Exploiting Localization in Matrix Computations
III. Localization in Matrix Functions

Michele Benzi

Department of Mathematics and Computer Science

Emory University

Atlanta, Georgia, USA

CIME-EMS Summer School in Applied Mathematics

Exploiting Hidden Structure in Matrix Computations:
Theory and Applications

Cetraro, 22-26 June, 2015

Outline

Outline

Decay properties of matrix functions

Generally speaking, functions of banded or sparse matrices, such as the inverse, the exponential, the logarithm, roots, trig functions, etc. are full: all entries are nonzero.

Indeed, suppose that $f(z)$ can be expanded as a power series around z_0:

$$f(z) = a_0 + a_1(z - z_0) + a_2(z - z_0)^2 + \cdots = \sum_{k=0}^{\infty} a_k (z - z_0)^k$$

and that A is a $n \times n$ matrix such that $f(A)$ is defined:

$$f(A) = a_0 I_n + a_1(A - z_0 I_n) + a_2(A - z_0 I_n)^2 + \cdots = \sum_{k=0}^{\infty} a_k (A - z_0 I_n)^k.$$

If A is irreducible then $(A - z_0 I_n)^k$ is structurally full for all $k \geq n - 1$.

Example: $\mathrm{e}^A = \sum_{k=0}^{\infty} \frac{1}{k!} A^k$ with tridiagonal A

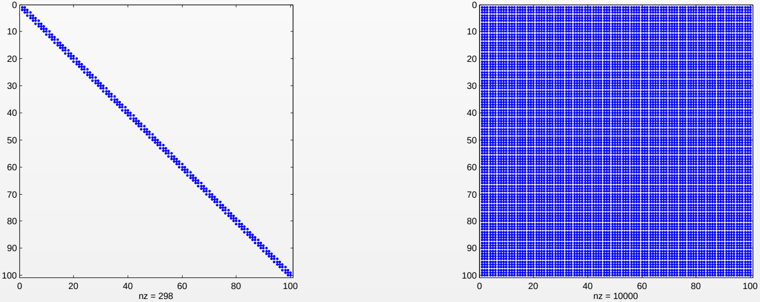

Sparsity pattern of $A = \text{trid}(-1, 2, -1)$ and $\mathrm{e}^A = \text{expm}(A)$.

Decay properties of matrix functions (cont.)

Hence, barring fortuitous cancellation, $f(A)$ is a full matrix.

Simple numerical experiments, however, show that when A is a banded matrix and $f(z)$ is a smooth function for which $f(A)$ is defined, the entries of $f(A)$ usually decay rapidly as one moves away from the diagonal.

The same property is often (but not always!) satisfied by more general sparse matrices: in this case the decay is away from the support (nonzero pattern) of A.

In other words, non-negligible entries of $f(A)$ tend to be strongly localized near the positions (i, j) for which $a_{ij} \neq 0$.

This observation opens up the possibility of developing fast algorithms for approximating functions of sparse matrices.

We are especially interested in the possibility of developing approximation methods with optimal computational complexity, i.e., $\mathcal{O}(n)$ methods.

Example: $\mathrm{e}^A = \sum_{k=0}^{\infty} \frac{1}{k!} A^k$ with tridiagonal A

$|[\mathrm{e}^A]_{ij}|$

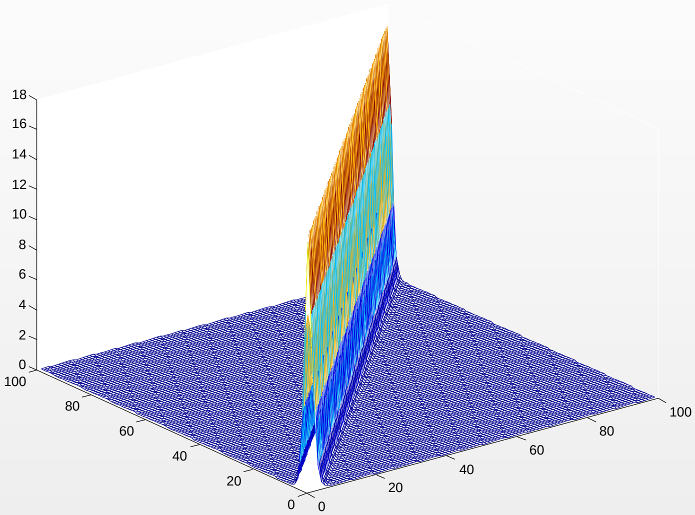

Example: Decay behavior in $\log(A)$, with A sparse, SPD

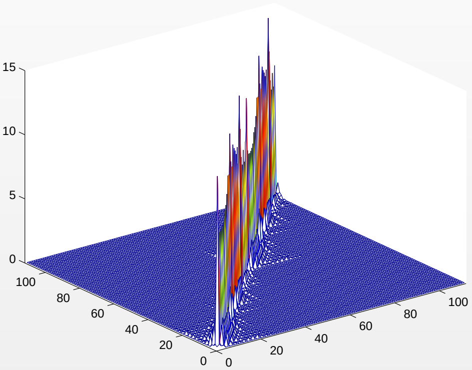

Decay behavior in the logarithm of matrix `bcsstk03` from Matrix Market.

Example: $\cos(A)$ with A banded, nonsymmetric

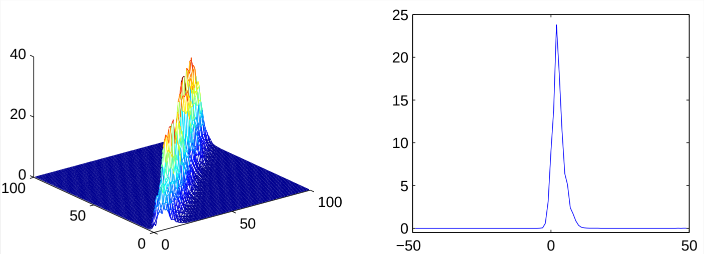

The decay property is not limited to functions of Hermitian matrices.

Example: Decay behavior in e^A, with A Hamiltonian

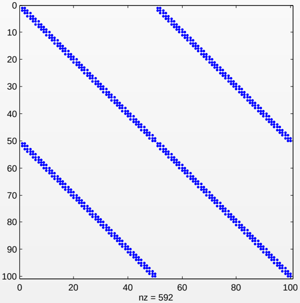

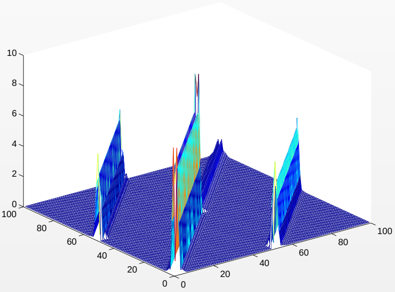

Sparsity pattern of matrix A and decay in e^A. Note that e^A is symplectic.

Some basic questions

These examples suggest the following natural questions:

1. Under which conditions can we expect decay in $f(A)$?
2. Can we obtain sharp bounds on the entries of $f(A)$?
3. Can we characterize the rate of decay in $f(A)$ in terms of
    - the bandwidth/sparsity of A?
    - the spectral properties of A?
    - the location of possible singularities of $f(z)$ in relation to the spectrum of A?
4. What if $f(z)$ is an entire function?
5. When is the rate of decay independent of the matrix size n?

The last point is crucial if we want to develop $\mathcal{O}(n)$ algorithms for approximating functions of sparse matrices.

Some applications

Decay results for functions of sparse matrices have important applications in various areas, including

- Numerical analysis
- Wavelet analysis
- Quantum chemistry (electronic structure)
- Solid state physics
- Quantum information theory (entanglement entropy)
- High-dimensional statistics, random matrix theory
- Control Theory
- ...

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

J. Eisert, M. Cramer, and M. B. Plenio, *Colloquium: Area laws for the entanglement entropy*, Rev. Modern Phys., 82 (2010), pp. 277–306.

The case of A^{-1}

In 1984, Demko, Moss and Smith showed that the entries of A^{-1}, where A is Hermitian positive definite and m-banded ($a_{ij} = 0$ if $|i-j| > m$), satisfy the following exponential decay bound:

$$|[A^{-1}]_{ij}| \leq K\,\lambda^{|i-j|}, \quad \forall i, j$$

where $[a, b]$ is the smallest interval containing the spectrum $\sigma(A)$ of A, $K = \max\{a^{-1}, K_0\}$, $K_0 = (1+\sqrt{\kappa})/2b$, $\lambda = q^{1/m}$,

$$q = q(\kappa) = \frac{\sqrt{\kappa}-1}{\sqrt{k}+1}, \qquad \kappa = \frac{b}{a}\,.$$

The result holds for finite matrices as well as for bounded, infinite matrices acting on ℓ^2.

The case of A^{-1} (cont.)

The DMS result implies that the rate of decay is independent of the matrix size n if the condition number κ and the bandwidth m remain uniformly bounded with respect to n. The bound is sharp.

On the other hand, if either $a \to 0$ or $b \to \infty$ (or both), the (bound on the) decay rate deteriorates as $n \to \infty$.

The same happens if the bandwidth m increases.

The case of A^{-1} (cont.)

Similar results were proved, using different techniques, by Jaffard (1990) for $f(A) = A^{-1}$ and for $f(A) = A^{-1/2}$, where A is not sparse, but has entries that decay exponentially.

Improved bounds for the inverses of matrices of the form

$$A = T \otimes I + I \otimes T$$

(with T banded) are due to Canuto, Simoncini and Verani (2014).

S. Demko, W. F. Moss and P. W. Smith, *Decay rates for inverses of band matrices*, Math. Comp., 43 (1984), pp. 491–499.

S. Jaffard, *Propriétés des matrices "bien localisées" près de leur diagonale et quelques applications*, Ann. Inst. Henri Poincarè, 7 (1990), pp. 461–476.

C. Canuto, V. Simoncini and M. Verani, *On the decay of the inverse of matrices that are sum of Kronecker products*, Linear Algebra Appl., 452 (2014), pp. 21–39.

The case of A^{-1} (cont.)

The DMS estimate implies the following (asymptotic) uniform approximation result:

Theorem

Let A_n be a sequence of $n \times n$ matrices, all HPD and m-banded for $n \to \infty$ (with m independent of n). Assume further that there exists an interval $[a, b]$, $0 < a < b < \infty$, such that $\sigma(A_n) \subset [a, b]$, for all n. Then, for all $\varepsilon > 0$ and for all n there exist an integer $p = p(\varepsilon, m, a, b)$ (independent of n) and a matrix $B_n = B_n^*$ with bandwidth p such that $\|A_n^{-1} - B_n\|_2 < \varepsilon$.

Each approximate inverse $B_n \approx A_n^{-1}$ can be expressed as a polynomial in A_n. The DMS conditions imply that the degree of this polynomial does not depend on n, only on ε. Hence, the approximation requires only $\mathcal{O}(n)$ flops/storage, and is therefore optimal.

Decay bounds for general matrix functions

The proof of DMS is based on a result of Chebyshev on the best uniform approximation over $[a, b]$ of the function $f(x) = x^{-1}$ by polynomials, which shows that the error $\|p_k(x) - x^{-1}\|_\infty$ decays exponentially fast in the degree k. Combined with the Spectral Theorem (which allows to go from scalar functions to matrix functions, with the $\|\cdot\|_2$ matrix norm replacing the $\|\cdot\|_\infty$ norm), this result gives the exponential decay bound for $[A^{-1}]_{ij}$.

Replacing the result of Chebyshev with a more general result of Bernstein on the best uniform approximation of analytic functions by polynomials yields exponential decay bounds for analytic functions of band matrices.

M. Benzi and G. H. Golub, *Bounds for the entries of matrix functions with applications to preconditioning*, BIT, 39 (1999), pp. 417–418.

Decay bounds for general matrix functions (cont.)

As in the case of the DMS bounds for A^{-1}, the rate of decay is controlled by the distance between the poles of $f(z)$ and the spectral interval $[a, b]$, and by the bandwidth m of A.

WLOG, we assume that $A = A^*$ has been scaled and shifted so that $\sigma(A) \subset [-1, 1]$. Then there is a family of ellipses $\mathcal{E} = \mathcal{E}(\chi)$ with foci -1 and 1 and semiaxes $\kappa_1 > 1$, $\kappa_2 > 0$, parameterized by $\chi = \kappa_1 + \kappa_2 > 1$, such that $f(z)$ is analytic inside each $\mathcal{E}$ and continuous on $\mathcal{E}$.

We have a whole family of exponential decay bounds:

Theorem (B., Golub)

Let $M = M(\chi) = \max_{z \in \mathcal{E}} |f(z)|$ and $\lambda = \lambda(\chi) = \chi^{-\frac{1}{m}}$. Then

$$|[f(A)]_{ij}| \le K\, \lambda^{|i-j|}, \quad \forall i, j = 1, 2, \dots$$

where $K = \max\{K_0, \|f(A)\|_2\}$ and $K_0 = \frac{2\chi M}{\chi - 1}$.

Decay bounds for general matrix functions (cont.)

Proof

Bernstein's Theorem states that

$$E_k(f, [-1, 1]) := \|f - p_k\|_\infty \leq K_0\, q^k, \quad k = 0, 1, \ldots,$$

where $q = \chi^{-1}$ and p_k is the polynomial of degree k for which the approximation error $\|f - p_k\|_\infty$ is minimum. Next, observe that if A is m-banded then A^k (and therefore $p_k(A)$) is km-banded: $[p_k(A)]_{ij} = 0$ if $|i - j| > km$. For $i \neq j$ write $|i - j| = km + l$, $l = 1, 2, \ldots, m$, hence $k < |i - j|/m$ and $q^k < q^{\frac{|i-j|}{m}} = \lambda^{|i-j|}$. Therefore, for all $i \neq j$ we have

$$|[f(A)]_{ij}| = |[f(A)]_{ij} - [p_k(A)]_{ij}| \leq \|f(A) - p_k(A)\|_2 \leq \|f - p_k\|_\infty < K_0 \lambda^{|i-j|}.$$

For $i = j$ we have $|[f(A)]_{ii}| \leq \|f(A)\|_2$, therefore the bound holds $\forall i, j$.

Decay bounds for general matrix functions (cont.)

Letting $\alpha = -\ln \lambda > 0$ the bounds can be written as

$$|[f(A)]_{ij}| \le K\,\mathrm{e}^{-\alpha\,|i-j|}, \quad \forall i, j.$$

Note that the bounds can be "optimized", for each i, j, by choosing the ellipse $\mathcal{E}(\chi)$ that yields the minimum value of $K\,\mathrm{e}^{-\alpha\,|i-j|}$.

Decay bounds for general matrix functions (cont.)

There is a clear trade-off in the choice of χ: a larger value of χ yields a larger value of α, and thus a faster decay rate in the exponential, but also a larger prefactor K (through M). If $f(z)$ has poles near $[-1,1]$, then $\chi \approx 1$ and decay may be slow.

What if the function is entire, i.e, analytic everywhere? Example: $f(A) = \mathrm{e}^A$.

In this case, χ can be taken arbitrarily large and decay is now superexponential. In this case, however, different techniques give sharper bounds, see

A. Iserles, *How large is the exponential of a banded matrix?*, New Zealand J. Math., 29:177–1992, 2000.

M. Benzi an V. Simoncini, *Decay bounds for functions of Hermitian matrices with banded or Kronecker structure*, to appear in SIAM J. on Matrix Analysis and Applications (2015).

Decay bounds for general matrix functions (cont.)

We have the following uniform approximation result:

Theorem

Let A_n be a sequence of $n \times n$ Hermitian matrices, all m-banded (with m independent of n). Assume further that there exists an interval $[a, b]$ with $-\infty < a < b < \infty$ such that $\sigma(A_n) \subset [a, b]$, for all n. Let $f(z)$ be analytic on an open set $\Omega \subseteq \mathbb{C}$ containing $[a, b]$, and such that $f(x)$ is real for real x. Then, for all $\varepsilon > 0$ and for all n there exist an integer $p = p(\varepsilon, m, a, b)$ (which does not depend on n) and a matrix $B_n = B_n^*$ with bandwidth p such that $\|f(A_n) - B_n\|_2 < \varepsilon$.

Again, B_n can be expressed as a polynomial $p_k(A_n)$ with fixed degree $k = k(\varepsilon)$ (independent of n). Thus, optimal $\mathcal{O}(n)$ approximation is possible.

Outline

Sparse approximate inverses

Consider a large linear system of equations $A\mathbf{x} = \mathbf{b}$.

If we have an approximation $M \approx A^{-1}$, with M sparse, then we can use M as a preconditioner in a Krylov subspace method, such as CG or GMRES.

For example, if we know that A^{-1} decays rapidly away from the main diagonal, then one could use a banded approximation M.

More generally, if we know that A^{-1} decays rapidly away from some nonzero pattern $\mathcal{S} = \{(i,j) \,|\, 1 \leq i,j \leq n\}$, then we can approximate A^{-1} by a sparse matrix with nonzero pattern in $\mathcal{S}$, perhaps by means of Frobenius norm minimization:

$$\|I - AM\|_F \to \text{min}, \quad \text{s. t.} \quad \text{supp}(M) \subseteq \mathcal{S}.$$

Information on the decay pattern of A^{-1} can be used to determine $\mathcal{S}$. The main appeal of this approach is its high potential for parallelization.

M. Benzi, *Preconditioning techniques for large linear systems: a survey*, J. Comput. Phys., 182 (2002), pp. 418–477.

Sparse approximate inverses (cont.)

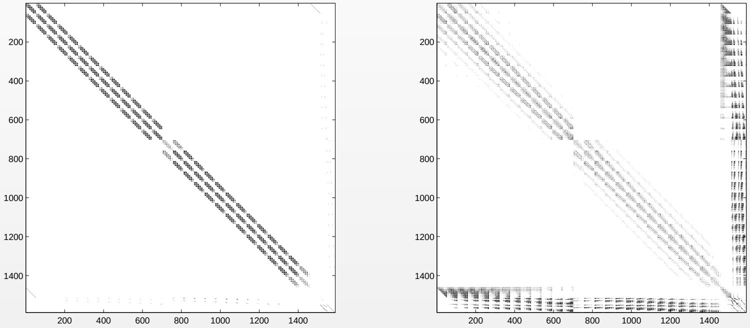

Coefficient matrix and sparse approximate inverse: FEM for fluid-structure interaction problem. Preconditioned Bi-CGSTAB converges in 39 iterations. Here $nnz(M)/nnz(A) \approx 1.56$.

Perturbation theory for eigenvalues

Assume $A = A^*$, and consider the eigenvalue problem $A\mathbf{x} = \lambda\mathbf{x}$. For simplicity we assume that λ is a simple eigenvalue.

If $\mathbf{x}$ is normalized ($\|\mathbf{x}\|_2 = 1$), then

$$\lambda = \mathbf{x}^* A\mathbf{x} = \sum_{i=1}^{n}\sum_{j=1}^{n} a_{ij}\bar{x}_i x_j.$$

Suppose now we want to estimate how sensitive λ is to small changes in the entries of A. We have $(i \neq j)$

$$\frac{\partial\lambda}{\partial a_{ij}} = \bar{x}_i x_j + \bar{x}_j x_i.$$

Now, the (i, j) entry of the spectral projector $P = \mathbf{x}\mathbf{x}^*$ is $P_{ij} = x_i\bar{x}_j$. Therefore,

$$|P_{ij}| \approx 0 \Rightarrow \lambda \quad \text{is insensitive to small changes in } a_{ij},\ a_{ji} = \bar{a}_{ij}.$$

Perturbation theory for eigenvalues (cont.)

The spectral projector P is a "delta" function in A. It can be written as a Lagrange interpolation polynomial in A:

$$P = L(A), \quad L(x) = \prod_{\mu \in \sigma(A) \setminus \{\lambda\}} \frac{x - \mu}{\lambda - \mu}.$$

This function can be approximated with arbitrary accuracy by a Gaussian:

$$P \approx \exp(-\beta(A - \lambda I)^2)$$

for $\beta > 0$ sufficiently large. The size of β depends on the gap between λ and the nearest eigenvalue $\mu \neq \lambda$ and on the desired accuracy.

Assume now that A is banded, e.g., tridiagonal. Since P is an entire function of A, its entries must decay super-exponentially away from the main diagonal. Hence, $|i - j| \gg 1 \Rightarrow |P_{ij}| \approx 0$.

Therefore, replacing an off-diagonal zero in position (i, j) with $|i - j| \gg 1$ with a small nonzero will have almost no effect on λ.

Perturbation theory for eigenvalues (cont.)

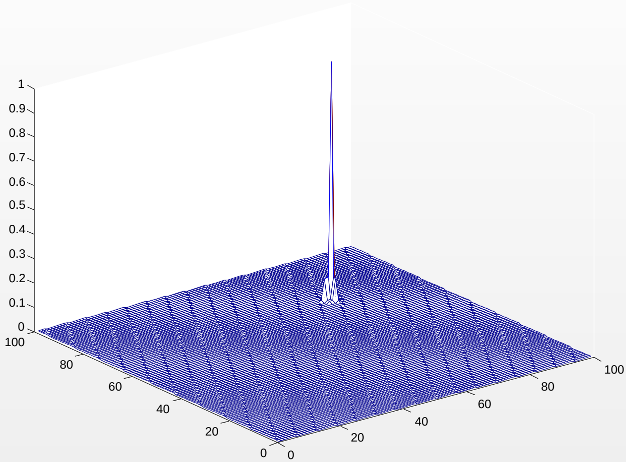

The spectral projector onto the invariant subspace corresponding to an isolated eigenvalue of a banded matrix.

Perturbation theory for eigenvalues (cont.)

More generally, suppose we are interested in computing the quantity ("total energy")

$$\mathrm{Tr}(PA) = \lambda_1 + \lambda_2 + \cdots + \lambda_k, \tag{1}$$

where P is the orthogonal projector onto the invariant subspace spanned by the eigenvectors corresponding to the k smallest eigenvalues of A, a sparse matrix.

If the gap $\lambda_{k+1} - \lambda_k$ is "large", then the entries of P can be shown to decay exponentially away from the sparsity pattern of A (see Lecture IV).

Differentiating (1) with respect to a_{ij} shows again that the total energy is insensitive to small perturbations in positions of A that are far from the nonzero pattern supp(A).

This fact has important consequences in quantum chemistry.

Matrix factorizations

Localization is not limited to matrix functions such as the inverse, the exponential, the square root, and so forth. It is also present in many standard matrix factorizations.

For instance, if A is SPD and sparse, and $A = LL^T$ is its Cholesky factorization, then L may be much less sparse than A, but its entries will usually decay rapidly outside of the nonzero pattern of the lower triangular part of A.

Under appropriate conditions, similar remarks also apply to other matrix factorizations, such as LU and QR, and to the inverse factors L^{-1}, U^{-1}, R^{-1}, etc. This fact is the main reason behind the success of incomplete factorization preconditioners.

M. Benzi and M. Tuma, *Orderings for factorized sparse approximate inverse preconditioners*, SIAM J. Sci. Comput., 21 (2000), pp. 1851–1868.

I. Krishtal, T. Strohmer, and T. Wertz, *Localization of matrix factorizations*, Found. Comput. Math., to appear (2015).

Network communicability

Let $G = (V, E)$ be a graph with adjacency matrix A (sparse).

Recall that the communicability between two nodes $i, j \in V$ is given by

$$C(i,j) = [\mathrm{e}^A]_{ij} = \sum_{k=0}^{\infty} \frac{[A^k]_{ij}}{k!}.$$

Bounds on $C(i,j)$ can be very useful in many situations. Low values of $C(i,j)$ indicate poor communication, for example due to bottlenecks, between the two nodes. Conversely, high values of $C(i,j)$ indicate that two nodes, even though they are not directly connected ($a_{ij} = 0$), are nevertheless strongly connected. This can be used for instance for community detection, or for network robustness studies.

E. Estrada, *Community detection based on network communicability*, Chaos, 21 (2011), 016103.

F. Arrigo and M. Benzi, *Updating and downdating techniques for optimizing network communicability*, arXiv:1410.5303v2 (2015).

Network communicability (cont.)

In general, communicability values are low for distant pairs of nodes if the graph G has large diameter, for example, if G corresponds to a path, or a Cartesian product of paths (so-called grid graphs), and generally for spatial networks, like road networks.

On the other hand, for small-world graphs (such as social networks, collaboration networks, protein-protein interaction networks, and airline connection networks) the communicability is usually very high between most pairs of nodes.

Hence, communicability is localized for the first type of graph, and delocalized for the second. This correspons to the intuition that information diffuses slowly in the first case, and rapidly in the second case.

Sexual contact network of injecting drug users

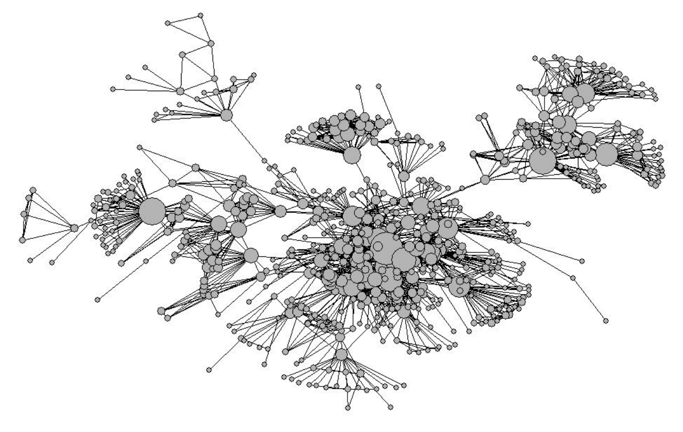

Figure courtesy of Ernesto Estrada.

Sexual contact network of injecting drug users (cont.)

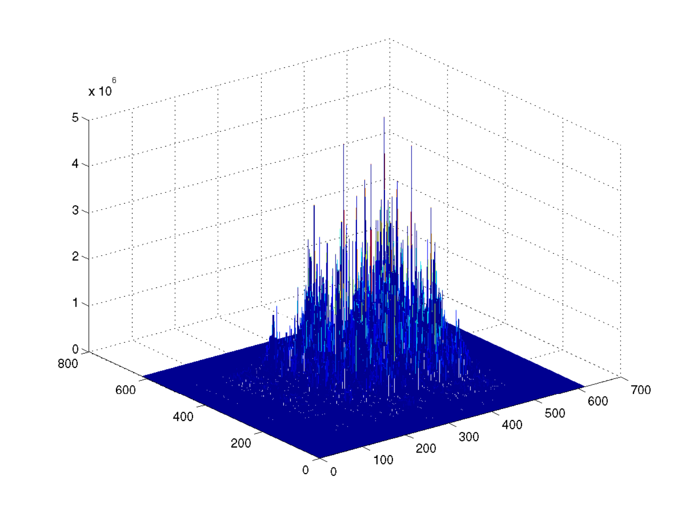

Communicability matrix e^A for the previous network.

Outline

Extension 1: General sparsity patterns

In most applications one deals with sparse matrices rather than banded ones. We can associate to every $n \times n$ matrix A a graph $G = (V, E)$ with $|V| = n$ such that there is an edge $(i, j) \in E$ if and only if $a_{ij} \neq 0$; if A has a symmetric sparsity pattern, then G is undirected.

A path joining nodes i, j is a list of edges $(i, i_1), (i_1, i_2), \ldots, (i_k, j)$ in E such that the nodes $i, i_1, \ldots, i_k, j$ are all distinct. The distance $d(i, j)$ of two nodes i and j is defined as the length of the shortest path in G connecting i and j.

If the graph is undirected, this is an actual distance (i.e., a metric). Also, if A is irreducible then G is connected and $d(i, j)$ is well defined for any pair of nodes.

These notions can be easily adapted to the directed case.

Extension 1: General sparsity patterns (cont.)

Theorem

Let $A = A^*$ with $\sigma(A) \subset [-1, 1]$ and let $f(z)$ be analytic in an open region containing $[-1, 1]$ and such that $f(x)$ is real for real x. Then there exist constants $K > 0$ and $\alpha > 0$ such that

$$|[f(A)]_{ij}| \leq K\, \mathrm{e}^{-\alpha\, d(i,j)}, \quad \forall i, j.$$

The constants K and α can be explicitly determined as before.

If $\{A_n\}$ is a sequence of $n \times n$ matrices with a bounded number of nonzeros per row and $\sigma(A_n) \subset [-1, 1]$ for all n, then α and K can be chosen independent of n and $\mathcal{O}(n)$ approximation is again possible. The sparsity assumption ensures that the distance $d(i, j)$ can grow to infinity as $|i - j| \to \infty$. In other words, the graph diameter $\mathrm{diam}(G(A_n))$ must be unbounded for $n \to \infty$ if $d_{\mathsf{max}}((G(A_n)) \leq D_{\mathsf{max}}$ for all n.

Extension 2: Kronecker structure

Recall that the Kronecker product (or tensor product) of two matrices A and B of sizes $n_a \times m_a$ and $n_b \times m_b$, respectively, is defined as

$$A \otimes B := \begin{bmatrix} a_{11}B & a_{12}B & \cdots & a_{1m_a}B \\ a_{21}B & a_{22}B & \cdots & a_{2m_a}B \\ \vdots & \vdots & \ddots & \vdots \\ a_{n_a1}B & a_{n_a2}B & \cdots & a_{n_am_a}B \end{bmatrix} \in \mathbb{C}^{n_an_b \times m_am_b}.$$

The Kronecker sum of two matrices $M_1 \in \mathbb{C}^{n_1 \times n_1}$, $M_2 \in \mathbb{C}^{n_2 \times n_2}$ is the matrix

$$\mathcal{A} = M_1 \oplus M_2 := M_1 \otimes I_{n_2} + I_{n_1} \otimes M_2 \in \mathbb{C}^{N \times N},$$

where $N = n_1 n_2$. These definitions are easily extended to the case of more than two factors or summands.

Matrices with Kronecker structure arise in a number of important applications.

Extension 2: Kronecker structure (cont.)

The simplest example is the 2D discrete Laplacian, which is the Kronecker sum of two 1D discrete Laplacians.

More generally, the graph Laplacian of the Cartesian product of two graphs is the Kronecker sum of of the Laplacians of the component graphs.

The matrix exponential is especially well-behaved with respect to Kronecker operations, since it satisfies

$$\exp(M_1 \oplus M_2) = \exp(M_1) \otimes \exp(M_2)$$

Exploiting this property allows us to get very good bounds for the entries of $\exp(\mathcal{A})$, where $\mathcal{A} = M_1 \oplus M_2$, in terms of bounds for the entries of $\exp(M_1)$ and $\exp(M_2)$. These bounds are generally better than those in terms of the graph distance.

In particular, the oscillatory behavior of the entries of $\exp(\mathcal{A})$ is perfectly captured by these bounds.

Extension 2: Kronecker structure (cont.)

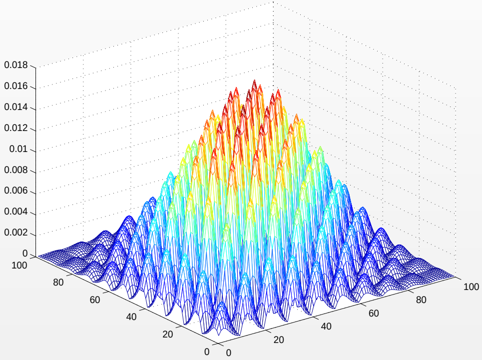

Decay plot for $[\exp(-5\mathcal{A})]_{ij}$ where $\mathcal{A}$ is the 5-point finite difference discretization of the negative Laplacian on the unit square on a 10×10 uniform grid with zero Dirichlet boundary conditions.

Extension 2: Kronecker structure (cont.)

The decay bounds for $\exp(\mathcal{A})$ can be extended to other matrix functions by means of integral representations, as follows. Assume that $f(x)$ can be represented as the Laplace–Stieltjes transform of a positive measure $d\alpha(t)$:

$$f(x) = \int_0^\infty \mathrm{e}^{-tx} d\alpha(t), \quad x \in (a,b).$$

Many important functions have this property: for example,

$$x^{-1/2} = \int_0^\infty \mathrm{e}^{-tx} \frac{dt}{\sqrt{\pi t}}, \quad x > 0.$$

Hence, if $\mathcal{A}$ is SPD we can write

$$[\mathcal{A}^{-1/2}]_{ij} = \int_0^\infty \left[\mathrm{e}^{-t\mathcal{A}}\right]_{ij} \frac{dt}{\sqrt{\pi t}}, \quad 1 \leq i,j \leq n$$

and by evaluating or bounding these integrals we obtain decay bounds on the entries of $\mathcal{A}^{-1/2}$.

M. Benzi an V. Simoncini, *Decay bounds for functions of Hermitian matrices with banded or Kronecker structure*, to appear in SIMAX (2015).

Extension 2: Kronecker structure (cont.)

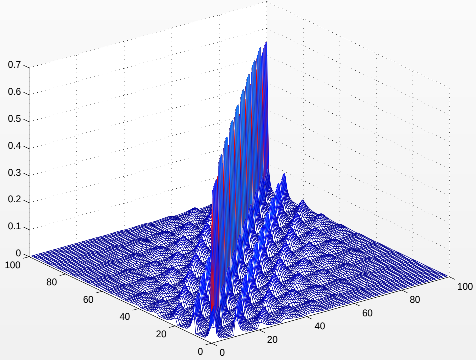

Three-dimensional decay plot for $[\mathcal{A}^{-1/2}]_{ij}$ where $\mathcal{A}$ is the 5-point finite difference discretization of the negative Laplacian on the unit square on a 10×10 uniform grid with zero Dirichlet boundary conditions.

Extension 3: Algebras of matrices with decay

Suppose now that A is not necessarily banded or sparse, but its entries decay according to some law, e.g.,

$$|a_{ij}| \leq C_0\, \phi(|i-j|), \quad \text{where } \phi(r) \to 0 \quad \text{for } r \to \infty,$$

and C_0 is a constant. The decay could be exponential, or a power law. Here A is assumed to be an infinite matrix (bounded linear operator on the sequence space ℓ^p, for some $p \geq 1$).

Using the general theory of Banach algebras, Jaffard (1990) has shown that for matrices of this type, the inverse (if it exists) must satisfy a similar decay law:

$$|[A^{-1}]_{ij}| \leq C_1 \psi(|i-j|),$$

where C_1 is a constant and ψ is of the same kind as ϕ (in fact $\phi = \psi$ in the case of a power law).

Extension 3: Algebras of matrices with decay (cont.)

Using the contour integral representation formula,

$$f(A) = \int_\Gamma f(z)(zI - A)^{-1} dz,$$

it is possible to extend this result to more general matrix functions.

Let $\mathcal{A}$ denote the algebra of matrices with decay. The fact that $A \in \mathcal{A} \Rightarrow A^{-1} \in \mathcal{A}$ (if A is invertible) is a noncommutative analogue of Wiener's Lemma, which states that if f is a non-vanishing function and the Fourier expansion $f(x) = \sum c_n \mathrm{e}^{inx}$ is absolutely convergent (that is, $\sum |c_n| < \infty$), then the function $g(x) = 1/f(x)$ has the same property.

K. Gröchenig and M. Leinert, *Symmetry and inverse-closedness of matrix algebras and functional calculus for infinite matrices*, Trans. Amer. Math. Soc., 358 (2006), pp. 2695–2711.

K. Gröchenig and A. Klotz, *Noncommutative approximation: inverse-closed subalgebras and off-diagonal decay of matrices*, Constr. Approx., 32 (2010), pp. 429–466.

Extension 4: Non-Hermitian matrices

There are several ways to extend these results to the non-Hermitian case.

Things are straightforward in the case of normal matrices ($AA^* = A^*A$), using the theory of Faber polynomials (a generalization of Chebyshev polynomials to the complex domain; Bernstein's Theorem continues to hold).

If A is diagonalizable, $A = XDX^{-1}$, then similar bounds hold but now the prefactor K contains the additional term $\kappa_2(X) = \|X\|_2 \|X^{-1}\|_2$. If this is very large the bounds can be terrible, even useless (compare this with the situation for the standard GMRES error bound). Note that $\kappa_2(X) = 1$ if A is normal, and we recover the known bounds.

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.

Extension 4: Non-Hermitian matrices (cont.)

A more elegant alternative is to obtain decay bounds for general A by applying Crouzeix's Theorem. Recall that the numerical range, or field of values of a matrix, or linear operator, A is the set

$$\mathcal{W}(A) := \{\langle Ax, x\rangle \,|\, \langle x, x\rangle = 1\} \subset \mathbb{C}.$$

Theorem (Crouzeix)

Let A be a bounded linear operator on a Hilbert space $\mathcal{H}$ and let $g(z)$ be analytic in an open region $\Omega \subseteq \mathbb{C}$ containing the closure $\overline{\mathcal{W}(A)}$ of $\mathcal{W}(A)$. Then

$$\|g(A)\| \leq \mathcal{Q}\, \|g\|_{\infty, \overline{\mathcal{W}(A)}},$$

where $2 \leq \mathcal{Q} \leq 11.08$.

M. Crouzeix, *Numerical range and functional calculus in Hilbert space*, J. Funct. Anal., 244 (2007), pp. 668–690.

Extension 4: Non-Hermitian matrices (cont.)

Using this result, we can prove the following general decay result:

Theorem (B., Boito)

Let $\{A_n\}$ be a sequence of $n \times n$ matrices and let $f(z)$ be analytic on an open region $\Omega \subseteq \mathbb{C}$ containing a compact set $\mathcal{C}$ such that $\mathcal{W}(A_n) \subseteq \mathcal{C}$, for all n. Then there exist explicitly computable constants $K > 0$, $\alpha > 0$ such that

$$|[f(A)]_{ij}| \leq K\, \mathrm{e}^{-\alpha\, d(i,j)}, \quad \forall i, j.$$

Corollary

If A_n is a sequence of sparse matrices of order n such that each row contains at most m nonzeros (with m independent of n) and if the poles of $f(z)$ remain bounded away from $\mathcal{C}$ as $n \to \infty$, we can approximate $f(A_n)$ with a polynomial $p_k(A_n)$ of fixed degree $k = k(\varepsilon, m, \mathcal{C})$, in $\mathcal{O}(n)$ operations and storage.

Bounds for non-Hermitian matrices

A simple approach to the computation of constants K and α $(=-\ln\lambda)$ in the decay bound goes as follows. The set $\mathcal{C}$ in the last theorem can be be chosen as as a disk of sufficiently large radius r. Assume that $f(z)$ is analytic on an open neighborhood of the disk of center 0 and radius $R > r$. The standard theory of complex Taylor series gives the following estimate for the Taylor approximation error:

$$\|f - T_k\|_{\infty,\mathcal{C}} \le \frac{M(R)}{1-\frac{r}{R}} \left(\frac{r}{R}\right)^{k+1},$$

where $M(R) = \max_{|z|=R} |f(z)|$. Therefore we can choose

$$K = \max\left\{\|f(A)\|_{\infty,\mathcal{C}},\ \mathcal{Q}\, M(R)\frac{r}{R-r}\right\}, \qquad \lambda = \left(\frac{r}{R}\right)^{1/m}.$$

Note that R is somewhat arbitary; any value of R will do, as long as $r < R < \min|\zeta|$, where ζ varies over the poles of f (if f is entire, we let $\min|\zeta| = \infty$).

Bounds for non-Hermitian matrices (cont.)

Choosing as large a value of R as possible gives a better asymptotic decay rate, but also a potentially large constant K. For practical purposes, one may therefore want to pick a value of R that ensures a good trade-off between the magnitude of K and the decay rate.

In summary, use of Crouzeix's Theorem allows us to obtain decay bounds for nonnormal matrices without the diagonalizability assumption, and to replace the factor $\kappa_2(X)$ in the bounds with a small, universal constant.

The drawback is that f must be analytic on a set containing the field of values, and that $\|f(A)\|_{\infty,\mathcal{W}(A)}$ could be very large.

In practice, if A is diagonalizable with an ill-conditioned eigenbasis, the decay bounds based on Crouzeix's Theorem are often tighter than those containing $\kappa_2(X)$.

Bounds for non-Hermitian matrices (cont.)

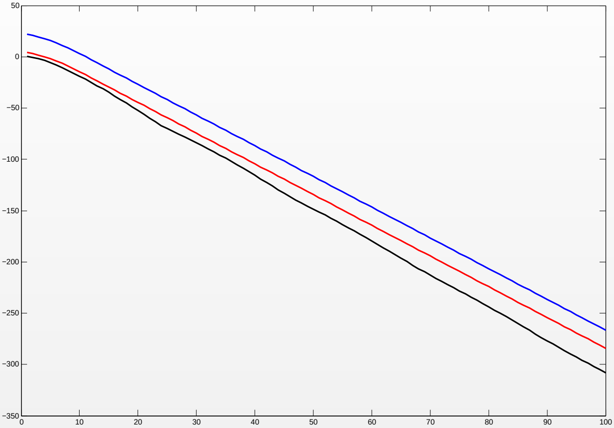

Example: A is the 100×100 Toeplitz matrix generated by the symbol $\phi(t) = 2t^{-1} + 1 + 3t$. We have $\kappa_2(X) = 5.26 \cdot 10^8$. In this logarithmic plot: in black the first row of e^A, in blue the κ_2-bound, in red the Crouzeix bound.

Extension 5: Matrices over C^*-algebras

On September 5, 2011, I received the following email:

```
Dear Sir,

I am currently doing research in the University of Oxford about
functions of matrices with non-commuting matrix elements.

When reading your article ``Decay Bounds and O(n) algorithms for
approximating functions of sparse matrices", I was wondering if
your results hold for non-commuting elements as well?  It looks
to me that your results should hold so long as we use the
appropriate norm to evaluate the decay.  Could you tell if [you]
know whether this is indeed the case or not?

Thank you very much,
Pierre-Louis Giscard,
Atomic & Laser Physics
University of Oxford
```

Extension 5: Matrices over C^*-algebras (cont.)

As it turns out, the answer is yes: decay holds not only for functions of matrices with entries in $\mathbb{R}$ or in $\mathbb{C}$, but also in rather general normed algebras.

For instance, the matrix entries could be other matrices, or quaternions, or bounded linear operators on a Hilbert space.

In the process of finding the answer to Giscard's question, we also realized that decay results even hold for functions of matrices whose entries are (continuous) functions of one or more variables—i.e., for matrix functions of matrix-valued functions.

As a special case, one can also define functions of certain tensors and prove decay results for these.

M. Benzi and P. Boito, *Decay properties for functions of matrices over C^*-algebras*, Linear Algebra Appl., 456 (2014), pp. 174–198.

C^*-algebras

The natural setting for the desired extension is provided by the theory of C^*-algebras. These algebras play an important role in functional analysis (operator theory, spectral theory), harmonic analysis, control theory, and quantum physics. They also have applications in numerical analysis (see References).

Recall that a *Banach algebra* is an algebra $\mathcal{A}_0$ with a norm making $\mathcal{A}_0$ into a Banach space and satisfying

$$\|ab\| \leq \|a\|\|b\|$$

for all $a, b \in \mathcal{A}_0$. Here we consider only *unital* complex Banach algebras, i.e., algebras with a multiplicative unit I with $\|I\| = 1$.

C^*-algebras (cont.)

An *involution* on a Banach algebra $\mathcal{A}_0$ is a map $a \to a^*$ of $\mathcal{A}_0$ into itself satisfying

1. $(a^*)^* = a$
2. $(ab)^* = b^* a^*$
3. $(\lambda a + b)^* = \overline{\lambda} a^* + b^*$

for all $a, b \in \mathcal{A}_0$ and $\lambda \in \mathbb{C}$ (a^* is called the *adjoint* of a).

A C^**-algebra* is a Banach algebra with an involution such that the following C^**-identity*,

$$\|a^* a\| = \|a\|^2 ,$$

holds for all $a \in \mathcal{A}_0$. Note that we do not make any assumption on whether $\mathcal{A}_0$ is commutative or not.

There can be at most one norm $\| \cdot \|$ which makes $\mathcal{A}_0$ into a C^*-algebra.

Examples of C^*-algebras

1. The commutative algebra $C(X)$ of all continuous complex-valued functions on a compact Hausdorff space X. Here the addition and multiplication operations are defined pointwise, and the norm is given by
$$\|f\|_\infty = \max_{t \in X} |f(t)|.$$
The involution on $C(X)$ maps each function f to its complex conjugate f^*, defined by $f^*(t) = \overline{f(t)}$ for all $t \in X$.
2. The algebra $\mathcal{B}(\mathcal{H})$ of all bounded linear operators on a complex Hilbert space $\mathcal{H}$, with the operator norm
$$\|T\| = \sup_{x \neq 0} \frac{\|Tx\|_{\mathcal{H}}}{\|x\|_{\mathcal{H}}}.$$
The involution on $\mathcal{B}(\mathcal{H})$ maps each bounded linear operator T on $\mathcal{H}$ to its adjoint, T^*.

C^*-algebras (cont.)

The second example contains as a special case the algebra $\mathbb{C}^{n\times n}$ of all $n \times n$ matrices with complex entries, with the norm being the usual spectral norm and the involution mapping each matrix $A = [a_{ij}]$ to its conjugate transpose $A^* = [\,\overline{a_{ji}}\,]$.

Examples 1 and 2 above provide, in a precise sense, the "only" examples of C^*-algebras. Indeed, every (unital) commutative C^*-algebra admits a faithful representation as an algebra of the form $C(X)$ for a suitable (and essentially unique) compact Hausdorff space X; and, similarly, every unital (possibly noncommutative) C^*-algebra can be represented faithfully onto a norm-closed $*$-subalgebra of $\mathcal{B}(\mathcal{H})$ for a suitable complex Hilbert space $\mathcal{H}$.

Note: a subalgebra of a C^*-algebra is called a $*$-*subalgebra* if it is closed under involution.

C^*-algebras (cont.)

It turns out that the theory of analytic functions, including Cauchy's integral formula, can be generalized *verbatim* from the scalar case

$$f : \Omega \subseteq \mathbb{C} \longrightarrow \mathbb{C}$$

to the case of functions taking values in a Banach algebra:

$$f : \Omega \subseteq \mathbb{C} \longrightarrow \mathcal{A}.$$

Moreover, since $\mathcal{A}^{n \times n}$ is also a Banach algebra, we may generalize the theory to analytic functions

$$f : \Omega \subseteq \mathbb{C} \longrightarrow \mathcal{A}^{n \times n}.$$

The case of matrix functions is just a special case of this so-called holomorphic functional calculus, corresponding to $\mathcal{A} = \mathbb{C}$.

The C^*-algebra case corresponds to using the spectral norm $\|A\|_2$.

Hence, all our results generalize almost without changes to the case of functions of matrices over $\mathcal{A}$.

Example: Hermitian matrix over $\mathcal{A}_0 = C([0,1])$

Consider the tridiagonal Hermitian Toeplitz matrix-valued function of size 20×20

$$A = A(t) = \begin{bmatrix} 1 & \mathrm{e}^{-t} & & \\ \mathrm{e}^{-t} & 1 & \ddots & \\ & \ddots & \ddots & \mathrm{e}^{-t} \\ & & \mathrm{e}^{-t} & 1 \end{bmatrix}, \quad t \in [0,1].$$

- Scale $A(t)$ so that $\sigma(A(t)) \subset [-1,1]$.
- Compute (symbolically!) the Chebyshev approximation $\mathrm{e}^{A(t)} \approx \sum_{k=0}^{12} c_k A(t)^k$.
- The approximation error is bounded in norm by $3.9913 \cdot 10^{-14}$.

Example: Hermitian matrix over $\mathcal{A}_0 = C([0,1])$ (cont.)

Decay behavior in $\|[e^{A(t)}]_{ij}\|_\infty$. Comparison with bounds (first row, $\log_{10}$-plot):

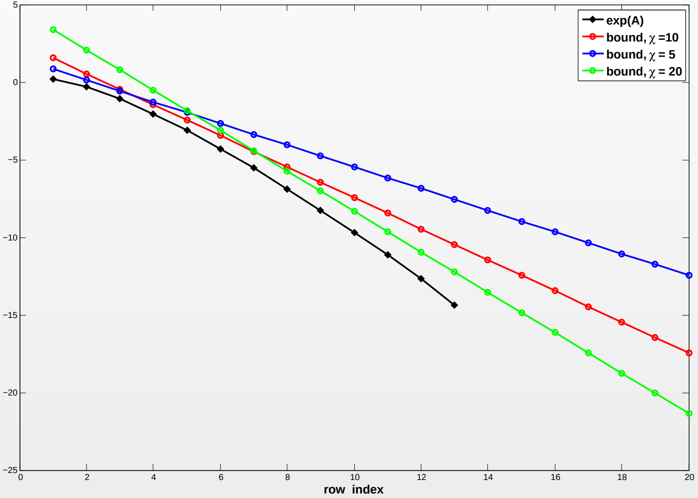

Example: inverse of non-Hermitian matrix over $\mathcal{A}_0 = C([1,2])$

$$B(t) = \begin{bmatrix} 1 & e^{-t} & & \\ & \ddots & \ddots & \\ & & 1 & e^{-t} \\ & & & 1 \end{bmatrix}, \quad C(t) = \begin{bmatrix} 1 & & & \\ e^{\frac{1}{3}-\frac{t}{2}} & \ddots & & \\ & \ddots & 1 & \\ & & e^{\frac{1}{3}-\frac{t}{2}} & 1 \end{bmatrix}.$$

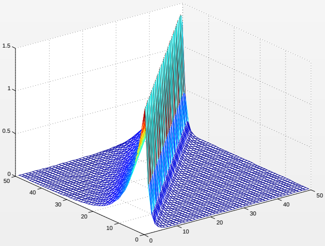

Decay behavior of $\|[A(t)^{-1}]_{ij}\|_\infty$ where $A(t) = B(t)C(t)$

The case of quaternions

What about functions of matrices with entries in $\mathbb{H}$, the algebra of quaternions? Note that $\mathbb{H}$ is *not* a complex C^*-algebra (it is a real one).

However, quaternions can be represented as 2×2 matrices over $\mathbb{C}$, and this allows us to extend our decay theorems to functions of matrices over $\mathbb{H}$, with one restriction: f must be expressible as a power series with real coefficients. Indeed, $\mathbb{H}^{n\times n}$ can be identified with a real subalgebra of $\mathbb{C}^{2n\times 2n}$.

$$\mathbb{H} = \{q = a + b\mathbf{i} + c\mathbf{j} + d\mathbf{k} \,:\, a,b,c,d \in \mathbb{R}\} \cong \left\{ Q = \begin{pmatrix} a + b\mathbf{i} & c + d\mathbf{i} \\ -c + d\mathbf{i} & a - b\mathbf{i} \end{pmatrix} \in \mathbb{C}^{2\times 2} \right\}$$

Note: $\|q\| = \sqrt{a^2 + b^2 + c^2 + d^2} = \|Q\|_2$.

Decay in functions of matrices over $\mathbb{H}$

Recall:

$$\mathbb{H} = \{a + b\mathbf{i} + c\mathbf{j} + d\mathbf{k} : a, b, c, d \in \mathbb{R}\} \cong \left\{ \begin{pmatrix} a + b\,\mathbf{i} & c + d\,\mathbf{i} \\ -c + d\,\mathbf{i} & a - b\,\mathbf{i} \end{pmatrix} \in \mathbb{C}^{2\times 2} \right\}$$

- $A \in \mathbb{H}^{50\times 50}$ is an Hermitian Toeplitz tridiagonal matrix with quaternionic entries.
- $\widetilde{A} \in \mathbb{C}^{100\times 100}$ is the associated complex block matrix.
- Compute $f(\widetilde{A})$ in Matlab, convert it back to a matrix with quaternionic entries, plot norms of entries.

Example: matrix over the quaternion algebra $\mathbb{H}$

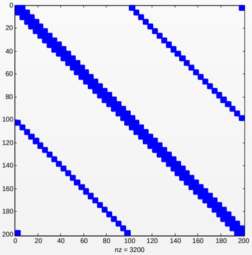

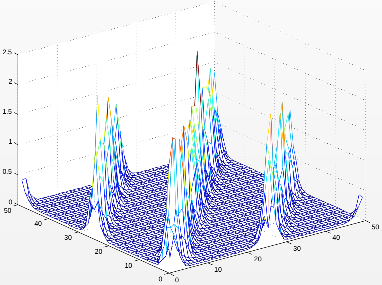

Left plot: sparsity pattern of $A \in \mathbb{H}^{50\times 50}$

Right plot: decay in the norm of the entries of e^A

Example over $\mathbb{H}$: decay in the exponential

Norms of entries of e^A:

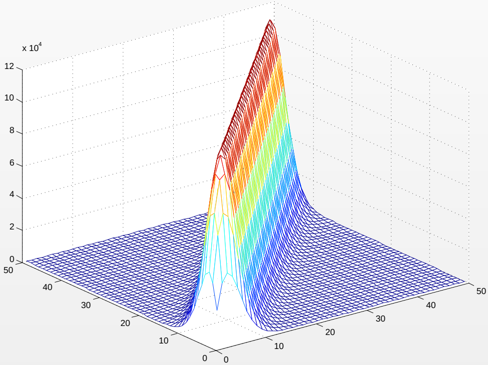

Example over $\mathbb{H}$: decay in the logarithm

Norms of entries of $\log(A)$:

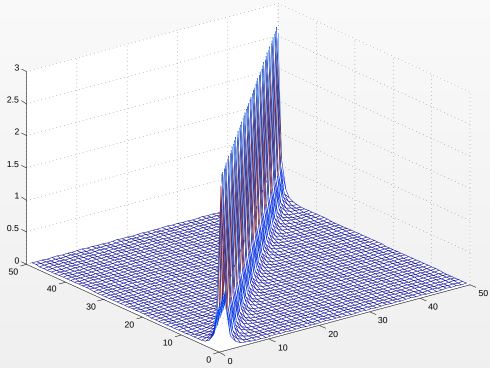

Example over $\mathbb{H}$: decay in the inverse square root

Norms of entries of $A^{-\frac{1}{2}}$:

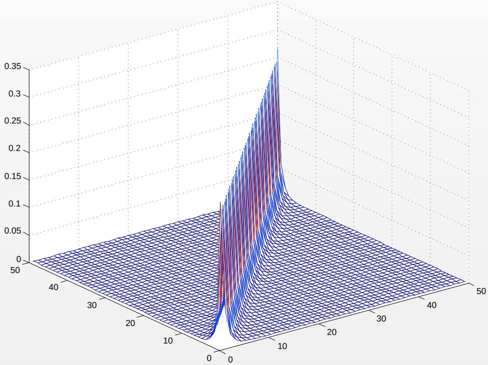

Outline

Conclusions

In conclusion:

- Decay results for matrix functions are important in several areas of mathematics and in applications
- When present, localization can be exploited to produce $\mathcal{O}(n)$ approximations of otherwise apparently dense matrix problems
- Results can be extended to non-normal matrices and to different types of sparsity and decay
- Such results hold for functions of matrices over rather general normed algebras, not just $\mathbb{R}$ or $\mathbb{C}$

Outline

Monographs

For a very good introduction to approximation theory, see

- G. Meinardus, *Approximation of Functions: Theory and Numerical Methods*, Springer, NY, 1967.

Recommended books on the basics of C^*-algebras:

- R. S. Doran, ed., *C^*-Algebras: 1943–1993. A Fifty Year Celebration*, Contemporary Mathematics, 167, American Mathematical Society, Providence, RI, 1994.
- R. Kadison and J. Ringrose, *Fundamentals of the Theory of Operator Algebras*, 2 Volls., Academic Press, Orlando, FL, 1983.
- W. Rudin, *Functional Analysis*, McGraw-Hill, NY, 1983.

Papers

Selected papers on decay:

- 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 rates and $O(n)$ algorithms for approximating functions of sparse matrices*, Electr. Trans. Numer. Anal., 28 (2007), pp. 16–39.
- S. Demko, W. S. Moss and P. W. Smith, *Decay rates for inverses of band matrices*, Math. Comp., 43 (1984), pp. 491–499.
- A. Iserles, *How large is the exponential of a banded matrix?*, New Zealand J. Math., 29 (2000), pp. 177–192.
- S. Jaffard, *Propriétés des matrices "bien localisées" près de leur diagonale et quelques applications*, Ann. Inst. Henri Poincarè, 7 (1990), pp. 461–476.

Papers (cont.)

Some more recent papers on decay:

- M. Benzi and P. Boito, *Decay properties for functions of matrices over C^*-algebras*, Linear Algebra Appl., 456 (2014), pp. 174–198.
- M. Benzi, P. Boito and N. Razouk, *Decay properties of spectral projectors with applications to electronic structure*, SIAM Review, 55 (2013), pp. 3–64.
- M. Benzi an V. Simoncini, *Decay bounds for functions of Hermitian matrices with banded or Kronecker structure*, to appear in SIAM J. on Matrix Analysis and Applications (2015).
- C. Canuto, V. Simoncini and M. Verani, *On the decay of the inverse of matrices that are sum of Kronecker products*, Linear Algebra Appl., 452 (2014), pp. 21–39.
- I. Kryshtal, T. Strohmer, and T. Wertz, *Localization of matrix factorizations*, Found. Comput. Math, in press (2015).
- M. Shao, *On the finite section method for computing exponentials of doubly-infinite skew-Hermitian matrices*, Linear Algebra App., 451 (2014), pp. 65–96.

Papers (cont.)

For a taste of Banach and C^*-algebras in numerical analysis:

- W. Arveson, *C^*-algebras and numerical linear algebra*, J. Funct. Anal., 122 (1994), pp. 333–360.
- A. Böttcher, *C^*-algebras in numerical analysis*, Irish. Math. Soc. Bulletin, 45 (2000), pp. 57–133.
- K. Gröchenig, *Wiener's Lemma: Theme and Variations. An Introduction to Spectral Invariance and its Applications*, Four Short Courses on Harmonic Analysis. Applied and Numerical Harmonic Analysis, Birkhäuser, Boston, 2010.
- K. Gröchenig and A. Klotz, *Noncommutative approximation: inverse-closed subalgebras and off-diagonal decay of matrices*, Constr. Approx., 32 (2010), pp. 429–466.