

# A Fast Shadowing Algorithm for High Dimensional ODE Systems

WAYNE HAYES

wayne@ics.uci.edu

Computer Science Department, University of California, Irvine  
Irvine, CA, 92697-3425, USA.

KENNETH R. JACKSON

krj@cs.toronto.edu

Computer Science Department, University of Toronto  
Toronto, Ontario, Canada M5S 3G4.

## Abstract

Numerical solutions to chaotic dynamical systems are suspect because the chaotic nature of the problem implies that numerical errors are magnified exponentially with time. To bolster confidence in the reliability of such solutions, we turn to the study of *shadowing*. An exact trajectory of a chaotic dynamical system lying near an approximate trajectory is called a *shadow*. Finding shadows of numerical trajectories of systems of ordinary differential equations is very compute intensive, and until recently it has been infeasible to study shadows of higher-dimensional systems. This paper introduces several optimizations to previous algorithms. The optimizations collectively achieve an average speedup of almost 2 orders of magnitude with no detectable loss in effectiveness. The algorithm is tested on systems with up to 180 dimensions. Its application to large gravitational  $N$ -body integrations has already led to a deeper understanding of the reliability of galaxy simulations.

## 1 Introduction

### 1.1 Background

Computer simulation is a popular tool in the modern physical scientist's study of complex dynamical systems. These systems often display *sensitive dependence on initial conditions*: small changes in initial conditions produce solutions that exponentially diverge from each other. Since a numerical solution introduces small perturbations like roundoff and truncation error, it will generally diverge exponentially away from the exact solution having the same initial conditions. The *Lyapunov exponent*  $\lambda$  of a system defines how quickly nearby trajectories diverge: a small distance  $d_0$  between two nearby solutions increases with time  $t$  approximately as  $d_0 e^{\lambda t}$ . Thus even a single numerical error of size  $d_0$  is magnified exponentially with the passage of time.

These physical systems are usually modelled by a system of *ordinary differential equations* (ODEs). Although there exist other methods of global error analysis for ODEs (eg., Corless (1994)), they always involve demonstrating that the numerical solution exactly solves a problem whose defining equation is slightly perturbed from the one we tried to solve. Shadowing is unique in that it demands that the equation governing the ODE remains fixed. Instead, we allow perturbations only to the initial conditions, in an effort to find a nearby solution which exactly solves the desired ODE while remaining close to the computed solution. The measures of error are the distance between the numerical trajectory and the shadow, and the duration of time over which the two remain close to each other. Although the more traditional backward error approach is useful in some contexts, it is not clear that changing the governing equation is always the best way to measure error, especially if the governing equation satisfies special conditions that, if violated, could produce qualitatively different solutions. For example, a Hamiltonian system may not remain Hamiltonian if small, arbitrary, time-varying perturbations are allowed in its governing equation. An extended discussion of the

motivation for shadowing can be found elsewhere (Hayes 1995; Hayes 2001; Hayes and Jackson 2003; Hayes and Jackson 2004).

The primary difficulty with using shadowing as a measure of error of an ODE integration is that it is *extremely* expensive, requiring expensive, highly accurate integrations of the system, its Jacobian, and its variational equation, as well as the inversion of large matrices. For example, the first study of shadows of the gravitational  $N$ -body problem occurred in the early 1990s (Quinlan and Tremaine 1992). At that time, state-of-the-art galaxy simulations contained up to  $10^5$  particles. Due to the expense of shadowing, Quinlan and Tremaine (1992) were able only to attempt shadowing of a very simplified system of just *one* particle moving in a fixed potential created by 99 particles whose positions were fixed in space. Furthermore, the algorithm described by Quinlan and Tremaine (1992) scales as  $O(M^3)$ , where  $M$  is the number of particles which move ( $M = 1$  in their shadowing simulations). Thus, shadowing can not be used as a practical tool in the study of numerical errors without significant algorithmic improvements.

In this paper we introduce several improvements to the algorithm of Quinlan and Tremaine (1992), called *refinement*, resulting in an average speedup of about two orders of magnitude. This optimized algorithm has been tested on a smooth, almost hyperbolic Hamiltonian system with 180 dimensions whose physical size was scaled to be of order unity. Starting with a numerical trajectory with local error of about  $10^{-4}$ , shadows were successfully found to last about 50 Lyapunov times. Considering that a numerical error of size  $10^{-4}$  is magnified to be comparable to unity after about 9 Lyapunov times (since  $e^9 \approx 10^4$ ), this demonstrates that the algorithm is capable of finding long, high-dimensional shadows. The algorithm has also been successfully applied to gravitational  $N$ -body systems with over 200 phase-space dimensions, leading to a deeper understanding of the reliability of galaxy and cosmological simulations (Hayes 2003b; Hayes 2003a).

## 1.2 Definitions and previous work

We start with some definitions. The terms *orbit*, *trajectory*, and *solution* are used interchangeably. For numerical discussions, we assume a machine precision of approximately 16 decimal digits, and a well-scaled problem where all macroscopic quantities of interest are of order unity. Throughout this paper,  $|\cdot|$  denotes Euclidian norm, and  $\|\cdot\|$  denotes the max norm, although we do not expect the choice of norm to qualitatively change our results.

*Definition.* An *exact trajectory*  $\{\mathbf{x}_i\}_{i=a}^b$  of  $\varphi$  satisfies  $\mathbf{x}_{i+1} = \varphi(\mathbf{x}_i)$  for  $a \leq i < b$ .

We are interested in the case that  $a$  and  $b$  are finite integers. For a chaotic map,  $\varphi$  may be a simple equation, such as the logistic equation  $\varphi(x) = 1 - 2x^2$ , which always maps the interval  $[-1, 1]$  onto itself. For an ODE system like the  $N$ -body problem,  $\varphi(\mathbf{x})$  represents the exact solution passing through  $\mathbf{x}$ , one timestep after  $\mathbf{x}$ .

*Definition.*  $\{\mathbf{p}_i\}_{i=a}^b$  is a  $\delta$ -*pseudo-trajectory*, also called a *noisy orbit*, for  $\varphi$  if  $\|\mathbf{p}_{i+1} - \varphi(\mathbf{p}_i)\| < \delta$  for  $a \leq i < b$ , where  $\delta$  is the noise amplitude.

For a map,  $\delta$  is often the machine epsilon; for an ODE system that has position  $\mathbf{p}_i$  at time  $t_i$ , it is the error per step, also called the *local error* of the numerical integration.

*Definition.* For  $a \leq i < b$ , the *1-step error* made between step  $i$  and step  $i + 1$  of the pseudo-trajectory  $\{\mathbf{p}_i\}_{i=a}^b$  is the vector difference  $\mathbf{e}_{i+1} = \mathbf{p}_{i+1} - \varphi(\mathbf{p}_i)$ .

Thus, an exact trajectory is one whose 1-step errors are identically zero, and a  $\delta$ -pseudo-trajectory is one whose 1-step errors satisfy  $\|\mathbf{e}_i\| < \delta$  for  $a < i \leq b$ .

*Definition of shadowing.* An exact trajectory  $\{\mathbf{x}_i\}_{i=a}^b$   $\varepsilon$ -*shadows*  $\{\mathbf{p}_i\}_{i=a}^b$  if  $\|\mathbf{x}_i - \mathbf{p}_i\| < \varepsilon$  for  $a \leq i \leq b$ .

*Definition.* The pseudo-trajectory  $\{\mathbf{p}_i\}_{i=a}^b$  has a *glitch* at point  $i = G_0 < b$  if for some relevant  $\varepsilon$  there exists

an exact trajectory that  $\varepsilon$ -shadows  $\{\mathbf{p}_i\}_{i=a}^{G_0}$ , but no exact trajectory that  $\varepsilon$ -shadows  $\{\mathbf{p}_i\}_{i=a}^G$ , for  $G > G_0$ .

*Definition.* A *shadow step* is an interval, that can be larger than the internal timestep of a numerical integrator, across which a 1-step error is computed.

*Definition.* *Refinement* is a process, similar to Newton’s Method, that iteratively takes a noisy orbit and tries to produce a nearby orbit with less noise, *i.e.*, one with smaller 1-step errors. A refinement iteration is *successful* if before the iteration the trajectory has noise  $\delta^0$ , after the iteration it has noise  $\delta^1$ , and

$$\delta^1 < \mu\delta^0 \text{ for some reasonable } \mu \in (0, 1). \quad (1)$$

Otherwise the refinement iteration is *unsuccessful*.

Shadowing was first discussed in relation to *hyperbolic* systems by Anosov (1967) and Bowen (1975). In a 2-dimensional hyperbolic system, there are two special directions called the *unstable* (or *expanding*) and the *stable* (or *contracting*) directions, which may vary with time and generally are not orthogonal. Small perturbations along the stable direction decay exponentially in forward time, while small perturbations in the unstable direction grow exponentially in forward time. The two directions reverse rôles in backward time. In addition, the angle between the stable and unstable directions is uniformly bounded away from 0. In such a system, Anosov and Bowen proved that, for sufficiently small noise amplitude, shadows exist for arbitrarily long times.

For non-hyperbolic systems, it seems that we must be satisfied with finite-length shadows. The first studies of shadows for non-hyperbolic systems appear to be by Beyn (1987) and Hammel, Yorke, and Grebogi (1987). In the 1990s, there was a flurry of activity in the study of numerical methods for finding shadows (Chow and Palmer 1991; Chow and Palmer 1992; Chow and Van Vleck 1993; Chow and Van Vleck 1994; Sanz-Serna and Larsson 1993; Sauer and Yorke 1991; Dawson, Grebogi, Sauer, and Yorke 1994; Van Vleck 1995; Quinlan and Tremaine 1992; Hayes 1995). In this paper, we concern ourselves with the refinement procedure of Quinlan and Tremaine (1992), which is an extension of the two dimensional procedure of Grebogi, Hammel, Yorke, and Sauer (1990) to arbitrary-dimensional Hamiltonian systems. These two papers will be referred to as QT and GHYS, respectively.

## 2 The refinement procedure of GHYS and QT

### 2.1 Outline

The refinement procedure of GHYS and QT is similar to Newton’s method for finding a zero of a function. Here we provide a heuristic explanation. Let  $\mathbf{P} = \{\mathbf{p}_i\}_{i=0}^S$  be a trajectory with  $S$  steps that has noise  $\delta > \eta\varepsilon_{mach}$ , where  $\varepsilon_{mach}$  is the machine precision, and  $\eta$  is some constant significantly greater than 1 that allows room for improvement towards the machine precision. Let  $\mathbf{e}_{i+1} = \mathbf{p}_{i+1} - \varphi(\mathbf{p}_i)$  be the 1-step error at step  $i + 1$ , where  $\|\mathbf{e}_{i+1}\| < \delta$  for all  $i$ . The set of all 1-step errors is represented by  $\mathbf{E} = \{\mathbf{e}_i\}_{i=1}^S$ , and is estimated by a numerical integration technique that has higher accuracy than that used to compute  $\mathbf{P}$ . This describes a function, which we call  $g$ , which takes as input the entire orbit  $\mathbf{P}$  and whose output is the set of 1-step errors  $\mathbf{E}$ , *i.e.*,  $g(\mathbf{P}) = \mathbf{E}$ . Since the 1-step errors are assumed to be small,  $\|\mathbf{E}\|$  is small. That is,  $\mathbf{P}$  is close to a zero of  $g$ , if one exists. A zero of the function would represent an orbit in which the 1-step errors are identically zero, *i.e.*, an exact orbit. This is an ideal situation in which to run Newton’s method. Refinement adds a few technical complications that disallow us from directly using Newton’s method, but the ideas are clearly similar, as we shall see below. Just as with Newton’s method, refinement is *iterative*, in that we repeatedly compute and apply a correction, which (ideally) gives an orbit with less noise with each Newton iteration. In practice, we continue iteration until the 1-step errors are close to the machine precision, which we call convergence to a *numerical shadow*. The existence of a numerical shadow supports

the existence of an exact shadow of comparable duration (Quinlan and Tremaine 1992; Sauer and Yorke 1991), although rigorously proving the existence of an exact shadow (Sauer and Yorke 1991; Coomes, Koçak, and Palmer 1994; Hayes and Jackson 2003) is even more expensive than refinement. We do not pursue rigorous proofs here.

## 2.2 The GHYS refinement algorithm

Assume we have a noisy orbit  $\mathbf{P} = \{\mathbf{p}_i\}_{i=0}^S$ , and it has a shadow  $\{\mathbf{x}_i\}_{i=0}^S$ . Then

$$\mathbf{x}_{i+1} = \varphi(\mathbf{x}_i) \text{ and } \mathbf{p}_{i+1} = \tilde{\varphi}(\mathbf{p}_i),$$

where  $\tilde{\varphi}$  is an approximation to  $\varphi$  with a typical accuracy or *noise* of  $\delta$ . Here,  $\delta$  should be the accuracy typically used by practitioners of the problem being studied. Now suppose we compute the 1-step errors

$$\mathbf{e}_{i+1} = \mathbf{p}_{i+1} - \varphi(\mathbf{p}_i),$$

although numerically we use  $\hat{\varphi}$ , a better approximation to  $\varphi$  than  $\tilde{\varphi}$ . Then  $\hat{\mathbf{c}}_i \equiv \mathbf{x}_i - \mathbf{p}_i$  represents a correction term that perturbs point  $\mathbf{p}_i$  towards the shadow. So,

$$\hat{\mathbf{c}}_{i+1} \equiv \mathbf{x}_{i+1} - \mathbf{p}_{i+1} = \varphi(\mathbf{x}_i) - \varphi(\mathbf{p}_i) - \mathbf{e}_{i+1} = \frac{\partial \varphi(\mathbf{p}_i)}{\partial \mathbf{p}_i} \hat{\mathbf{c}}_i - \mathbf{e}_{i+1} + O(|\hat{\mathbf{c}}_i|^2) \quad (2)$$

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

$$\mathbf{c}_{i+1} = L_i \mathbf{c}_i - \mathbf{e}_{i+1}, \quad (3)$$

where  $L_i \equiv \frac{\partial \varphi(\mathbf{p}_i)}{\partial \mathbf{p}_i}$  is the linearized perturbation map. For a discrete map,  $L_i$  is just the Jacobian of the map at step  $i$ . For a system of ODEs,  $L_i$  is the Jacobian of the integral of the ODE from step  $i$  to step  $i+1$ .<sup>1</sup> It is the computation of the linear maps  $L_i$ , called *resolvents* for an ODE system, that takes most of the CPU time during a refinement. If we let  $D$  be the dimension of the problem, then the resolvent has  $O(D^2)$  elements in it, and it takes  $O(D^3)$  time to compute, assuming each element takes time  $O(D)$  to compute, as is the case in the  $N$ -body problem (Hayes 1995). Presumably, if one is interested in studying simpler high-dimensional systems, a chaotic map would be a better choice than an ODE system, because no variational equation integration is needed.

For simplicity of explanation, we assume  $D = 2$  for the remainder of this subsection, deferring discussion of the higher dimensional case to the next subsection. If the problem were not chaotic, the correction terms  $\mathbf{c}_i$  could be computed directly from (3). But since  $L_i$  will amplify any errors in  $\mathbf{c}_i$  not lying in the stable direction, computing the  $\mathbf{c}_i$ 's by iterating (3) forward will amplify errors and typically produce nothing but noise. Therefore, we split the error and correction terms into components in the stable ( $\mathbf{s}_i$ ) and unstable ( $\mathbf{u}_i$ ) directions at each timestep:

$$\mathbf{e}_i = e_{u_i} \mathbf{u}_i + e_{s_i} \mathbf{s}_i, \quad \mathbf{c}_i = c_{u_i} \mathbf{u}_i + c_{s_i} \mathbf{s}_i. \quad (5)$$

---

<sup>1</sup> In other words, let

$$\dot{\mathbf{p}} = \mathbf{h}(t, \mathbf{p}) \quad (4)$$

be the first-order ODE. Note that  $\mathbf{p}_{i+1} = \varphi(\mathbf{p}_i)$  is the solution of (4) using  $\mathbf{p}_i$  as the initial condition and integrating  $\mathbf{h}$  to time  $i+1$ . Then  $J = \frac{\partial \mathbf{h}}{\partial \mathbf{p}}$  is the Jacobian of  $\mathbf{h}(t, \mathbf{p})$ . The Jacobian measures how  $\dot{\mathbf{p}}$  changes if  $\mathbf{p}$  is changed by a small amount. The *resolvent*  $R(t_{i+1}, t_i)$  is the integral of  $J(t)$  along the path  $\mathbf{p}(t)$ , and describes how a small perturbation  $\delta \mathbf{p}$  from  $\mathbf{p}_i$  at time  $t_i$  gets mapped to a perturbation from  $\mathbf{p}_{i+1}$  at time  $t_{i+1}$ . That is,  $R(t_{i+1}, t_i)$  is the solution of the *variational equation*

$$\frac{\partial R}{\partial t} = J(t)R(t, t_i), \quad R(t_i, t_i) = I,$$

where  $I$  is the identity matrix. The reason the arguments to  $R$  seem reversed is for notational convenience: they satisfy the identity  $R(t_2, t_0) = R(t_2, t_1)R(t_1, t_0)$ , and so a perturbation  $\delta \mathbf{p}$  at time  $t_0$  gets mapped to a perturbation at time  $t_2$  by the matrix-matrix and matrix-vector multiplication  $R_2 \delta \mathbf{p} = R_1 R_0 \delta \mathbf{p}$  (Hairer, Nørsett, and Wanner 1993). Finally, the linear map in the GHYS refinement procedure is  $L_i = R(t_{i+1}, t_i)$ .

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

$$\bar{\mathbf{u}}_{i+1} = L_i \mathbf{u}_i, \quad \mathbf{u}_{i+1} = \bar{\mathbf{u}}_{i+1}/|\bar{\mathbf{u}}_{i+1}|. \quad (6)$$

Since  $L_i$  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 Lyapunov times  $\mathbf{u}_i$  will point roughly in the unstable direction at  $t_i$ . Similarly, the stable unit direction vectors  $\mathbf{s}_i$  are computed by initializing  $\mathbf{s}_S$  to an arbitrary unit vector and iterating backward,

$$\bar{\mathbf{s}}_i = L_i^{-1} \mathbf{s}_{i+1}, \quad \mathbf{s}_i = \bar{\mathbf{s}}_i/|\bar{\mathbf{s}}_i|. \quad (7)$$

Substituting (5) into (3) yields

$$c_{u_{i+1}} \mathbf{u}_{i+1} + c_{s_{i+1}} \mathbf{s}_{i+1} = L_i (c_{u_i} \mathbf{u}_i + c_{s_i} \mathbf{s}_i) - (e_{u_{i+1}} \mathbf{u}_{i+1} + e_{s_{i+1}} \mathbf{s}_{i+1}). \quad (8)$$

While  $L_i$  magnifies errors in the unstable direction, it damps them in the stable direction. Likewise,  $L_i^{-1}$  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 (8) in the unstable direction at step  $i+1$ , we iterate backward on

$$c_{u_i} = (c_{u_{i+1}} + e_{u_{i+1}})/|\bar{\mathbf{u}}_{i+1}|, \quad (9)$$

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

$$c_{s_{i+1}} = |L_i \mathbf{s}_i| c_{s_i} - e_{s_{i+1}}. \quad (10)$$

The initial choices for  $c_{s_0}$  and  $c_{u_S}$  are arbitrary as long as they are small — smaller than the maximum shadowing distance — because (10) damps initial conditions and (9) damps final conditions. GHYS and QT choose them both as 0, as do we for simplicity. 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.<sup>2</sup> Another way of looking at these initial choices for  $c_{s_0}$  and  $c_{u_S}$  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 (3) numerically stable.

Note that, counter-intuitively, these boundary conditions demand that the initial condition for the shadow and noisy orbits differ along the unstable direction. In fact, this *must* occur if the change in initial conditions is to have any effect. That is, when looking for a shadow, if perturbations were only allowed in the stable direction, those perturbations would die out, leading the perturbed exact orbit to follow the original exact orbit that passes through the initial conditions — the one that is already known to diverge exponentially from the noisy orbit.

### 2.3 QT’s generalization to arbitrary Hamiltonian systems

If the configuration space is  $D$  dimensional, then there are  $2D$  dimensions in the phase space. It can be shown that in a Hamiltonian system, the number  $B$  of stable and unstable directions is equal, although  $B < D$

---

<sup>2</sup>For any system, *even a chaotic one*, given any exact orbit of fixed length, a small enough perturbation in the initial condition in any direction produces a small change in the final condition, although for chaotic systems this perturbation must be exponentially small in the length of the orbit. (If the perturbation is restricted to the stable subspace, then obviously a similar solution will be obtained.) Thus given any exact orbit that  $\varepsilon$ -shadows a noisy orbit, there exist infinitely many exact orbits nearby that also shadow it. However, it may be that all the exact orbits are packed into a space unresolved by the machine precision.

is possible.<sup>3</sup> At time  $t_i$ , let  $\{\mathbf{u}_i^j\}_{j=1}^D$  represent  $D$  unstable unit vectors, and let  $\{\mathbf{s}_i^j\}_{j=1}^D$  represent  $D$  stable unit vectors. For any particular timestep, it will be convenient if the unstable vectors are orthonormal to each other, and the stable vectors are orthonormal to each other. However, the stable and unstable vectors together will not in general form an orthogonal system.

The vectors are evolved as in the two dimensional case, except using Gram-Schmidt orthonormalization to produce two sets of  $D$ -orthonormal vectors at each timestep. Arbitrary orthonormal bases are chosen,  $\mathbf{u}_0^j$  at time  $t_0$ , and  $\mathbf{s}_S^j$  at time  $t_S$ , and then evolved according to

$$\bar{\mathbf{u}}_{i+1}^j = L_i \mathbf{u}_i^j, \quad \bar{\mathbf{s}}_i^j = L_i^{-1} \mathbf{s}_{i+1}^j.$$

At each step  $i$ , two Gram-Schmidt orthonormalizations are done: one on  $\{\bar{\mathbf{u}}_{i+1}^j\}_{j=1}^D$  to produce  $\{\mathbf{u}_{i+1}^j\}_{j=1}^D$ , and another on  $\{\bar{\mathbf{s}}_i^j\}_{j=1}^D$  to produce  $\{\mathbf{s}_i^j\}_{j=1}^D$ . After a few Lyapunov times,  $\mathbf{u}_i^1$  lies close to the most unstable direction at step  $i$ ,  $\mathbf{u}_i^2$  lies close to the second most unstable direction, *etc.* Likewise,  $\mathbf{s}_i^1$  lies close to the most stable direction at time  $i$ ,  $\mathbf{s}_i^2$  lies close to the second most stable direction, *etc.*

The multidimensional generalization of (5) is the obvious

$$\mathbf{e}_i = \sum_{j=1}^D (e_{u_i^j} \mathbf{u}_i^j + e_{s_i^j} \mathbf{s}_i^j), \quad \mathbf{c}_i = \sum_{j=1}^D (c_{u_i^j} \mathbf{u}_i^j + c_{s_i^j} \mathbf{s}_i^j). \quad (11)$$

To convert the 1-step error at step  $i$  from phase-space co-ordinates  $\mathbf{e}'_i = \{e_i^k\}_{k=1}^{2D}$  to the stable and unstable basis  $\mathbf{e}_i = \left\{ \{e_{u_i^j}\}_{j=1}^D, \{e_{s_i^j}\}_{j=1}^D \right\}$ , one constructs the matrix  $V_i$  whose columns are the unstable and stable unit vectors,  $V_i = (\mathbf{u}_i^1 \mathbf{u}_i^2 \cdots \mathbf{u}_i^D \mathbf{s}_i^1 \mathbf{s}_i^2 \cdots \mathbf{s}_i^D)$ , and solves the system  $V_i \mathbf{e}_i = \mathbf{e}'_i$ .

From (3) and (11) the equation for the correction co-efficients in the unstable subspace at step  $i+1$  is

$$\sum_{j=1}^D (c_{u_{i+1}^j} + e_{u_{i+1}^j}) \mathbf{u}_{i+1}^j = \sum_{j=1}^D c_{u_i^j} L_i \mathbf{u}_i^j$$

which are projected out along  $\mathbf{u}_{i+1}^k$  producing

$$c_{u_{i+1}^k} + e_{u_{i+1}^k} = \sum_{j=k}^D c_{u_i^j} U_i^{kj}$$

where the scalar  $U_i^{kj} = \mathbf{u}_{i+1}^k \cdot L_i \mathbf{u}_i^j = \mathbf{u}_{i+1}^k \cdot \bar{\mathbf{u}}_{i+1}^j$ , and the Gram-Schmidt process ensures  $U_i^{kj} = 0$  if  $j < k$ . As stated previously, the the boundary condition requires the unstable component of the corrections to be small at timestep  $S$ , and their stable components to be small at timestep 0. For simplicity we take  $\{c_{u_S^j} = 0\}_{j=1}^D$  as did QT, and the co-efficients are computed backward using

$$c_{u_i^k} = \frac{1}{U_i^{kk}} \left( c_{u_{i+1}^k} + e_{u_{i+1}^k} - \sum_{j=k+1}^D c_{u_i^j} U_i^{kj} \right).$$

---

<sup>3</sup>It is not hard to show that the Jacobian of the  $N$ -body problem has the form

$$J = \begin{pmatrix} \mathbf{0} & I \\ J_a & \mathbf{0} \end{pmatrix},$$

where  $J_a = \frac{\partial \mathbf{a}}{\partial \mathbf{r}}$  is symmetrical, and has real but not necessarily positive eigenvalues. If  $\mu$  is an eigenvalue of  $J_a$ , then  $\pm\sqrt{\mu}$  are eigenvalues of  $J$ . If  $\mu$  is positive, this gives one expanding and one contracting direction; if  $\mu$  is negative, it gives directions that neither expand nor contract, but "rotate" in some sense.

We first solve for  $\{c_{u^D}\}_{i=0}^S$ , which does not require knowledge of the other  $c_u$  co-efficients, then solve for  $\{c_{u^{D-1}}\}_{i=0}^S$ , *etc.*

Again from (3) and (11), the equation for the correction co-efficients in the stable subspace at timestep  $i$  is

$$\sum_{j=1}^D (c_{s_{i+1}^j} + e_{s_{i+1}^j}) L_i^{-1} \mathbf{s}_{i+1}^j = \sum_{j=1}^D c_{s_i^j} \mathbf{s}_i^j$$

which are projected out along  $\mathbf{s}_i^k$  producing

$$c_{s_i^k} = \sum_{j=k}^D (c_{s_{i+1}^j} + e_{s_{i+1}^j}) S_i^{kj}$$

where  $S_i^{kj} = \mathbf{s}_i^k \cdot L_i^{-1} \mathbf{s}_{i+1}^j = \mathbf{s}_i^k \cdot \bar{\mathbf{s}}_i^j$ , and the Gram-Schmidt process ensures  $S_i^{kj} = 0$  if  $j < k$ . The boundary condition at step 0 is  $\{c_{s_0^j} = 0\}_{j=1}^D$ , and the co-efficients are computed forward using

$$c_{s_{i+1}^k} = \frac{1}{S_i^{kk}} \left( c_{s_i^k} - \sum_{j=k+1}^D (c_{s_{i+1}^j} + e_{s_{i+1}^j}) S_i^{kj} \right) - e_{s_{i+1}^k}.$$

As with the unstable corrections, we first compute  $\{c_{s_i^D}\}_{i=0}^S$  which does not require knowledge of the other  $c_s$  co-efficients, then we compute  $\{c_{s_i^{D-1}}\}$ , *etc.*

Applying the above computed corrections to the noisy orbit, producing a nearby orbit with (ideally) less noise, concludes one refinement iteration.

### 3 Optimizations to the GHYS/QT refinement algorithm

#### 3.1 Brief overview of the optimizations

This paper introduces 8 major optimizations to the GHYS/QT refinement procedure. Each provides a constant speedup of roughly between 1.5 and 8, with a cumulative average speedup of about a factor of 100. We briefly introduce all 8 optimizations here. The first three are trivial and are not discussed further, while the remaining 5 are discussed in more detail in the following subsections. Most of the effort is focussed on reducing both the number of resolvents to recompute, and the time spent recomputing each. This necessitates handling the ensuing complications of when a resolvent *must* be recomputed, and when it would be worthless to do so. Furthermore, each of the optimizations is independent of all the others. It is conceivable that some optimizations may significantly decrease the reliability of refinement for some systems, although this was not observed for the systems we tested.

Each optimization (except the first) has a name, and a single-letter identification for later discussion. We use the acronym *RUS* to mean the set of all *R*esolvents, *U*nstable, and *S*table vectors for the entire trajectory. The *RUS* at a particular timestep  $i$  is referred to as  $\text{RUS}_i$ .

*Definition.* If program  $A$  has a *speedup* of  $\alpha$  over program  $B$ , then the execution time of  $A$  is a fraction  $1/\alpha$  of  $B$ 's time; in other words  $A$  runs  $\alpha$  times faster than  $B$ .

**The 3M Oversight.** QT re-computed a new resolvent for each stable and unstable vector, *i.e.*, at each timestep they computed  $6M$  resolvents. However, a resolvent  $R(t_1, t_0)$  is a matrix operator that applies

to *all* small perturbations of the orbit from time  $t_0$  to  $t_1$ . Thus only 2 resolvents per timestep need to be computed: one for evolving perturbations forward in time, and another to evolve them backward. Thus QT would have been required to perform  $O(M^4)$  work for  $M$  moving particles. Since most of the time is spent computing the resolvents, fixing this oversight provides an immediate speedup of a factor of about  $3M$ . The effect of fixing this oversight is *not* included in the speedup values of Table I below, but it partly explains why QT thought that shadowing is wholly impractical for large  $N$  systems.

**Compute Backward Resolvent by Inverting Forward Resolvent (I).** The above optimization can be extended further by noting that the resolvent matrix  $L_i$  that maps perturbations from time  $t_i$  to  $t_{i+1}$  is precisely the inverse matrix of the resolvent that maps small perturbations from  $t_{i+1}$  to  $t_i$ . Thus, rather than doing an integration from  $t_{i+1}$  backward to  $t_i$  to compute  $L_i^{-1}$ ,  $L_i$  can be inverted using matrix operations, or the system  $L_i \mathbf{x} = \mathbf{b}$  can be solved, whenever necessary, using standard elimination schemes. This gives a speedup of about 2.

**Fixed Tolerance Resolvent (F).** For smooth problems like the  $N$ -body problem, it appears that when a resolvent is recomputed at all, it does not need to be computed very accurately, even in late refinements where the 1-step errors approach the machine precision. This is analogous to computing only an approximate Jacobian when applying Newton’s Method. In these  $N$ -body experiments, when this optimization was used, resolvents were always computed to a fixed tolerance of  $10^{-6}$ . It seems that it is the movement of the orbit during refinement that dictates recomputation of a resolvent, not the numerical error in its construction, just as it is the movement of the estimated zero that dictates recomputation of the Jacobian in Newton’s Method.

**Cheaper “Accurate” Integrator (C).** When computing the resolvent and 1-step errors during early refinements, it is not necessary to compute them to extremely high accuracy. Since the initial 1-step errors may have magnitudes of, say,  $10^{-4}$ , the resolvents and 1-step errors need only be computed to a tolerance of, say,  $10^{-7}$ , *i.e.*, 3 extra digits. QT always computed the 1-step errors and resolvents to a tolerance of  $10^{-13}$ . Of course, if the *Fixed Tolerance Resolvent* optimization is used, this only effects the 1-step error tolerances.

**Constant Resolvent, Unstable, and Stable Directions (R).** If a shadow exists, then by definition it cannot be too far away from the original noisy orbit. It is reasonable to assume that sometimes the resolvents, unstable and stable directions (abbreviated RUS) will change little as the orbit is perturbed towards the shadow. Thus, it may not be necessary to recompute the RUS for every refinement iteration. This is probably the biggest savings, because in combination with the cheaper accurate integrator, it means that the RUS often needs only to be computed once to a lax tolerance. This is analogous to a Newton’s method that does not recompute the Jacobian at every iteration.

**Re-use RUS From a Previous Successful Shadow (P).** To find the longest possible shadow of a noisy orbit, we attempt to shadow for longer and longer times until shadowing fails. Assume shadowing for  $S$  shadow steps produces a shadow  $A$ . When attempting to shadow  $S + k$  timesteps for some integer  $k$ , assume a successful shadow  $B$  will be found. Since  $A$  and the first  $S$  steps of  $B$  both shadow the same noisy orbit, they must be close to one another. By the same argument as the previous paragraph, the RUS that was computed for  $A$  can probably be re-used for the first  $S$  steps of  $B$ . Thus only  $k$  new RUS $_i$ ’s need to be computed.

**Recompute Specific RUS $_i$ ’s (i).** If it can be shown that a particular RUS $_i$ , or a small set of them, is causing Constant RUS refinements to fail, then perhaps it is sufficient only to recompute the “bad” RUS $_i$ ’s,



rather than all of them along the entire orbit. It is not trivial to decide which ones need recomputing, though.

**Large shadow steps ( $L$ ).** QT used every internal timestep of their integrator’s noisy orbit as a shadow step, but this is not necessary. It is reasonable to skip timesteps in the original orbit to build larger shadow steps. This means that fewer resolvents and stable/unstable directions need to be computed and stored. This optimization alone produced speedups of more than an order of magnitude in some cases.

## 3.2 Constant resolvent, unstable and stable directions

For a smooth, continuous, well-behaved function, the Jacobian is also smooth and continuous. Most Hamiltonian systems (that we have seen) are at least piecewise well-behaved in this sense. The  $N$ -body problem, in particular, is of this class, assuming no 0-distance collisions occur, or if the forces are artificially bounded through “softening”. This implies that the resolvent along a particular path will not change much if the path is perturbed slightly. Since the unstable and stable unit vectors are computed solely from the resolvents, they will also change only slightly when the path is perturbed. Thus, the RUS computed on one refinement iteration may be re-usable on some subsequent iterations. Computing the new 1-step errors is then the only significant work to be performed on each refinement iteration after the first.

We have observed that computing the RUS to a tolerance of between  $10^{-6}$  and  $10^{-9}$  usually suffices to refine the orbit down to 1-step errors near  $10^{-15}$ . If refinement fails using Constant RUS, then perhaps there is a spot in the orbit where some  $RUS_i$ ’s change quickly with small perturbations in the orbit. In that case, the RUS can be re-computed for a refinement or two, after which Constant RUS may be tried again. However, when Constant RUS fails it often suggests that there is a glitch somewhere in the portion of the orbit which is currently being tested for a shadow.

### 3.2.1 Switching between Constant and Recomputed RUS

The algorithm used to switch between Constant RUS and Recomputed RUS uses a running-average “memory” scheme to keep track of the progress of the norm of the maximum 1-step error over the previous few refinements. Originally, we used a simple algorithm (like QT’s) that would signal failure when a certain number  $K$  of successive refinements had failed, for  $\mu = 0.9$  (see (1)) and  $K = 3$ . However, refinement would sometimes get into a loop in which there would be a few successful refinements, followed by one or more that would “undo” the improvements of the previous ones. For example, the largest 1-step error may cycle as (approximately)  $\{8, 4, 2, 1, 8, 4, 2, 1, 8, \dots\} \times 10^{-13}$ . Cases were seen in which the number of refinements in these loops was 3, 4, and even up to 8. Clearly a simple “die when  $K$  successive failures are seen” is not general enough to catch this kind of loop.

The *arithmetic running average*

$$A_{j+1} = \alpha A_j + (1 - \alpha)\nu, \quad \alpha \in (0, 1) \tag{12}$$

is useful here, where  $\nu$  is the newest element added to the set, and  $\alpha$  is the *memory* constant. The higher the memory constant, the longer the memory — *i.e.*, the less the effect of each new element. A rule of thumb is that  $A$  is roughly an average of the most recent  $2/(1 - \alpha)$  elements. However, (12) is not suited for measuring errors such as those in refinement that can change by many orders of magnitude, and is especially unsuited when smaller means better. For example, a string of errors of  $10^{-4}$  followed by an error of  $10^{-7}$  would average to about  $10^{-4}$ , whereas a change from  $10^{-4}$  to  $10^{-7}$  is an indication that refinement is succeeding. A *geometric* equivalent of (12) is more appropriate, to allow values differing by orders of magnitude to average

meaningfully. We define the *geometric running average*  $G$  as

$$G_{j+1} = \sqrt[n+1]{(G_j)^{n\nu}} \quad (13)$$

where  $\nu$  is the new element, and  $\frac{n}{n+1}$  is analogous to  $\alpha$  in (12), so higher  $N$  implies longer memory. Both (12) and (13) require initialization to some reasonable starting value  $A_0$  and  $G_0$ , respectively.  $N$  should not be too large; we found 2 or 3 worked best.

Finally, we want to measure the *improvement* that is being made by successive refinements, rather than the absolute error, because it is the *change* in 1-step errors that indicates whether current refinements are succeeding, not their absolute size. Note that the 1-step errors may stop decreasing for two reasons: (a) 1-step errors have reached the machine precision, or (b) refinement is failing to find a shadow. Using the improvement rather than the absolute value allows us to use the same algorithm to halt refinement in either case. We use the geometric running average improvement to measure the progress of refinement, where the improvement  $I$  is

$$I = \left| \frac{\text{the maximum 1-step error of this refinement}}{\text{maximum 1-step error of previous refinement}} \right|.$$

If  $I$  is close to or greater than 1, then the current refinement did not improve on the previous one; if  $I < \mu$  ( $\mu$  from (1)) then the current refinement is successful.

Here is how we use the improvement  $I$  to control whether the RUS is recomputed at each refinement iteration. We have two refinement methods. In order of decreasing cost, they are Recomputed RUS, and Constant RUS. The general idea is: if refinement is working well with an expensive method, then switch to a cheaper one; if a cheaper one seems to be failing, switch to a more expensive one; if the most expensive one is failing, then give up refinement entirely. Here is the heuristic we have built over many months of trial and error. Let  $\mu_R$  be the success value (from (1)) when the RUS is being recomputed, and  $\mu_C$  the success value when using Constant RUS. We used  $\mu_R = 0.1$ ,  $\mu_C = 0.5$ , because we expect a Constant RUS iteration not to be as successful as one that recomputes the RUS, but they are much cheaper so we can afford a lower expected improvement per Constant RUS refinement.

- When using recomputed RUS, it is safe to switch to Constant RUS when the geometric running average improvement becomes  $< \mu_R$ . At this time the current orbit and all its statistics must be saved in case Constant RUS refinement fails.
- When using Constant RUS, it is necessary to switch back to recomputed RUS when the geometric running average improvement stays  $> \mu_C$  for 3 successive refinements. We then discard all progress made by Constant RUS and revert to the previous orbit that used recomputed RUS. We believe it is advantageous to discard all progress made by Constant RUS, even if significant progress was made, for the following reasons:
  - A refinement that computes the RUS is far more expensive than one that does not, so discarding all progress made by Constant RUS refinements usually makes little difference, percentage-wise, in the final run-time.
  - Future Constant RUS refinements, that will be performed after the RUS is recomputed, will converge faster than the current ones that just failed, because the RUS will be more accurate.
  - Recomputing the RUS with small 1-step errors would mean, using the simple “3 extra digits of accuracy” heuristic, recomputing it at a much higher accuracy than is necessary. (Unless the *Fixed Tolerance Resolvent* optimization is used.)
- Finally, after having switched back to recomputed RUS, it is time to admit failure if the geometric running average improvement stays  $> \mu_R$  for 3 successive refinements; and it is safe to switch again to Constant RUS when the geometric running average improvement becomes less than  $\mu_R$ .

We found that there are usually no half-measures when using recomputed RUS — refinement either succeeds geometrically or fails miserably. Finally, there may be a place for other noise-reduction procedures in the universe of expensive and cheap refinement algorithms.

### 3.2.2 Re-use RUS from a previous successful shadow

This idea is based on the same observation as Constant RUS, and is only useful if Constant RUS is employed. It takes advantage of the RUS for an orbit  $B$  being near to the RUS for a nearby orbit  $A$ . So, when trying to find a shadow for a segment  $B$  which is an extension to segment  $A$ , the RUS of  $A$  can be re-used on the segment of  $B$  that overlaps  $A$ .

One interesting question is which of  $A$ 's RUS's to use, if  $A$ 's RUS was ever recomputed. An argument in favour of using the first is that the noisy orbit is exactly the same, since  $B$  is just an extension of  $A$ . However, if  $A$  had to recompute the RUS, then probably so will  $B$  if  $A$ 's first RUS is used. Recall that the RUS needed for the early refinements need not be as accurate as the ones used later. Furthermore, we assume that the segment of  $B$ 's shadow that overlaps  $A$ 's shadow will be *closer* to  $A$ 's shadow than the first few iterations of  $B$  is to the first few iterations of  $A$ . Thus, if we use the last RUS that  $A$  used, we are less likely to recompute the RUS for  $B$ . If  $B$  is a *subset* of  $A$  rather than an extension, then obviously no new RUS $_i$ 's need to be computed.

### 3.2.3 Recompute specific RUS $_i$ 's

When Constant RUS refinements fail, it may be because some of the RUS $_i$ 's are invalid due to the current trajectory drifting too far from the trajectory for which the RUS $_i$  was computed. If the set of RUS $_i$ 's that are invalid is small, then it would be advantageous to recompute only those RUS $_i$ 's that are causing the failure.

A difficulty arises, however, when trying to pinpoint which RUS $_i$ 's are causing the problem. Note that refinement using the GHYS/QT algorithm is a *global* method: the correction  $\mathbf{c}_j$  is a function of the resolvents and 1-step errors of *all* the timesteps, not just  $L_j$  and  $\mathbf{e}_j$ . Thus, it is not trivial to decide which RUS $_i$ 's are causing a particular  $\mathbf{c}_j$  to be invalid.

What we would like is some measure  $M_i$  of how bad a RUS $_i$  is. Then we can insert RUS $_i$ 's into a priority queue based on  $M_i$ , and extract the RUS $_i$ 's that are worse than some criteria, and recompute them. Some obvious measures for  $M_i$  are:

- the 1-step error at step  $i$ .
- the distance  $\mathbf{p}_i$  has moved since the last time its RUS $_i$  was computed.
- the norm  $\|L_i\|$ , which is a measure of how sensitive that area of the orbit is to small changes.

Once the  $M_i$ 's are ordered by some measure, a criterion is needed to decide where to put the dividing line between those RUS $_i$ 's that need recomputing and those that don't. Some obvious criteria are:

- recompute all RUS $_i$ 's for which  $M_i > M'$  for some constant  $M'$ .
- recompute some constant number of the RUS $_i$ 's.
- recompute some constant fraction of the RUS $_i$ 's.
- recompute all whose  $M_i > w(\bar{M})$  for some function  $w$  of the geometric mean measure  $\bar{M}$ .

None of these measures taken alone with any of the above criteria taken alone seem sufficient, since there were cases in which recomputing the entire RUS resulted in a successful refinement, but a single-measure-single-criterion method for recomputing particular  $\text{RUS}_i$ 's resulted in a long sequence of unsuccessful refinements.

After much trial and error, a measure was found which gives a speedup no better than a factor of 2, which is much less than we expected. Let us focus on a particular timestep  $i$ . Let  $k$  be the current refinement iteration, and  $k'$  be the refinement at which the  $\text{RUS}_i$  was last computed. Let  $L_i^{k'}$  be the resolvent for step  $i$  computed during refinement  $k'$ . Let  $\mathbf{p}_i^{k'}$  be the orbit point at step  $i$  from refinement  $k'$ , and let  $\mathbf{p}_i^k$  be the orbit point at step  $i$  for the current refinement  $k$ . When trying to decide if we need to compute  $L_i^k$ , the measure we devised is

$$M_i = \|L_i^{k'}(\mathbf{p}_i^{k'} - \mathbf{p}_i^k)\|.$$

The intent is to combine the effects of the norm of  $L_i$ , and how much the orbit has moved at that point since  $L_i$  was last computed.  $L_i^k$  is recomputed if it satisfies *any* of the following criteria:

- always recompute the worst  $K$   $\text{RUS}_i$ 's, for some integer  $K$ ;
- always recompute the worst  $F$   $\text{RUS}_i$ 's, for some fraction  $F$  of all of them;
- always recompute  $\text{RUS}_i$  if  $M_i > M'$  for some constant  $M'$ .
- for each  $\text{RUS}_i$  satisfying one of the above, also recompute  $\text{RUS}_{i-1}$ , since  $L_i$  starts at  $\mathbf{p}_i$ , while  $L_{i-1}$  ends at  $\mathbf{p}_i$ .

We used  $K = 1$ ,  $F = \frac{1}{8}$ , and  $M' = 10^{-3}$ .

Unfortunately, this set of measures and criteria eventually recomputes most of the  $\text{RUS}_i$ 's except in the final few refinements of a failing search. In other words, this set of criteria gives the best speedups only during shadow searches that will eventually fail to find a shadow. This is a useful speedup, because a large fraction of the time spent trying to find the longest shadow is spent in the failing search for a shadow longer than the longest eventually found. However, it would be better if some criteria could be found that also aids in searches that are eventually successful. It seems plausible that some such measure and criterion should exist. Since the benefits are potentially enormous for orbits with many timesteps, more study seems appropriate. There is also the possibility that no significantly better measure exists.

### 3.3 Large Shadow Steps

A numerical solution of an ODE is represented by an ordered, discrete sequence of points. Assume we construct a continuous solution  $\mathbf{p}(t)$  from these points, for example by using a suitably smooth and accurate interpolant. Then we could extend the definition of  $\varepsilon$ -shadowing to continuous systems as follows: if  $\mathbf{x}(t)$  is an exact solution then  $\mathbf{x}(t)$   $\varepsilon$ -shadows  $\mathbf{p}(t)$  on  $t_0 \leq t \leq t_1$  if  $\forall t \in [t_0, t_1]$ ,  $\|\mathbf{x}(t) - \mathbf{p}(t)\| < \varepsilon$ .

Now, it should be clear that we can choose *any* representative sequence of points along  $\mathbf{p}(t)$  to use in the refinement algorithm; we need not choose the points from which  $\mathbf{p}(t)$  was originally constructed. In particular, we can choose a subset of the original set of points, as long as enough points are chosen to be “representative” of  $\mathbf{p}(t)$ . The steps that are finally chosen are called *shadow steps*.

There are at least two reasons to desire as few shadow steps as possible. First, if Constant RUS is being used, then the RUS needs to be stored. Each  $\text{RUS}_i$  requires a matrix (the resolvent), and two sets of basis vectors.<sup>4</sup> Second, if the “accurate” integrator has a large startup cost, then we wish to minimize the number

---

<sup>4</sup>For example, with  $M = 25$  moving particles, a RUS of length  $S = 128$  requires about 100 megabytes of memory, and this space scales as  $O(SM^2)$ .  $((25 \text{ particles} \times 6 \text{ dimensions per particle})^2 \times 8 \text{ bytes per double precision number} \times 4 \text{ (2 resolvents, 2 sets of basis vectors each covering half the space)} \times 128 \text{ steps})$ .

of startups. For example, the Adams’s method we used takes extremely small internal timesteps at the beginning of each shadow step.

For the  $N$ -body simulations reported in this paper, a set of “standardized” units (Heggie and Mathieu 1986) was used such that the scale of the system in all macroscopic units of interest was of order 1 — *i.e.*, the system had a diameter of order 1, the crossing time and total energy were of order 1, and the  $N$  particles each had mass  $1/N$ . In such a system, the timesteps of QT’s integrator averaged about 0.01, and they used each timestep as a shadow step. We found empirically that, in this system, using shadow steps about ten times longer, of size 0.1, works well. Smaller shadow steps use more time and memory but could not find shadows any better; shadow steps of 0.2 were slightly less reliable, and steps of size 0.5 were unreliable. A significant area of further work would be to dynamically determine in an algorithmic fashion what size shadow step is appropriate in the general case.

It might be reasonable to assume that as refinement proceeds to decrease the 1-step errors, it would be advantageous to construct larger-and-larger shadow steps. However, we found that this does not give much of a speedup, because (1) if refinement is working well, then Constant RUS will be invoked, in which case there is no need to recompute the RUS at all; (2) conversely, if refinement is failing, then there is no justification to increase the size of shadow steps — in fact, it may be advantageous to decrease them. Note that, if refinement *is* succeeding, it may be advantageous to use larger-and-larger shadow steps in the computation of the 1-step errors, even though the RUS is not being recomputed. We did not do this.

Finally, implementing dynamic shadow step sizes may help in another area. As described above, we currently discard all progress made by Constant RUS when it begins to fail, reverting to the orbit and RUS of the most recent refinement that computed the RUS. However, if Constant RUS works for several refinements, but then starts to fail, perhaps it would be advantageous to use the orbit of the most recent successful refinement to build new shadow steps using the above dynamic algorithm, regardless of whether the most recent successful refinement computed the RUS or not.

### 3.4 Cheaper accurate integrator

When computing the 1-step errors of a noisy orbit, an integrator must be used that has smaller 1-step errors than the integrator used to compute the noisy orbit. The question is, how much more accurate does it need to be? During early refinements when the 1-step errors are large, the errors in the correction factors may be dominated by the first-order approximation of the resolvents, so the 1-step errors do not need to be computed to machine precision. In practice, we found that it is sufficient to estimate the 1-step errors of the noisy trajectory to an accuracy of  $10^{-3}$  times that of the size of the maximum 1-step error of the previous refinement; computing them any more accurately does not speed refinement, although computing them only  $10^{-2}$  times more accurately sometimes results in more iterations of the refinement procedure.<sup>5</sup> In addition, a factor of only  $10^{-2}$  is not enough because some integrators have sufficiently gross control of their errors that requesting only 2 extra digits will sometimes result in the integrator returning *exactly* the same result. This is because it did too much work in the less accurate case, so that the solution it computed had at least 2 more accurate digits than were requested. This is less likely to happen with a difference of  $10^{-3}$ . Note it is necessary to loosen this criterion when the 1-step errors are within  $10^{-3}$  of the machine precision, in order not to request more accuracy than the machine precision allows.

The only methods we have tried as our accurate integrator are Adams methods. It may be wise to try others. In particular, if the shadow steps are small, it may pay to use a less sophisticated routine, such as a high-order explicit Runge-Kutta method. If the shadow steps are large, an Adams or a Bulirsch-Stoer method is probably apt, because even though they both have high startup cost, they can eventually move quickly along the solution if they are not restarted too often. It may also be interesting to attempt using

---

<sup>5</sup>More iterations, but each iteration takes less time. There is clearly a trade-off here that is probably problem dependent.

extended precision integration to test the accuracy of the standard (double precision) routines.

Finally, an important factor in problems, like the  $N$ -body problem, that have a large variance of scales all occurring at once (*i.e.*, some pairs of stars are several orders of magnitude closer to each other than other pairs) is a concept called *individual time steps*. The modern hand-coded  $N$ -body integrator gives each star a personalized integration timestep; stars that have high accelerations are integrated with smaller timesteps than those with lower accelerations. Perhaps it would be fruitful to attempt shadowing with an integration routine that uses individual timesteps for each star, because currently the timestep of the entire system is restricted to the smallest required timestep.

## 4 Numerical results of the optimizations

To quantify the speedups of the optimizations, 32 noisy orbits were chosen randomly from a simplified gravitational  $N$ -body system described in detail elsewhere (Quinlan and Tremaine 1992; Hayes 2003b; Hayes 2003a). Each was shadowed using several different combinations of the optimizations. In each case, the system consisted of 99 fixed particles and one moving particle (*i.e.*,  $N = 100, M = 1$ ), identical to the shadowing experiments of QT. Forces were not softened. The 32 orbits were chosen by generating random 3-dimensional positions for all particles from the uniform distribution on  $(0, 1)$ ; a 3-dimensional random initial velocity for the moving particle was also chosen uniformly on  $(0, 1)$ .<sup>6</sup> The standardized units of Heggie and Mathieu (Heggie and Mathieu 1986) were used, in which each particle has mass  $1/N$ . The pseudo-random number generator was the popular *Unix* 48-bit *drand48()*, with seeds 1 through 32.

Once the initial conditions were set, each noisy orbit was generated by integrating for 1.28 standard time units (about 1 crossing time) using LSODE (Hindmarsh 1980) with pure relative error control of  $10^{-6}$ . This figure agreed well with the magnitude of the initial 1-step errors computed during the first refinement. Although one crossing time sounds short, it is long enough that 5 of the orbits contained trouble spots.<sup>7</sup> The number of unoptimized QT refinements required to find a shadow that has no trouble spots seems independent of the length of the orbit — each refinement takes much longer, but the convergence is still geometric per refinement. Thus, the results below should be independent of the length of the orbit. This is what has been observed with the longer orbits we have shadowed, although we do not document them here.

For short shadow steps, a constant shadow step size of 0.01 was used, which approximates the average sized shadow step in QT. This results in 128 shadow steps. For large shadow steps, 16 shadow steps of size 0.08 were used. As in a usual shadowing attempt, longer and longer segments are shadowed in an attempt to find the longest shadow. Here, each successive segment had twice as many shadow steps as the previous one, up to 16 and 128, for long and short shadow steps, respectively. No attempt was made to isolate glitches more accurately than this factor of two.

The highest tolerance requested of the accurate integrator was  $10^{-15}$ . A successful numerical shadow was defined to be one whose 1-step errors were all less than  $2 \times 10^{-14}$ . This number was chosen simply because it was the smallest number that geometrically converging refinement sequences could consistently achieve;  $10^{-14}$  was too small, because many refinement sequences would proceed geometrically until their maximum 1-step error was about  $1.5 \times 10^{-14}$ , and then they would “bounce around” for many refinements until, apparently by chance, the maximum 1-step error would fall below  $10^{-14}$ . The maximum allowed shadow distance was 0.1, although none over 0.0066 were observed with successful shadows.

---

<sup>6</sup>Astronomers will note that this does not correspond to any realistic astronomical system; however, it seems unlikely that the precise particle distribution will effect shadowing results. This is supported by the close correspondence of our results with those of QT, even though they used a more realistic “Plummer” distribution.

<sup>7</sup>If the orbit has a trouble spot, then *Constant Resolvent* and *Re-use RUS from previous successful shadow* will have little effect, because the RUS will be re-computed in an attempt to find a shadow. The other optimizations — *Cheaper Accurate Integrator*, *Large Shadow Steps*, *Fixed Tolerance Resolvent*, and *Backwards Resolvent by Inverting Forward Resolvent* — still offer significant performance improvement.

Table I displays the speedups of the various optimizations. All ratios are speedup factors with respect to the time column. There are several interesting observations to make about this table. First, note that the run times of the original QT algorithm (after correction of the  $3M$  oversight) are comparable to the run times of our unoptimized version (differing on average by a factor of 1.07), as would be expected. Now looking at each optimization acting alone, we see that using Large Shadow Steps (column  $L$ ) 8 times longer gives an average speedup of about 5.5. Large shadow steps are trivial to implement, so we would recommend that, even if no other optimizations are adopted, large shadow steps be used in any new shadowing implementation. Third, inverting the forward resolvent to produce the inverse resolvent (column  $I$ ) produces the expected average speedup of 2.0. Fourth, the cheap accurate integrator (column  $C$ ) gives an average speedup of 2.36. This is about what is expected, because both the QT and  $C$  algorithms on average require just over 3 refinements to converge, and all  $C$  refinements are cheaper than a QT refinement except the last one, which is about equal in expense. Finally, using Constant RUS gives a speedup of almost 3. Again this is about what is expected because QT requires about 3 refinements to converge while  $R$  has one RUS computation followed by several cheap Constant RUS refinements. The  $P$ ,  $F$  and  $i$  optimizations were not tested alone.

Next we look at combinations of optimizations. First, it is interesting to note that combining the cheap accurate integrator  $C$  with Constant RUS  $R$  results in an average speedup that is greater than the product of the two individual speedups ( $2.36 \times 2.68 = 6.32 < 7.55$ ). This is because, when using  $CR$ , only *one* RUS is computed, and it is computed cheaply. Using  $R$  alone requires computing the one RUS to high accuracy; using  $C$  alone requires computing at least one RUS (the final one) to high accuracy. Second, re-using the RUS of a previous successful shadow ( $P$ ) makes sense only when  $R$  is also used. It gives an average speedup of 43% over  $CR$ . The next three columns show other combinations. An interesting point to note is that, except for  $CR$ , the combinations give a speedup less than the product of appropriate previous columns. Perhaps this is at least partially owing to an effect similar to Ahmdal's Law: there still exist parts of the algorithm that have not been sped up, and as the remainder of the program that *is* being optimized speeds up, these unoptimized sections take a greater proportion of the time.

Finally, it can be seen that the optimizations in the last two columns,  $F$  and  $i$ , contribute little to the average. This is because most of the systems in this table were shadowable. In fact,  $F$  slightly slows down refinement in one case. However, orbits 5, 6, 13, 19, and 29 obtained speedups significantly less than the average in the  $KRI$  column, but gain was made in the final two columns of up to a factor of 5. These were precisely the orbits that were most difficult to shadow (as the time column shows), and in fact some of them were not shadowable for their entire lengths. For these orbits, significant time was spent in the failed search to find a shadow lasting 1.28 time units. This indicates that the  $F$  and  $i$  optimizations are most effective in speeding up failing searches.

## 5 Discussion and conclusions

QT argue that if the refinement algorithm fails, there is good reason to believe that no shadow exists. This seems reasonable for several 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. Secondly, QT's results are consistent with a conjecture by GHYS on the frequency of glitches. Finally, Sauer and Yorke (1991) find that 1-step errors close to the machine precision are often enough to rigorously imply the existence of an exact shadow (assuming sufficient hyperbolicity).

However, in a non-negligible number of cases, a shadow is not found for shadow steps  $0..S$ , but a shadow is found for the *superset*  $0..2S$ . In other words, the algorithm failed to find a shadow of length  $S$  even though one exists. This also occurs when all optimizations are disabled, so it is not merely that the optimizations reduce reliability. Thus, in the search for the longest shadow, if our algorithm finds a shadow of length  $S$  but none of length  $2S$ , we also try  $4S$  before giving up. This is expensive, but necessary to lessen the chance of failing to find a shadow that exists.

TABLE I — SPEEDUPS OF THE OPTIMIZATIONS

seed	time(s)	QT	L	I	C	R	CR	CPR	CLR	KR	KRI	KRIF	KRIFi
1	1442	1.06	4.54	2.01	1.96	3.00	8.87	14.2	30.8	37.2	47.4	53.53	57.37
2	9379	0.70	6.08	1.82	3.09	3.48	11.2	15.7	52.3	68.3	75.7	79.83	79.67
3	7847	0.72	6.30	1.88	2.36	2.56	9.29	13.0	44.0	59.7	65.5	87.61	86.78
4	7011	0.86	11.1	1.98	2.74	2.84	7.88	12.3	40.5	55.8	74.3	79.19	78.52
5	12204	0.93	12.2	2.13	2.42	1.08	1.66	1.54	5.31	4.96	8.0	22.14	37.55
6	9443	1.14	7.24	2.01	3.18	3.98	10.8	13.5	5.36	56.8	29.8	80.13	137.64
7	6277	1.13	5.51	1.85	2.63	2.71	8.50	11.6	46.9	63.7	82.9	91.60	90.58
8	6582	1.13	4.79	2.04	1.91	2.79	7.09	10.6	33.8	45.4	56.9	72.76	67.95
9	7154	1.14	4.33	1.92	1.96	2.70	7.74	11.5	30.5	39.6	51.9	54.12	53.78
10	5832	1.13	4.32	1.87	1.54	2.72	6.46	10.7	29.1	45.5	57.2	63.89	63.59
11	6659	1.13	5.10	1.93	2.34	2.65	7.40	10.7	37.1	50.7	62.9	67.15	62.76
12	8623	0.86	5.00	2.44	2.57	3.37	9.64	13.7	44.8	52.4	75.4	79.77	79.19
13	14566	0.45	4.34	1.00	4.52	0.78	0.93	0.93	5.67	5.15	7.6	23.63	35.86
14	6976	0.74	4.54	2.00	2.17	2.85	7.23	11.1	32.9	45.2	49.4	58.08	57.66
15	5664	0.72	4.79	2.02	1.89	2.71	6.83	10.2	33.5	45.4	62.9	67.03	66.43
16	6634	1.15	4.89	1.97	1.82	2.64	7.35	10.6	33.8	44.8	58.7	62.65	62.08
17	6854	0.94	4.96	2.06	2.16	2.88	8.97	12.7	47.2	60.8	83.2	89.45	88.57
18	4932	0.62	2.88	1.74	1.69	2.17	5.38	7.84	25.7	34.1	45.7	49.70	49.27
19	8470	1.06	5.16	2.98	2.96	1.32	4.37	4.02	31.1	27.5	53.5	64.21	151.73
20	4566	1.12	5.59	1.91	2.04	2.78	6.21	9.50	33.7	48.6	64.3	62.95	66.32
21	7409	1.07	6.01	2.00	2.92	2.86	9.59	13.4	46.6	61.3	82.2	84.44	86.69
22	6988	1.22	4.89	2.17	2.00	3.03	7.34	11.4	36.1	44.0	57.2	61.00	60.52
23	7699	0.75	4.66	1.93	1.91	3.20	8.27	10.4	35.2	46.3	51.1	62.82	62.32
24	8851	1.18	2.53	1.79	2.78	2.89	10.2	15.4	52.9	71.1	93.1	95.29	100.39
25	7287	1.13	3.58	1.96	2.67	2.53	7.58	11.5	34.1	47.4	59.3	62.81	62.34
26	10736	0.85	7.74	2.84	2.79	4.33	12.4	16.9	58.2	76.5	81.1	63.96	63.50
27	7001	1.13	5.07	2.01	2.24	2.79	8.35	12.1	38.9	54.5	70.1	76.20	75.58
28	7324	1.13	4.49	2.06	1.93	2.72	8.01	11.2	33.2	49.3	55.0	58.98	55.21
29	13779	0.80	4.01	1.89	2.01	0.89	1.16	1.15	4.03	3.99	6.5	17.00	29.02
30	7932	1.14	10.9	2.01	2.17	2.81	8.82	13.0	43.3	67.5	86.7	91.38	90.70
31	6989	1.20	5.57	2.26	2.37	2.81	8.44	12.3	47.4	57.2	77.6	84.06	83.29
32	6720	1.08	4.80	1.89	1.90	2.78	7.61	11.5	34.0	50.5	65.5	70.33	69.67
avg	7682	1.07	5.56	2.01	2.36	2.68	7.55	10.8	34.6	47.5	59.3	66.8	72.3

LEGEND

- 
- time(s): time in seconds for unoptimized version.
  - QT: original code of Quinlan & Tremaine (Quinlan and Tremaine 1992)
  - L: *Large shadow steps*
  - I: backward resolvent by *Inverting* forward resolvent
  - C: *Cheaper accurate integrator*
  - R: *Constant RUS*
  - P: re-use RUS from *Previous successful shadow* (appears only in combinations)
  - K: *CPL together.*
  - F: *Fixed Tolerance Resolvent.*
  - i: *Recompute specific RUS<sub>i</sub>'s.*
-



The GHYS/QT refinement algorithm is trivially parallelizable, since the computation of each  $RUS_i$  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.

There may be further optimizations possible. For example, in the  $N$ -body problem, the phase-space dimensions are split into sets of 6, one set for each particle in the system. Instead of computing the  $6M \times 6M$  resolvent in  $O((6M)^3)$  time, perhaps it may be possible to refine the trajectories of each particle separately, requiring  $M$  resolvents each of dimension  $6 \times 6$ , and  $O(6^3M)$  time to compute.

Source code is available from the first author.

## Acknowledgements

We thank Gerry Quinlan and Scott Tremaine for helpful discussions throughout this research, and for providing their original source code. The University of Toronto undergraduate Computing Disciplines Facility generously provided workstation clusters on which we consumed several CPU years during this research. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada, and the Information Technology Research Centre of Ontario.

## References

- Anosov, D. V. (1967). Geodesic Flows and Closed Riemannian Manifolds with Negative Curvature. *Proc. Steklov Inst. Math* 90, 1.
- Beyn, W.-J. (1987). On Invariant Closed Curves for One-Step Methods. *Numer. Math.* 51, 103–122.
- Bowen, R. (1975).  $\omega$ -Limit Sets for Axiom A Diffeomorphisms. *Journal of Differential Equations* 18, 333.
- 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 (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 (1994). Shadowing of lattice maps. *Contemporary Mathematics* 172, 97–113.
- Coomes, B. A., H. Koçak, and K. J. Palmer (1994). Shadowing orbits of ordinary differential equations. *Journal of Computational and Applied Mathematics* 52, 35–43.
- Corless, R. M. (1994). Error Backward. *Contemporary Mathematics* 172, 31–62.
- 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.
- Grebogi, C., S. M. Hammel, J. A. Yorke, and T. Sauer (1990). Shadowing of Physical Trajectories in Chaotic Dynamics: Containment and Refinement. *Physical Review Letters* 65(13), 1527–1530.
- 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.
- Hayes, W. (1995). 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. (2003a). Shadowing-based reliability decay in softened  $n$ -body simulations. *Astrophical Journal Letters* 587, L59–L62.
- Hayes, W. (2003b, 7 February). Shadowing high-dimensional Hamiltonian systems: the gravitational  $n$ -body problem. *Physical Review Letters* 90(5).
- Hayes, W. and K. Jackson (2004). A survey of shadowing methods for numerical solutions of ordinary differential equations. *Applied Numerical Mathematics*. Accepted for publication.
- 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.* 41:5, 1948–1973.
- Heggie, D. C. and R. D. Mathieu (1986). Standardized Units and Time Scales. pp. 233–235. Springer-Verlag.
- Hindmarsh, A. C. (1980). LSODE and LSODI, two new initial value ordinary differential equation solvers. *ACM-SIGNUM Newsletter* 15(4), 10–11.
- 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. and S. Larsson (1993). Shadows, Chaos, and Saddles. *Appld. Numer. Math.* 13, 181–190.
- Sauer, T. and J. A. Yorke (1991). Rigorous Verification of Trajectories for the Computer Simulation of Dynamical Systems. *Nonlinearity* 4, 961–979.
- Van Vleck, E. S. (1995). Numerical Shadowing Near Hyperbolic Trajectories. *SIAM Journal on Scientific Computing* 16(5), 1172–1189.