



HAL
open science

Asymptotic distribution of global errors in the numerical computations of dynamical systems

Giorgio Turchetti, Sandro Vaienti, Francesco Zanlungo

► **To cite this version:**

Giorgio Turchetti, Sandro Vaienti, Francesco Zanlungo. Asymptotic distribution of global errors in the numerical computations of dynamical systems. *Physica A: Statistical Mechanics and its Applications*, 2010, 389 (21), pp.4994-5006. 10.1016/j.physa.2010.06.060 . hal-00476240

HAL Id: hal-00476240

<https://hal.science/hal-00476240>

Submitted on 25 Apr 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Asymptotic distribution of global errors in the numerical computations of dynamical systems

G. Turchetti^{*}, S. Vaienti[†], F. Zanlungo[‡]

March 7, 2010

Abstract

We propose an analysis of the effects introduced by finite accuracy and round off arithmetic on numerical computations of discrete dynamical systems. Our method, that uses the statistical tool of the decay of fidelity, computes the error comparing directly the numerical orbit with the exact one (or, more precisely, with another numerical orbit computed with a much higher accuracy). Furthermore, as a model of the effects of round off arithmetic on the map, we also consider a random perturbation of the exact orbit with an additive noise, for which exact results can be obtained for some prototype maps. We investigate the decay laws of fidelity and their relationship with the error probability distribution for regular and chaotic maps, both for additive and numerical noise. In particular, for regular maps we find an exponential decay for additive noise, and a power law decay for numerical noise. For chaotic maps numerical noise is equivalent to additive noise, and our method is suitable to identify a threshold for the reliability of numerical results, i.e. a number of iterations below which global errors can be ignored. This threshold grows linearly with the number of bits used to represent real numbers.

1 Introduction

The reliability of numerical computations in dynamical systems is a relevant long standing question. For continuous systems the global error on the orbit is the result of the accumulation of two types of errors: the first one is the local error due to the discretization algorithms, the second one is due to the finite precision representation of reals in any computational device and to the round off arithmetic. In this paper we propose a statistical analysis of the global errors due to finite accuracy and round off in discrete dynamical systems. Even

^{*}Department of Physics, University of Bologna

[†]Centre de Physique Théorique, CNRS, Universités d'Aix-Marseille I, II, Université du Sud, Toulon-Var and FRUMAM

[‡]Department of Physics, University of Bologna and ATR Intelligent Robotics and Communication Laboratories, Kyoto, Japan

though our results on numerical maps are based on numerical computations, the comparison with randomly perturbed maps, for which analytical results are available, allows to understand the key features. For maps with hyperbolic attractors the existence of true orbits close to pseudo-orbits is of great theoretical importance but does not provide any statistical information on the discrepancies between a pseudo-orbit and a true orbit with the same initial condition. We show that for chaotic maps there is a sharp transition between two regions, one in which the error is negligible and one in which its variance is comparable with the size of the attractor, both for the numerical and the randomly perturbed map. For regular or quasi regular anisochronous maps the transition is smooth and the accuracy decay follows a power law for the numerical map and exponential law for the randomly perturbed map.

In a previous work [1], we compared the value of an observable computed along the orbit of a given map and the orbit of the same map with a random perturbation for the same initial condition. The difference is a random process and, for large iteration times, its probability density function (pdf) converges toward a widespread distribution, which depends on the choice of the observable. For chaotic maps, the relaxation is sharp, in the sense that there exists a threshold time below which the width of the pdf is very small approximating a δ function and above which it has already reached its asymptotic value. These results were obtained analytically and numerically by computing the inverse Fourier transform of a suitable correlation integral between the two observables, called *fidelity*. For chaotic systems the fidelity decays at least exponentially fast and this allows to recover the average of the observables with respect to the invariant and the stationary measures. For the Bernoulli maps the decay is super-exponential and there is numerical evidence for this type of decay for most chaotic maps numerically investigated. For regular maps such as the translations on the torus the decay is smoother and follows an exponential law. Numerical evidence was found that the same exponential decay persists for quasi integrable maps.

In the subsequent letter [2] we applied the previous theoretical framework to compare the exact orbit of a map with its numerical computation (on a machine). The latter could be considered as a pseudo-random orbit where the noise is given by the round off errors.

In the present paper we develop with more details and formal considerations the results announced in [2] and provide some new results.

The difference between the values of an observable computed along the exact and the numerical orbits shows again the existence of an abrupt transition time n^* for the corresponding pdf, which has an interpretation of practical interest. We could in fact argue that below the transition time the numerical orbits is *faithful* to the exact orbit. By the way, this approach requires the knowledge of the exact map which is not accessible. This difficulty is overcome by comparing the numerical map computed for a given accuracy (simple floating point precision) with the same map evaluated numerically with a higher accuracy (double or higher floating point precision), that we consider as the *reference* map. We show that the results of the single-higher precision comparison are equivalent, at least below a significant time scale, to the results we would obtain by comparing

a single precision and the exact map. We show that the threshold n^* grows like $-\ln \varepsilon$, where ε is the accuracy specified by the least significant bit used to represent a real number. As a consequence the fidelity threshold n^* grows linearly with the length of the string of bits used to represent real numbers. Beyond this threshold the error distribution spreads quickly over its accessible range. The behavior just described is typical of chaotic systems. For regular maps the situation is slightly different: for isochronous maps the fidelity error does not decay, whereas for anisochronous maps the decay follows a power law. For these systems it is not possible to identify a sharp transition from the delta function to the asymptotic one; the transition is a gradual process whose length depends on ε , i.e. gets longer with the number of bits used to represent real numbers. A similar decay occurs for the correlations where it is due to the local mixing properties, a property typical of anisochronous maps.

2 Additive noise

We begin by recalling the main results obtained applying the fidelity to random perturbations of dynamical systems. We consider a map T defined on a phase space X which is a subset of \mathbb{R}^d , endowed with an invariant *physical* measure μ defined by

$$\lim_{n \rightarrow \infty} \int_X \Psi(T^n(x)) dm(x) = \int_X \Psi(x) d\mu(x) \quad (1)$$

where m denotes the Lebesgue measure and Ψ a continuous observable. Let us then consider a sequence of independent and identically distributed random variables ξ_i with values in the probability space Ξ and with probability density distribution $\eta(\xi)$ such that $Tx + \varepsilon \xi$ still maps X into itself. The iteration of the map T is therefore replaced by a composition of maps chosen randomly close to it (note that T itself is recovered when $\varepsilon = 0$): $T_\varepsilon^n(x) = (T + \varepsilon \xi_n) \circ (T + \varepsilon \xi_{n-1}) \circ \dots \circ (T + \varepsilon \xi_1)(x)$ and the stationary measure μ_ε of the process is defined by [3]

$$\lim_{n \rightarrow \infty} \int_{X, \Xi} dm(x) \prod_i \eta(\xi_i) d\xi_i \Psi(T_\varepsilon^n(x)) = \int_X \Psi(x) d\mu_\varepsilon(x) \quad (2)$$

We want to study the statistical properties of the error at the n -th iteration, defined as $\Delta_\varepsilon^n(x) = f(T^n x) - f(T_\varepsilon^n x)$, where f is a smooth observable; in the following we take $f(x) = x$ for one dimensional maps and $f(\mathbf{x}) = x_i$, $i = 1, \dots, d$, $\mathbf{x} \equiv (x_1, \dots, x_d)$, for multidimensional maps. We are in particular interested in studying ρ_ε^n , the probability distribution function of $\Delta_\varepsilon^n(x)$. This probability distribution can be studied directly, through a Monte Carlo sampling over initial conditions x and random perturbations ξ , or indirectly, using the fidelity.

Fidelity is defined through the following integral:

$$F_\varepsilon^n = \int_{X, \Xi} dm(x) \prod_i \eta(\xi_i) d\xi_i \Psi(T^n(x)) \Phi(T_\varepsilon^n(x)) \quad (3)$$

For a large class of maps which mix exponentially fast, it can be shown that the fidelity converges to $\int_X \Psi(x)d\mu(x) \int_X \Phi(x)d\mu_\varepsilon(x)$, and the absolute value of the difference between the integral (3) and its limiting value in terms of the invariant and stationary measure will be called the *fidelity error* and denoted with δF_ε^n . Note that by the asymptotic characterization of the invariant and stationary measures, the fidelity error could be equivalently defined as:

$$\delta F_\varepsilon^n = F_\varepsilon^n - \int_X \Psi(T^n(x))dm(x) \int_{X,\Xi} \prod_i \eta(\xi_i) d\xi_i dm(x) \Phi(T_\varepsilon^n(x)) \quad (4)$$

which emphasizes the role of the Lebesgue measure in the numerical computations. This quantity is relevant because the inverse Fourier transform of the fidelity $F_\varepsilon^n(u)$, computed by choosing $\Psi(x) = e^{iuf(x)}$, $\Phi(x) = e^{-iuf(x)}$ is the probability error distribution $\rho_\varepsilon^n(x)$. If the fidelity error converges to zero, then the fidelity converges to the product of the Fourier transforms of the measures $\int_X e^{iux} d\mu(x) \int_X e^{-iux} d\mu_\varepsilon(x)$ and its inverse Fourier transform is just ρ_ε^∞ . The initial error distribution is the Dirac distribution $\rho^0(s) = \delta(s)$, whereas it can be shown that if the invariant and stationary measures are Lebesgue the asymptotic distribution is the triangular function $\rho_\infty(s) = (1 - |s|)\vartheta(1 - |s|)$. Since the decay time scales of the fidelity error depend very weakly on the choice of the observables Ψ and Φ , the study of δF_ε^n using $\Psi(x) = \Phi(x) = x$ is usually the most efficient tool to investigate the convergence to the asymptotic error distribution ρ_ε^∞ .

For a couple of systems, translations on the torus (a regular system) and the Bernoulli map (a chaotic system), analytical results are available also for the transient. Choosing the probability distribution $\eta(\xi) = \frac{1}{2}\chi_{[-1,1]}(\xi)$ we found that the fidelity for translations is given by [1]

$$F_\varepsilon^n = \sum_{k \in \mathbb{Z}} \Psi_k \Phi_{-k} S^n(k\varepsilon) \quad S(x) = \frac{\sin(2\pi x)}{2\pi x} \quad (5)$$

where Ψ_k and Φ_k are the Fourier components of functions Ψ and Φ , while for the Bernoulli map $Tx = qx \bmod 1$, with integer $q \geq 2$ we have

$$F_\varepsilon^n = \sum_{k \in \mathbb{Z}} \Psi_k \Phi_{-k} S_{n,q}(k\varepsilon) \quad S_{n,q}(x) = \prod_{j=0}^{n-1} S(q^j x) \quad (6)$$

In the first case the decay of fidelity is exponential, with time scale ε^{-2} , while in the second case we have a plateau of length $n_* \propto -\ln \varepsilon$, below which fidelity is almost constant, followed by an ε independent super-exponential decay. Below the threshold n_* the error probability distribution can be approximated by a δ function, and the perturbed system considered as equivalent to the unperturbed one. The asymptotic error distribution is the same for the two systems (since the physical and stationary measure coincide with Lebesgue), and results to be the triangular function (figures 1 and 2). The transition between the δ and the triangular functions is sharp and ε independent for the Bernoulli system, while




Figure 1: Comparison between analytical results and Monte Carlo integrals for the decay of fidelity. Left: translations, $\varepsilon = 0.1$ (analytical result: blue dashed line; Monte Carlo: circles). Right: $3x \bmod 1$. $\varepsilon = 10^{-4}$ in black and $\varepsilon = 10^{-8}$ in green (analytical result: dashed line and crosses; Monte Carlo: circles). We used $N = 10^7$ integration points in the Monte Carlo method. The test functions $\Psi = \Phi^*$ were defined in Fourier space as $\Psi_k = k^{-1}$, truncated at $k = 30$.




Figure 2: ρ_ε^n for $3x \bmod 1$, $\varepsilon = 10^{-8}$. Left: $n = 13$; center: $n = 15$, right: $n = 19$. Compare the transition times with figure 1, right.

it is gradual and ε dependent for the translations (see [1]).

Our numerical study of maps for which analytical results are not available (Hénon, Baker's, Intermittent, Logistic and Standard maps) [1] shows that the behavior of translations and of Bernoulli maps can be considered as a prototype of, respectively, regular and chaotic maps. In particular for all the studied chaotic maps is possible to identify a threshold $n_* \propto -\ln \varepsilon$ below which the perturbed system can be considered as faithful to the unperturbed one; a threshold followed by an ε independent, super-exponential decay. To this threshold corresponds a sharp transition of the error probability distribution from a δ function to the asymptotic distribution (that in general is not a triangular function but depends on the invariant and stationary measures).

3 Numerical noise

Since real numbers have to be represented as strings of bits on a computer, to each discrete map T corresponds a numerical map T_* . The action of the numerical map depends on the length of bit strings used to represent real numbers and on the details of round off algebra, which are hardware dependent [4], nevertheless some general results can be stated. In our computational device any real number $x \in \mathbb{R}$ is represented by a floating point number x^{*p} , that is a binary string with p significant bits

$$x^{*p} = f 2^e \quad f = \pm 0.f_1 f_2 \dots f_p = \sum_{k=1}^p f_k 2^{-k} \quad e = \pm \sum_{k=1}^q e_k 2^k$$

where f_k, e_k are 0 or 1. The suffix p will be neglected whenever there is no ambiguity. The result of any arithmetic operation such as $x^{*p} \oplus y^{*p}$ corresponding to $x + y$ is a floating point number $z^{*p} \in \mathbb{F}_p$ and sum \oplus implies a round off. The relative error r_p is defined by

$$x^{*p} = x(1 + r_p) \quad |r_p| \leq \epsilon = 2^{-p}$$

For instance in floating point simple precision we have $\epsilon = 2^{-25}$. We can write, using the notation introduced for additive noise, $T_* x = T_\epsilon x = Tx + \varepsilon \xi(x)$ where ξ is now determined in a deterministic way by the initial condition and ε is a constant whose magnitude is of the order of the last significant bit used to represent x . The iterated map T_*^n can be written as T_ϵ^n using the previously introduced notation, but in this case the single step errors ξ_i will be n different functions of the initial condition x . Notice that T_* is defined by the round off rules of the computer, and the notation T_*^n correspond to apply n times the map as defined by those round off rules; while

$$\varepsilon \xi_n = x_n^* - T(x_{n-1}^*) = T_*^n(x) - T(T_*^{n-1}(x)) \quad (7)$$

is introduced to compare the round off results with the previous results for additive noise. We introduce also the global error Δ_*^n which can be written as

$$\Delta_*^n = x_n^* - x_n = T_*^n(x) - T^n(x)$$

The error Δ_*^n is the cumulative result of the local error for n iterations of the map. It is interesting to compare these definitions with the local and global errors when the perturbation to the map $T(x)$ is given by the additive noise. In this case the map $T_*(x)$ is replaced by $T_\epsilon = T(x) + \epsilon \xi$ so that letting $x_*^n = T_\epsilon^n(x) = (T(x) + \epsilon \xi_n) \circ \dots \circ (T(x) + \epsilon \xi_1)$ we have

$$\epsilon \xi_n = T_\epsilon(x_*^{n-1}) - T(x_*^{n-1}) = T_\epsilon^n(x) - T(T_\epsilon^{n-1}(x))$$

In this case we can write an explicit expression for the global error, which at first order in ϵ reads

$$\Delta_\epsilon^n = \epsilon \sum_{k=1}^n DT^{n-k}(x_1) \xi_k + O(\epsilon^2) \quad (8)$$

where $DT(x)$ is the jacobian matrix of the map T , the ξ_k are random vectors and the chain rule applies to compute DT^n . For the translations on the torus $DT = 1$ and we can write (8) as

$$\Delta_\epsilon^n = \epsilon(n\bar{\xi} + w_n) \quad (9)$$

where $\bar{\xi} \equiv \langle \xi \rangle$ is the average of the stochastic process and w_n are the fluctuations whose absolute value is of order 1. As a consequence in this case we have $\lim_{n \rightarrow \infty} \Delta_\epsilon^n/n = \epsilon\bar{\xi}$. The random perturbations we consider have usually zero mean so that $\bar{\xi} = 0$; this may not be the case of round off errors whose average can be nonzero. In the case of the Bernoulli maps DT is a constant and if the process ξ has zero mean again $\bar{\xi} = 0$ unlike the case of round errors discussed below. We also outline that when dealing with round off errors, the summation in eq. 8 starts with $k = 0$ because also the initial error $x^0 - x_*^0$ is present. Here x_*^0 stands for the finite accuracy representation of the initial condition x^0 , which is in general different from x^0 . As we will see, this may have relevant consequences.

4 Generalizations of Fidelity

This observable is used to characterize statistically the spread between the orbit and its stochastic perturbation. As previously introduced (3) fidelity was defined by an integral over all the possible single step error realizations ξ_i . Nevertheless, from a numerical point of view, integrals were performed with a Monte Carlo method, i.e. choosing N representative random vectors (x, ξ_i) , a procedure that led to a relative error of order $N^{-1/2}$ (see figure 1).

We suppose that if the deterministic $\xi_i(x)$ functional dependence of single step errors on the initial condition is complex enough, the vector $(x, \xi_i(x))$ can be considered as equivalent to a random sequence, i.e. the Monte Carlo integral over initial conditions and noise can be substituted for round off by an integral over the only initial conditions.

Corresponding to this *ansatz*, whose validity we are going to verify, we can compute the fidelity error for a system perturbed with numerical noise as

$$F_*^n = \int_X \Psi(T^n x) \Phi(T_*^n x) dm(x) \quad (10)$$

$$\delta F_*^n = F_*^n - \left(\int_X \Psi(T^n x) dm(x) \right) \left(\int_X \Phi(T_*^n x) dm(x) \right) \quad (11)$$

This definition requires the knowledge of the exact map T , which is in general not available. This problem can be solved by comparing the round off map T_* , realized with a given precision, i.e. as a string of bits of a given length, with a map realized at a higher precision, T_{\dagger} , that we call the “reference” map. For example T_* could be a single precision (8 digits) map, and T_{\dagger} a double precision (16 digits) map. The numerically computed fidelity will thus be

$$F_*^n = \int_X \Psi(T_{\dagger}^n x) \Phi(T_*^n x) dm(x) \quad (12)$$

To check the relevance of these results we can compare them with those obtained substituting T_{\dagger} with T_{\ddagger} (for example a map realized using 24 or 32 significant digits). If the results do not depend on the precision of the reference map, we can assume that they are equivalent to those that could be obtained if we had access to the exact map. We have applied this procedure to the results shown in this paper. Typically, the results obtained using as reference map a double precision T_{\dagger} or a 24 (32) digit map T_{\ddagger} are equivalent below a given time scale. This time scale, that corresponds to the time scale under which T_{\dagger} can be considered equivalent to T_{\ddagger} , as can be checked through a direct comparison between T_{\dagger} and T_{\ddagger} , is considerably longer than the time scale at which the error probability distribution of T_* has reached its asymptotic form and thus the results of the single-double precision comparison can be considered equivalent to those that would be obtained by a comparison between a single precision and an exact map.

Another way to avoid the problems related to the inaccessibility of the exact map would be to rely, for invertible maps, on a different definition of fidelity as

$$\tilde{F}_*^n = \int_X \Psi(x) \Phi(T_*^n T_*^{-n} x) dm(x) \quad (13)$$

where T_*^{-1} is the numerical realization of the inverse map, which is in general different from the inverse of T_* . Note that in the deterministic setting and with invertible maps T_* preserving the measure m , eq. (13) is equivalent to (3) [5]. We generalize it by integrating on given different realizations respectively of T_*^{-n} and T_*^n , and by supposing that this integral converges to $\int \Psi dm \int \Phi dm$ with the same rate as the fidelity error (11). For the numerical realization of an invertible map as the standard map, the equivalence between the two definitions has been checked with good results, at least in the chaotic regime (figure 3).

We have performed our study using three different architectures: a processor




Figure 3: Decay of fidelity for the standard map with $K = 10$, $\varepsilon = 10^{-8}$, Monte Carlo integrals using $N = 10^6$; red and crosses: comparison between single and double precision, black and circles: single precision using inversion.

2GHz Intel Core Duo on a Mac OSX v10.5 (Darwin 9) operating system, a processor Intel Pentium 4 3.00 GHz on a Linux Fedora 4 OS, and a processor 1 500 MHZ IP35 on a IRIX 6.5 OS (when not specified, the results shown in figures are obtained on the Darwin architecture). The details of the round off process depend on the architecture, nevertheless our studies show that some general rules about the error distribution can be stated.

5 Regular maps

For regular maps the behavior is significantly different from additive noise, showing that the integral over the only initial conditions is not equivalent for these systems to an integral also on the noise. For translations on the torus, $Tx = x + \omega \pmod{1}$, which are the prototypes for integrable maps, we have found that is possible to write the global error $\Delta_*^n(x, \omega) = T^n x - T_*^n x$ as in (9)

$$\Delta_*^n(x, \omega) = \varepsilon(\bar{\xi}(\omega)\phi(n) + w_n(x)), \quad \phi(n) = n + \phi_0 + \phi_1 n^{-1} + \dots \quad (14)$$

Here ε is a constant that represents the last significant bit accuracy, $\bar{\xi}$ is a constant that depends in a non trivial (and machine dependent) way on ω , which as we will see is equivalent to an average of the single step errors ξ_i , while w_n is a bounded, periodic function which depends on the initial condition x and has zero average with respect to the initial conditions. Notice that in this case $\bar{\xi}$ is not zero in contrast with the case of random noise. We thus have

$$\lim_{n \rightarrow \infty} \frac{\Delta_*^n(x, \omega)}{n} = \varepsilon \bar{\xi}(\omega) \quad (15)$$




Figure 4: $\Delta_*^n(x, \omega)$ in translations on the Torus for two different, randomly chosen, initial conditions x (red and black), compared with an average over 10^4 values of x (green). $\omega = \sqrt{2}/10$. The two figures show the same results on different time scales.

Figure 4 shows the global error $\Delta_*^n(x, \omega)$ for two different initial conditions x and the same ω , on two different time scales, showing that the contributions of w_n goes to zero in the limit. The periodic nature of the single step errors ξ_i (and thus of w_n) is shown by figure 5 (periodicity emerges after some non periodic iterations). The “asymptotic average error” $\bar{\xi}(\omega)$ is clearly correlated to the round off error $\delta\omega = \omega - \omega_*$ of the translation parameter as shown by figure 6. This figure shows at left the distribution of $\bar{\xi}(\delta\omega)$ and at right the $\bar{\xi}$ and $\delta\omega$ distributions obtained using a Monte Carlo sampling of the parameter ω . From the figure at left we can see that $\bar{\xi}(\delta\omega)$ is given by a distribution law centred around $\bar{\xi} = \delta\omega$. The distribution law for the round off error on the translation parameter

$$\eta(\delta\omega) = 2^{p-1} \sum_{k=p}^{\infty} \chi_{[-2^{-k}, 2^{-k}]}(\delta\omega)$$

where $\chi_I(x)$ denotes the characteristic function of the interval I , is stepwise as a consequence of the exponent in the finite precision representations of reals, with accuracy $\epsilon = 2^{-p}$. The stepwise shape of $\eta(\delta\omega)$ is clearly reflected also in the distribution law for $\bar{\xi}$, as shown in figure 6. These results suggest that an uniform integration over ω implies an integration over $\bar{\xi}$. This integration over $\bar{\xi}$ should correspond more or less to a random or pseudo-random process, since is related to a truncation of the less significative digits. The probability distribution for $\bar{\xi}$ following an uniform integration for ω is not uniform as shown in figure 6, but in the following discussions we will often approximate it with a uniform one.

The presence of a non zero average contribution of noise (15) leads to a quite different situation with respect to that obtained with random white noise, in which this contribution is not present.

Even if these results do not show a qualitative dependence on the hardware architecture used to implement the algorithm (while a quantitative dependence




Figure 5: Single step errors ξ_i , for two different, randomly chosen, initial conditions x . Translations on the Torus, $\omega = \sqrt{2}/10$.




Figure 6: Left: $\bar{\xi}(\delta\omega)$ distribution (translation on the Torus). Right: comparison between the $\bar{\xi}$ (red) and $\delta\omega$ (black) distributions. Both figures are obtained with a Monte Carlo sampling over ω (10^5 values for the first figure, 10^7 for the second one)




Figure 7: Left: $\Delta_*^n(x, \omega)$ for two different, randomly chosen, x values (black and red), compared with an average over 10^4 values of x (green). $\omega = \sqrt{2}/5$. Map realized as a 2D rotation. Right: variance of w_n for $\omega = \sqrt{2}/10$, average over 10^4 different initial conditions, map realized as a 2D rotation. Black and dashed corresponds to the results obtained on the Darwin architecture, while red and continuous to those obtained on the IRIX architecture.

can be found see the right frame of figure 7), they can depend strongly on the algorithm. The results just shown, for which the period of w_n is of order 100 and its magnitude of order 10, where obtained by implementing the map as a translation on the torus $x' = x + \omega \pmod{1}$. From an analytical point of view, this is equivalent to a 2D rotation on a circle, but from a numerical point of view the latter leads to a sequence w_n which has initially, for a number of iterations of order 10^5 , properties similar to a random sequence (i.e. its variance grows approximately with a linear law), see figure 7. For the translation on the torus $\xi(\omega)$ is given by $\langle \Delta_*^n(x, \omega) \rangle_x / n$ with a nice accuracy even for low values of n .

Concerning the decay of fidelity, for regular maps we distinguish between isochronous maps, for which the frequency is fixed, and anisochronous maps, such as the skew map on the cylinder $x' = x + \omega(y)$, $y' = y$, for which the frequency depends on the initial conditions. In the former case fidelity does not decay, while in the latter it has a power law decay. For the isochronous maps the system is basically equivalent to a deterministic rotation with a frequency $\omega + \varepsilon \bar{\xi}$, and it can be easily shown that in this case fidelity does not decay, but oscillates with constant amplitude. Indeed, ignoring the contribution of the term w_n , we can write the action of n iterations of the numerical map as $x_*^n = x^0 + n(\omega + \varepsilon \bar{\xi})$. Substituting in the Fourier analysis of appendix B in [1] and keeping $\bar{\xi}$ fixed (i.e., not performing the integral in the noise), this leads to an oscillating phase term but not to a decay term (actually the presence of the oscillating term w_n can cause a decay of fidelity for the highest Fourier modes, those corresponding to distances of the order of the magnitude of w_n).

For the anisochronous map with $\omega(y) = y$ (to which any map with monotonic $\omega(y)$ can be reduced), a power law decay for $F_n - F_0$ is actually observed. Indeed the map in this case is $T(x, y)$ defined on the cylinder $\mathbb{T} \times [a, b]$ or the torus \mathbb{T}^2

according to

$$T_x = x + y \bmod 1 \quad T_y = y$$

and the fidelity with respect to perturbed map is computed by integrating over x, y . If T_* denotes the map with numerical noise, the map becomes $T_*(x, y)$ whose iterate can be written

$$x_*^n = x^0 + \epsilon \xi_x(x^0) + n(y^0 + \epsilon \xi_y(y^0)) \quad y_*^n = y^0 + \epsilon \xi_y(y^0)$$

In writing this equation we use the previous results for translations on the Torus, were the role of $\bar{\xi}$ is now played by ξ_y and that of ω by y^0 . We have already ignored the bounded term w_n but we have explicitly wrote the round off error of x^0 for symmetry with respect to y^0 . Let us consider Φ and Phi to be functions of the only variable x , and use the usual Fourier analysis approach. The integrals can be computed if we replace, following again the previous discussion, $\xi_x(x^0)$ and $\xi_y(y^0)$ with $\bar{\xi}_x$ and $\bar{\xi}_y$ assuming that they are random variables ranging in $[-1, 1]$ so that the integration over x_0, y_0 is replaced by an integration over x_0, y_0 and $\bar{\xi}_x, \bar{\xi}_y$. The fidelity F_n is given by

$$F_*^n = \sum_{k \in \mathbb{Z}} \Psi_k \Phi_{-k} \frac{\sin(2\pi n k \epsilon)}{2\pi n k \epsilon} \frac{\sin(2\pi k \epsilon)}{2\pi k \epsilon} \quad (16)$$

This rather strong assumption provides a result which is in good agreement with the numerical computations see figure 8. A map where at any step both x and y where affected by a random error

$$T_{\epsilon x} = x + y + \epsilon \xi_x \quad T_{\epsilon y} = y + \epsilon \xi_y$$

would give a completely different result. Indeed taking into account that the iterate of order n is

$$x^n = x^0 + n y^0 + \epsilon \left(\sum_{j=1}^n \xi_x^j + \sum_{j=1}^{n-1} (n-j) \xi_y^j \right) \quad y^n = y^0 + \epsilon \sum_{j=1}^n \xi_y^j$$

the following result for the fidelity is found

$$F_*^n = \sum_{k \in \mathbb{Z}} \Psi_k \Phi_{-k} \left(\frac{\sin(2\pi k \epsilon)}{2\pi k \epsilon} \right)^n \prod_{j=1}^n \frac{\sin(2\pi j k \epsilon)}{2\pi j k \epsilon} \quad (17)$$

Figure 8 shows that this result agrees with the fidelity computed numerically for a randomly perturbed map and differs drastically from the fidelity of the map with round off noise.

6 Chaotic maps

Fidelity decay can be easily studied for chaotic numerical maps (for example Bernoulli, Hènon, Logistic, Intermittent, Baker's map and the Standard map




Figure 8: Left: Decay of fidelity for the skew map. Black: additive noise, compared with its analytical prediction eq 17 (red); green round off noise compared with its analytical prediction eq 16 (blue). Right: comparison of round off noise and analytical prediction on a longer time scale. $N = 10^4$

with $K \gg 1$) and shows always a good qualitative agreement with the results obtained for additive noise, i.e. the presence of a threshold below which fidelity is constant and the error function is qualitatively a δ function (its support is many order of magnitude smaller than the size of the phase space, as it happened in figure 2). Beyond this threshold, that we can call n_* and that grows as $-\ln \varepsilon$, i.e. linearly in the number of bits used to represent real numbers (figure 9), the error distribution spreads quickly over its accessible range, as can be checked using also a Monte Carlo sampling of the error distribution. The results for chaotic maps do not show a significant dependence on the choice of the architecture.

Stating that the decay law of fidelity is qualitatively similar to the one observed for additive noise is not enough to prove that the sequence of single step errors ξ_i (see equation 7) can be considered equivalent to a random one, as we claimed in our *ansatz*. In order to better understand the nature of this sequence we are going to use two approaches. The first approach would be to study the distribution of ξ_i for different i values. In the case of additive noise we would have a continuous uniform distribution identical for each i , and thus the more different is the actual ξ_i distribution from a continuous uniform one, the less valid is our assumption. The second approach consists in studying in greater detail the decay law, comparing it with the one for additive noise, to see if some difference emerges. In order to do that, can be useful to study the decay of fidelity for *quenched* noise. By quenched noise we mean a perturbation that uses the same sequence ξ_i for each initial condition (the sequence is obtained randomly choosing the values ξ_i but without integrating on them). If our assumption is valid, the fidelity for round off noise should be more similar to that obtained using additive noise than to the one obtained using quenched noise. Let us start with the map $3x \bmod 1$. From a detailed observation of figure 9 is possible to see, in particular for single precision, that some difference between




Figure 9: Decay of fidelity for $3x \bmod 1$ represented in single precision (black, circles) and double precision (red, diamonds) compared to a reference map T_{\dagger} using 32 digits; the results are compared to the decay of fidelity for random (additive) noise with $\varepsilon = 2^{-25}$ (blue, crosses) and $\varepsilon = 2^{-53}$ (green, squares), the value of single and double precision last significant bits for reals on the torus. Monte Carlo integrals with $N = 10^7$

round off and additive noise is present. For this map, we have found that the global error can be described as

$$\Delta_*^n(x) = \sum_{i=0}^n 3^{n-i} \varepsilon \xi_i(x) \quad (18)$$

(this equation, as eq. (14), does not take in account boundary effects) where the initial condition round off ξ_0 has a step-wise continuum spectrum $\eta(\varepsilon\xi) = 2^{p-1} \sum_{k=p}^{\infty} \chi_{[-2^{-k}, 2^{-k}]}(\varepsilon\xi)$ distribution once sampled over the space of initial configurations (due to the exponent in the floating point representation, 2^{-p} is the value of the least significant bit, that we can also consider as equivalent to the constant ε), while the ξ_i with $i \geq 1$ have a discrete spectrum, which results to be almost completely reduced to zero for $i \geq 2$ (i.e. no relevant errors are made after the first iteration). (figure 10). Due to this effect, which is caused by the extremely simple algorithmic nature of the map, the sequence of single step errors cannot be used as a representative sequence for an integral over the noise.

It is than evident that only the initial round off has a continuous (even if not uniform) distribution, and can be in some way considered as equivalent to the additive noise. In the discussion of the skew map, we have seen that the initial round off determined the decay law of fidelity, and this is true also for the map $3x \bmod 1$. Let us assume for simplicity's sake, as we have previously done for




Figure 10: Single step error distribution for the map $3x \bmod 1$ in single precision. Left: ξ_0 ; center ξ_1 ; right ξ_2 . Monte Carlo sampling with $N = 10^7$

the skew map, that the integral on the initial condition x^0 is equivalent to an uniform integral on the initial round off error, and let us ignore the contribution of the rest of the single step error sequence, ξ_i with $i \leq 1$. . It is possible to show using Fourier analysis that the decay law for the fidelity of a $qx \bmod 1$ map perturbed at the only initial step, with an uniform integral over noise is

$$\delta F_\varepsilon^n = \sum_{k \neq 0} \Psi_k \Phi_{-k} \frac{\sin(2\pi \varepsilon k q^n)}{2\pi \varepsilon k q^n} \quad (19)$$

(this result is trivially obtained following the Fourier analysis procedure in appendix B of [1] and performing only the first integral over noise). The decay law (19) describes better the decay law for fidelity due to round off noise than to (6), showing that the main contribution to decay fidelity is due to the initial, continuous spectrum perturbation (figure 11). To analyse the contribution of ξ_i with $i \leq 1$ let us remove the initial round off (choosing rational initial conditions that can be represented exactly with the prescribed precision). In this case fidelity presents a threshold, but does not decay to zero (after the threshold the error distribution expands to the whole phase space but does not converge to an asymptotic distribution, figure 12). In this case, i.e. when only initial condition points that can be represented exactly using the numerical precision under examination and thus the initial round off is removed, the fidelity oscillates after the threshold without going to zero (figure 13). Is not surprising that this behaviour is quite similar to that obtained studying the quenched noise perturbation of this map. Indeed it can be shown, (generalizing a result in [6]), that the fidelity for quenched noise oscillates but does not decay to zero. We have found that this result is particularly difficult to verify using Monte Carlo numerical integration. As it is shown in figure 13, a the map $3x \bmod 1$ studied with numerical precision ε_1 and perturbed with numerical noise ε_2 , starts oscillating in correspondence with the fidelity threshold for ε_2 (in agreement with the analytical result) but then drops to zero in correspondence with the fidelity threshold for ε_1 (see figure 13, comparing with figure 9 to know the single and double precision thresholds). According to us this result supports our conclusion




Figure 11: Decay of fidelity for $3x \bmod 1$. Black, circles: analytical result (6), using $\varepsilon = 2^{-25}$. Red, squares: analytical result (19). Green, diamonds: round off noise (single precision compared to double precision), $N = 10^7$




Figure 12: Error distribution ρ^n obtained comparing a single precision $3x \bmod 1$ map with a double precision map, using initial conditions in single precision (no initial round off). Black, $n = 30$; red, $n = 31$; green $n = 32$; compared to the triangular function (blue) obtained using initial condition in double precision (initial round off). Monte Carlo samplings using $N = 10^7$ initial conditions.




Figure 13: Decay of fidelity for $3x \bmod 1$. Black, diamonds: quenched noise computed with 32 digits ($\varepsilon = 2^{-25}$). Red, squares: quenched noise computed with double precision ($\varepsilon = 2^{-25}$). Green, circles: decay of fidelity obtained comparing a single precision map with a double precision one (round off noise) using initial conditions that can be exactly represented in single precision, i.e. removing the initial round off. Monte Carlo integrals, $N = 10^6$

that the fidelity decay threshold is a good measure for the maximum number of iterations under which the numerical computations are completely reliable.

For the other, (slightly) more algorithmically complex maps, the sequence of error ξ_i has a continuous spectrum for each value of i , and converges quickly to an asymptotic distribution (figure 14). For all these systems the asymptotic error distribution results to be almost indistinguishable from that obtained using additive noise (figure 15), suggesting that for these maps the assumption of equivalence between round off and additive noise could be valid (for these maps the initial round off error plays no special role).

In order to better verify this assumption, we have studied the decay of fidelity for maps perturbed with quenched noise, i.e. without performing the integral over noise. Nevertheless, for all the tested chaotic maps (with the exception of the Baker's map) the decay laws for quenched and additive noise were not distinguishable, at least using the computational precision allowed by our Monte Carlo method. For the Baker's map, it is possible to see that the fidelity decay is more irregular for quenched noise (due to the absence of the smoothing integral). In this case the decay using round off noise is more similar to the additive one, suggesting the validity of our assumption for this map (figure 16).




Figure 14: Distribution of single step errors ξ_i for the Hénon map, x observable, $N = 10^7$. Black, $i = 0$. Red, $i = 1$. Green, $i = 20$. Blue $i = 50$ (green and blue are very difficult to distinguish)




Figure 15: Asymptotic error distribution for the Hénon map, x observable. Black, additive noise ($\varepsilon = 10^{-8}$), red comparison between single and double precision (the two functions are almost identical). Monte Carlo samplings using $N = 10^7$ initial conditions.




Figure 16: Decay of fidelity for the Baker's map. Black, circles: the sequence errors of order 2^{-25} is the same for each initial condition (quenched noise), showing an irregular decay law; green, squares: single precision round off noise, showing the same smooth decay of additive noise with $\varepsilon = 2^{-25}$ (red, crosses). Monte Carlo integrals with $N = 10^7$

7 Conclusions

We have used the results of a previous work on additive noise to study the effects of round off noise on discrete systems. For regular maps the behavior depends on the algorithmic realization and on its character: for isochronous maps the fidelity error does not decay, for anisochronous maps it has a power law decay, whereas in presence of an additive noise the decay is always exponential. For chaotic systems round off and noise are almost equivalent. We also showed that the fidelity is an efficient tool for the study of the time scales for the convergence to the asymptotic error distribution, since, at least for chaotic maps, it allows us to find a threshold value below which the numerical system can be considered as equivalent to the exact one. This threshold linearly grows as the number of bits used to represent real numbers.

Acknowledgments

G. Turchetti acknowledges financial support from the PRIN 2007 grant *Simulazione della produzione di elettroni dalla interazione laser-(pre)plasma nel contesto della ignizione veloce* coordinated by D. Batani

References

- [1] P. Marie, G. Turchetti, S. Vaienti, F. Zanlungo, *Chaos* **19**, 043118 (2009)
- [2] G. Turchetti, S. Vaienti and F. Zanlungo, accepted on Europhysics Letters
- [3] Y. Kifer, Ergodic Theory of Random Transformations, Birkhäuser, 1986
- [4] D. Knuth *The art of computer programming* Vol 2, Addison-Wesley, (1969)
- [5] G. Benenti, G. Casati, G. Velbe, *Phys. Rev. E*, **670**, 055202-1, (2003)
- [6] C. Liverani, Ph. Marie, S. Vaienti, Random Classical Fidelity, *Journal of Statistical Physics*, **128**,4, 1079, (2007)
- [7] C. Grebogi, S. Hammel, J. Yorke, *J. Complexity*, **3**, 136, (1987)
- [8] T. Sauer, *Physical Review E*, **65**, (2002), 036220
- [9] G. Turchetti in *CHAOS, Systèmes Dynamiques*, p.239, Hermann Éditeurs, (2007)