Applications of Matrix Functions
Part III: Quantum Chemistry

Michele Benzi

Emory University
Department of Mathematics and Computer Science
Atlanta, GA 30322, USA

Prologue

The main purpose of this lecture is to present a rigorous mathematical theory for a class of methods, called $O(N)$ methods, that are being developed by computational physicists and chemists for the solution of the electronic structure problem, which is fundamental to quantum chemistry, solid state physics, biology, etc.

Prologue

The main purpose of this lecture is to present a rigorous mathematical theory for a class of methods, called $O(N)$ methods, that are being developed by computational physicists and chemists for the solution of the electronic structure problem, which is fundamental to quantum chemistry, solid state physics, biology, etc.

The theory is based on general results on the decay in the entries of functions of sparse matrices. In particular, one needs to study the asymptotic behavior of the off-diagonal matrix elements for $N \to \infty$.

Prologue

The main purpose of this lecture is to present a rigorous mathematical theory for a class of methods, called $O(N)$ methods, that are being developed by computational physicists and chemists for the solution of the electronic structure problem, which is fundamental to quantum chemistry, solid state physics, biology, etc.

The theory is based on general results on the decay in the entries of functions of sparse matrices. In particular, one needs to study the asymptotic behavior of the off-diagonal matrix elements for $N \longrightarrow \infty$.

While this work is primarily theoretical, the theory can be used to construct better algorithms for electronic structure computations. Some of our techniques are already used in a code called FreeON developed by a group headed by Matt Challacombe at Los Alamos National Laboratory.

References

This lecture is based in parts on the results contained in the following papers:

M. Benzi and G. H. Golub, *Bounds for the entries of matrix functions with applications to preconditioning*, BIT, 39 (1999), pp. 417–438.

M. Benzi and N. Razouk, *Decay bounds and $O(n)$ algorithms for approximating functions of sparse matrices*, Electr. Trans. Numer. Anal., 28 (2007), pp. 16–39.

M. Benzi, P. Boito and N. Razouk, *Decay properties of spectral projectors with applications to electronic structure*, SIAM Review, 55 (2013), pp. 3–64.

Overview

Overview

Overview

Overview

1. The electronic structure problem
2. Density matrices
3. $O(N)$ methods
4. A mathematical foundation for $O(N)$ methods

Overview

1. The electronic structure problem
2. Density matrices
3. $O(N)$ methods
4. A mathematical foundation for $O(N)$ methods
5. $O(N)$ approximation of functions of sparse matrices

Overview

Overview

Overview

The Electronic Structure Problem

A fundamental problem in quantum chemistry and solid state physics is to determine the electronic structure of (possibly large) atomic and molecular systems.

The Electronic Structure Problem

A fundamental problem in quantum chemistry and solid state physics is to determine the electronic structure of (possibly large) atomic and molecular systems.

The problem amounts to computing the ground state (smallest eigenvalue and corresponding eigenfunction) of the many-body quantum-mechanical Hamiltonian (Schrödinger operator), $\mathcal{H}$.

The Electronic Structure Problem

A fundamental problem in quantum chemistry and solid state physics is to determine the electronic structure of (possibly large) atomic and molecular systems.

The problem amounts to computing the ground state (smallest eigenvalue and corresponding eigenfunction) of the many-body quantum-mechanical Hamiltonian (Schrödinger operator), $\mathcal{H}$.

Variationally, we want to minimize the Rayleigh quotient:

$$E_0 = \min_{\Psi \neq 0} \frac{\langle \mathcal{H}\Psi, \Psi \rangle}{\langle \Psi, \Psi \rangle} \quad \text{and} \quad \Psi_0 = \text{argmin}_{\Psi \neq 0} \frac{\langle \mathcal{H}\Psi, \Psi \rangle}{\langle \Psi, \Psi \rangle}$$

where $\langle \cdot, \cdot \rangle$ denotes the L^2 inner product.

The Electronic Structure Problem

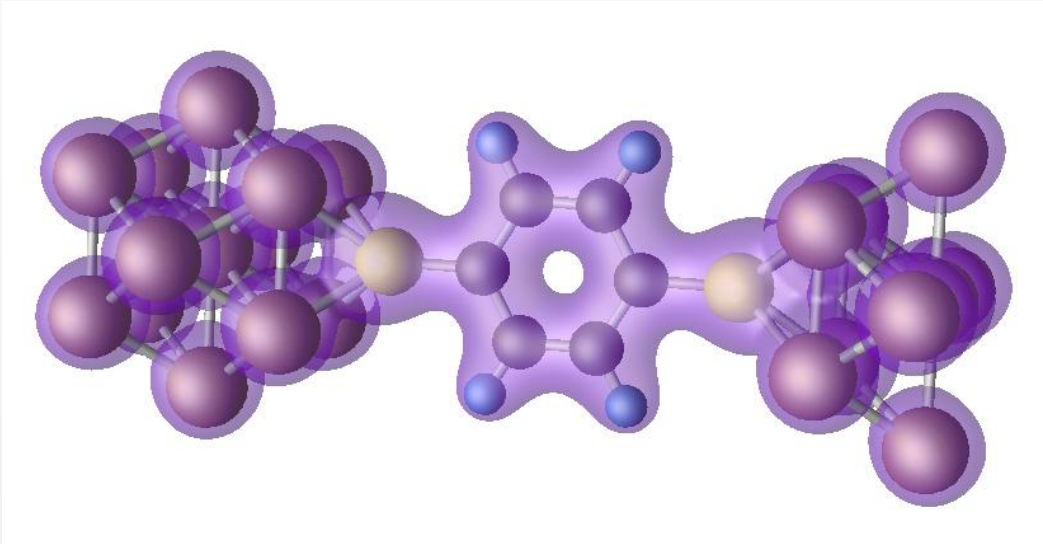

The Electronic Structure Problem

In the Born-Oppenheimer approximation, the many-body Hamiltonian (in atomic units) is given by

$$\mathcal{H} = \sum_{i=1}^{N} \left(-\frac{1}{2}\Delta_i - \sum_{j=1}^{M} \frac{Z_j}{|\mathbf{x}_i - \mathbf{r}_j|} + \sum_{j \neq i}^{N} \frac{1}{|\mathbf{x}_i - \mathbf{x}_j|} \right)$$

where N = number of electrons and M = number of nuclei in the system.

The Electronic Structure Problem

In the Born-Oppenheimer approximation, the many-body Hamiltonian (in atomic units) is given by

$$\mathcal{H} = \sum_{i=1}^{N} \left(-\frac{1}{2}\Delta_i - \sum_{j=1}^{M} \frac{Z_j}{|\mathbf{x}_i - \mathbf{r}_j|} + \sum_{j \neq i}^{N} \frac{1}{|\mathbf{x}_i - \mathbf{x}_j|} \right)$$

where $N =$ number of electrons and $M =$ number of nuclei in the system. The operator $\mathcal{H}$ acts on a suitable subspace $D(\mathcal{H}) \subset L^2(\mathbb{R}^{3N})$, the antisymmetrized tensor product of N copies of $H^1(\mathbb{R}^3)$:

$$D(\mathcal{H}) = H^1(\mathbb{R}^3) \wedge \ldots \wedge H^1(\mathbb{R}^3)\,.$$

The Electronic Structure Problem

In the Born-Oppenheimer approximation, the many-body Hamiltonian (in atomic units) is given by

$$\mathcal{H} = \sum_{i=1}^{N} \left(-\frac{1}{2}\Delta_i - \sum_{j=1}^{M} \frac{Z_j}{|\mathbf{x}_i - \mathbf{r}_j|} + \sum_{j \neq i}^{N} \frac{1}{|\mathbf{x}_i - \mathbf{x}_j|} \right)$$

where $N =$ number of electrons and $M =$ number of nuclei in the system. The operator $\mathcal{H}$ acts on a suitable subspace $D(\mathcal{H}) \subset L^2(\mathbb{R}^{3N})$, the antisymmetrized tensor product of N copies of $H^1(\mathbb{R}^3)$:

$$D(\mathcal{H}) = H^1(\mathbb{R}^3) \wedge \ldots \wedge H^1(\mathbb{R}^3)\,.$$

This is because electrons are Fermions and therefore subject to Pauli's Exclusion Principle. Hence, the wavefunction must be antisymmetric:

$$\Psi(\mathbf{x}_1, \ldots, \mathbf{x}_i, \ldots, \mathbf{x}_j, \ldots, \mathbf{x}_N) = -\Psi(\mathbf{x}_1, \ldots, \mathbf{x}_j, \ldots, \mathbf{x}_i, \ldots, \mathbf{x}_N).$$

The Electronic Structure Problem

In the Born-Oppenheimer approximation, the many-body Hamiltonian (in atomic units) is given by

$$\mathcal{H} = \sum_{i=1}^{N} \left(-\frac{1}{2}\Delta_i - \sum_{j=1}^{M} \frac{Z_j}{|\mathbf{x}_i - \mathbf{r}_j|} + \sum_{j \neq i}^{N} \frac{1}{|\mathbf{x}_i - \mathbf{x}_j|} \right)$$

where $N =$ number of electrons and $M =$ number of nuclei in the system. The operator $\mathcal{H}$ acts on a suitable subspace $D(\mathcal{H}) \subset L^2(\mathbb{R}^{3N})$, the antisymmetrized tensor product of N copies of $H^1(\mathbb{R}^3)$:

$$D(\mathcal{H}) = H^1(\mathbb{R}^3) \wedge \ldots \wedge H^1(\mathbb{R}^3)\,.$$

This is because electrons are Fermions and therefore subject to Pauli's Exclusion Principle. Hence, the wavefunction must be antisymmetric:

$$\Psi(\mathbf{x}_1, \ldots, \mathbf{x}_i, \ldots, \mathbf{x}_j, \ldots, \mathbf{x}_N) = -\Psi(\mathbf{x}_1, \ldots, \mathbf{x}_j, \ldots, \mathbf{x}_i, \ldots, \mathbf{x}_N).$$

NOTE: To simplify notation, spin is ignored here.

The Electronic Structure Problem

Unless N is very small, the curse of dimensionality makes this problem intractable.

The Electronic Structure Problem

Unless N is very small, the curse of dimensionality makes this problem intractable.

In order to make the problem tractable various approximations have been introduced, most notably:

The Electronic Structure Problem

Unless N is very small, the curse of dimensionality makes this problem intractable.

In order to make the problem tractable various approximations have been introduced, most notably:

- Wavefunction methods (e.g., Hartree-Fock)

The Electronic Structure Problem

Unless N is very small, the curse of dimensionality makes this problem intractable.

In order to make the problem tractable various approximations have been introduced, most notably:

- Wavefunction methods (e.g., Hartree-Fock)
- Density Functional Theory (e.g., Kohn-Sham; Nobel Prize, 1998)

The Electronic Structure Problem

Unless N is very small, the curse of dimensionality makes this problem intractable.

In order to make the problem tractable various approximations have been introduced, most notably:

- Wavefunction methods (e.g., Hartree-Fock)
- Density Functional Theory (e.g., Kohn-Sham; Nobel Prize, 1998)
- Hybrid methods

The Electronic Structure Problem

Unless N is very small, the curse of dimensionality makes this problem intractable.
In order to make the problem tractable various approximations have been introduced, most notably:

- Wavefunction methods (e.g., Hartree-Fock)
- Density Functional Theory (e.g., Kohn-Sham; Nobel Prize, 1998)
- Hybrid methods

In these approximations the original, linear eigenproblem $H\Psi = E\Psi$ for the many-electrons Hamiltonian is replaced by a nonlinear one-particle eigenproblem of the form

$$F(\psi_i) = \lambda_i \psi_i, \quad \langle \psi_i, \psi_j \rangle = \delta_{ij}, \quad 1 \leq i, j \leq N$$

where $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_N$. This problem is nonlinear because the operator F depends nonlinearly on the ψ_i.

The Electronic Structure Problem

Roughly speaking, in DFT the idea is to consider a single electron moving in the electric field generated by the nuclei and by some average distribution of the other electrons. Starting with an initial guess of the charge density, a potential is formed and the corresponding one-particle eigenproblem is solved; the resulting charge density is used to define the new potential, and so on until the charge density no longer changes appreciably.

The Electronic Structure Problem

Roughly speaking, in DFT the idea is to consider a single electron moving in the electric field generated by the nuclei and by some average distribution of the other electrons. Starting with an initial guess of the charge density, a potential is formed and the corresponding one-particle eigenproblem is solved; the resulting charge density is used to define the new potential, and so on until the charge density no longer changes appreciably.

More formally, DFT reformulates the problem so that the unknown function is the electronic density

$$\rho(\mathbf{x}) = N \int_{\mathbb{R}^{3(N-1)}} |\Psi(\mathbf{x}, \mathbf{x}_2, \ldots, \mathbf{x}_N)|^2 d\mathbf{x}_2 \cdots d\mathbf{x}_N \,,$$

a scalar field on $\mathbb{R}^3$.

The Electronic Structure Problem

Roughly speaking, in DFT the idea is to consider a single electron moving in the electric field generated by the nuclei and by some average distribution of the other electrons. Starting with an initial guess of the charge density, a potential is formed and the corresponding one-particle eigenproblem is solved; the resulting charge density is used to define the new potential, and so on until the charge density no longer changes appreciably.

More formally, DFT reformulates the problem so that the unknown function is the electronic density

$$\rho(\mathbf{x}) = N \int_{\mathbb{R}^{3(N-1)}} |\Psi(\mathbf{x}, \mathbf{x}_2, \ldots, \mathbf{x}_N)|^2 d\mathbf{x}_2 \cdots d\mathbf{x}_N \,,$$

a scalar field on $\mathbb{R}^3$.

The function ρ minimizes a certain functional, the form of which is not known explicitly.

The Electronic Structure Problem

Various forms of the density functional have been proposed, the most successful being the Kohn-Sham model:

$$I_{KS}(\rho) = \inf \left\{ T_{KS} + \int_{\mathbb{R}^3} \rho V \, \mathrm{d}\mathbf{x} + \frac{1}{2} \int_{\mathbb{R}^3} \int_{\mathbb{R}^3} \frac{\rho(\mathbf{x})\rho(\mathbf{y})}{|\mathbf{x}-\mathbf{y}|} \mathrm{d}\mathbf{x}\mathrm{d}\mathbf{y} + E_{xc}(\rho) \right\},$$

where $\rho(\mathbf{x}) = \sum_{i=1}^{N} |\psi_i(\mathbf{x})|^2$, $T_{KS} = \frac{1}{2} \sum_{i=1}^{N} \int_{\mathbb{R}^3} |\nabla \psi_i|^2 \, \mathrm{d}\mathbf{x}$ is the kinetic energy term, V denotes the Coulomb potential, and E_{xc} denotes the exchange term that takes into account the interaction between electrons. The infimum above is taken over all functions $\psi_i \in H^1(\mathbb{R}^3)$ such that $\langle \psi_i, \psi_j \rangle = \delta_{ij}$, where $1 \leq i, j \leq N$ and $\sum_{i=1}^{N} |\psi_i(\mathbf{x})|^2 = \rho$.

The Electronic Structure Problem

Various forms of the density functional have been proposed, the most successful being the Kohn-Sham model:

$$I_{KS}(\rho) = \inf \left\{ T_{KS} + \int_{\mathbb{R}^3} \rho V \, \mathrm{d}\mathbf{x} + \frac{1}{2} \int_{\mathbb{R}^3} \int_{\mathbb{R}^3} \frac{\rho(\mathbf{x})\rho(\mathbf{y})}{|\mathbf{x}-\mathbf{y}|} \mathrm{d}\mathbf{x}\mathrm{d}\mathbf{y} + E_{xc}(\rho) \right\},$$

where $\rho(\mathbf{x}) = \sum_{i=1}^{N} |\psi_i(\mathbf{x})|^2$, $T_{KS} = \frac{1}{2} \sum_{i=1}^{N} \int_{\mathbb{R}^3} |\nabla \psi_i|^2 \, \mathrm{d}\mathbf{x}$ is the kinetic energy term, V denotes the Coulomb potential, and E_{xc} denotes the exchange term that takes into account the interaction between electrons. The infimum above is taken over all functions $\psi_i \in H^1(\mathbb{R}^3)$ such that $\langle \psi_i, \psi_j \rangle = \delta_{ij}$, where $1 \leq i, j \leq N$ and $\sum_{i=1}^{N} |\psi_i(\mathbf{x})|^2 = \rho$.

This I_{KS} is minimized with respect to ρ. Note that ρ, being the electron density, must satisfy $\rho > 0$ and $\int_{\mathbb{R}^3} \rho \, \mathrm{d}\mathbf{x} = N$.

The Electronic Structure Problem

The Euler–Lagrange equations for this variational problem are the Kohn–Sham equations:

$$F(\rho)\psi_i = \lambda_i \psi_i, \quad \langle \psi_i, \psi_j \rangle = \delta_{ij} \quad (1 \leq i, j \leq N)$$

with

$$F(\rho) = -\frac{1}{2}\Delta + U(\mathbf{x}, \rho)$$

where U denotes a (complicated) potential, and $\rho = \sum_{i=1}^{N} |\psi_i(\mathbf{x})|^2$.

The Electronic Structure Problem

The Euler–Lagrange equations for this variational problem are the Kohn–Sham equations:

$$F(\rho)\psi_i = \lambda_i \psi_i, \quad \langle \psi_i, \psi_j \rangle = \delta_{ij} \quad (1 \leq i, j \leq N)$$

with

$$F(\rho) = -\frac{1}{2}\Delta + U(\mathbf{x}, \rho)$$

where U denotes a (complicated) potential, and $\rho = \sum_{i=1}^{N} |\psi_i(\mathbf{x})|^2$.

Hence, the original intractable linear eigenvalue problem for the many-body Hamiltonian is reduced to a tractable nonlinear eigenvalue problem for a single-particle Hamiltonian.

The Electronic Structure Problem

The nonlinear problem can be solved by a 'self-consistent field' (SCF) iteration, leading to a sequence of linear eigenproblems

$$F^{(k)}\psi_i^{(k)} = \lambda_i^{(k)}\psi_i^{(k)}, \quad \langle \psi_i^{(k)}, \psi_j^{(k)} \rangle = \delta_{ij}, \quad k = 1, 2, \ldots$$

$(1 \leq i, j \leq N)$, where each $F^{(k)} = -\frac{1}{2}\Delta + U^{(k)}$ is a one-electron linearized Hamiltonian:

$$U^{(k)} = U^{(k)}(\mathbf{x}, \rho^{(k-1)}), \quad \rho^{(k-1)} = \sum_{i=1}^{N} |\psi_i^{(k-1)}(\mathbf{x})|^2.$$

The Electronic Structure Problem

The nonlinear problem can be solved by a 'self-consistent field' (SCF) iteration, leading to a sequence of linear eigenproblems

$$F^{(k)}\psi_i^{(k)} = \lambda_i^{(k)}\psi_i^{(k)}, \quad \langle \psi_i^{(k)}, \psi_j^{(k)} \rangle = \delta_{ij}, \quad k = 1, 2, \ldots$$

$(1 \leq i, j \leq N)$, where each $F^{(k)} = -\frac{1}{2}\Delta + U^{(k)}$ is a one-electron linearized Hamiltonian:

$$U^{(k)} = U^{(k)}(\mathbf{x}, \rho^{(k-1)}), \quad \rho^{(k-1)} = \sum_{i=1}^{N} |\psi_i^{(k-1)}(\mathbf{x})|^2.$$

Solution of each of the (discretized) linear eigenproblems above leads to a typical $O(N^3)$ cost per SCF iteration. However, the actual eigenpairs $(\psi_i^{(k)}, \lambda_i^{(k)})$ are unnecessary, and diagonalization of the one-particle Hamiltonians can be avoided!

The Electronic Structure Problem

The individual eigenfunctions ψ_i are not needed. All one needs is the orthogonal projector P onto the occupied subspace

$$V_{occ} = \text{span}\{\psi_1, \ldots, \psi_N\}$$

corresponding to the N lowest eigenvalues $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_N$.

The Electronic Structure Problem

The individual eigenfunctions ψ_i are not needed. All one needs is the orthogonal projector P onto the occupied subspace

$$V_{occ} = \text{span}\{\psi_1, \ldots, \psi_N\}$$

corresponding to the N lowest eigenvalues $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_N$.

At the kth SCF cycle, an approximation to the orthogonal projector $P^{(k)}$ onto the occupied subspace $V_{occ}^{(k)}$ needs to be computed.

The Electronic Structure Problem

The individual eigenfunctions ψ_i are not needed. All one needs is the orthogonal projector P onto the occupied subspace

$$V_{occ} = \text{span}\{\psi_1, \ldots, \psi_N\}$$

corresponding to the N lowest eigenvalues $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_N$.

At the kth SCF cycle, an approximation to the orthogonal projector $P^{(k)}$ onto the occupied subspace $V_{occ}^{(k)}$ needs to be computed.

All quantities of interest in electronic structure theory can be computed from P. For example, the expected value for the total energy of the system can be computed once P is known, as are the forces.

The Electronic Structure Problem

The individual eigenfunctions ψ_i are not needed. All one needs is the orthogonal projector P onto the occupied subspace

$$V_{occ} = \text{span}\{\psi_1, \ldots, \psi_N\}$$

corresponding to the N lowest eigenvalues $\lambda_1 \leq \lambda_2 \leq \cdots \leq \lambda_N$.

At the kth SCF cycle, an approximation to the orthogonal projector $P^{(k)}$ onto the occupied subspace $V_{occ}^{(k)}$ needs to be computed.

All quantities of interest in electronic structure theory can be computed from P. For example, the expected value for the total energy of the system can be computed once P is known, as are the forces.

Higher moments of observables can also be computed once P is known.

Overview

Density matrices

P is called the density operator of the system (von Neumann, 1927). For systems in equilibrium, $[H, P] = 0$ and P is a function of the Hamiltonian: $P = f(H)$.

Density matrices

P is called the density operator of the system (von Neumann, 1927). For systems in equilibrium, $[H, P] = 0$ and P is a function of the Hamiltonian: $P = f(H)$.

Density operators are sometimes normalized so that $\mathrm{Tr}(P) = 1$.

Density matrices

P is called the density operator of the system (von Neumann, 1927). For systems in equilibrium, $[H, P] = 0$ and P is a function of the Hamiltonian: $P = f(H)$.

Density operators are sometimes normalized so that $\mathrm{Tr}(P) = 1$.

In practice, the operators are replaced by matrices by Rayleigh-Ritz projection onto a finite-dimensional subspace spanned by a set of basis functions $\{\phi_i\}_{i=1}^{N_b}$, where N_b is a multiple of N. Typically, $N = C \cdot N_e$ where $C \geq 2$ is a moderate integer when linear combinations of GTOs (Gaussian-type orbitals) are used.

Density matrices

P is called the density operator of the system (von Neumann, 1927). For systems in equilibrium, $[H, P] = 0$ and P is a function of the Hamiltonian: $P = f(H)$.

Density operators are sometimes normalized so that $\mathrm{Tr}(P) = 1$.

In practice, the operators are replaced by matrices by Rayleigh-Ritz projection onto a finite-dimensional subspace spanned by a set of basis functions $\{\phi_i\}_{i=1}^{N_b}$, where N_b is a multiple of N. Typically, $N = C \cdot N_e$ where $C \geq 2$ is a moderate integer when linear combinations of GTOs (Gaussian-type orbitals) are used.

In some codes, plane waves are used. Finite difference and finite element ("real space") methods, while also used, are less popular.

Density matrices

The Gramian matrix $S = (S_{ij})$ where $S_{ij} = \langle \phi_i, \phi_j \rangle$ is called the overlap matrix in electronic structure. It is dense, but its entries fall off very rapidly for increasing separation. In the case of GTOs:

$$|S_{ij}| \approx \mathrm{e}^{-|i-j|^2}, \quad 1 \leq i, j \leq N_b \,.$$

Neglecting tiny entries leads to banded overlap matrices.

Density matrices

The Gramian matrix $S = (S_{ij})$ where $S_{ij} = \langle \phi_i, \phi_j \rangle$ is called the overlap matrix in electronic structure. It is dense, but its entries fall off very rapidly for increasing separation. In the case of GTOs:

$$|S_{ij}| \approx \mathrm{e}^{-|i-j|^2}, \quad 1 \leq i, j \leq N_b \,.$$

Neglecting tiny entries leads to banded overlap matrices.

The density matrix P is the S-orthogonal projector onto the occupied subspace. It is convenient to make a change of basis, from non-orthogonal to orthogonal.

Density matrices

The Gramian matrix $S = (S_{ij})$ where $S_{ij} = \langle \phi_i, \phi_j \rangle$ is called the overlap matrix in electronic structure. It is dense, but its entries fall off very rapidly for increasing separation. In the case of GTOs:

$$|S_{ij}| \approx \mathrm{e}^{-|i-j|^2}, \quad 1 \leq i, j \leq N_b \,.$$

Neglecting tiny entries leads to banded overlap matrices.

The density matrix P is the S-orthogonal projector onto the occupied subspace. It is convenient to make a change of basis, from non-orthogonal to orthogonal.

Trasforming the Hamiltonian H to an orthogonal basis means performing a congruence trasformation: $\hat{H} = ZHZ^T$, where Z is such that $ZZ^T = S^{-1}$. Common choices for Z are the inverse Cholesky factor of S or the inverse square root $S^{-1/2}$. This can be done efficiently (AINV algorithm).

Example: Hamiltonian for $C_{52}H_{106}$, AO basis.

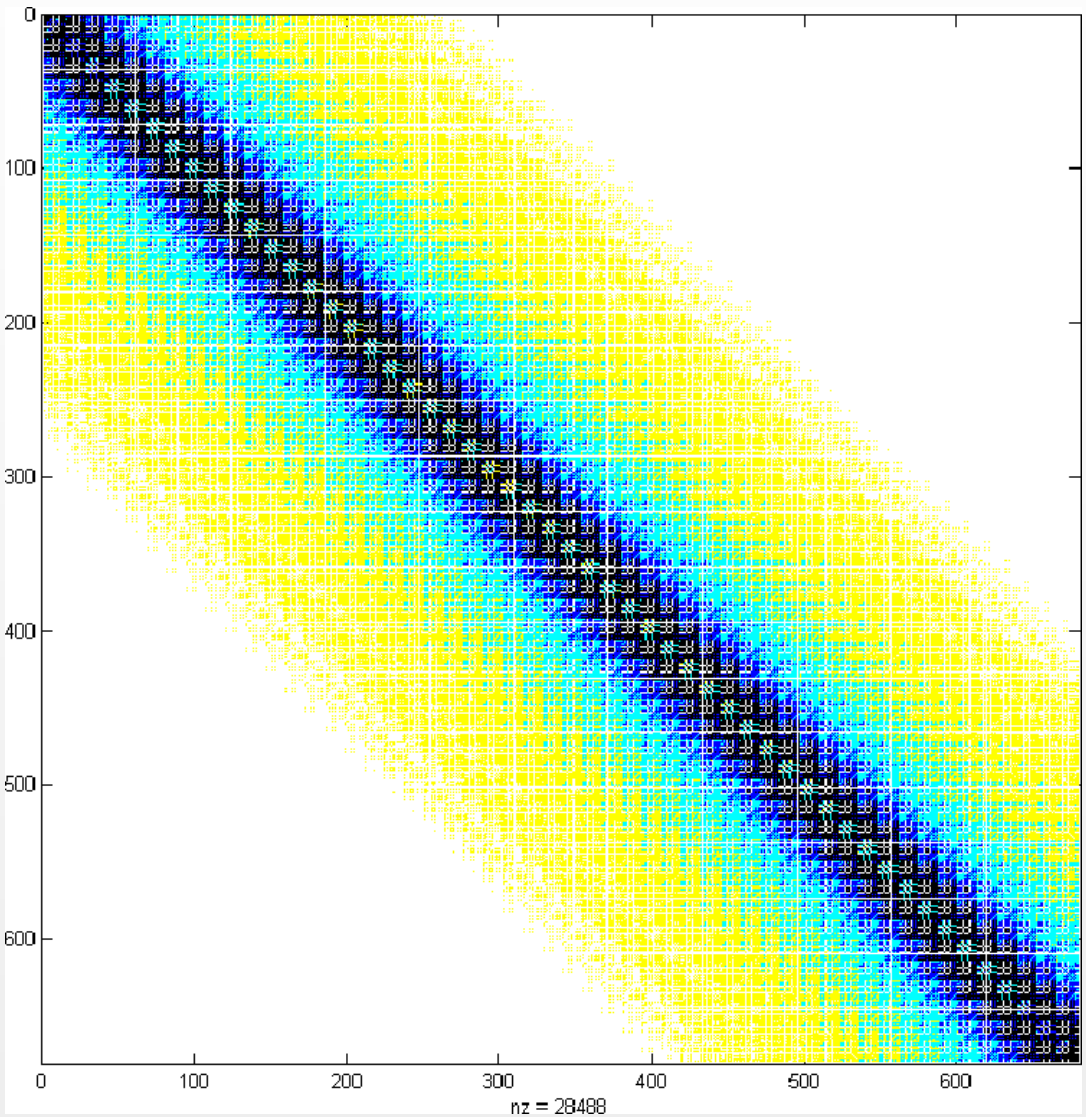

Example: Hamiltonian for $C_{52}H_{106}$, orthogonal basis

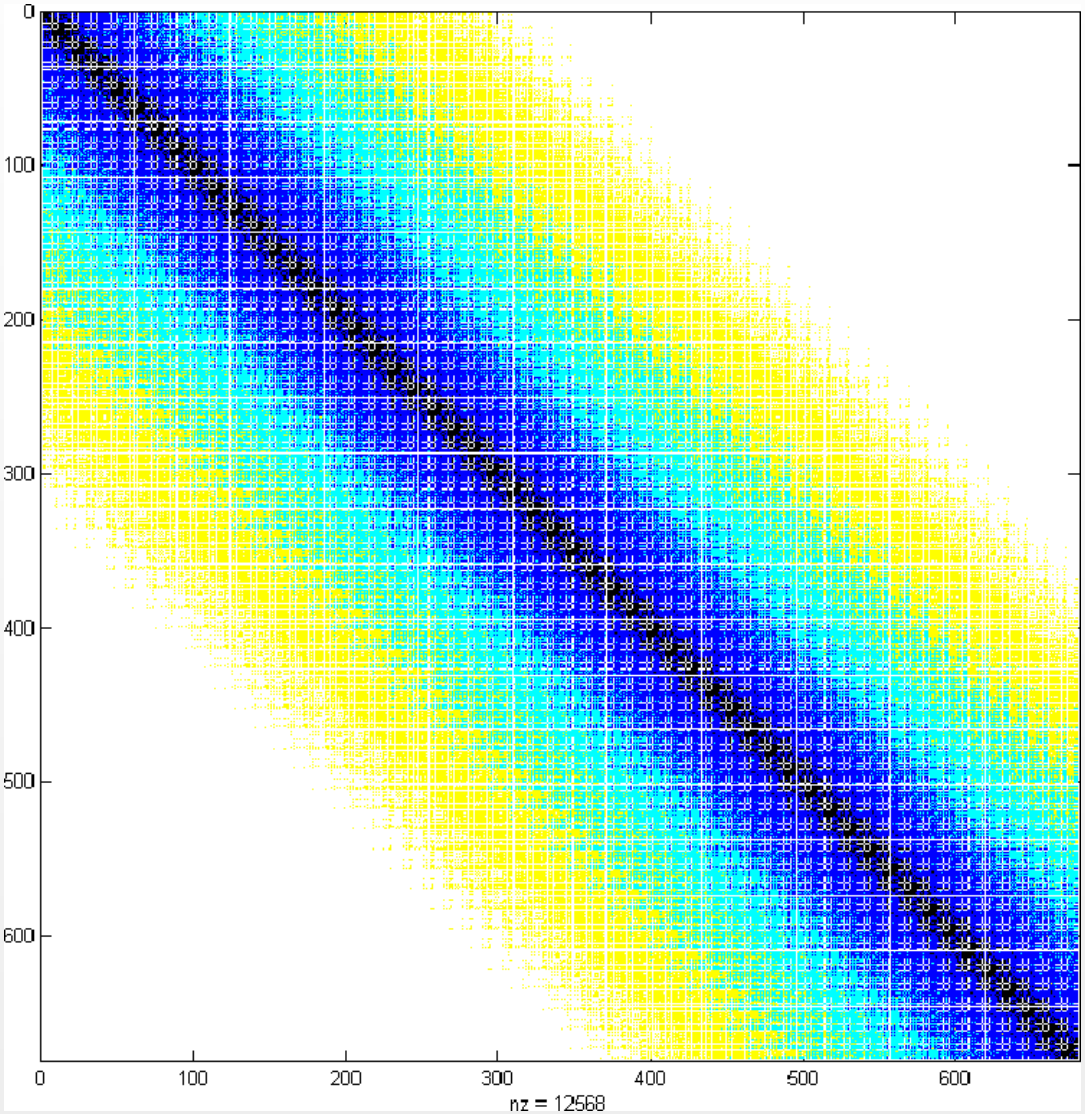

Density matrices

Summarizing, in electronic structure theory we need to compute P, the spectral projector onto the subspace spanned by the N lowest eigenfunctions of H (occupied states):

$$P = \psi_1 \otimes \psi_1 + \cdots + \psi_N \otimes \psi_N = |\psi_1\rangle\langle\psi_1| + \cdots + |\psi_N\rangle\langle\psi_N|$$

where N is the number of electrons and $H\psi_i = \lambda_i \psi_i, \ i = 1, \ldots, N.$

Density matrices

Summarizing, in electronic structure theory we need to compute P, the spectral projector onto the subspace spanned by the N lowest eigenfunctions of H (occupied states):

$$P = \psi_1 \otimes \psi_1 + \cdots + \psi_N \otimes \psi_N = |\psi_1\rangle\langle\psi_1| + \cdots + |\psi_N\rangle\langle\psi_N|$$

where N is the number of electrons and $H\psi_i = \lambda_i \psi_i, \ i = 1, \ldots, N$.

Note that we can write $P = f(H)$ where f is the step function

$$f(x) = \begin{cases} 1 & \text{if} \quad x < \mu \\ \frac{1}{2} & \text{if} \quad x = \mu \\ 0 & \text{if} \quad x > \mu \end{cases}$$

with $\lambda_N < \mu < \lambda_{N+1}$ (μ is the "Fermi level").

Overview

$O(N)$ methods

Physicists have observed that the entries of the density matrix P decay away from the main diagonal. The decay rate is algebraic for metallic systems, and exponential (or faster) for insulators and semiconductors. Hence, the density matrix is localized.

$O(N)$ methods

Physicists have observed that the entries of the density matrix P decay away from the main diagonal. The decay rate is algebraic for metallic systems, and exponential (or faster) for insulators and semiconductors. Hence, the density matrix is localized.

This property has been called nearsightedness by W. Kohn.

$O(N)$ methods

Physicists have observed that the entries of the density matrix P decay away from the main diagonal. The decay rate is algebraic for metallic systems, and exponential (or faster) for insulators and semiconductors. Hence, the density matrix is localized.

This property has been called nearsightedness by W. Kohn.

This property implies that in the thermodynamic limit ($N \to \infty$ while keeping the particle density constant) the number of entries P_{ij} with $|P_{ij}| > \varepsilon$ grows only linearly with N, for any prescribed $\varepsilon > 0$.

$O(N)$ methods

Physicists have observed that the entries of the density matrix P decay away from the main diagonal. The decay rate is algebraic for metallic systems, and exponential (or faster) for insulators and semiconductors. Hence, the density matrix is localized.

This property has been called nearsightedness by W. Kohn.

This property implies that in the thermodynamic limit ($N \to \infty$ while keeping the particle density constant) the number of entries P_{ij} with $|P_{ij}| > \varepsilon$ grows only linearly with N, for any prescribed $\varepsilon > 0$.

For insulators and semiconductors, this makes $O(N)$ methods possible.

Example: Density matrix for $C_{52}H_{106}$, orthogonal basis

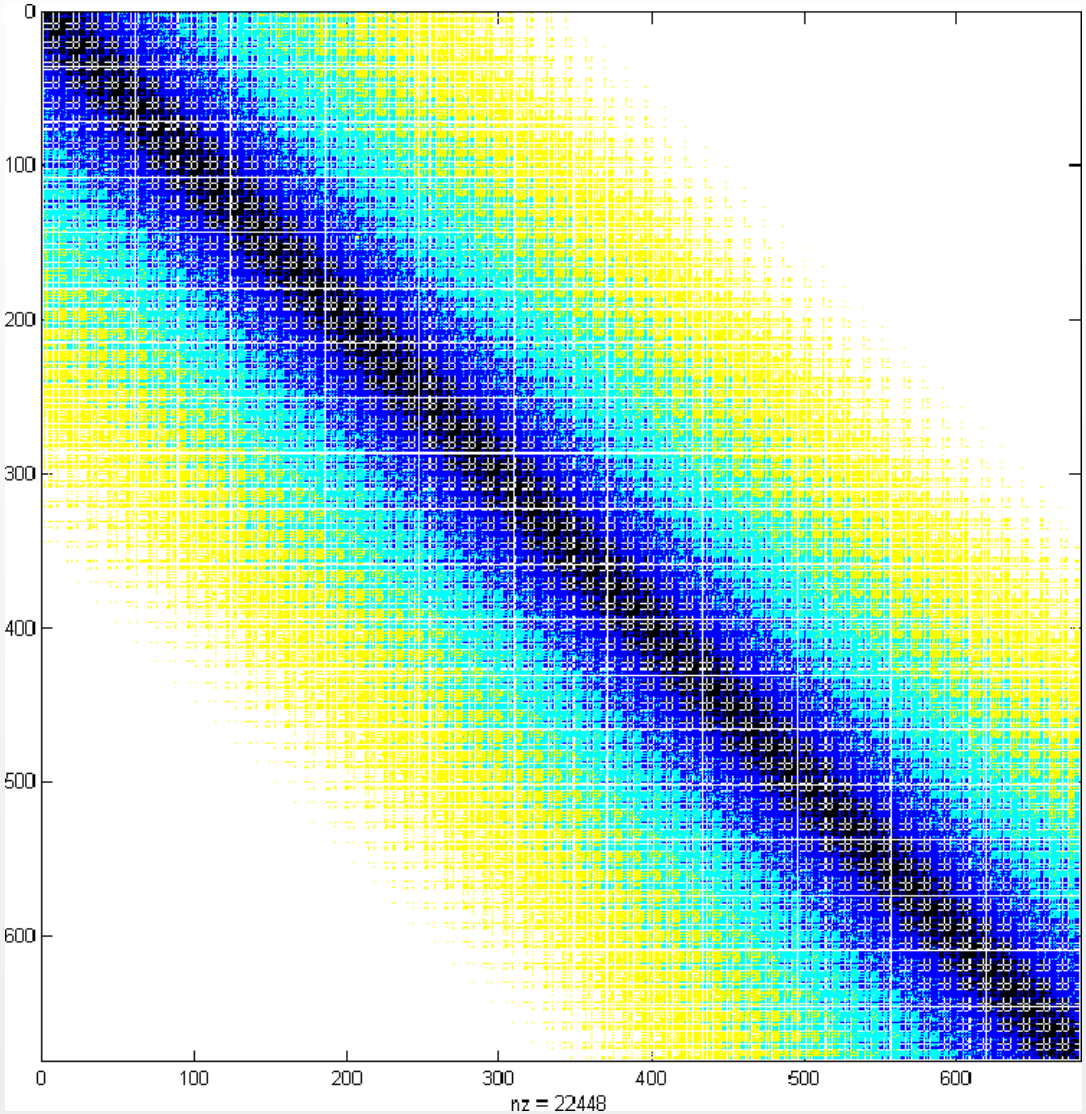

Example: Density matrix for $H = -\frac{1}{2}\Delta + V$, random V, finite differences (2D lattice), $N = 10$

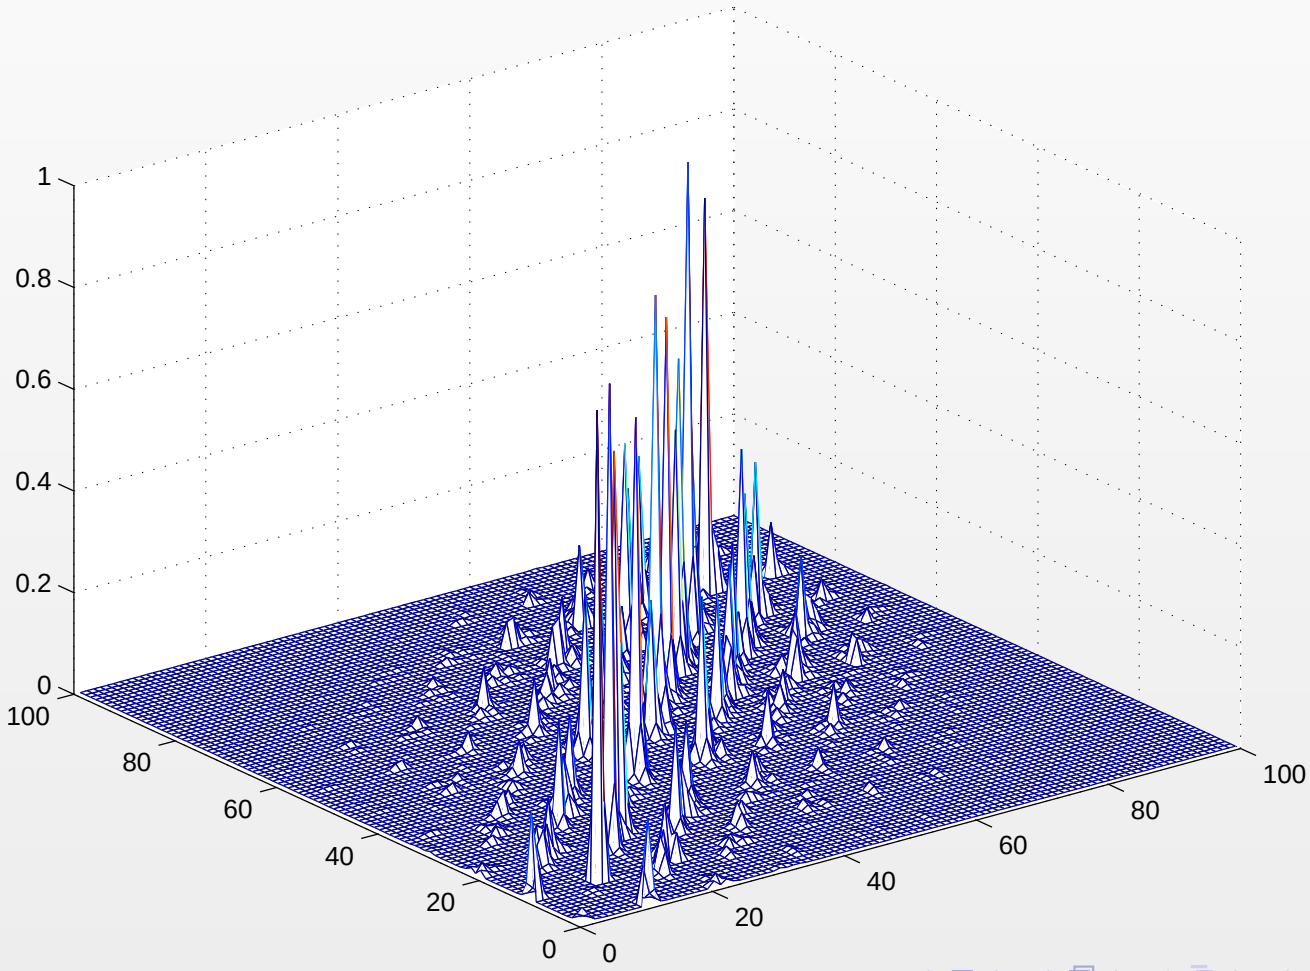

$O(N)$ methods

The fact that P is very nearly sparse allows the development of $O(N)$ approximation methods. Popular methods include

$O(N)$ methods

The fact that P is very nearly sparse allows the development of $O(N)$ approximation methods. Popular methods include

1. Chebyshev expansion: $P = f(H) \approx \frac{c_0}{2} + \sum_{k=1}^{n} c_k T_k(H)$
2. Rational expansions based on contour integration:

$$P = \frac{1}{2\pi i} \int_{\Gamma} (zI - H)^{-1} \mathrm{d}z \approx \sum_{k=1}^{q} w_k (z_k I - H)^{-1}$$

3. Density matrix minimization:

$$\mathrm{Tr}(PH) = \min, \quad \text{subject to} \quad P = P^* = P^2 \quad \text{and} \quad \mathrm{rank}(P) = N$$

4. Etc. (see C. Le Bris, *Acta Numerica*, 2005; Bowler & Miyazaki, 2011)

$O(N)$ methods

The fact that P is very nearly sparse allows the development of $O(N)$ approximation methods. Popular methods include

1. Chebyshev expansion: $P = f(H) \approx \frac{c_0}{2} + \sum_{k=1}^{n} c_k T_k(H)$
2. Rational expansions based on contour integration:

$$P = \frac{1}{2\pi i} \int_{\Gamma} (zI - H)^{-1} \mathrm{d}z \approx \sum_{k=1}^{q} w_k (z_k I - H)^{-1}$$

3. Density matrix minimization:

$$\mathrm{Tr}(PH) = \min, \quad \text{subject to} \quad P = P^* = P^2 \quad \text{and} \quad \mathrm{rank}(P) = N$$

4. Etc. (see C. Le Bris, *Acta Numerica*, 2005; Bowler & Miyazaki, 2011)

All these methods can achieve $O(N)$ scaling by exploiting "sparsity."

Overview

A mathematical foundation for $O(N)$ methods

In literature one finds no rigorously proved, general results on the decay behavior in P. Decay estimates are either non-rigorous (at least for a mathematician!), or valid only for very special cases.

A mathematical foundation for $O(N)$ methods

In literature one finds no rigorously proved, general results on the decay behavior in P. Decay estimates are either non-rigorous (at least for a mathematician!), or valid only for very special cases.

Goal: To establish rigorous estimates for the rate of decay in the off-diagonal entries of P in the form of upper bounds on the entries of the density matrix, in the thermodynamic limit $N \to \infty$.

A mathematical foundation for $O(N)$ methods

In literature one finds no rigorously proved, general results on the decay behavior in P. Decay estimates are either non-rigorous (at least for a mathematician!), or valid only for very special cases.

Goal: To establish rigorous estimates for the rate of decay in the off-diagonal entries of P in the form of upper bounds on the entries of the density matrix, in the thermodynamic limit $N \to \infty$.

Ideally, these bounds can be used to provide cut-off rules, that is, to decide which entries in P need not be computed, given a prescribed error tolerance.

A mathematical foundation for $O(N)$ methods

In literature one finds no rigorously proved, general results on the decay behavior in P. Decay estimates are either non-rigorous (at least for a mathematician!), or valid only for very special cases.

Goal: To establish rigorous estimates for the rate of decay in the off-diagonal entries of P in the form of upper bounds on the entries of the density matrix, in the thermodynamic limit $N \to \infty$.

Ideally, these bounds can be used to provide cut-off rules, that is, to decide which entries in P need not be computed, given a prescribed error tolerance.

To achieve this goal we first need to find an appropriate formalization for the type of problems encountered by physicists.

A mathematical foundation for $O(N)$ methods

In literature one finds no rigorously proved, general results on the decay behavior in P. Decay estimates are either non-rigorous (at least for a mathematician!), or valid only for very special cases.

Goal: To establish rigorous estimates for the rate of decay in the off-diagonal entries of P in the form of upper bounds on the entries of the density matrix, in the thermodynamic limit $N \to \infty$.

Ideally, these bounds can be used to provide cut-off rules, that is, to decide which entries in P need not be computed, given a prescribed error tolerance.

To achieve this goal we first need to find an appropriate formalization for the type of problems encountered by physicists.

The main tool used in our analysis, besides linear algebra, is classical approximation theory.

A mathematical foundation for $O(N)$ methods

Let C be a fixed positive integer and let $N_b := C \cdot N$, where $N \longrightarrow \infty$. We think of C as the number of basis functions per electron while N is the number of electrons.

A mathematical foundation for $O(N)$ methods

Let C be a fixed positive integer and let $N_b := C \cdot N$, where $N \to \infty$. We think of C as the number of basis functions per electron while N is the number of electrons.

We call a sequence of discrete Hamiltonians a sequence of matrices $\{H_N\}$ of order N_b, such that

1. The matrices H_N have spectra uniformly bounded w.r.t. N: up to shifting and scaling, we can assume $\sigma(H_N) \subset [-1, 1]$ for all N.

A mathematical foundation for $O(N)$ methods

Let C be a fixed positive integer and let $N_b := C \cdot N$, where $N \to \infty$. We think of C as the number of basis functions per electron while N is the number of electrons.

We call a sequence of discrete Hamiltonians a sequence of matrices $\{H_N\}$ of order N_b, such that

1. The matrices H_N have spectra uniformly bounded w.r.t. N: up to shifting and scaling, we can assume $\sigma(H_N) \subset [-1, 1]$ for all N.
2. The H_N are banded, with uniformly bounded bandwidth as $N \to \infty$. More generally, the H_N are sparse with maximum number of nonzeros per row bounded w.r.t. N.

A mathematical foundation for $O(N)$ methods

Let C be a fixed positive integer and let $N_b := C \cdot N$, where $N \to \infty$. We think of C as the number of basis functions per electron while N is the number of electrons.

We call a sequence of discrete Hamiltonians a sequence of matrices $\{H_N\}$ of order N_b, such that

1. The matrices H_N have spectra uniformly bounded w.r.t. N: up to shifting and scaling, we can assume $\sigma(H_N) \subset [-1, 1]$ for all N.
2. The H_N are banded, with uniformly bounded bandwidth as $N \to \infty$. More generally, the H_N are sparse with maximum number of nonzeros per row bounded w.r.t. N.

These assumptions model the fact that the Hamiltonians have finite interaction range, which remains bounded in the thermodynamic limit $N \to \infty$.

A mathematical foundation for $O(N)$ methods

Next, let H be a Hamiltonian of order N_b. Denote the eigenvalues of H as

$$-1 \leq \lambda_1 \leq \ldots \leq \lambda_N < \lambda_{N+1} \leq \ldots \leq \lambda_{N_b} \leq 1\,.$$

The spectral gap is then $\gamma_N = \lambda_{N+1} - \lambda_N$. In quantum chemistry this is known as the HOMO-LUMO gap, in solid state physics as the band gap.

A mathematical foundation for $O(N)$ methods

Next, let H be a Hamiltonian of order N_b. Denote the eigenvalues of H as

$$-1 \leq \lambda_1 \leq \ldots \leq \lambda_N < \lambda_{N+1} \leq \ldots \leq \lambda_{N_b} \leq 1\,.$$

The spectral gap is then $\gamma_N = \lambda_{N+1} - \lambda_N$. In quantum chemistry this is known as the HOMO-LUMO gap, in solid state physics as the band gap.

Two cases are possible:

1. There exists $\gamma > 0$ such that $\gamma_N \geq \gamma$ for all N.

A mathematical foundation for $O(N)$ methods

Next, let H be a Hamiltonian of order N_b. Denote the eigenvalues of H as

$$-1 \leq \lambda_1 \leq \ldots \leq \lambda_N < \lambda_{N+1} \leq \ldots \leq \lambda_{N_b} \leq 1\,.$$

The spectral gap is then $\gamma_N = \lambda_{N+1} - \lambda_N$. In quantum chemistry this is known as the HOMO-LUMO gap, in solid state physics as the band gap.

Two cases are possible:

1. There exists $\gamma > 0$ such that $\gamma_N \geq \gamma$ for all N.
2. $\inf_N \gamma_N = 0$.

A mathematical foundation for $O(N)$ methods

Next, let H be a Hamiltonian of order N_b. Denote the eigenvalues of H as

$$-1 \leq \lambda_1 \leq \ldots \leq \lambda_N < \lambda_{N+1} \leq \ldots \leq \lambda_{N_b} \leq 1\,.$$

The spectral gap is then $\gamma_N = \lambda_{N+1} - \lambda_N$. In quantum chemistry this is known as the HOMO-LUMO gap, in solid state physics as the band gap.

Two cases are possible:

1. There exists $\gamma > 0$ such that $\gamma_N \geq \gamma$ for all N.
2. $\inf_N \gamma_N = 0$.

The first case corresponds to insulators and semiconductors, the second one to metallic systems.

Example: Spectrum of the Hamiltonian for $C_{52}H_{106}$.

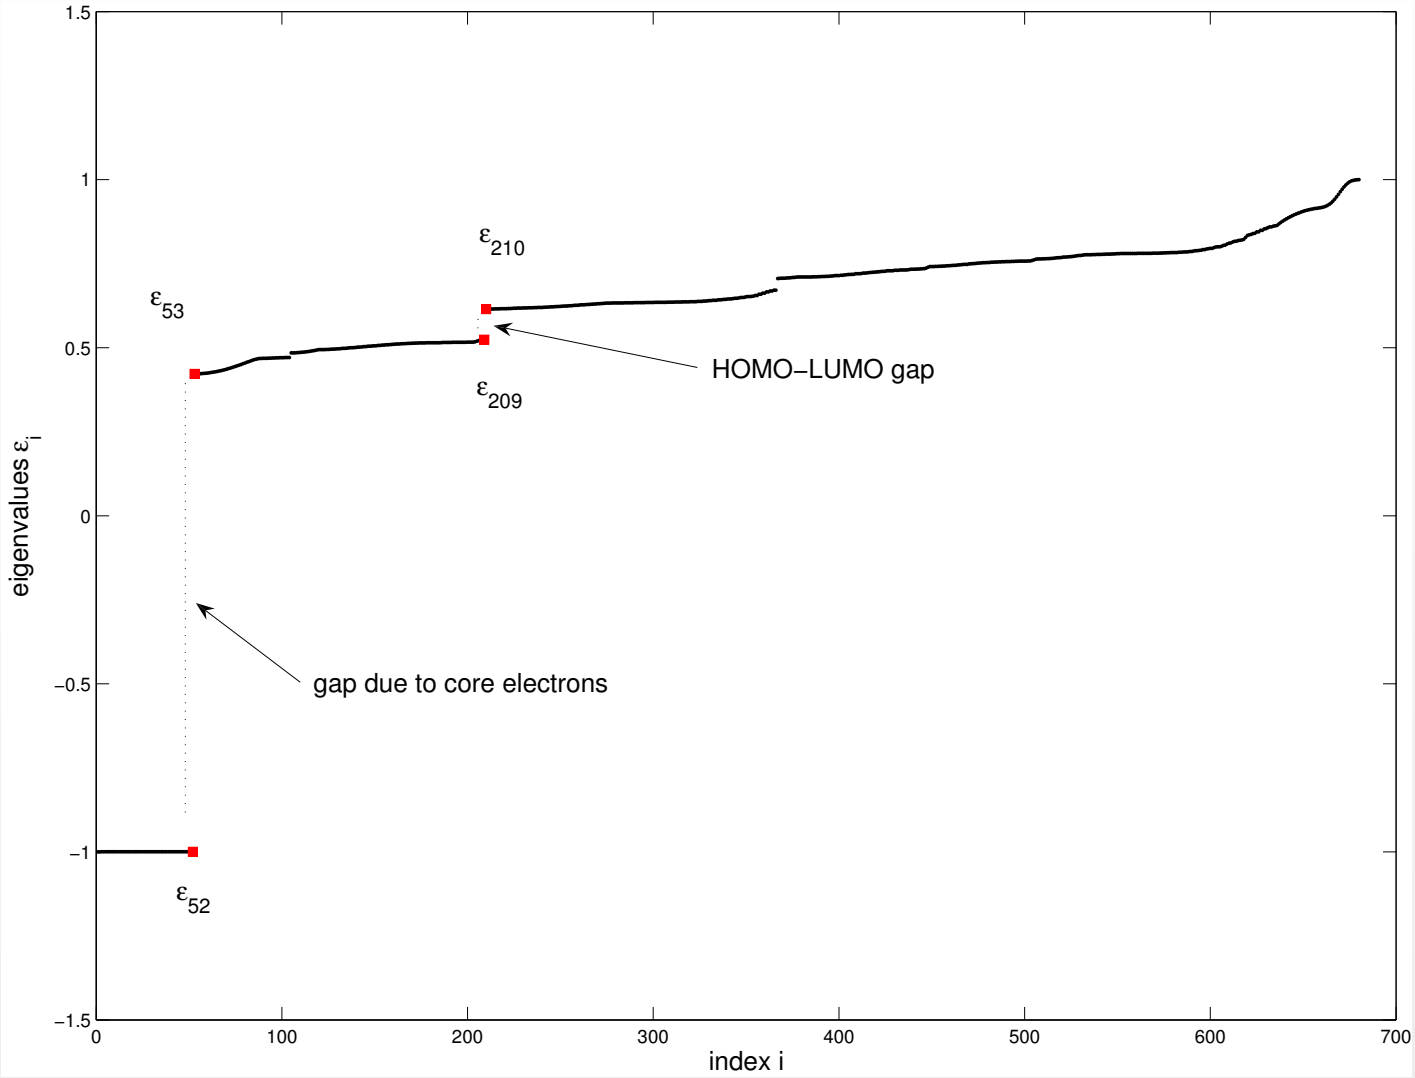

Exponential decay in the density matrix for gapped systems

Theorem

Let $\{H_N\}$ be a sequence of discrete Hamiltonians of size $N_b = C \cdot N$, with C constant and $N \to \infty$. Let P_N denote the spectral projector onto the N occupied states associated with H_N. If there exists $\gamma > 0$ such that the gaps $\gamma_N \geq \gamma$ for all N, then there exists constants K and α such that

$$|[P_N]_{ij}| \leq K \, \mathrm{e}^{-\alpha d_N(i,j)} \quad (1 \leq i, j \leq N),$$

where $d_N(i,j)$ denotes the geodetic distance between node i and node j in the graph G_N associated with H_N. The constants K and α depend only on the gap γ (not on N) and are easily computable.

Note: The graph $G_N = (V_N, E_N)$ is the graph with N_b vertices such that there is an edge $(i,j) \in E_N$ if and only if $[H_N]_{ij} \neq 0$.

Exponential decay in the density matrix for gapped systems

Sketch of the proof. There are a few steps involved:

1. Recall that $P_N = f(H_N)$ where f is a step function. For gapped systems, f can be approximated arbitrarily well in the sup norm by an analytic function, for example, the Fermi-Dirac function f_{FD}.

Exponential decay in the density matrix for gapped systems

Sketch of the proof. There are a few steps involved:

1. Recall that $P_N = f(H_N)$ where f is a step function. For gapped systems, f can be approximated arbitrarily well in the sup norm by an analytic function, for example, the Fermi-Dirac function f_{FD}.
2. Use best approximation polynomials to approximate f_{FD} (in the sup norm).

Exponential decay in the density matrix for gapped systems

Sketch of the proof. There are a few steps involved:

1. Recall that $P_N = f(H_N)$ where f is a step function. For gapped systems, f can be approximated arbitrarily well in the sup norm by an analytic function, for example, the Fermi-Dirac function f_{FD}.
2. Use best approximation polynomials to approximate f_{FD} (in the sup norm).
3. By Bernstein's Theorem, for any analytic function g defined on a neighborhood of $[-1, 1]$, the best nth degree polynomial error satisfies

$$E_n(g) := \|p_n - g\|_\infty \leq c \cdot \mathrm{e}^{-\xi n}$$

for suitable constants c and ξ which depend only on g.

Exponential decay in the density matrix for gapped systems

Sketch of the proof. There are a few steps involved:

1. Recall that $P_N = f(H_N)$ where f is a step function. For gapped systems, f can be approximated arbitrarily well in the sup norm by an analytic function, for example, the Fermi-Dirac function f_{FD}.
2. Use best approximation polynomials to approximate f_{FD} (in the sup norm).
3. By Bernstein's Theorem, for any analytic function g defined on a neighborhood of $[-1, 1]$, the best nth degree polynomial error satisfies

$$E_n(g) := \|p_n - g\|_\infty \leq c \cdot \mathrm{e}^{-\xi n}$$

   for suitable constants c and ξ which depend only on g.
4. Apply Bernstein's result to $g = f_{FD}$, compute the decay constants, and use the spectral theorem to go from scalars to matrices.

Analytic approximations of the step function

If μ (the "Fermi level") is in the gap, $\lambda_N < \mu < \lambda_{N+1}$, the step function can be approximated by the Fermi-Dirac function: $f(x) = \lim_{\beta\to\infty} f_{FD}(x)$, where

$$f_{FD}(x) = \frac{1}{1 + e^{\beta(x-\mu)}}.$$

Here β can be interpreted as an inverse temperature.

Analytic approximations of the step function

If μ (the "Fermi level") is in the gap, $\lambda_N < \mu < \lambda_{N+1}$, the step function can be approximated by the Fermi-Dirac function: $f(x) = \lim_{\beta\to\infty} f_{FD}(x)$, where

$$f_{FD}(x) = \frac{1}{1 + e^{\beta(x-\mu)}}.$$

Here β can be interpreted as an inverse temperature.

Other approximations of the step function are also in use, such as

$$f(x) = \lim_{\beta\to\infty} \left[\frac{1}{2} + \frac{1}{\pi}\tan^{-1}(\beta\pi(x-\mu))\right],$$

$$f(x) = \lim_{\beta\to\infty} \operatorname{erfc}\left(-\beta(x-\mu)\right),$$

or

$$f(x) = \lim_{\beta\to\infty} \left[1 + \tanh(\beta((x-\mu))\right].$$

Fermi-Dirac approximation of step function ($\mu = 0$)

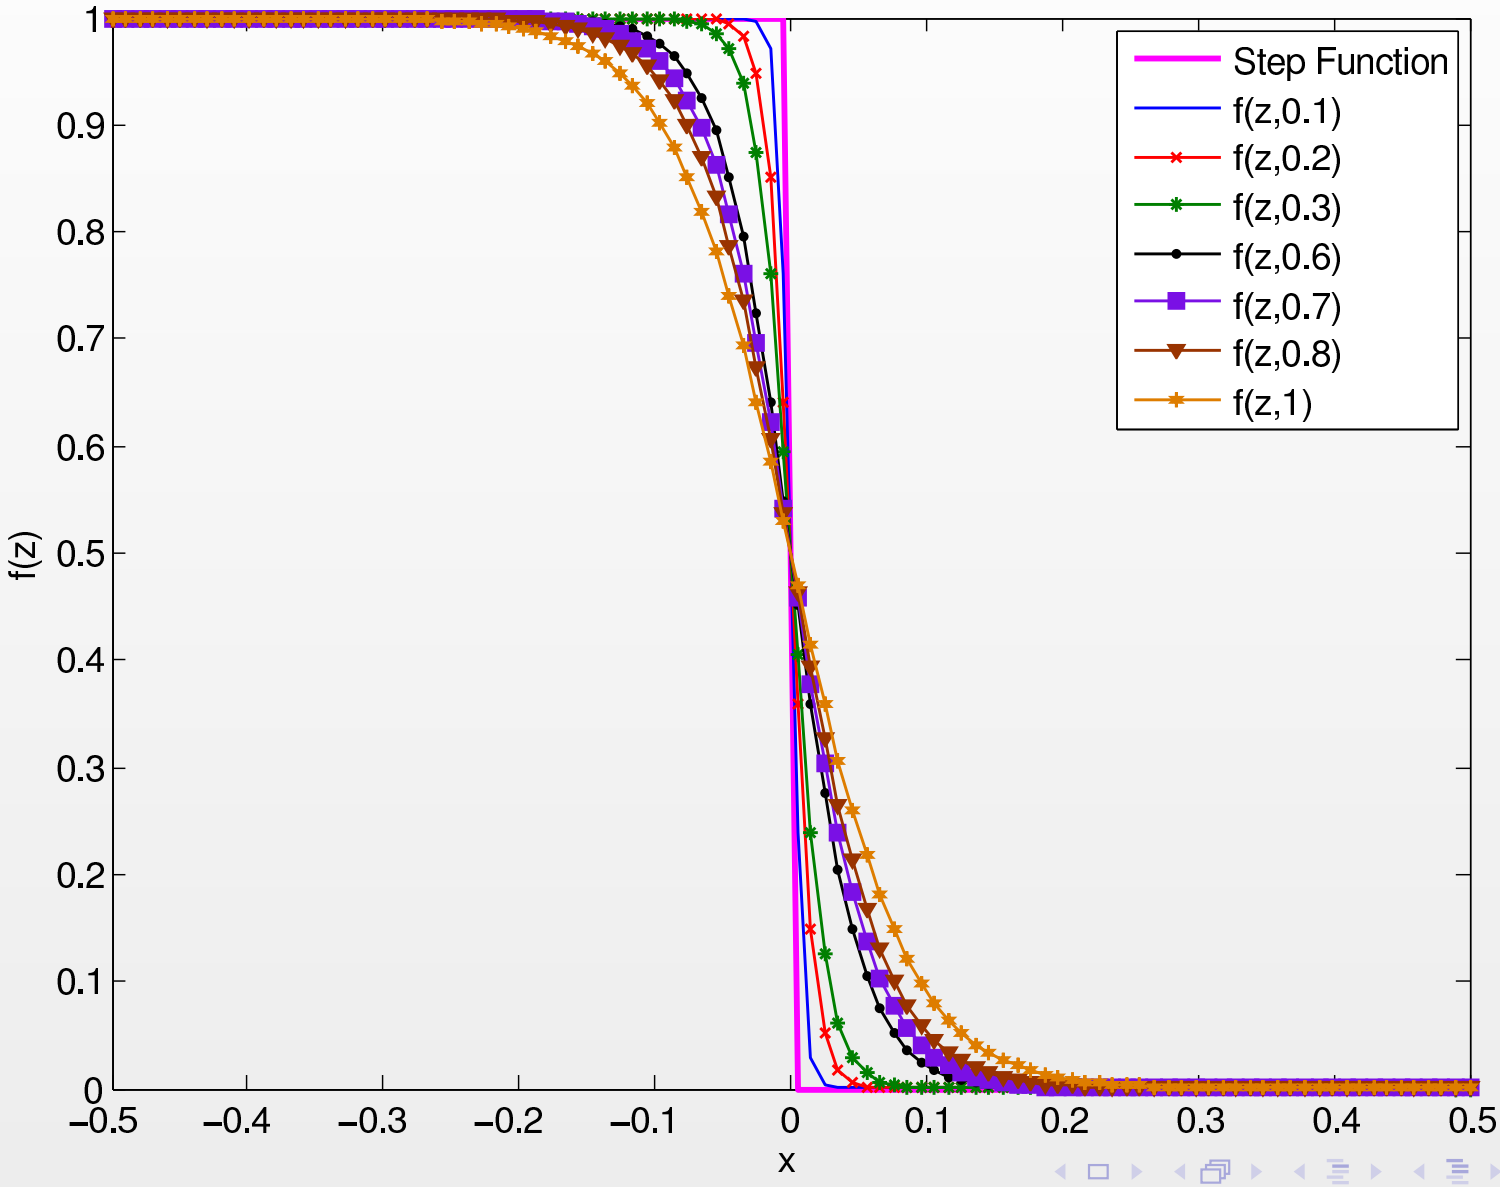

Fermi-Dirac approximation of step function

The actual choice of β in the Fermi-Dirac function is dictated by the size of the gap γ and by the approximation error.

Fermi-Dirac approximation of step function

The actual choice of β in the Fermi-Dirac function is dictated by the size of the gap γ and by the approximation error.

What is important is that for any prescribed error, there is a maximum value of β that achieves the error for all N, provided the system has non-vanishing gap.

Fermi-Dirac approximation of step function

The actual choice of β in the Fermi-Dirac function is dictated by the size of the gap γ and by the approximation error.

What is important is that for any prescribed error, there is a maximum value of β that achieves the error for all N, provided the system has non-vanishing gap.

On the other hand, for $\gamma \to 0$ we have $\beta \to \infty$ and the bounds blow up. This happens for metals (at zero temperature).

Fermi-Dirac approximation of step function

The actual choice of β in the Fermi-Dirac function is dictated by the size of the gap γ and by the approximation error.

What is important is that for any prescribed error, there is a maximum value of β that achieves the error for all N, provided the system has non-vanishing gap.

On the other hand, for $\gamma \to 0$ we have $\beta \to \infty$ and the bounds blow up. This happens for metals (at zero temperature).

There is still decay in the density matrix, but it may be *algebraic* rather than exponential. Simple examples show it can be as slow as $O(|i-j|^{-1})$.

Fermi-Dirac approximation of step function

The actual choice of β in the Fermi-Dirac function is dictated by the size of the gap γ and by the approximation error.

What is important is that for any prescribed error, there is a maximum value of β that achieves the error for all N, provided the system has non-vanishing gap.

On the other hand, for $\gamma \to 0$ we have $\beta \to \infty$ and the bounds blow up. This happens for metals (at zero temperature).

There is still decay in the density matrix, but it may be *algebraic* rather than exponential. Simple examples show it can be as slow as $O(|i-j|^{-1})$.

In practice, γ is either known experimentally or can be estimated by computing the eigenvalues of a moderate-size Hamiltonian.

Dependence of decay rate on the spectral gap and on the temperature

In the physics literature, there has been some controversy on the precise dependence of the inverse correlation length α in the decay estimate

$$|[P_N]_{ij}| \leq c \cdot \mathrm{e}^{-\alpha d_N(i,j)}$$

on the spectral gap γ (for insulators) and on the electronic temperature T (for metals at positive temperature).

Our theory gives the following results:

1. $\alpha = c\gamma + O(\gamma^3)$, for $\gamma \to 0+$ and $T = 0$;
2. $\alpha = \pi\kappa_B T + O(T^3)$, for $T \to 0+$ (indep. of γ).

These asymptotics are in agreement with experimental and numerical results, as well as with physical intuition.

Decay bounds for the Fermi-Dirac approximation

Assume that H is m-banded and has spectrum in $[-1, 1]$, then

$$\left| \left[\left(I + \mathrm{e}^{\beta(H-\mu I)} \right)^{-1} \right]_{ij} \right| \leq K \mathrm{e}^{-\alpha|i-j|} \equiv K \lambda^{\frac{|i-j|}{m}}.$$

Decay bounds for the Fermi-Dirac approximation

Assume that H is m-banded and has spectrum in $[-1, 1]$, then

$$\left| \left[\left(I + \mathrm{e}^{\beta(H-\mu I)} \right)^{-1} \right]_{ij} \right| \leq K \mathrm{e}^{-\alpha|i-j|} \equiv K \lambda^{\frac{|i-j|}{m}}.$$

Note that K, λ depend only on β. In turn, β depends on γ and on the desired accuracy.
We have

$$\gamma \to 0^+ \quad \Rightarrow \quad \lambda \to 1^-$$

and

$$\gamma \to 1 \quad \Rightarrow \quad \lambda \to 0.872.$$

We choose β and $\hat{m}$ so as to guarantee an accuracy $\|P - f(H)\|_2 < 10^{-6}$.

Decay bounds for the Fermi-Dirac approximation

Assume that H is m-banded and has spectrum in $[-1, 1]$, then

$$\left| \left[\left(I + \mathrm{e}^{\beta(H-\mu I)} \right)^{-1} \right]_{ij} \right| \leq K \mathrm{e}^{-\alpha|i-j|} \equiv K \lambda^{\frac{|i-j|}{m}}.$$

Note that K, λ depend only on β. In turn, β depends on γ and on the desired accuracy.
We have

$$\gamma \to 0^+ \quad \Rightarrow \quad \lambda \to 1^-$$

and

$$\gamma \to 1 \quad \Rightarrow \quad \lambda \to 0.872.$$

We choose β and $\hat{m}$ so as to guarantee an accuracy $\|P - f(H)\|_2 < 10^{-6}$.

We can regard γ^{-1} as the condition number of the problem.

Computed bandwidth for approximations of P

$$f(x) = \frac{1}{1 + e^{\beta(x-\mu)}}$$

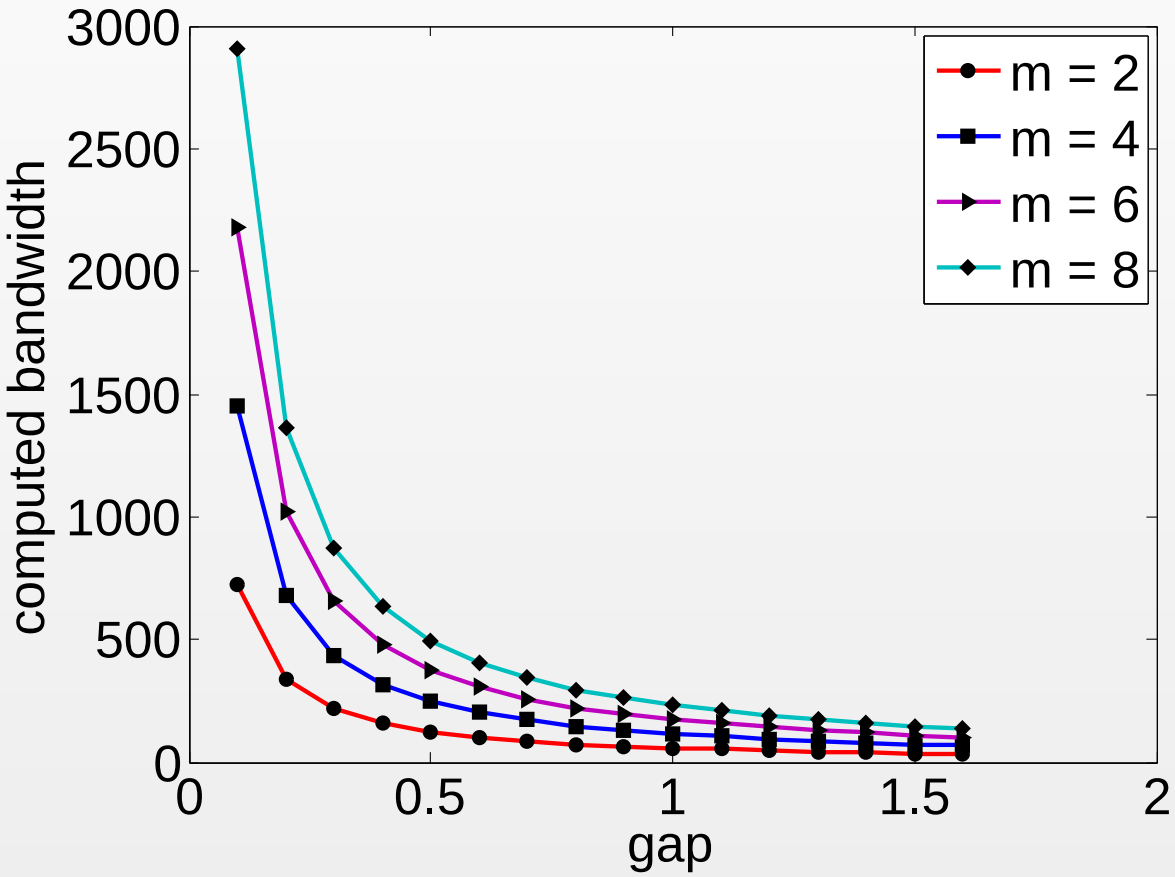

Density matrix, $N_e = 30$, relative gap $\gamma = 0.6$

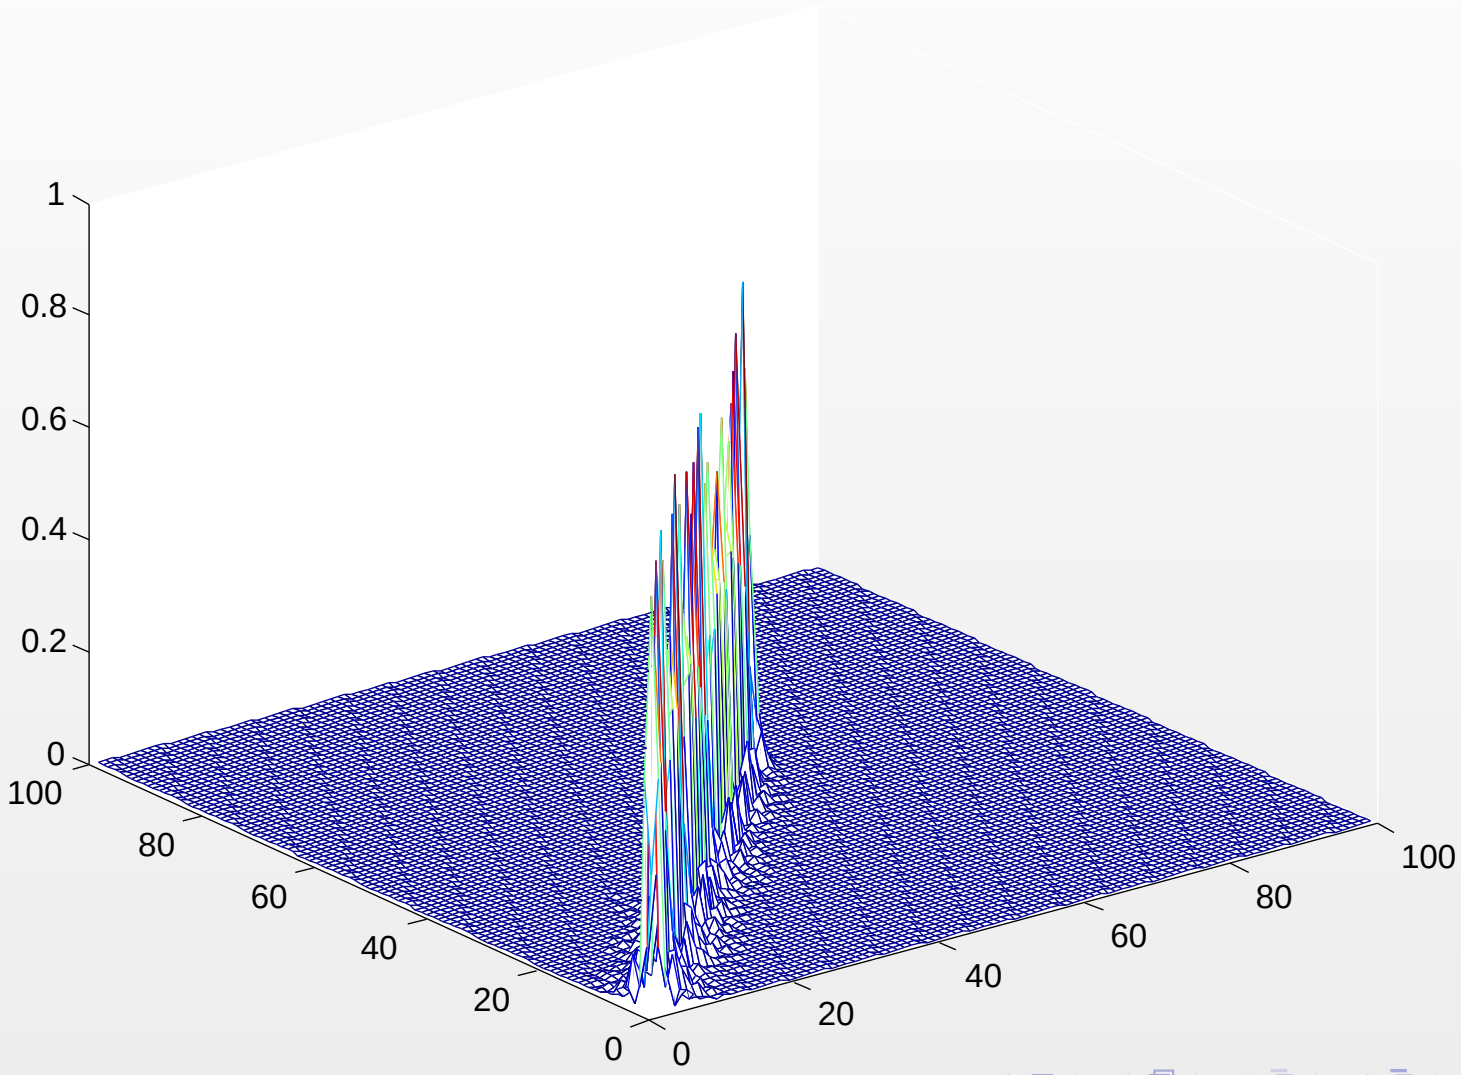

Density matrix, $N_e = 30$, relative gap $\gamma = 0.2$

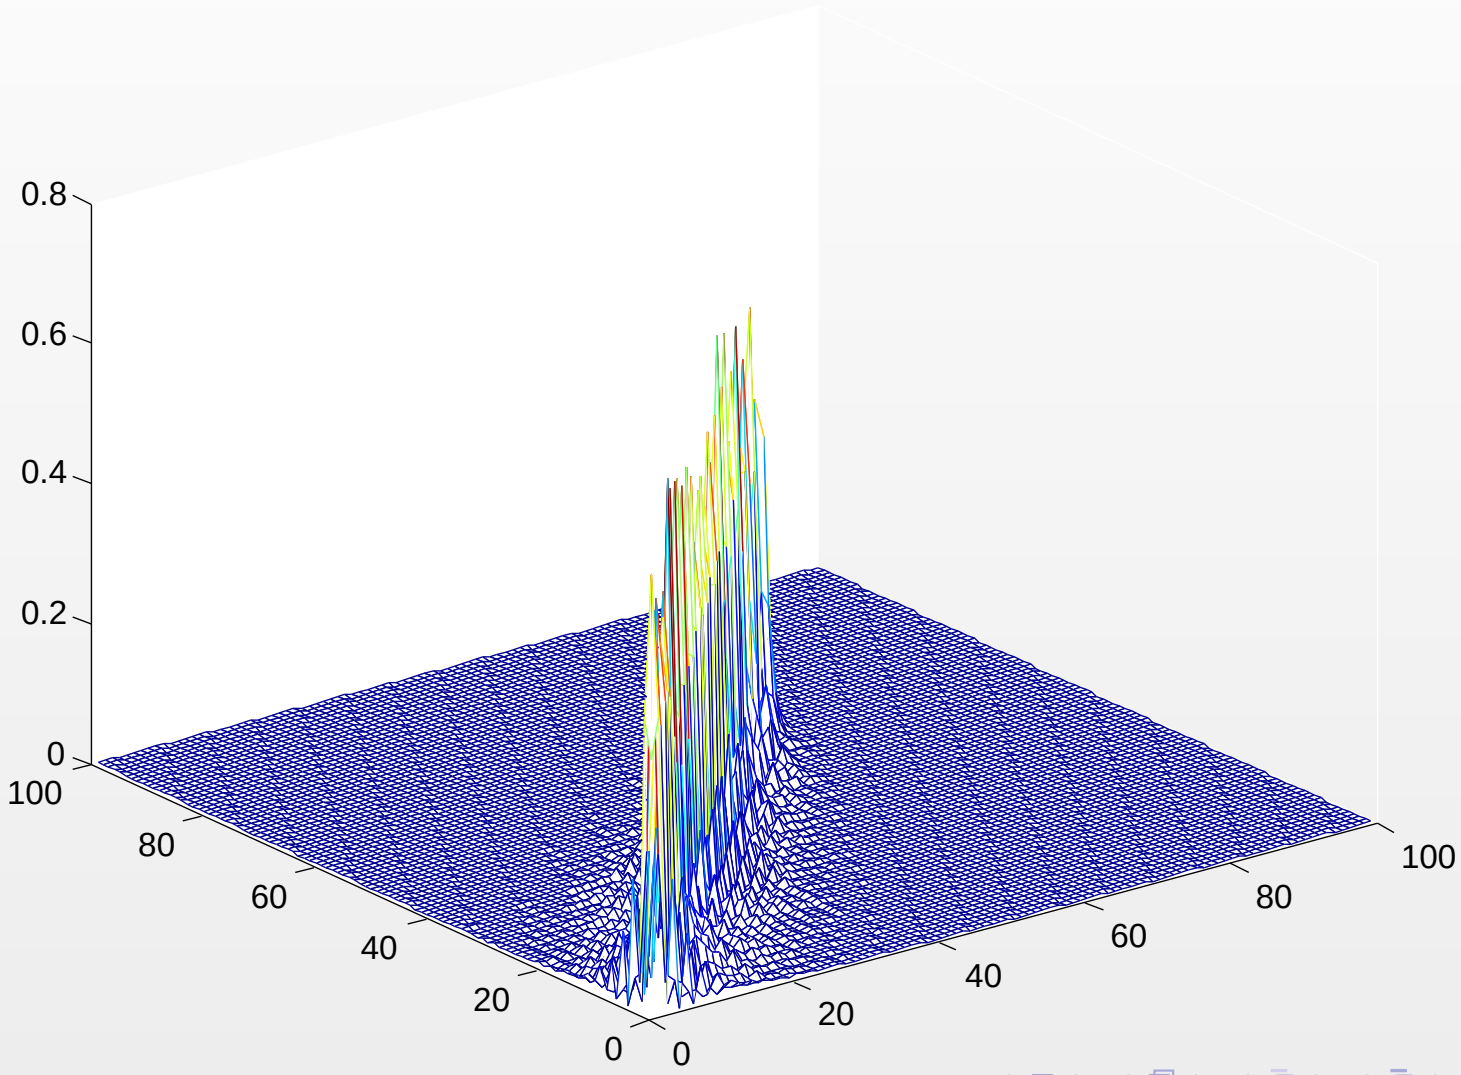

Density matrix, $N_e = 30$, relative gap $\gamma = 0.0001$

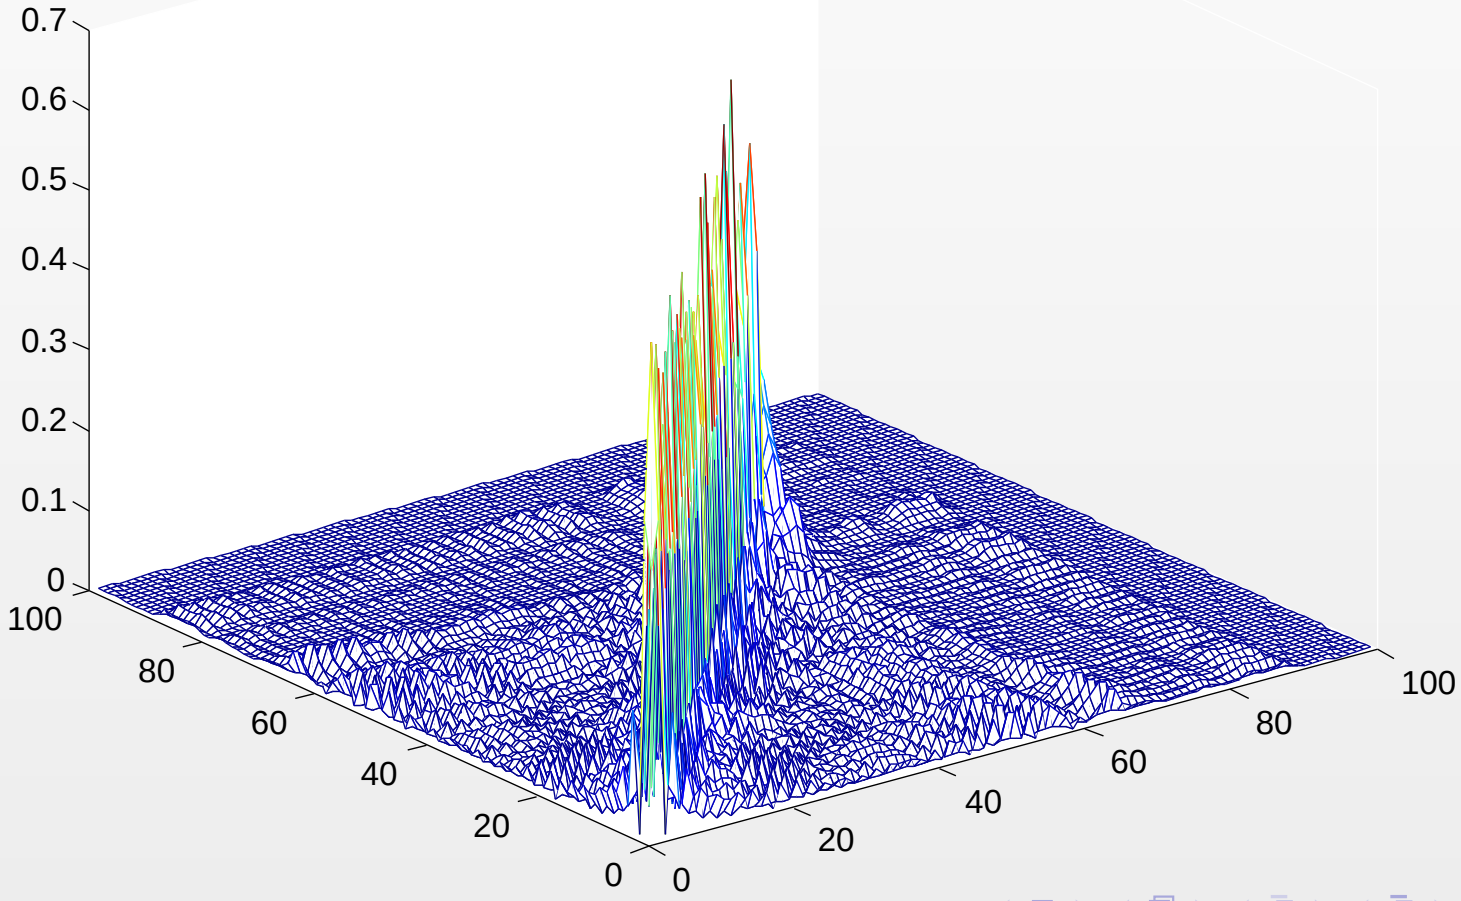

Overview

Decay in analytic functions of sparse matrices

Our approach can be used to establish exponential decay bounds for arbitrary analytic functions of sparse matrices.

Decay in analytic functions of sparse matrices

Our approach can be used to establish exponential decay bounds for arbitrary analytic functions of sparse matrices.

Sparsity pattern of a $2n \times 2n$ Hamiltonian matrix A and decay in e^A.

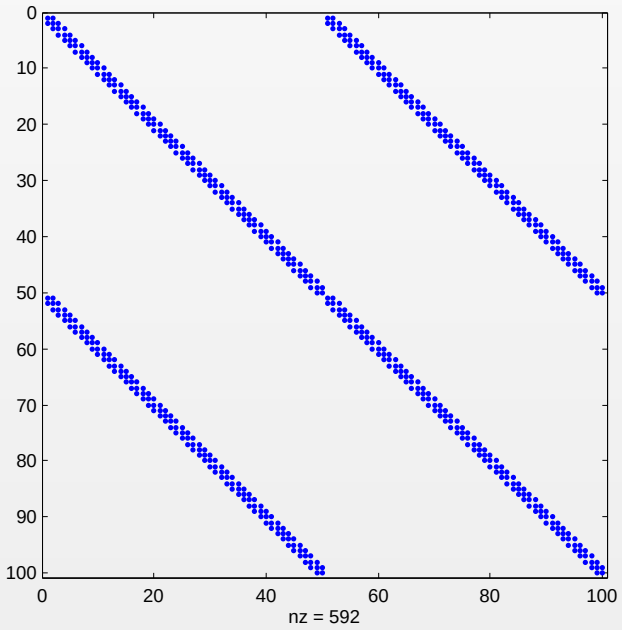

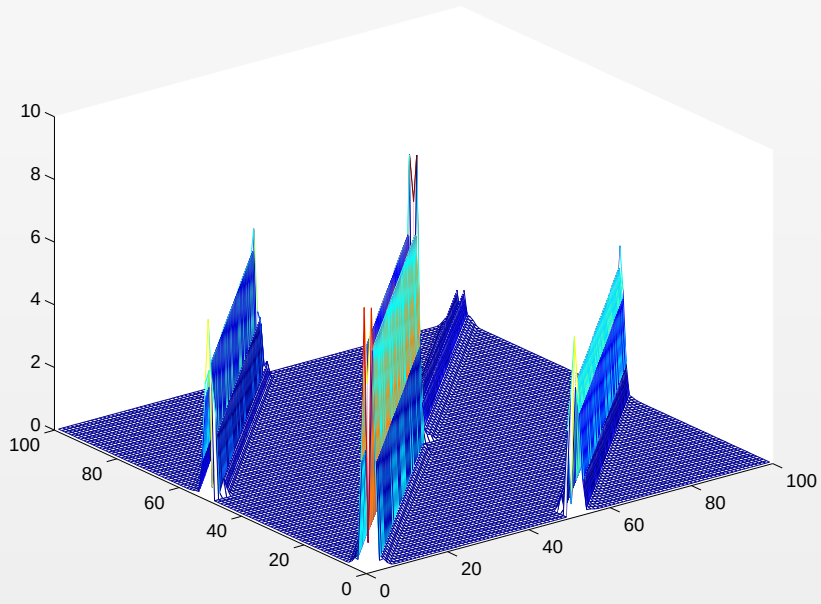

Note that e^A is symplectic.

Decay for logarithm of a sparse matrix

Sparsity pattern of $H = \texttt{mesh3e1}$ (from NASA) and decay in $\log(H)$.

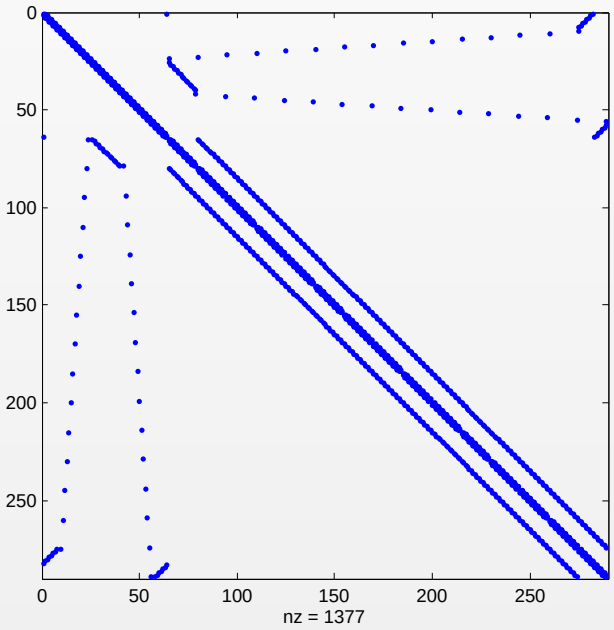

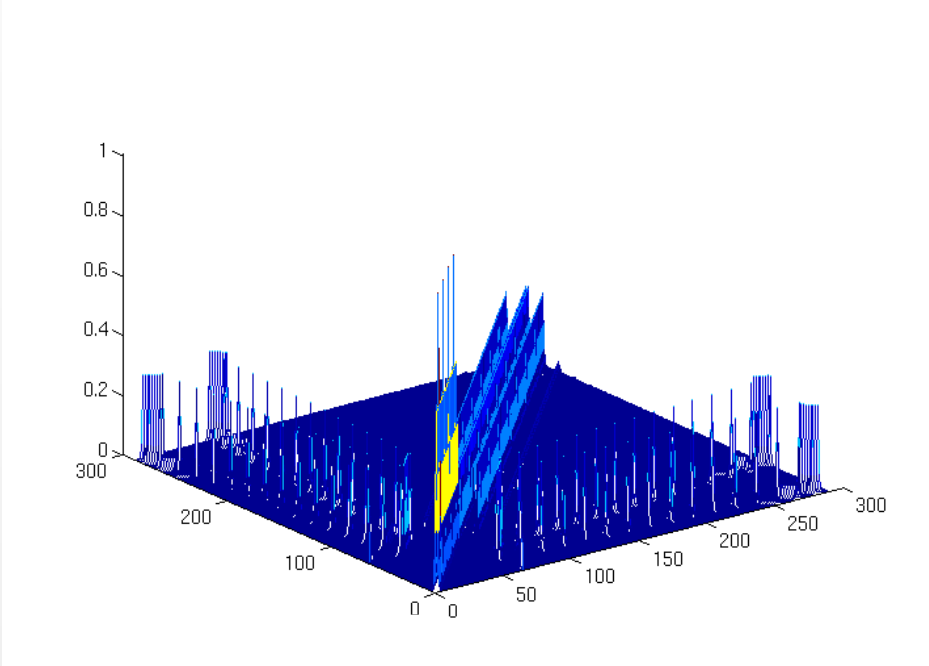

Here H is symmetric positive definite.

Sufficient conditions for $O(N)$ approximation of $f(H)$

- Let $\{H_N\}$ be a sequence of $N \times N$ Hermitian matrices such that there is a closed interval $\mathcal{I}$ with the property that $\sigma(H_N) \subset \mathcal{I}$ for all N

Sufficient conditions for $O(N)$ approximation of $f(H)$

- Let $\{H_N\}$ be a sequence of $N \times N$ Hermitian matrices such that there is a closed interval $\mathcal{I}$ with the property that $\sigma(H_N) \subset \mathcal{I}$ for all N
- Assume that $\{H_N\}$ has bandwidth m independent of N

Sufficient conditions for $O(N)$ approximation of $f(H)$

- Let $\{H_N\}$ be a sequence of $N \times N$ Hermitian matrices such that there is a closed interval $\mathcal{I}$ with the property that $\sigma(H_N) \subset \mathcal{I}$ for all N
- Assume that $\{H_N\}$ has bandwidth m independent of N
- Let f be a function analytic on a neighborhood of $\mathcal{I}$ such that the spectra $\sigma(H_N)$ remain bounded away from any singularities of f as $N \to \infty$

Sufficient conditions for $O(N)$ approximation of $f(H)$

- Let $\{H_N\}$ be a sequence of $N \times N$ Hermitian matrices such that there is a closed interval $\mathcal{I}$ with the property that $\sigma(H_N) \subset \mathcal{I}$ for all N
- Assume that $\{H_N\}$ has bandwidth m independent of N
- Let f be a function analytic on a neighborhood of $\mathcal{I}$ such that the spectra $\sigma(H_N)$ remain bounded away from any singularities of f as $N \to \infty$
- Then for all N and for any $\varepsilon > 0$ there is a matrix B of bandwidth $\hat{m}$ independent of N such that $\|f(H_N) - B\|_2 < \varepsilon$

Sufficient conditions for $O(N)$ approximation of $f(H)$

- Let $\{H_N\}$ be a sequence of $N \times N$ Hermitian matrices such that there is a closed interval $\mathcal{I}$ with the property that $\sigma(H_N) \subset \mathcal{I}$ for all N
- Assume that $\{H_N\}$ has bandwidth m independent of N
- Let f be a function analytic on a neighborhood of $\mathcal{I}$ such that the spectra $\sigma(H_N)$ remain bounded away from any singularities of f as $N \to \infty$
- Then for all N and for any $\varepsilon > 0$ there is a matrix B of bandwidth $\hat{m}$ independent of N such that $\|f(H_N) - B\|_2 < \varepsilon$

- The result can be extended to general sparsity patterns of $\{H_N\}$ (independent of N)

Sufficient conditions for $O(N)$ approximation of $f(H)$

- Let $\{H_N\}$ be a sequence of $N \times N$ Hermitian matrices such that there is a closed interval $\mathcal{I}$ with the property that $\sigma(H_N) \subset \mathcal{I}$ for all N
- Assume that $\{H_N\}$ has bandwidth m independent of N
- Let f be a function analytic on a neighborhood of $\mathcal{I}$ such that the spectra $\sigma(H_N)$ remain bounded away from any singularities of f as $N \to \infty$
- Then for all N and for any $\varepsilon > 0$ there is a matrix B of bandwidth $\hat{m}$ independent of N such that $\|f(H_N) - B\|_2 < \varepsilon$
- The result can be extended to general sparsity patterns of $\{H_N\}$ (independent of N)
- Generalizations to non-normal matrices are possible, e.g., using Crouzeix's theorem.

Overview

Approximation of $f(H)$ by Chebyshev polynomials

Algorithm (Goedecker & Colombo, 1994) ▸ More

- We compute approximations of $f(H)$ using Chebyshev polynomials
  - The degree of the polynomial can be estimated a priori
  - The coefficients of the polynomial can be pre-computed (indep. of N)
  - Estimates for the extreme eigenvalues of H are required
- The polynomial expansion is combined with a procedure that a priori determines a bandwidth or sparsity pattern for $f(H)$ outside which the elements are so small that they can be neglected

Approximation of $f(H)$ by Chebyshev polynomials

Algorithm (Goedecker & Colombo, 1994) ▸ More

- We compute approximations of $f(H)$ using Chebyshev polynomials
  - The degree of the polynomial can be estimated a priori
  - The coefficients of the polynomial can be pre-computed (indep. of N)
  - Estimates for the extreme eigenvalues of H are required
- The polynomial expansion is combined with a procedure that a priori determines a bandwidth or sparsity pattern for $f(H)$ outside which the elements are so small that they can be neglected

Cost

This method is multiplication-rich; the matrices are kept sparse throughout the computation, hence $O(N)$ arithmetic and storage requirements. Matrix polynomials can be efficiently evaluated by the Paterson-Stockmeyer algorithm.

Chebyshev expansion of Fermi-Dirac function

The bandwidth was computed prior to the calculation to be ≈ 20; here H is tridiagonal (1D Anderson model).

Table: Results for $f(x) = \frac{1}{1+\mathrm{e}^{(\beta(x-\mu))}}$

	$\mu = 2,\ \beta = 2.13$			$\mu = 0.5,\ \beta = 1.84$		
N	error	k	$\hat{m}$	error	k	$\hat{m}$
100	9e−06	18	20	6e−06	18	22
200	4e−06	19	20	9e−06	18	22
300	4e−06	19	20	5e−06	20	22
400	6e−06	19	20	8e−06	20	22
500	8e−06	19	20	8e−06	20	22

Computation of Fermi-Dirac function

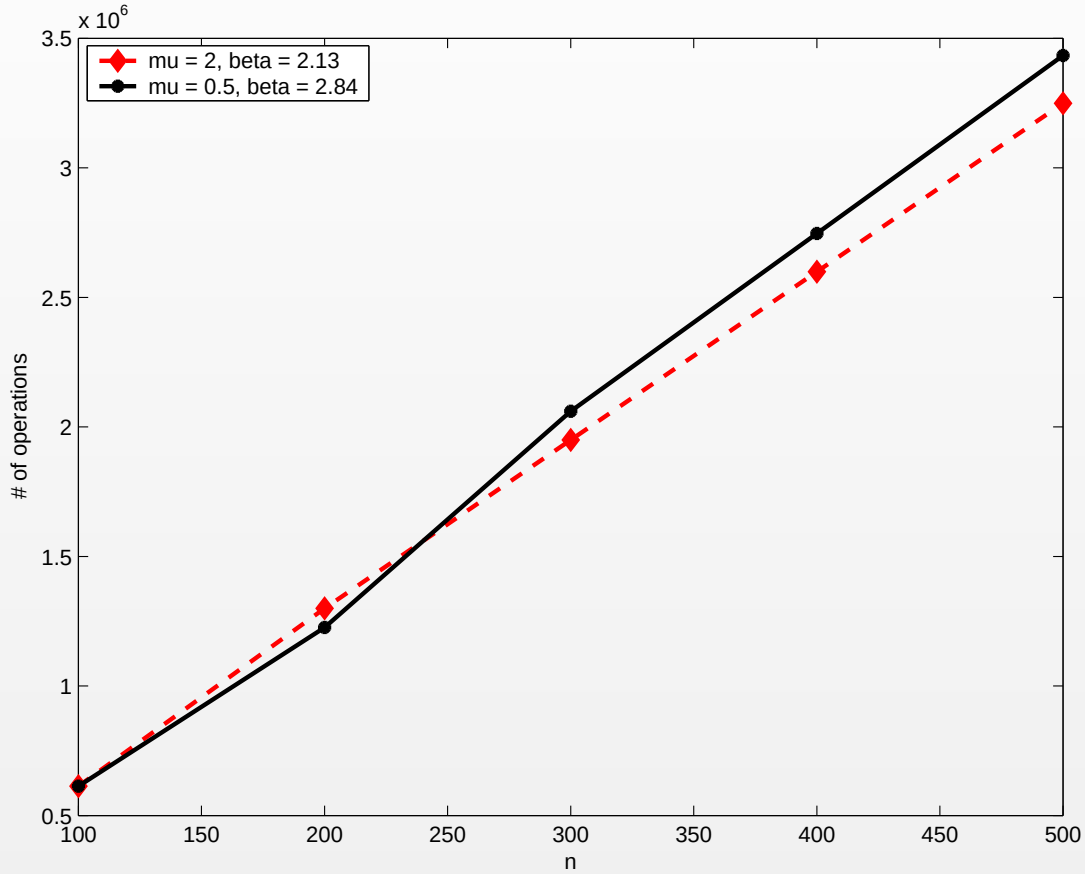

The $O(N)$ behavior of Chebyshev's approximation to the Fermi–Dirac function $f(H) = (\exp(\beta(H - \mu I)) + I)^{-1}$.

Chebyshev expansion of entropy-like function

Some results for $H = H_N$ tridiagonal, SPD, $f(x) = x \log(x)$

	$H \log(H)$	$\mathrm{Tr}[H \log(H)]$		
N	rel. error	error	$\hat{m}$	k
100	$5e{-}07$	$3e{-}04$	20	9
200	$6e{-}07$	$8e{-}04$	20	9
300	$1e{-}07$	$3e{-}04$	20	10
500	$2e{-}07$	$5e{-}04$	20	10

In the Table, $\hat{m}$ is the estimated bandwidth and k is the number of terms in the Chebyshev expansion. Note the $O(N)$ behavior in terms of cost.

Summary

- 'Gapped' systems, like insulators, exhibit strong localization

Summary

- 'Gapped' systems, like insulators, exhibit strong localization
- Localization in $P = f(H)$, when strong enough, can lead to fast approximation algorithms

Summary

- 'Gapped' systems, like insulators, exhibit strong localization
- Localization in $P = f(H)$, when strong enough, can lead to fast approximation algorithms
- Our decay bounds for density matrices depend only on the gap γ and on the sparsity of H; they are, therefore, universal

Summary

- 'Gapped' systems, like insulators, exhibit strong localization
- Localization in $P = f(H)$, when strong enough, can lead to fast approximation algorithms
- Our decay bounds for density matrices depend only on the gap γ and on the sparsity of H; they are, therefore, universal
- These bounds can be useful in determining appropriate sparsity patterns (or bandwidths) that capture the nonnegligible entries in $P = f(H)$

Summary

- 'Gapped' systems, like insulators, exhibit strong localization
- Localization in $P = f(H)$, when strong enough, can lead to fast approximation algorithms
- Our decay bounds for density matrices depend only on the gap γ and on the sparsity of H; they are, therefore, universal
- These bounds can be useful in determining appropriate sparsity patterns (or bandwidths) that capture the nonnegligible entries in $P = f(H)$
- Constants in $O(N)$ algorithms can be large. In practice, to beat $O(N^p)$ algorithms ($p = 2, 3$) on real problems, N must be large enough.

Summary

- 'Gapped' systems, like insulators, exhibit strong localization
- Localization in $P = f(H)$, when strong enough, can lead to fast approximation algorithms
- Our decay bounds for density matrices depend only on the gap γ and on the sparsity of H; they are, therefore, universal
- These bounds can be useful in determining appropriate sparsity patterns (or bandwidths) that capture the nonnegligible entries in $P = f(H)$
- Constants in $O(N)$ algorithms can be large. In practice, to beat $O(N^p)$ algorithms ($p = 2, 3$) on real problems, N must be large enough.
- Extensions to non-Hermitian case possible (applications??)

Overview

Some open problems

- The bounds represent the worst-case scenario and can be pessimistic. To improve them, one must make use of additional information, such as the eigenvalue distribution, which may not be readily available.

Some open problems

- The bounds represent the worst-case scenario and can be pessimistic. To improve them, one must make use of additional information, such as the eigenvalue distribution, which may not be readily available.
- How to deal with metallic systems ($\gamma \to 0$ as $N \to \infty$)?

Some open problems

- The bounds represent the worst-case scenario and can be pessimistic. To improve them, one must make use of additional information, such as the eigenvalue distribution, which may not be readily available.
- How to deal with metallic systems ($\gamma \to 0$ as $N \to \infty$)?
  1. Wavelets?

Some open problems

- The bounds represent the worst-case scenario and can be pessimistic. To improve them, one must make use of additional information, such as the eigenvalue distribution, which may not be readily available.
- How to deal with metallic systems ($\gamma \to 0$ as $N \to \infty$)?
  1. Wavelets?
  2. $\mathcal{H}$-matrices? Multipole methods? Domain Decomposition...?

Some open problems

- The bounds represent the worst-case scenario and can be pessimistic. To improve them, one must make use of additional information, such as the eigenvalue distribution, which may not be readily available.
- How to deal with metallic systems ($\gamma \to 0$ as $N \to \infty$)?
  1. Wavelets?
  2. $\mathcal{H}$-matrices? Multipole methods? Domain Decomposition...?
  3. Optimally localized bases? $P = XX^*$, $X = [u_1, \ldots, u_N]$

Some open problems

- The bounds represent the worst-case scenario and can be pessimistic. To improve them, one must make use of additional information, such as the eigenvalue distribution, which may not be readily available.
- How to deal with metallic systems ($\gamma \to 0$ as $N \to \infty$)?
  1. Wavelets?
  2. $\mathcal{H}$-matrices? Multipole methods? Domain Decomposition...?
  3. Optimally localized bases? $P = XX^*$, $X = [u_1, \ldots, u_N]$
- What is the effect of errors in estimating $\lambda_{\min}$, $\lambda_{\max}$, and γ?

Some open problems

- The bounds represent the worst-case scenario and can be pessimistic. To improve them, one must make use of additional information, such as the eigenvalue distribution, which may not be readily available.
- How to deal with metallic systems ($\gamma \to 0$ as $N \to \infty$)?
  1. Wavelets?
  2. $\mathcal{H}$-matrices? Multipole methods? Domain Decomposition...?
  3. Optimally localized bases? $P = XX^*$, $X = [u_1, \ldots, u_N]$
- What is the effect of errors in estimating $\lambda_{\min}$, $\lambda_{\max}$, and γ?
- What about other asymptotics? E.g., scaling for $C \to \infty$?

Some open problems

- The bounds represent the worst-case scenario and can be pessimistic. To improve them, one must make use of additional information, such as the eigenvalue distribution, which may not be readily available.
- How to deal with metallic systems ($\gamma \to 0$ as $N \to \infty$)?
  1. Wavelets?
  2. $\mathcal{H}$-matrices? Multipole methods? Domain Decomposition...?
  3. Optimally localized bases? $P = XX^*$, $X = [u_1, \ldots, u_N]$
- What is the effect of errors in estimating $\lambda_{\min}$, $\lambda_{\max}$, and γ?
- What about other asymptotics? E.g., scaling for $C \to \infty$?
- Which are the best $O(N)$ algorithms? There are many of them...

Some open problems

- The bounds represent the worst-case scenario and can be pessimistic. To improve them, one must make use of additional information, such as the eigenvalue distribution, which may not be readily available.
- How to deal with metallic systems ($\gamma \to 0$ as $N \to \infty$)?
  1. Wavelets?
  2. $\mathcal{H}$-matrices? Multipole methods? Domain Decomposition...?
  3. Optimally localized bases? $P = XX^*$, $X = [u_1, \ldots, u_N]$
- What is the effect of errors in estimating $\lambda_{\min}$, $\lambda_{\max}$, and γ?
- What about other asymptotics? E.g., scaling for $C \to \infty$?
- Which are the best $O(N)$ algorithms? There are many of them...

An excellent introduction: C. Le Bris, *Computational Chemistry from the Perspective of Numerical Analysis*, Acta Numerica 14 (2005), 363-444.

Localization in spectral projectors: small relative gap

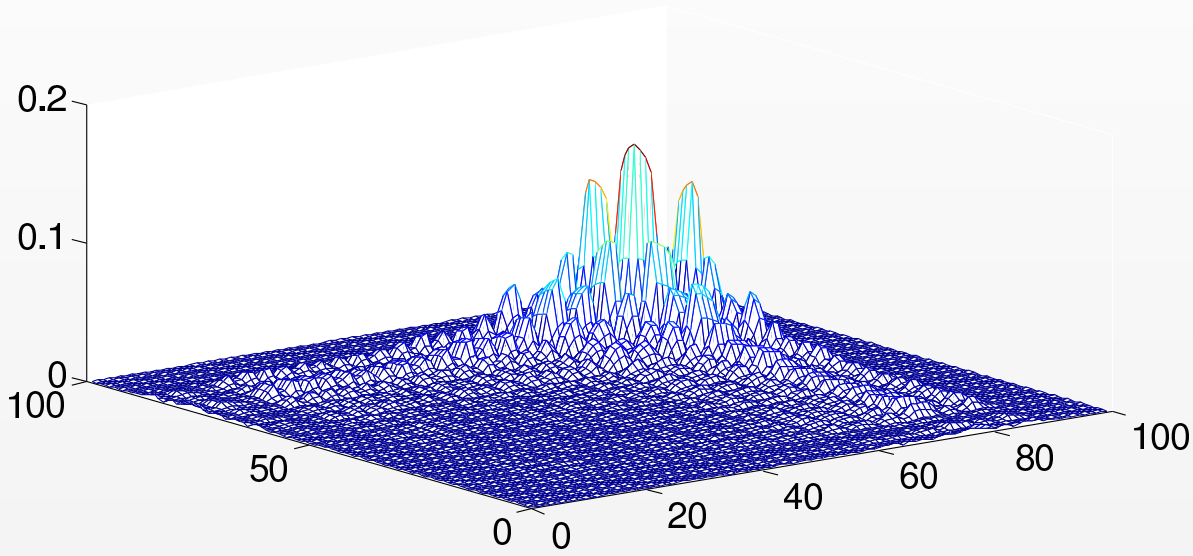

Rank-one spectral projector for $H = H^*$ tridiagonal. Relative gap $\gamma = 10^{-3}$. Note the slow decay and oscillatory behavior of $|P_{ij}|$.

Localization in spectral projectors: large relative gap

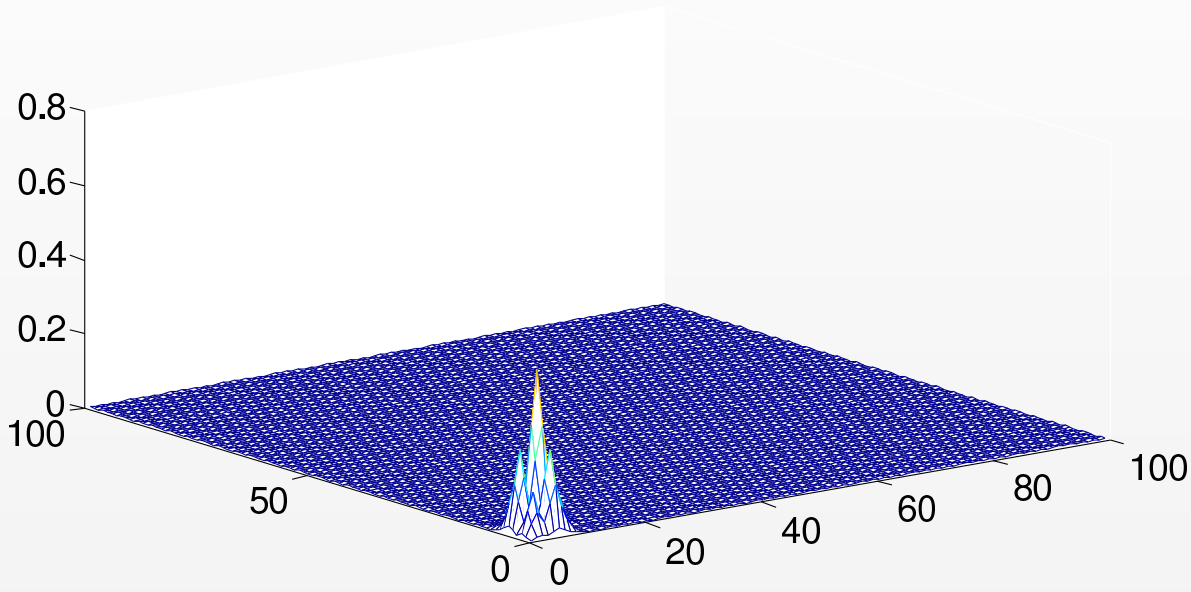

Rank-one spectral projector for $H = H^*$ tridiagonal. Relative gap $\gamma = 0.5$.

▸ Back

Chebyshev approximation

For H with $\sigma(H) \subset [-1, 1]$ the Chebyshev polynomials are given by

$$T_{k+1}(H) = 2HT_k(H) - T_{k-1}(H), \; T_1(H) = H, \; T_0(H) = I.$$

Then $f(H)$ can be represented in a series of the form

$$f(H) = \sum_{k=0}^{\infty} c_k T_k(H).$$

The coefficients of the expansion are given by

$$c_k \approx \frac{2}{M} \sum_{j=1}^{M} f(\cos(\theta_j)) \cos((k-1)\theta_j),$$

where $\theta_j = \pi(j - \frac{1}{2})/M$. Back

The N-independence of the error

The mth truncation error without dropping can be written as

$$\|e_m(H)\| = \|f(H) - \sum_{k=0}^{m} c_k T_k(H)\|.$$

For x in $[-1,1]$ we have that $|T_k(x)| \leq 1$ for $k = 1, 2, \ldots$. Then

$$\|e_m(H)\| = \| \sum_{k=m+1}^{\infty} c_k T_k(H)\| \leq \sum_{k=m+1}^{\infty} |c_k|.$$

Back

A Theorem of Bernstein

The set of Faber polynomials can be used to obtain a uniform approximation to an analytic function f with a sequence of polynomials of bounded degree, i.e.,

$$|f(z) - \Pi_N(z)| < cq^N \quad (0 < q < 1)$$

for all $z \in F$, where c and q depend on the analytic properties of f.

A Theorem of Bernstein

The set of Faber polynomials can be used to obtain a uniform approximation to an analytic function f with a sequence of polynomials of bounded degree, i.e.,

$$|f(z) - \Pi_N(z)| < cq^N \quad (0 < q < 1)$$

for all $z \in F$, where c and q depend on the analytic properties of f.

Example – Disk

If the region is a disk of radius ρ centered at z_0, then for any function f analytic on the disk of radius ρ/q centered at z_0, where $0 < q < 1$, there exists a polynomial Π_N of degree at most N and a positive constant c such that

$$|f(z) - \Pi_N(z)| < cq^N,$$

for all $z \in F$.