Multilevel domain decomposition for electronic structure calculations

G. Bencteux, M. Barrault (EDF),
W. Hager (Univ. of Florida),
E. Cancès, C. Le Bris (CERMICS)

PMAA'06, Rennes, September 2006

Outline

Outline

The electronic problem

A molecular system made of N electrons and M nuclei.
Force field and energy are given by the solution of the *electronic problem* $\mathcal{P}$:

$$\inf \left\{ E^{RHF,KS}(\psi_1, \ldots, \psi_N),\ \psi_i \in H^1(\mathbf{R}^3), \right.$$
$$\left. \int_{\mathbf{R}^3} \psi_i \psi_j = \delta_{ij},\ \forall\ 1 \leq i,j \leq N \right\}$$

- $(\psi_i)_{1 \leq i \leq N}$: *orbitals* of the system
- $E^{RHF,KS}(\psi_1, \ldots, \psi_N)$: energy functional depending on the model (RHF, KS) and on the nuclei

Electronic density computation

$\{\chi_\mu\}_{1\leq\mu\leq N_b}$ a Galerkin basis : $\psi_i(x) = \sum_{\mu=1}^{N_b} C_{\mu i}\ \chi_\mu(x)$

$$(\mathcal{P}) \Rightarrow (\mathcal{P}_{dis}) \left\{ \begin{array}{ccc} F(D)C & = & SC\Lambda \\ C^t SC & = & I_N \\ D & = & CC^t \end{array} \right. , \ C \in \mathcal{M}^{N_b,N}(\mathbf{R}),\ \Lambda \in \mathcal{M}_S^N(\mathbf{R})$$

- D : *density matrix*, S : *overlapping matrix*
- Λ : *Lagrange multipliers matrix*
- $\rho(x) = \sum_{k=1}^{N} |\psi_k(x)|^2 = \sum_{1\leq i,j\leq N_b} D_{ij}\chi_i(x)\chi_j(x)$

  ρ : electronic density

Electronic density computation

- $\mathcal{P}_{dis}$ is solved with a fixed-point strategy
- at each step, we seek $D = CC^t$, with :

$$\left\{ \begin{array}{rcl} FC & = & SC\Lambda \\ C^tSC & = & I_N \end{array} \right. , \; C \in \mathcal{M}^{N_b,N}(\mathbf{R}), \; \Lambda \in \mathcal{M}_S^N(\mathbf{R})$$

- Λ : can be chosen diagonal, with the N lowest eigenvalues of F
- At each step, find orthogonal projector on N lowest eigenvectors of F

Methods that scale linearly with N (localized basis set)

1. Linear subproblem rewriting
   - Projection methods (FOE, ...)
   - Variational methods (DMM, ...)

   $\Rightarrow$ matrix products
2. Reaching linear complexity
   - Sparsity of F
   - Sparsity of D and of some C $\Rightarrow$ sparse matrix products
   - Truncation of C and D during the algorithm

- Inadequacies
  - Sensitivity to initial guess, truncation strategy
  - Fermi energy computation
  - Additional approximation if $S \neq I_{N_b}$
  - No existing rigorous domain decomposition method

Goal

To develop a method in the frame of localized Galerkin basis that :

- scales linearly with the number of electrons (memory and CPU)
- is strongly parallel

$\Longrightarrow$ Domain decomposition method

Outline

Linear subproblem re-formulation

If $S = I_{n_b}$, linear subproblem is equivalent to :

$$\inf \left\{ \mathrm{Tr}\left(FCC^t\right),\ C^tC = I_N,\ C \in \mathcal{M}^{N_b,N}(\mathbf{R}) \right\}.$$

If we decompose C in p blocks $C_i \in \mathcal{M}^{N_b,m_i}(\mathbf{R})$ such that $C = |C_1 \dots C_p|$:

$$\mathrm{Tr}\left(FCC^t\right) = \sum_{i=1}^{p} \mathrm{Tr}\left(FC_iC_i^t\right)$$

Electron locality = sparsity of C

$\Rightarrow$ Decomposition on the electrons

An approximating problem

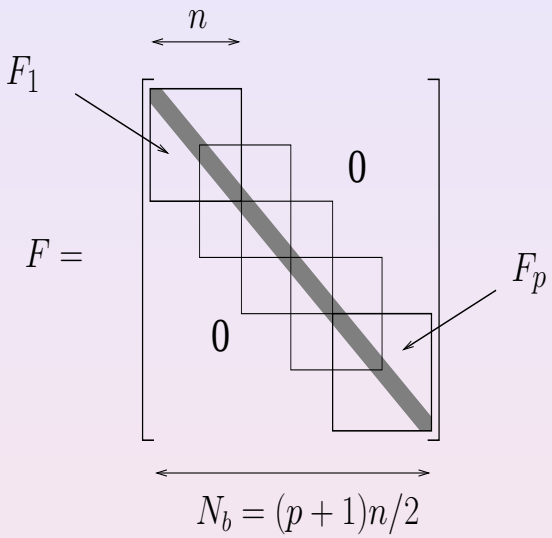

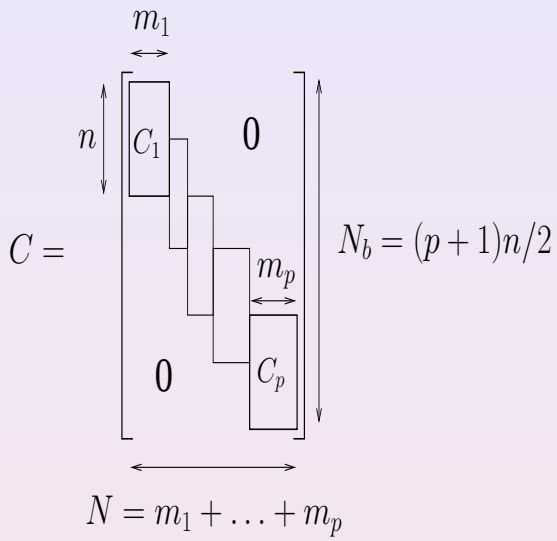

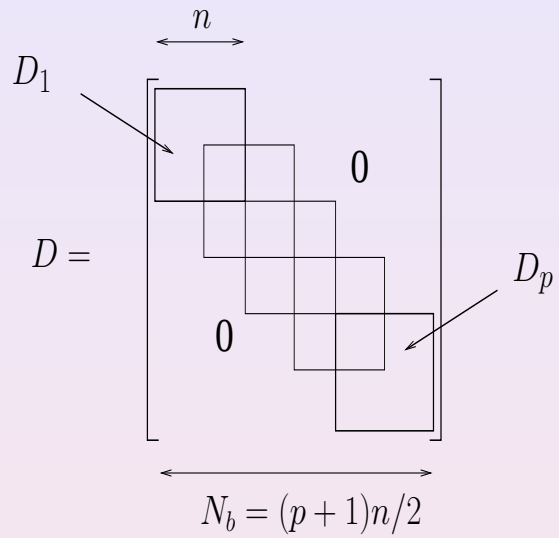

$$\inf \left\{ \sum_{i=1}^{p} \mathrm{Tr}\left(F_i C_i C_i^t\right),\ C_i \in \mathcal{M}^{n,m_i}(\mathbf{R}),\ m_i \in \mathbf{N}^\star,\ C_i^t C_i = I_{m_i} \quad \forall\, 1 \le i \le p, \right.$$

$$\left. C_i^t T C_{i+1} = 0 \quad \forall\, 1 \le i \le p-1, \quad \sum_{i=1}^{p} m_i = N \right\}.$$

Parallel local step

$$C_i = \text{arginf}\left\{ \text{Tr}\left(F_i C C^t\right),\ C^t C = I_{m_i},\ C^t T C_{i+1} = 0 \ \text{ and } C_{i-1}^t T C = 0 \right\}$$

1. compute P_i the orthogonal projector on
   $$V_i = \left\{ C\ , C^t T C_{i+1} = 0 \ \text{ and } C_{i-1}^t T C = 0 \right\}$$
2. diagonalize $P_i^T F_i P_i$

Computations of C_i and C_{i+2} are independant

A global step is needed

Local step gives $\tilde{C}_1$ et $\tilde{C}_2$ such that ($n = 2$, $m_1 = m_2 = 1$)

$$\left| \begin{array}{l} \tilde{C}_1 = \mathrm{Arginf}\Big\{ \mathrm{Tr}\big(F_1 C C^t\big),\ C^t C = 1,\ C^t C_2^n = 0 \Big\}, \\ \tilde{C}_2 = \mathrm{Arginf}\Big\{ \mathrm{Tr}\big(F_2 C C^t\big),\ C^t C = 1,\ \tilde{C}_1^t C = 0 \Big\}. \end{array} \right.$$

Each step generates Lagrange multipliers : $\Lambda_{21} \neq \Lambda_{12}^t$

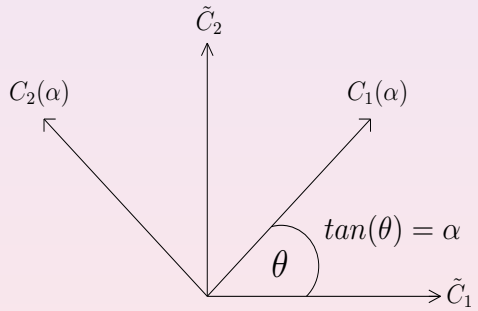

Mix $\tilde{C}_1$ and $\tilde{C}_2$ and preserve orthogonality
$\rightarrow$ rotation of $\{\tilde{C}_1, \tilde{C}_2\}$

$$C_1(\alpha) = \frac{\tilde{C}_1 + \alpha \tilde{C}_2}{\sqrt{1+\alpha^2}},\ C_2(\alpha) = \frac{\tilde{C}_2 - \alpha \tilde{C}_1}{\sqrt{1+\alpha^2}}.$$

- Global step : find α that minimizes
  $\mathrm{Tr}\big(F_1 C_1(\alpha) C_1^t(\alpha)\big) + \mathrm{Tr}\big(F_2 C_2(\alpha) C_2^t(\alpha)\big)$

Global step

- Property

$$\text{Let } f(\alpha) = \sum_{i=1}^{2} \text{Tr}\Big(F_i C_i(\alpha) C_i^t(\alpha)\Big), \quad \frac{\partial f}{\partial \alpha}(\alpha = 0) = 0 \Rightarrow \Lambda_{21} = \Lambda_{12}^t$$

- Implementation : Newton or Conjugated Gradient strategy
- Parallelism
  - blocks are gathered in overlapping groups

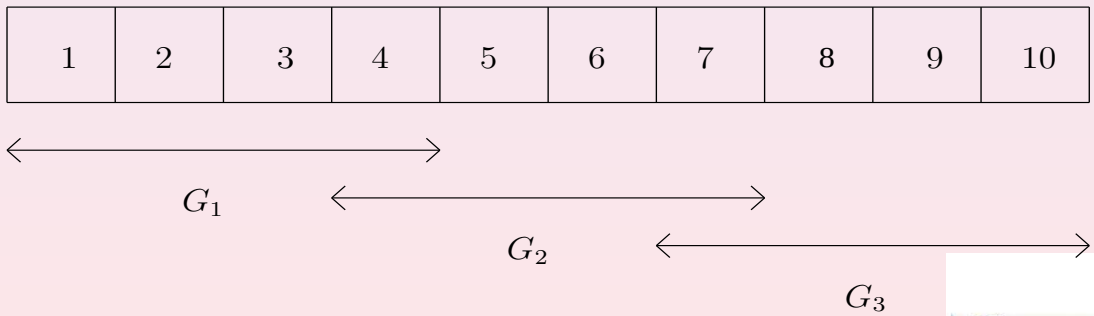

  - non overlapping groups are independant

The complete algorithm

- Repeat until convergence :
  - local steps for blocks 1, 3, 5, ...
  - local steps for blocks 2, 4, 6, ...
  - global steps for groups $\{1, 2\}$, $\{3, 4\}$, ...
  - global steps for groups $\{2, 3\}$, $\{4, 5\}$, ...

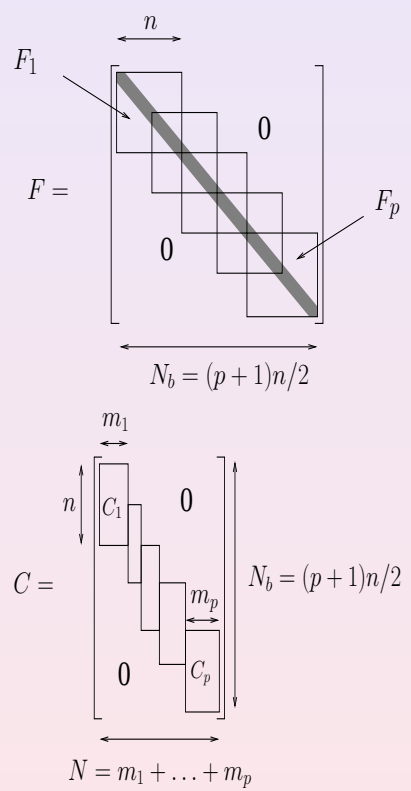

Outline

Studied systems

Insulating linear polymers . Restricted Hartree-Fock.

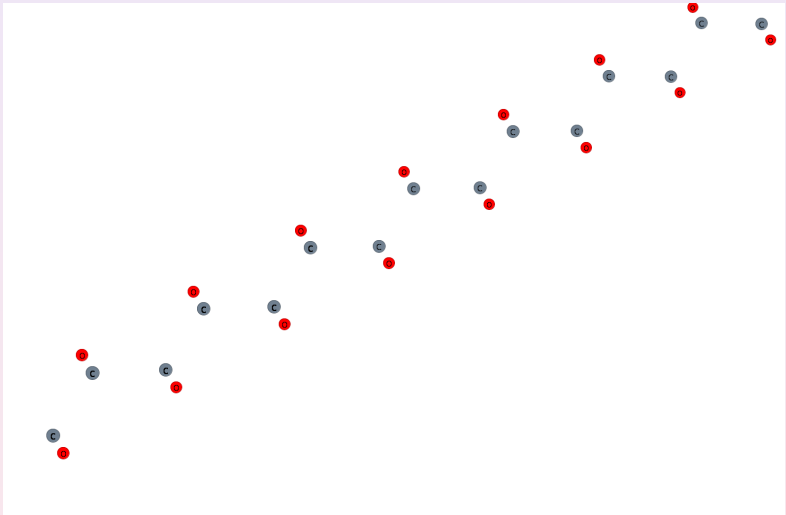

$H(CO)_nH$

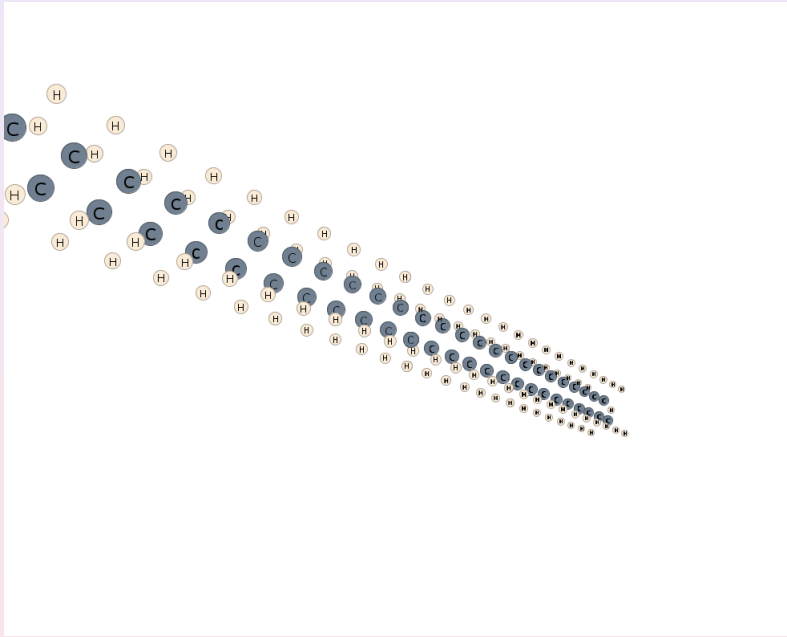

polyethylen $CH_3(CH_2)_nCH_3$

Linear scaling is nearly achieved

$H(CO)_nH$

DMM = Density Matrix Minimisation [Li, Nunes, Vanderbilt, 1993]

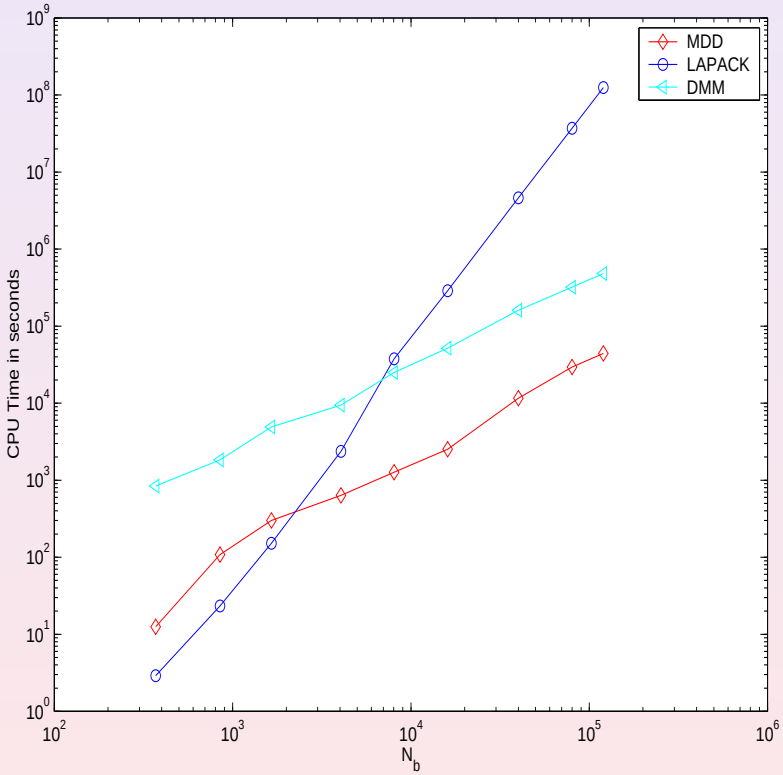

CPU

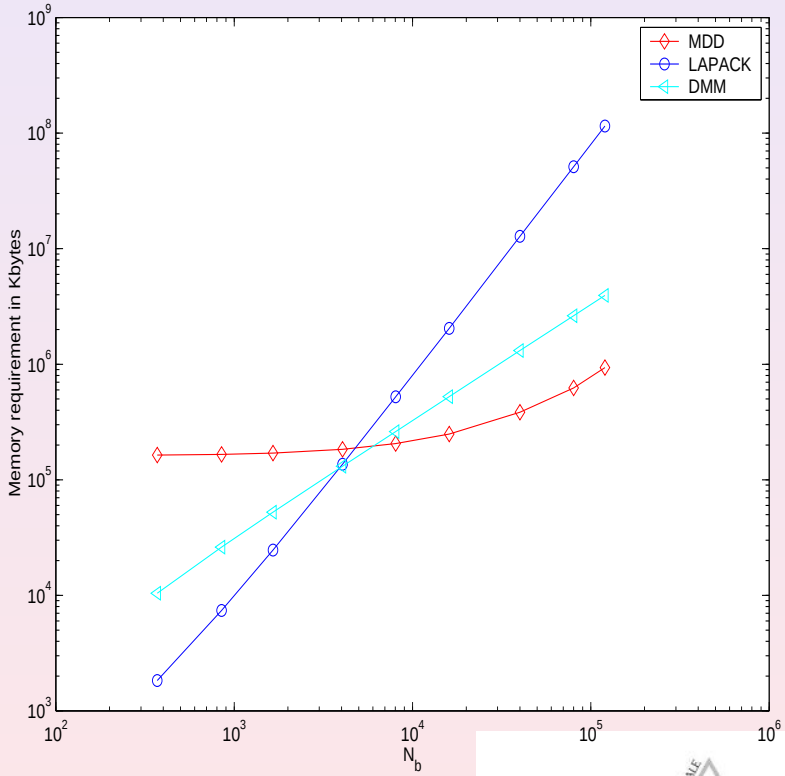

memory

Scalability

Preliminary results on a Xeon cluster : bi-processors, 2.8 GHz, ethernet 100 Mbs

$H(CO)_{1027}H$, 128 blocks

Procs	1	8	16	32
CPU (s)	1661	254	165	86
Sp. Up	1	6.5	10	19

$CH_3(CH_2)_{3342}CH_3$, 128 bl.

Procs	1	8	32
CPU (s)	28275	3661	1694
Sp. Up	1	7.7	17

- message passing parallelism (MPI)
- no work on load balancing
- memory conflicts in the case of polyethylen

Outline

Conclusion and future developments

A method has been developped to replace diagonalization step in electronic structure computations with localized basis sets.

This method has been tested on matrices coming from Hartree-Fock simulation of linear insulating polymers. It shows :

- linear scaling of CPU and memory with the number of atoms
- the possibility to treat large systems on parallel machines

Next developments :

- interaction with non linear SCF loop
- extension of the software to 2D and 3D molecular systems
- extension to conducting materials