NUMERICAL SHADOWING NEAR HYPERBOLIC TRAJECTORIES*

ERIK S. VAN VLECK†

Abstract. Shadowing is a means of characterizing global errors in the numerical solution of initial value ordinary differential equations by allowing for a small perturbation in the initial condition. The method presented in this paper allows for a perturbation in the initial condition and a reparameterization of time in order to compute the shadowing distance in the neighborhood of a periodic orbit or more generally in the neighborhood of an attractor. The method is formulated for one-step methods and both a serial and parallel implementation are applied to the forced van der Pol equation, the Lorenz equation and to the approximation of a periodic orbit.

Key words. global error analysis

AMS subject classification. 65L

1. Introduction. Of practical importance when solving differential equations numerically is to determine a bound or estimate of the global error. There are several approaches to this problem. The classical approach is forward error analysis where one asks that there be a solution of the original problem “close” to the numerically computed solution with the same initial condition. This is the ideal situation, but it does not appear that forward error analysis is applicable to a wide range of physically relevant problems. Another approach is backward error analysis where one asks that there be a solution to a “nearby” problem, i.e., a problem whose vector field is a small perturbation of the vector field of the original problem, “close” to the numerically computed solution. For conservative systems this seems to be the best possible error analysis statement one can hope to obtain. A third approach is to perturb the initial condition while leaving the vector field unperturbed. This is the idea behind shadowing error analysis. The difficulty with this approach is that it does not appear applicable to non-hyperbolic problems such as the case near a periodic orbit or general nontrivial attracting set, i.e., problems for which numerical methods typically rescale time. In this paper we present a method of shadowing error analysis that is applicable in the case of a hyperbolic periodic orbit and general hyperbolic attracting sets. The approach taken in this paper is to allow for reparameterization of time, which is effectively allowing for a special perturbation of the vector field as well as a perturbation of the initial condition. Thus, the method presented here may be thought of as a combination of standard shadowing error analysis and backward error analysis.

Our main contribution is in developing a simple, practical, automatic, parallel shadowing error analysis method while obtaining very accurate approximations of the amplification of the local error that gives the global error. To determine this amplification factor we find and bound the norm of a right inverse of a linear operator that has the form of a “multiple shooting matrix” but without boundary conditions. We choose as a right inverse the pseudo inverse. This choice is independent of the dynamics of the problem and allows for the automatic determination of the norm of a right inverse (the pseudo inverse). We solve directly for the norm of the pseudo inverse of an approximation to the exact linear operator by forming the approximate pseudo inverse a row or rows at a time. The rows of the pseudo inverse can be found quite efficiently using parallel numerical linear algebra techniques. We determine the difference in the norm

---

*Received by the editors February 15, 1994; accepted for publication (in revised form) July 26, 1994. This work was supported in part by the Army Research Office contract DAALO3-89-C-0038 with the University of Minnesota Army High Performance Computing Research Center (AHPARC) and the DoD Shared Resource Center at the AHPARC.

†Department of Mathematical and Computer Sciences, Colorado School of Mines, Golden, Colorado 80401 (evanvlec@mines.colorado.edu).

of the approximate pseudo inverse with respect to the exact pseudo inverse using standard perturbation techniques for the solution of linear systems.

The original shadowing work appears to be due to Anosov and Bowen (see [A] and [Bo]). The first use of numerical shadowing techniques is due to Beyn [Be2] in the numerical analysis community and by Hammel, Yorke, and Grebogi [HYG1], [HYG2] in the dynamical systems community. Subsequently there has been a wealth of study on different ways to do numerical shadowing (see [CP1], [CP2], [CVV1], [CVV2], [SSL], [SY]). These numerical techniques were applied to problems that are hyperbolic, i.e., for problems where there is a splitting into exponentially stable and unstable components. Earlier a shadowing lemma was proved by Franke and Selgrade in 1977 [FS1] for the case near a hyperbolic invariant set and in [FS2] this technique was used to numerically prove the existence of hyperbolic periodic orbits. Recently, Coomes, Kocak, and Palmer [CKP] have proved a shadowing lemma in the spirit of [FS1] using techniques similar to those used in [CLP]. The work of [FS1] and [CKP] applies to the case in which the system is hyperbolic except in the direction of the vector field as is the case near a hyperbolic periodic orbit or hyperbolic attracting set.

The methods that appear in this work may be used to show that numerical approximations of hyperbolic periodic orbits are in fact “close” to an actual periodic orbit of the original system. This subject is considered in [Be2] and [E2] and for more general hyperbolic attracting sets in [KL1], [KL2], and [HLR]. A general error analysis statement that includes both shadowing and backward error analysis is given by Eirola in [E3]. A concept related to backward error analysis is the method of modified equations (see [GSS]). Results in the spirit of those obtained here appear in [E1] and [K] where a posteriori error bounds for two-point boundary value problems are obtained. Reparameterization in time is equivalent to allowing for perturbations in the time steps, so shadowing with reparameterization may facilitate the study of more effective time stepping strategies (see, for example, [SH]).

The outline of this paper is as follows. In §2 the theoretical foundations of our approach is outlined. This includes a formulation for shadowing with reparameterization of time for one-step methods. In §3 the algorithmic details are presented. In §4 we outline the details of our implementations. Section 5 is devoted to examples. We consider an equation with a hyperbolic periodic orbit, the forced van der Pol equation and the Lorenz equation. In §6 we discuss some avenues for future work and the conclusions we have reached from this work.

2. Theoretical foundations. Consider the initial value problem

$$(2.1) \quad \begin{cases} \tilde{x} = f(\tilde{x}), \\ \tilde{x}(0) = x_0 \end{cases}$$

where $\tilde{x}(t) \in \mathbb{R}^N$, $\tilde{x} = \frac{d\tilde{x}}{dt}$, and $f \in C^k(\mathbb{R}^N; \mathbb{R}^N)$ for some $k \geq 2$. Let $\phi : \mathbb{R} \times \mathbb{R}^N \rightarrow \mathbb{R}^N$ be the associated solution operator so that $\phi(0, x_0) = x_0$ and $\phi(t, x_0)$ is the solution at time t with initial condition x_0 . For ease of exposition we consider the formulation of the shadowing technique for autonomous problems, but this can be extended to nonautonomous problems with little modification (see for example [CVV2]).

To solve (2.1) numerically we consider one-step methods of the form

$$(2.2) \quad \begin{cases} x_{n+1} = \Theta(h_n, x_n), \\ x_n \quad \text{given,} \end{cases}$$

where h_n is the current time step.

Given an orbit $x := \{x_n\}_0^M$ and time steps $h := \{h_n\}_0^{M-1}$ produced by a one-step method and a scalar value $\theta \geq 0$ let $z = x \in \mathbb{R}^{(M+1)N} \equiv X$ for $\theta = 0$ and let $z = (x, h) \in$

$\mathbb{R}^{(M+1)N} \times \mathbb{R}^M \equiv X$ for $\theta > 0$, with norms $\|z\| = \sup_n \|x_n\|$ and $\|z\| = \|(x, h)\| = \max\{\sup_n \|x_n\|, \sup_n |\theta^{-1} \cdot h_n|\}$, respectively. The norm of x_n may be with respect to any norm in $\mathbb{R}^N$ although in subsequent sections we will consider the supremum norm explicitly. We set $\theta = 0$ when no reparameterization of time is desired and $\theta > 0$ when time reparameterization is desired. Varying θ for $\theta > 0$ allows for different scales in the time steps versus the values of the orbit x .

Let $Y = \mathbb{R}^{MN}$ with the norm $\|y\| = \sup_n \|y_n\|$ for $y = \{y_n\}_0^{M-1}$. Suppose that $\|z\| < \infty$ and consider the function $F : X \rightarrow Y$ where the n th iterate or component of $F(z)$ is defined for any $z \in X$ to be

$$(2.3) \quad (F(z))_n = x_{n+1} - \phi(h_n, x_n), n = 0, \dots, M-1$$

so that F measures the local error at each iterate. We wish to find a solution $w \in X$ of $F(w) = 0$ in a closed ϵ -neighborhood of our numerical solution z , i.e., a solution of the original initial value problem (IVP) (2.1) but possibly with a slightly different initial condition and a slightly different time step sequence when $\theta > 0$.

Consider the first variation $DF(z) : X \rightarrow Y$ of $F(z)$ defined by

$$(2.4) \quad \begin{aligned} (DF(z)\Delta z)_n &= \Delta x_{n+1} - \phi_{x_n}(h_n, x_n)\Delta x_n - \theta \phi_{h_n}(h_n, x_n)\Delta h_n \\ &\equiv \Delta x_{n+1} - \phi_{x_n}(h_n, x_n)\Delta x_n - \theta f(\phi(h_n, x_n))\Delta h_n \end{aligned}$$

for $n = 0, \dots, M-1$ where $\Delta z = (\Delta x, \Delta h)$, $\phi_{x_n} \equiv \frac{\partial \phi}{\partial x_n}$ and $\phi_{h_n} \equiv \frac{\partial \phi}{\partial h_n}$.

A shadowing theorem. The following theorem is an approximate implicit function theorem (Newton's method) that we use to find a zero of the function F defined in (2.3) given a sufficiently small local error and a bounded right inverse for $DF(z)$ defined in (2.4).

THEOREM 2.1. *Let $F : X \rightarrow Y$ be a C^2 map. Let z be a point in X such that $DF(z)$ has a bounded right inverse $DF(z)^\dagger$ and let $\epsilon_0 > 0$ be chosen so that*

$$(2.5) \quad \|DF(z) - DF(w)\| \leq 1/(2\|DF(z)^\dagger\|)$$

for $\|w - z\| \leq \epsilon_0$. If $0 < \epsilon \leq \epsilon_0$ and

$$(2.6) \quad \|F(z)\| \leq \epsilon/(2\|DF(z)^\dagger\|)$$

then the equation $F(w) = 0$ has a solution w such that $\|w - z\| \leq \epsilon$.

Proof. For the proof see [CLP]. $\square$

We now state a shadowing type theorem that we use in subsequent sections to show that there are true solutions near those we compute numerically.

THEOREM 2.2. *Given constants $\delta, c > 0$ and $\eta \geq 0$ suppose L is an approximation to $DF(z)$ such that*

- (i) *a right inverse $L^\dagger$ of L satisfies $\|L^\dagger\| \leq c$;*
- (ii) *$\|L^\dagger - DF(z)^\dagger\| \leq \eta$ for some right inverse $DF(z)^\dagger$ of $DF(z)$.*

Assume that $\|F(z)\| \leq \delta$ and let $\epsilon := 2\delta(\eta + c)$.

If

$$(iii) \quad \|DF(z) - DF(w)\| \leq 1/(2(\eta + c))$$

for $\|w - z\| \leq \epsilon$, then F defined by (2.3) has a solution w of $F(w) = 0$ such that $\|w - z\| \leq \epsilon$.

Proof. The proof is immediate from Theorem 2.1 with $\epsilon_0 := \epsilon$. $\square$

If the hypotheses of Theorem 2.2 are satisfied then given an orbit $x = \{x_n\}_0^M$ and a corresponding sequence of time steps $\{h_n\}_0^{M-1}$ there exists an orbit $y = \{y_n\}_0^M$ and a sequence of time steps $\{\tau_n\}_0^{M-1}$ such that

$$\|x_n - y_n\| \leq \epsilon \quad \text{and} \quad |h_n - \tau_n| \leq \theta \cdot \epsilon$$

for all n . Furthermore, $y_{n+1} = \phi(\sum_{i=0}^n \tau_i, y_0)$ for $n = 0, 1, \dots, M-1$. For $\theta = 0$ we expect that Theorem 2.2 is applicable when there is a splitting of solution components into stable and unstable modes (see for example [CLP]) or more generally when there is a splitting and the number of stable modes is not decreasing (see [CVV2]). For $\theta > 0$ we expect that Theorem 2.2 is applicable when there is a splitting into stable and unstable modes except in the direction of the vector field (see [FS1] and [CKP]). All that we require to apply the theorem is the existence of a bounded right inverse $L^\dagger$ and values of c , η and δ so that (i)-(iii) in Theorem 2.2 are satisfied. We note that Theorem 2.2 guarantees the existence but not necessarily the uniqueness of a solution w in an ϵ -neighborhood of z .

3. Algorithmic details. In order to apply Theorem 3.2 we must determine an approximation L of $DF(z)$, the norm of a right inverse of L , a bound, η , on the difference in the norms of $DF(z)^\dagger$ and $L^\dagger$, and a bound on the local error δ . In our computation we approximate the local error δ using the local error control mechanism of the numerical integration scheme. We concentrate our efforts on the magnification of the local error, $\|L^\dagger\| + \eta$, that gives the global error. Our goal is to obtain accurate approximations of the norm of a right inverse in the supremum norm. We choose as our right inverse the pseudo inverse. This choice has the advantage that no a priori information about the dynamics, i.e., number of stable modes, etc., is necessary to determine the right inverse. The disadvantage of the choice of the right inverse is that it may be computationally expensive to bound its norm. As such, we form the pseudo inverse of L explicitly, rows at a time, using parallel numerical linear algebra techniques. In this way we are able to obtain the supremum norm of the approximate right inverse very accurately so that any crude bounds or estimates appear in the perturbation term, η .

Basic algorithm. Our basic algorithm is the following one.

Algorithm

Step 1. Use a numerical integrator to simultaneously integrate

$$(3.1) \quad \begin{cases} \dot{x} = f(x), \\ \dot{u} = Df(x(t))u \end{cases}$$

from t_j to t_{j+1} with initial data x_j and $u_j = I$ for $j = 0, \dots, M-1$ to obtain x_{j+1} an approximation of $\phi(h_j, x_j)$ and A_j an approximation to $\phi_{x_j}(h_j, x_j)$.

Step 2. Find c such that $\|L^\dagger\| \leq c$.

Step 3. Find a bound η such that $\|DF(x)^\dagger - L^\dagger\| \leq \eta$.

Step 4. If (iii) in Theorem 2.2 is satisfied, then set the global shadowing error to $\epsilon = 2\delta(c + \eta)$.

Pseudo-inverse method. We find an approximation L of the operator $DF(x)$ by numerically integrating the linear variational equation simultaneously with the original nonlinear problem. To form L we employ the approximation A_j of $\phi_{x_j}(h_j, x_j)$ and the approximation $f(x_{j+1})$ of $\phi_{h_j}(h_j, x_j) \equiv f(\phi(h_j, x_j))$, so analogous to the definition of $DF(z)$ in (2.4) we may write the n th iterate of L as

$$(L\Delta z)_n = \Delta x_{n+1} - A_n \Delta x_n - \theta f(x_{n+1}) \Delta h_n,$$

where $\Delta z = (\Delta x, \Delta h)$. Note that L may be thought of as an $MN \times ((M+1)N)$ -dimensional matrix for $\theta = 0$ and a $MN \times ((M+1)N + M)$ -dimensional matrix for $\theta > 0$.

Now let $A = LL^T$ and write $A + \Delta A \equiv DF(z)DF(z)^T$. The matrix A is block tridiagonal for one-step methods, block pentadiagonal for two-step methods, etc. To form $L^\dagger$, the pseudo inverse of L , we solve the equation $A(L^\dagger)^T = L$ since $L^\dagger = L^T(LL^T)^{-1}$. In this way we solve for the pseudo inverse by calculating one or more rows of $L^\dagger$ at a time; i.e. we need not form all of $L^\dagger$ since we only require the supremum norm of $L^\dagger$, but we may solve with multiple columns of L as right-hand sides in parallel. An alternative approach to the approach considered here would be to form $L^\dagger$ as $L^\dagger = QR^{-T}$ where $L^T = QR$ with R upper triangular and Q satisfying $Q^T Q = I_{MN}$, the $MN \times MN$ identity matrix. We have not employed this approach primarily because it does not appear to be as memory efficient.

We employ the following lemma, which is a standard result for perturbation in linear systems (see [GvL, p. 59]).

LEMMA 3.1. *If A is a nonsingular matrix and $\|A^{-1}\Delta A\|_1 = r < 1$, then $A + \Delta A$ is nonsingular.*

Note that $\|A^{-1}\Delta A\|_1 \leq \|A^{-1}\|_1 \cdot \|\Delta A\|_1$, so the hypothesis of Lemma 3.1 is satisfied if $\|A^{-1}\|_1 \cdot \|\Delta A\|_1 < 1$. If x solves $Ax = b$ and y solves $(A + \Delta A)y = (b + \Delta b)$ then

$$\|x - y\|_1 \leq \frac{\|A^{-1}\|_1(\|\Delta b\|_1 + \|\Delta A\|_1 \cdot \|x\|_1)}{1 - \|A^{-1}\|_1 \|\Delta A\|_1}$$

provided $\|A^{-1}\|_1 \|\Delta A\|_1 < 1$. If we have

$$(3.2) \quad \|\Delta A\|_1 \leq \xi_{\Delta A} \quad \text{and} \quad \|\Delta b\|_1 \leq \xi_{\Delta b}$$

for some $\xi_{\Delta A}, \xi_{\Delta b} > 0$ where b takes on the role of the columns of L ; i.e., $b = L_i$ for $i = 1, \dots, (M+1)N + M$, then we can bound η in Theorem 2.2 as

$$(3.3) \quad \eta \leq \frac{\|A^{-1}\|_1(\xi_{\Delta b} + \xi_{\Delta A}\|L^\dagger\|_\infty)}{1 - \|A^{-1}\|_1 \xi_{\Delta A}}.$$

We estimate $\|A^{-1}\|_1$ numerically using the condition number estimation code of Higham (see [Hi]) and find $\|L^\dagger\|_\infty$ during our computations, so the only quantities we must find bounds for are $\xi_{\Delta A}$ and $\xi_{\Delta b}$. We do so in the following calculations.

Estimates for one-step methods. We make repeated use of the following lemma.

LEMMA 3.2. *If $\|x_i - y_i\| \leq \epsilon$, then*

$$(3.4) \quad \|\phi(t, x_i) - \phi(t, y_i)\| \leq \epsilon \cdot \exp(t \cdot L_f)$$

and

$$(3.5) \quad \|\phi_{x_i}(t, x_i) - \phi_{y_i}(t, y_i)\| \leq \epsilon \frac{L_{Df}}{L_f} (\exp(t \cdot L_f) - 1)$$

where L_f and L_{Df} are local Lipschitz constants for f and Df in a $\epsilon \exp(t \cdot L_f)$ -neighborhood of the local solution $\phi(t, x_i)$.

Proof. This is a simple application of Gronwall's inequality (see [Ha, p. 36]).

Using Lemma 3.2 we replace (iii) in Theorem 2.2 by

$$(3.6) \quad \epsilon \cdot \sup_n \left\{ \theta [L_f^{(n)}(C^{(n)} + 1) + \theta \cdot B_f^{(n)}] + \frac{L_{Df}^{(n)}}{L_f^{(n)}} (C^{(n)} - 1) \right\} \leq \frac{1}{2(\eta + \|L^\dagger\|)},$$

where $C^{(n)} = \exp(h_n L_f^{(n)})$ and $L_f^{(n)}$, $L_{Df}^{(n)}$ and $B_f^{(n)}$ are the Lipschitz constants for f , Df and a bound for f in a $\epsilon \exp(L_f^{(n)} t)$ -neighborhood of $\phi(t, x_n)$ for $0 \leq t \leq h_n$, respectively. In our computations the constants and the norm of $L^\dagger$ in (3.6) will be with respect to the supremum norm. The left-hand side of (3.6) bounds $\|DF(z) - DF(w)\|$ in (iii) of Theorem 2.2 for $\|w - z\| \leq \epsilon$.

Let $L_{f,p}$ denote the Lipschitz constant for f in an appropriate neighborhood of $\phi(t, x_n)$ with respect to the p -norm. We make use of the following bounds to determine $\xi_{\Delta A}$ and $\xi_{\Delta b}$. For $p = 1, \infty$ we assume

$$(3.7) \quad \|f(\phi(h_n, x_n)) - f(x_{n+1})\|_p \leq L_{f,p} \cdot \delta \quad \text{and} \quad \|\phi_{x_n}(h_n, x_n) - A_n\|_p \leq \delta,$$

where $L_{f,p} = \sup_n L_{f,p}^{(n)}$ so that $\xi_{\Delta b} = \max\{\theta \cdot L_{f,1}\delta, \delta\}$.

Using (3.7) one obtains the bounds

$$(3.8) \quad \begin{aligned} & \|\phi_{x_n}(h_n, x_n)\phi_{x_n}(h_n, x_n)^T - A_n A_n^T\|_1 \\ &= \|\phi_{x_n}(h_n, x_n)\phi_{x_n}(h_n, x_n)^T - [(A_n - \phi_{x_n}(h_n, x_n)) + \phi_{x_n}(h_n, x_n)]A_n^T\|_\infty \\ &\leq \|\phi_{x_n}(h_n, x_n)(\phi_{x_n}(h_n, x_n)^T - A_n^T)\|_\infty + \|(\phi_{x_n}(h_n, x_n) - A_n)A_n^T\|_\infty \\ &\leq \delta\{\|A_n\|_1 + \|A_n\|_\infty + \delta\} =: C_1^{(n)} \end{aligned}$$

and

$$(3.9) \quad \begin{aligned} & \|f(\phi(h_n, x_n))f(\phi(h_n, x_n))^T - f(x_{n+1})f(x_{n+1})^T\|_1 \\ &= \|f(\phi(h_n, x_n))f(\phi(h_n, x_n))^T - [(f(x_{n+1}) - f(\phi(h_n, x_n))) + f(\phi(h_n, x_n))]f(x_{n+1})^T\|_\infty \\ &\leq \|f(\phi(h_n, x_n))(f(\phi(h_n, x_n))^T - f(x_{n+1})^T)\|_\infty + \|(f(\phi(h_n, x_n)) - f(x_{n+1}))f(x_{n+1})^T\|_\infty \\ &\leq L_{f,1}^{(n)}\delta\{\|f(x_{n+1})\|_\infty + L_{f,\infty}\delta\} + L_{f,\infty}\delta\|f(x_{n+1})\|_1 =: C_2^{(n)}. \end{aligned}$$

Then we set

$$(3.10) \quad \xi_{\Delta A} = \sup_n \{2\delta + C_1^{(n)} + \theta^2 \cdot C_2^{(n)}\}.$$

Periodic orbits. To show that there exists a periodic orbit near a computed solution we must consider the operator $F(z)$ in (2.3) with periodic boundary conditions, i.e., with the extra equation $(F(z))_M = x_M - x_0$ added to (2.3). Thus, the value of δ is the maximum of the local error and the difference between x_0 and x_M . Additionally, we add to (2.4) the term $(DF(z)\Delta z)_M = \Delta x_M - \Delta x_0$. In this case the matrix $A = LL^T$ is no longer block tridiagonal for one-step methods, but has the form

$$A = \begin{pmatrix} X & X & 0 & \ddots & X \\ X & X & X & \ddots & \ddots \\ 0 & \ddots & \ddots & \ddots & 0 \\ \ddots & \ddots & X & X & X \\ X & \ddots & 0 & X & X \end{pmatrix}.$$

We employ the Sherman–Morrison–Woodbury formula (see [GvL]) to solve for $L^\dagger$ when using periodic boundary conditions.

4. Implementation details. We have implemented both a serial and a parallel version of our algorithm for computing a numerical solution of an initial value problem combined with an a posteriori estimate of the shadowing global error. Both codes employ the variable order, variable stepsize, extrapolation code ODEX (see [HNW]) that is based upon the explicit Euler and explicit midpoint rules. We have modified the local error control so that the error test is a one-norm test for the system (3.1). Thus, we are able to assume the error bounds in (3.7).

Our serial code was developed and all experiments using it were run on a Silicon Graphics Indigo while our parallel code runs on a Connection Machine CM-5. To determine $\|L^\dagger\|_\infty$ in serial we employ the LAPACK codes DPBTRF and DPBTRS for the Cholesky factorization of a positive definite band matrix and the subsequent solves, respectively. For our parallel code we employed the Connection Machine Scientific Software Library (CMSSL) routines BLOCK_TRIDIAG_FACTOR and BLOCK_TRIDIAG_SOLVE_FACTORED for factoring and solving block tridiagonal linear systems, respectively. Several methods for factoring and solving block tridiagonal linear systems are available and our experiments were done using block cyclic reduction with substructuring. The CMSSL routines allow for the solution of multiple right-hand sides in parallel and in this way we are able to perform the estimate of the global error in parallel in $O(N \cdot M)$ operations (see §5) where M is the number of iterates produced by our initial value solver and N is the number of equations in the original problem. In our implementation we solve for $\lceil \log(N \cdot M)/N \rceil \cdot (N + 1)$ right-hand sides simultaneously (see the “Cost” results in the tables in §5). This is based on the assumption that the cost of solving for multiple right-hand sides of a block tridiagonal system in parallel for an $N \cdot M$ -dimensional system is $O(\log(N \cdot M))$ combined with the need to solve approximately $(N + 1) \cdot M$ right-hand sides.

In our code we provide subroutines to evaluate the original problem and the linear variational equation simultaneously, the one and supremum norm of the first derivative of the vector field at a point, and the supremum norm of the second derivative of the vector field. We only evaluate the first and second derivatives pointwise, but routines to evaluate these derivatives in a neighborhood to obtain rigorous bounds on the quantities used in (3.6)–(3.10) may be provided. These could be made into rigorous bounds using the dense output option of ODEX and then obtaining bounds in a neighborhood of each of these points using Gronwall type estimates. Additional input values are the desired initial condition of the original problem, the selected local error tolerance δ and the value of θ . We have chosen the values of θ used in our numerical results in §5 through experimentation. In practice one would choose $\theta = 0$ when no reparameterization is thought to be necessary and $\theta > 0$ when it was thought that reparameterization of time is necessary. Several strategies are possible for choosing an appropriate positive value of θ . If some knowledge of the ratio of the error in the time step versus the error in the solution values is available, then θ may be chosen based on this. Alternatively, one could choose θ based entirely on providing a well-conditioned matrix A (see for example [Sk]) and the value of θ could depend on the particular block of the diagonal of A that is being computed, i.e., set $\theta = \{\theta_n\}_1^M$. We note that if the chosen value of θ results in a matrix A where $\|A^{-1}\|_1$ is deemed too large, then the matrix A can easily be recomputed for a small cost with a different value of θ since the major cost in our procedure is in the determination of the rows of $L^\dagger$.

5. Numerical results. In this section we present results of our algorithm applied to several nonlinear problems. Since our local error estimates are not rigorous we concentrate on the amplification factor $\|L^\dagger\|$ or $\|L^\dagger\| + \eta$ which is the amplification of twice the local error that gives the global error. Throughout this section “ δ ” denotes our local error tolerance and the global error estimate is given by “ ϵ ” where $\epsilon := 2\delta(\|L^\dagger\|_\infty + \eta)$. In the tables below, if the inequality (3.6) is not satisfied, then we will denote this by putting “-” in the ϵ column and

TABLE 1
Forced van der Pol equation ($\delta = 2.5D - 5$).

Method	θ	T	M	$\ A^{-1}\ _1$	$\ L^\dagger\ _\infty$	η	ϵ	Cost
Serial	0.D0	852.9	1.D+3	2.21D+2	21.83	1.6D0	-	37.57
Serial	1.D-2	852.9	1.D+3	2.20D+2	21.74	1.6D0	-	43.33
Serial	1.D-1	852.9	1.D+3	1.65D+2	15.93	8.7D-1	-	43.33
Serial	1.D0	852.9	1.D+3	2.43d+1	5.37	1.3D-1	2.75D-4	42.69

TABLE 2
Forced van der Pol equation ($\delta = 1.D - 5$).

Method	θ	T	M	$\ A^{-1}\ _1$	$\ L^\dagger\ _\infty$	η	ϵ	Cost
Serial	0.D0	851.9	1.D+3	2.68D+2	21.94	1.0D0	4.59D-4	37.01
Serial	1.D0	851.9	1.D+3	2.33D+1	6.15	5.9D-2	1.24D-4	43.33
Serial	1.D-1	851.9	1.D+x3	1.97D+2	16.74	5.9D-1	3.47D-4	43.58
Parallel	1.D-1	851.9	1.D+3	1.97D+2	16.74	5.9D-1	3.47D-4	3.57
Serial	0.D0	8564.3	1.D+4	2.68D+2	21.94	1.0D0	4.60D-4	343.02
Serial	1.D-1	8564.3	1.D+4	1.97D+2	16.74	5.9D-1	3.47D-4	349.79
Parallel	1.D-1	8546.3	1.D+4	1.97D+2	16.74	5.9D-1	3.47D-4	4.27

a “-” is put in the “ η ” column if $\|A^{-1}\|_1 \cdot \xi_{\Delta A} \geq 1$. The value $\|L^\dagger\|_\infty$ is the norm of the pseudo inverse defined by the mapping corresponding to the approximation of the operator $DF(z)$ while “ η ” measures the difference between $\|L^\dagger\|_\infty$ and $\|DF(z)^\dagger\|_\infty$, the norm of the pseudo inverse corresponding to the local solution operator. Additionally, in our tables we record the values of “ θ ” which describes the norm used in the time rescaling, “ $\|A^{-1}\|_1$,” the norm of the inverse of the symmetric matrix A , and “Cost” which measure the computational cost of determining the error estimate; i.e., $\text{Cost} = (\text{Total CPU Time})/(\text{CPU Time for solving (3.1)})$. Here Cost is measuring total time in parallel computations versus time in parallel computations to solve (3.1) and total time in serial computations versus time in serial computations to solve (3.1) depending on whether the serial or parallel version is being used. We measure the cost in solving (3.1) instead of just the original equation even though the second component of (3.1) is necessary only to obtain our error estimates. The value “ T ” is the final time computed, while “ M ” denotes the number of time steps taken.

The results on the CM-5 are based upon a beta version of the software and, consequently, is not necessarily representative of the performance of the full version of this software. All experiments on the CM-5 were performed using 256 nodes.

Example 5.1. The first example we consider is the forced van der Pol equation

$$\ddot{x} + \alpha(x^2 - 1)\dot{x} + x = \beta \cos(\omega t).$$

We employ the parameter values $\alpha = k = \sigma = 2/5$ where $k = \beta/(2\alpha)$ and $\sigma = (1 - \omega^2)/\alpha$ and the initial condition $(x(0), \dot{x}(0)) = (0, 0)$.

We record the results of our experiments in Tables 1 and 2. From Table 1 we see that the value of θ that is used has a big impact on the norm of A^{-1} and the norm of $L^\dagger$ and thus ultimately on whether or not the inequality (3.6) is satisfied. For $\delta = 2.5D - 5$ the inequality is only satisfied for $\theta = 1$. On the other hand, for $\delta = 1.D - 5$ the inequality (3.6) is satisfied for all of the θ values we tested, although $\|A^{-1}\|_1$ and $\|L^\dagger\|_\infty$ increased as θ decreased. The cost of performing the global error analysis increases with the number of time steps for the serial version of our code but grows very slowly in the parallel version. Note that we are able to successfully shadow even though the average stepsize T/M is relatively large.

FIG. 1. $x - \dot{x}$ plot of computed trajectory for $\delta = 1.D - 5$ and $T = 851.9$.

Figure 1 contains the computed trajectory that was successfully shadowed for $\delta = 1.D - 5$ and $T = 851.9$.

Example 5.2. The Lorenz equation [Lo] is given by

$$\begin{pmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{pmatrix} = \begin{pmatrix} \sigma(y - x) \\ \rho x - xz - y \\ xy - \beta z \end{pmatrix}.$$

We consider the parameter values $\sigma = 10$, $\rho = 28$ and $\beta = 8/3$ and the initial condition is chosen to be $(x_0, y_0, z_0) = (0, 1, 0)$. We choose these values for historical reasons but have obtained similar results for other parameter/initial values.

Table 3 contains some numerical results we obtained for the Lorenz equation. The procedure quickly breaks down when no time reparameterization is requested (i.e., for $\theta = 0$). For this problem reparameterizing time is essential. For 1000 time steps we see that $\theta = 1$ is too large and $\theta = 1.D - 4$ is too small, but $\theta = 1.D - 2$ works quite well and suggests the global error is little more than 10 times the local error. Again we see that the norms of A^{-1} and $L^\dagger$ increase as θ decreases, although for $\theta = 1$ we are not able to satisfy the inequality (3.6). For $\theta = 5.D - 2$ we note that the value of $\|L^\dagger\|_\infty$ stays relatively constant as the number of time steps increases. These results were obtained for a reasonably sized local error tolerance of $\delta = 1.D - 6$. The differences in “ T ” for the serial versus the parallel runs are due to differences in the machine precisions on the two machines. Our results suggest that there exists a true trajectory of the Lorenz equation that visits an ϵ -neighborhood of the computed orbit pictured in Fig. 2 but possibly at a slightly different sequence of times.

Example 5.3. As our final example we wish to use shadowing to show the existence of a periodic orbit of the following problem which has a hyperbolic periodic orbit

$$\begin{pmatrix} \dot{x} \\ \dot{y} \\ \dot{z} \end{pmatrix} = \begin{pmatrix} (1-r)x - y \\ x + (1-r)y \\ z \end{pmatrix},$$

TABLE 3
Lorenz equation ($\delta = 1.D - 6$).

Method	θ	T	M	$\ A^{-1}\ _1$	$\ L^\dagger\ _\infty$	η	ϵ	Cost
Serial	0.D0	1.04	1.D+1	1.91D+3	23.03	8.4D-1	4.78D-5	2.0
Serial	0.D0	1.38	1.2D+1	1.56D+5	525.42	-	-	2.0
Serial	1.D0	117.5	1.D+3	8.75D0	4.89	3.5D0	-	56.15
Serial	1.D-2	117.5	1.D+3	5.17D+1	12.51	1.9D-2	2.51D-5	52.50
Serial	1.D-4	117.5	1.D+3	1.29D+5	553.71	-	-	49.39
Serial	1.D-6	117.5	1.D+3	3.32D+5	876.23	-	-	52.48
Serial	0.D0	117.5	1.D+3	3.32D+5	876.39	-	-	45.19
Serial	5.D-2	117.5	1.D+3	1.26D+1	6.28	1.1D-2	1.25D-5	54.66
Parallel	5.D-2	118.5	1.D+3	1.26D+1	6.28	1.1D-2	1.25D-5	2.96
Serial	5.D-2	1126.3	1.D+4	1.26D+1	6.28	1.1D-2	1.25D-5	413.9
Parallel	5.D-2	1133.8	1.D+4	1.55D+1	6.28	1.3D-2	1.25D-5	3.41
Parallel	5.D-2	11258.9	1.D+5	2.34D+1	6.89	2.2D-2	1.38D-5	6.10

FIG. 2. $x - y$ plot of computed trajectory for $T = 117.5$.

where $r = (x^2 + y^2)^{1/2}$. We consider the initial condition $(x_0, y_0, z_0) = (1, 0, 0)$. We use as a final time for the numerical integration $T = 6.283186$ which is six-digit accurate approximation to the actual period 2π . For this problem we restricted the maximum possible order of ODEX to order eight.

The results for showing the existence of an actual periodic orbit near the numerically computed orbit in Fig. 3 are tabulated in Table 4. These numerical results suggest that there exists a periodic solution of the original IVP with a slightly different period that is within ϵ of the computed orbit in Fig. 3. Time reparameterization is necessary as one would expect when approximating a periodic orbit numerically with a standard method.

Fig. 3. $x - y$ plot of computed trajectory for $\delta = 1.D - 6$ and $T = 6.283186$.

TABLE 4
Periodic orbit ($\delta = 1.D - 6$, $T = 6.283186$, $M = 65$).

Method	θ	$\ A^{-1}\ _1$	$\ L^\dagger\ _\infty$	η	ϵ
Serial	0.D0	3.76D+2	1407.22	2.3D0	-
Serial	1.D-2	3.71D+2	1382.14	2.2D0	-
Serial	1.D-1	1.62D+2	520.12	3.6D-1	1.04D-3
Serial	5.D-1	8.43D+1	181.91	8.18D-2	3.64D-4
Serial	1.D0	8.43D+1	172.62	1.26D-1	3.45D-4

6. Conclusions. The method presented here is a simple, accurate method for determining the shadowing global error with or without time reparameterization. The amplification factor $\|L^\dagger\|_\infty$ is computed quite accurately and all gross estimates are contained in the perturbation term, η . The perturbation term is derived from simple estimates for errors in the solution of linear systems. Because of our choice of the pseudo inverse as our right inverse of choice the method is very automatic and the rescaling of time can easily be modified by changing the value of θ . The method is applicable to initial value problems that are expansive or even non-hyperbolic in some direction; i.e., for problems where forward error analysis is not applicable. The method can be made rigorous by obtaining rigorous bounds on the local error, the errors in the factor and solves used to obtain $L^\dagger$, the norm of A^{-1} and the quantities in (3.6)–(3.10) and roundoff errors. As can be seen from our experiments in §5 the parallel version provides a global error estimate in 2–5 times the cost of integrating the original problem coupled with the linear variational equation. The pseudo inverse method presented in this paper could be made much more efficient if we were able to determine which right-hand side produced the row of the pseudo inverse that gives us the norm of the pseudo inverse.

In future work we intend to explore a practical formulation for shadowing when using multistep methods. To be practical a formulation must not consider a multistep method as a one-step method in a higher dimension. Additionally, we plan to consider shadowing as a means of a posteriori error analysis for time dependent partial differential equations. This will involve developing new techniques since the simultaneous solution of a fundamental matrix solution of the linear variational equation coupled with the original problem is impractical for many partial differential equations.

Acknowledgments. The author is grateful to Luca Dieci and the referees for helpful remarks on an earlier version of this paper.

REFERENCES

- [A] D. V. ANOSOV, *Geodesic flows and closed Riemannian manifolds with negative curvature*, Proc. Steklov Inst. Math., 90 (1967).
- [Be1] W.-J. BEYN, *On the numerical approximation of phase portraits near stationary points*, SIAM J. Numer. Anal., 24 (1987), pp. 1095–1113.
- [Be2] ———, *On invariant closed curves for one-step methods*, Numer. Math., 51 (1987), pp. 103–122.
- [Bo] R. BOWEN, *ω -limit sets for axiom A diffeomorphisms*, J. Differential Equations, 18 (1975), pp. 333–339.
- [CKP] B. A. COOMES, H. KOCÁK, AND K. J. PALMER, *A Shadowing Theorem for Ordinary Differential Equations*, preprint, 1993.
- [CLP] S. N. CHOW, X. B. LIN, AND K. J. PALMER, *A shadowing lemma maps in infinite dimensions*, in Differential Equations: Proceedings of the EQUADIFF Conference, C. M. Daffermos, G. Ladas, and G. Papanicolaou, eds., Marcel Dekker, New York, 1989, pp. 127–136.
- [CP1] S. N. CHOW AND K. J. PALMER, *On the numerical computation of orbits of dynamical systems: The one-dimensional case*, Dynamics Differential Equations, 3 (1991), pp. 361–380.
- [CP2] ———, *On the numerical computation of orbits of dynamical systems: The higher dimensional case*, J. Complexity, 8 (1992), pp. 398–423.
- [CVV1] S. N. CHOW AND E. S. VAN VLECK, *Shadowing of lattice maps*, Contemp. Math., 172 (1994), pp. 97–113.
- [CVV2] ———, *A shadowing lemma approach to global error analysis for initial value ODEs*, SIAM J. Sci. Comput., 15 (1994), pp. 959–976.
- [E1] T. EIROLA, *A posteriori error bounds for the back and forth shooting method*, Numer. Funct. Anal. Optim., 9 (1987), pp. 1–18.
- [E2] ———, *Invariant curves of one-step methods*, BIT, 28 (1988), pp. 113–122.
- [E3] ———, *Aspects of backward error analysis in numerical ODEs*, J. Comput. Appl. Math., 45 (1993), pp. 65–74.
- [FS1] J. E. FRANKE AND J. F. SELGRADE, *Hyperbolicity and chain recurrence*, J. Differential Equations, 26 (1977), pp. 27–36.
- [FS2] ———, *A computer method for verification of asymptotically stable periodic orbits*, SIAM J. Math. Anal., 10 (1979), pp. 614–628.
- [GvL] G. GOLUB AND C. VAN LOAN, *Matrix Computations*, Johns Hopkins University Press, Baltimore, MD, 1988.
- [GSS] D. F. GRIFFITHS AND J. M. SANZ-SERNA, *On the scope of the method of modified equations*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 994–1008.
- [Ha] J. K. HALE, *Ordinary Differential Equations*, Krieger, Melbourne, FL, 1980.
- [HLR] J. K. HALE, X.-B. LIN, AND G. RAUGEL, *Upper semicontinuity of attractors of semigroups and partial differential equations*, Math. Comput., 50 (1988), pp. 89–123.
- [HYG1] S. HAMMEL, J. A. YORKE, AND C. GREBOGI, *Do numerical orbits of chaotic dynamical processes represent true orbits?*, J. Complexity, 3 (1987), pp. 136–145.
- [HYG2] ———, *Numerical orbits of chaotic processes represent true orbits*, Bull. Amer. Math. Soc., 19 (1988), pp. 465–470.
- [HNW] E. HAIRER, S. P. NORSETT, AND G. WANNER, *Solving Ordinary Differential Equations I: Nonstiff Problems*, Springer-Verlag, New York, Berlin, 1987.
- [Hi] N. J. HIGHAM, *Algorithm 674: Fortran codes for estimating the one-norm of a real or complex matrix, with applications to condition estimation*, ACM Trans. Math. Software, 14 (1989), pp. 381–396.
- [K] G. KEDEM, *A posteriori error bounds for two-point boundary value problems*, SIAM J. Numer. Anal., 18 (1981), pp. 431–448.

- [KL1] P. E. KLOEDEN AND J. LORENZ, *Stable attracting sets in dynamical systems and in their one-step discretization*, SIAM J. Numer. Anal., 23 (1986), pp. 986–995.
- [KL2] ———, *A note on multistep methods and attracting sets of dynamical systems*, Numer. Math., 56 (1990), pp. 667–673.
- [SSL] J. M. SANZ-SERNA AND S. LARSSON, *Shadows, chaos and saddles*, Appl. Numer. Math., 13 (1993), pp. 181–190.
- [SY] T. SAUER AND J. A. YORKE, *Rigorous verification of trajectories for the computer simulation of dynamical systems*, Nonlinearity, 4 (1991), pp. 961–979.
- [Sk] R. D. SKEEL, *Scaling for numerical stability in Gaussian elimination*, J. Assoc. Comput. Mach., 26 (1979), pp. 494–526.
- [SH] A. M. STUART AND A. R. HUMPHRIES, *An Analysis of Local Error Control for Dissipative, Contractive and Gradient Dynamical Systems*, preprint, 1992.