A Survey of Shadowing Methods for Numerical Solutions of Ordinary Differential Equations ⋆

Wayne Hayes and Kenneth R. Jackson *

Computer Science Department, University of Toronto, Toronto, Ontario, M5S 2E4, Canada

Abstract

A *shadow* is an exact solution to a set of equations that remains close to a numerical solution for a long time. Shadowing can thus be used as a form of backward error analysis for numerical solutions to ordinary differential equations. This survey introduces the reader to shadowing with a detailed tour of shadowing algorithms and practical results obtained over the last 15 years.

Key words: shadowing, validated computing, ODEs, IVPs

1 Introduction

An *initial value problem* (IVP) for an *ordinary differential equation* (ODE) is

$$\mathbf{y}'(t) = \mathbf{f}(\mathbf{y}(t)), \tag{1}$$
$$\mathbf{y}(t_0) = \mathbf{y}_0. \tag{2}$$

An *autonomous* ODE such as (1) contains no explicit dependence on t. If $\mathbf{y}(t; t_0, \mathbf{y}_0)$ is the solution of (1,2), we let the *time-h solution operator* φ_h be

$$\varphi_h(\mathbf{u}) \equiv \mathbf{y}(h; 0, \mathbf{u}). \tag{3}$$

Let $\tilde{\varphi}_h$ be a numerical approximation to φ_h computed by some numerical method for small h, and let $\mathbf{y}_{i+1} = \tilde{\varphi}_{h_i}(\mathbf{y}_i)$ define a sequence of discrete points

⋆ The authors thank NSERC Canada for support.

* *E-mail:* {wayne,krj}@cs.toronto.edu

representing approximations to $\mathbf{y}(t_{i+1}; t_0, \mathbf{y}_0)$ where $t_{i+1} = t_i + h_i$. We call this sequence a *pseudo-trajectory*. The natural first question to ask about pseudo-trajectories is how accurately they approximate the exact solution. Since the *forward error* $\|\mathbf{y}_i - \mathbf{y}(t_i; t_0, \mathbf{y}_0)\|$ can in general grow exponentially (Dahlquist and Björck, 1974, §§8.1.2, 8.3.6), more sophisticated methods of error analysis must be used to gain insight into the value of the numerical solution.

Backward error analysis is a general term applied to methods of error analysis that relate the pseudo-trajectory to the exact solution of a nearby problem (Corless, 1994a). *Defect based* backward error analysis requires a piecewise differentiable interpolant $\mathbf{u}(t)$ of the pseudo-trajectory, and then defines the *defect* as $\delta(t) = \mathbf{u}'(t) - \mathbf{f}(\mathbf{u}(t))$. If for some input tolerance ε we can show that $\|\delta(t)\| \leq \varepsilon$ wherever $\mathbf{u}'(t)$ is defined over the whole of the interpolated solution, then the interpolated solution is the exact solution to an ε-close problem (see for example (Corless and Corliss, 1992; Corless, 1994a)), namely $\mathbf{u}'(t) = \mathbf{f}(\mathbf{u}(t)) + \delta(t)$.

Defect-based and other backward error analysis methods modify (1) but leave (2) untouched. In contrast, shadowing (Pilyugin, 1999; Palmer, 2000) is a method of backward error analysis in which (1) remains fixed while (2) is allowed to change. In other words, a *shadow* $\bar{\mathbf{y}}(t)$ is an exact solution to (1) remaining close to the pseudo-trajectory $\mathbf{u}(t)$, but having slightly different initial conditions:

$$\bar{\mathbf{y}}'(t) = \mathbf{f}(\bar{\mathbf{y}}(t)), \qquad \|\bar{\mathbf{y}}(t) - \mathbf{u}(t)\| < \varepsilon,$$

for a nontrivial duration of time t including $t = t_0$. Shadowing is thus best applied to systems in which the governing equations are extremely well-known, and virtually all error is introduced by imprecise knowledge of initial conditions and/or by numerical error in the computation of the solution. It is less applicable to systems in which the mathematics only approximately model the truth. For example, shadowing is an appropriate measure of error for the gravitational n-body problem, because the equations of motion are extremely well-understood, and an exact solution of the model very closely approximates the behaviour of a real system under the assumed conditions of the model, whereas the initial conditions for the system are often not known precisely. Conversely, shadowing is an inappropriate measure of error for a weather simulation, because the models are known to be only rough approximations of the real processes involved, and so having a numerical solution that closely follows an exact solution of such an approximate model is of dubious value.

1.1 Motivation

Many physical systems under active study today can be modelled using ODEs; however, many of them display *sensitive dependence on initial conditions*, which means that two solutions that are initially close to each other tend to diverge exponentially with time. Since numerical methods introduce small errors that produce a pseudo-trajectory rather than an exact solution, it is virtually guaranteed that a pseudo-trajectory of such an ODE will diverge exponentially away from the exact solution with the same initial conditions. Although this is widely recognized, its impact on the qualitative properties of a pseudo-trajectory is not well understood.

Even when the mathematical model corresponds extremely closely to reality, real systems undergo external perturbations, so one could argue that numerical errors can be grouped into the same category as external perturbations (Corless, 1992b, 1994a). However, numerical errors may be biased in qualitatively different ways than natural perturbations, and may introduce biases into the numerical solution that cause it to behave in a nonphysical manner. For example, (i) many physical perturbations do not appreciably change the energy of the system, whereas spurious energy dissipation can be a major problem in long numerical integrations of conservative systems. Although symplectic integrators (Channell and Scovel, 1990; Sanz-Serna, 1992) and other types of conservative integrators (Shadwick et al., 1999) may conserve certain quantities, it is not clear that they do not introduce other biases, such as nonphysical energy transport. (ii) Physical systems often satisfy properties such as conservation of energy, phase-space volume and various types of momentum, many of which are well-conserved in real systems that experience perturbations, but are not well-conserved by an otherwise well-behaved numerical method. (iii) The effect on simulations of numerical error can be much greater than actual perturbations, even if those perturbations are larger. For example, nearby stars and the Galaxy at large exert forces on the Solar System that are about 10^{-12} times smaller than the forces from our Sun. It is not difficult to create integrations with local numerical errors several orders of magnitude smaller than this, and yet unless these integrations somehow account for symplectic structure or energy conservation, they produce an integration of the Solar System which quickly and clearly diverges from the behaviour of the real Solar System. (iv) In general, we want to ensure that modifying the system of equations represented by (1) does not affect any quantities of interest (Skeel, 1996, 1999).

Despite the above issues, numerical solutions often appear to mimic with astounding accuracy the phenomena they purport to simulate. Simulations of galaxies often closely resemble real galaxies (see almost any paper on galaxy simulation in Clarke and West (1997) or Merritt et al. (1999)). Even galaxy collisions can be modelled in a convincing manner (Struck, 1997). More generally, exponential divergence of nearby trajectories implies that an initially dense ensemble of points will disperse into a uniform distribution in a relatively short time (Skeel, 1996, 1999). This effect is also seen in numerical simulations of chaotic systems (Merritt and Valluri, 1996; Merritt, 1999). The natural question to ask is whether these simulations are behaving in a fashion similar to real systems, or if they only superficially mimic real systems but are in fact behaving incorrectly at a more fundamental level. If this were the case, then we could be lulled into a false sense of security whilst our understanding of these systems becomes compromised.

Since shadowing disallows changes in the model (1), some of these kinds of insidious errors can be ruled out. Furthermore, if the problem (1) arises from a purely mathematical context and not a physical one, we may be earnestly interested in the properties of exact solutions of (1), in which case a recourse to shadowing may be the only option. The only remaining question would then be whether shadows are typical of exact solutions chosen at random. This question is explored more deeply later in this survey. On the other hand, shadowing is extremely expensive. Whereas defect controlled methods (Enright and Hayes, 2003) are of roughly equal expense compared to more traditional integration methods, shadowing typically requires integration of the variational equation, which for a large system can be orders of magnitude more expensive to compute than the numerical solution.

1.2 Definitions

An **orbit** is a discrete sequence of points, a **solution** is a continuous curve, and a **trajectory** more generally refers to either an orbit or a solution depending upon context. We assume a well-scaled problem where all macroscopic quantities of interest are of order unity; $|\cdot|$ denotes the magnitude of a scalar, while $\|\cdot\|$ denotes a norm of a vector or matrix. We use the max norm unless otherwise noted.

Let $\varphi(\mathbf{x})$ be a diffeomorphism. If φ is just a map, then it may be a simple equation, such as the "logistic" equation $\varphi(x) = 1 - 2x^2$, which maps the interval $[-1, 1]$ onto itself. The case we are more interested in here is where φ represents the one-timestep flow through $\mathbf{x}$ of the solution to an ODE. If the timestep is fixed, then φ_h is the same function on each step, and for simplicity we refer to φ_h simply as φ for constant h.

The **iterated map** $\varphi^i(\mathbf{x})$ is the result of repeatedly composing φ with itself i times, *e.g.,* $\varphi^3(\mathbf{x}) = \varphi(\varphi(\varphi(\mathbf{x})))$. An **exact orbit** $\{\mathbf{x}_i\}_{i=j}^k$ of φ satisfies $\mathbf{x}_{i+1} = \varphi(\mathbf{x}_i)$, *i.e.,* $\mathbf{x}_i = \varphi^{i-j}(\mathbf{x}_j)$, for $j \le i < k$. We allow $j = -\infty$ and $k = \infty$;

$\{\mathbf{y}_i\}_{i=j}^k$ is a δ-**pseudo-orbit** or **noisy orbit** for φ if $\|\mathbf{y}_{i+1} - \varphi(\mathbf{y}_i)\| \le \delta$ for $j \le i < k$, where δ is called the **noise amplitude**. For a discrete map, δ can be as small as the machine epsilon; for both discrete maps and ODE systems, it is a bound on the one-step error. For $j \le i < k$, the **one-step error** made between step i and step $i+1$ of the pseudo-orbit $\{\mathbf{y}_i\}_{i=j}^k$ is $\mathbf{e}_{i+1} = \mathbf{y}_{i+1} - \varphi(\mathbf{y}_i)$. Thus, an exact trajectory is one whose one-step errors are identically zero, and a δ-pseudo-orbit is one whose one-step errors satisfy $\|\mathbf{e}_i\| \le \delta$ for $j < i \le k$. An exact trajectory $\{\mathbf{x}_i\}_{i=j}^k$ ε-**shadows** a pseudo-trajectory $\{\mathbf{y}_i\}_{i=j}^k$ if $\|\mathbf{y}_i - \mathbf{x}_i\| \le \varepsilon$ for $j \le i \le k$. The δ'-pseudo-trajectory $\mathbf{Z} = \{\mathbf{z}_i\}_{i=j}^k$ is a **numerical shadow** of the δ-pseudo-trajectory $\mathbf{Y} = \{\mathbf{y}_i\}_{i=j}^k$ if their one-step errors are tightly bounded by δ' and δ, respectively, and $\delta' \ll \delta$. In practice, a numerical shadow usually only has smaller error *estimates* than the original noisy orbit, because in most cases neither orbit has rigorously computed error bounds. To have confidence in the value of a numerical shadow, we like its noise to be as small as possible. For a map, the noise should ideally be the machine precision. For an ODE solution, the noise is "as small as possible" using some accurate integrator with its error tolerance set very stringently. The pseudo-orbit $\{\mathbf{y}_i\}_{i=j}^k$ has a **glitch** at point $i = G_0 < k$ if for some relevant ε there exists an exact trajectory that ε-shadows $\{\mathbf{y}_i\}_{i=j}^{G_0}$, but no exact trajectory exists that ε-shadows $\{\mathbf{y}_i\}_{i=j}^{G}$ for $G > G_0$ (Grebogi et al., 1990).

Although rigorously disproving the existence of shadows of particular numerical trajectories is a virtually untouched area of research, the failure of a particular method to find a shadow is often cited as evidence that an actual glitch occurs somewhere in the vicinity of the computed end-of-shadow (Grebogi et al., 1990; Sauer and Yorke, 1991; Dawson et al., 1994; Sauer et al., 1997; Quinlan and Tremaine, 1992; Hayes, 1995). This conclusion is not always valid, however (see the discussion following Theorem 3.1, in this survey), and so Hayes (2001) proposes two different terms. The term glitch, or **hard glitch**, should be reserved for the case in which the above definition can be verified, *i.e.,* non-existence of shadows can be *proved.* For example, a function $\varphi : X \to X$ which maps an interval onto itself may produce a numerically generated orbit of the iterated map which lies outside this interval. If a numerically generated point, say x_i, moves more than ε away from the interval X then a hard glitch has occurred. However, the failure of a particular method to find a shadow is a different matter, and for this case Hayes (2001) proposes the term **soft glitch**. For systems such as the n-body problem, the notion of a hard glitch cannot be used without proof because the set of unphysical points in phase space has measure zero. That is, in a Newtonian, Euclidian phase space, particles can have any finite position and velocity. Furthermore, small numerical errors are constantly occuring, and, if the system is integrated carefully and local errors remain small, there is no obvious point at which one can say, "this behaviour is nonphysical". One can arbitrarily decide, for example, that when the total computed energy of the numerical solution has diverged from the known energy of the system by some chosen amount, the solution is no longer valid. But this is not the spirit of the term "glitch". The spirit of the term is "a point at which all exact trajectories diverge from a numerical one", and currently this can only be proved for simple systems.

1.3 Tutorial

A simple example of a shadow is provided by Quinlan and Tremaine (1992), hereafter referred to as QT. Let $y'' = y$, which can be re-written as a pair of first order equations as $y' = v, \;\; v' = y$, where v is velocity. If $y(t_0) = v(t_0) = 0$ for any t_0, then the exact solution is $y = v = 0 \; \forall t$. Now, assume that $y = v = 0$ for $t < 0$, and assume that the system is solved exactly for all $t \neq 0$. Introducing a perturbation of size $\Delta v = \varepsilon$ at $t = 0$ gives the following "noisy" solution:

$$y(t) = \begin{cases} 0, & t < 0, \\ \varepsilon \frac{e^t - e^{-t}}{2}, & t > 0. \end{cases}$$

A shadow of this noisy solution is $x(t) = \varepsilon e^t/2$, which remains within $\varepsilon/\sqrt{2}$ (in phase space) from $y(t)$ for all t.

Next, purely for pedagogical purposes, we offer a proof of an almost "trivial" theorem: if a map is contracting, then noisy orbits are shadowed.

Theorem 1.1 (Contracting map shadowing theorem) *Let X be a metric space and let $\varphi : X \to X$ be a continuous, uniformly contracting map, i.e., $\exists \rho < 1$ s.t. $\forall \mathbf{x}, \mathbf{y} \in X, \;\; \|\varphi(\mathbf{x}) - \varphi(\mathbf{y})\| \leq \rho \|\mathbf{x} - \mathbf{y}\|$. Then for every $\varepsilon > 0$ there exists $\delta > 0$ such that every δ-pseudo orbit remaining in X is ε-shadowed.*

Proof 1.1 *Assume we are given $\varepsilon > 0$. Let $\delta = \varepsilon(1 - \rho)$. Suppose $\{\mathbf{y}_i\}_{i=j}^{\infty}$ is a δ-pseudo-orbit that remains in X. Let $\mathbf{x}_j = \mathbf{y}_j$ and let $\mathbf{x}_{i+1} = \varphi(\mathbf{x}_i)$ for $i \geq j$, i.e., $\{\mathbf{x}_i\}_{i=j}^{\infty}$ is an exact orbit. We will show by induction on i that $\|\mathbf{x}_i - \mathbf{y}_i\| \leq \varepsilon$ for $i \geq j$.*

Base case: $\|\mathbf{x}_j - \mathbf{y}_j\| = 0 < \varepsilon$, by our choice of $\mathbf{x}_j$.

Induction step: Assume $\|\mathbf{x}_i - \mathbf{y}_i\| \leq \varepsilon$ for $i \geq j$. Then

$$\begin{aligned} \|\mathbf{x}_{i+1} - \mathbf{y}_{i+1}\| &= \|\varphi(\mathbf{x}_i) - \varphi(\mathbf{y}_i) + \delta_1\|, \qquad \|\delta_1\| \leq \delta \\ &\leq \|\varphi(\mathbf{x}_i) - \varphi(\mathbf{y}_i)\| + \delta \\ &\leq \rho\varepsilon + \delta \\ &= \varepsilon \end{aligned}$$

Remark: If φ is expanding, then φ^{-1} is contracting, and we can apply the above theorem in reverse time, as long as $\mathbf{y}_i \in X\ \forall i \geq j$.

Another instructive way to look at shadowing is in terms of its relation to finding the zero of a function. To wit, let $\mathbf{Y} = \{\mathbf{y}_i\}_{i=0}^N$ be a δ-pseudo-trajectory in $\mathbf{R}^n$, and let $\mathbf{E} = \{\mathbf{e}_i\}_{i=1}^N$ be the set of one-step errors. Let $\mathbf{g} : \mathbf{R}^{(N+1)n} \to \mathbf{R}^{Nn}$ be a function that takes as input the entire orbit $\mathbf{Y}$ and produces as output the set of one-step errors $\mathbf{E}$, *i.e.,* $\mathbf{g}(\mathbf{Y}) = \mathbf{E}$. Since the one-step errors are assumed to be small, $\|\mathbf{E}\|$ is small. That is, $\mathbf{Y}$ may be close to a zero of $\mathbf{g}$, if one exists. A zero of $\mathbf{g}$ would represent an orbit with zero one-step error, *i.e.,* an exact orbit. This is an ideal situation in which to apply a zero-finding method such as Newton's method. If the method converges to an orbit $\mathbf{X}$ which is ε-close to $\mathbf{Y}$, then $\mathbf{X}$ ε-shadows $\mathbf{Y}$. This is the idea behind *refinement* (Grebogi et al., 1990; Quinlan and Tremaine, 1992) which will be discussed in more detail below.

2 Hyperbolicity

One of the most important concepts in shadowing is that of *hyperbolicity.* The following definitions are commonly used in the shadowing literature. See for example Palmer (1988), on which the following description is based. In this section, we will concentrate on *maps*, keeping in mind that we can translate between maps and solutions of ODEs by looking at the time-h solution operator $\varphi_h(x)$ defined in equation (3). Let $\varphi : \mathbf{R}^n \to \mathbf{R}^n$ be a diffeomorphism. Let $D\varphi(\mathbf{x})$ be the Jacobian of $\varphi(\mathbf{x})$, which exists, is unique, and is invertible since φ is a diffeomorphism. Every orbit of φ has associated with it a linear difference equation called the *linear variational equation,*

$$\mathbf{z}_{i+1} = D\varphi(\mathbf{x}_i)\mathbf{z}_i. \tag{4}$$

A sequence of Jacobians along an orbit can be multiplied together to produce a Jacobian of the corresponding sequence of applications of the map, $\Phi(i, j)$. The linear variational equation (4) is said to have an *exponential dichotomy* if there are positive constants K, α and a family of projections P_i such that

$$\begin{aligned} P_{i+1} D\varphi(\mathbf{x}_i) &= D\varphi(\mathbf{x}_i) P_i \text{ for all } i, & (5)\\ \|\Phi(i,j)P_j\| &\leq K e^{-\alpha(i-j)} \text{ for } i \geq j, & (6)\\ \|\Phi(i,j)(I - P_j)\| &\leq K e^{-\alpha(j-i)} \text{ for } j \geq i. & (7) \end{aligned}$$

By repeated application of (5) we obtain the identity $P_i\Phi(i,j) = \Phi(i,j)P_j$. This means that the projections P_i are invariant with respect to equation (4). That is, if $\{\mathbf{z}_i\}_{i=j'}^{k}$ is a solution to (4) such that $\mathbf{z}_j$ is in the range (resp. nullspace) of P_j for some j then $\mathbf{z}_i$ is in the range (resp. nullspace) of P_i for all i. Inequalities (6–7) say firstly, that the P_i are bounded (proof: set $i = j$ in (6)) and secondly, that the solutions $\mathbf{z}_i$ of equation (4) which lie in the range of P_i decay exponentially in forward time (the *stable subspace*), while those in the nullspace of P_i (the *unstable subspace*) decay exponentially in backward time (Palmer, 1988).

Definitions. A trajectory $\mathbf{X} = \{\mathbf{x}_i = \varphi^i(\mathbf{x})\}_{i=j'}^{k}$ for some $\mathbf{x}$ is said to be **hyperbolic under** φ if the linear variational equation $\mathbf{z}_{i+1} = D\varphi(\mathbf{x}_i)\mathbf{z}_i$ along $\mathbf{X}$ has an exponential dichotomy. Equivalently, we say that φ is **hyperbolic along X**. A set $S \subset \mathbf{R}^n$ is said to be **invariant** under φ if $\varphi(S) = S$. A compact invariant set S is said to be **hyperbolic under** φ if every trajectory $\mathbf{X}$ in S is hyperbolic with the same constants K, α, and the projection matrices P_i have a rank which is independent of $\mathbf{X}$. Equivalently, we say that φ **is hyperbolic on** S, or that S and φ form a **hyperbolic system**. If a system is hyperbolic, then the angle between the stable and unstable subspaces is always bounded away from 0 (Grebogi et al., 1990).

2.1 Hyperbolic systems

Shadowing was first discussed by Anosov (1967) and Bowen (1975), in relation to hyperbolic systems. Let S and φ be the invariant set and the map of a hyperbolic system, respectively. In such systems, Anosov (1967) proved that $\forall \varepsilon > 0$, $\exists\, \delta > 0$ such that every infinite-length δ-pseudo orbit remaining in S is ε-shadowed by a true trajectory in S. Bowen (1975) proved that the same result holds if the map is required to be hyperbolic only along trajectories in the vicinity of the pseudo-orbit. Palmer (1988) proved a similar theorem along the way towards using the theory of exponential dichotomies to prove Smale's Theorem (Smale 1965; 1967):

Theorem 2.1 (Hyperbolic set shadowing theorem) *Let S be a compact hyperbolic set for the C^1 diffeomorphism $\varphi : \mathbf{R}^n \to \mathbf{R}^n$. Then given any $\varepsilon > 0$ sufficiently small there exists $\delta > 0$ such that every doubly-infinite δ-pseudo-orbit in S has a unique ε-shadowing orbit.*

Chow and Van Vleck (1992) proved a similar theorem in the case that the function φ is allowed to change at each step. We omit the (rather long and involved) specifications of the hyperbolicity conditions of the following theorem, except to note that when the conditions hold, the difference equation $\mathbf{z}_{i+1} = D\phi_i(\mathbf{x}_i)\mathbf{z}_i$ has an exponential dichotomy for *all* sequences of functions $\{\phi_i\}_{i=0}^{\infty}$ if $\mathbf{x}_{i+1} = \phi_i(\mathbf{x}_i)$. These conditions, of course, tightly restrict the classes of sequences of functions whose orbits can be shadowed; otherwise, shadowing of numerical solutions of ODEs would be trivial!

Theorem 2.2 (Random Diffeomorphism Shadowing Lemma) *Let M be* *a smooth compact k-dimensional Riemannian manifold and let Diff(M) represent the set of all diffeomorphisms from M to M. Assume further that the [omitted] hyperbolicity conditions are satisfied. Let $\{\mathbf{y}_i\}_{i=0}^{\infty}$ be a sequence of points in M. Then for all $\varepsilon > 0$ sufficiently small $\exists\delta > 0$ such that if there exists a sequence of functions $\{\phi_i \in \text{Diff}(M)\}_{i=0}^{\infty}$ satisfying $\|\mathbf{y}_{i+1} - \phi_i(\mathbf{y}_i)\| \leq \delta$ then there exists a unique sequence $\{\mathbf{x}_i\}_{i=0}^{\infty}$ such that $\mathbf{x}_{i+1} = \phi_i(\mathbf{x}_i)$ and $\|\mathbf{x}_i - \mathbf{y}_i\| \leq \varepsilon$ for all i.*

3 Algorithms for finding shadows of maps

Although our primary purpose in this survey is to shadow numerical ODE integrations, we start with maps because they are simpler, and shadowing ODEs is a direct extension of shadowing maps.

3.1 Containment

This survey deals not with hyperbolic systems, but with systems whose pseudo-trajectories are shadowable for finite but nontrivial lengths of time even though they are not hyperbolic. For this to occur, a system must display pseudo-hyperbolicity (Hayes, 2001). We say that a system is **pseudo-hyperbolic** if trajectories of the system tend to have solutions to the variational equation which can be split into two classes, one of which tends to expand exponentially, while the other tends to contract exponentially, both simultaneously and for nontrivial lengths of time. This notion could be made more formal by, for example, attempting to find the two classes of solutions using the common methods described below, and then performing least-squares fits of these solutions to exponential curves.

For systems that are not hyperbolic, but whose trajectories display pseudo-hyperbolicity for a finite number of iterations of φ, we must be satisfied with proving the existence of finite-length shadows. The first study of shadows for non-hyperbolic systems appears to be Hammel et al. (1987). Hammel et al. (1988) and Grebogi et al. (1990) (hereafter GHYS) provide the first proof of the existence of a shadow for a non-hyperbolic system over a non-trivial length of time. Their method consists of two parts. First, they *refine* a noisy trajectory using an iterative method that produces a nearby trajectory with less noise. Refinement will be discussed in more detail below. When refinement converges to the point that the noise is of order the machine precision, they invoke *containment*, which can prove the existence of a nearby exact trajectory.

Let $\{\mathbf{y}_i\}_{i=a}^{b} \subset \mathbf{R}^2$ be a two-dimensional δ-pseudo-orbit of φ for integers a and

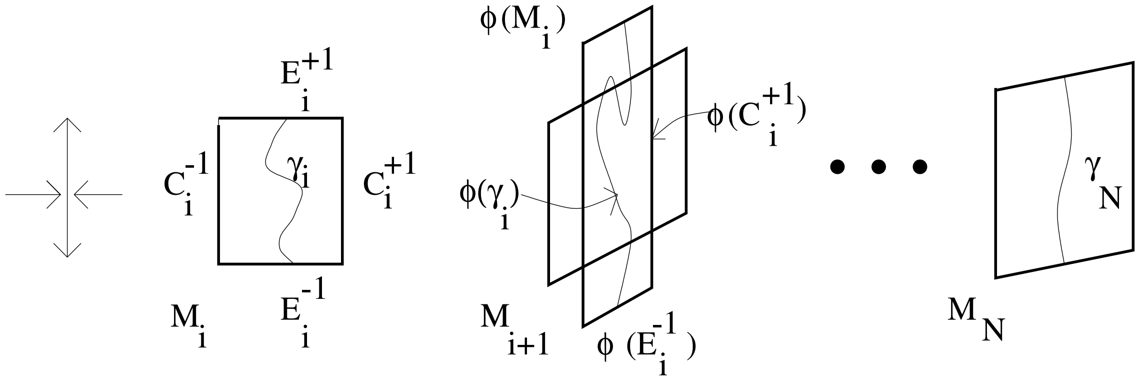

Fig. 1. Containment in two dimensions, reproduced from GHYS. The horizontal direction is contracting, and the vertical direction is expanding.

b. As i increases, orbits separated from each other by a small distance along the expanding direction diverge on average away from each other, while orbits separated by a small distance along the contracting direction approach each other, on average. The containment process consists of building a parallelogram M_i around each point $\mathbf{y}_i$ of the pseudo-orbit such that two sides $C_i^{\pm 1}$ are separated from each other along the contracting direction, while the other two sides $E_i^{\pm 1}$ are separated along the expanding direction. In order to prove the existence of a shadow, the image of M_i under φ must intersect M_{i+1} such that $\varphi(M_i)$ makes a "plus sign" with M_{i+1} (Figure 1). The property that GHYS define as a "plus sign" is

$$\varphi(E_i^j) \cap M_{i+1} = \emptyset, \qquad \varphi(M_i) \cap C_{i+1}^j = \emptyset, \;\; j = \pm 1. \tag{8}$$

To ensure this occurs, GHYS require a bound on the second derivative of φ, and the expansion and contraction amounts need to be resolvable by the machine precision. The proof of the existence of a shadow then relies on the following argument. Let γ_0 be a continuous curve in M_0 connecting the expanding sides E_0^{-1} and E_0^{+1}. Its image $\varphi(\gamma_0)$ is then stretched such that there is a subsection of $\varphi(\gamma_0)$ lying wholly within M_1, and in particular $\varphi(\gamma_0)$ leaves M_1 through the expanding sides $E_1^{\pm 1}$ at both ends. Let γ_1 be a subsection of $\varphi(\gamma_0)$ lying wholly within M_1. Now look at $\varphi(\gamma_1)$ in M_2. Repeat this process along the orbit, producing γ_N lying wholly within the final parallelogram M_N. Then any point lying along γ_N, traced backwards, represents an exact orbit that stays within M_i, $\;i = N, N-1, \ldots, 1, 0$, and we are done (Grebogi et al., 1990).

With this picture, there is a nice geometric interpretation of the requirement that the angle between the stable and unstable directions be bounded away from 0: if the angle gets too small, then the parallelogram essentially loses a dimension, and $\varphi(M_i)$ can not make a "plus sign" with M_{i+1}. Practically speaking, this occurs when the angle becomes comparable with the noise amplitude of the refined orbit. Hence, the more accurate the orbit, the longer it can be shadowed (Grebogi et al., 1990; Quinlan and Tremaine, 1992).

Extending the above argument to three dimensions is fairly straightforward (Hayes, 2001), in which there are precisely two interesting cases:

(i) 1 expanding direction, and 2 contracting (Figure 2). Assume that the z direction is expanding, while the x and y directions are contracting. (We assume, for simplicity of exposition and for ease of drawing, that these three directions are roughly orthogonal, although in practice they need only be resolvable from each other.) Then, analogous to the 2 dimensional argument, we draw *cubes* M_i around the noisy points $\mathbf{y}_i$, and require that $\varphi(M_i)$ maps over M_{i+1} so that φ stretches M_i into a long, thin tube, a segment of which lies wholly in M_{i+1}. Then, precisely as in the 2-dimensional case,

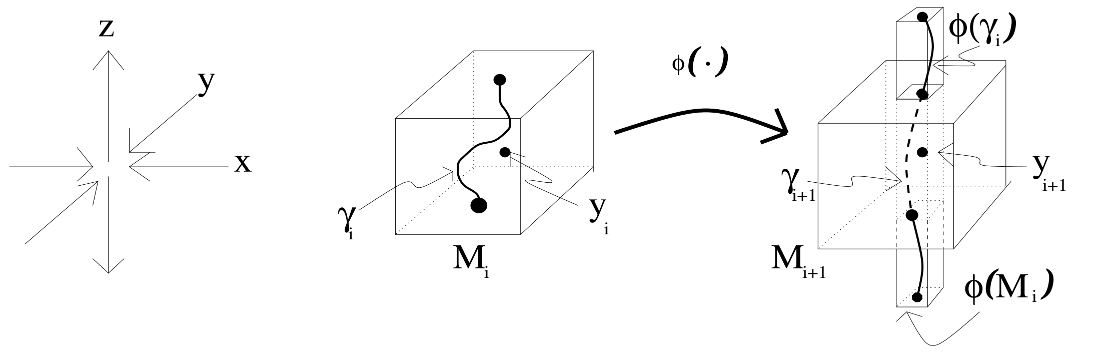

Fig. 2. Containment in 3D, case (i): 1 expanding direction and 2 contracting.

we introduce a curve γ_i that runs approximately along the expanding (vertical) direction from any point on the top of M_i to its bottom. If $\varphi(M_i)$ maps over M_{i+1} as in Figure 2, then we are guaranteed that a contiguous section of $\varphi(\gamma_i)$ lies inside M_{i+1}, connecting its top and bottom along the expanding direction. This becomes γ_{i+1}, and by induction γ_N lies inside M_N, and any point $\mathbf{x}_N$ on it can be traced backwards to a point $\mathbf{x}_i \in M_i$ for $i = 0, 1, \ldots, N-1$.

(ii) 2 expanding and 1 contracting direction (Figure 3). Assume now that the z (vertical) direction is contracting, while the x and y directions are expanding. We again draw a cube M_i around each noisy point $\mathbf{y}_i$, except now $\varphi(M_i)$ maps over M_{i+1} so that φ flattens M_i into a thin slice, cutting M_{i+1} into 3 pieces, the middle piece of which contains a contiguous section of $\varphi(M_i)$. Now, γ_i must be a *surface*, whose boundary connects all of the expanding sides, so that under the mapping, $\varphi(\gamma_i)$ is stretched in all the directions it has extent (both horizontal directions), and is "compressed" along the direction it has measure zero (vertical). Then, we are guaranteed that there is a contiguous segment of $\varphi(\gamma_i)$ lying wholly in M_{i+1} and connecting all of its expanding sides. We call this surface γ_{i+1}, and by induction γ_N lies wholly within M_N, and any point on γ_N, traced backwards to a point $\mathbf{x}_i \in M_i$ for $i = 0, 1, \ldots, N-1$.

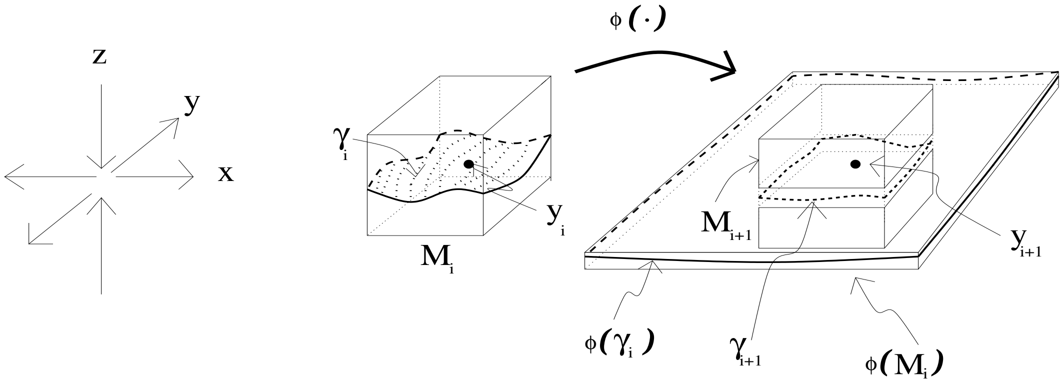

Fig. 3. Containment in 3D, case (ii): 2 expanding directions and 1 contracting.

It seems intuitively clear that we can replace "cube" with "n-cube", "surface" with "manifold", and the above argument still applies in arbitrarily high dimension, although the proof has been elusive (Hayes, 2001; Hayes and Jackson, 2003). The crucial points appear to be that γ has dimension equal to the number of expanding directions, and that its border must "wrap around" all the expanding sides of M_i.

3.2 Refinement

Refinement (Hammel et al. 1987; 1988; Grebogi et al. 1990; Quinlan and Tremaine 1991; 1992; Hayes 1995) is a numerical procedure similar to Newton's Method (and also analogous to iterative improvement methods for solving linear systems (Golub and Van Loan, 1991)) that takes a noisy orbit as input and attempts to produce a nearby orbit with less noise, *i.e.,* one with smaller one-step errors. A refinement iteration is *successful* if before the iteration the trajectory has noise tightly bounded by δ^0, after the iteration it has noise tightly bounded by δ^1, and $\delta^1 < \mu\delta^0$ for some practical $\mu \in [0, 1)$. Otherwise the refinement iteration is *unsuccessful.* Here, a "practical" μ is one that will allow a noisy trajectory to be refined to noise levels near the machine precision in a small number of refinement iterations.

The refinement procedure of GHYS is analagous to Newton's method for finding a zero of a function. GHYS presented their method for the two-dimensional case. (The basic idea was described immediately following Theorem 1.1.) Assume we have a noisy n-dimensional orbit $\mathbf{Y} = \{\mathbf{y}_i\}_{i=0}^N$, $\mathbf{y}_i \in \mathbf{R}^n$, and it has a shadow $\{\mathbf{x}_i\}_{i=0}^N$, $\mathbf{x}_i \in \mathbf{R}^n$. Then $\mathbf{x}_{i+1} = \varphi(\mathbf{x}_i)$ and $\mathbf{y}_{i+1} = \tilde{\varphi}(\mathbf{y}_i) = \varphi(\mathbf{y}_i) + \mathbf{e}_{i+1}$, where $\tilde{\varphi}$ is an approximation to φ with noise bounded by δ. Now suppose we approximate the one-step errors $\mathbf{e}_{i+1} = \mathbf{y}_{i+1} - \varphi(\mathbf{y}_i)$ using a method with noise significantly less than δ. Let $\hat{\mathbf{c}}_i \equiv \mathbf{x}_i - \mathbf{y}_i$ represent a correction term that perturbs $\mathbf{y}_i$ towards $\mathbf{x}_i$. Then

$$\hat{\mathbf{c}}_{i+1} = \mathbf{x}_{i+1} - \mathbf{y}_{i+1} = \varphi(\mathbf{x}_i) - \varphi(\mathbf{y}_i) - \mathbf{e}_{i+1} = D\varphi(\mathbf{y}_i)\hat{\mathbf{c}}_i - \mathbf{e}_{i+1} + O(\|\hat{\mathbf{c}}_i\|^2). \quad (9)$$

In the spirit of Newton's method, we ignore the $O(\|\hat{\mathbf{c}}_i\|^2)$ term, and so one refinement iteration defines the corrections along the entire orbit:

$$\mathbf{c}_{i+1} := D\varphi(\mathbf{y}_i)\mathbf{c}_i - \mathbf{e}_{i+1}. \quad (10)$$

For a discrete map, $D\varphi(\mathbf{y}_i)$ is just the Jacobian of the map at step i. For a system of ODEs, $D\varphi(\mathbf{y}_i)$ is the Jacobian of the solution of the ODE from step i to step $i+1$. For simplicity of explanation, we assume an $n = 2$ dimensional problem for the remainder of this subsubsection. For a generalization to arbitrary n, see Quinlan and Tremaine (1992) or Hayes (1995).

If the problem did not display pseudo-hyperbolicity, then the correction terms $\mathbf{c}_i$ could be computed directly from (10). But since $D\varphi$ displays an approximate exponential dichotomy, it tends to amplify any numerical errors in $\mathbf{c}_i$ not lying in the stable direction. Thus computing the $\mathbf{c}_i$'s by iterating (10) forward will amplify errors and typically produce nothing but noise; iterating backwards suffers the same problem. Therefore, GHYS split the error and correction terms into components in the stable ($\mathbf{s}_i$) and unstable ($\mathbf{u}_i$) directions at each timestep:

$$\mathbf{e}_i = e_{u_i}\mathbf{u}_i + e_{s_i}\mathbf{s}_i, \qquad \mathbf{c}_i = c_{u_i}\mathbf{u}_i + c_{s_i}\mathbf{s}_i. \quad (11)$$

Since it is not known *a priori* which direction is unstable at each timestep, the unstable vector $\mathbf{u}_0$ at time t_0 is initialized to an arbitrary unit vector. The linearized map is then iterated forward with

$$\bar{\mathbf{u}}_{i+1} = D\varphi(\mathbf{y}_i)\mathbf{u}_i, \qquad \mathbf{u}_{i+1} = \bar{\mathbf{u}}_{i+1}/\|\bar{\mathbf{u}}_{i+1}\|. \quad (12)$$

Since $D\varphi(\mathbf{y}_i)$ magnifies any component that lies in the unstable direction, and assuming we are not so unlucky to choose a $\mathbf{u}_0$ that lies too close to the stable direction, then after a few iterations $\mathbf{u}_i$ will point roughly in the unstable direction at t_i. Similarly, the stable unit direction vectors $\mathbf{s}_i$ are computed by initializing $\mathbf{s}_N$ to an arbitrary unit vector and iterating backward with $\bar{\mathbf{s}}_i = D\varphi(\mathbf{y}_i)^{-1}\mathbf{s}_{i+1}$ and $\mathbf{s}_i = \bar{\mathbf{s}}_i/\|\bar{\mathbf{s}}_i\|$. Substituting (11) into (10) yields

$$c_{u_{i+1}}\mathbf{u}_{i+1} + c_{s_{i+1}}\mathbf{s}_{i+1} = D\varphi(\mathbf{y}_i)(c_{u_i}\mathbf{u}_i + c_{s_i}\mathbf{s}_i) - (e_{u_{i+1}}\mathbf{u}_{i+1} + e_{s_{i+1}}\mathbf{s}_{i+1}). \quad (13)$$

While $D\varphi(\mathbf{y}_i)$ magnifies errors in the unstable direction, it damps them in the stable direction. Likewise, $D\varphi(\mathbf{y}_i)^{-1}$ damps errors in the unstable direction and magnifies errors in the stable direction. Thus the c_u terms should be computed backward, and the c_s terms forward. Taking components of (13) in the unstable direction at step $i+1$, we iterate backward on

$$c_{u_i} = (c_{u_{i+1}} + e_{u_{i+1}})/\|\bar{\mathbf{u}}_{i+1}\|, \tag{14}$$

and taking components in the stable direction, we iterate forward on

$$c_{s_{i+1}} = \|D\varphi(\mathbf{y}_i)\mathbf{s}_i\| c_{s_i} - e_{s_{i+1}}. \tag{15}$$

The initial choices for c_{s_0} and c_{u_N} are arbitrary as long as they are small — smaller than the maximum shadowing distance — because (15) damps initial conditions and (14) damps final conditions. GHYS and QT choose them both as 0. This choice is probably as good as any, but it can be seen here that, if one shadow exists, there are infinitely many of them (because a small enough perturbation of the initial condition of one shadow produces another). Another way of looking at these initial choices for c_{s_0} and c_{u_N} is that they "pinch" the growing components at the end point, and the backward-growing components at the initial point, to be small. That is, *boundary conditions* are being forced on the problem so that the exponential divergence is forcibly masked, if possible, making the solution of (10) numerically stable.

The refinement algorithm of GHYS as originally presented (Hammel et al., 1987, 1988; Grebogi et al., 1990) was not rigorous; if it worked at all, it only produced a new pseudo-trajectory with less noise than the original. Refinement was made rigorous by Sauer and Yorke (1991) with the following theorem:

Theorem 3.1 (Sauer and Yorke (1991)) *Let $\mathbf{Y} = \{\mathbf{y}_i\}_{i=0}^N$ be an $n \geq 2$ dimensional δ-pseudo-orbit of the map φ. Assume further that the local stable and unstable subspaces, $\mathbf{S}_i$ and $\mathbf{U}_i$, respectively, at each step are known to a tolerance of δ. Let θ_i be the angle between the stable and unstable subspaces at step i. Let $\|D\varphi(\mathbf{z})\| \leq r_i\|\mathbf{z}\|$ for $\mathbf{z} \in \mathbf{S}_i$, and let $\|D\varphi(\mathbf{z})^{-1}\| \leq t_i\|\mathbf{z}\|$ for $\mathbf{z} \in \mathbf{U}_{i+1}$. Let $C_0 = D_N = 0$, and recursively define $C_{i+1} = \csc\theta_{i+1} + r_i C_i$ for $i = 0, \ldots, N-1$ and $D_{i-1} = \csc\theta_{i-1} + t_{i-1}D_i$ for $i = 1, \ldots, N$. Let B be a bound on $D\varphi, D\varphi^{-1}, D^2\varphi$, and $D^2\varphi^{-1}$. If $\delta < \frac{1}{20n^2}$ and $\max\{C_i, D_i\} \leq \left(n^{5/2}B^2\sqrt{\delta}\right)^{-1}$ for all $i = 0, \ldots, N$, then $\mathbf{Y}$ has an ε-shadow of φ such that $\varepsilon = \sqrt{\delta}$.*

The proof of the theorem (see Sauer and Yorke (1991), Theorem 3.3) is constructive, in the sense that it uses the procedure for refining noisy orbits originally given in Hammel et al. (1988). The essential point of the proof is to show that under the conditions of the theorem, the iterated application of the refinement procedure beginning with the pseudo-orbit results in a sequence of refined pseudo-orbits with decreasing noise level whose limit is an exact orbit. Furthermore, the exact orbit is not too far from the original pseudo-orbit.

Sauer and Yorke (1991) considered this theorem as a justification for the non-rigorous refinement procedure. Conversely, QT argued that if the refinement algorithm fails then there is good reason to believe that no shadow exists, for two reasons. First, from the more rigorous study of simpler systems, glitches are known to exist and are not just a failure of any particular refinement algorithm. Second, QT's results are consistent with a conjecture by GHYS on the frequency of glitches. However, there is no guarantee that refinement converges towards an exact orbit. In fact, even if some refinements are successful, numerical refinement alone does not prove that an *exact* shadow exists; it only proves the existence of a numerical shadow, *i.e.,* a trajectory with less noise than the original. Hayes (1995) frequently saw cases in which the refinement algorithm failed to find a numerical shadow for noisy orbits of length N, but succeeded in finding a numerical shadow for the superset of length $2N$. Hence, the algorithm failed to find a numerical shadow of length N, even though one clearly exists. On the other hand, refinement can iterate indefinitely without converging and without "blowing up" (Hayes, 1995), implying the local errors cannot be decreased and that no shadow exists. This leads us to ask the question of whether convergence to machine precision is enough: is it possible that refinement, if continued in higher precision, would stop before converging to an exact orbit (Hayes, 1995)? Despite these objections, we believe that refinement to machine precision implies with reasonable probability that a shadow exists whose length is comparable to that of the numerical shadow, although this evidence should not be taken as conclusive.

If one is simply interested in studying high-dimensional systems, then a chaotic map would be a more efficient test problem than an ODE system, because no variational equation integration is needed. Hayes (1995) and Hayes and Jackson (1996) list several optimizations to the procedure that increase the speed of GHYS/QT refinement by about two orders of magnitude. We note that the GHYS/QT refinement algorithm is trivially parallelizable, since the computation of each $D\varphi(\mathbf{y}_i)$ is completely independent of all the others. For the same reason, it also has excellent locality of reference in a serial implementation, so virtual memory paging is minimized. Finally, we note that $D^2\varphi$ has $O(n^3)$ elements so, unless significant sparsity is present, actually applying Theorem 3.1 is impractical for any but small n.

3.3 Results by bounding non-hyperbolicity

The procedures of containment and refinement do not make explicit use of the hyperbolicity of the system, although they work only if some measure of hyperbolicity is present (Chow and Palmer, 1991). In contrast, Chow and Palmer 1991; 1992 make explicit use of the hyperbolicity of the system, and use the ideas of the traditional Shadowing Lemma (Anosov, 1967; Bowen, 1975; Palmer, 1988) to estimate how far a shadow is from a pseudo-orbit. Chow and Palmer (1991) discussed the one-dimensional case, and Hadeler (1996) made explicit the relationship between the one-dimensional case and Kantorovich's

Theorem, which lays out conditions under which Newton's method will converge. We omit detailed discussion of the one-dimensional case because later work by the same authors (Chow and Palmer, 1992) subsumes it, except to note one very interesting fact: Chow and Palmer (1991) proved that in the one-dimensional case, the shadowing distance not only has an upper bound, but a *lower* bound as well. That is, they proved that the shadow must maintain a minimum distance from the noisy orbit; it cannot approach the noisy orbit arbitrarily closely. It is not clear if this result is extendible to higher dimensions, nor is it clear exactly what the significance of this result is; however, it is certainly interesting. Their multi-dimensional theorem follows.

Let $\varphi : \mathbf{R}^n \to \mathbf{R}^n$ be a C^2 function and let $\{\mathbf{y}_i\}_{i=0}^N$ be a δ-pseudo-orbit of φ. Given any sequence $(\mathbf{h}_i)_{i=0}^{N-1}$ in $\mathbf{R}^{nN}$, the difference equation $\mathbf{z}_{i+1} = D\varphi(\mathbf{y}_i)\mathbf{z}_i + \mathbf{h}_i$ has many solutions. So the linear operator $L : \mathbf{R}^{n(N+1)} \to \mathbf{R}^{nN}$ defined for $\mathbf{Z} = \{\mathbf{z}_i\}_{i=0}^N$ by $(L\mathbf{Z})_i = \mathbf{z}_{i+1} - D\varphi(\mathbf{y}_i)\mathbf{z}_i$ is onto and so has right inverses. For the following theorem, we choose any such right inverse.

Theorem 3.2 *Let $\varphi : \mathbf{R}^n \to \mathbf{R}^n$ be a C^2 function and let $M = \sup\{\|D^2\varphi(\mathbf{x})\| : \mathbf{x} \in \mathbf{R}^n\}$. Let $\{\mathbf{y}_i\}_{i=0}^N$ be a δ-pseudo-orbit of φ with $2M\|L^{-1}\|^2\delta \leq 1$, where L^{-1} is a right inverse of L. Then there is an exact orbit $\{\mathbf{x}_i\}_{i=0}^N$ of φ such that* $\|\mathbf{x}_i - \mathbf{y}_i\| \leq \frac{2\|L^{-1}\|\delta}{1+\sqrt{1-2M\|L^{-1}\|^2\delta}}, \qquad i = 0, \ldots, N.$

Remark: It is not necessary to assume that $D^2\varphi(\mathbf{x})$ is bounded over $\mathbf{R}^n$ because usually $\mathbf{y}_i$ would be restricted to a bounded set and M could be replaced by a bound for $\|D^2\varphi(\mathbf{x})\|$ over that set (Chow and Palmer, 1992). Although the proof of the theorem (Chow and Palmer, 1992) is short and elegant, the result is unfortunately of little practical use for large problems if no closed form for $D^2\varphi$ is available, because M can be extremely expensive to compute numerically.

If δ is the local error made in computing the orbit, then $\|L^{-1}\|$ is called th "magnification factor", because $\|L^{-1}\|\delta$ is approximately the distance to the shadow. Thus, the next step is to choose L^{-1} in such a way that $\|L^{-1}\|$ is minimized. Not surprisingly, the best L^{-1} to choose is one whose components are as aligned as possible with the stable and unstable subspaces at each step, computed in a fashion similar to the refinement algorithm. Finally, computing an upper bound for $\|L^{-1}\|$ involves noting that even though the orbit $\{\mathbf{x}_i\}$ is not hyperbolic under φ, it may be hyperbolic under φ^p (φ iterated p times) for some integer $p > 1$. If such a p is found, it allows explicit bounds to be computed on the hyperbolicity constants for the orbit $\{\mathbf{x}_i\}$ under φ^p using the ideas of the traditional Shadowing Lemma (Anosov, 1967; Bowen, 1975; Palmer, 1988), leading to an upper bound on $\|L^{-1}\|$. Chow and Palmer demonstrate their method on a δ-pseudo-orbit of the Hénon map with $\delta = 2^{-54} \approx 10^{-15.5}$. For a particular orbit of $N = 333,000$ iterates of the map, they find that $p = 40$ guarantees hyperbolicity of the orbit under φ^p and that

$\|L^{-1}\| \leq 113277 \approx 10^5$. This means that the shadowing distance is about 10^5 times the size of the one-step errors, giving a shadow distance of about $10^{-10.5}$.

This "magnification factor", the ratio between the shadow distance and the local error, is termed the "brittleness" of an orbit by Dawson et al. (1994) (also called "modulus of continuity" and "condition number"). If the brittleness is of order the inverse of the machine epsilon or larger, then all accuracy is lost as the shadowing error is comparable to the size of the variables themselves. They show that if the number of positive and negative Lyapunov exponents changes, or if a Lyapunov exponent fluctuates about zero, then the brittleness can blow up. The effect of a fluctuating exponent is depicted in Figure 4. However,

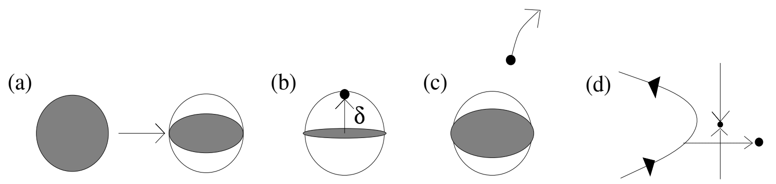

Fig. 4. Fluctuating Lyapunov exponent in the "vertical" direction, reproduced from Dawson *et al.* (1994). δ is the local error, and the vertical direction is initially contracting, but then becomes expanding. (a) an ensemble of trajectories that starts off in an ε-ball is first compressed into a sheet. (b) If the local error steps outside this sheet, and then the direction becomes an expanding direction, then (c) the numerical trajectory diverges away from all exact trajectories that started in the original ε-ball, (d) possibly entering regions of phase space with qualitatively different behaviour than the exact trajectories.

Dawson et al. (1994) make the strong claim that they believe this kind of fluctuating Lyapunov exponent is "common" in high dimensional systems, with the only justification apparently being that there are so many dimensions that there must be a fluctuating exponent *somewhere*. Although this argument is not formally compelling, it may have some merit. On the other hand, Hayes and Jackson (1996) demonstrated numerical shadowing of a 180-dimensional non-hyperbolic system, although that system was artificially constructed to have pseudo-hyperbolicity.

Systems which possess such fluctuating Lyapunov exponents are termed *hyperchaotic* by Sauer et al. (1997). Let $\mathbf{z}_i$ be the displacement from the pseudo-orbit to the shadow at step i. Sauer et al. (1997) observe that the evolution of $\mathbf{z}_i$ with i is similar to a biased random walk. A glitch occurs when the random walk pushes the numerical orbit further away from the shadow than the hyperbolicity can correct for. They model the random walk formally as a *Kolmogorov diffusion process* and demonstrate that the distribution of shadowing distances using this model closely resembles actual shadowing distance distributions. Furthermore, they compute how often glitches occur, based on the behaviour of the fluctuating Lyapunov exponent which is closest to zero. They show that the expected time $\langle \tau \rangle$ for the shadowing distance to become the same size as the variables is proportional to $\langle\tau\rangle \sim \delta^{-2\lambda_0/\sigma_0^2}$, where λ_0 and σ_0 are the mean and standard deviation, respectively, of the fluctuating Lyapunov exponent closest to zero. Finally, they demonstrate that when the fluctuations are sufficiently badly behaved, the length of the shadow is virtually independent of the local error — in other words, in a sufficiently badly behaved system, the shadow length will never get very long for any practical local error.

Methods have also been developed to shadow one-dimensional lattice maps, typically discretizations of partial differential equations (Chow and Van Vleck 1993; 1994b), and for problems that are piecewise hyperbolic in which the number of stable dimensions is monotonically increasing with time (Chow and Van Vleck, 1994a).

4 Shadowing algorithms designed explicitly for ODE systems

There is a fundamental difference between a discrete map, whose shadows were discussed above, and a discrete solution to an ODE. Local errors of the former are restricted to being "space like" — there is no notion of the passage of time between iterations of the map. The latter, however, can have errors in space *as well as time*. The numerical error in the length of each timestep can accumulate, leading the numerical solution to have a slightly different time scale than the real system. In the integration of periodic or almost periodic systems like the solar system, this is also known as phase error, because the numerical solution may have a slightly different period than the exact solution. Thus, although the orbit of a planet may be reproduced correctly by the numerical trajectory, the time at which a real and simulated planet pass through a fixed plane perpendicular to the orbit may differ. This is the case even if the integrator is symplectic (Gonzalez and Stuart, 1999). Thus, when attempting to shadow a numerical solution of an ODE, it may be necessary to "rescale" time (Coomes et al. 1994b; 1995a; 1995b; Van Vleck (1995); Hayes (2001); Hayes and Jackson (2003)). To take this into account, we redefine a shadow of an ODE system as follows:

Definition of ODE shadowing: A pseudo-trajectory $\mathbf{Y} = \{\mathbf{y}_i\}_{i=0}^{N}$ with timesteps $\{h_i\}_{i=0}^{N-1}$ is ε-shadowed by an exact trajectory $\mathbf{X} = \{\mathbf{x}_i\}_{i=0}^{N}$ with timesteps $\{\tau_i\}_{i=0}^{N-1}$ if $\mathbf{x}_{i+1} = \varphi_{\tau_i}(\mathbf{x}_i)$, where $\|\mathbf{y}_i - \mathbf{x}_i\| \leq \varepsilon$, and $|h_i - \tau_i| \leq \varepsilon$.

Remark: In the above definition, we assume that $\varepsilon \ll h_i$, that is, the shadowing distance is significantly smaller than the timesteps. In practice, this appears sufficient for the systems we have studied. If this were not the case, the above definition could be modified to include some notion of global time error per-unit-step.

In other words, the numerical trajectory is shadowed if it closely follows the *path* of an exact solution, but at time t it is allowed to be a little ahead of or behind the exact solution. This linear growth of time errors is due to a lack of hyperbolicity in the direction of the flow in phase space (Van Vleck, 1995). For large $|t - t_0|$ this can be a significant difference, so a shadowing method which does not take the rescaling of time into account is likely to grossly underestimate the length of the shadow. Coomes et al. 1994b; 1995a; 1995b dramatically demonstrate this when they show that a rescaling of time allows the Lorenz equations to be shadowed for almost 10^5 time units, while the *map method*, which does not rescale time, finds shadows lasting only 10 time units—an astounding increase in shadow length of a factor of 10^4!

Finally, note that the non-shadowable example given in the tutorial ($y'' = 0$) *is* shadowable if time is rescaled. This matches what our intuition would say: as long as we care only about qualitative properties of the solution, it should not matter if the numerical trajectory traverses the path at a slightly different velocity than the exact solution, as long as the trajectories, taken as a whole, remain near to each other.

4.1 Explicitly rescaling time in Newton's method

Errors in time manifest themselves as errors directed along the direction of $\mathbf{y}'$, and so one way to account for these errors is to explicitly perturb the noisy solution along the $\mathbf{y}'$ direction. These perturbations translate back into a rescaling of time. To this end, Van Vleck (1995) proves a theorem similar to that of Chow and Van Vleck 1993; 1994b in which time is explicitly added to the variational equation of the one-step error function. To wit, if $\mathbf{Y} = \{\mathbf{y}_i\}_{i=0}^{N}$ is a δ-pseudo trajectory with associated timesteps $\{h_i\}_{i=0}^{N-1}$, then let $\mathbf{z}_i = (\mathbf{y}_i, h_i)$ and $\mathbf{Z} = \{\mathbf{z}_i\}_{i=0}^{N}$ and compute the one-step error by $\mathbf{g}(\mathbf{Z})_i = \mathbf{y}_{i+1} - \varphi_{h_i}(\mathbf{y}_i)$. Then the first variational equation $D\mathbf{g}(Z) : \mathbf{R}^{n(N+1)} \times \mathbf{R}^N \to \mathbf{R}^{nN}$ including the effects of time is

$$
\begin{aligned}
(D\mathbf{g}(\mathbf{Z})\Delta\mathbf{Z})_i &= \Delta\mathbf{y}_{i+1} - \frac{\partial\varphi_{h_i}(\mathbf{y}_i)}{\partial\mathbf{y}_i}\Delta\mathbf{y}_i - \theta\frac{\partial\varphi_{h_i}(\mathbf{y}_i)}{\partial h_i}\Delta h_i \\
&\equiv \Delta\mathbf{y}_{i+1} - \frac{\partial\varphi_{h_i}(\mathbf{y}_i)}{\partial\mathbf{y}_i}\Delta\mathbf{y}_i - \theta f(\varphi_{h_i}(\mathbf{y}_i))\Delta h_i
\end{aligned}
$$

where θ is a user-input parameter controlling the amount of time rescaling which is allowed. (This is the only place in this survey where $D\varphi$ includes a differentiation with respect to h_i. It is this term which allows a rescaling of time by allowing an adjustment along $\mathbf{y}'$.) More formally, we are changing the norm with respect to which the variation is performed: $\theta = 0$ corresponds to the norm in which variations with respect to time are not considered at all, whereas $\theta = 1$ corresponds to the norm in which variations with respect to time are fully considered. Choosing $\theta \in [0, 1]$ allows the scale of variations in time to be *different* from the scale of variations in space, which is precisely what we need in order to perform a rescaling of time. Then we have the following theorem.

Theorem 4.1 (Van Vleck (1995)) *Given constants $\delta, c > 0$ and $\eta \geq 0$ suppose L is an approximation to $D\mathbf{g}(\mathbf{Z})$ such that*

(i) a right inverse L^{-1} of L satisfies $\|L^{-1}\| \leq c$.
(ii) $\|L^{-1} - D\mathbf{g}(\mathbf{Z})^{-1}\| \leq \eta$ for some right inverse $D\mathbf{g}(\mathbf{Z})^{-1}$ of $D\mathbf{g}(\mathbf{Z})$.

Assume that $\|\mathbf{g}(\mathbf{Z})\| \leq \delta$ and let $\varepsilon := 2\delta(\eta + c)$. If $\|D\mathbf{g}(\mathbf{Z}) - D\mathbf{g}(\mathbf{W})\| \leq 1/(2(\eta + c))$ for $\|W - Z\| \leq \varepsilon$, then $\mathbf{g}$ has a solution $\mathbf{W}$ of $\mathbf{g}(\mathbf{W}) = 0$ such that $\|\mathbf{W} - \mathbf{Z}\| \leq \varepsilon$.

For problems that lack hyperbolicity in the direction of motion, Van Vleck (1995) demonstrates that non-zero values of θ are capable of finding shadows between 10 and 100 times longer than if $\theta = 0$, with shadow lengths for the Lorenz system lasting up to about 10^4 time units. However, good values for θ must be found by trial and error, and forcing θ to be constant for all timesteps—which forces the rescaling to be similar across all timesteps—may also restrict its applicability.

4.2 Implicitly rescaling time

Coomes et al. 1994b; 1995a build a hyperplane $\mathbf{H}_i$ perpendicular to $\mathbf{f}(\mathbf{y}_i)$ and containing $\mathbf{y}_i$, and then find a sequence of points $\mathbf{x}_i \in \mathbf{H}_i$ such that $\mathbf{x}_{i+1} = \varphi_{\tau_i}(\mathbf{x}_i)$ and $|\tau_i - h_i| < \varepsilon$. (See Figure 5.) In this way, they avoid having to find τ_i explicitly, as opposed to Van Vleck (1995) who computes

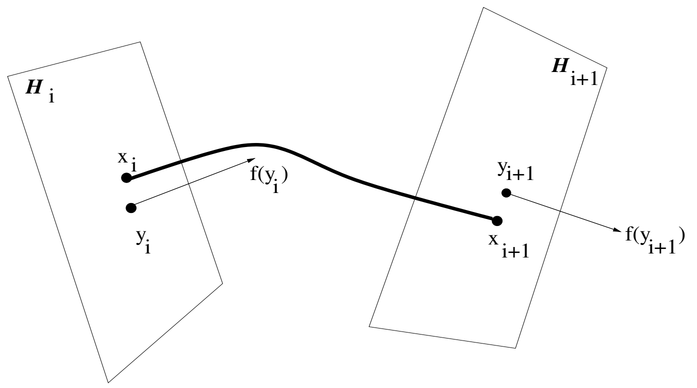

Fig. 5. Pseudo-orbit $\mathbf{y}_i$ and the shadowing orbit $\mathbf{x}_i$ in hyperplane $\mathbf{H}_i$ (Coomes, Koçak and Palmer 1994).

the τ_i explicitly. The statement of their theorem requires some introductory notation.

Let $\mathbf{Y} = \{\mathbf{y}_i\}_{i=0}^N$ be a δ-pseudo orbit with associated stepsizes $\{h_i\}_{i=0}^{N-1}$. Also suppose that we have a sequence $\{Y_i\}_{i=0}^{N-1}$ of $n \times n$ matrices such that $\|Y_i - D\varphi_{h_i}(\mathbf{y}_i)\| \leq \delta, \ \ i = 0, \ldots, N-1$. Now, let S_i be an $n \times (n-1)$ matrix chosen so that its columns form an almost-orthonormal basis for the subspace orthogonal to $\mathbf{f}(\mathbf{y}_i)$, $\|S_i^T \mathbf{f}(\mathbf{y}_i)\| \leq \delta_1, \ \ \|S_i^T S_i - I\| \leq \delta_1$, for some positive number δ_1. Now, we compute $(n-1) \times (n-1)$ matrices A_i satisfying $\|A_i - S_{i+1}^T Y_i S_i\| \leq \delta_1, \ \ i = 0, \ldots, N-1$. Geometrically, A_i is Y_i restricted to the subspace orthogonal to $\mathbf{f}(\mathbf{y}_i)$ and then projected to the subspace orthogonal to $\mathbf{f}(\mathbf{y}_{i+1})$. Next, define a linear operator $L : (\mathbf{R}^{(n-1)})^{(N+1)} \to (\mathbf{R}^{(n-1)})^N$ in the following way: If $\mathbf{\Xi} = \{\xi_i\}_{i=0}^N$ is in $(\mathbf{R}^{(n-1)})^{(N+1)}$, then we take $L\mathbf{\Xi} = \{(L\Xi)_i\}_{i=0}^{N-1}$ where $(L\Xi)_i = \xi_{i+1} - A_i\xi_i, \ \ i = 0, \ldots, N-1$. The operator L has right inverses and we choose one such right inverse L^{-1}. Let U be a convex subset of $\mathbf{R}^n$ containing $\{\mathbf{y}_i\}_{i=0}^N$ in its interior. For such U, we define

$$M_0 = \sup_{\mathbf{x} \in U} \|\mathbf{f}(\mathbf{x})\|, \qquad M_1 = \sup_{\mathbf{x} \in U} \|D\mathbf{f}(\mathbf{x})\|, \qquad M_2 = \sup_{\mathbf{x} \in U} \|D^2\mathbf{f}(\mathbf{x})\|.$$

Then we define

$$\overline{h} = \sup_{0 \leq i \leq N-1} h_i, \qquad \underline{h} = \inf_{0 \leq i \leq N-1} h_i.$$

Next, we choose a positive number $\varepsilon_0 \leq \underline{h}$ such that for $i = 0, \ldots, N-1$ and all $\mathbf{x}$ satisfying $\|\mathbf{x} - \mathbf{y}_i\| \leq \varepsilon_0$, the solution $\varphi_t(\mathbf{x})$ is defined and remains in U for $0 \leq t \leq h_i + \varepsilon_0$. Finally, we define

$$\underline{M}_0 = \inf_{0 \leq i \leq N} \|\mathbf{f}(\mathbf{y}_i)\|, \qquad \overline{M}_0 = \sup_{0 \leq i \leq N} \|\mathbf{f}(\mathbf{y}_i)\|, \qquad \overline{M}_1 = \sup_{0 \leq i \leq N} \|D\mathbf{f}(\mathbf{y}_i)\|, \qquad \Theta = \sup_{0 \leq i \leq N-1} \|Y_i\|.$$

Then, we have the following theorem.

Theorem 4.2 (Coomes et al. (1994b)) *Let*

$$C = \max\left\{\underline{M}_0^{-1}(\Theta\ \|L^{-1}\|\ (1+\delta_1) + 1), \ \ \|L^{-1}\|\sqrt{1+\delta_1}\right\},$$

$$\delta_{\mathbf{H}} = C((M_1 + \sqrt{1+\delta_1})\delta + (3\delta_1(\sqrt{1+\delta_1} + \underline{M}_0^{-1}))/(1 - \delta_1(1 + \underline{M}_0^{-2})),$$

$$\overline{M} = (\overline{M}_1 + M_2\nu\delta)\left(\overline{M}_0 + M_1\nu\delta + 2e^{M_1(\overline{h}+\varepsilon_0)}\sqrt{1+\delta_1}\right) + M_2(\overline{h}+\varepsilon_0)(1+\delta_1)e^{2M_1(\overline{h}+\varepsilon_0)},$$

where $\nu = 2C(e^{M_1(\overline{h}+\varepsilon_0)}\sqrt{1+\delta_1} + M_0)(1-\delta_{\mathbf{H}})^{-1} + 1$. *If these quantities together with* δ, δ_1 *and* ε_0 *satisfy the inequalities (i)* $\delta_1(1 + \underline{M}_0^{-2}) < 1$ *(ii)* $\delta_{\mathbf{H}} < 1$ *(iii)* $2C(1-\delta_{\mathbf{H}})^{-1}\delta\sqrt{1+\delta_1} < \varepsilon_0$ *(iv)* $2\overline{M}C^2(1-\delta_{\mathbf{H}})^{-2}\delta \leq 1$, *then* $\mathbf{Y}$ *is* ε*-shadowed with shadowing distance* $\varepsilon \leq 2C(1-\delta_{\mathbf{H}})^{-1}\delta\sqrt{1+\delta_1}$.

Coomes et al. use a Taylor series integration method with interval arithmetic (see, for example, Nedialkov (1999)) to produce a rigorously bounded local error of their numerical trajectory, and also require the computation of an integer p identical to the p in Chow and Palmer (1992).

As Coomes et al. state, “Admittedly, the statement of the theorem seems rather imposing.” The requirement that the bounds M_0, M_1, and M_2 be over the *entire* convex set U containing the pseudo-trajectory can probably be weakened to be over a “tube” surrounding the trajectory, making it more applicable to problems like the n-body problem that contains poles. The bound on the second derivative of $\mathbf{f}$ over U could be very expensive to compute if a closed form bound is not available. However, requiring bounds on the first and second derivatives of $\mathbf{f}$ is a significant improvement over requiring bounds on the first and second derivatives of φ, as required by Theorem 3.1. For a local error of about 10^{-13}, Coomes et al. were able to find shadows for the Lorenz system lasting 10^5 time units, with a shadowing distance of about 10^{-9}.

4.3 Rescaling time in containment

Hayes (2001) and Hayes and Jackson (2003) use an idea similar to Coomes et al. 1994b; 1995a as depicted in Figure 5 (although the proofs are profoundly different), inspired by the idea of the *Poincaré section*, also known as a *Poincaré map* or *return map*. To rescale time for containment, the idea is to place a plane $\mathbf{H}_i$ in the vicinity of the solution at time t_i, placed so that $\mathbf{H}_i$ is approximately perpendicular to $\mathbf{y}'(t_i)$, and prove that the solution passes through $\mathbf{H}_i$ precisely once during some short time interval surrounding t_i.

To facilitate containment, the idea of the Poincaré section must be generalized to encompass a small ensemble of solutions. To that effect, take a set $M_{i-1} \subset \mathbf{H}_{i-1}$, where the diameter of M_{i-1} is small, and place a plane $\mathbf{H}_i$ in the vicinity of $\varphi_{h_{i-1}}(M_{i-1})$. Then define the Poincaré section of the set $\varphi_{h_{i-1}}(M_{i-1})$ pointwise as follows. Let Δh_{i-1} bound the time interval over which the ensemble $\varphi_{h_{i-1}}(M_{i-1})$ crosses $\mathbf{H}_i$; *i.e.,*

$$\forall \mathbf{x} \in M_{i-1} \exists h \in [h_{i-1} - \Delta h_{i-1}, h_{i-1} + \Delta h_{i-1}] \text{ s.t. } \varphi_h(\mathbf{x}) \in \mathbf{H}_i,$$

assuming that for each $\mathbf{x}$, the h chosen is unique. That is, take the point-by-point Poincaré section of the points in M_{i-1} with respect to the plane $\mathbf{H}_i$. This is called a *splash* operation (Hayes, 2001; Hayes and Jackson, 2003), because we imagine that the points in M_{i-1}, evolving via φ_h for $h \in [h_{i-1} - \Delta h_{i-1}, h_{i-1} + \Delta h_{i-1}]$, “splash” through $\mathbf{H}_i$ approximately simultaneously, and we assume that each trajectory intersects $\mathbf{H}_i$ precisely once during that interval. See Figure 6.

This allows us to build the parallelepipeds M_i inside $\mathbf{H}_i$, and then show that the point-by-point Poincaré section at $\mathbf{H}_i$, *i.e.,* the splash operation, is a home-

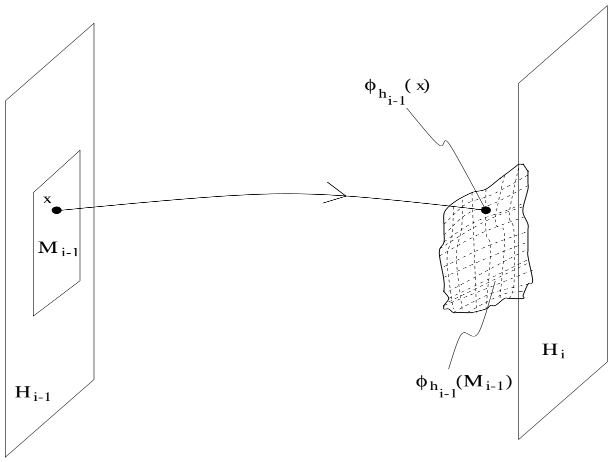

Fig. 6. The "splash" operation depicted for a two-dimensional ensemble evolving in a three-dimensional configuration space. M_{i-1} is embedded in the plane $\mathbf{H}_{i-1}$, and evolves through one timestep to $\varphi_{h_{i-1}}(M_{i-1})$. As depicted, the ensemble is about to splash through $\mathbf{H}_i$.

omorphism. This then allows direct application of the containment theorems detailed earlier. Namely, we apply the containment theorems to the $n-1$-dimensional M_i's which are each contained in the $n-1$-dimensional hyperplane $\mathbf{H}_i$, for an ODE system of n equations.

4.4 Periodic shadowing

The problem of errors in time is exacerbated when attempting to shadow periodic solutions of ODEs, because any non-zero error in time is repeated *ad infinitum*. Thus, a rescaling of time is absolutely necessary to shadow periodic solutions of ODEs. The idea for shadowing periodic solutions is simple. Given a pseudo-trajectory $\{\mathbf{y}_i\}_{i=0}^N$ with timesteps $\{h_i\}_{i=0}^{N-1}$, we require not only that the local error $\|\mathbf{y}_{i+1} - \varphi_{h_i}(\mathbf{y}_i)\|$ is small, but also that $\|\mathbf{y}_0 - \varphi_{h_N}(\mathbf{y}_N)\|$ is small. This gives a periodic pseudo-orbit. Then, only minor modifications are required to non-periodic shadowing theorems to produce a periodic shadowing theorem (Beyn, 1987; Van Vleck, 1995; Coomes et al., 1994a). It is also possible to use refinement-like algorithms to produce accurate pseudo-trajectories from remarkably *in*accurate ones, allowing one to prove the existence of very long periodic trajectories (Coomes et al., 1997).

4.5 Shadowing conservative integrations

Much attention has recently been devoted to integrators that preserve various quantities such as symplectic structure (Channell and Scovel, 1990; Sanz-Serna, 1992) and energy (Gonzalez and Stuart, 1999; Shadwick et al., 1999).

Coomes (1997) demonstrates that such integrations are often shadowable. In particular, if $\mathbf{M}$ is the submanifold of interest (*e.g.,* symplectic manifold or energy surface) on which the initial condition $\mathbf{y}_0$ lies, then a shadow of the pseudo-orbit $\mathbf{Y} = \{\mathbf{y}_i\}_{i=0}^N$ exists in $\mathbf{M}$ if $\mathbf{Y}$ has sufficiently small local error, remains close to $\mathbf{M}$, avoids the neighborhood of fixed points of $\mathbf{f}$, and the variational equation along $\mathbf{Y}$ exhibits sufficient hyperbolicity. This is a very significant result for problems in which such submanifolds occur, most notably Hamiltonian systems.

5 Example application: the gravitational n-body problem

As mentioned previously, a numerical simulation of the gravitational n-body problem is an ideal candidate system to attempt to shadow, because the mathematical model is so well understood and is a virtually perfect description of the real system. In such a system, the existence of a shadow is a strong property: it asserts that a numerical solution can be viewed as an *experimental observation* of an exact solution. As such, within the "observational" error, the dynamics observed in a numerical solution that has a shadow represent the dynamics of an exact solution. Furthermore, most studies of dynamical systems do not aim to predict the precise evolution of a particular choice of initial conditions. Instead, the dynamics of the system is *sampled* in order to study its general behaviour. In such cases, initial conditions are typically choosen from a random distribution and we would be happy if a numerical solution exhibited behaviour typical of *any* valid choice of initial conditions from the distribution. In particular, we may be satisfied if the numerical solution closely follows some exact solution whose initial conditions are close to those that were chosen.

However, as mentioned previously, shadowing is extremely expensive. Thus, shadow searches can only be performed on simplified systems. Quinlan and Tremaine (1992) performed the first experiments on shadowing a single particle moving in a potential of 99 fixed particles, and demonstrated the particle could orbit the system as many times as a full simulation would rotate, while still having a shadow. Hayes (2003) studied how this result changes as the number of moving particles M increases, and found that if the gravitational potential is "softened", shadow durations tend to decrease with M slowly enough that shadows of large N-body simulations may exist for nontrivial durations of time. Hayes (2003) also demonstrated that M *one*-moving-particle systems can be used to approximately predict the shadow duration of an M-moving-particle system. This was done by postulating that a glitch in the large system occurs when *one* of its moving particles encounters a glitch local to its own trajectory, and then showing that the same process can be approximated by superimposing M one-moving-particle trajectories and taking the *minimum* shadow duration of those trajectories. The approximation is excellent for unsoftened systems, and reasonable for softened ones. This reduces the amount of computation required to study shadows of large systems from $O(M^3)$ to $O(M)$, which *greatly* facilitates the study of large systems. (Hayes, 2002) then used this conclusion to perform an extensive study of the dependence of shadow lengths on various simulation parameters, and demonstrated an explicit statistical relationship between simulation accuracy A, the number of particles N being simulated, and average shadow length L. Thus, given N and a desired shadow length L, one can pre-compute the accuracy required for that simulation to have an expected shadow length of L. To this author's knowledge, the n-body problem for large n represents by far the largest system of equations to which shadowing methods have been successfully applied.

6 Are shadows typical of true orbits chosen at random?

The presence of a shadowing orbit does not imply that the statistical properties of a numerical orbit having a shadow are typical of those of true orbits chosen at random; the shadowing orbit might be atypical (Quinlan and Tremaine, 1992). This is a key open question remaining regarding the use of shadows as a measure of error.

For example, consider the binary shift map $x_{i+1} = 2x_i \bmod 1$. Iteration on a computer that uses binary floating point arithmetic always results in $x_i = 0$ after a finite (and relatively small) number of iterations. Although $\{x_i = 0\}_{i=m}^{\infty}$ for some m is a valid exact orbit, it is highly atypical, with misleading statistical properties (Farmer and Sidorowich, 1991). Fryska and Zohdy (1992) proved that numerical simulation of a simple piecewise linear ODE sometimes produces solutions with substantially different statistical properties than the closed-form solution. This idea is taken further by Corless (1994b) (see also Corless (1992a)), who studies the Gauss map,

$$G(x) = \begin{cases} 0, & \text{if } x = 0, \\ x^{-1} \bmod 1, & \text{otherwise.} \end{cases} \tag{16}$$

This well-known map has several properties which make it very interesting, especially from the shadowing viewpoint (Corless, 1994b):

1. The orbit $\{x_i\}$ (where $x_{i+1} = G(x_i),\ i = 0 \ldots \infty$) of every rational initial point x_0 goes to zero in a finite number of iterations. The rationals are dense in [0,1].
2. An orbit is ultimately periodic if and only if it starts from a *quadratic irrational* or, trivially, a rational initial point. Quadratic irrationals are roots

of quadratics with integer coefficients, and are dense in [0,1]. Like the rationals, they are countable, and hence of measure zero. There are an infinite number of orbits with each period.

3. The map is ergodic, meaning almost all initial points have orbits that are dense in [0,1].
4. The Lyapunov exponent of this map is, for almost all initial points, $\pi^2/(6\log 2) \approx 2.3731$, but is *undefined* for rational initial points and is *different* for each quadratic irrational initial point.

As Corless (1994b) states, we see that there are "formidable numerical difficulties in simulating this map." From point #1, we see that unless our numerical orbit goes to zero in a finite number of iterations, it is not representing the properties of the exact orbit starting at our (numerically represented) initial point. Since any numerical orbit must ultimately be periodic, and if our numerical orbit does *not* converge to zero, we see from point #2 that we can only shadow periodic solutions whose initial points are unrepresentable. From point #4 we see that a numerically computed Lyapunov exponent may be completely unrepresentative of almost all orbits. Paradoxically, the numerically computed Lyapunov exponent *does* give a good approximation to the almost-sure value. In fact, a very strong shadowing result can be proved (Corless, 1992a, 1997). However, from point #2, we see that, ultimately, we can shadow only periodic orbits, and thus the shadow that follows our numerical solution has a quadratic irrational initial point, and thus does not have a dense orbit (point #3) or the "correct" Lyapunov exponent. The final resolution of this paradox must account for the fact that the true shadowing orbit behaves like a typical orbit, even though it is not. An analysis of this behaviour is provided by Corless (1994b), based upon Góra and Boyarsky (1988).

On the other hand, Góra and Boyarsky (1988) showed that long pseudo-trajectories of a one-dimensional map τ satisfying some special properties have densities which approach that of τ itself. This is an exciting result, and if it can be generalized to continuous systems of arbitrary dimension, it may go a long way towards answering the question of whether shadows are typical of true orbits chosen at random.

A weak result concerning this question can be abstracted from Coomes et al. (1997). The paper is chiefly concerned with shadowing long periodic orbits, and they use the Lorenz equations as their example. Long-term solutions to the Lorenz system are confined approximately to two disks in three-space, and solutions generally jump between the two disks chaotically. If a revolution around one disk is labelled '0' and a revolution around the other is labelled '1', Coomes et al. (1997) demonstrated that they were able to build pseudo-trajectories with an arbitrary sequence of '0's and '1's, and then prove the existence of periodic shadows for these pseudo-trajectories. This eliminates at least one simple kind of bias: if we assume that true periodic orbits of the Lorenz system chosen at random can produce arbitrary sequences of 0s and 1s, it appears that we can build pseudo-trajectories that possess each sequence, and so shadows of the Lorenz system are not biased in such a way as to disallow certain sequences. PalmerStoffer99 (1999) demonstrate a similar result for the Hénon map.

Note that if shadows are generally atypical of true orbits chosen at random, then the properties of the original pseudo-trajectories that produce the shadows are also atypical. This conclusion would have grave implications for the vast quantities of literature over the past several decades that have studied problems numerically. If otherwise reliable-looking pseudo-trajectories *are* atypical, they must be atypical in an extremely subtle way, because researchers have been making apparently reliable, self-consistent, peer-reviewed conclusions based on numerical simulations for decades. Considering that shadowing is only one of many available methods of error analysis, it would be very surprising (to say the least!) if shadows and their otherwise reliable-looking parent pseudo-trajectories were atypical in a substantial way. This does not mean that the problem should not be studied, of course; the apparently small chance that pseudo-trajectories are substantially atypical is balanced by the importance of proving that they are not.

Finally, we would like to point out that similar criticisms can be levelled against *all* forms of backward error analysis. For example, defect analysis says that the solution obtained by a defect-controlled method is the exact solution to a nearby problem in which the right-hand-side of the ODE suffers a small time-varying perturbation. We can then ask, "Is this slightly perturbed problem typical of nearby problems chosen at random?" Or even more pointedly, we can ask if the perturbations are typical of perturbations suffered by a real-life system? We argue in subsection 1.1 that the answer is sometimes "no". This criticism can also be levelled at the method of modified equations. Even symplectic integrations, which have received much attention recently, suffer the same problem: a solution to a Hamiltonian problem integrated with a symplectic integrator is guaranteed to be exponentially close to the exact solution of a nearby Hamiltonian problem; but is that nearby Hamiltonian problem typical of (pertinent) nearby Hamiltonian problems chosen at random?

This discussion illustrates that answering the question, "Are shadows typical of exact solutions chosen at random?" may be very difficult, and that to be fair, we must ask similar questions of other forms of backward error analysis.

References

Anosov, D. V. (1967). Geodesic Flows and Closed Riemannian Manifolds with Negative Curvature. *Proc. Steklov Inst. Math 90*, 1.

Ascher, U. M., R. M. M. Mattheij, and R. D. Russell (1988). *Numerical Solution of Boundary Value Problems for Ordinary Differential Equations.* Prentice-Hall Series in Computational Mathematics. Prentice-Hall.

Beyn, W.-J. (1987). On Invariant Closed Curves for One-Step Methods. *Numer. Math. 51*, 103–122.

Bowen, R. (1975). ω-Limit Sets for Axiom A Diffeomorphisms. *Journal of Differential Equations 18*, 333.

Braun, M. (1983). *Differential Equations and Their Applications* (3rd ed.). Springer-Verlag.

Channell, P. J. and C. Scovel (1990). Symplectic integration of Hamiltonian Systems. *Nonlinearity 3*, 231–259.

Chow, S. N., X. B. Lin, and K. J. Palmer (1989). A shadowing lemma for maps in infinite dimensions. In C. M. Daffermos, G. Ladas, and G. Papanicolaou (Eds.), *Differential Equations: Proceedings of the EQUADIFF Conference*, pp. 127–136. Marcel Dekker, New York.

Chow, S.-N. and K. J. Palmer (1991). On the numerical computation of orbits of dynamical systems: the one-dimensional case. *Dynamics and Differential Equations 3*, 361–380.

Chow, S.-N. and K. J. Palmer (1992). On the numerical computation of orbits of dynamical systems: the higher dimensional case. *Journal of Complexity 8*, 398–423.

Chow, S.-N. and E. S. Van Vleck (1992). A shadowing lemma for random diffeomorphisms. *Random & Computational Dynamics 1*(2), 197–218.

Chow, S. N. and E. S. Van Vleck (1993). Shadowing of Lattice Maps. In P. E. Kloeden and K. J. Palmer (Eds.), *Chaotic Numerics*, pp. 97–113. American Mathematical Society.

Chow, S.-N. and E. S. Van Vleck (1994a, July). A Shadowing Lemma Approach to Global Error Analysis for Initial Value ODEs. *SIAM Journal on Scientific Computing 15*(4), 959–976.

Chow, S. N. and E. S. Van Vleck (1994b). Shadowing of lattice maps. *Contemporary Mathematics 172*, 97–113.

Clarke, D. A. and M. J. West (Eds.) (1997). *The 12th "Kingston Meeting": Computational Astrophysics*, Volume 123 of *ASP Conference Series.* Astron. Soc. of the Pacific.

Coomes, B. A. (1997, January). Shadowing orbits of ordinary differential equations on invariant submanifolds. *Transactions of the American Mathematical Society 349*(1), 203–216.

Coomes, B. A., H. Koçak, and K. J. Palmer (1994a). Periodic shadowing. In P. Kloeden and K. Palmer (Eds.), *Chaotic Numerics*, Volume 172 of *Contemporary Mathematics*, pp. 115–130. Amer. Math. Soc., Providence, RI.

Coomes, B. A., H. Koçak, and K. J. Palmer (1994b). Shadowing orbits of ordinary differential equations. *Journal of Computational and Applied Mathematics 52*, 35–43.

Coomes, B. A., H. Koçak, and K. J. Palmer (1995a). Rigorous computational shadowing of orbits of ordinary differential equations. *Numerische Mathematik 69*, 401–421.

Coomes, B. A., H. Koçak, and K. J. Palmer (1995b). A shadowing theorem for ordinary differential equations. *Z angew Math Phys (ZAMP) 46*, 85–106.

Coomes, B. A., H. Koçak, and K. J. Palmer (1997). Long periodic shadowing. *Numerical Algorithms 14*, 55–78.

Corless, R. M. (1992a). Continued fractions and chaos. *American Mathematical Monthly 99*(3), 203–215.

Corless, R. M. (1992b). Defect-controlled numerical methods and shadowing for chaotic differential equations. *Physica D 60*, 323–334.

Corless, R. M. (1994a). Error Backward. *Contemporary Mathematics 172*, 31–62.

Corless, R. M. (1994b). What good are numerical simulations of chaotic dynamical systems? *Computers Math. Applic. 28*(10–12), 107–121.

Corless, R. M. (1997). Continued fractions and chaos. In *Canadian Mathematical Society Conference Proceedings*, Volume 20. Reprinted from *American Mathematical Monthly* 99 (1992), no. 3, 203–215.

Corless, R. M. and G. F. Corliss (1992). Rationale for guaranteed ODE defect control. In L. Atanassova and J. Herzberger (Eds.), *Computer Arithmetic and Enclosure Methods*. Elsevier Science Publishers B. V. (North-Holland). 1992 IMACS.

Dahlquist, G. and Å. Björck (1974). *Numerical Methods*. Prentice-Hall series in Automatic Computation. Prentice-Hall.

Dawson, S., C. Grebogi, T. Sauer, and J. A. Yorke (3 Oct 1994). Obstructions to Shadowing When a Lyapunov Exponent Fluctuates about Zero. *Physical Review Letters 73*(14), 1927–1930.

Enright, W. H. and W. B. Hayes (2003). Robust defect control for continuous rk-methods with high-order interpolants. Submitted.

Farmer, J. D. and J. J. Sidorowich (1991). Optimal shadowing and noise reduction. *Physica D 47*, 373–392.

Fryska, S. T. and M. A. Zohdy (1992). Computer dynamics and the shadowing of chaotic orbits. *Physics Letters A 166*, 340–346.

Golub, G. H. and C. F. Van Loan (1991). *Matrix computations*. Johns Hopkins University Press.

Gonzalez, O. and D.J. Higham and A. M. Stuart (1999). On the Qualitative Properties of Modified Equations. IMA J. Num. Anal. **19**, 169-190.

Góra, P. and A. Boyarsky (1988). Why computers like Lebesgue measure. *Comput. Math. Applic. 16*(4), 321–329.

Grebogi, C., S. M. Hammel, J. A. Yorke, and T. Sauer (1990, 24 September). Shadowing of Physical Trajectories in Chaotic Dynamics: Containment and Refinement. *Physical Review Letters 65*(13), 1527–1530.

Hadeler, K. P. (1996). Shadowing orbits and Kantorovich's theorem. *Numerische Mathematik 73*, 65–73.

Hairer, E., S. P. Nørsett, and G. Wanner (1993). *Solving Ordinary Differential Equations* (2nd ed.). Springer-Verlag. Two volumes.

Hammel, S. M., J. A. Yorke, and C. Grebogi (1987). Do Numerical Orbits of Chaotic Dynamical Processes Represent True Orbits? *Journal of Complexity 3*, 136–145.

Hammel, S. M., J. A. Yorke, and C. Grebogi (1988). Numerical Orbits of Chaotic Dynamical Processes Represent True Orbits. *Bull. Am. Math. Soc. 19*, 465–470.

Hayes, W. (1995, January). Efficient shadowing of high dimensional chaotic systems with the large astrophysical n-body problem as an example. Master's thesis, Dept. of Computer Science, University of Toronto.

Hayes, W. (2002). A shadowing-based timestep criterion for softened gravitational n-body simulations. In preparation.

Hayes, W. (2003, 7 February). Shadowing high-dimensional Hamiltonian systems: the gravitational n-body problem. *Physical Review Letters 90*(5).

Hayes, W. and K. R. Jackson (1996). A Fast Shadowing Algorithm for High Dimensional ODE Systems. Unpublished.

Hayes, W. B. (2001). *Rigorous shadowing of numerical solutions of ordinary differential equations by containment.* Ph. D. thesis, Department of Computer Science, University of Toronto. Available on the web as `http://www.cs.toronto.edu/NA/reports.html#hayes-01-phd`.

Hayes, W. B. and K. R. Jackson (2003). Rigorous shadowing of numerical solutions of ordinary differential equations by containment. *SIAM J. Numer. Anal. 00*, 00–00. In press.

Kahaner, D., C. Moler, and S. Nash (1989). *Numerical Methods and Software.* Prentice-Hall series in Computational Mathematics. Prentice-Hall.

Merritt, D. (1999, February). Elliptical galaxy dynamics. *Publications of the Astronomical Society of the Pacific 111*(756), 129–168.

Merritt, D., J. A. Sellwood, and M. Valluri (Eds.) (1999). *Galaxy Dynamics: a Rutgers Symposium*, Volume 182 of *ASP Conference Series.* Astron. Soc. of the Pacific.

Merritt, D. and M. Valluri (1996). Chaos and mixing in triaxial stellar systems. *The Astrophysical Journal 471*, 82–105.

Murdock, J. (1995). Shadowing multiple elbow orbits: an application of dynamical systems to perturbation theory. *J. Diff. Eq. 119*(1), 224–247.

Nedialkov, N. S. (1999). *Computing rigorous bounds on the solution of an initial value problem for an ordinary differential equation.* Ph. D. thesis, Department of Computer Science, University of Toronto.

Palmer, K. J. (1988). Exponential Dichotomies, the Shadowing Lemma and Transversal Homoclinic Points. In U. Kirchgraber and H. O. Walther (Eds.), *Dynamics Reported*, Volume 1. Wiley and Teubner.

Palmer, K. J. (2000). *Shadowing in Dynamical Systems - Theory and Applications.* Kluwer Academic.

Palmer, K. J. and D. Stoffer (1999). Rigorous verification of Chaotic Behaviour of Maps using Validated Shadowing . *Nonlinearity 12*, 1683–1698.

Pilyugin, S. Y. (1999). *Shadowing in dynamical systems.*

Quinlan, G. D. and S. Tremaine (1991). Shadow orbits and the gravitational N-body problem. In B. Sundelius (Ed.), *Dynamics of Disc Galaxies*, pp. 143–148. Göteborg, Sweden.

Quinlan, G. D. and S. Tremaine (1992). On the reliability of gravitational N-body integrations. *Monthly Notices of the Royal Astronomical Society 259*, 505–518.

Sanz-Serna, J. M. (1992). Symplectic integrators for Hamiltonian problems: an overview. In A. Iserles (Ed.), *Acta Numerica 1992*, pp. 243–286. Cambridge University Press.

Sauer, T., C. Grebogi, and J. A. Yorke (7 July 1997). How long do numerical chaotic solutions remain valid? *Physical Review Letters 79*(1), 59–62.

Sauer, T. and J. A. Yorke (1991). Rigorous Verification of Trajectories for the Computer Simulation of Dynamical Systems. *Nonlinearity 4*, 961–979.

Shadwick, B. A., J. C. Bowman, and P. J. Morrison (1999). Exactly conservative integrators. *SIAM Journal on Applied Mathematics 59*(3), 1112–1133.

Skeel, R. (1996). The Meaning of Molecular dynamics. Unpublished, e-mail skeel@@cs.uiuc.edu.

Skeel, R. D. (1999). Integration schemes for molecular dynamics and related applications. In M. Ainsworth, J. Levesley, and M. Marletta (Eds.), *The Graduate Student's Guide to Numerical Analysis*, SSCM, pp. 119–176. Springer-Verlag.

Smale, S. (1965). Diffeomorphisms with many periodic points. In S. Cairns (Ed.), *Differential and Combinatorial Topology*, pp. 63–80. Princeton Univ. Press.

Smale, S. (1967). Differentiable dynamical systems. *Bull. Amer. Math. Soc. 73*, 747–817.

Struck, C. (1997). Galaxy splashes: The effects of collisions between gas-rich galaxy disks. See Clarke and West (1997), pp. 225–230.

Van Vleck, E. S. (1995). Numerical Shadowing Near Hyperbolic Trajectories. *SIAM Journal on Scientific Computing 16*(5), 1172–1189.