3. Blanc, X., Bris, C.L.: Periodicity of the infinite-volume ground state of a one-dimensional quantum model. Nonlinear Anal. **48**(6, Ser. A: Theory Methods), 791–803 (2002)
4. Cancès, É., Ehrlacher, V.: Local defects are always neutral in the Thomas–Fermi–von Weiszäcker model for crystals. Arch. Ration. Mech. Anal. **202**(3), 933–973 (2011)
5. Cancès, E., Defranceschi, M., Kutzelnigg, W., Le Bris, C., Maday, Y.: Computational quantum chemistry: a primer. In: Ciarlet, P.G. (ed.) Handbook of Numerical Analysis, vol. X, pp. 3–270. North-Holland, Amsterdam (2003)
6. Cancès, É., Deleurence, A., Lewin, M.: A new approach to the modelling of local defects in crystals: the reduced Hartree-Fock case. Commun. Math. Phys. **281**(1), 129–177 (2008)
7. Cancès, É., Deleurence, A., Lewin, M.: Non-perturbative embedding of local defects in crystalline materials. J. Phys. **20**, 294, 213 (2008)
8. Catto, I., Le Bris, C., Lions, P.L.: The Mathematical Theory of Thermodynamic Limits: Thomas-Fermi Type Models. Oxford Mathematical Monographs. Clarendon/Oxford University Press, New York (1998)
9. Catto, I., Le Bris, C., Lions, P.L.: On the thermodynamic limit for Hartree-Fock type models. Ann. Inst. H. Poincaré Anal. Non Linéaire **18**(6), 687–760 (2001)
10. Fefferman, C.: The thermodynamic limit for a crystal. Commun. Math. Phys. **98**(3), 289–311 (1985)
11. Ghimenti, M., Lewin, M.: Properties of periodic Hartree-Fock minimizers. Calc. Var. Partial Differ. Equ. **35**(1), 39–56 (2009)
12. Hainzl, C., Lewin, M., Solovej, J.P.: The thermodynamic limit of quantum Coulomb systems. Part I. General theory. Adv. Math. **221**, 454–487 (2009)
13. Hainzl, C., Lewin, M., Solovej, J.P.: The thermodynamic limit of quantum Coulomb systems. Part II. Applications. Adv. Math. **221**, 488–546 (2009)
14. Kittel, C.: Introduction to Solid State Physics, 8th edn. Wiley, New York (2004)
15. Lieb, E.H.: Variational principle for many-fermion systems. Phys. Rev. Lett. **46**, 457–459 (1981)
16. Lieb, E.H., Lebowitz, J.L.: The constitution of matter: existence of thermodynamics for systems composed of electrons and nuclei. Adv. Math. **9**, 316–398 (1972)
17. Lieb, E.H., Simon, B.: The Thomas-Fermi theory of atoms, molecules and solids. Adv. Math. **23**(1), 22–116 (1977)
18. Pisani, C.: Quantum-mechanical ab-initio calculation of the properties of crystalline materials. In: Pisani, C. (ed.) Quantum-Mechanical Ab-initio Calculation of the Properties of Crystalline Materials. Lecture Notes in Chemistry, vol. 67, Springer, Berlin (1996)
19. Radin, C.: Classical ground states in one dimension. J. Stat. Phys. **35**(1–2), 109–117 (1984)
20. Radin, C., Schulman, L.S.: Periodicity of classical ground states. Phys. Rev. Lett. **51**(8), 621–622 (1983)
21. Reed, M., Simon, B.: Methods of Modern Mathematical Physics. IV. Analysis of Operators. Academic, New York (1978)
22. Ruelle, D.: Statistical Mechanics. Rigorous Results. World Scientific, Singapore/Imperial College Press, London (1999)
23. Theil, F.: A proof of crystallization in two dimensions. Commun. Math. Phys. **262**(1), 209–236 (2006)

Matrix Functions: Computation

Nicholas J. Higham
School of Mathematics, The University of Manchester, Manchester, UK

Synonyms

Function of a matrix

Definition

A matrix function is a map from the set of complex $n \times n$ matrices to itself defined in terms of a given scalar function in one of various, equivalent ways. For example, if the scalar function has a power series expansion $f(x) = \sum_{i=1}^{\infty} a_i x^i$, then $f(A) = \sum_{i=1}^{\infty} a_i A^i$ for any $n \times n$ matrix A whose eigenvalues lie within the radius of convergence of the power series. Other definitions apply more generally without restrictions on the spectrum [6].

Description

Transformation Methods

Let A be an $n \times n$ matrix. A basic property of matrix functions is that $f(X^{-1}AX) = X^{-1}f(A)X$ for any nonsingular matrix X. Hence, if A is diagonalizable, so that $A = XDX^{-1}$ for some diagonal matrix $D = \mathrm{diag}(d_i)$ and nonsingular X, then $f(A) = Xf(D)X^{-1} = X\mathrm{diag}(f(d_i))X^{-1}$. The task of computing $f(A)$ is therefore trivial when A has a complete set of eigenvectors and the eigendecomposition is known. However, in general the diagonalizing matrix X can be arbitrarily ill conditioned and the evaluation in floating point arithmetic can therefore be inaccurate, so this approach is recommended only for matrices for which X can be assured to be well conditioned. For Hermitian, symmetric, or more generally normal matrices (those satisfying $AA^* = A^*A$), X can be taken unitary and evaluation by diagonalization is an excellent approach.

For general matrices, it is natural to restrict to unitary similarity transformations, in which case the Schur decomposition $A = QTQ^*$ can be exploited, where

Q is unitary and T is upper triangular. Now $f(A) = Qf(T)Q^*$ and the problem reduces to computing a function of a triangular matrix. In the 2×2 case there is an explicit formula:

$$f\left(\begin{bmatrix} \lambda_1 & t_{12} \\ 0 & \lambda_2 \end{bmatrix}\right) = \begin{bmatrix} f(\lambda_1) & t_{12} f[\lambda_2, \lambda_1] \\ 0 & f(\lambda_2) \end{bmatrix}, \qquad (1)$$

where $f[\lambda_2, \lambda_1]$ is a first-order divided difference and the notation reflects that $\lambda_i = t_{ii}$ is an eigenvalue of A. More generally, when the eigenvalues are distinct, $f(T)$ can be computed by an elegant recurrence due to Parlett [10]. This recurrence breaks down for repeated eigenvalues and can be inaccurate when two eigenvalues are close. These problems can be avoided by employing a block form of the recurrence, in which $T = (T_{ij})$ is partitioned into a block $m \times m$ matrix with square diagonal blocks T_{ii}. The *Schur–Parlett algorithm* of Davies and Higham [4] uses a unitary similarity to reorder the blocks of T so that no two distinct diagonal blocks have close eigenvalues while within every diagonal block the eigenvalues are close, then applies a block form of Parlett's recurrence. Some other method must be used to compute the diagonal blocks $f(T_{ii})$, such as a Taylor series taken about the mean of the eigenvalues of the block. The Schur–Parlett algorithm is the best general-purpose algorithm for evaluating matrix functions and is implemented in the MATLAB function `funm`.

For the square root function, $f(T)$ can be computed by a different approach: the equation $U^2 = T$ can be solved for the upper triangular matrix U by a recurrence of Björck and Hammarling [3] that runs to completion even if A has repeated eigenvalues. A generalization of this recurrence can be used to compute pth roots [11].

Approximation Methods

Another class of methods is based on approximations to the underlying scalar function. Suppose that for some rational function r, $r(A)$ approximates $f(A)$ well for A within some ball. Then we can consider transforming a general A to a matrix B lying in the ball, approximating $f(B) \approx r(B)$, then recovering an approximation to $f(A)$ from $r(B)$. The most important example of this approach is the *scaling and squaring method* for the matrix exponential, which approximates $e^A \approx r_m(A/2^s)^{2^s}$, where m and s are nonnegative integers and r_m is the $[m/m]$ Padé approximant to e^x. Backward error analysis can be used to determine a choice of the parameters s and m that achieves a given backward error (in exact arithmetic) at minimal computational cost [1, 7].

The analogue for the matrix logarithm is the *inverse scaling and squaring method*, which uses the approximation $\log(A) \approx 2^s r_m(A^{1/2^s} - I)$, where $r_m(x)$ is the $[m/m]$ Padé approximant to $\log(1 + x)$. Here, amongst the many logarithms of a matrix, log denotes the principal logarithm: the one whose eigenvalues have imaginary parts lying in $(-\pi, \pi)$; there is a unique such logarithm for any A having no eigenvalues on the closed negative real axis. Again, backward error analysis can be used to determine an optimal choice of the parameters s and m [2].

The derivation of (inverse) scaling and squaring algorithms requires attention to many details, such as how to evaluate a Padé approximant at a matrix argument, how to obtain the sharpest possible error bounds while using only norms, and how to avoid unnecessary loss of accuracy due to rounding errors.

Approximation methods can be effectively used in conjunction with a Schur decomposition, in which case the triangularity can be exploited [1, 2, 8].

Matrix Iterations

For functions that satisfy an algebraic equation, matrix iterations can be set up that, under appropriate conditions, converge to the matrix function. Many different derivations are possible, one of which is to apply Newton's method to the relevant equation. For example, for the equation $X^2 = A$, Newton's method can be put in the form

$$X_{k+1} = \frac{1}{2}(X_k + X_k^{-1} A), \qquad (2)$$

under the assumption that X_0 commutes with A. This iteration does not always converge. But if A has no eigenvalues on the closed negative real axis and we take $X_0 = A$, then X_k converges quadratically to $A^{1/2}$, the unique square root of A whose spectrum lies in the open right half-plane. Matrix iterations potentially suffer from two problems: they may be slow to converge initially, before the asymptotic fast convergence (in practice of quadratic or higher rate) sets in, and they may be unstable in finite precision arithmetic. Iteration (2) suffers from both these

problems. However, (2) is mathematically equivalent to the coupled iteration

$$X_{k+1} = \frac{1}{2}\left(X_k + Y_k^{-1}\right), \qquad X_0 = A,$$
$$Y_{k+1} = \frac{1}{2}\left(Y_k + X_k^{-1}\right), \qquad Y_0 = I, \tag{3}$$

of Denman and Beavers [5]: the X_k from (3) are identical to those from (2) with $X_0 = I$ and $Y_k \equiv A^{-1}X_k$. This iteration is numerically stable. Various other equivalent and practically useful forms of (2) are available [6, Chap. 6].

The convergence of matrix iterations in the early stages can be accelerated by including scaling parameters. Consider the Newton iteration

$$X_{k+1} = \frac{1}{2}(X_k + X_k^{-1}), \qquad X_0 = A. \tag{4}$$

Assuming that A has no pure imaginary eigenvalues, X_k converges quadratically to $\mathrm{sign}(A)$, which is the matrix function corresponding to the scalar sign function that maps points in the open right half-plane to 1 and points in the open left half-plane to -1. Although the iteration converges at a quadratic rate, convergence can be extremely slow initially. To accelerate the iteration we can introduce a positive scaling parameter μ_k:

$$X_{k+1} = \frac{1}{2}\left(\mu_k X_k + \mu_k^{-1} X_k^{-1}\right), \qquad X_0 = A.$$

Various choices of μ_k are available, with differing motivations. One is the determinantal scaling $\mu_k = |\det(X_k)|^{-1/n}$, which tries to bring the eigenvalues of μX_k close to the unit circle.

The number of iterations required for convergence to double precision accuracy (unit roundoff about 10^{-16}) varies with the iteration (and function) and the scaling but in some cases can be strictly bounded. For certain scaled iterations for computing the unitary polar factor of a matrix, it can be proved that less than ten iterations are needed for matrices with condition number less than 10^{16} (e.g., [9]). Moreover, for these iterations only one or two iterations might be needed if the starting matrix is nearly unitary.

References

1. Al-Mohy, A.H., Higham, N.J.: A new scaling and squaring algorithm for the matrix exponential. SIAM J. Matrix Anal. Appl. **31**(3), 970–989 (2009). doi:http://dx.doi.org/10.1137/09074721X
2. Al-Mohy A.H., Higham, N.J.: Improved inverse scaling and squaring algorithms for the matrix logarithm. SIAM J. Sci. Comput. **34**(4), C152–C169, (2012)
3. Björck, Å., Hammarling, S.: A Schur method for the square root of a matrix. Linear Algebra Appl. **52/53**, 127–140 (1983)
4. Davies, P.I., Higham, N.J.: A Schur–Parlett algorithm for computing matrix functions. SIAM J. Matrix Anal. Appl. **25**(2), 464–485 (2003). doi:http://dx.doi.org/10.1137/S0895479802410815
5. Denman, E.D., Beavers, A.N., Jr.: The matrix sign function and computations in systems. Appl. Math. Comput. **2**, 63–94 (1976)
6. Higham, N.J.: Functions of Matrices: Theory and Computation. Society for Industrial and Applied Mathematics, Philadelphia (2008)
7. Higham, N.J.: The scaling and squaring method for the matrix exponential revisited. SIAM Rev. **51**(4), 747–764 (2009). doi:http://dx.doi.org/10.1137/090768539
8. Higham, N.J., Lin, L.: A Schur–Padé algorithm for fractional powers of a matrix. SIAM J. Matrix Anal. Appl. **32**(3), 1056–1078 (2011). doi:http://dx.doi.org/10.1137/10081232X
9. Nakatsukasa, Y., Bai, Z., Gygi, F.: Optimizing Halley's iteration for computing the matrix polar decomposition. SIAM J. Matrix Anal. Appl. **31**(5), 2700–2720 (2010)
10. Parlett, B.N.: A recurrence among the elements of functions of triangular matrices. Linear Algebra Appl. **14**, 117–121 (1976)
11. Smith, M.I.: A Schur algorithm for computing matrix pth roots. SIAM J. Matrix Anal. Appl. **24**(4), 971–989 (2003)

M

Mechanical Systems

Bernd Simeon
Department of Mathematics,
Felix-Klein-Zentrum, TU Kaiserslautern,
Kaiserslautern, Germany

Synonyms

Constrained mechanical system; Differential-algebraic equation (DAE); Euler–Lagrange equations; Multibody system (MBS); Time integration methods