

Wider contours and adaptive contours

Shev MacNamara, William McLean and Kevin Burrage

Abstract Contour integrals in the complex plane are the basis of effective numerical methods for computing matrix functions, such as the matrix exponential and the Mittag-Leffler function. These methods provide successful ways to solve partial differential equations, such as convection–diffusion models. Part of the success of these methods comes from exploiting the freedom to choose the contour, by appealing to Cauchy’s theorem. However, the pseudospectra of non-normal matrices or operators present a challenge for these methods: if the contour is too close to regions where the norm of the resolvent matrix is large, then the accuracy suffers. Important applications that involve non-normal matrices or operators include the Black–Scholes equation of finance, and Fokker–Planck equations for stochastic models arising in biology. Consequently, it is crucial to choose the contour carefully. As a remedy, we discuss choosing a contour that is wider than it might otherwise have been for a normal matrix or operator. We also suggest a semi-analytic approach to adapting the contour, in the form of a parabolic bound that is derived by estimating the field of values.

To demonstrate the utility of the approaches that we advocate, we study three models in biology: a monomolecular reaction, a bimolecular reaction and a trimolecular reaction. Modelling and simulation of these reactions is done within the framework of Markov processes. We also consider non-Markov generalisations that

Shev MacNamara

Australian Research Council Centre of Excellence for Mathematical and Statistical Frontiers (ACEMS), University of Technology Sydney, e-mail: shev.macnamara@uts.edu.au

William McLean

The School of Mathematics and Statistics, University of New South Wales, Sydney, e-mail: w.mclean@unsw.edu.au

Kevin Burrage

Australian Research Council Centre of Excellence for Mathematical and Statistical Frontiers (ACEMS), School of Mathematical Sciences, Queensland University of Technology, Brisbane, Australia and Visiting Professor to the University of Oxford, U.K., e-mail: kevin.burrage@qut.edu.au

have Mittag-Leffler waiting times instead of the usual exponential waiting times of a Markov process.

1 Introduction

We begin with the Chapman–Kolmogorov forward equation, associated with a Markov process on discrete states, for the evolution in continuous time of the probability of being in state j at time t :

$$\frac{d}{dt}p(j,t) = -|a_{jj}|p(j,t) + \sum_{i \neq j} a_{j,i}p(i,t). \quad (1)$$

Here for $j \neq i$, $a_{i,j} \geq 0$, and $a_{i,j}dt$ is approximately the probability to transition from state j to state i in a small time dt . The diagonal entry of an associated matrix $\mathbb{A} = \{a_{i,j}\}$ is defined by the requirement that the matrix has columns that sum to zero, namely $a_{jj} = -\sum_{i \neq j} a_{i,j}$. This equation in the Markov setting (1) can be generalised to a non-Markovian form:

$$\frac{d}{dt}p(j,t) = -\int_0^t K(j,t-u)p(j,u)du + \sum_{i \neq j} \frac{a_{j,i}}{|a_{jj}|} \int_0^t K(i,t-u)p(i,u)du. \quad (2)$$

Here the so-called memory function $K(i,t-u)$ is defined via its Laplace transform $\hat{K}(j,s)$, as the ratio of the Laplace transform of the waiting time to the Laplace transform of the survival time. The process is a Markov process if and only if the waiting times are exponential random variables, in which case the $K(j,t)$ appearing in the convolutions in (2) are Dirac delta distributions and (2) reduces to the usual equation in (1). In the special case of Mittag-Leffler waiting times, (2) can be rewritten as [12, 13]:

$$D_t^\alpha \mathbf{p} = \mathbb{A} \mathbf{p} \quad \text{with solution} \quad \mathbf{p}(t) = E_\alpha(\mathbb{A}t^\alpha) \mathbf{p}(0). \quad (3)$$

Here D_t^α denotes the Caputo fractional derivative, and where the Mittag-Leffler function is

$$E_\alpha(z) = \sum_{k=0}^{\infty} \frac{z^k}{\Gamma(\alpha k + 1)}. \quad (4)$$

Here $0 \leq \alpha \leq 1$. In Section 4, we provide a MATLAB code to simulate sample paths of this stochastic process with Mittag-Leffler waiting times. When $\alpha = 1$ the series (4) reduces to the usual exponential function, and the fractional equation in (3) reduces to the original Markov process (1). In Section 5, we provide a MATLAB code to compute the solution $E_\alpha(\mathbb{A}t^\alpha) \mathbf{p}(0)$ in (3) directly, via a contour integral.

Next, we introduce the (Markovian) Fokker–Planck partial differential equation (sometimes also known as the forward Kolmogorov equation)

$$\frac{\partial}{\partial t} p(x,t) = -\frac{\partial}{\partial x} (a(x)p(x,t)) + \frac{1}{2} \frac{\partial^2}{\partial x^2} (b(x)p(x,t)) \quad (5)$$

for a probability density $p(x,t)$ in one space dimension and time, with suitable boundary conditions, initial conditions, and smoothness assumptions on the coefficients $a(x)$, and on $b(x) \geq 0$. Later in the work we will use complex-variable methods, so it is worth noting at the outset that our coefficients $a(x)$ and $b(x)$ are always real-valued. It is also worth noting that $a(x)$ here in (5) is not the same as $a_{i,j}$ appearing in the matrix above in (3), although there is a close relationship that allows one to be deduced from the other. This PDE for the density corresponds to a stochastic differential equation for sample paths

$$dX = a(X)dt + \sqrt{b(X)}dW. \quad (6)$$

Both the PDE (5) and SDE (6) are continuous in the spatial variable. An introduction to these models, and their connections with discrete versions, can be found in [1, 5]. Our discrete models do respect important properties such as non-negativity. However, there are issues with the Fokker–Planck PDE model, the Chemical Langevin SDE model, and other diffusion approximations: often these approximations do not maintain positivity. These issues are discussed by Williams in the Kolmogorov Lecture and accompanying paper, where a method that maintains non-negativity is proposed [11].

We do not simulate or solve either of the PDE (5) or the SDE (6). We do however simulate and solve the closely related models that are discrete in space and that are governed by master equations (1) or generalized master equations (2), which can be thought of as finite difference discretizations of (5). In particular, the PDE (5) can be thought of as a continuous-in-space analogue of the discrete process in (1) that involves a matrix \mathbb{A} . The utility of the PDE is that it is easier to find an estimate of the field of values of the spatial differential operator on the right hand side of (5) than of the corresponding matrix \mathbb{A} . We can then use an estimate of one as an approximation for the other.

Developing appropriate methods for modelling and simulation of these important stochastic processes is a necessary first step for more advanced scientific endeavors. A natural next step is an inverse problem, although we do not pursue that in this article. For example, it is of great interest to estimate the rate constants in the models described in Section 2, typically based on limited observations of samples resembling the simulation displayed in Figure 4. It is more challenging to estimate the parameter α in fractional models, or indeed to address the question of whether or not a single α (the only case considered in this article) is appropriate. Interest in fractional models is growing fast as they are finding wide applications including in cardiac modelling [2], and this includes exciting new applications of Mittag-Leffler functions [3].

Section 2 introduces three models that can be cast in the form of (3) and Section 3 uses these models as vignettes to exhibit manifestations of pseudospectra. Next we introduce two methods of simulation for these models. A Monte Carlo approach is presented in Section 4. Section 5 presents an alternative method that directly com-

puts the solution of (3) as a probability vector via a contour integral. Finally we suggest a bound on the field of values that is useful when designing contour methods.

To our knowledge Figure 3 is the first visualization of the pseudospectra of the Schlögl reactions. The estimate in (24) is also new. In fact (24) is the specialization of our more general estimates appearing in (22) and (23) to the monomolecular model, but it should be possible to likewise adapt our more general estimates to other models such as bimolecular models.

2 Three fundamental models

All three models that we present are represented by *tri-diagonal matrices*: $j \notin \{i-1, i, i+1\} \Rightarrow \mathbb{A}_{i,j} = 0$. Since all other entries are zero, below, we only specify the non-zero entries on the three main diagonals. In fact, an entry on the main diagonal is determined by the requirement that columns sum to zero (which corresponds to conservation of probability), so it would suffice to specify only the two non-zero diagonals immediately below and above the main diagonal.

2.1 Monomolecular, bimolecular and trimolecular models

A model of a monomolecular reaction



such as arises in chemical isomerisation [5], or in models of ion channels in cardiac-electrophysiology, or neuro-electrophysiology, can be represented by the following $N \times N$ matrix.

Monomolecular model matrix

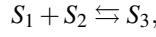
$$\mathbb{A}_{i,j} = \begin{cases} j-1, & i = j-1, \\ -m, & i = j, \\ m-j+1, & i = j+1. \end{cases} \quad (7)$$

Here we have assumed that the rate constants are equal to unity, $c_1 = c_2 = 1$, and we have assumed that $m = N - 1$ where m is the maximum number of molecules. More details, including exact solutions, can be found in [8]. An instance of this matrix when $m = 5$ is

$$\mathbb{A} = \begin{pmatrix} -5 & 1 & & & & \\ & 5 & -5 & 2 & & \\ & & 4 & -5 & 3 & \\ & & & 3 & -5 & 4 \\ & & & & 2 & -5 & 5 \\ & & & & & 1 & -5 \end{pmatrix}. \quad (8)$$

The model is two-dimensional, but the conservation law allows it to be treated as effectively one-dimensional, by letting x represent S_1 , say. Then one possible corresponding continuous model (5) has drift coefficient $a(x) = -c_1x + c_2(m-x)$ and diffusion coefficient $b(x) = c_1x + c_2(m-x)$, for $0 < x < m$. In the discrete Markov case, the exact solution is binomial. When $c_1 = c_2$ the stationary probability density of the continuous Markov model is given by Gillespie [5] as a Gaussian, which highlights the issues associated with the continuous models such as choosing the domain, boundary conditions, and respecting positivity. To enforce positivity, we might instead pose the PDE on the restricted domain $0 \leq x \leq m$.

Next we introduce a model for the bimolecular reaction



via the following $N \times N$ matrix.

Bimolecular model matrix

$$\mathbb{A}_{i,j} = \begin{cases} j-1, & i = j-1, \\ -j+1 - (m-j+1)^2 & i = j, \\ (m-j+1)^2, & i = j+1. \end{cases} \quad (9)$$

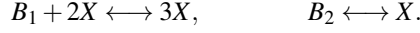
This matrix and the model are also introduced in [12], where more details can be found. A small example with $m = 5$ is

$$\mathbb{A} = \begin{pmatrix} -25 & 1 & & & & \\ & 25 & -17 & 2 & & \\ & & 16 & -11 & 3 & \\ & & & 9 & -7 & 4 \\ & & & & 4 & -5 & 5 \\ & & & & & 1 & -5 \end{pmatrix}. \quad (10)$$

Here we have assumed that the rate constants are equal to unity, $c_1 = c_2 = 1$, and that the initial condition is $[S_1, S_2, S_3] = [m, m, 0]$, so m is the maximum number of molecules, and $m = N - 1$. The model is three-dimensional, but the conservation law together with this initial condition allow it to be treated as effectively one-dimensional, by letting x represent S_3 , say. Then a possible corresponding continu-

ous model (5) has drift coefficient $a(x) = -c_1x + c_2(m-x)^2$ and diffusion coefficient $b(x) = c_1x + c_2(m-x)^2$.

Finally, we introduce the Schlögl model [13], which consists of two reversible reactions



Here $B_1 = 1 \times 10^5$, $B_2 = 2 \times 10^5$. The associated matrix is given below, where $k_1 = 3 \times 10^{-7}$, $k_2 = 1 \times 10^{-4}$, $k_3 = 1 \times 10^{-3}$, and $k_4 = 3.5$. In theory this matrix is infinite but we truncate to a finite section for numerical computation.

Schlögl model matrix (an example of a **trimolecular model** scheme)

$$\mathbb{A}_{i,j} = \begin{cases} \frac{1}{6}k_2(j-1)(j-2)(j-3) + k_4(j-1), & i = j-1, \\ k_3B_2 + \frac{1}{2}k_1B_1(j-1)(j-2), & i = j+1. \end{cases} \quad (11)$$

For $i = j$, the diagonal entry is $-(\frac{1}{6}k_2(j-1)(j-2)(j-3) + k_4(j-1) + k_3B_2 + \frac{1}{2}k_1B_1(j-1)(j-2))$. The first column is indexed by $j = 1$ and corresponds to a state with $0 = j - 1$ molecules. The corresponding continuous model (5) has drift coefficient $a(x) = k_3B_2 + \frac{1}{2}k_1B_1x(x-1) - \frac{1}{6}k_2x(x-1)(x-2) - k_4x$ and diffusion coefficient $b(x) = k_3B_2 + \frac{1}{2}k_1B_1x(x-1) + \frac{1}{6}k_2x(x-1)(x-2) + k_4x$.

3 Pseudospectra are important for stochastic processes

All three matrices introduced in the previous section exhibit interesting pseudospectra. As an illustration of the way that the pseudospectra manifest themselves, we will now consider numerically computing eigenvalues of the three matrices. This is a first demonstration that one must respect the pseudospectrum when crafting a numerical method. That issue will be important again when we use numerical methods based on contour integrals to compute Mittag-Leffler matrix functions.

The reader can readily verify that using any standard eigenvalue solver, such as `eig` in MATLAB, leads to numerically computed eigenvalues that are complex numbers. However, these numerically computed complex eigenvalues are *wrong*: the *true eigenvalues are purely real*. It is the same story for all three models. See the numerically computed eigenvalues displayed here in Figure 1 and Figure 2, for example. We suggest this effect happens much more widely for models arising in computational biology.

Figure 1 and Figure 2 make use of the following method to create a diagonal scaling matrix. Here is a MATLAB listing to create a diagonal matrix D that symmetrizes a tridiagonal matrix of dimension N of the form described in any of the three models considered here:

```

d(1) = 1;
for i = 1:N-1
    d(i+1) = sqrt(A(i,i+1)/A(i+1,i)) *d(i);
end
D = diag(d);    Asym = D*A*inv(D);

```

This symmetrization by a diagonal matrix in a similarity transform is known to physicists as a gauge transformation, and it is described by Trefethen and Embree [15, Section 12]. A real symmetric matrix has real eigenvalues so the eigenvalues of DAD^{-1} are purely real. The matrix DAD^{-1} and the matrix A share the same eigenvalues because they are similar. This is one way to confirm that the true eigenvalues of A are purely real. Numerical algorithms typically perform well on real symmetric matrices, so it is better to numerically compute the eigenvalues of DAD^{-1} than to compute eigenvalues of A directly.

The ε -pseudospectra [15] can be defined as the region of the complex plane where the norm of the resolvent matrix is large: the set $z \in \mathbb{C}$ such that

$$\|(z\mathbb{I} - \mathbb{A})^{-1}\| > \frac{1}{\varepsilon}. \quad (12)$$

We use the 2-norm in this article. Equivalently, this is the region of the complex plane for which z is an eigenvalue of $(\mathbb{A} + \mathbb{E})$ for some perturbing matrix that has a small norm: $\|\mathbb{E}\| < \varepsilon$. Thus the wrong complex ‘eigenvalues’ appearing as crosses in Figure 1 and Figure 2 offer a way to (crudely) visualise the pseudospectrum. Numerical visualizations of the pseudospectra of the family of monomolecular matrices defined by (7) can be found in [8], and the pseudospectra of the bimolecular reaction defined by (9) can be found in [12]. Here, as an example of a trimolecular scheme, we present in Figure 3 the pseudospectra for the Schlögl reactions (11), as computed by `EigTool`. The resolvent norm is largest near the negative real axis, as we would expect because that is where the true eigenvalues are located. The level curves are not equally spaced in the figure, becoming bunched up together, suggesting a more interesting structure that remains to be elucidated as the dimension of the matrix grows. An interesting experiment is to vary the parameters, namely the product k_3B_2 , as is done in [14, Figure 10]. In experiments not shown here, we find that when k_3B_2 is varied, the numerical computation of the eigenvalues becomes less reliable (for example, when $k_3 = 1.4 \times 10^{-3}$ and B_2 is unchanged).

Figure 1 and Figure 2 also display the *numerical range* of a matrix. That is also known as the *field of values*, which we denote by W , and it is defined for a matrix \mathbb{A} , as the set of complex numbers that come from a quadratic form with the matrix

$$W(\mathbb{A}) \equiv \{x^* \mathbb{A} x \in \mathbb{C} : \|x\|_2 = 1\}. \quad (13)$$

The field of values always contains the eigenvalues. A simple algorithm [9] for computing $W(\mathbb{A})$ involves repeatedly ‘rotating’ the matrix and then finding the *nu-*

merical abscissa

$$\frac{1}{2} \max(\text{eig}(\mathbb{A} + \mathbb{A}^*)). \quad (14)$$

The ε -pseudospectrum of a matrix is contained in an ε -neighbourhood of the field of values, in a sense that can be made a precise theorem [15]. We see an example of this in Figure 1 and Figure 2.

4 A Mittag-Leffler Stochastic Simulation Algorithm

In this short section we provide a method for Monte Carlo simulation of the solutions of the stochastic processes that we describe. This Monte Carlo method can be considered an alternative to contour integral methods, that we describe later. As the Monte Carlo methods do not seem to fail in the presence of non-normality, they can be a useful cross-check on contour integral methods.

Here is a `MATLAB` listing to simulate Monte Carlo sample paths of the monomolecular model of isomerization corresponding to the matrix in (7), with Mittag-Leffler waiting times:

```
t_final = 100; m=10; c1=1; c2=1; v = [-1, 1; 1, -1];
initial_state = [m,0]'; alpha = 0.9; all = 1/alpha;
alpi=alpha*pi; sinalpi=sin(alpi); cosalpi=cos(alpi);
t = 0; x = initial_state; T = [0]; X = [initial_state];
while (t < t_final)
    a(1) = c1*x(1); a(2) = c2*x(2); asum = sum(a);
    r1=rand; r2=rand; z=sinalpi/tan(alpi*r2)-cosalpi;
    tau = -(z/asum)^(all)*log(r1);
    r = rand*asum; j = find(r<cumsum(a),1,'first');
    x = x + v(:,j); t =t+tau; T = [T t]; X = [X x];
end
if (t > t_final)
    T(end) = t_final; X(:,end) = X(:,end-1);
end
figure(1); stairs(T, X(1,:), 'LineWidth', 2);
xlabel('Time'); ylabel('molecules of S_1');
title({'Isomerisation: a monomolecular model';...
'Mittag-Leffler SSA'; ['\alpha == ', num2str(alpha)]})
```

This is a Markov process with exponential waiting times when $\alpha = 1$, in which case the program reduces to a version of the Gillespie Stochastic Simulation Algorithm. Figure 4 shows a sample path simulated with this `MATLAB` listing. A histogram of many such Monte Carlo samples is one way to approximate the solution

of (3). A different way to compute that solution is via a contour integral, as we describe in the next section, and as displayed in Figure 5.

A time-fractional diffusion equation in the form of (3) in two space dimensions is solved by such a contour integral method in [16, Figure 16.2]. As an aside, we can modify the code above to offer a method for Monte Carlo simulation of sample paths of a closely related stochastic process, as follows. Simplify to one space dimension (it is also straight forward to simulate in two dimensions), and by supposing that the governing $(m+1) \times (m+1)$ matrix is

$$\mathbb{A} = \begin{pmatrix} -1 & 1 & & & \\ & 1 & -2 & 1 & \\ & & \dots & & \\ & & & 1 & -2 & 1 \\ & & & & 1 & -1 \end{pmatrix} \quad (15)$$

with initial vector $p(0)$ being zero everywhere except the middle entry, i.e. the $\text{round}(m/2)$ th entry is one. Then (3) corresponds to a random walk on the integers $0, 1, \dots, m$, beginning in the middle. To simulate, modify the first few lines of the above code segment to

```
t_final = 1; m=10; v = [-1, 1];
initial_state = round(m/2); ...
```

and also the first line in the while loop, to

```
a(1) = (x>0); a(2) = (x<m);
```

leaving the rest unchanged.

5 Computing a Mittag-Leffler matrix function

In this section we first establish the utility of contour integral methods by directly computing the desired solution of (3). This is motivation for exploring bounds on the pseudospectrum, so that informed choices can be made when designing contour methods.

5.1 Computing contour integrals

Start with (3). Take the Laplace transform. Then take the inverse Laplace transform. Of course those two steps arrive at the same solution we started with. The advantage is the desired solution of (3) represented as a contour integral [12, 13, 16]:

$$\mathbf{p}(t) = E_{\alpha}(\mathbb{A}t^{\alpha})\mathbf{p}(0) = \frac{1}{2\pi i} \int_{\Gamma} \exp(zt) (z\mathbb{I} - z^{1-\alpha}\mathbb{A})^{-1} \mathbf{p}(0) dz. \quad (16)$$

The eigenvalues of \mathbb{A} must lie to the left of the contour Γ . In all our examples, they lie along the negative real axis. Also, we exploit symmetry when \mathbb{A} is real. Apart from those requirements, we are free to choose the contour. Typical choices are displayed in Figure 6. The $\exp(zt)$ factor is nearly zero when the real part of z is sufficiently negative. This motivates choosing Γ to go into the left-half plane because then we can neglect the infinite part of the contour that lies to the left of say, -30 , and still obtain high accuracy. To evaluate the integral on the remaining finite part of the contour, we use quadrature. The trapezoidal rule is a good choice. For a quadrature rule with M nodes z_k on the contour, we approximate (16) by

$$\mathbf{p}(t) \approx \sum_{k=1}^M w_k(t) \mathbf{u}_k \quad (17)$$

where we solve a linear system for the vector \mathbf{u}_k at each node k

$$(z\mathbb{I} - z^{1-\alpha}\mathbb{A}) \mathbf{u}_k = \mathbf{p}(0).$$

The values $w_k(t)$ (and note that we allow these numbers to depend on t) incorporate both the usual weights coming from the quadrature rule on a line, and also any scalings coming from parameterizing the contour. There are variations of this quadrature scheme. When $\alpha = 1$, this procedure simplifies to compute the familiar matrix exponential solutions of the usual Markov systems.

Here is a MATLAB listing, adapted from Le, Mclean & Lamichhane [10], to compute the Mittag-Leffler solution by applying the quadrature recipe (17) on a *hyperbolic* contour.

```
function p = hyperbola_MittagLeffler(t,A,v,alpha,M)
alpha1 = 1-alpha; n = length(v); Dxi = 1.08179214/M;
xi = [-M:M]*Dxi; delta=1.17210423; mu =4.49207528*M/t;
z = mu * ( 1 - sin(complex(delta,-xi)) );
dz = 1i * mu * cos(complex(delta,-xi)) ;
c = Dxi * dz .* exp(z*t) / (2*pi*1i);
I = speye(n); p = zeros(size(v));
for k = 1:M
    p = p + c(k) * ((z(k)*I - z(k)^(alpha1)*A)\v);
end
p = 2*real(p); k = M+1;
p = p + real(c(k) * ((z(k)*I - z(k)^(alpha1)*A)\v));
```

We choose $M = 16$ quadrature points. An example of the hyperbolic contour being used here is displayed in Figure 6. Figure 5 shows a solution of (3) computed with this MATLAB listing.

5.2 Estimating the field of values

One strategy to choose the contour, Γ , is to first compute a pseudospectrum of the matrix with `EigTool` [17] (as in Figure 3 for example). Then choose Γ so that $\|(z\mathbb{I} - \mathbb{A})^{-1}\|$ is never too large for $z \in \Gamma$. For example, we might choose the contour so that the bound $\|(z\mathbb{I} - \mathbb{A})^{-1}\| < 10^3$ holds (which corresponds to $\varepsilon = 10^{-3}$ in (12)). Arguably, the value 10^3 could be replaced by something smaller, say $\mathcal{O}(1)$. The particular value would depend on the application. In numerical experiments with these models, the issue seems to matter only when $\|(z\mathbb{I} - \mathbb{A})^{-1}\|$ is significantly larger than, say, 100. Such an unfavourable situation can certainly arise. It is demonstrated to happen on the parabolic contour displayed here in Figure 6 for bimolecular reactions [12]. The issue is also addressed by In't Hout and Weideman [7] for Black–Scholes models. Figure 6 compares a parabolic contour used for a bimolecular model [12] with a hyperbolic contour used here. A similar comparison of these contours and discussion can be found in a survey article by Trefethen and Weideman [16, Figure 15.2].

However, a drawback of this procedure is that it is expensive to first compute the pseudospectrum. It would therefore be preferable to instead find an estimate of the region where the resolvent is large by some more efficient means. We find one such estimate next.

Trefethen and Embree [15] discuss various ways to bound the pseudospectrum. One approach uses the fact that if z is not inside the field of values of a matrix, then the norm of the resolvent is bounded by the distance to the field of values:

$$\|(z\mathbb{I} - \mathbb{A})^{-1}\|_2 < \frac{1}{\text{dist}(z, W(\mathbb{A}))}.$$

This result suggests an idea for adapting the contour: choose the contour to be outside of the field of values.

We now estimate the field of values of the spatial operator in the Fokker–Planck PDE (5); a similar approach has been applied to the Black–Scholes equation of finance [7, Theorem 3.1]. We focus on the particular example of the monomolecular model, but we believe this approach will be extended to the other models in future work.

Begin by writing the Fokker–Planck equation in the form of a conservation law,

$$u_t + Au = 0 \quad \text{for } 0 < x < m \text{ and } t > 0,$$

where, using a dash for $\partial/\partial x$,

$$Au = -\frac{1}{2}(bu)'' + (au)' = -\left(\frac{1}{2}bu' + Bu\right)' \quad \text{and} \quad B(x) = -a(x) + \frac{1}{2}b'(x). \quad (18)$$

For the monomolecular model, the coefficients are

$$b(x) = c_1x + c_2(m - x) \quad \text{and} \quad a(x) = -c_1x + c_2(m - x).$$

Here, c_1 , c_2 and m are positive constants. Note that A is uniformly elliptic because $b(x) \geq \min(c_1, mc_2)$ for $0 < x < m$.

We impose either homogeneous Dirichlet boundary conditions,

$$u(0) = 0 = u(m), \quad (19)$$

or else zero-flux boundary conditions,

$$\frac{1}{2}bu' + Bu = 0 \quad \text{for } x \in \{0, m\}. \quad (20)$$

The domain of A is then the complex vector space $D(A)$ of C^2 functions $v : [0, m] \rightarrow \mathbb{C}$ satisfying the chosen boundary conditions.

Denote the numerical range (field of values) of A by

$$W(A) = \{ \langle Au, u \rangle : u \in D(A) \text{ with } \langle u, u \rangle = 1 \}, \quad (21)$$

where $\langle u, v \rangle = \int_0^m u \bar{v}$. Compare this definition of $W(A)$ in the continuous setting with the definition (13) in the discrete setting. For either (19) or (20), integration by parts gives

$$\langle Au, u \rangle = \frac{1}{2} \int_0^m b|u'|^2 + \int_0^m Bu\bar{u}'.$$

We write

$$X + iY \equiv \langle Au, u \rangle = \frac{1}{2}P - Q \quad \text{where} \quad P = \int_0^m b|u'|^2 \quad \text{and} \quad Q = - \int_0^m Bu\bar{u}',$$

and assume that

$$\frac{B(x)^2}{b(x)} \leq K \quad \text{and} \quad 0 < \beta_0 \leq \frac{B'(x)}{2} \leq \beta_1 \quad \text{for } 0 \leq x \leq m.$$

In the case of zero-flux boundary conditions (20), we require the additional assumption

$$B(0) \leq 0 \leq B(m).$$

Then we claim that for Dirichlet boundary conditions (19),

$$Y^2 \leq 2K(X + \beta_1) - \beta_0^2, \quad (22)$$

whereas for zero-flux boundary conditions (20),

$$Y^2 \leq 2K(X + \beta_1). \quad (23)$$

To derive these estimates, first observe that

$$2\Re Q = Q + \bar{Q} = - \int_0^m B(u\bar{u}' + \bar{u}u') = - \int_0^m B(u\bar{u})' = -[B|u|^2]_0^m + \int_0^m B'|u|^2.$$

Assume that $\langle u, u \rangle = 1$. The bounds on B' give

$$\beta_0 - [\tfrac{1}{2}B|u|^2]_0^m \leq \Re Q \leq \beta_1 - [\tfrac{1}{2}B|u|^2]_0^m,$$

and by the Cauchy–Schwarz inequality,

$$|Q|^2 \leq \left(\int_0^m B^2 |u'|^2 \right) \left(\int_0^m |u|^2 \right) = \int_0^m \frac{B^2}{b} b |u'|^2 \leq \left(\max_{[0,m]} \frac{B^2}{b} \right) \int_0^m b |u'|^2 \leq KP;$$

thus,

$$Y^2 = (-\Im Q)^2 = |Q|^2 - (\Re Q)^2 \leq KP - (\Re Q)^2.$$

For Dirichlet boundary conditions we have $[B|u|^2]_0^m = 0$, so $\beta_0 \leq \Re Q \leq \beta_1$ and hence

$$P = 2X + 2\Re Q \leq 2X + 2\beta_1 \quad \text{and} \quad Y^2 \leq KP - \beta_0^2,$$

implying (22). For zero-flux boundary conditions, the extra assumptions on B ensure that $[B|u|^2]_0^m \geq 0$, so $\Re Q \leq \beta_1$ and hence

$$P = 2X + 2\Re Q \leq 2X + 2\beta_1 \quad \text{and} \quad Y^2 \leq KP,$$

implying (23).

In the simple case $c_1 = c_2 = c$ we have

$$b(x) = cm \quad \text{and} \quad B(x) = -a(x) = c(2x - m)$$

and therefore

$$\frac{B(x)^2}{b(x)} = \frac{c}{m} (2x - m)^2 \leq cm \quad \text{and} \quad B'(x) = 2c \quad \text{for } 0 \leq x \leq m,$$

giving $K = cm$ and $\beta_0 = \beta_1 = c$. In addition, $B(0) = -cm \leq 0$ and $B(m) = cm \geq 0$. Thus, for zero-flux boundary conditions,

$$Y^2 \leq 2cm(X + c).$$

Note that the sign convention in the notation in (18) makes the operator A positive definite (though not symmetric), whereas our model matrices are negative definite, so $W(A)$ provides an approximation for $W(-\mathbb{A}) = -W(\mathbb{A})$. We therefore have to flip signs in (22) and (23) to get estimates for $W(\mathbb{A})$, as in the following bound.

Parabolic estimate of the field of values for the **monomolecular model** matrix when $c_1 = c_2 = c$ in (7):

$$Y^2 \leq 2cm(c - X). \tag{24}$$

This bound is displayed in Figure 7. As in the Figure, our matrix examples typically have a unique zero eigenvalue, with all other eigenvalues having negative real part, and the numerical abscissa (14) is typically strictly positive. Having derived a bound in the continuous setting, we can only regard it as an approximation in the discrete setting. Nonetheless, in these numerical experiments it does indeed seem to offer a useful bound for the discrete case.

Discussion. A natural next step is to incorporate this bound into the design of the contour for methods such as the MATLAB listing provided in Section 5. That will be pursued elsewhere. This article has thus laid the foundations for such attractive future research.

It is worth commenting that if a matrix is real symmetric (unlike the examples in this article) then the pseudospectrum is not an issue so it would not be good to make the contour wider (and thus also longer), and instead previously proposed (and often optimized and shorter) contours such as surveyed in [16] would be good choices. Making the contour wider and thus longer does impact the numerical method, but this is unavoidable in applications where the behaviour of the pseudospectra is an issue, such as the examples discussed here. It is also worth commenting on numerical methods for computing the field of values. Here we used `Chebfun`, which in turn is based on Johnson’s algorithm [9]. It is conceivable that a different algorithm might be devised to estimate the field of values more cheaply for the purposes of contour integrals, but we do not explore such an approach here. However, we do observe in these numerical experiments, as displayed in the figures for example, that the field of values — and therefore also the estimate that we derive from it — seems to be too conservative. That is especially noticeable in regions where the real part is large and negative. It might therefore also be worth considering methods of more cheaply, directly, estimating the pseudospectra. We began with the continuous PDE so using a coarse mesh in a PDE-solver, such as a finite difference method, might be one way to obtain a cheap estimate of the pseudospectra.

6 Conclusion

We have discussed, and illustrated by example, the significance of the pseudospectra for stochastic models in biology, including both Markovian and non-Markovian generalisations. Although we focused exclusively on the 2-norm, the pseudospectra of these stochastic processes are perhaps more naturally addressed in the 1-norm. In that regard, future work will explore ways to incorporate the 1-norm methods described by Higham and Tisseur [6], to adaptively choosing the contour.

Numerical methods to compute the matrix exponential functions and matrix Mittag-Leffler functions via contour integrals must take the pseudospectrum of the matrix into account. In particular, such methods must choose contours that are wide enough, and that adapt to avoid regions of the complex plane where the norm of the resolvent matrix is too large. We have derived a simple, parabolic bound on the field of values of an associated Fokker–Planck PDE, which can be used as an approxima-

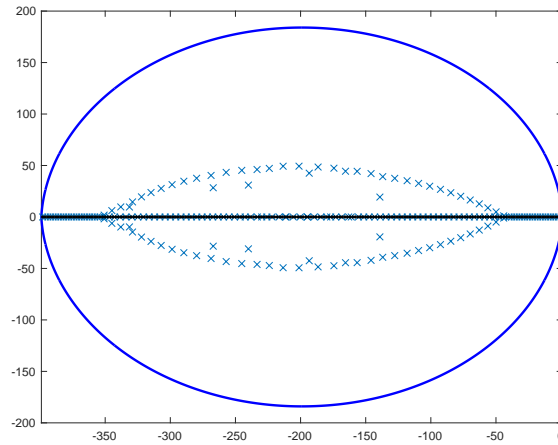


Fig. 1 The field of values (solid line resembling an oval-like shape) for the discrete monomolecular matrix in (7) when $N = 200$ ($m = N - 1$), as computed by `Chebfun` [4]. The crosses mark numerically computed eigenvalues via `eig`, but complex eigenvalues are *wrong*. The true eigenvalues are purely real and are marked on the real axis by dots (note that the dots are so close together that they may seem to be almost a solid line). These correct values can come by instead computing the eigenvalues of the symmetrized matrix, after using the diagonal scaling matrix created by the MATLAB listing provided here.

tion to the field of values of the corresponding discrete matrix model. We believe this bound can help inform us when making a good choice for the contour. Ultimately, how to devise contours in a truly adaptive fashion, so that they can be efficiently computed, automatically and without a priori analysis, remains an important open question.

Acknowledgements The authors thank the organisers of the Computational Inverse Problems theme at the MATRIX, 2017.

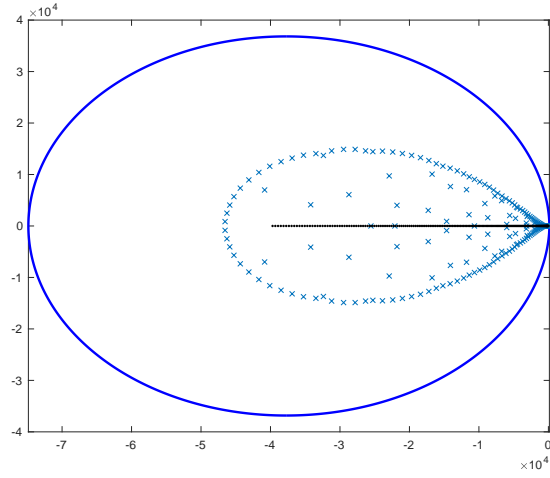


Fig. 2 Same as Figure 1, for the bimolecular model. The field of values for the discrete bimolecular matrix in (9) when $N = 200$ is the solid line resembling an oval-like shape, as computed by Chebfun [4].

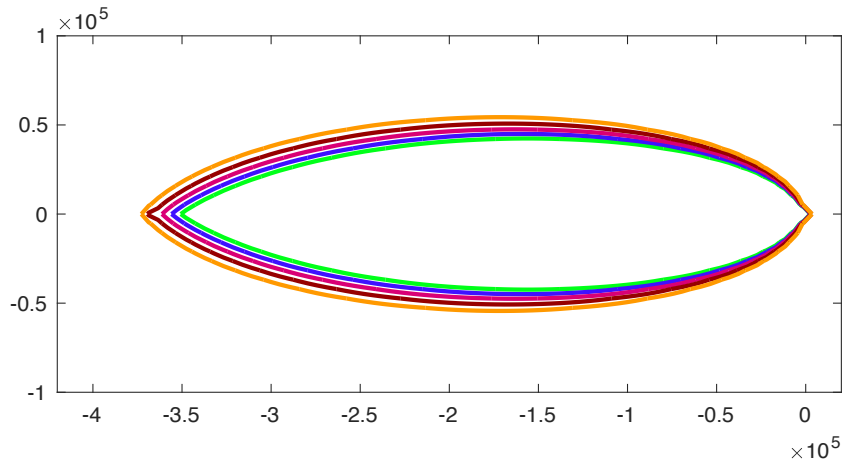


Fig. 3 Pseudospectrum of a 2000×2000 finite section of the matrix in (11) representing the Schlögl reactions, as computed by EigTool. Contours correspond to $\varepsilon = 10^{-2}, 10^{-4}, 10^{-6}, 10^{-8}, 10^{-10}$ in (12). The contour closest to the real axis corresponds to 10^{-10} .

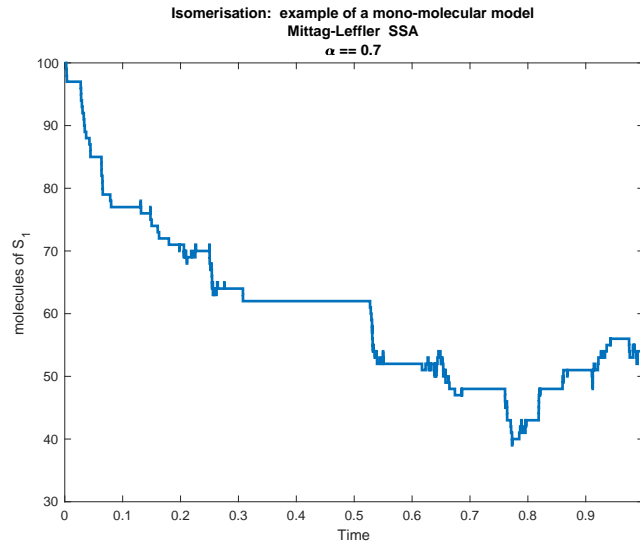


Fig. 4 A sample simulation of the monomolecular reaction, with the Mittag-Leffler SSA provided here in the MATLAB listing. Parameters: $\alpha = 0.7$, $t = 1$, and initial condition a Dirac delta distribution on the state $[S_1, S_2] = [m, 0] = [100, 0]$. Compare with Figure 5.

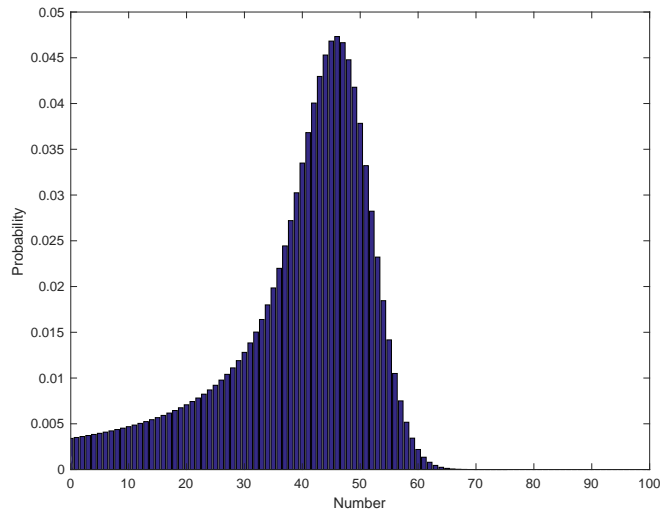


Fig. 5 A discrete probability mass function is the solution of a Mittag-Leffler master equation (3) for the monomolecular reaction, and can be computed with the MATLAB listing provided here that uses a hyperbolic contour. Parameters: $\alpha = 0.7$, $t = 1$, and initial condition a Dirac delta distribution on the state $[S_1, S_2] = [m, 0] = [100, 0]$. The x -axis shows the number of molecules of S_2 . Compare with Figure 4.

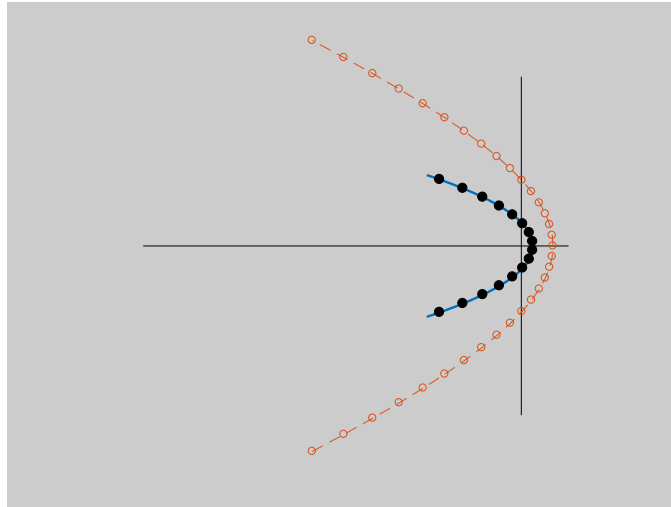


Fig. 6 A parabolic contour (solid) as described in [16], and a hyperbolic contour (dashed, with $M = 16$) used in the `MATLAB` listing provided here. Nodes for a quadrature rule as in (17) are also marked on the contours.

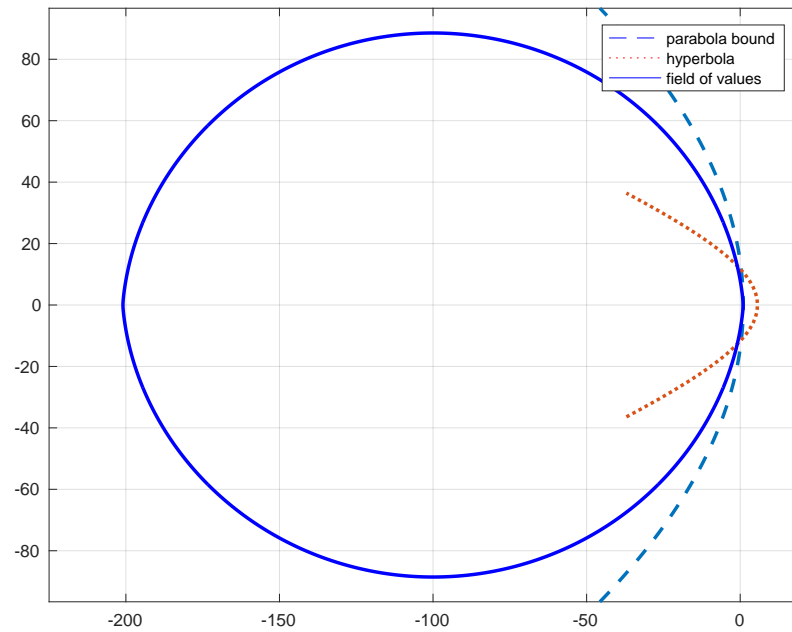


Fig. 7 The field of values for the discrete matrix in (7) when $c_1 = c_2 = 1$ and $m = 100$, as computed by `Chebfun`. The dashed line is a parabolic bound in (24) for the field of values of the corresponding continuous PDE (5). This parabolic bound is semi-analytic and requires some analysis of the equation ‘by hand.’ Also displayed is a hyperbolic contour, as in Figure 6 as used in the MATLAB listing provided here.

References

1. Anderson, D.F., Kurtz, T.G.: Continuous time Markov chain models for chemical reaction networks. In: Design and analysis of biomolecular circuits. Springer (2011)
2. Bueno-Orovio, A., Kay, D., Grau, V., Rodriguez, B., Burrage, K.: Fractional diffusion models of cardiac electrical propagation: role of structural heterogeneity in dispersion of repolarization. *J R Soc Interface* **11**(97), 20140,352 (2014)
3. Bueno-Orovio, A., Teh, I., Schneider, J.E., Burrage, K., Grau, V.: Anomalous diffusion in cardiac tissue as an index of myocardial microstructure. *IEEE TRANSACTIONS ON MEDICAL IMAGING* **35**(9), 2200–2207 (2016)
4. Driscoll, T.A., Hale, N., Trefethen, L.N.: Chebfun Guide. Pafnuty Publications (2014). URL <http://www.chebfun.org/docs/guide/>
5. Gillespie, D.T.: The chemical Langevin and Fokker–Planck equations for the reversible isomerization reaction. *The Journal of Physical Chemistry A* **106**(20), 5063–5071 (2002). DOI 10.1021/jp0128832
6. Higham, N.J., Tisseur, F.: A block algorithm for matrix 1-norm estimation, with an application to 1-norm pseudospectra. *SIAM J. MATRIX ANAL. APPL.* (2000)
7. in 't Hout, K.J., Weideman, J.A.C.: A contour integral method for the Black–Scholes and Heston equations. *SIAM J. Sci. Comput.* **33**, 763–785 (2011). DOI 10.1137/090776081
8. Iserles, A., MacNamara, S.: Magnus expansions and pseudospectra of Markov processes. <https://arxiv.org/abs/1701.02522> (2017)
9. Johnson, C.: Numerical determination of the field of values of a general complex matrix. *SIAM Journal on Numerical Analysis* **15**, 595–602 (1978)
10. Le, K.N., Mclean, W., Lamichhane, B.: Finite element approximation of a time-fractional diffusion problem for a domain with a re-entrant corner. *ANZIAM Journal* pp. 1–22 (2017)
11. Leite, S.C., Williams, R.J.: A constrained Langevin approximation for chemical reaction networks. Kolmogorov Lecture, Ninth World Congress In Probability and Statistics, Toronto, <http://www.math.ucsd.edu/~williams/biochem/biochem.pdf> (2016). URL <http://www.math.ucsd.edu/~williams/biochem/biochem.pdf>
12. Macnamara, S.: Cauchy integrals for computational solutions of master equations. *ANZIAM Journal* **56**, 32–51 (2015). DOI 10.21914/anziamj.v56i0.9345
13. MacNamara, S., Henry, B.I., McLean, W.: Fractional Euler limits and their applications. *SIAM Journal on Applied Mathematics* (2016)
14. Sargsyan, K., Debusschere, B., Najm, H., Marzouk, Y.: Bayesian inference of spectral expansions for predictability assessment in stochastic reaction networks. *Journal of Computational and Theoretical Nanoscience* **6**, 2283–2297 (2009)
15. Trefethen, L.N., Embree, M.: Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators. Princeton University Press (2005)
16. Trefethen, L.N., Weideman, J.A.C.: The exponentially convergent trapezoidal rule. *SIAM Review* (2014)
17. Wright, T.G.: Eigtool (2002). URL <http://www.comlab.ox.ac.uk/pseudospectra/eigtool/>