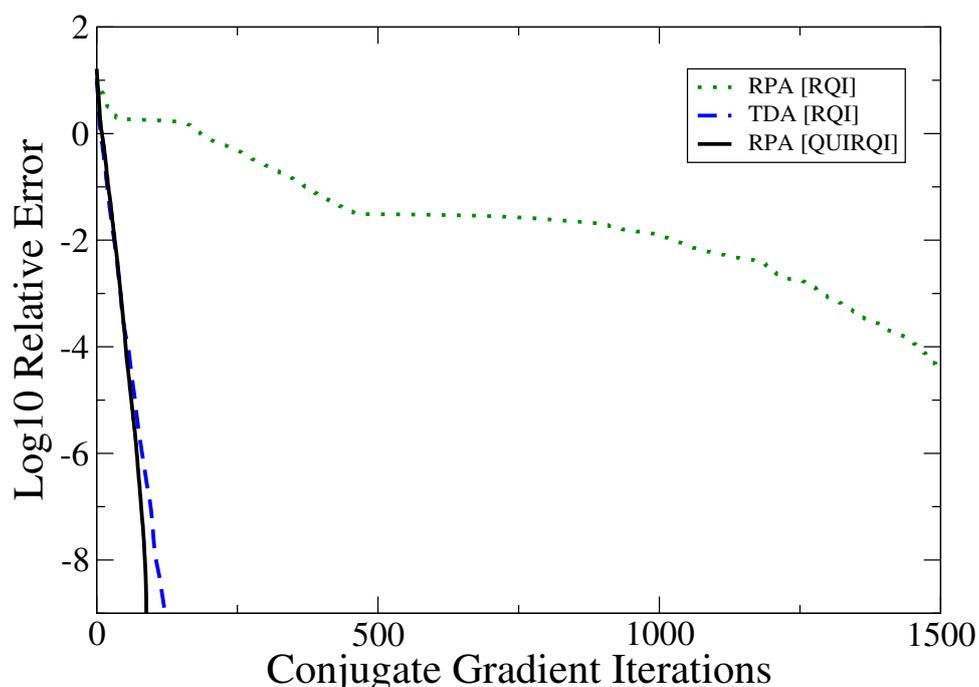
with an analogous update for the q -channel obtained by swapping subscripts. As the solution decouples (S_{pq} , T_{pq} and U_{pq} become small) the steps are found independently.

3. Results and Discussion

QUIRQI has been implemented in FreeON [22], which employs the linear scaling Coulomb and Hartree-Fock exchange kernels QCTC and ONX with cost and accuracy controlled by the two-electron screening threshold τ_{2e} [18]. N -scaling solution of the QUIRQI matrix equations is achieved through “sparsification” (In previous works [15], this process has been loosely referred to as SpAMM, involving both truncation and dynamic dropping of small row-column contributions from the sparse matrix multiply based on the BCSR data structure. See Section IVA3 of Reference [15]. In this work, only truncation in the BCSR data structure has been used, but was incorrectly referred to as SpAMM in a previous instance [21]. In more recent developments, SpAMM refers to recursive, hierarchical truncation in the product space [23,24], rather than the row-column approach outlined in Reference [15].), of the underlying vector space, where the *filter* operation is applied to drop atom-blocks with norm smaller than a drop tolerance τ_{mtx} in a block-CSR data structure (BCSR) [15–17]. All calculations were carried out with version 4.3 of the gcc/gfortran compiler under version 8.04 of the Ubuntu Linux distribution and run on a 2GHz AMD Quad Opteron 8350.

For systems studied to date, QUIRQI is found to converge monotonically with rates comparable to the TDA as shown in Figure 1. Based on the comparative performance presented by TINC, the TDA rate of convergence appears to be a lower bound for RPA solvers. In addition to the convergence rate, performance is strongly determined by the initial guess. The following results have been obtained using the polarization response density along the polymer axis [12], which can be computed in $\mathcal{O}(N)$ by Perturbed Projection [25]. Also, a relative precision of 4 digits in the excitation energy is targeted with the convergence parameters $\epsilon = 10^{-4}$ and $\gamma = 10^{-3}$, with exit from the optimization loop on violation of monotonic convergence ($\omega > \omega^{\text{old}}$ due to precision limitations associated with linear scaling approximations).

Figure 1. Convergence of RHF/3-21G Tamm-Dancoff approximation (TDA) and Random Phase Approximation (RPA) with the Rayleigh Quotient Iteration (RQI) and QUasi-Independent RQI (QUIRQI) algorithms for linear decaene ($C_{10}H_2$). Calculations were started from the same random guess, and tight numerical thresholds were used throughout. In the representation independent scheme, the cost per iteration is the same for TDA and RPA.



In Figure 2, linear scaling and convergence to the bulk limit are demonstrated for a series of polyphenylene vinylene (PPV) oligomers at the RHF/6-31G** level of theory for the threshold combinations $\{\tau_{\text{mtx}}, \tau_{2e}\} = \{10^{-4}, 10^{-5}\}$ and $\{10^{-5}, 10^{-6}\}$. Significantly more conservative thresholds have been used for the Coulomb sums, which incur only minor cost. Convergence is reached in 24–25 iterations, with the cost of Coulomb summation via QCTC comparable to the cost of BCSR ($\tau_{\text{mtx}} = 10^{-4}$). In Figure 3, linear scaling and convergence to the bulk limit are demonstrated for a series of (4,3) carbon nanotube segments at the RHF/3-21G level of theory for the same threshold combinations, again with convergence achieved in about 24–25 cycles. In both cases, tightening the pair $\{\tau_{\text{mtx}}, \tau_{2e}\}$ leads to a systematically improved result. While the $\{10^{-4}, 10^{-5}\}$ thresholds that work well for PPV lead to a non-monotone behavior with respect to extent for the nanotube series, dropping one more decade to $\{10^{-5}, 10^{-6}\}$ leads to a sharply improved behavior. Dropping thresholds further to $\{10^{-6}, 10^{-7}\}$ yields identical results to within the convergence criteria (\sim four digits) across the series, also scaling with N but at roughly twice the cost.

These results demonstrate that QUIRQI can achieve both systematic error control and linear scaling in solution of the RPA eigenproblem for systems with extended conjugation. Relative to PPV, the greater numerical sensitivity encountered with the nanotube series is consistent with the ground state problem, where a smaller band gap and greater atomic connectivity typically demand tighter thresholds.

Figure 2. Approach to the bulk limit of the polyphenylene vinylene (PPV) first excited state at the 6-31G**/RPA level of theory, with inset showing linear scaling cost for Hartree-Fock (HF) exchange (ONX) and sparse linear algebra (BCSR). The cost of Coulomb sums with much tighter thresholds are comparable to those for the BCSR.

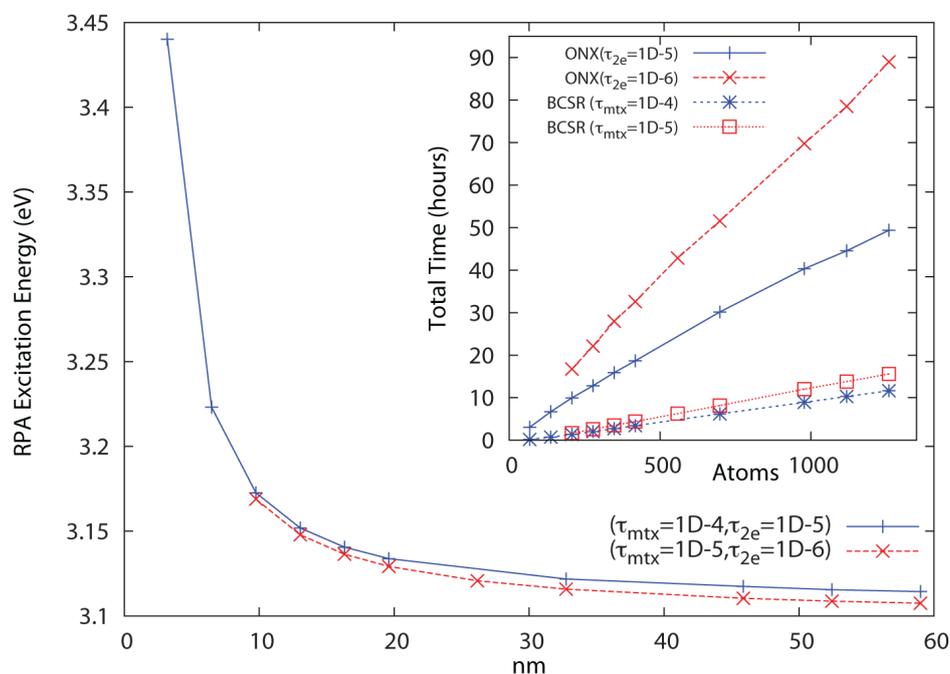
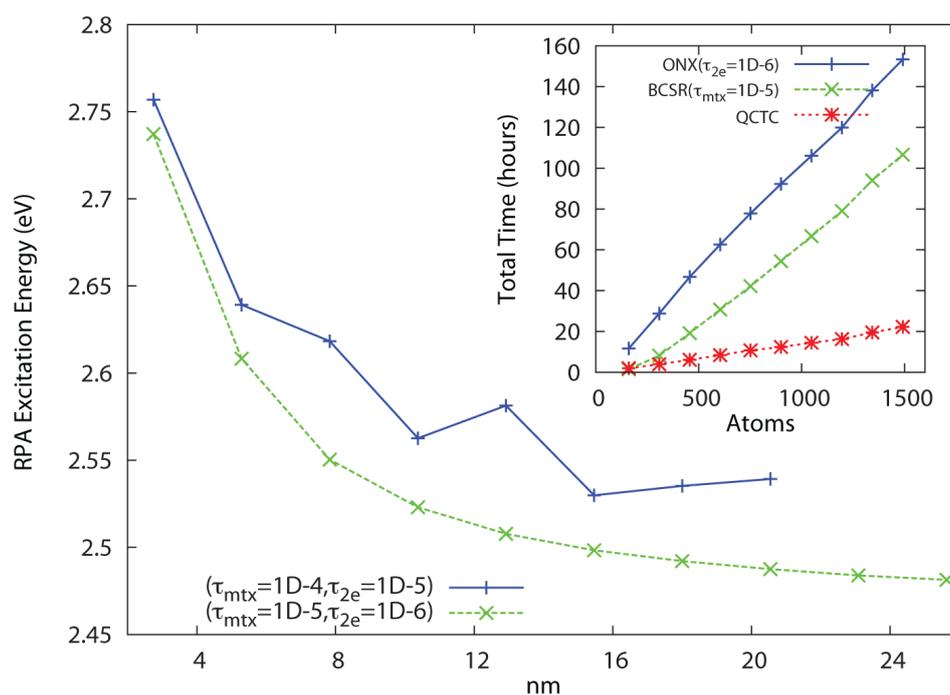


Figure 3. Approach to the bulk limit of the first excited state of the (4,3) carbon nanotube segment at the 3-21G/RPA level of theory, with inset showing linear scaling cost for HF exchange (ONX), sparse linear algebra (BCSR) and Coulomb sums (QCTC).



4. Conclusions

Since this note appeared some time ago in arXiv [21], several related efforts have appeared that deserve comment: (A) single channel optimization with radial cutoffs [26] and (B) conventional algebra with a four channel line search [27,28]. In the first instance, the ONETEP group have implemented a single channel quotient scheme for the TDA and demonstrated linear scaling for a number of systems using the radial cutoff approach to achieve reduced complexity. In the radial cutoff approach, portions of the vector space are eliminated from the linear algebra when the Cartesian distance between associated atoms becomes greater than some cutoff radius (Going a step further, new technologies are emerging that achieve reduced complexity without truncation in the vector space [23,24].). The ONETEP paper is recommended by their careful discussion of radial cutoffs leading to artificial truncation in cases of long range charge transfer, e.g., Reference [26] Figure 5. In the current implementation, the *filter* operation eliminates elements of the vector space that are numerically small; in the case of long range charge transfer, extended conjugation, *etc.*, an unphysical truncation does not occur. For problems without long range charge transfer or extended conjugation, for example large problems with well localized chromophores as in References [29,30], the complexity of QUIRQI with respect to system size becomes $\mathcal{O}(1)$.

In the second instance, the QUIRQI method has been extended to include two additional channels in the line search [27,28]; $\operatorname{argmin}_{\{\alpha,\beta,\lambda_p,\lambda_q\}} \omega [\alpha \mathbf{p} + \lambda_p \mathbf{h}_p, \beta \mathbf{q} + \lambda_q \mathbf{h}_q]$. The authors claim without elaboration that “the solution by our 4D search is and can be much better (than the dual channel approach)” [28]. As shown in Figure (1), QUIRQI decouples in the first few steps achieving convergence equivalent to TDA[RQI], $R_{pq} \rightarrow 1, S_{pq} \rightarrow 0, T_{pq} \rightarrow 0, U_{pq} \rightarrow 0$ in Equation (3), so it is hard to understand this unsupported claim, considering also the imperative that $\alpha * \beta = 1$ to maintain normalization. These claims are also undercut by apparently slow rates of convergence; compare for example Figure (1) of this work with Reference [27], especially Figure 3. These authors further claim without explanation that “dual channel optimization are not readily extensible to the subspace search” [28]. Again, its hard to understand how the dual channel case isn’t extensible to subspace schemes for finding multiple eigenvalues; the equivalent of (single channel) block RQI has been demonstrated in the ONETEP paper [26], and no obvious problems are foreseen with more sophisticated methods such as LOBPCG [31] for single or dual channel approaches.

To summarize, the QUIRQI method is characterized by two innovations: First, dual channel optimization separates the Tsipper functional into two, nearly independent quotients that cannot be further improved by additional channels in the line search. Reflecting this separation, convergence of the TD-SCF matrix eigenproblem with QUIRQI is found to be equivalent to the single quotient matrix eigenproblem in the Tamm-Dancoff approximation, as shown in Figure 1. Second, the method is variational with respect to an incomplete linear algebra, controlled in this work through the *filter* threshold τ_{mtx} [15–17], as shown in Figures 2 and 3. While QUIRQI is not variational with respect to the screening parameter τ_{2e} , the solution can be systematically improved by tightening τ_{2e} [18], in comparison to nested subspace methods that encounter an iterative accumulation of errors [19,20]. Indeed, eigensolution posed as optimization provides considerable flexibility in choosing a path to

solution, offering opportunities for mixed precision GPU acceleration [32] and variable thresholding (tightening the parameter τ_{2e} during convergence).

Acknowledgments

This work was supported by the U. S. Department of Energy under Contract No. DE-AC52-06NA25396 and LDRD-ER grant 20070256ER. Special acknowledgments go the International Ten Bar Café for scientific libations, and to Sergei Tretiak and Nicolas Bock for valuable input.

Conflicts of Interest

The authors declare no conflict of interest.

References

1. Autschbach, J. Charge-transfer excitations and time-dependent density functional theory: Problems and some proposed solutions. *ChemPhysChem* **2009**, *10*, 1757–1760.
2. Song, J.W.; Watson, M.A.; Hirao, K. An improved long-range corrected hybrid functional with vanishing Hartree-Fock exchange at zero interelectronic distance (LC2gau-BOP). *J. Chem. Phys.* **2009**, *131*, 144108.
3. Igumenshchev, K.I.; Tretiak, S.; Chernyak, V.Y. Excitonic effects in a time-dependent density functional theory. *J. Chem. Phys.* **2007**, *127*, 114902.
4. Magyar, R.J.; Tretiak, S. Dependence of spurious charge-transfer excited states on orbital exchange in TDDFT: Large molecules and clusters. *J. Chem. Theor. Comp.* **2007**, *3*, 976–987.
5. Lu, D.; Li, Y.; Rocca, D.; Galli, G. *Ab initio* Calculation of van der Waals Bonded Molecular Crystals. *Phys. Rev. Lett.* **2009**, *102*, 206411.
6. Furche, F. Developing the random phase approximation into a practical post-Kohn-Sham correlation model. *J. Chem. Phys.* **2008**, *129*, 114105.
7. Harl, J.; Kresse, G. Accurate bulk properties from approximate many-body techniques. *Phys. Rev. Lett.* **2009**, *103*, 056401.
8. Toulouse, J.; Gerber, I.C.; Jansen, G.; Savin, A.; Ángyán, J.G. Adiabatic-connection fluctuation-dissipation density-functional theory based on range separation. *Phys. Rev. Lett.* **2009**, *102*, 096404.
9. Tamura, T.; Udagawa, T. Higher random phase approximation applied to the description of the spherical vibrational nuclei. *Nuc. Phys.* **1964**, *53*, 33–55.
10. Narita, S.; Shibuya, T. New scheme for large-scale eigenvalue problems of the RPA-type matrix equation. *Can. J. Chem.* **1992**, *70*, 296.
11. Tsiper, E.V. A classical mechanics technique for quantum linear response. *J. Phys. B* **2001**, *34*, L401–L407.
12. Tretiak, S.; Isborn, C.M.; Niklasson, A.M.N.; Challacombe, M. Representation independent algorithms for molecular response calculations in time-dependent self-consistent field theories. *J. Chem. Phys.* **2009**, *130*, 054111.

13. Yokojima, S.; Chen, G.H. Linear-scaling localized-density-matrix method for the ground and excited states of one-dimensional molecular systems. *Chem. Phys. Lett.* **1999**, *300*, 540–4.
14. Yam, C.Y.; Yokojima, S.; Chen, G. Localized-density-matrix implementation of time-dependent density-functional theory. *J. Chem. Phys.* **2003**, *119*, 8794–8803.
15. Challacombe, M. A simplified density matrix minimization for linear scaling self-consistent field theory. *J. Chem. Phys.* **1999**, *110*, 2332–2342.
16. Challacombe, M. A general parallel sparse-blocked matrix multiply for linear scaling SCF theory. *Comput. Phys. Commun.* **2000**, *128*, 93.
17. Niklasson, A.M.N.; Tymczak, C.J.; Challacombe, M. Trace resetting density matrix purification in self-consistent-field theory *J. Chem. Phys.* **2003**, *118*, 8611.
18. Schwegler, E.; Challacombe, M.; Head-Gordon, M. Linear scaling computation of the fock matrix. 2. Rigorous bounds on exchange integrals and incremental fock build. *J. Chem. Phys.* **1997**, *106*, 9708–9717.
19. Simoncini, V. Variable accuracy of matrix-vector products in projection methods for eigencomputation. *SIAM J. Numer. Anal.* **2005**, *43*, 1155–1174.
20. Simoncini, V.; Szyld, D.B. Recent computational developments in Krylov subspace methods for linear systems. *Numer. Linear Algebra Appl.* **2007**, *14*, 1–59.
21. Challacombe, M. *Linear Scaling Solution of the Time-Dependent Self-Consistent-Field Equations*; Cornell University Library: Ithaca, NY, USA, 2010; p. 4.
22. Bock, N.; Challacombe, M.; Gan, C.K.; Henkelman, G.; Nemeth, K.; Niklasson, A.M.N.; Odell, A.; Schwegler, E.; Tymczak, C.J.; Weber, V. FreeON (Sep-18-2009). Los Alamos National Laboratory (LA-CC-04-086), University of California, CA, USA, 2009.
23. Challacombe, M.; Bock, N. *Fast Multiplication of Matrices with Decay*. Cornell University Library: Ithaca, NY, USA, 2010.
24. Bock, N.; Challacombe, M. An optimized sparse approximate matrix multiply for matrices with decay. *SIAM J. Sci. Comput.* **2013**, *35*, C72–C98.
25. Weber, V.; Niklasson, A.M.N.; Challacombe, M. *Ab initio* linear scaling response theory: Electric polarizability by perturbed projection. *Phys. Rev. Lett.* **2004**, *92*, 193002.
26. Zuehlsdorff, T.J.; Hine, N.D.M.; Spencer, J.S.; Harrison, N.M.; Riley, D.J.; Haynes, P.D. Linear-scaling time-dependent density-functional theory in the linear response formalism. *J. Chem. Phys.* **2013**, *139*, 064104.
27. Rocca, D.; Bai, Z.; Li, R.C.; Galli, G. A block variational procedure for the iterative diagonalization of non-Hermitian random-phase approximation matrices. *J. Chem. Phys.* **2012**, *136*, 034111.
28. Bai, Z.; Li, R.C. Minimization Principles for the Linear Response Eigenvalue Problem II: Computation. *SIAM J. Matrix Anal. Appl.* **2013**, *34*, 392–416.
29. Isborn, C.M.; Götz, A.W.; Clark, M.A.; Walker, R.C.; Martínez, T.J. Electronic absorption spectra from MM and *ab initio* QM/MM molecular dynamics: Environmental effects on the absorption spectrum of photoactive yellow protein. *J. Chem. Theory Comput.* **2012**, *8*, 5092–5106.

30. Isborn, C.M.; Mar, B.D.; Curchod, B.F.E.; Tavernelli, I.; Martínez, T.J. The charge transfer problem in density functional theory calculations of aqueously solvated molecules. *J. Phys. Chem. B* **2013**, *117*, 12189–121201.
31. Knyazev, A.V. Toward the optimal preconditioned eigensolver: Locally optimal block preconditioned conjugate gradient method. *SIAM J. Sci. Comput.* **2001**, *23*, 517–541.
32. Luehr, N.; Ufimtsev, I.S.; Martinez, T.J. Dynamic precision for electron repulsion integral evaluation on graphical processing units (GPUs). *J. Chem. Theory Comput.* **2011**, *7*, 949–954.

© 2014 by the author; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution license (<http://creativecommons.org/licenses/by/3.0/>).