

Strategies for large-scale molecular self-consistent field calculations

Trygve Helgaker

Centre for Theoretical and Computational Chemistry, University of Oslo, Norway

Filip Pawłowski, Simen Reine, Erik Tellgren

Stinne Høst, Branislav Jansík, Poul Jørgensen, Jeppe Olsen, University of Aarhus, Denmark

Sonia Coriani, University of Trieste, Italy

Paweł Sałek, Royal Institute of Technology, Sweden

Teoretisk og beregningsorientert kjemi

Det 18. landsmøte i kjemi

Norsk Kjemisk Selskap

Norges Varemesse, Lillestrøm, Norway

October 17, 2007

## The one-electron density matrix

- Traditional SCF theory rests on the concept of molecular orbitals (MOs):

$$\mathbf{FC} = \mathbf{SC}\epsilon \leftarrow \text{Roothaan-Hall equations}$$

- cost of diagonalization is cubic—expensive for large systems
- ill suited to parallelization

- Fortunately, MOs are not needed: all information is contained in the density matrix

$$\mathbf{D} = \mathbf{C}_{\text{occ}}^{\text{T}} \mathbf{C}_{\text{occ}} \leftarrow \text{density matrix}$$

- it is sufficient to work in terms of the one-electron density matrix
- this is particularly advantageous if it can be compactly (sparsely) represented

- This approach is taken in our work, but then several questions arise:

- Q1: In what basis do we represent the density matrix?
- Q2: How do we optimize the density matrix, without recourse to orbitals?
- Q3: How do we calculate its response to perturbations?

- All tasks should be accomplished in terms of elementary matrix manipulations

- additions, multiplications, trace operations—no diagonalizations
- with sparsity, linear scaling is then achievable for large systems
- diagonalization-free methods are simpler to parallelize

## Overview

- the generation of a sparse, orthonormalized atomic-orbital (OAO) basis
  - diagonalization-free evaluation of  $\mathbf{S}^{\pm 1/2}$
  - multiplication-rich globally convergent iterative method
  - Jansik *et al.*, *J. Chem. Phys.* **126**, 124104 (2007)
- energy optimization for large systems
  - direct optimization of density matrix  $\mathbf{D}$  in OAO basis
  - multiplication-rich diagonalization-free self-consistent field (SCF) theory
  - Salek *et al.*, *J. Chem. Phys.* **126**, 114110 (2007)
- linear response theory for large systems
  - direct perturbation of density matrix  $\mathbf{D}$  in OAO basis
  - calculation of polarizabilities and excitation energies
  - Coriani *et al.* *J. Chem. Phys.* **126**, 154108 (2007)

## Part I: Orthonormalization of the AO basis

- For large systems, the **atomic-orbital (AO) basis** is ideal in the sense that it is sparse
  - unfortunately, many computational tasks are ill-conditioned in this basis
- This may be avoided by transforming to an **orthonormalized AO (OAO) basis**

$$\mathbf{Z}^T \mathbf{S} \mathbf{Z} = \mathbf{I}, \quad \mathbf{Z} = \begin{cases} \mathbf{S}^{-1/2} & \text{Löwdin} \\ \mathbf{U}^{-1} \text{ (upper triangular)} & \text{Cholesky} \end{cases}$$

- Millam and Scuseria, JCP **106**, 5569 (1997); Challacombe, JCP **110**, 2332 (1999)
- The **Löwdin basis** resembles the AO basis most closely (Carlson and Keller, 1957)
  - it should therefore retain sparsity to the greatest possible extent
  - however, its construction traditionally involves diagonalization
- The **Cholesky basis** has therefore been preferred instead
  - many uses in quantum chemistry (rank reduction techniques)
- However, since the Löwdin basis is (presumably) the sparsest of all OAO bases, it would be surprising if it could not be generated in linear time!
- We have recently presented a robust linear-scaling algorithm for Löwdin factorization
  - Jansik *et al.*, J. Chem. Phys. **126**, 124104 (2007)

## A globally convergent Newton–Schulz method for the inverse square root

- Application of Newton’s elementary method for root finding gives

$$Z^{-2} - S = 0 \quad \Rightarrow \quad Z_{n+1} = \frac{3}{2}Z_n - \frac{1}{2}Z_n^3 S \quad \leftarrow \text{Newton–Schulz iteration}$$

- $Z_n$  converges quadratically but not globally to  $S^{-1/2}$ ; no inversion!

- In 2004, Niklasson proposed the corresponding matrix iteration

$$\mathbf{Z}_{n+1} = \frac{1}{2}\mathbf{Z}_n \left( 3\mathbf{I} - \mathbf{Z}_n^T \mathbf{S} \mathbf{Z}_n \right) \quad \leftarrow \text{stabilized Newton–Schulz matrix iteration}$$

- higher-order formulas may also be derived
- for  $\mathbf{Z}_0 = \mathbf{I}$ , there are severe convergence restrictions  $\|\mathbf{S} - \mathbf{I}\|_2 < 1$  for such methods
- intended for refinement of approximate inverse square roots

- However, convergence is guaranteed with the scaled  $\lambda_{\min} \mathbf{S}$  that minimizes  $\|\lambda \mathbf{S} - \mathbf{I}\|_2$ :

$$\lambda_{\min} = \frac{2}{\epsilon_{\min} + \epsilon_{\max}} \quad \leftarrow \text{expensive}$$

- to avoid eigenvalues, we replace the 2-norm by a ratio of Frobenius norms

$$f(\lambda) = \sqrt{\frac{\text{Tr}(\lambda \mathbf{S} - \mathbf{I})^4}{\text{Tr}(\lambda \mathbf{S} - \mathbf{I})^2}} \quad \leftarrow \text{cheap lower bound to the 2-norm}$$

- the evaluation and minimization of  $f(\lambda)$  costs only **one** matrix multiplication

Number of iterations for convergence to  $10^{-10}$

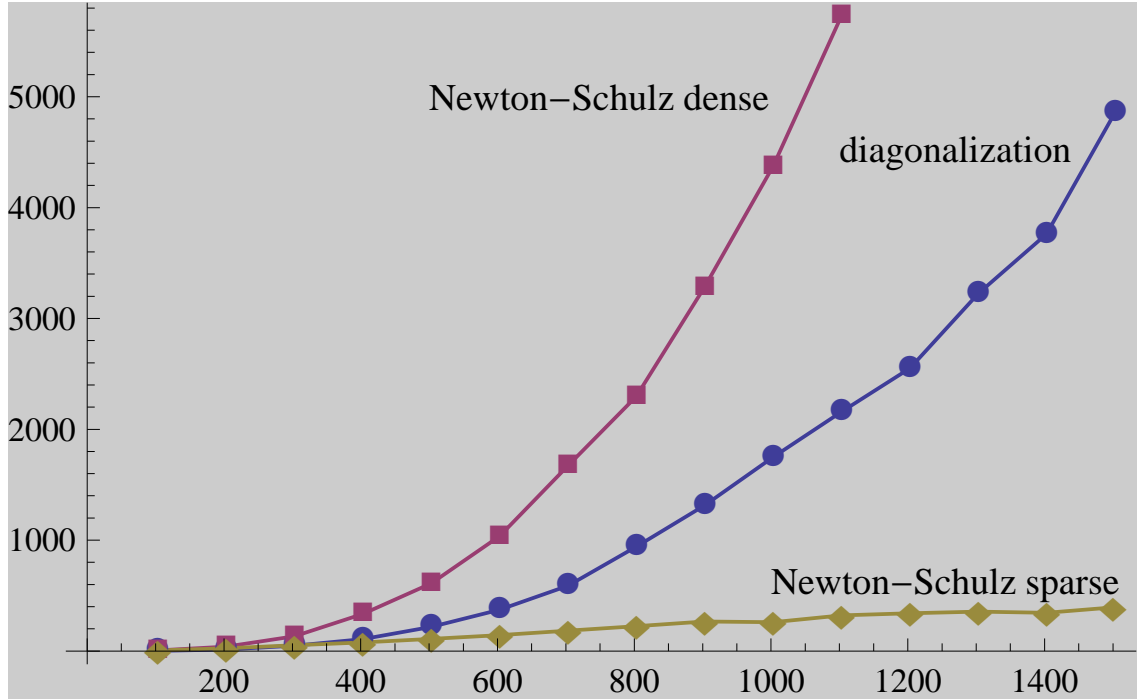
	$\ \mathbf{S} - \mathbf{I}\ _2$	$N^{(2)}$	$N^{(3)}$
<b>Fullerene C<sub>60</sub></b>	6.0	17	11
<b>Polysaccharide (438)</b>	5.2	14	10
<b>Polysaccharide (875)</b>	5.2	14	10
<b>DNA fragment (583)</b>	4.9	15	11
<b>PRC model (1082)</b>	5.7	16	11
<b>Red phosphorus (506)</b>	1.9	8	6
<b>Red phosphorus (810)</b>	4.7	14	10
<b>Crystobalite (984)</b>	4.4	13	9
<b>Tridymite (1314)</b>	5.2	15	11
<b>Random molecule (64)</b>	16.0	23	27

## Performance of the scaled Newton–Schulz method

- About 15 scaled Newton–Schulz iterations needed for convergence to  $10^{-10}$ 
  - less than 50 matrix multiplies
  - example: polyalanine peptide residues (6-31G) (cost vs. number of atoms)

- In dense-matrix algebra:
  - of cubic complexity
  - two–three times more expensive than diagonalization
  - easy to parallelize
  - beats diagonalization on 4–6 processors

- In sparse-matrix algebra:
  - of linear complexity
  - beats diagonalization for a few hundred atoms



- Transformation to orthonormal basis is an essential preparatory step to optimization

## Part II: Energy optimization

- In traditional SCF theory, we solve the **Roothaan–Hall (RH) eigenvalue problem**

$$\mathbf{FC} = \mathbf{SC}\epsilon, \quad \mathbf{D}_{\text{new}} = \mathbf{C}_{\text{occ}}\mathbf{C}_{\text{occ}}^{\text{T}}$$

- Its solution is equivalent to minimizing the sum of the occupied orbital energies

$$\epsilon_{\text{min}} = \min \sum_i \epsilon_i = \min_{\mathbf{D}} \text{Tr } \mathbf{D}\mathbf{F} \leftarrow \text{not all } \mathbf{D} \text{ variations valid!}$$

- The valid variations may be represented by an antisymmetric matrix  $\mathbf{X}$ :

$$\mathbf{D}(\mathbf{X}) = \exp(-\mathbf{X}\mathbf{S})\mathbf{D}\exp(\mathbf{S}\mathbf{X}), \quad \mathbf{X}^{\text{T}} = -\mathbf{X} \leftarrow \text{exponential parametrization}$$

- these transformations satisfy the symmetry, trace and idempotency conditions

$$\mathbf{D}(\mathbf{X})^{\text{T}} = \mathbf{D}(\mathbf{X}), \quad \text{Tr } \mathbf{D}(\mathbf{X})\mathbf{S} = N_{\text{el}}, \quad \mathbf{D}(\mathbf{X})\mathbf{S}\mathbf{D}(\mathbf{X}) = \mathbf{D}(\mathbf{X})$$

- Helgaker *et al.*, CPL **327**, 397 (2000); Head-Gordon *et al.*, MP **101**, 37 (2003)

- Each RH diagonalization is therefore equivalent to the **trace minimization problem**

$$\epsilon_{\text{min}} = \min_{\mathbf{X}} \text{Tr } \mathbf{D}(\mathbf{X})\mathbf{F}, \quad \mathbf{D}_{\text{new}} = \mathbf{D}(\mathbf{X}_{\text{min}})$$

- minimization performed with Newton's method



## The Roothaan–Hall Newton method

- To carry out the trace minimization, we expand the energy in powers of  $\mathbf{X}$ :

$$\begin{aligned}\text{Tr } \mathbf{F} \mathbf{D}(\mathbf{X}) &= \text{Tr } \mathbf{F} \exp(-\mathbf{X}) \mathbf{D} \exp(\mathbf{X}) \\ &= \text{Tr } \mathbf{F} \mathbf{D} + \text{Tr } \mathbf{F} [\mathbf{D}, \mathbf{X}] + \frac{1}{2} \text{Tr } \mathbf{F} [[\mathbf{D}, \mathbf{X}], \mathbf{X}] + \dots\end{aligned}$$

- Truncating at second order and setting the gradient to zero, we obtain the Newton step:

$$(\mathbf{F}^{\text{vv}} - \mathbf{F}^{\text{oo}}) \mathbf{X} + \mathbf{X} (\mathbf{F}^{\text{vv}} - \mathbf{F}^{\text{oo}}) = \mathbf{F}^{\text{vo}} - \mathbf{F}^{\text{ov}} \leftarrow \text{Roothaan–Hall Newton equations}$$


- upper indices indicate projections onto the occupied and virtual spaces

$$\mathbf{F} = \mathbf{F}^{\text{oo}} + \mathbf{F}^{\text{ov}} + \mathbf{F}^{\text{vo}} + \mathbf{F}^{\text{vv}}, \quad (\mathbf{P}_o = \mathbf{D}, \mathbf{P}_v = \mathbf{I} - \mathbf{D})$$

- Because of their large dimensions, the Newton equations cannot be solved directly
  - solution by **the preconditioned conjugate-gradient method** (typically 10 iterations)
  - elementary (sparse) matrix manipulations (typically less than 100 multiplications)
  - use of an orthonormal basis essential for rapid convergence
- A RH diagonalization corresponds to an exact minimization (many Newton steps)
  - however, a **partial minimization** will do (one RH Newton step is sufficient)
- Larsen *et al.*, J. Chem. Phys. **113** 8908 (2000); Shao *et al.*, *ibid.* **118** 6144 (2003)  
Sąłek *et al.*, *ibid.* **126** 114110 (2007)


## Performance of the Roothaan–Hall Newton method

- We have successfully avoided Roothaan–Hall diagonalization
  - minimization rather than the solution of a generalized eigenvalue problem
  - rapidly convergent: 50–100 sparse matrix multiplications needed
  - well suited to parallelization
- Linear scaling is obtained by employing sparse-matrix algebra
  - compressed sparse-row (CSR) representation of few-atom blocks



## SCF optimizations in small and large molecules

- Diagonalization can be avoided by solving Newton equations
- However, SCF convergence is typically more difficult in larger systems
  - small (or negative) HOMO–LUMO gaps and small Hessian eigenvalues in DFT
  - 6-31G HOMO–LUMO gap and lowest Hessian eigenvalue in alanine residue peptides



- The standard SCF scheme must be modified, making it more robust.

## The augmented Roothaan–Hall method

- SCF optimizations have two ingredients
  1. Roothaan–Hall minimization (diagonalization):  $\min_{\mathbf{X}} \text{Tr } \mathbf{D}(\mathbf{X})\mathbf{F}$
  2. DIIS-type averaging of density matrices:  $\bar{\mathbf{D}} = \sum_{i=0}^n c_i \mathbf{D}_i, \sum_i c_i = 1$
- This two-step procedure is prone to oscillations and does not always converge reliably
- A better approach is to minimize the total energy at each step  $\min_{\mathbf{X}} \text{Tr } E[\mathbf{D}(\mathbf{X})]$ 
  - this gives rise to an additional term in the Roothaan–Hall Newton equations

$$(\mathbf{F}^{\text{vv}} - \mathbf{F}^{\text{oo}}) \mathbf{X} + \mathbf{X} (\mathbf{F}^{\text{vv}} - \mathbf{F}^{\text{oo}}) + \frac{d(\mathbf{F}^{\text{ov}} - \mathbf{F}^{\text{vo}})}{d\mathbf{D}} [\mathbf{D}, \mathbf{X}] = \mathbf{F}^{\text{vo}} - \mathbf{F}^{\text{ov}}$$



- the new terms are very expensive (many recalculations of Fock/KS matrices)
- derivatives of the Fock/KS matrix (approximately) known in the old directions

$$\frac{d\mathbf{F}(\mathbf{D})}{d\mathbf{D}} \mathbf{D}_i \approx \mathbf{F}(\mathbf{D}_i) \leftarrow \text{quasi-Newton condition}$$

- we approximate the new term in the Roothaan–Hall equations in this manner
- The resulting method is called **the augmented Roothaan–Hall method (ARH)**
  - the ARH method is vastly superior to the standard SCF scheme



## The ARH method compared with RH/DIIS (I)

- Convergence of B3LYP calculations for a variety of molecules
  - cadmium–imidazole complex (circles) B3LYP/3-21G
  - 29-residue polyalanine peptide (crosses) B3LYP/6-31G
  - model B12 vitamin (triangles) HF/AhlrichsVDZ
- Energy-error on a logarithmic scale (ARH left, RH/DIIS right)



## The ARH method compared with RH/DIIS (II)

- Convergence of B3LYP calculations for a variety of molecules
  - cluster of 50 water molecules (full triangles) B3LYP/cc-pVTZ
  - insulin (empty circles) B3LYP/6-21G
  - 438-atom polysaccharide (full squares) LDA/6-31G
- Energy-error on a logarithmic scale (ARH left, RH/DIIS right)



### Part III: Response theory

- We consider a system described by the one-electron density matrix

$$\mathbf{D}(\mathbf{X}) = \exp(-\mathbf{X})\mathbf{D}\exp(\mathbf{X}) \quad (\mathbf{X} = \mathbf{0} \text{ for unperturbed system})$$

and define the Hessian and metric operators in terms of their transformations

$$\mathbf{E}^{[2]}(\mathbf{X}) = (\mathbf{F}^{\text{vv}} - \mathbf{F}^{\text{oo}})\mathbf{X} + \mathbf{X}(\mathbf{F}^{\text{vv}} - \mathbf{F}^{\text{oo}}) + \mathbf{G}^{\text{vo}}([\mathbf{D}, \mathbf{X}]) - \mathbf{G}^{\text{ov}}([\mathbf{D}, \mathbf{X}])$$

$$\mathbf{S}^{[2]}(\mathbf{X}) = \mathbf{X}^{\text{ov}} - \mathbf{X}^{\text{vo}}$$

- When perturbed by  $\mathbf{V}_\omega$  of frequency  $\omega$ , the system responds to first order as

$$\mathbf{E}^{[2]}(\mathbf{X}_\omega) - \omega\mathbf{S}^{[2]}(\mathbf{X}_\omega) = [\mathbf{D}, \mathbf{V}_\omega] \quad \leftarrow \text{linear response matrix equation}$$

- perturbed density matrix:  $\mathbf{D}_\omega = [\mathbf{D}, \mathbf{X}_\omega]$
- perturbed expectation values:  $\langle\langle \hat{A}; \hat{V}_\omega \rangle\rangle_\omega = \text{Tr } \mathbf{A} [\mathbf{D}, \mathbf{X}_\omega]$

- In the absence of a perturbation  $\mathbf{V}_\omega = \mathbf{0}$ , we obtain an eigenvalue equation

$$\mathbf{E}^{[2]}(\mathbf{X}_n) = \omega_n\mathbf{S}^{[2]}(\mathbf{X}_n) \quad \leftarrow \text{RPA matrix eigenvalue equation}$$

- transition density matrix:  $\mathbf{D}_{0n} = [\mathbf{D}, \mathbf{X}_n]$
- transition moments:  $\langle 0 | \hat{A} | n \rangle = \text{Tr } \mathbf{A} [\mathbf{D}, \mathbf{X}_n]$

- Coriani *et al.*, J. Chem. Phys. **126**, 154108 (2007)

- Ochsenfeld, Head-Gordon, Weber, Niklasson, and Challacombe (static properties)

## Solution of the response equations

- For the solution, it is useful to consider two response equations

$$\mathbf{E}^{(2)}(\mathbf{X}) - \omega \mathbf{S}^{(2)}(\mathbf{X}) = \mathbf{B} \leftarrow \text{full response equation}$$

$$\mathbf{E}_{\mathbf{F}}^{(2)}(\mathbf{X}) - \omega \mathbf{S}^{(2)}(\mathbf{X}) = \mathbf{B} \leftarrow \text{simplified response equation}$$

where the simplified Hessian is a good but cheap approximation to the full Hessian:

$$\mathbf{E}_{\mathbf{F}}^{(2)}(\mathbf{X}) = (\mathbf{F}^{\text{vv}} - \mathbf{F}^{\text{oo}})\mathbf{X} + \mathbf{X}(\mathbf{F}^{\text{vv}} - \mathbf{F}^{\text{oo}}) \leftarrow \text{no two-electron part}$$

– its eigenvalues are the orbital energy differences  $\epsilon_a - \epsilon_i$

- We wish solve the full response equations iteratively

$$\mathbf{R}_i = \mathbf{E}^{(2)}(\mathbf{X}_i) - \omega \mathbf{S}^{(2)}(\mathbf{X}_i) - \mathbf{B} \leftarrow \text{residual}$$

– new trial vectors are generated from the residual until it is sufficiently small

- For fast convergence, we precondition with the simplified response equation

$$\mathbf{E}_{\mathbf{F}}^{(2)}(\tilde{\mathbf{R}}_i) - \omega \mathbf{S}^{(2)}(\tilde{\mathbf{R}}_i) = \mathbf{R}_i$$

– in the MO basis,  $\mathbf{E}_{\mathbf{F}}^{(2)}$  is diagonal (orbital-energy differences) and solution is trivial

– in the OAO basis,  $\mathbf{E}_{\mathbf{F}}^{(2)}$  is nondiagonal and 5–20 iterations are required for solution

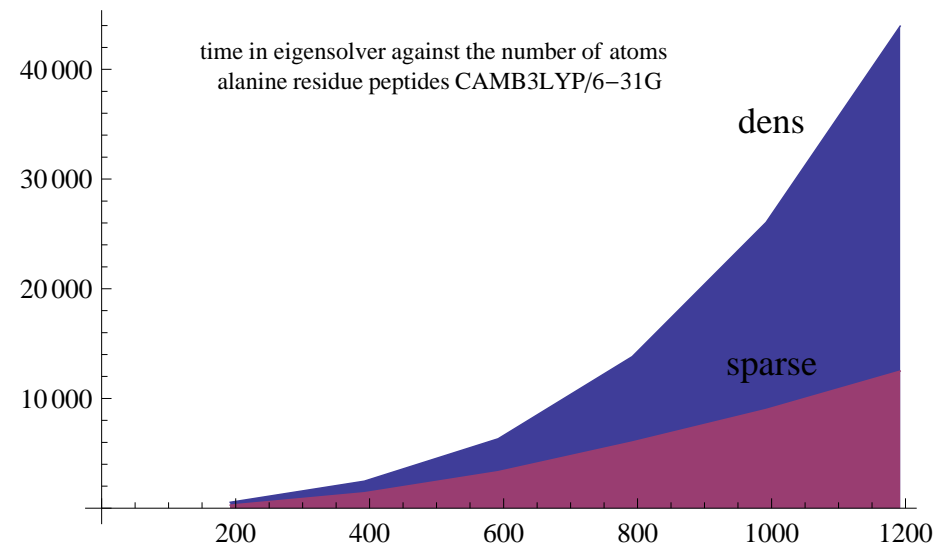


## Excitation energies


- The calculation of excitation energies (RPA) more difficult than that of polarizabilities
  - the eigenmatrices of the simplified Hessian provide good starting guesses:

$$\mathbf{E}_F^{(2)}(\mathbf{X}_{ia}) = (\epsilon_a - \epsilon_i)\mathbf{S}^{(2)}(\mathbf{X}_{ia}) \quad \text{where} \quad \mathbf{X}_{ia} = \mathbf{C}_i \mathbf{C}_a^T, \quad \begin{cases} \mathbf{F}^{oo} \mathbf{C}_i = \epsilon_i \mathbf{C}_i \\ \mathbf{F}^{vv} \mathbf{C}_a = \epsilon_a \mathbf{C}_a \end{cases}$$

- slow convergence of preconditioning equations (subspace problem nearly singular)
- CAMB3LYP/6-31G alanine residues
- cubic complexity with dense-matrix algebra
- linear complexity with sparse-matrix algebra
- preconditioning part dominates
- Fock/KS matrix construction dominates (not shown here)




CAMB3LYP/6-31G excitation energies of alanine residue peptides



- lowest excitation energy almost identical with lowest Hessian eigenvalue
- CAMB3LYP has 19% short-range and 65% long-range exact exchange

## B3LYP/6-31G excitation energies of alanine residue peptides


- B3LYP (with 20% exact exchange) behaves differently, in an unphysical manner
  - beyond 27 alanine residues, the lowest excitation energy becomes very small
  - DIIS with a small level shift converges to an excited state (first-order saddle point)




- for LDA and GGA functionals, the crossing occurs even earlier
  - local DFT functionals unable to account for long-range exchange in large molecules

## B3LYP/6-31G polyaniline HOMO–LUMO gap and lowest Hessian eigenvalue

- The B3LYP energy minimum behaves in a strange manner
  - the lowest Hessian eigenvalue goes through a minimum for 27 residues
  - the level-shifted DIIS solution behaves smoothly with increasing chain length




B3LYP/6-31G polyaniline polarizabilities



## Static polarizabilities

- To illustrate, we have calculated longitudinal polarizabilities of linear alkene chains
  - HF and DFT  $\alpha/N$  in 6-31G basis, plotted against the number of carbons  $N$



- poor long-range exchange leads to overpolarization


## Asymptotic behaviour of group polarizabilities

- How does the group polarizability converge towards the infinite limit?

$$\bar{\alpha}_{\infty} - \bar{\alpha}_N = eN^{-1} + \mathcal{O}(N^{-2}) \quad \text{Kudin *et al.*, JCP **122**, 134907 (2005)}$$

– this behaviour is universal, holding at all levels of theory

- Log–log plots of  $\bar{\alpha}_{\infty} - \bar{\alpha}_N$  for alkanes and alkenes:



– limit obtained by extrapolation  $\bar{\alpha}_{\infty} = (\alpha_N - \alpha_M)/(N - M)$

– straight lines of slope  $-1$  superimposed through the points at  $N = 350$

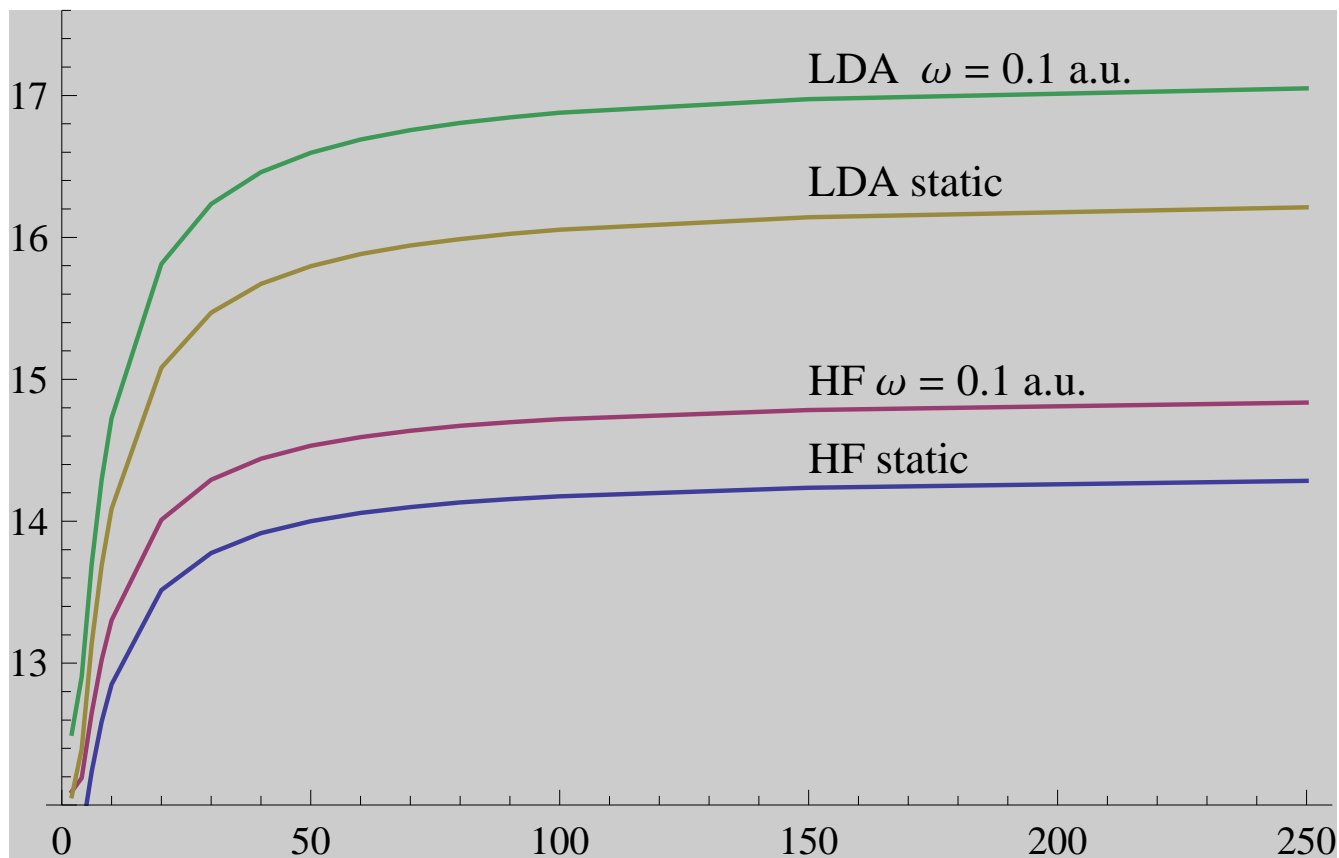
- The asymptotic region is reached with  $\text{C}_{30}\text{H}_{62}$  (alkanes) and  $\text{C}_{60}\text{H}_{62}$  (alkenes)

– alkane  $\bar{\alpha}_{\infty}$  predicted to within 1% from  $\text{C}_{30}\text{H}_{62}$

– alkene  $\bar{\alpha}_{\infty}$  predicted to within 1% from  $\text{C}_{60}\text{H}_{62}$  for HF and from  $\text{C}_{150}\text{H}_{152}$  for LDA

## Frequency-dependent polarizabilities

- HF and LDA longitudinal polarizabilities in linear alkane chains
  - $\alpha/N$  in 6-31G basis, plotted against the number of carbons  $N$



- long-range exchange is less important for the saturated alkanes
- LDA overestimates dispersion as well as the static limit



## Summary

- We have considered several computational tasks in self-consistent field theory
  - orthogonalization of the atomic orbitals
  - solution of the Roothaan–Hall eigenvalue problem
  - linear response theory
- For sufficiently sparse AO matrices, all these tasks can be carried out in linear time
  - the Fock/KS matrix is reevaluated no more often than in MO theory
  - the remaining work is rich in matrix multiplication
  - well suited to parallelization
- We have not considered the global convergence of the SCF method
  - standard methods may not converge to the ground state
  - better global optimization schemes must be developed
  - standard functionals often perform poorly for large systems
- Standard functionals may not be suitable for large systems
  - local functionals do not recover long-range correlation and exchange effects
  - the inclusion of exact exchange important