

In contrast to the abundance of monographs on the numerical treatment of linear systems and eigenvalue problems, there has not previously been a book devoted to matrix functions. For that reason alone, the author should be thanked for his immense efforts that must have gone into the production of the research monograph under review. But this would not be a book by Higham if there was not much more to be grateful for.

Chapter 1 sets the stage by defining the function of a matrix. The following noteworthy convention, attributed to Gantmacher [MR0065520. Gantmaher, F. R. Teoriya matric. Gosudarstv. Izdat. Tehn.-Teor. Lit., Moscow, 1953. 491 pp.], is used. Given a square matrix A , let n_i denote the size of the largest Jordan block belonging to an eigenvalue λ_i of A . A function $f : \Omega \subseteq \mathbb{C} \rightarrow \mathbb{C}$ is said to be *defined on the spectrum of A* , if at each eigenvalue λ_i the value of f and all its derivatives up to order $n_i - 1$ exist. This allows to define the Hermite polynomial p interpolating these values and derivatives of f . The function f of A is then simply set to $f(A) := p(A)$, where $p(A)$ is evaluated by substituting A for the scalar variable. Apart from this definition, two equivalent characterizations of $f(A)$ are given: the classical one based on the Jordan canonical form of A and the Cauchy integral formula, additionally assuming f to be analytic in the vicinity of the eigenvalues. Some basic properties of matrix functions are summarized, including several less known and even new results. One of the gems is Theorem 1.35, a natural generalization of the Sherman-Morrison-Woodbury formula:

$$f(\alpha I + AB) = f(\alpha)I + A(BA)^{-1}(f(\alpha I + BA) - f(\alpha)I)B,$$

where α is a scalar and I denotes the identity matrix of appropriate size.

The square root and logarithm are multi-valued functions; choosing a different branch for each Jordan block belonging to an eigenvalue gives rise to so called non-primary matrix functions. Not fitting into any of the definitions of $f(A)$ mentioned above, they can formally be defined as solutions to the nonlinear matrix equations $X^2 = A$ and $e^X = A$, respectively. In Sections 1.5 and 1.6, the classification of all possible (real) solutions to these equations are given. A brief history of matrix functions concludes the chapter.

Chapter 2 describes many examples of applications where functions of matrices play an important role. Applications described in some detail include exponential integrators for differential equations, generators of Markov models, (non)linear equations in control theory, divide-and-conquer algorithms for nonsymmetric eigenvalue problems, orthogonalization, and general Riccati equations. Even though most of the applications are only briefly touched, the extensive list of references offers good starting points for learning more. If one additional application was to be included in Chapter 2, it would probably be density functional theory.

Chapter 3 is concerned with the sensitivity of $f(A)$ under perturbations of A . For this purpose, the Fréchet and Gâteaux derivatives of a matrix function are introduced. The relation

$$f\left(\begin{bmatrix} A & E \\ 0 & A \end{bmatrix}\right) = \begin{bmatrix} A & \left.\frac{d}{dt}\right|_{t=0} f(A + tE) \\ 0 & A \end{bmatrix}.$$

offers a general approach to compute these derivatives for any sufficiently smooth function f . (For a particular function f , this might not be the most efficient way.) The induced norm of the Fréchet derivative of f at A defines the absolute condition number of $f(A)$. Bounds for this number are obtained from a novel result in Theorem 3.9, which states that the eigenvalues of the Fréchet derivative are given as the divided differences $f[\lambda_i, \lambda_j]$ for all possible pairs of eigenvalues λ_i, λ_j of A . Alternatively, to compute reliable estimates for the condition number, several variants of the power method are proposed.

Chapter 4 provides an arsenal of techniques useful for tackling arbitrary functions. First, the efficient evaluation of a matrix polynomial of degree m is discussed. While a straightforward extension of Horner's method requires $m - 1$ matrix-matrix multiplications, it is shown that a method of Paterson and Stockmeyer requires only about $2\sqrt{m}$ matrix-matrix multiplications. In Section 4.3, a convergence result when inserting A into the Taylor series of f is given along with a general bound on the truncation error. This is followed by a brief introduction to rational, in particular Padé, approximation. Section 4.5 is concerned with methods based on diagonalizing A . The obvious disadvantages of diagonalization for strongly non-normal matrices are pointed out and block diagonalization is proposed as an alternative. Closely related to block diagonalization is the block Parlett-Schur method, which first reduces A to a Schur form with clusters of nearby eigenvalues collected in diagonal blocks and then applies a block recurrence to compute the function of this block triangular matrix. The details of this method are described later on, in Chapter 9. Section 4.9, by far the longest in Chapter 4, provides general results on one-step iterations having the form

$$X_{k+1} = g(X_k) \tag{1}$$

with some (preferably simpler) function g . After a discussion of suitable termination criteria, the novel Theorem 4.15 contains a general convergence result for this iteration. Essentially it states that (1) converges to $f(A)$ if the corresponding scalar iteration converges for every eigenvalue of A to an *attracting* fixed point. The general discussion on the stability of (1) is another important novel contribution of this section. A fixed point of (1) is defined to be stable if the Fréchet derivative of g at this fixed point has bounded powers. In this context, a significant difference between iterations for scalars and matrices is observed. While a superlinearly convergent scalar iteration is always stable, the same is not necessarily true for matrix iterations. However, it is shown that for an idempotent function f , superlinear convergence

does imply stability. After a section on preprocessing techniques (shifting and scaling) to decrease the norm of A before computing $f(A)$, Section 4.11 provides bounds on $\|f(A)\|$ based on the Jordan canonical form, the pseudospectrum or the Schur decomposition of A .

Chapter 5 is devoted to the sign function, which maps a complex number to the sign of its real part. The matrix sign function $\text{sign}(A)$ for a matrix A with no eigenvalues on the imaginary axis is less trivial to compute and has a range of important properties. For example, $(I - \text{sign}(A))/2$ is the spectral projector belonging to all eigenvalues of A in the left half-plane. Based on the results from Chapter 3, expressions and useful bounds for the condition number of the matrix sign function are given. Also, a specialization of the Parlett-Schur method is briefly described. It is worth to mention, however, that the main purpose of the matrix sign function is often the avoidance of having to compute a Schur decomposition. Discussed in Section 5.3, the Newton method $X_{k+1} = (X_k + X_k^{-1})/2$ with $X_0 = A$ is probably the most popular method for computing $\text{sign}(A)$, as it always converges quadratically and yet it requires only matrix inversion and addition. In contrast, the seemingly more attractive Newton-Schultz iteration $X_{k+1} = X_k(3I - X_k^2)/2$, derived from the Newton method by applying one step of Newton's method for the matrix inverse, is only locally convergent. Both iterations belong to a whole family of iterations derived in Section 5.4 from Padé approximations to $(1 - \xi)^{-1/2}$. Theorem 5.8, a result by Kenney and Laub, shows that only the diagonal (m, m) and superdiagonal $(m - 1, m)$ entries of the corresponding Padé table yield global convergence. A detailed discussion is provided on the somewhat technical but practically important aspects of scaling strategies to accelerate initial convergence and termination criteria for the Newton iteration. The chapter is concluded by Zolotarev's result on the best rational approximation to the sign function on symmetric intervals.

Chapter 6 is about matrix square roots. A section on the sensitivity of the matrix square root is followed by a specialization of the Schur method, which is shown to have essentially optimal stability. Without any modification, Newton's method for computing the principal matrix square root of a matrix A requires the costly solution of a matrix Sylvester equation in each iteration. This can be avoided by exploiting the fact that the iterates commute with A , leading to the iteration $X_{k+1} = (X_k + X_k^{-1}A)/2$, $X_0 = A$, which – at least in theory – produces the same iterates as the unmodified Newton method. In finite-precision arithmetic, however, commutativity is getting gradually lost, leading to numerical instabilities unless all eigenvalues of A are very closely clustered. Fortunately, coupled iterations provide viable alternatives at a somewhat increased computational cost. A consequence of the general results from Section 4.9, several of these coupled iterations are found to be unconditionally stable and to have optimal limiting accuracy. After a discussion of suitable scaling strategies, Section 6.7

elaborates on the identity

$$\text{sign} \left(\begin{bmatrix} 0 & A \\ I & 0 \end{bmatrix} \right) = \begin{bmatrix} 0 & A^{1/2} \\ A^{-1/2} & 0 \end{bmatrix}, \quad (2)$$

implying that iterations for the matrix sign function can be extended to the matrix square root. In Section 6.8, methods for several special classes of matrices are discussed, such as the binomial iteration that converges slowly but monotonically to the square root of $I - C$, where C is a nonnegative matrix. The chapter is concluded by brief treatments of the more specialized questions of computing small-normed square roots and of finding integer roots of the identity matrix.

Chapter 7 deals with a natural but less frequent generalization of matrix square roots: the p th root of a matrix A . After an extension of the Schur method, the (modified) Newton method, and the coupled iterations from Chapter 6, a hybrid algorithm is proposed. It uses the Schur decomposition of A for finding a choice of parameters that yields quick initial convergence in a certain coupled iteration due to Lakić. Provided that p is even and not a multiple of 4, the relation (2) can be generalized in the sense that the p th root of A can be determined from the sign function of a $pn \times pn$ block cyclic matrix. This allows the application of, e.g., the sign function iteration.

Chapter 8 is concerned with the polar decomposition $A = UH$, where U has orthonormal columns and H is Hermitian positive semi-definite. Note that A may be rectangular, and – unless A has full column rank – U is not uniquely determined. It is shown that this nonuniqueness can be eradicated by allowing U to be only a partial isometry but requiring the ranges of U^* and H to be identical. Section 8.1 gives an overview of approximation problems that can be addressed with polar decompositions, including the classical result of Fan and Hoffmann, which states that there is no matrix with orthonormal columns closer in any unitarily invariant norm to A than U . This is followed by results on the sensitivities of U and H with respect to perturbations in A . If A is square and nonsingular then (2) has a polar brother,

$$\text{sign} \left(\begin{bmatrix} 0 & A \\ A^* & 0 \end{bmatrix} \right) = \begin{bmatrix} 0 & U \\ U^* & 0 \end{bmatrix}.$$

Not surprisingly, this relation can be used to translate the Newton iteration and other Padé-type iterations from Section 5.8 into iterations for computing the polar factor U . Once again, a comprehensive discussion of numerical stability, scaling strategies and termination criteria is given.

Chapter 9 provides the details of the block Schur-Parlett algorithm already introduced in Chapter 4, such as heuristic blocking/clustering strategies and the treatment of diagonal blocks by means of truncated Taylor series.

Chapter 10 is devoted to the most prominent and most studied matrix function, the matrix exponential e^A . A section on basic properties is followed

by results on the Fréchet derivative and the conditioning of e^A . One of the best methods for computing e^A , scaling+squaring is detailed in Section 10.3. The basic idea is as follows. Choose an integer s such that the norm of $A/2^s$ is of order 1, use a diagonal Padé approximation to evaluate the exponential of the scaled matrix, and square the result s times to undo the scaling. A thorough investigation of the approximation error leads to the conclusion that a Padé approximant of order at most 13 is sufficient to attain double precision if $\|A/2^s\| \leq 5.4$. Numerical experiments reported in Section 10.5 reveal no clear winner between scaling+squaring and the Schur-Parlett algorithm with respect finite-precision accuracy. On average, however, the former seems to perform better. Section 10.6 suggests two methods for evaluating the Fréchet derivative $L(A, E)$ of e^A : one based on the numerical quadrature of an integral representation for $L(A, E)$ and one based on a Kronecker product representation requiring the solution of a sequence of matrix Sylvester equations. Both methods can be plugged into a power method to estimate the norm of $L(A, \cdot)$ and hence the conditioning of e^A .

Chapter 11 shows how to extend the ideas of scaling+squaring to the matrix logarithm. After repeatedly taking the square root of A , a diagonal Padé approximation is used to approximate the logarithm of the modified matrix. The “square rooting” is undone by an appropriate scaling. To reduce the costs for computing square roots, it is proposed to replace A by its Schur form. The chapter is concluded by an extension of Section 10.6 to the evaluation of the Fréchet derivative for the matrix logarithm.

Chapter 12 is concerned with the matrix cosine and sine. Besides a discussion of basic properties and conditioning, an analogue of scaling+squaring is developed. Choose an integer s such that the norm of $A/2^s$ is of order one and use a diagonal Padé approximation of even degree to evaluate $C_0 = \cos(A/2^s)$. Then $\cos(A) = C_s$ can be obtained after s steps of the recurrence $C_{i+1} = 2C_i^2 - I$. It is noted that the sine of a matrix A can be evaluated using $\sin(A) = \cos(A - \pi/2 I)$. Furthermore, an efficient algorithm for simultaneously computing both $\sin(A)$ and $\cos(A)$ is given.

Chapter 13 offers a brief discussion on Krylov subspace methods for approximating the function of a matrix times a vector, $f(A)b$. Furthermore, it is shown how an integral representation for $f(A)$ can be turned – by means of numerical quadrature – into an efficient method for approximating $f(A)b$.

Chapter 14 represents a collection of miscellaneous topics, including structure-preserving algorithms for computing the matrix function of a structured matrix.

The book has a number of appendices: An appendix on basic background material in linear algebra and numerical analysis, an appendix on the number of elementary operations needed by typical operations occurring in algorithms for matrix functions, and an appendix on the matrix function toolbox, a collection of educational MATLAB programs implementing many

of the algorithms described in the book.

The book is concluded by 34 pages of solution sketches to almost all problems posed at the end of each chapter, a rich bibliography of 625 references, and a comprehensive index.

There is not much to be said about the quality of the book under review. It is without any doubt one of the most carefully prepared and most monumental contemporary research monographs in numerical analysis. Researcher getting in touch with matrix functions have to get in touch with this book.