View metadata, citation and similar papers at core.ac.uk
brought to you by CORE
provided by Elsevier - Publisher Connector

Journal of Computational and Applied Mathematics 46 (1993) 405–413
North-Holland

CAM 1229

Computing the exponential of an intensity matrix

Brian J. Melloy
Department of Industrial Engineering, Clemson University, Clemson, SC, United States

G. Kemble Bennett
Department of Industrial Engineering, Texas A&M University, College Station, TX, United States

Received 12 August 1991

Abstract

Melloy, B.J. and G.K. Bennett, Computing the exponential of an intensity matrix, Journal of Computational and Applied Mathematics 46 (1993) 405–413.

A scaling and squaring procedure for computing the exponential of an intensity matrix is developed in this paper. Intensity matrices occur naturally in inventory, reliability and queueing systems. It will be seen that when these matrices are properly transformed, they are well-conditioned for the matrix exponential function, and particularly well-suited for the scaling and squaring approach. As a result, many of the reliability problems associated with the standard procedure have been circumvented.

Keywords: Matrix exponential; intensity matrix; scaling and squaring.

1. Introduction

A discrete-state continuous-time stationary Markov process may be represented as a system of linear, first-order differential equations. For time-homogeneous processes, the probability state vector may be found directly as

$$\pi(t) = e^{Qt}\pi(0), \tag{1}$$

where t, $t \geqslant 0$, is a time scalar, $\pi(0)$ is the initial-state probability vector, e^{Qt} is the transition matrix, and Q is the intensity matrix.

The intensity matrix Q has a very specialized structure which is characterized by [4]

$$q_{ij} \geqslant 0, \tag{2}$$

Correspondence to: Prof. B.J. Melloy, Department of Industrial Engineering, Clemson University, Clemson, SC 29634-0920, United States.

0377-0427/93/$06.00 © 1993 – Elsevier Science Publishers B.V. All rights reserved

for all i and j, $i \neq j$, and

$$\sum_{j=1}^{n} q_{ij} = 0, \tag{3}$$

for all i. These matrices arise naturally, for example, in reliability, inventory and queueing systems [3]. Moreover, the exponential of an intensity matrix is (row) stochastic [4].

A methodology for computing the matrix exponential will be developed in Section 2 which exploits both the special structure of intensity and transition matrices. Upper and lower error bounds will be constructed for this procedure in Section 3. In Section 4, the reliability of the algorithm will be examined. In particular, it will be seen (in Section 5) that when intensity matrices are properly preconditioned, they exhibit many of the same desirable properties as do normal matrices for the matrix exponential function. Finally, concluding remarks will follow in Section 6.

2. Algorithm development

While there are many methods available for computing the matrix exponential, the scaling and squaring technique is among the most highly regarded [6]. Scaling and squaring methods employ the identity

$$e^{Qt} = \left[e^{2^{-m}Qt}\right]^{2^m}, \tag{4}$$

where m is a positive integer. Squaring the matrix is both more efficient and less prone to round-off error than performing 2^{m-1} successive matrix multiplications. The motivation for scaling the matrix is to reduce the round-off error difficulties and operation count, which are proportional to $\|2^{-m}Qt\|$. The matrix exponential may then be computed satisfactorily using either diagonal Padé or Taylor approximants [6].

Algorithms based on the scaling and squaring algorithm have been developed for both Taylor and diagonal Padé approximants in [8,12], respectively. These two algorithms are fundamentally identical, as the former was adapted from the latter. Prior to scaling the intensity matrix, the algorithm begins with a series of two inexpensive preconditioning steps, translation and balancing, which serve to reduce the norm of the matrix.

First, the intensity matrix is translated by shifting the origin in the manner indicated below:

$$Q' = Q + \omega I. \tag{5}$$

The objective is to choose the shift operator ω in such a way that

$$\|Q'\| < \|Q\|, \tag{6}$$

where the norm will be defined as

$$\|Q\| = \max_{1 \leq i \leq n} \sum_{j=1}^{n} |q_{ij}|. \tag{7}$$

Clearly, this operation is less expensive (and more stable [10]) than scaling/squaring steps. (The norm that was used by the aforementioned authors was actually L_1, rather than L_∞.

Nevertheless, if the alternate column stochastic convention was employed here (e.g., [5,9]), the identical results would be obtained for L_1.)

The value of the shift operator which minimizes the norm of an arbitrary matrix over all possible translations is difficult to obtain in practice, and consequently it is normally approximated [12]. However, if the original matrix is an intensity matrix, no such approximation is necessary; the exact shift operator can be shown to be the absolute value of the most negative diagonal element. This is demonstrated in the following theorem.

Theorem 1. *Let $Q' = Q + \omega I$, where Q is an intensity matrix. Then the value of ω which minimizes $\| Q' \|$ is*

$$\omega = q = \max_{1 \leqslant i \leqslant n} |q_{ii}|.$$

Proof.

$$\min_{\omega} \| Q' \| = \min_{\omega} \left[\max_{1 \leqslant i \leqslant n} \left[\sum_{j=1}^{n} |q'_{ij}| \right] \right] = \min_{\omega} \left[\max_{1 \leqslant i \leqslant n} \left[|q_{ii} + \omega| + \sum_{\substack{j=1 \\ j \neq i}}^{n} |q_{ij}| \right] \right]$$

$$= \min_{\omega} \left[\max_{1 \leqslant i \leqslant n} \left[|q_{ii} + \omega| - q_{ii} \right] \right].$$

Let $q = \max_{1 \leqslant i \leqslant n} |q_{ii}|$. Then, two cases must be distinguished: $\omega < q$ and $\omega \geqslant q$. First suppose that $\omega < q$ and let i denote the row where q occurs. Then

$$\sum_{j=1}^{n} |q'_{ij}| = |q + \omega| + q.$$

Since $|q + \omega| > 0$ for i, it follows that

$$\sum_{j=1}^{n} |q'_{ij}| > q.$$

and hence for this case $\| Q' \|$ is strictly greater than q.

Next consider the case when $\omega \geqslant q$. Then

$$\min_{\omega} \| Q' \| = \min_{\omega} \left[\max_{1 \leqslant i \leqslant n} \left[q_{ii} + \omega - q_{ii} \right] \right] = \omega.$$

Therefore $\| Q' \|$ is minimized when $\omega = q$. □

Finally, note that the translation of Theorem 1 yields a matrix with a nonnegative diagonal and identical row norms. Therefore, Q' is a positive-scalar multiple of a stochastic matrix. Moreover, since

$$\| Q \| = \max_{1 \leqslant i \leqslant n} \left[|q_{ii}| + \sum_{\substack{j=1 \\ j \neq i}}^{n} q_{ij} \right] = 2q, \tag{8}$$

the optimal translation of the matrix has reduced the norm by exactly one-half.

The second preconditioning step, that of balancing the translated matrix, attempts to minimize the norm over all possible diagonal similarity transformations. That is

$$\min_{D \in \Psi} \| D^{-1} Q' D \|, \tag{9}$$

where Ψ is the set of all $n \times n$ nonsingular diagonal matrices. (The set Ψ is restricted to integer powers of the machine base in order that rounding errors are not introduced.) This preconditioning step is also superior to scaling/squaring operations with respect to stability and efficiency. It will now be demonstrated, however, that this step may be omitted as the translated intensity matrix is already balanced.

Theorem 2. *Let $Q' = Q + qI$, where Q is an intensity matrix. If $q = \max_{1 \leqslant i \leqslant n} | q_{ii} |$, then*

$$\min_{D \in \Psi} \| D^{-1} Q' D \| = \| Q' \|.$$

Proof. Let $Q'' = D^{-1} Q' D$. The ith row norm of the balanced matrix Q'' may then be expressed as

$$\sum_{j=1}^{n} q''_{ij} = \sum_{j=1}^{n} \left| \frac{d_j}{d_i} \right| q'_{ij},$$

for all i. Next define

$$d_i = \min_{1 \leqslant i \leqslant n} | d_i |,$$

with i denoting the row where this entry occurs. Now from Theorem 1 it is known that

$$\sum_{j=1}^{n} q'_{ij} = q,$$

for all i. Hence for column i of the balanced matrix Q'' it follows that

$$\sum_{j=1}^{n} \left| \frac{d_j}{d_i} \right| q'_{ij} \geqslant q,$$

since

$$\left| \frac{d_j}{d_i} \right| \geqslant 1,$$

for all j. Therefore the norm of the balanced matrix Q'' will be greater than the norm of the original matrix Q', unless

$$d_j = d_i,$$

for all j. □

Thus there is no diagonal similarity transformation which will reduce the norm of the translated intensity matrix.

Now, as a result of (4), (5) and Theorem 1, the transition matrix is expressed as

$$e^{Qt} = e^{-qt} \left[e^{2^{-m} Q' t} \right]^{2^m}. \tag{10}$$

The positive integer m is chosen such that

$$2^{m-1} \leqslant \| Q't \| < 2^m, \tag{11}$$

which ensures that the norm of the scaled matrix is bounded by unity. (Observe, therefore, that as a result of (8), the number of scaling/squaring steps required has been reduced by one.) This serves both to reduce the round-off error and to increase the accuracy of the approximants.

Alternately, the matrix exponential may be represented as

$$e^{Qt} = \left[e^{-2^{-m}qt} e^{2^{-m}Q't}\right]^{2^m}. \tag{12}$$

This modification will preclude the possibility of a machine overflow error, since

$$\left\| \left[e^{-2^{-m}qt} e^{2^{-m}Q't}\right]^{2^i} \right\| = \left\| \left[e^{2^{-m}Qt}\right]^{2^i} \right\| = 1, \tag{13}$$

for $i = 1, 2, \ldots, m$. Equation (13) follows from the fact that integer powers of a stochastic matrix are also stochastic [4].

Subsequently, the scaled matrix exponential is approximated using Padé rational functions:

$$e^{sQ'} \approx \left[D_{kl}(-sQ')\right]^{-1} N_{kl}(sQ'), \tag{14}$$

where

$$D_{kl}(-sQ') = \sum_{i=0}^{l} \frac{(k+l-i)!\, l!}{(k+l)!\, i!\, (l-i)!} (-sQ')^i, \tag{15}$$

$$N_{kl}(sQ') = \sum_{i=0}^{k} \frac{(k+l-i)!\, k!}{(k+l)!\, i!\, (k-i)!} (sQ')^i, \tag{16}$$

and, for convenience, $s = 2^{-m}t$. A Taylor approximant is obtained when l is set equal to zero. Diagonal Padé approximants are those which have the same degree polynomial in the numerator and denominator, i.e., when $l = k$. In both cases, the value of k is selected such that a predetermined truncation error criterion is satisfied. The diagonal form of the Padé approximants is generally preferred because it has the smallest truncation error of those Padé approximants requiring $k - l$ matrix multiplications [12].

3. Error bounds

Error bounds are commonly constructed in order to provide a measure of accuracy for numerical procedures. Three bounds will be described in this section: a lower bound, a recursive bound and an upper error bound. First, it is not possible to formulate an expression for a lower bound ordinarily, because the final form of the resultant matrix is unknown. However, here, due to the well-known structure of the (stochastic) transition matrix, the derivation of this bound is straightforward:

$$\epsilon_l = \frac{\| e^{Qt} - e_c^{Qt} \|}{\| e^{Qt} \|} = \max_{1 \leqslant i \leqslant n} \left[\sum_{j=1}^{n} \left| \left[e^{Qt}\right]_{ij} - \left[e_c^{Qt}\right]_{ij} \right| \right] \geqslant \max_{1 \leqslant i \leqslant n} \left| \sum_{j=1}^{n} \left| \left[e^{Qt}\right]_{ij} \right| - \sum_{j=1}^{n} \left| \left[e_c^{Qt}\right]_{ij} \right| \right|$$

$$= \max_{1 \leqslant i \leqslant n} \left| 1 - \sum_{j=1}^{n} \left| \left[e_c^{Qt}\right]_{ij} \right| \right|, \tag{17}$$

where the subscript c denotes a computed quantity. Observe that the absolute and relative error are equivalent, which is also due to the fact that the transition matrix is stochastic.

Second, in order to detect algorithm failure, it is necessary to monitor the error during the course of the procedure. This is accomplished by recursively calculating the error at each stage of computation. The recursive equation, originally developed in [12], is modified using (10), (12) and (13) to obtain

$$\|\xi_i\| \leqslant 2\|\xi_{i-1}\| + \|\xi_{i-1}\|^2 + n\alpha^{-\beta}\left\|\left[\left[\mathrm{e}^{-sq}\mathrm{e}^{sQ'}\right]^{2^{i-1}}\right]_{\mathrm{c}}\right\|^2, \tag{18}$$

for $i = 1, 2, \ldots, m$, where β denotes the number of digits used to represent the mantissa of a floating-point number, with

$$\|\xi_0\| \leqslant \mathrm{e}^{-sq}[\|T_k(sQ')\| + \|R_k(sQ')\|], \tag{19}$$

where $T_k(sQ')$ and $R_k(sQ')$ denote the truncation and round-off error, respectively. (Of course, the equations for these errors are specific to both the Taylor and Padé approximants; the former are detailed in [5,8], and the latter in [12].) The purpose of this bound is to indicate the minimum number of accurate digits in the norm of the computed matrix exponential; a prohibitively large value implies algorithm failure. Lastly, the upper error bound is simply

$$\epsilon_u = \xi_m, \tag{20}$$

due to (18).

4. Algorithm reliability

Recall that the computational error will be monitored using the recursive formula of (18). A large error estimate, which will render the results unusable, may be the result of any of three causes: (i) the problem is inherently sensitive; (ii) the algorithm is unstable; or (iii) the error bound has been severely overestimated [6]. Thus, although the error bound precludes the possibility of accepting erroneous results, it is not entirely satisfactory since (in general) it cannot discriminate between the above three situations. These issues will now be re-examined with regard to the specific matrix exponential problem described herein.

Van Loan [11] has developed an exponential condition number which serves to measure the sensitivity of the map $A \to \mathrm{e}^{At}$, where A denotes an arbitrary matrix. Naturally, an ill-conditioned problem may result in algorithm failure. However, the condition number $\nu(Q', t)$ is as small as it can be for the translated intensity matrix Q'; that is,

$$\begin{aligned}\nu(Q', t) &\leqslant \int_0^t \|\mathrm{e}^{Q'(t-x)}\| \cdot \|\mathrm{e}^{Q'x}\| \,\mathrm{d}x \frac{\|Q'\|}{\|\mathrm{e}^{Q't}\|} \\ &= \int_0^t \mathrm{e}^{\|Q'\|(t-x)}\mathrm{e}^{\|Q'\|x}\,\mathrm{d}x \frac{\|Q'\|}{\mathrm{e}^{\|Q'\|t}} = \|Q'\| \int_0^t \mathrm{d}x = \|Q'\|\, t,\end{aligned} \tag{21}$$

due to Lemma A.1 (see the Appendix). Therefore, the translation step yields a matrix which is well-conditioned for the matrix exponential function.

Algorithm instability may be the result of either of two causes; the "hump" phenomenon, or catastrophic subtractive cancellation [6]. The algorithm becomes unstable when the elements of

the matrix exponential grow before they decay, which gives rise to a "hump" in the norm. The difficulty occurs when $2^{-m}t$ is under the "hump" and t is beyond it, for then

$$\|e^{2^{-m}At}\|^{2^m} \gg \|e^{At}\| \tag{22}$$

[6]. Since the rounding error ϵ is of order [1]

$$\epsilon \approx \boldsymbol{u}\,\|e^{2^{-m}At}\|^{2^m}, \tag{23}$$

where $\boldsymbol{u}$ represents the unit round-off, the round-off error is usually small relative to $\|e^{2^{-m}At}\|^{2^m}$ rather than to $\|e^{At}\|$. However, for this particular problem,

$$\epsilon \approx \boldsymbol{u}\,\|e^{2^{-m}Qt}\|^{2^m} = \boldsymbol{u}\,\|e^{Qt}\| = \boldsymbol{u} \tag{24}$$

(since $e^{2^{-m}Qt}$ is obviously also a stochastic matrix), which precludes the possibility of a hump.

Subtractive cancellation, on the other hand, typically occurs when relatively large values of the same order of magnitude are subtracted. In this event, absolute errors may be created which are larger than the final result, causing irreversible contamination. It has been conjectured that this catastrophic cancellation may only occur in the presence of a large "hump", but the relationship of round-off error to the sensitivity of the matrix exponential function remains an open issue [6].

The diagonal Padé approximants may be susceptible to subtractive cancellation because the Taylor expansion of the diagonal Padé approximant, beyond $2n$, has negative coefficients. In addition, $D_{kk}(-sQ')$ may be ill-conditioned with respect to inversion [6]. In contrast, the Taylor approximant $N_{k0}(sQ')$ is a function of matrices which are strictly nonnegative [2], and requires no matrix inversion. Thus although the Padé approximant is more efficient, the Taylor approximant is arguably more stable.

Ideally, a large error estimate would occur only when the algorithm has failed. Unfortunately, this is not the case; such an indication may be spurious. Statistical bounds have been developed to address this problem [13], but they are not completely satisfactory. However, for this particular problem, in the absence of numerical difficulties, the actual computational error will be approximately equal to

$$\|\epsilon_i\| \approx 2\,\|\epsilon_{i-1}\| + \|\epsilon_{i-1}\|^2 + n\alpha^{-\beta}, \tag{25}$$

for $i = 1, 2, \ldots, m$, due to (13) and (18). Observe that this equation reflects only the truncation and round-off error, which precludes the possibility of a "false alarm". Thus, the upper bound of (20) will be extremely sharp. As a result, when this upper bound is used in conjunction with the lower bound of (17), the characteristic uncertainty surrounding the error estimate is minimized.

Finally, it is noteworthy that the magnitude of the error realized when employing the Taylor polynomial is (normally) roughly equal to that of the *lower* bound (17). This is the case because, typically, $[e_c^{Qt}]_{ij} \in [0, 1]$, for all i and j, due to the fact that the Taylor polynomial is a monotonically increasing function (assuming negligible round-off error). In other words,

$$\sum_{l=1}^{n} |[\epsilon_l]_i| \approx \sum_{l=1}^{n} [\epsilon_l]_i, \tag{26}$$

due to the tendency of round-off error to cancel.

5. Discussion

It is noteworthy that Q', which is a positive-scalar multiple of a stochastic matrix, displays many characteristics with respect to L_∞ that normal matrices exhibit with respect to L_2 for the matrix exponential problem. Specifically, (i) the matrices are inherently balanced (Theorem 2 and [7]), (ii) the "hump" does not exist ((24) and [6]), and (iii) the exponential condition number is as small as it can be ((21) and [11]). Thus it may be conjectured that these matrices are "L_∞-normal". However, since the condition number for the inversion of a positive-scalar multiple of a stochastic matrix is not equal to 1, this supposition is not correct.

6. Conclusion

In this paper, an algorithm was developed expressly for the purpose of computing the exponential of an intensity matrix. The procedure is based on the scaling and squaring procedure, which is one of the most effective procedures known for computing the matrix exponential. First, the intensity matrix is preconditioned in a single step which yields both the optimal translation (with respect to L_∞) and a balanced matrix. The scaled matrix exponential is then computed using either Taylor [8] or diagonal Padé [12] functions. Once the scaled approximant is obtained, it is normalized, which precludes the possibility of a program-terminating overflow error during the squaring process. Lastly, the normalized scaled approximant is squared, which yields the transition matrix.

Computing the matrix exponential in this manner resulted in an extremely reliable procedure. First, the translated matrix was shown to be ideally conditioned for the matrix exponential problem. Moreover, since the translated matrix is nonnegative, employing the Taylor polynomial would also circumvent any possibility of catastrophic cancellation. Further, no hump exists, due to the inherent structure of the transition matrix. Finally, an error interval was developed (with a sharp upper bound) which minimizes the characteristic uncertainty surrounding the error estimate (thereby avoiding the need for a supplemental statistical bound).

Appendix

Lemma A.1. *Let $U = cP$, where P is a stochastic matrix and c is a nonnegative real scalar. Then*

$$\|e^U\| = e^{\|U\|}.$$

Proof. It is not difficult to show that the Cauchy–Schwarz and triangle inequalities are strict equalities for positive-scalar multiples of row stochastic matrices in L_∞. Therefore it follows directly that

$$\|e^U\| = \left\| \sum_{i=0}^{\infty} \frac{1}{i!} U^i \right\| = \sum_{i=0}^{\infty} \frac{1}{i!} \|U^i\| = \sum_{i=0}^{\infty} \frac{1}{i!} \|U\|^i = e^{\|U\|}. \qquad \square$$

Acknowledgement

The authors wish to acknowledge the contributions of Dr. Thomas E. Price, Professor of Mathematical Sciences at the University of Akron, Ohio.

References

[1] G.H. Golub and C.F. Van Loan, *Matrix Computations* (Johns Hopkins Univ. Press, Baltimore, MD, 1983).
[2] W.K. Grassman, Transient solutions in Markovian queueing systems, *Comput. Oper. Res.* **4** (1977) 47–53.
[3] D. Gross, D.R. Miller and R.M. Soland, On some common interests among reliability, inventory and queueing, *IEEE Trans. Reliability* **R-34** (3) (1985) 204–208.
[4] D.P. Heyman and M.J. Sobel, *Stochastic Models in Operations Research, Vol. I* (McGraw-Hill, New York, 1982).
[5] B.J. Melloy and G.K. Bennett, The exact truncation error norm for the transition matrix of a continuous-time Markov chain, Dept. Indust. Management Systems Engrg. Working Paper Ser., Pennsylvania State Univ., 1991.
[6] C.B. Moler and C.F. Van Loan, Nineteen dubious ways to compute the exponential of a matrix, *SIAM Rev.* **20** (4) (1978) 801–836.
[7] B.N. Parlett and C. Reinsch, Balancing a matrix for calculation of eigenvalues and eigenvectors, *Numer. Math.* **13** (1969) 293–304.
[8] R.K. Saha and J.R. Boriotti, Truncation and round-off errors in computation of matrix exponentials, *Internat. J. Control* **33** (1) (1981) 137–147.
[9] C.J. Standish, Truncated Taylor series approximation to the state transition matrix of a continuous parameter finite Markov chain, *Linear Algebra Appl.* **12** (1975) 179–183.
[10] C.F. Van Loan, On the limitation and application of the Padé approximation to the matrix exponential, in: E.B. Saff and R.S. Varga, Eds., *Padé and Rational Approximation: Theory and Applications* (Academic Press, New York, 1977) 439–448.
[11] C.F. Van Loan, The sensitivity of the matrix exponential, *SIAM J. Numer. Anal.* **14** (6) (1977) 971–981.
[12] R.C. Ward, Numerical computation of the matrix exponential with accuracy estimate, *SIAM J. Numer. Anal.* **14** (4) (1977) 600–610.
[13] R.C. Ward, Statistical round-off error analysis of a Padé algorithm for computing the matrix exponential, in: E.B. Saff and R.S. Varga, Eds., *Padé and Rational Approximation: Theory and Applications* (Academic Press, New York, 1977) 449–460.