Beyond the classical theory of computational ordinary differential equations

Arieh Iserles

University of Cambridge

1 The classical problem and beyond

It is appropriate to commence this brief survey of the numerical analysis of 'strange' ordinary differential equations by quoting verbatim from Jack Lambert's informative survey in the last *State of the Art in Numerical Analysis* proceedings [22]: "By a *problem* we shall mean the initial value problem

$$\dot{\boldsymbol{y}} = \boldsymbol{f}(t, \boldsymbol{y}), \quad \boldsymbol{y}(t_0) = \boldsymbol{y}_0, \quad t_0 \le t \le T, \quad \boldsymbol{f} : \mathbb{R} \times \mathbb{R}^m \to \mathbb{R}^m$$

and by a *method* we shall mean a numerical method of the class defined by

$$\begin{aligned} \boldsymbol{y}_\mu &= \boldsymbol{\eta}_\mu(h), \qquad \mu = 0, 1, \cdots, k-1 \\ \sum_{j=0}^{k} \alpha_j \boldsymbol{y}_{n+j} &= h\boldsymbol{\phi}_f(t_n, \boldsymbol{y}_{n+k}, \boldsymbol{y}_{n+k-1}, \cdots, \boldsymbol{y}_n) \end{aligned}$$

where h is the steplength and k the stepnumber." Everything is here, in this single sentence, both the object of study and the tools for algorithmic development and mathematical analysis. Indeed, in many respects the computational work on this 'plain-vanilla ODE problem' was one of the great success stories of modern numerical mathematics. Guided by the seminal work of Germund Dahlquist, a whole generation of first-class numerical analysts have established a profound understanding of the subject on theoretical and practical levels alike. Arguably, computational ODE theory became the gold standard of the numerical analysis of differential equations and a subject for emulation by the more elusive theory of computational PDEs.

Thus, a decade ago matters appeared to settle down to the leisurely course of incremental development: speeding up a programme, tightening an error estimate.... As often, however, as soon as a subject seemingly reaches a plateau, something is bound to happen. Specifically, building on the aforementioned understanding, numerical community becomes equipped to pose new questions in a meaningful fashion. And, as soon as new answers arrive, it transpires that, beyond their mathematical or æsthetical virtues, they are of relevance to a range of applications. The plateau is abandoned for a climb up new and uncharted summits.

The critical step in abandoning the old paradigm is that by a 'problem' we no longer mean a single, standard ordinary differential system with an anonymous function $\boldsymbol{f}$. Different families of ODEs share certain qualitative attributes and they deserve special numerical treatment that does not trade these attributes away. Moreover, many problems that seemingly look like ODEs are in reality completely different mathematical constructs. Although their algorithmic analysis draws heavily on existing theory of computational ODEs, it poses different questions and seeks different answers.

Actually, our designation of the classical theory as dealing with just one kind of an object, a 'general ODE', is somewhat unfair, because of the long-standing distinction between *stiff* and *nonstiff* ODEs. Yet, it is a measure of the great advances of the last decade that stiff ODEs, once considered to be the most difficult computational nut to crack, are recognised at present as, basically, quite simple. The seeming difficulty of stiff ODEs, namely that trajectories corresponding to different initial values approach each other at an exponential speed, is ultimately also their redeeming grace. As long as the underlying equations are discretized by time-stepping algorithms with decent stability properties, bunching of trajectories means that numerical perturbations are bound to go away.

Two of the themes of this survey can be visualised as limiting cases of stiff ODEs. Firstly, in a well-defined sense, allowing the stiffness ratio to become infinite results in *differential-algebraic equations (DAEs),* the subject of Section 2. A good starting point for DAEs is the system

$$\begin{aligned} \boldsymbol{x}' &= \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{y}), \\ \boldsymbol{0} &= \boldsymbol{g}(\boldsymbol{x}, \boldsymbol{y}), \end{aligned} \qquad t \geq t_0, \tag{1.1}$$

which, as its name implies, combines elements of differential and algebraic equations.

Another limiting case of stiff ODEs occurs when, instead of decaying modes, one contemplates solutions that neither die out nor necessarily tend to a fixed point but evolve forever, whilst conserving certain invariants. Since such invariants often possess important physical interpretation or correspond to the outcome of qualitative mathematical analysis, there are obvious advantages to their retention under discretization. A familiar (and perhaps the most important) example is presented by Hamiltonian systems

$$\begin{aligned} \boldsymbol{p}' &= -\frac{\partial H(\boldsymbol{p}, \boldsymbol{q})}{\partial \boldsymbol{q}}, \\ \boldsymbol{q}' &= \frac{\partial H(\boldsymbol{p}, \boldsymbol{q})}{\partial \boldsymbol{p}}, \end{aligned} \qquad t \geq t_0, \tag{1.2}$$

which are surveyed at great depth elsewhere in this volume, in Chus Sanz-Serna's paper [27]. Another example (about which more later) is the ODE

$$Y' = B(Y)Y, \quad t \geq t_0, \qquad Y(t_0) = Y_0, \tag{1.3}$$

where the unknown Y is a $d \times d$ real matrix, the initial value Y_0 is an orthogonal matrix and the function B maps orthogonal matrices to skew-symmetric matrices. It is easy to prove that the invariant is *orthogonality:* the matrices $Y(t)$ are orthogonal for all $t \geq t_0$ [11, 19].

It is useful to express the retention of invariants in the language of differential manifolds. In Section 3 we review numerical methods on manifolds, a theme that has already led to significant new results and which is a subject of intensive and exciting current research.

Section 4 is devoted to yet other constructs which, while resembling standard ODEs, are quite different: *functional differential equations (FDEs).* A convenient starting point for their discussion is the equation

$$\boldsymbol{y}'(t) = \boldsymbol{f}(t, \boldsymbol{y}(t), \boldsymbol{y}(t-\tau)), \qquad t \geq 0, \tag{1.4}$$

given in tandem with the initial condition $\boldsymbol{y}(t) = \boldsymbol{\phi}(t)$, $-\tau < t \leq 0$.

A common thread running through the numerical analysis of 'nonclassical ODEs' (a misnomer, since many are not ODEs at all!) is that a convenient algorithmic point of departure is the familiar arsenal of standard numerical methods for the ODE

$$\boldsymbol{y}' = \boldsymbol{f}(t, \boldsymbol{y}), \quad t \geq t_0, \qquad \boldsymbol{y}(t_0) = \boldsymbol{y}_0,$$

namely *multistep* methods

$$\sum_{k=0}^{s} \rho_k \boldsymbol{y}_{n+k} = h \sum_{k=0}^{s} \sigma_k \boldsymbol{f}(t_{n+k}, \boldsymbol{y}_{n+k}), \qquad n \in \mathbb{Z}^+, \tag{1.5}$$

and *Runge–Kutta* schemes

$$\begin{array}{l} \left.\begin{array}{rcl} \boldsymbol{\varphi}_\ell &=& \boldsymbol{y}_n + h \displaystyle\sum_{j=1}^{\nu} a_{\ell,j} \boldsymbol{k}_j, \\ \boldsymbol{k}_\ell &=& \boldsymbol{f}(t_n + c_\ell h, \boldsymbol{\varphi}_\ell), \end{array}\right\} \quad \ell = 1, 2, \ldots, \nu, \\ \boldsymbol{y}_{n+1} = \boldsymbol{y}_n + h \displaystyle\sum_{\ell=1}^{\nu} b_\ell \boldsymbol{k}_\ell, \end{array} \qquad n \in \mathbb{Z}^+ \tag{1.6}$$

Classical methods can be occasionally extended fairly easily, e.g. by the provision of continuous output, although sometimes considerably more effort is required. Their analysis, however, is difficult and often calls for unfamiliar tools: differential topology, Lie group theory, symplectic geometry, harmonic analysis. . . .

Having quoted from Jack Lambert's survey in a spirit of mild criticism (an easy commodity when aided by hindsight), this introduction will not be complete without another quote from an earlier *State of the Art in Numerical Analysis* survey by the same author, which admirably (and presciently) sums up the new paradigm: "The initial value problem is a lively mathematical creature (compared with, for example, the linear algebraic problem); where there is structure, we ought to be able to use it in the numerical method." [21]

2 Differential-algebraic equations

DAEs occur naturally in a surprisingly large number of applications: optimal control, the behaviour of constrained mechanical systems (e.g. in robotics) and singular perturbation problems. Given their ubiquity, it is surprising that they have not received comprehensive attention in the numerical community until the last decade.

As often in an exposition of mathematical material, it is preferable to commence by considering the simplest possible model which is rich enough to expose the underlying structure. Thus, let us consider the equation (1.1), where $\boldsymbol{x} \in \mathbb{R}^{m_1}$, $\boldsymbol{y} \in \mathbb{R}^{m_2}$, the functions $\boldsymbol{f} : \mathbb{R}^{m_1+m_2} \to \mathbb{R}^{m_1}$ and $\boldsymbol{g} : \mathbb{R}^{m_1+m_2} \to \mathbb{R}^{m_2}$ are sufficiently smooth and the initial condition $(\boldsymbol{x}_0, \boldsymbol{y}_0)$ obeys the equation $\boldsymbol{g}(\boldsymbol{x}_0, \boldsymbol{y}_0) = \boldsymbol{0}$. In addition, we stipulate that the Jacobian matrix

$$\frac{\partial \boldsymbol{g}(\boldsymbol{x}_0, \boldsymbol{y}_0)}{\partial \boldsymbol{y}}$$

is nonsingular. It is then easy to prove, invoking the implicit function theorem, that there exists a function $\boldsymbol{q} : \boldsymbol{R}^{m_1} \to \boldsymbol{R}^{m_2}$, defined in a non-empty neighbourhood of $\boldsymbol{x}_0$, such that $\boldsymbol{y}(t) = \boldsymbol{q}(\boldsymbol{x}(t))$ for all $t \in [t_0, t^*]$. Substituting $\boldsymbol{q}$ into (1.1) we derive the *state space form* of the DAE,

$$\boldsymbol{x}' = \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{q}(\boldsymbol{x})), \qquad t \geq t_0. \tag{2.1}$$

A very fruitful interpretation of the DAE (1.1) is as a limiting case of a *singular perturbation problem*

$$\begin{aligned} \boldsymbol{x}' &= \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{y}), \\ \varepsilon \boldsymbol{y}' &= \boldsymbol{g}(\boldsymbol{x}, \boldsymbol{y}), \end{aligned} \qquad t \geq t_0, \tag{2.2}$$

as $\varepsilon \downarrow 0$. In particular, it allows us to extend standard ODE discretization methods to this framework. Matters are simple with regard to multistep schemes: we apply the method (1.5) to the singularly-perturbed problem (2.2), let $\varepsilon \downarrow 0$ and the outcome is

$$\begin{aligned} \sum_{k=0}^{s} \rho_k \boldsymbol{x}_{n+k} &= h \sum_{k=0}^{s} \sigma_k \boldsymbol{f}(\boldsymbol{x}_{n+k}, \boldsymbol{y}_{n+k}), \\ \boldsymbol{0} &= \sum_{k=0}^{s} \sigma_k \boldsymbol{g}(\boldsymbol{x}_{n+k}, \boldsymbol{y}_{n+k}), \end{aligned} \qquad n \in \mathbb{Z}^+. \tag{2.3}$$

We hasten to remark that this simplicity is misleading. As soon as we add stability considerations, simple reasoning singles out one family of multistep methods as a natural choice for the task in hand.

The generalization of a Runge–Kutta scheme (1.6) to DAEs requires more effort. For reasons that will become apparent soon, we commence by stipulating that the $\nu \times \nu$ *Runge–Kutta matrix* $A = (a_{i,j})$ is nonsingular and denote by $(\omega_{i,j})$ the components of A^{-1}. We let $\varepsilon \downarrow 0$ in the Runge–Kutta scheme

$$\left.\begin{aligned} \boldsymbol{\varphi}_\ell &= \boldsymbol{x}_n + h\sum_{j=1}^{\nu} a_{\ell,j}\boldsymbol{k}_j, \\ \varepsilon\boldsymbol{\psi}_\ell &= \varepsilon\boldsymbol{y}_n + h\sum_{j=1}^{\nu} a_{\ell,j}\boldsymbol{p}_j, \\ \boldsymbol{k}_\ell &= \boldsymbol{f}(\boldsymbol{\varphi}_\ell, \boldsymbol{\psi}_\ell) \\ \boldsymbol{p}_\ell &= \boldsymbol{g}(\boldsymbol{\varphi}_\ell, \boldsymbol{\psi}_\ell) \end{aligned}\right\} \qquad \ell = 1, 2, \ldots, \nu$$

$$\boldsymbol{x}_{n+1} = \boldsymbol{x}_n + h\sum_{\ell=1}^{\nu} b_\ell \boldsymbol{k}_\ell,$$

$$\varepsilon\boldsymbol{y}_{n+1} = \varepsilon\boldsymbol{y}_n + h\sum_{\ell=1}^{\nu} b_\ell \boldsymbol{p}_\ell,$$

whereby nonsingularity of A implies that

$$\boldsymbol{g}(\boldsymbol{\varphi}_\ell, \boldsymbol{\psi}_\ell) = \boldsymbol{0}, \qquad \ell = 1, 2, \ldots, \nu.$$

This, of course, is an eminently sensible outcome: the algebraic condition is satisfied exactly at each Runge–Kutta node. It also simplifies the scheme a great deal. Easy calculation affirms that it assumes the form

$$\begin{aligned} &\left.\begin{aligned} \boldsymbol{\varphi}_\ell &= \boldsymbol{x}_n + h\sum_{j=1}^{\nu} a_{\ell,j}\boldsymbol{f}(\boldsymbol{\varphi}_j, \boldsymbol{\psi}_j), \\ \boldsymbol{0} &= \boldsymbol{g}(\boldsymbol{\varphi}_\ell, \boldsymbol{\psi}_\ell), \end{aligned}\right\} \qquad \ell = 1, 2, \ldots, \nu, \\ &\boldsymbol{x}_{n+1} = \boldsymbol{x}_n + h\sum_{\ell=1}^{\nu} b_\ell \boldsymbol{k}_\ell, \\ &\boldsymbol{y}_{n+1} = \left(1 - \sum_{j,\ell=1}^{\nu} b_j\omega_{j,\ell}\right)\boldsymbol{y}_n + h\sum_{j,\ell=1}^{\nu} b_j\omega_{j,\ell}\boldsymbol{\psi}_\ell. \end{aligned} \tag{2.4}$$

Had this been all, numerical treatment of DAEs would have not been much more than a mildly singular perturbation of the standard computational ODE theory. That this is not the case becomes apparent upon the consideration of a scalar linear model of (2.2),

$$\begin{aligned} x' &= \alpha_{1,1}x + \alpha_{1,2}y, \\ \varepsilon y' &= \alpha_{2,1}x + \alpha_{2,2}y. \end{aligned} \tag{2.5}$$

This is also the place to mention another helpful interpretation of DAEs, namely as ODEs with *infinite stiffness.* Writing (2.5) in a vector form, $\boldsymbol{z}' = B\boldsymbol{z}$, we obtain the matrix

$$B = \begin{bmatrix} \alpha_{1,1} & \alpha_{1,2} \\ \varepsilon^{-1}\alpha_{2,1} & \varepsilon^{-1}\alpha_{2,2} \end{bmatrix}.$$

It is easy to verify that the *stiffness ratio* (the quotient between the eigenvalues of the largest and the least magnitude, the usual, albeit sometimes unreliable, measuring rod of stiffness [13, 18]) becomes infinite when $\varepsilon \downarrow 0$.

Consider again multistep methods. The standard condition for the ODE time-stepping scheme (1.5) to be convergent and of order $p \geq 1$ can be formulated in the language of the polynomials

$$\rho(w) = \sum_{k=0}^{s} \rho_k w^k \qquad \text{and} \qquad \sigma(w) = \sum_{k=0}^{s} \sigma_k w^k.$$

Thus, ρ must obey the *root condition* (all its zeros must reside in $|w| \leq 1$ and the zeros on $|w| = 1$ must be simple), whereby

$$\rho(w) - \log w \sigma(w) = c(w-1)^{p+1} + \mathcal{O}(|w-1|^{p+2}), \quad w \to 1, \qquad c \neq 0,$$

means that the method is of order p [12, 18]. However, to make any sense, we need to restrict the step size $h > 0$ so that $h\lambda$ lies in the *linear stability domain* $\mathcal{D}$ of the method for every $\lambda \in \sigma(B)$. Letting $\varepsilon \downarrow 0$ we conclude an additional necessary condition, specific to DAEs, namely that $\infty \in \mathcal{D}$. In that case it is possible to prove that

$$\boldsymbol{x}_n = \boldsymbol{x}(t_n) + \mathcal{O}(h^p), \quad \boldsymbol{y}_n = \boldsymbol{y}(t_n) + \mathcal{O}(h^p), \qquad h \downarrow 0,$$

uniformly for $t_n = t_0 + nh$ in a compact interval, provided that the starting values $\boldsymbol{x}_n, \boldsymbol{y}_n$, $n = 0, 1, \ldots, s$, are adequately chosen [13, 24].

The condition $\infty \in \mathcal{D}$ is tantamount to the requirement that all the zeros of σ reside in $|w| < 1$. The natural choice is placing all its zeros at the origin: the outcome is the familiar *backward differentiation formula (BDF).*

Order conditions for Runge–Kutta methods are, as always, more intricate. Insofar as the ODE method (1.6) is concerned, the reader is referred to [12] for graph-theoretical techniques which proved themselves useful in error analysis. In the case of DAEs there are two further considerations, both implied by the aforementioned interpretation of such systems as 'infinitely stiff'. We say that the method (1.6) has *stage order* $q \geq 1$ if each $\boldsymbol{\varphi}_\ell$ approximates $\boldsymbol{y}$ at $t_n + c_\ell h$ to order q. Moreover, we recall from the ODE theory that the linear stability domain of the Runge–Kutta method (1.6) is

$$\mathcal{D} = \{z \in \mathbb{C} \,:\, |r(z)| < 1\},$$

where

$$r(z) = 1 + z\boldsymbol{b}^T(I - zA)^{-1}\mathbf{1}, \qquad z \in \mathbb{C},$$

and $\boldsymbol{b}$ is the vector of the Runge–Kutta weights. Suppose that (1.6) is of order p (as an ODE method). Then the DEA method (2.4) is of order $\tilde{p}$, where

$$\begin{array}{ll} \tilde{p} = p & \text{if } r(\infty) = 0, \\ \tilde{p} = \min\{p, q+1\} & \text{if } r(\infty) \in [-1,1) \setminus \{0\}, \\ \tilde{p} = \min\{p-1, q\} & \text{if } r(\infty) = 1. \end{array}$$

If $|r(\infty)| > 1$ then the method diverges.

Simple algebra readily verifies that $a_{\nu,j} = b_j$, $j = 1, 2, \ldots, \nu$, implies $r(\infty) = 0$. Such *stiffly accurate* Runge–Kutta schemes possess obvious appeal insofar as DAEs (or very stiff ODEs) are concerned [13].

The analysis of general DAEs is more intricate and, as will be soon apparent, it imposes further restrictions on numerical methods. A general algebraic-differential equation can be written in the form

$$\boldsymbol{p}(\boldsymbol{u}, \boldsymbol{u}') = \boldsymbol{0}. \tag{2.6}$$

We say that (2.6) is of *differential index m* if $m \geq 0$ is the least integer such that, perhaps after some algebraic manipulation,

$$\begin{aligned} \boldsymbol{p}(\boldsymbol{u}, \boldsymbol{u}') &= \boldsymbol{0}, \\ \frac{\mathrm{d}}{\mathrm{d}t}\boldsymbol{p}(\boldsymbol{u}, \boldsymbol{u}') &= \boldsymbol{0}, \\ &\vdots \\ \frac{\mathrm{d}^m}{\mathrm{d}t^m}\boldsymbol{p}(\boldsymbol{u}, \boldsymbol{u}') &= \boldsymbol{0} \end{aligned}$$

is a conventional ODE system.

To illustrate this important concept let us reexamine the DAE (1.1). Differentiating once the second equation, we have

$$\boldsymbol{0} = \frac{\partial \boldsymbol{g}}{\partial \boldsymbol{x}}\boldsymbol{x}' + \frac{\partial \boldsymbol{g}}{\partial \boldsymbol{y}}\boldsymbol{y}',$$

therefore

$$\boldsymbol{y}' = -\left[\frac{\partial \boldsymbol{g}}{\partial \boldsymbol{y}}\right]^{-1} \frac{\partial \boldsymbol{g}}{\partial \boldsymbol{x}}\boldsymbol{f}$$

(recall that we have stipulated nonsingularity of the Jacobian matrix) and substitution into the derivative of the first equation yields

$$\boldsymbol{x}'' = \left(\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{x}} - \frac{\partial \boldsymbol{f}}{\partial \boldsymbol{y}}\left[\frac{\partial \boldsymbol{g}}{\partial \boldsymbol{y}}\right]^{-1} \frac{\partial \boldsymbol{g}}{\partial \boldsymbol{x}}\right)\boldsymbol{f}.$$

Hence the DAE (1.1) is of index $m = 1$.

An example of an index-two DAE is provided by

$$\begin{aligned} \boldsymbol{x}' &= \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{y}), \\ \boldsymbol{0} &= \boldsymbol{g}(\boldsymbol{x}), \end{aligned} \tag{2.7}$$

where we require that both Jacobian matrices

$$\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{y}} \quad \text{and} \quad \frac{\partial \boldsymbol{g}}{\partial \boldsymbol{x}}$$

are nonsingular. Differentiating once the second equation in (2.7), we obtain

$$\begin{aligned} \boldsymbol{x}' &= \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{y}), \\ \boldsymbol{0} &= \frac{\partial \boldsymbol{g}(\boldsymbol{x}, \boldsymbol{y})}{\partial \boldsymbol{x}} \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{y}), \end{aligned}$$

an index-one system. Hence (2.7) is of index two. This reduction to a lower-index system by differentiation is typical to index analysis.

It is left to the reader to verify that the DAE

$$\begin{aligned} \boldsymbol{x}' &= \boldsymbol{f}(\boldsymbol{x}, \boldsymbol{y}), \\ \boldsymbol{y}' &= \boldsymbol{q}(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{z}), \\ \boldsymbol{0} &= \boldsymbol{g}(\boldsymbol{x}), \end{aligned} \tag{2.8}$$

where the Jacobian matrices

$$\frac{\partial \boldsymbol{f}}{\partial \boldsymbol{y}}, \quad \frac{\partial \boldsymbol{q}}{\partial \boldsymbol{z}} \quad \text{and} \quad \frac{\partial \boldsymbol{g}}{\partial \boldsymbol{x}}$$

are all nonsingular, is of index three.

Numerical solution of high-index DAE systems (that is, with $m \geq 2$) is associated with a reduction in the order of the underlying ODE method. Although an in-depth consideration of this important phenomenon is well outside the scope of this survey, we present a simple illustration. The reader is referred to [2, 13, 24] for a more comprehensive exposition.

Let us consider the solution of the innocently-looking scalar, linear index-two system

$$x' + y = s_1(t), \qquad x = s_2(t)$$

by the BDF method

$$\sum_{k=0}^{s} \rho_k \boldsymbol{y}_{n+k} = h\sigma_s \boldsymbol{f}(t_{n+s}, \boldsymbol{y}_{n+s}).$$

The outcome is

$$\begin{aligned} \sum_{k=0}^{s} \rho_k x_{n+k} + h\sigma_s y_{n+s} &= h s_1(t_{n+s}), \\ x_{n+s} &= s_2(t_{n+s}). \end{aligned}$$

Substituting the second equation into the first, we solve explicitly for y_{n+s},

$$y_{n+s} = \frac{1}{\sigma_s}\left[s_1(t_{n+s}) - \frac{1}{h}\sum_{k=0}^{s} \rho_k s_2(t_{n+k}) \right].$$

Therefore the magnitude of the error is multiplied by $\mathcal{O}(h^{-1})$ and the order is $s-1$, rather than the nominal order s of the s-step BDF method for standard ODEs. This order reduction should cause no surprise since it originates in numerical differentiation.

The general rule is that the order is degraded by $m-1$ in numerical discretization of index-m DAEs.

3 Differential equations on manifolds

Numerical analysis is perhaps the most important tool in the investigation of differential equations. Having said this, many other techniques – functional analysis, nonlinear dynamical systems, perturbation theory, differential geometry, degree theory, Lie group theory, special functions' theory… – can be used to a highly beneficial effect. A growing concensus is that there is no conflict between computational and 'qualitative' approaches and that all means are fair. Indeed, harnessing computation in tandem with mathematical analysis is bound to lead to synergy and transform both for the better [31].

Qualitative information about a differential system being known, it is often worthwhile to incorporate it into the discretized solution. This has not just the obvious virtue of retaining important features of the solution under discretization but it also frequently leads to better behaviour of the global error [28]. Qualitative information, though, comes in two basic varieties: asymptotic information and invariants. Insofar as the recovery of correct ODE asymptotics is concerned, the important developments of the last decade are covered elsewhere in this volume by Andrew Stuart [29] and we shall not dwell upon them in this article.

Let us suppose that the initial-value ODE system

$$\boldsymbol{y}' = \boldsymbol{f}(t,\boldsymbol{y}), \quad t \geq t_0, \qquad \boldsymbol{y}(t_0) = \boldsymbol{y}_0 \in \mathbb{R}^d, \tag{3.1}$$

possesses an invariant. In other words, there exists a function $\rho : \mathbb{R}^d \to \mathbb{R}$ such that $\rho(\boldsymbol{y}(t)) \equiv \rho(\boldsymbol{y}_0)$, $t \geq t_0$. (The problem (3.1) might make sense only for initial values in a subset of $\mathbb{R}^d$, requiring a minor and obvious amendment to our argument.) In the language of differential topology we might say that the ODE (3.1) induces a foliation of $\mathbb{R}^d$ into *differentiable manifolds* $\mathcal{M}(\boldsymbol{v})$, $\boldsymbol{v} \in \mathbb{R}^d$, such that $\boldsymbol{y}(t) \in \mathcal{M}(\boldsymbol{y}_0)$, $t \geq t_0$. Letting

$$\mathcal{M} := \{\mathcal{M}(\boldsymbol{v}) \,:\, \boldsymbol{v} \in \mathbb{R}^d\},$$

we say that (3.1) is $\mathcal{M}$-*invariant.* Without any danger of confusion, we adopt similar terminology with regard to numerical methods: a discretization of the $\mathcal{M}$-invariant ODE (3.1) is said to be $\mathcal{M}$-invariant if $\boldsymbol{y}_n \in \mathcal{M}(\boldsymbol{y}_0)$ for all $n \in \mathbb{Z}^+$ and *all* $\mathcal{M}$-invariant ODEs (3.1).

Classical ODE methods are doing very poorly indeed when it comes to $\mathcal{M}$-invariance. Before formulating this statement in a more precise form, we remind the reader that each differentiable manifold is a level set of some C^1 function [14] and that the Runge–Kutta method (1.6) is said to be *nonconfluent* if $c_k \neq c_\ell$ for every $k \neq \ell$.

Theorem 1. (The Cheesecutter Theorem) *Suppose that there exists a two-dimensional 'slice' $\mathcal{K}(\boldsymbol{v})$ through the manifold $\mathcal{M}(\boldsymbol{v})$. Then $\mathcal{K}(\boldsymbol{v})$ is itself a manifold, embedded in $\mathbb{R}^2$. Let it be the level set of a function $\varrho : C^3[\mathbb{R}^2 \to \mathbb{R}]$ and suppose that there exists $\boldsymbol{v} \in \mathbb{R}^d$ such that ϱ does* not *obey the partial differential equation*

$$\varrho_y^3 \varrho_{xxx} - 3\varrho_x \varrho_y^2 \varrho_{xxy} + 3\varrho_x^2 \varrho_y \varrho_{xyy} - \varrho_x^3 \varrho_{yyy} = 0. \tag{3.2}$$

Then no nonconfluent Runge–Kutta method may be $\mathcal{M}$-invariant [7].

The situation is, if at all, even worse with regard to multistep methods. In greater generality, suppose that the function $\boldsymbol{f}$ is analytic. In that case we can repeatedly differentiate (3.1) to obtain expressions for progressively higher derivatives of $\boldsymbol{y}$, e.g.

$$\boldsymbol{y}'' = \frac{\partial \boldsymbol{f}(t,\boldsymbol{y})}{\partial t} + \frac{\partial \boldsymbol{f}(t,\boldsymbol{y})}{\partial \boldsymbol{y}} \boldsymbol{f}(t,\boldsymbol{y}).$$

In general, we have $\boldsymbol{y}^{(r)} = \boldsymbol{g}_r(t,\boldsymbol{y})$, $r \in \mathbb{Z}^+$. Many ODE methods, inclusive of multistep and multiderivative schemes (but, in general, no Runge–Kutta!), can be expanded (at least for some starting values) in the form

$$\boldsymbol{y}_{n+1} = \sum_{r=0}^{\infty} c_r h^r \boldsymbol{g}_r(t_n, \boldsymbol{y}_n), \tag{3.3}$$

where the coefficients $\{c_r\}$ depend only on the underlying method.

Theorem 2. *Let $\mathcal{M}(\boldsymbol{v})$ be an arbitrary analytic nonlinear manifold. The only $\mathcal{M}$-invariant choice of the coefficients $\{c_r\}$ in* (3.3) *is $c_r = 1/r!$, $r \in \mathbb{Z}^+$, i.e. the exact solution of the ODE* (3.1) *[7].*

The two theorems indicate that classical numerical methods are unlikely to retain invariants. The one exception is when the equation (3.2) is obeyed by every two-dimensional section of every manifold in the foliation. This takes place in the very important case of *quadratic* manifolds,

$$\mathcal{M}(\boldsymbol{v}) = \{\boldsymbol{x} \in \mathbb{R}^d \,:\, \boldsymbol{\beta}^T \boldsymbol{x} + \boldsymbol{x}^T \Gamma \boldsymbol{x} = \boldsymbol{\beta}^T \boldsymbol{v} + \boldsymbol{v}^T \Gamma \boldsymbol{v}\}$$

where $\boldsymbol{\beta} \in \mathbb{R}^d$ and Γ is a $d \times d$ real matrix.

Theorem 3. *The Runge–Kutta method* (1.6) *is $\mathcal{M}$-invariant for all quadratic manifolds $\mathcal{M}$ subject to the identity*

$$m_{k,\ell} := b_k a_{k,\ell} + b_\ell a_{\ell,k} - b_k b_\ell = 0, \qquad k,\ell = 1,2,\ldots,\nu$$

[9]. Moreover, if the above condition fails and the method is nonconfluent, it is possible to find an ODE which is invariant on such a manifold but for which invariance fails under discretization [5].

Specific cases of quadratic invariance are the retention of mechanical energy, conservation of orthogonality [11] and invariance on Stiefel and Grassmann manifolds. Yet, perhaps the most important instance of a quadratic manifold is the symplectic group, which represents the main invariant of Hamiltonian systems (1.2) [27, 28].

Special manifolds require special algorithmic remedies and this is true even for quadratic manifolds. Hamiltonian equations are a case in point but they are surveyed elsewhere in this volume [27]. Instead, let us examine briefly orthogonal flows (1.3). Any method (1.6) with $m_{k,\ell} \equiv 0$ retains invariance, but all such methods are implicit, hence expensive. A little bit of algorithmic ingenuity, however, can do wonders! Thus, suppose that (1.3) is solved by an arbitrary Runge–Kutta method of order $p \geq 1$. In that case the *orthogonality error* $e_n := \|Y_n^T Y_n - I\|$ will be also of order p. Together with the equation (1.3), we consider the *adjoint equation*

$$Z' = -ZB(Z^T), \quad t \geq t_0, \qquad Z(t_0) = Y_0^{-1}. \tag{3.4}$$

Clearly, the exact solution of (3.4) is $Z(t) = Y^T(t)$. As long as we are employing an *orthogonal* Runge–Kutta method, there is no difference between using (1.3) or (3.4). This, however, is no longer the case for arbitrary methods. Thus, suppose that in each *even* step we solve (1.3) and in every *odd* step we use the adjoint equation (3.4) (of course, taking appropriate initial conditions). In that case, although the order of the combined method remains p, it is possible to prove that $e_n = \mathcal{O}(h^{p+2})$, hence the orthogonality error is of order $p+1$ [6].

We can, as a matter of fact, do much better! Thus, let us designate each step with (1.3) as '0' and each step of (3.4) as '1'. No matter what the combination of zeros and ones, the order of the method remains p. Yet, the orthogonality error is

$00000000\cdots:$ $\quad p,$
$01010101\cdots:$ $\quad p+1,$
$01100110\cdots:$ $\quad p+2,$
$01101001\cdots:$ $\quad p+3.$

In general, suppose that, to advance from t_n to t_{n+1} (with constant step sizes) we use '0' if the number of ones in the binary expansion of n is even, '1' otherwise. Then (in exact arithmetic) $e_{2^r n} = \mathcal{O}(h^{p+r+1})$, $r, n \in \mathbb{Z}^+$, and we can make the orthogonality error arbitrarily small for very little extra cost.

Figure 1 displays the orthogonality error in solving an *orthogonal flow* (1.3) with forward Euler for different sequences of '0's and '1's. The attenuation of orthogonality error is self-evident.

Orthogonal flows (1.3) are interesting in their own right but perhaps their most fascinating application is implicit in their connection with *isospectral flows*

$$X' = F(X)X - XF(X), \quad t \geq t_0, \qquad X(t_0) = X_0. \tag{3.5}$$

Again, the unknown $X(t)$ is a $d \times d$ matrix, while the function F maps $d \times d$ matrices to themselves. The invariants of (3.5) are the eigenvalues of the matrix

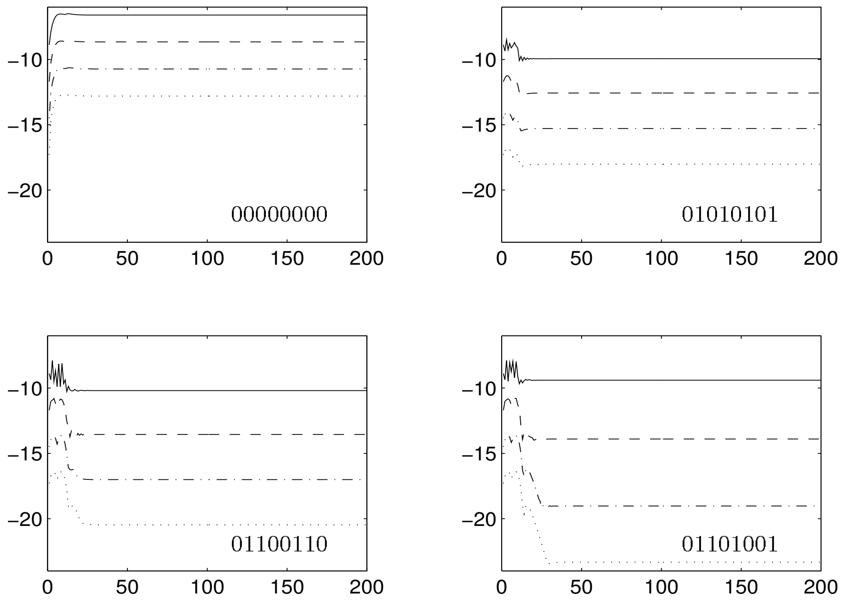

Figure 1. Orthogonality error (on a logarithmic scale) for the orthogonal flow, solved by the forward Euler method. Consecutive lines correspond to step-size refinement.

X_0 [30]. Isospectral flows are ubiquitous in many applications, e.g. in mechanics, quantum chemistry, micromagnetics and molecular dynamics, because of their connection with Toda lattice equations [30]. However, applications nearest to a numerical analyst's heart are in numerical algebra, *inter alia* in the determination of eigenvalues and in inverse eigenvalue problems [8]. In most such applications the matrix X_0 is symmetric and F maps $d \times d$ matrices into skew-symmetric matrices – this ensures that $X(t)$ is symmetric for all $t \geq t_0$. We henceforth stipulate this to be the case.

It is possible to prove, along the lines of Theorem 1, that no nonconfluent Runge–Kutta method may conserve eigenvalues for $d \geq 3$ [5]. However, retention of eigenvalues is possible by exploiting a link between (3.5) and orthogonal flows (1.3). Thus, letting

$$B(Y) := YX_0Y^{-1}, \qquad Y_0 = I,$$

it is easy to prove that $X(t) = Y(t)X_0Y^{-1}(t)$, where Y is the solution of (1.3) (incidentally, this proves that the eigenvalues are retained by X!). Subject to this reformulation, isospectral flows can be solved by quadratically-invariant Runge–Kutta methods [5].

Note that eigenvalues are retained no matter how (1.3) is solved. However, unless Y_n is orthogonal, X_n will not be symmetric. Similarly to the use of the adjoint equation to decrease the orthogonality error, we can resort to a similar device to solve isospectral flows while minimizing symmetry breakdown. This is

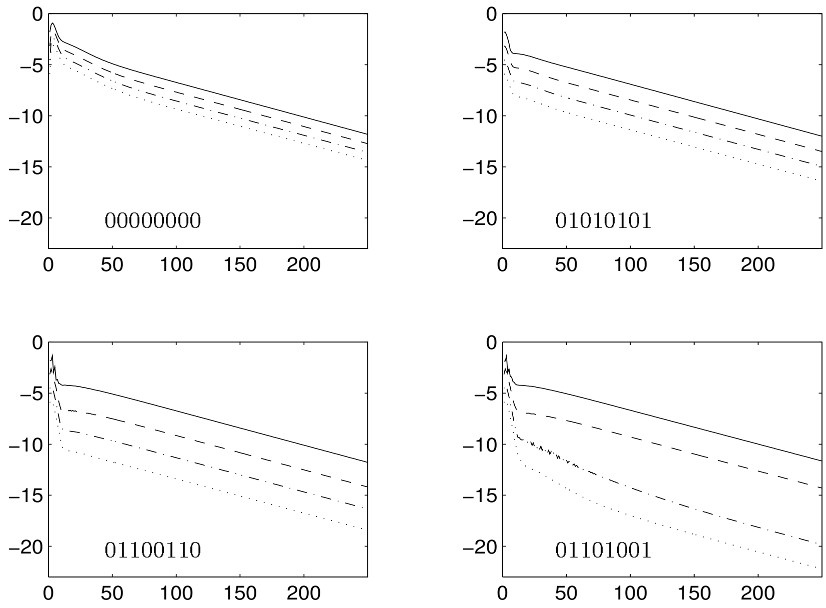

Figure 2. Symmetry error (on a logarithmic scale) for the isospectral flow, solved by the forward Euler method.

illustrated in Figure 2.

Simultaneously with the work on special families of ODEs that live on 'interesting' manifolds, there is a growing interest in new types of methods that are crafted to retain invariants of a general type. Although many of them are of a tentative and preliminary nature, it is perhaps worthwhile to mention emerging trends that are likely to inform future research. Four approaches have been considered in the last few years:

1. Perhaps the simplest and most natural technique is to project the numerical solution (originating in an arbitrary time-stepping solver) on the underlying manifold. Such *projective methods* have obvious connections with DAEs – indeed, one can interpret differential-algebraic equations as ODEs on manifolds. A conversion to DAEs, followed by, say, the use of BDF or of stiffly-accurate Runge–Kutta, constitutes a powerful approach to many ODEs with invariants [1]. Another projective approach, the *discrete gradient method,* is due to Reinout Quispel and his coworkers [26]. Projective methods, however, are more problematic when the underlying manifold cannot be easily described in a constructive fashion, e.g. in the case of isospectral flows.

2. Classical Runge–Kutta methods are formulated within the framework of Euclidean spaces. Peter Crouch and Robert Grossman have put forward an

alternative formulation, whereby derivatives are restricted to the *tangent space* of the underlying manifold [10]. Although this leads to exceedingly complicated order analysis (assisted a great deal by the recent interpretation of Butcher's graph theory-based order analysis in the terminology of Lie brackets by Hans Munthe-Kaas [25]), third and fourth-order methods have already been developed. The tangent space approach is very powerful because of its generality, although the latter is also its Achilles heel: in general, it is far from easy to produce numerically a good local basis of the tangent space. A special case whereby the tangent space is given for free in a constructive fashion, thereby rendering this approach very attractive, is presented by *Lie-type equations*

$$Y' = F(Y)Y, \quad t \geq t_0, \qquad Y(t_0) = Y_0, \tag{3.6}$$

where Y_0 lies in a Lie group $\mathfrak{A}$, while F maps $\mathfrak{A}$ into its Lie algebra $\mathfrak{a}$. In this case it is true that $Y(t) \in \mathfrak{A}$ for all $t \geq t_0$. (Of course, orthogonal matrices form the Lie group $\mathrm{O}(d, \mathbb{R})$, whose Lie algebra consists of $d \times d$ skew-symmetric matrices – orthogonal flows (1.3) are but an example of Lie-type equations (3.6).) In that case the tangent space is $\mathfrak{a}$ and above methods can be formulated with greater ease.

3. In a recent paper, Antonella Zanna has presented very interesting new methods that automatically stay on the Lie group $\mathfrak{A}$ whenever applied to the Lie-type equations (3.6) [31]. These methods are explicit and they can attain arbitrarily high orders (although high-order methods are expensive). They extensively exploit the closure of a Lie algebra $\mathfrak{a}$ under commutation and the fact that the exponential map takes $\mathfrak{a}$ to $\mathfrak{A}$.

4. Suppose that the ODE (3.1) is $\mathcal{M}$-invariant. The topology of the manifold $\mathcal{M}(\boldsymbol{y}_0)$ is often known. In particular, $\mathcal{M}(\boldsymbol{y}_0)$ can be often mapped diffeomorphically into another manifold $\mathcal{N}(\boldsymbol{y}_0)$ with simpler geometry. Provided that we know how to discretize invariantly on $\mathcal{N}(\boldsymbol{y}_0)$, we can subject the ODE to the same diffeomorphism, whereby it becomes $\mathcal{N}$-invariant, discretize it while retaining $\mathcal{N}$-invariance and map back. This simple idea has been exploited in [7] in a systematic fashion but many outstanding questions remain in its implementation.

4 Functional differential equations

A general functional differential equation is of the form

$$\boldsymbol{y}'(t) = \boldsymbol{f}(t, \boldsymbol{y}(t), \boldsymbol{y}(\vartheta(t))), \qquad t \geq 0, \tag{4.1}$$

with an initial condition given for $t \in (\inf_{\xi \geq 0} \vartheta(\xi), 0]$. Equations of this kind feature in a multitude of applications and, in author's opinion, would have probably been much more ubiquitous were their theory and computation based on a firmer footing.

It is deceptively simple to extend standard ODE methods, like (1.5) and (1.6), to an FDE setting. As long as the output is continuous (and this is always attainable, e.g. by polynomial interpolation), it can be exploited in an obvious fashion and *presto* an ODE method becomes an FDE method. We hasten to add that, needless to say, the actual algorithmic practice is *much* more complicated. The two obvious problems in a naive conversion of ODE to FDE methods are how to deal with derivative discontinuities (which are often present in the *exact* solution of (4.1)) and how to keep storage requirements at bay. However, the most demanding and least explored theme in the numerical analysis of FDEs is the question of stability.

Important strides have been made in the last decade with regard to two aspects of the stability theory of time-stepping methods for (4.1). Firstly, the linear theory for the constant-delay FDE (1.4) is at last reasonably well understood.

Much work has been published in the last few decades on the linear, scalar model equation

$$y'(t) = ay(t) + by(t-\tau), \qquad t \geq 0. \tag{4.2}$$

Unfortunately, all this has very little bearing on the vector equation

$$\boldsymbol{y}'(t) = P\boldsymbol{y}(t) + Q\boldsymbol{y}(t-\tau), \qquad t \geq 0. \tag{4.3}$$

In the ODE case one can generalize from scalar to vector linear equations since just a single matrix is present and standard Jordan factorization does the trick. However, unless A and B commute and the matrix pencil (A, B) can be factorized in unison, stability results for the equation (4.2) (whether exact or discretized) have little bearing on (4.3).

The understanding of the vector equation (4.3) in a numerical setting owes much to the work of Karel in 't Hout. First, however, we must state stability conditions for the exact equation. Standard FDE theory readily provides a condition for the asymptotic stability of the zero solution. Thus, let

$$\gamma(z, w) := \det[zI - (P + wQ)], \qquad z, w \in \mathbb{C}.$$

Then $\lim_{t\to\infty} \boldsymbol{y}(t) = \boldsymbol{0}$ for every initial value if and only if for every $z \in \mathbb{C}$

$$\gamma(z, \mathrm{e}^{-\tau z}) = 0 \qquad \Rightarrow \qquad \operatorname{Re} z < 0.$$

This, however, falls short of being a transparent asymptotic stability requirement, verifiable with minimal fuss from the matrices P and Q. Fortunately, Karel in 't Hout formulated asymptotic stability conditions in a more user-friendly fashion [15].

Theorem 4. *The zero solution of* (4.3) *is asymptotically stable for all initial conditions and for all delays* $\tau > 0$ *if and only if the following three conditions are satisfied.*

1. $\operatorname{Re} \lambda < 0$ *for all* $\lambda \in \sigma(P)$*;*

2. $\rho[(\mathrm{i}tI - P)^{-1}Q] < 1$ *for all* $t \in \mathbb{R} \setminus \{0\}$*;* *and*

3. $-1 \notin \sigma(P^{-1}Q)$.

We consider the solution of (4.3) by a *θ-method.* Applied to the standard ODE (3.1), such a method reads

$$\boldsymbol{y}_{n+1} = \boldsymbol{y}_n + h[\theta \boldsymbol{f}(t_{n+1}, \boldsymbol{y}_{n+1}) + (1-\theta)\boldsymbol{f}(t_n, \boldsymbol{y}_n)],$$

where $\theta \in [0,1]$ is a given constant. An extension of the method to the FDE (4.1) is

$$\boldsymbol{y}_{n+1} = \boldsymbol{y}_n + h[\theta \boldsymbol{f}(\boldsymbol{y}_{n+1}, \boldsymbol{v}_{n+1}) + (1-\theta)\boldsymbol{f}(\boldsymbol{y}_n, \boldsymbol{v}_n)],$$

where $\boldsymbol{v}_r$ is a linear interpolation to the numerical solution at $\vartheta(t_r)$. In the case of the linear system (4.3), letting the step h be commensurate with the delay τ, $h = \tau/m$, say, we obtain

$$(I - \theta h P)\boldsymbol{y}_{n+1} = [I + (1-\theta)hP]\boldsymbol{y}_n + hQ[\theta \boldsymbol{y}_{n-m+1} + (1-\theta)\boldsymbol{y}_{n-m}]. \tag{4.4}$$

The pencil (hP, hQ) is said to be *stable* if the matrix $(I - \theta hP)$ is invertible and $\lim_{n\to\infty} \boldsymbol{y}_n = \boldsymbol{0}$ in (4.4). Note that the linear stability domain of the θ-method (in the classical ODE context) is

$$\mathcal{D} = \left\{ z \in \mathbb{C} : \left| \frac{1 + (1-\theta)z}{1 - \theta z} \right| < 1 \right\}.$$

Theorem 5. *If*

$$h\sigma(P) \subset \mathcal{D} \qquad \textit{and} \qquad h \sup_{z \in \partial\mathcal{D}} \rho[(zI - hP)^{-1}Q] < 1$$

then (hP, hQ) *is stable. Moreover, if this pencil is stable then necessarily*

$$h\sigma(P) \subset \mathcal{D} \qquad \textit{and} \qquad h \sup_{z \in \partial\mathcal{D}} \rho[(zI - hP)^{-1}Q] \le 1.$$

Therefore, as in the ODE case, numerical stability follows by replacing the complex left half-plane with the linear stability domain $\mathcal{D}$ *[15].*

A θ-method is a simple instance of a two-stage Runge–Kutta, specifically,

$$\boldsymbol{c} = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \qquad A = \begin{bmatrix} 0 & 0 \\ \theta & 1-\theta \end{bmatrix}, \qquad \boldsymbol{b} = \begin{bmatrix} \theta \\ 1-\theta \end{bmatrix}.$$

Little wonder that in 't Hout's analysis has been already extended to more elaborate Runge–Kutta schemes [20].

A second development on the FDE front in the last decade is an initial foray into stability analysis of discretized FDEs with variable delays. The stability model is the *pantograph equation*

$$\boldsymbol{y}'(t) = P\boldsymbol{y}(t) + Q\boldsymbol{y}(qt), \quad t \ge 0, \qquad \boldsymbol{y}(0) = \boldsymbol{y}_0 \in \mathbb{C}^d, \tag{4.5}$$

where $q \in (0,1)$ is given. This equation is also of an independent interest in a number of applications [16].

The difference between constant and proportional delay – to wit, between (4.3) and (4.5) – cannot be greater. The pantograph equation (4.5) requires just a single initial condition and its solution resides in $C^\infty[\mathbb{C}^d]$. Its asymptotic behaviour is well-understood, yet highly nontrivial.

Theorem 6. *The zero solution of* (4.5) *is asymptotically stable if*

$$\operatorname{Re}\lambda < 0 \quad \forall \lambda \in \sigma(P) \qquad \textit{and} \qquad \rho(P^{-1}Q) < 1,$$

where $\rho(\,\cdot\,)$ is the spectral radius. The boundary of the asymptotic stability domain decomposes into three portions:

1. $\operatorname{Re}\lambda < 0$ *for all* $\lambda \in \sigma(P)$, $\rho(P^{-1}Q) = 1$: *Provided that the eigenvalues of $P^{-1}Q$ with unit modulus have the same algebraic and geometric multiplicity, there exists a* periodic *function $\boldsymbol{\phi}$ and for every $\varepsilon > 0$ there exists $t_\varepsilon > 0$ such that*
$$\|\boldsymbol{y}(t) - \boldsymbol{\phi}(q^{-t})\| < \varepsilon, \qquad t \ge t_\varepsilon.$$
Otherwise
$$\limsup_{t\to\infty} \|\boldsymbol{y}(t)\| = \infty \tag{4.6}$$
for some initial values.

2. $\max\{\operatorname{Re}\lambda \,:\, \lambda \in \sigma(P)\} = 0$, *$P$ is nonsingular and* $\rho(P^{-1}Q) < 1$: *The solution $\boldsymbol{y}$ tends to an* almost-periodic *function $\boldsymbol{\psi}$: for every $\varepsilon > 0$ there exists $s_\varepsilon > 0$ such that*
$$\|\boldsymbol{\psi}(t + s_\varepsilon) - \boldsymbol{\psi}(t)\| < \varepsilon, \qquad t \ge 0,$$
and $\liminf_{\varepsilon\downarrow 0} s_\varepsilon > 0$.

3. *P is singular: In that case* (4.6) *holds for some initial values and the zero solution is unstable [16, 23].*

Although stability of the discretized pantograph can be analysed by traditional techniques [4], this is inadequate. All such techniques are based, often implicitly, on Liapunov functions or similar estimates. Therefore, they are at their best if, at least in an asymptotic sense, the norm of the solution decays monotonically. This, however, is not always the case with regard to the pantograph and the equation calls for a more subtle approach.

Although the question of discretized pantograph stability for general methods remains wide-open, there already exists a body of interesting results about scalar equations and $q = 1/r$, where $r \ge 2$ is an integer. Let us restrict this brief debate to $q = \frac{1}{2}$, whereby, without loss of generality, the scalar pantograph becomes

$$y'(t) = ay(t) + by(\tfrac{1}{2}t), \quad t \ge 0, \qquad y(0) = 1. \tag{4.7}$$

It is easy to specialize Theorem 6 to the equation (4.7): the zero solution is asymptotically stable if and only if

$$\operatorname{Re} a < 0, \qquad |b| < |a|, \tag{4.8}$$

the function y tends (under a change of variable) to a periodic curve when $\operatorname{Re} a < 0$ and $|b| = |a|$ and it is almost-periodic when $\operatorname{Re} a = 0$, $a \neq 0$ and $|b| < |a|$. It is not difficult to verify that, as long as

$$\operatorname{Re} a + |b| < 0, \tag{4.9}$$

the function $|y(t)|$ tends to zero monotonically. This opens up the way to Liapunov function-based techniques and discretized stability analysis is relatively straightforward [4]. The set (4.9), however, is in general much smaller than the set (4.8).

Solving (4.7) with the familiar *trapezoidal rule* (a θ-method with $\theta = \frac{1}{2}$) results in the nonstationary recurrence

$$\begin{aligned} y_0 &= 1, \\ \left.\begin{aligned} y_{2n+1} &= r y_{2n} + \tfrac{1}{4}s(3y_n + y_{n+1}), \\ y_{2n+2} &= r y_{2n+1} + \tfrac{1}{4}s(y_n + 3y_{n+1}) \end{aligned}\right\} &\qquad n \in \mathbb{Z}^+, \end{aligned} \tag{4.10}$$

where

$$r = \frac{1 + \frac{1}{2}ha}{1 - \frac{1}{2}ha} \qquad \text{and} \qquad s = \frac{hb}{1 - \frac{1}{2}ha}.$$

To analyse the stability of (4.10), Martin Buhmann and the author proved that the Fourier transform $\hat{x}$ of the sequence $x_n = y_{n+1} - y_n$, $n \in \mathbb{Z}^+$, obeys the equation

$$\hat{x}(\xi) = \frac{1}{1 - re^{i\xi}} \left[(r + s - 1) + \tfrac{1}{4}s(1 + e^{i\xi})^2 \hat{x}(2\xi)\right] \pmod{2\pi}, \qquad \xi \in [0, 2\pi].$$

This, together with an application of the *mean ergodic theorem,* is the main step in proving that $\hat{x}$ is an L_2 function, therefore $\{x_n\}_{n\in\mathbb{Z}^+}$ is an ℓ_2 sequence, provided that (4.8) holds. A simple technical argument can be used to prove that also $\{y_n\}_{n\in\mathbb{Z}^+}$ is ℓ_2, whence $\lim_{n\to\infty} y_n = 0$ and the zero solution is asymptotically stable. Therefore the trapezoidal rule recovers faithfully the asymptotic stability domain of the exact pantograph equation (4.7).

The behaviour of (4.10) on the stability boundary is, if at all, even more complicated. Figure 3 displays the sequence $\{y_n\}_{n\in\mathbb{Z}^+}$ for two different choices of $a \in i\mathbb{R} \setminus \{0\}$ and $|b| < |a|$ in the 'phase plane' $(\operatorname{Re} y, \operatorname{Im} y)$. Both figures are suggestive of fractal sets – rightly so! It has been proved in [17] that, at least in specific simplified cases, the attractor of the underlying dynamical system is equivalent to a *probabilistic mixture of Julia sets* and is of a fractal Hausdorff dimension.

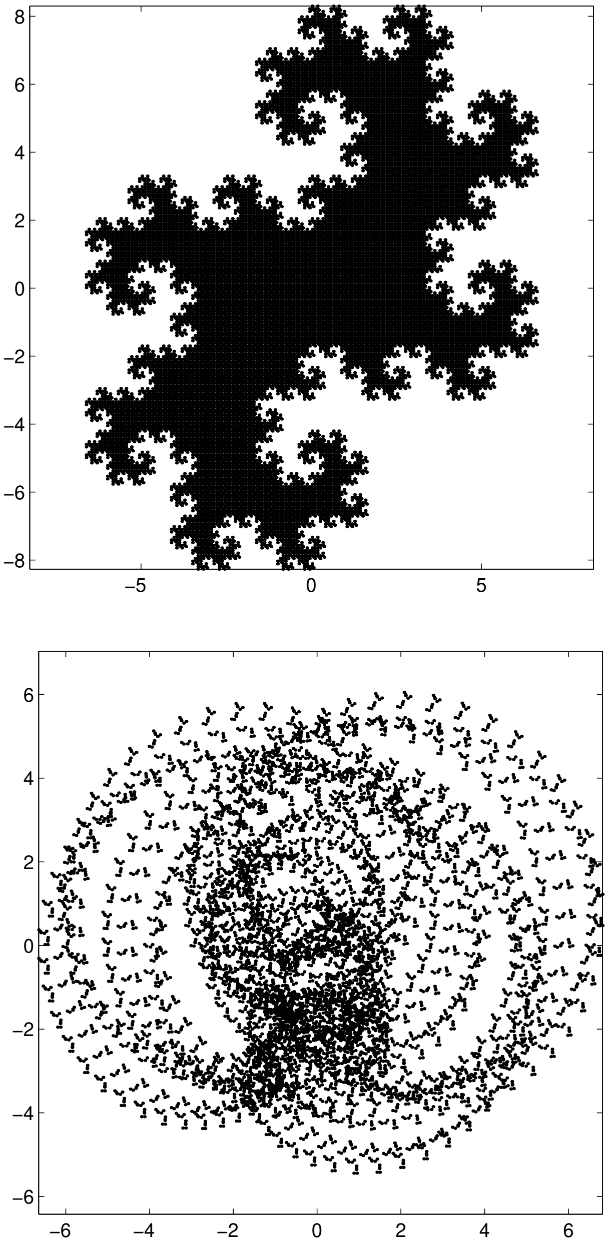

Figure 3. Examples of fractal attractors on the stability boundary of the discretized pantograph equation.

The pantograph equation (4.5) and its generalizations into neutral and integro-differential settings provide an enduring analytic and numerical challenge. And this is, in a sense, the simplest FDE with an unbounded delay! Having said this, the research of last decade has laid here, like elsewhere for different themes of this paper, firm mathematical foundations. It is up to the numerical community to take up the challenge. Many new and exciting research problems are presented by nonclassical breeds of ODEs: differential-algebraic equations, differential equations on manifolds, functional-differential equations.... Underlying problems are difficult and they often call for new tools and profound mathematics. Yet, the lesson of the last decade is that, difficulties notwithstanding, these problems are within reach of modern numerical analysis.

Bibliography

1. Ascher, U.M. (1996). Stabilization of invariants of discretized differential systems. Univ. of British Columbia Tech. Rep.

2. Brennan, K.E., Campbell, S.L. and Petzold, L.R. (1989). *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations.* North Holland, New York.

3. Buhmann, M.D. and Iserles, A. (1992). On the dynamics of a discretized neutral equation. *IMA J. Num. Anal. 12,* 339–363.

4. Buhmann, M.D. and Iserles, A. (1993). Stability of the discretized pantograph differential equation. *Maths Comp. 60,* 575–589.

5. Calvo, M.P., Iserles, A. and Zanna, A. (1995). Numerical solution of isospectral flows. DAMTP Tech. Rep. 1995/NA3, University of Cambridge.

6. Calvo, M.P., Iserles, A. and Zanna, A. (1996) Runge–Kutta methods for orthogonal and isospectral flows. To appear in *Appl. Num. Maths.*

7. Calvo, M.P., Iserles, A. and Zanna, A. (1996). Numerical methods on differentiable manifolds. In preparation.

8. Chu, M.T. (1994). A list of matrix flows with applications. In *Hamiltonian and Gradients Flows, Algorithms and Control* (A. Bloch, ed.), Fields Institute Communications, American Math. Soc., 87–97.

9. Cooper, G. (1987). Stability of Runge–Kutta methods for trajectory problems. *IMA J. Num. Anal. 7,* 1–13.

10. Crouch, P. and Grossman, R. (1993). Numerical integration of ordinary differential equations on manifolds, *J. Nonlinear Sci. 3,* 1–33.

11. Dieci, L., Russell, R.D. and van Vleck, E.S. (1994). Unitary integrators and applications to continuous orthonormalization techniques. *SIAM J. Numer. Anal. 31,* 261–281.

12. Hairer, E., Nørsett, S.P. and Wanner, G. (1991). *Solving Ordinary Differential Equations I: Nonstiff Problems* (2nd ed.). Springer-Verlag, Berlin.

13. Hairer, E. and Wanner, G. (1991). *Solving Ordinary Differential Equations II. Stiff and Differential-Algebraic Problems.* Springer-Verlag, Berlin.

14. Guillemin, V. and Pollack, A. (1974). *Differential Topology.* Prentice-Hall, Englewood Cliffs.

15. in 't Hout, K.J. (1994). The stability of θ-methods for systems of delay differential equations. *Annals of Numer. Maths 1,* 323–334.

16. Iserles, A. (1993). On the generalized pantograph functional-differential equation. *Europ. J. Appl. Maths 4,* 1–38.

17. Iserles, A. (1994). The asymptotic behaviour of certain difference equations with proportional delay. *Computers Math. Applics 28,* 141–152.

18. Iserles, A. (1996). *A First Course in the Numerical Analysis of Differential Equations.* Cambridge University Press, Cambridge.

19. Iserles, A. and Zanna, A. (1996). Qualitative numerical analysis of ordinary differential equations. To appear in *Lectures in Applied Mathematics* (J. Renegar, M. Shub and S. Smale, eds). American Mathematical Society, Providence RI.

20. Koto, T. (1994). A stability property of A-stable natural Runge–Kutta methods for systems of delay differential equations. *BIT 34,* 262–267.

21. Lambert, J.D. (1977). The initial value problem for ordinary differential equations. In *State of the Art in Numerical Analysis* (D.A.H. Jacobs, ed.), Academic Press, London, 451–500.

22. Lambert, J.D. (1987). Developments in stability theory for ordinary differential equations. In *State of the Art in Numerical Analysis* (A. Iserles and M.J.D. Powell, eds), Oxford University Press, Oxford, 409–431.

23. Liu, Y. (1996). On functional differential equations with proportional delays. Ph.D. dissertation. University of Cambridge.

24. März, R. (1992). Numerical methods for differential algebraic equations. *Acta Numerica 1,* 141–198.

25. Munthe-Kaas, H. (1995). Lie–Butcher theory for Runge–Kutta methods. *BIT 35,* 572–587.

26. Quispel, G.R.W. and Capel, H.W. (1996). Solving ODE's numerically while preserving a first integral. University of La Trobe Technical Report.

27. Sanz-Serna, J.M. (1996). Geometric integration. This volume.

28. Sanz-Serna, J.M. and Calvo, M.P. (1994), *Numerical Hamiltonian Problems.* Chapman & Hall, London.

29. Stuart, A.M. (1996). Convergence and stability in the numerical approximation of dynamical systems. This volume.

30. Toda, M. (1981). *Theory of Nonlinear Lattices.* Springer-Verlag, Berlin.

31. Zanna, A. (1996). The method of iterated commutators for ordinary differential equations on Lie groups. In preparation.