

APNUM 435

Shadows, chaos, and saddles

J.M. Sanz-Serna

Departamento de Matemática Aplicada y Computación, Facultad de Ciencias, Universidad de Valladolid, Valladolid, Spain

S. Larsson

Department of Mathematics, Chalmers University of Technology, S-412 96 Göteborg, Sweden

Received 21 September 1992

Revised 9 October 1992

Accepted 9 October 1992

Abstract

Sanz-Serna, J.M. and S. Larsson, Shadows, chaos, and saddles, Applied Numerical Mathematics 13 (1993) 181–190.

Shadowing provides a means for studying the behaviour of numerical methods in long-time integrations. We review the use of shadowing in the simulation of chaos and in the analysis of numerical methods near hyperbolic equilibrium points, such as saddles.

Keywords. Numerical methods; chaotic dynamics; shadowing; error estimation.

1. Introduction

Numerical methods are often used to investigate the *long-time* behaviour of solutions of ordinary or partial differential equations. Similarly mappings such as the well-known logistic map (see e.g. [11, Chapter 5])

$$u_{n+1} = \mu u_n(1 - u_n), \quad 0 \leq u_n \leq 1, \quad (1.1)$$

with μ a parameter in the range $0 \leq \mu \leq 4$, are often iterated many times in the computer in order to find the long-time behaviour of the corresponding orbits u_0, u_1, u_2, \dots . In both instances, i.e. both for continuous and for discrete dynamical systems, situations where the solutions of the system diverge exponentially from each other imply a catastrophic propagation of truncation and/or roundoff errors. As a consequence the global error (the distance between the computed and the true result of the initial value problem being simulated) may be enormous. *Shadowing* provides a means for the analysis of such situations. In this expository paper we review the use of shadowing in the study of chaotic iterations such as (1.1) (Section 4)

Correspondence to: J.M. Sanz-Serna, Departamento de Matemática Aplicada y Computación, Facultad de Ciencias, Universidad de Valladolid, Valladolid, Spain. E-mail: sanzserna@cpd.uva.es.

and in the analysis of numerical integrators in the neighbourhood of hyperbolic equilibrium points, such as saddles (Section 5). Sections 2 and 3 are respectively devoted to the presentation of the idea of shadowing and to the shadowing lemma. The final Section 6 contains some concluding remarks.

2. Shadowing: the basic idea

Let us begin with a very simple example. Consider the scalar initial value problem

$$dx/dt = -x, \quad 0 \leq t \leq T_{\max}, \quad x(0) = 1, \quad (2.1)$$

that we integrate by Euler's method

$$x_{n+1} = \psi x_n := (1 - \Delta t)x_n, \quad 0 \leq n \leq [T_{\max}/\Delta t], \quad x_0 = 1.$$

Figure 1 depicts the exact trajectory (solid line) and the computed points (stars) when $T_{\max} = 5$ and $\Delta t = 0.1$. The largest absolute value of the errors $e_n := x(n\Delta t) - x_n$ is 0.019 and occurs at $n = 10$ (i.e. at $t = 1$). Note that $|e_n|$ does not grow larger and larger as n increases. The reason for this is quite simple, see e.g. the venerable textbook by Isaacson and Keller [16, Chapter 8, Section 1, Corollary 2]. The error e_n can be written as

$$\begin{aligned} e_n &= x(n\Delta t) - x_n \\ &= [x(n\Delta t) - \psi x((n-1)\Delta t)] + \psi x((n-1)\Delta t) - \psi x_{n-1}, \end{aligned}$$

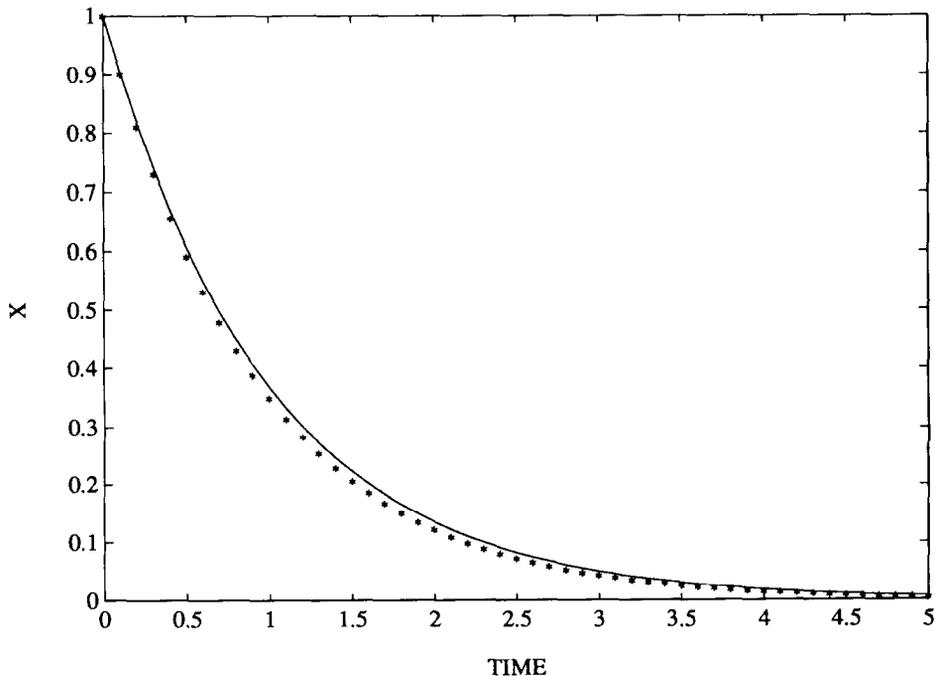


Fig. 1.

so that with $l_n := x(n\Delta t) - \psi x((n - 1)\Delta t)$ we have

$$e_n = \psi e_{n-1} + l_n. \tag{2.2}$$

The truncation error l_n represents the error in x_n had the approximation x_{n-1} at the previous time-level coincided with the exact $x((n - 1)\Delta t)$; the term $\psi e_{n-1} = (1 - \Delta t)e_{n-1}$ reflects the effect on e_n of the preceding error e_{n-1} . Since $|\psi| < 1$, there is a damping mechanism in the propagation of errors. From (2.2) we can write

$$|e_n| \leq |\psi| |e_{n-1}| + l,$$

where l is an upper bound for $|l_n|$. Iteration of this bound down to $n = 0$ results in

$$\begin{aligned} |e_n| &\leq |\psi| |e_{n-1}| + l \leq |\psi|^2 |e_{n-2}| + (|\psi| + 1)l \\ &\leq \dots \leq f \left(\sum_{i=0}^{n-1} |\psi|^i \right) l \leq l/\Delta t. \end{aligned}$$

Now $l_n = \frac{1}{2}\Delta t^2 \, d^2x(\xi_n)/dt^2$, with ξ_n an intermediate point, so that, for (2.1), $|l_n|$ can be bounded *independently* of $n \geq 0$ by $\frac{1}{2}\Delta t^2 |x(0)|$. Hence we finally conclude

$$|e_n| \leq \frac{1}{2}\Delta t |x(0)|, \quad n \geq 0. \tag{2.3}$$

For the case in Fig. 1 the bound in (2.3) is 0.050 and the maximum actual error 0.019.

We have thus obtained an error bound that holds uniformly for positive times t ; this has been possible because of two factors. First, the solutions of the equation $dx/dt = -x$ approach each other as $t \rightarrow \infty$; this contractive behaviour causes $|\psi|$ to be less than 1 for Δt small. Secondly, the local truncation error can be bounded uniformly for all positive t . Such long-time error estimates are studied in Stetter’s book [24, Chapters 3.5 and 4.6] for ordinary differential equations. For partial differential equations, Sanz-Serna and Stuart [22] discuss the situation for smooth solutions and Larsson [17] the more difficult case of nonsmooth solutions.

Let us now move to the unstable problem

$$dy/dt = y, \quad 0 \leq t \leq T_{\max}, \quad y(0) = 1, \tag{2.4}$$

again integrated with Euler’s method. Figure 2 shows the exact trajectory (solid line) and computed points (stars) when, as in Fig. 1, $T_{\max} = 5$ and $\Delta t = 0.1$. Now errors grow monotonically with n . The largest error 31.0 occurs at T_{\max} , where $y(5) = 148.4$ and $y_{50} = 117.4$. This is an accuracy of only 21%. Again this is easily explained. The errors e_n still satisfy the error recursion (2.2), but now $\psi = 1 + \Delta t$, leading to an exponential growth with n .

Shadowing provides a means to put to good use the bad numerical results in Fig. 2. Consider the solution $\tilde{y}(t)$ of $dy/dt = y$ that at $t = T_{\max}$ takes the same value as the computed numerical solution $y_{50} = 117.4$. In Fig. 2, $\tilde{y}(t)$ is given by a dashed line. The agreement between the computed points and $\tilde{y}(t)$ is quite good; now the maximum value $|\tilde{y}(n\Delta t) - y_n|$ is 2.1, more than an order of magnitude better than the maximum error $|y(n\Delta t) - y_n|$ which is 31.0. We say that the trajectory $\tilde{y}(t)$ *shadows* the numerical solution with an accuracy of 2.1. The conclusion is that the computed points y_n are poor approximations to the solution of the problem (2.4) which we integrated numerically; however they are good approximations to the shadowing solution: a solution of the correct differential equation with a “wrong” initial condition.

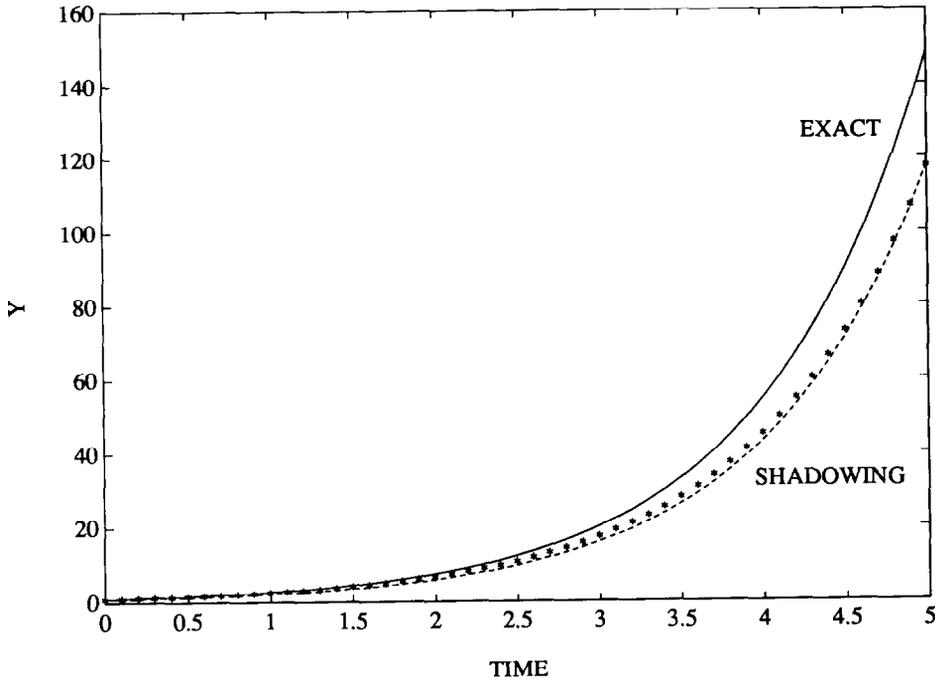


Fig. 2.

In the scenario studied in numerical analysis textbooks it is assumed that when integrating problems such as (2.4) the goal is to accurately compute the value of $y(t)$, $0 \leq t \leq T_{\max}$. Due to the exponential growth of the errors, this goal is not realistic when T_{\max} is not small. The idea of shadowing does not help in this scenario. On the other hand, a more realistic scenario may be envisaged. We could consider that the interest lies in finding the behaviour of the trajectories of the differential system being integrated. In this alternative scenario shadowing is certainly of use: in Fig. 2 we know that the computed points represent, except for a small error of 2.1, a *suitable* trajectory of the differential equation.

Before proceeding any further, let us examine why in Fig. 2 the errors $\tilde{e}_n := \tilde{y}(n\Delta t) - y_n$ turn out to be small. Instead of (2.2) we now have the error recursion

$$\tilde{e}_n = \psi \tilde{e}_{n-1} + \tilde{l}_n,$$

where $\tilde{l}_n := \tilde{y}(n\Delta t) - \psi \tilde{y}((n-1)\Delta t)$ and $\psi = 1 + \Delta t$. This formula is now used iteratively for n increasing up to $N = \lceil T_{\max}/\Delta t \rceil$. The result is

$$\begin{aligned} |\tilde{e}_n| &\leq |\psi^{-1}| (|\tilde{e}_{n+1}| + \tilde{l}) \leq |\psi^{-1}|^2 (|\tilde{e}_{n+2}| + \tilde{l}) + |\psi^{-1}| \tilde{l} \\ &\leq \dots \leq \left(\sum_{i=1}^{N-n} |\psi^{-1}|^i \right) \tilde{l} \leq (1 + \Delta t) \tilde{l} / \Delta t. \end{aligned}$$

After estimating \tilde{l} we conclude

$$|\tilde{e}_n| \leq \frac{1}{2} \Delta t (1 + \Delta t) |\tilde{y}(T_{\max})|, \quad n \leq N. \tag{2.5}$$

Note that now the error bound holds even for n negative and of arbitrarily large modulus. For Fig. 2 the bound in (2.5) equals 6.5; the maximum actual error (with respect to the shadowing trajectory) turned out to be 2.1. The main issue is that when estimating the errors e_n with respect to the true trajectory it is necessary to work with $|\psi| > 1$, while when estimating \tilde{e}_n it is the factor $|\psi^{-1}| < 1$ that appears. On the other hand note that the analyses leading to (2.3) and to (2.5) are very similar. In fact by changing the arrow of time we can obviously interpret the stars in Fig. 2 as numerical approximations to the *stable* initial value problem

$$dy/d\tau = -y, \quad 0 \leq \tau \leq 5, \quad y(0) = 117.4$$

and then apply the error bounds corresponding to the stable case. The point to note is that in this interpretation the numerical method for which error bounds have to be derived is really the *backward* Euler method: moving forward in τ with the backward Euler formula is moving forward in t with the explicit Euler formula, see e.g. [12, p. 215].

The examples above suggest that there are good shadowing trajectories when the solutions of the differential system either converge exponentially as in (2.1) or diverge exponentially as in (2.4). In the first case the shadowing trajectory is the trajectory that has the same *initial* value as the numerical solution, in the second case the shadowing trajectory has the same *final* value as the numerical solution. Of course, (2.1) and (2.4) can be put together into a system $du/dt = Du$, with $u = [x, y]^T$ and $D = \text{diag}(1, -1)$. Numerical solutions to this system by one-step numerical methods can be shadowed by prescribing $x(0)$ and $y(T_{\max})$. For a linear system $du/dt = Au$ with A a $d \times d$ matrix with no eigenvalue on the imaginary axis (i.e. A hyperbolic), the normal modes are either exponentially growing or exponentially decreasing; shadowing of numerical solution is possible by prescribing the initial value of the decreasing modes and the final value of the growing modes. This follows the spirit of the use of *exponential dichotomies* in the derivation of good error bounds in two-point boundary value problems, see [3, Section 3.4].

3. The shadowing lemma

We now look at processes of the form

$$u_{n+1} = f(u_n), \tag{3.1}$$

where f is a mapping in \mathbb{R}^d . A simple example is given by the logistic map (1.1). The case of autonomous systems of differential systems in \mathbb{R}^d is also included by considering f to be the Δt flow of the system, i.e. the mapping in \mathbb{R}^d that advances the numerical solution by Δt units of time; for instance, for (2.1), $f(x) = \exp(-\Delta t)x$.

Let us give some definitions. A *true orbit* is a sequence $\{u_n\}_{n=\nu}^{n=N}$ that satisfies (3.1) for $\nu \leq n \leq N$. In practical computation (3.1) is not satisfied exactly. Rather, the computed points $\{p_n\}_{n=\nu}^{n=N}$ satisfy $p_{n+1} = f(p_n) + \xi_n$, where ξ_n accounts for roundoff and/or other sources of error. For instance, if f is the flow of a differential system being integrated numerically, then ξ_n include roundoff and the local errors defined as in [12, Chapter II, Fig. 3.1]. The sequence $\{p_n\}_{n=\nu}^{n=N}$ is called a *pseudo-orbit*. Assume that there is a maximum noise δ such that

$$|\xi_n| = |p_{n+1} - f(p_n)| \leq \delta, \quad \nu \leq n \leq N.$$

We then say that $\{p_n\}_{n=\nu}^{n=N}$ is a δ -pseudo-orbit. Finally we say that a true orbit $\{\tilde{u}_n\}_{n=\nu}^{n=N}$ ε -shadows the pseudo-orbit $\{p_n\}_{n=\nu}^{n=N}$ if $|\tilde{u}_n - p_n| \leq \varepsilon$, $\nu \leq n \leq N$.

If the dynamics of f includes exponential divergence of orbits, for instance if it is chaotic, then very small perturbations ξ_n can have a big impact on the computed points. In (1.1) with $\mu = 3.8$ the orbit that starts at $u_0 = 0.4$ has $u_{30} = 0.1936$; when the initial datum is perturbed to 0.400001, the value of u at $n = 30$ changes to 0.9335, i.e. when the initial data are known with an uncertainty of one part in a million, the result at $n = 30$ is completely undetermined, because the size of the underlying errors is comparable to the size of the dynamic variable u . In such situations shadowing would be of much use; at least we could ensure that the computed points closely represent some true orbit. We have the following result due to Anosov [2] and Bowen [5,6] and usually known as the *shadowing lemma*.

Theorem 3.1. *If f is a hyperbolic diffeomorphism, then for every $\varepsilon > 0$ there is a $\delta > 0$ such that every δ -pseudo-orbit can be ε -shadowed.*

Note that the length $N - \nu$ of the pseudo-orbit may be arbitrarily large. Hyperbolicity means that at each point u the Jacobian matrix of f has eigenvalues that are in modulus either < 1 (damping) or > 1 (growth). Furthermore the angle between the stable and unstable directions should be bounded away from zero uniformly in u and the rates of damping and growth should be bounded away from 1 uniformly in u . Unfortunately, hyperbolicity is a very restrictive condition; for instance the logistic map (1.1) does not satisfy it. Shadowing results for nonhyperbolic situations can be seen in [7,19].

4. Checking shadowing with the computer

The results quoted at the end of the previous section are mainly of theoretical interest. Hammel, Yorke and Grebogi [14] devised a method whereby practical shadowing results can be obtained. In their approach once a numerical orbit has been generated by the computer, the same computer is used to calculate *rigorously* how long a true trajectory exists near the numerical orbit. In [14] it is assumed that f in (3.1) is one-dimensional, $d = 1$. Once a numerical orbit $\{p_n\}_{n=\nu}^{n=N}$ has been found, the computer is used to successively find for $n = N, \dots, \nu$ intervals I_n in such a way that $I_N = [p_N, p_N]$ and $f(I_{n-1}) \supseteq I_n$, $\nu \leq n \leq N$. The last condition guarantees that an exact orbit $\{\tilde{u}_n\}_{n=\nu}^{n=N}$ exists such that $\tilde{u}_n \in I_n$, $\nu \leq n \leq N$. Finally by measuring the distance between p_n and the endpoints of the interval I_n a bound for the distances $|\tilde{u}_n - p_n|$ is obtained. Note that the I_n are found in a recursion with n decreasing, in agreement with our treatment of the unstable problem (2.4). When working in single precision with a Cray X-MP computer, Hammel, Yorke and Grebogi find that for (1.1) with $\mu = 3.8$, $p_0 = 0.4$, $\nu = 0$, $N = 10^7$, the computed pseudo-orbit is ε -shadowed by a true orbit with $\varepsilon = 10^{-8}$. It is conjectured that with a noise amplitude δ , ε -shadowing with $\varepsilon \approx \sqrt{\delta}$ holds for orbits of length $N - \nu \approx 1/\sqrt{\delta}$.

In [9,15] similar results are derived for two-dimensional problems. Further developments may be seen in [23]. Chow and Palmer have considered a similar approach; this work is reported in [13].

5. Equilibrium points in differential equations

Another situation where the idea of shadowing is useful concerns the numerical solution of systems of differential equations in \mathbb{R}^d ,

$$du/dt = F(u), \tag{5.1}$$

in the neighbourhood of a hyperbolic equilibrium point. Without loss of generality we assume that the equilibrium is at $u = \mathbf{0}$ so that $F(\mathbf{0}) = (\mathbf{0})$. The hypothesis that the equilibrium is hyperbolic means that the Jacobian matrix A of F at $\mathbf{0}$ has no eigenvalue on the imaginary axis. Hence the modes of the linearized system

$$du/dt = Au, \tag{5.2}$$

are either exponentially growing or exponentially decreasing. Each vector u in \mathbb{R}^d may be uniquely decomposed as $u = x + y$ where x belongs to the subspace $X \subset \mathbb{R}^d$ associated with the decreasing modes and y belongs to the subspace $Y \subset \mathbb{R}^d$ associated with the increasing modes (see Fig. 3). If $X = \{\mathbf{0}\}$ and $Y = \mathbb{R}^d$, then the equilibrium is a sink; if $Y = \{\mathbf{0}\}$ and $X = \mathbb{R}^d$, then the equilibrium is a source; in any other case the equilibrium is a saddle [11]. The subspaces X and Y are invariant by the flow of (5.2); when going back to the original (5.1) we find, in a small neighbourhood W of $\mathbf{0}$ a local stable manifold M_S tangent to X at $\mathbf{0}$ and consisting of solutions that tend to the equilibrium as $t \rightarrow \infty$, along with a local unstable manifold M_U tangent to Y at $\mathbf{0}$ and consisting of solutions that tend to the equilibrium as $t \rightarrow -\infty$. Any solution of (5.1) not in M_S or M_U stays in W during a finite time interval. The length of this interval depends on the individual solution and may be arbitrarily large. By considering initial conditions on opposite sides of M_S it is easily seen that the distance between two trajectories may grow exponentially with t , even if they are initially very close. Therefore numerical methods cannot be expected to approximate individual trajectories very well over long time intervals and the idea of shadowing

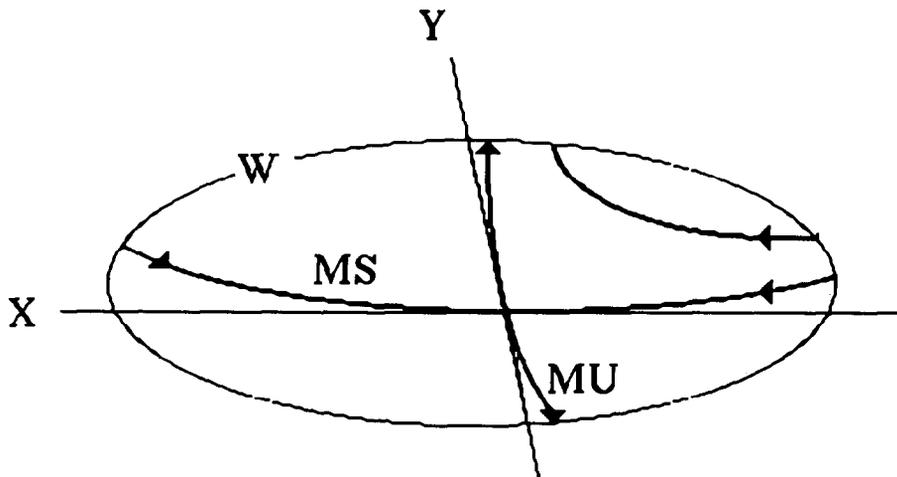


Fig. 3.

becomes useful as shown by Beyn [4]. In [4] both one-step and multi-step methods are considered, but for simplicity we only consider here the one-step case

$$\mathbf{u}_{n+1} = \mathbf{u}_n + \Delta t \mathbf{F}_{\Delta t}(\mathbf{u}_n). \tag{5.3}$$

The two following reasonable consistency hypotheses are made:

(H1) $\mathbf{F}_{\Delta t}(\mathbf{u})$ converges to $\mathbf{F}(\mathbf{u})$ as $\Delta t \rightarrow 0$, uniformly in \mathbf{u} , $\mathbf{u} \in W$. Furthermore the Jacobian matrix $\mathbf{F}'_{\Delta t}(\mathbf{u})$ of $\mathbf{F}_{\Delta t}(\mathbf{u})$ with respect to \mathbf{u} converges to the Jacobian matrix $\mathbf{F}'(\mathbf{u})$ as $\Delta t \rightarrow 0$, uniformly in \mathbf{u} , $\mathbf{u} \in W$.

(H2) The truncation errors $\mathbf{f}_{\Delta t}(\mathbf{u}) - [\mathbf{u} + \Delta t \mathbf{F}_{\Delta t}(\mathbf{u})]$ are $O(\Delta t^{p+1})$ uniformly in \mathbf{u} , $\mathbf{u} \in W$. Here $\mathbf{f}_{\Delta t}$ represents the Δt flow of (5.1) and p is the order of the numerical method.

Beyn shows that there are a neighbourhood W of the equilibrium and positive constants C and η such that, for $\Delta t \leq \eta$, to each numerical orbit $\{\mathbf{u}\}_{n=\nu}^{n=N}$ in W there corresponds a solution $\tilde{\mathbf{u}}(t)$ of (5.1) such that

$$|\tilde{\mathbf{u}}(n\Delta t) - \mathbf{u}_n| \leq C\Delta t^p, \quad \nu \leq n \leq N.$$

Note that N may be $\gg 1$ while ν may be $\ll -1$. The cases $\nu = -\infty$ and $N = \infty$ are allowed and respectively correspond to numerical orbits in the unstable and stable manifold of the mapping (5.3). The conclusion is that the numerical method faithfully reproduces the phase portrait of (5.1) near the equilibrium point.

As in Section 2, the shadowing solution $\tilde{\mathbf{u}}(t)$ is found by imposing that the stable X -component of $\tilde{\mathbf{u}}(\nu\Delta t)$ coincides with the X -component of \mathbf{u}_ν , while the unstable Y -component of $\tilde{\mathbf{u}}(N\Delta t)$ coincides with the Y -component of \mathbf{u}_N . To bound the errors $\tilde{\mathbf{e}}_n := \tilde{\mathbf{u}}(n\Delta t) - \mathbf{u}_n$, note that, as in Section 2, they satisfy

$$\tilde{\mathbf{e}}_{n+1} = \tilde{\mathbf{e}}_n + \Delta t [\mathbf{F}_{\Delta t}(\tilde{\mathbf{u}}(n\Delta t)) - \mathbf{F}_{\Delta t}(\mathbf{u}_n)] + \tilde{\mathbf{l}}_n,$$

where $\tilde{\mathbf{l}}_n$ denote the $O(\Delta t^{p+1})$ truncation errors (see (H2) above). Hence

$$\tilde{\mathbf{e}}_{n+1} = \tilde{\mathbf{e}}_n + \Delta t A \tilde{\mathbf{e}}_n + \tilde{\mathbf{l}}_n + \Delta t [\mathbf{F}_{\Delta t}(\tilde{\mathbf{u}}(n\Delta t)) - \mathbf{F}_{\Delta t}(\mathbf{u}_n) - A(\tilde{\mathbf{u}}(n\Delta t) - \mathbf{u}_n)].$$

Without the term in square brackets, this would be a linear recursion that is easily analyzed as in Section 2 by working forward in the X -component and backward in the Y -component. The term in brackets introduces a small perturbation because, due to (H1), it possesses a bound $L|\tilde{\mathbf{e}}_n|$ with a Lipschitz constant that may be rendered arbitrarily small by suitably reducing the size of the neighbourhood W .

The results in [4] have been extended to partial differential equations by Alouges and Debussche [1] and by Larsson and Sanz-Serna [18]. The paper [1] works in an abstract framework and only considers discretization of the time variable. The analysis provided is a direct extension of that in [4] and the error bounds obtained are not of optimal order. The paper [18] considers the piecewise linear Galerkin space discretization of the reaction–diffusion problem

$$\begin{aligned} u_t - \Delta u &= f(u), & x \in \Omega, \\ u &= 0, & x \in \partial\Omega, \end{aligned}$$

where Ω is a bounded convex polygonal domain in \mathbb{R}^d , $d = 1, 2$, or 3 and $f \in C^2(\mathbb{R})$. If $d = 2$ or 3 , then f is also supposed to satisfy a growth condition

$$f^{(j)}(u) \leq C(1 + |u|^{\delta-j}), \quad j = 0, 1, 2,$$

where $\delta = 3$ if $d = 3$ and $2 \leq \delta < \infty$ if $d = 2$. It is shown that there are positive numbers ρ_0 , h_0 , and C such that, for partitions of diameter $h < h_0$, any Galerkin solution $u^h(t)$ that for $\tau \leq t \leq T$ stays in the H^1 -ball of radius ρ_0 around a hyperbolic equilibrium can be shadowed by a true solution $\tilde{u}(t)$ in such a way that

$$\|u^h(t) - \tilde{u}(t)\|_{L_2} \leq C(1 + (t - \tau)^{-1/2})h^2, \quad \tau < t \leq T,$$

$$\|u^h(t) - \tilde{u}(t)\|_{H^1} \leq C(1 + (t - \tau)^{-1/2})h, \quad \tau < t \leq T.$$

Note that the exponents of h are optimal. The bounds deteriorate for t near τ , as we would expect in a situation where the solutions considered are not necessarily smooth.

6. Concluding remarks

For a survey of the use of numerical methods in long-time integrations see [20]. The idea of shadowing is related to the notion of *backward error analysis*; the numerical results are seen as approximations to perturbed problems. References on the backward error analysis of numerical integrators are [8,21]. A related idea is the method of *modified equations* [10], where the equation rather than the initial value is changed in order to shadow the numerical solution.

Acknowledgements

This research has been supported by “Junta de Castilla y León” under project 1031-89 and by “Dirección General de Investigación Científica y Técnica” under project PB89-0351.

References

- [1] F. Alouges and A. Debussche, On the qualitative behavior of the orbits of a parabolic partial differential equation and its discretization in the neighborhood of a hyperbolic fixed point, *Numer. Funct. Anal. Optimiz.* 12 (1991) 253–269.
- [2] D.V. Anosov, Geodesic flows and closed Riemannian manifolds with negative curvature, *Proc. Steklov Inst. Math.* 90 (1967).
- [3] U.M. Ascher, R.M.M. Mattheij and R.D. Russell, *Numerical Solutions of Boundary Value Problems for Ordinary Differential Equations* (Prentice Hall, Englewood Cliffs, NJ, 1988).
- [4] W.-J. Beyn, On the numerical approximation of phase portraits near stationary points, *SIAM J. Numer. Anal.* 24 (1987) 1095–1113.
- [5] R. Bowen, ω -limit sets for axiom A-diffeomorphisms, *J. Differential Equations* 18 (1975) 333–3339.
- [6] R. Bowen, *Equilibrium States and the Ergodic Theory of Anosov Diffeomorphisms* (Springer, Berlin, 1975).
- [7] E. Coven, I. Kan and J.A. Yorke, Pseudo-orbit shadowing in the family of tent maps, *Trans. Amer. Math. Soc.* 308 (1988) 227–241.
- [8] T. Eirola, Aspects of backward error analysis in numerical ODE’s, *J. Comput. Appl. Math.* 45 (1993) 65–73.
- [9] C. Grebogi, S.M. Hammel, J.A. Yorke and T. Sauer, Shadowing of physical trajectories in chaotic dynamics: containment and refinement, *Phys. Rev. Lett.* 65 (1990) 1527–1530.
- [10] D.F. Griffiths and J.M. Sanz-Serna, On the scope of the method of modified equations, *SIAM J. Sci. Statist. Comput.* 7 (1986) 994–1008.

- [11] J. Guckenheimer and P. Holmes, *Nonlinear Oscillations, Dynamical Systems and Bifurcation of Vector Fields* (Springer, New York, 1983).
- [12] E. Hairer, S.P. Nørsett and G. Wanner, *Solving Ordinary Differential Equations I: Nonstiff Problems* (Springer, Berlin, 1987).
- [13] J.K. Hale, Dynamics and numerics, in: D.S. Broomhead and A. Iserles, eds., *The Dynamics of Numerics and the Numerics of Dynamics* (Clarendon, Oxford, 1992) 243–253.
- [14] S.M. Hammel, J.A. Yorke and C. Grebogi, Do numerical orbits of chaotic dynamical processes represent true orbits?, *J. Complexity* 3 (1987) 136–145.
- [15] S.M. Hammel, J.A. Yorke and C. Grebogi, Numerical orbits of chaotic processes represent true orbits, *Bull. Amer. Math. Soc.* 19 (1988) 465–469.
- [16] E. Isaacson and H.B. Keller, *Analysis of Numerical Methods* (Wiley, New York, 1966).
- [17] S. Larsson, The long-time behavior of finite-element approximations of solutions to semilinear parabolic problems *SIAM J. Numer. Anal.* 26 (1989) 248–365.
- [18] S. Larsson and J.M. Sanz-Serna, The behavior of finite element solutions of semilinear parabolic problems near stationary points, *SIAM J. Numer. Anal.* (to appear).
- [19] H.E. Nusse and J.A. Yorke, Is every approximate trajectory of some process near an exact trajectory of a nearby process?, *Comm. Math. Phys.* 114 (1988) 381–347.
- [20] J.M. Sanz-Serna, Numerical ordinary differential equations vs. dynamical systems, in: D.S. Broomhead and A. Iserles, eds., *The Dynamics of Numerics and the Numerics of Dynamics* (Clarendon, Oxford, 1992) 81–106.
- [21] J.M. Sanz-Serna, Symplectic integrators for Hamiltonian problems: an overview, *Acta Numer.* 1 (1992) 243–286.
- [22] J.M. Sanz-Serna and A.M. Stuart, A note on uniform in time error estimates for approximations to reaction-diffusion equations, *IMA J. Numer. Anal.* 12 (1992) 457–462.
- [23] T. Sauer and J.A. Yorke, Rigorous verification of trajectories for the computer simulation of dynamical systems, *Nonlinearity* 4 (1991) 961–979.
- [24] H.J. Stetter, *Analysis of Discretization Methods for Ordinary Differential Equations* (Springer, New York, 1973).