

Filling Gaps in Chaotic Time Series

Francesco Paparella

Dipartimento di Matematica “E. de Giorgi”

Università di Lecce

*Lecce, Italy**

Abstract

We propose a method for filling arbitrarily wide gaps in deterministic time series. Crucial to the method is the ability to apply Takens’ theorem in order to reconstruct the dynamics underlying the time series. We introduce a functional to evaluate the degree of compatibility of a filling sequence of data with the reconstructed dynamics. An algorithm for finding highly compatible filling sequences with a reasonable computational effort is then discussed.

PACS numbers: 05.45.-a, 05.45.Ac, 05.45.Tp

Keywords: time series analysis, chaos, filling gaps

*Electronic address: francesco.paparella@unile.it

I. INTRODUCTION

One problem faced by many practitioners in the applied sciences is the presence of gaps (i.e. sequences of missing data) in observed time series, which makes hard or impossible any analysis. The problem is routinely solved by interpolation if the gap width is very short, but it becomes a formidable one if the gap width is larger than some time scale characterizing the predictability of the time series.

If the physical system under study is described by a small set of coupled ordinary differential equations, then a theorem by Takens [1, 2] suggests that from a single time series it is possible to build-up a mathematical model whose dynamics is diffeomorph to that of the system under examination. In this paper we leverage the dynamic reconstruction theorem of Takens for filling an arbitrarily wide gap in a time series.

It is important to stress that the goal of the method is not that of recovering a good approximation to the lost data. Sensitive dependence on initial conditions, and imperfections of the reconstructed dynamics, make this goal a practical impossibility, except for some special cases, such as small gap width, or periodic dynamics. We rather aim at giving one or more surrogate data which can be considered *compatible* with the observed dynamics, in a sense which will be made rigorous in the following.

We shall assume that an observable quantity s is a function of the state of a continuous-time, low-dimensional dynamical system, whose time evolution is confined on a strange attractor (that is, we explicitly discard transient behavior). Both the explicit form of the equations governing the dynamical system and the function which links its state to the signal $s(t)$ may be unknown. We also assume that an instrument samples $s(t)$ at regular intervals of length Δt , yielding an ordered set of \bar{N} data

$$s_i = s((i - 1)\Delta t), \quad i = 1, \dots, \bar{N}. \quad (1)$$

If, for any cause, the instrument is unable to record the value of s for a number of times, there will be some invalid entries in the time series $\{s_i\}$, for some values of the index i .

From the time series $\{s_i\}$ we reconstruct the underlying dynamics with the technique of delay coordinates. That is, we shall invoke Takens' theorem [1, 2] and claim that the m -dimensional vectors

$$\mathbf{x}_i = (s_i, s_{i+\tau}, \dots, s_{i+(m-1)\tau})$$

lie on a curve in \mathbb{R}^m which is diffeomorph to the curve followed in its (unknown) phase space by the state of the dynamical system which originated the signal $s(t)$. Here τ is a positive integer, and i now runs only up to $N = \bar{N} - (m - 1)\tau$. Several pitfalls have to be taken into account in order to choose the most appropriate values for m and τ . Strong constraints also come from the length of the time series, compared to the characteristic time scales of the dynamical system, and from the amount of instrumental noise which affects the data. We shall not review these issues here, but address the reader to references [3–5].

We note that gaps (that is, invalid entries) in the time series $\{s_i\}$ do not prevent a successful reconstruction of a set $\mathcal{R} = \{\mathbf{x}_i\}$ of state vectors, unless the total width of the gaps is comparable with \bar{N} . We simply mark as “missing” any reconstructed vector \mathbf{x}_i whose components are not all valid entries. If the gap in the signal s spans more than $(m - 1) * \tau$ data points, then it will be mapped into a contiguous gap in the sequence of reconstructed vectors.

If the valid vectors of \mathcal{R} sample well enough the underlying strange attractor embedded in \mathbb{R}^m , one may hope to find, by means of a suitable interpolation technique, a vector field $\mathbf{F} : U \rightarrow \mathbb{R}^m$, such that within an open set U of \mathbb{R}^m containing all the vectors \mathbf{x}_i , the observed dynamics can be approximated by

$$\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x}). \quad (2)$$

This very idea is at the base of several forecasting schemes, where one takes the last observed vector \mathbf{x}_N as the initial condition for equation (2), and integrates it forward in time (see e.g., [7, 8]).

The gap-filling problem was framed in terms of forecasts by Serre *et. al.* [9]. Their method, which amounts to a special form of the shooting algorithm for boundary value problems, is limited by the predictability properties of the dynamics, and cannot fill gaps of arbitrary width.

The rest of this note is organized as follows: in section II we cast the problem as a variational one, where a functional measures how well a candidate filling trajectory agrees with the vector field defining the observed dynamics. Then an algorithm is proposed for finding a filling trajectory. In section III we give an example of what can be obtained with this method. Finally, we discuss the algorithm and offer some speculations on future works in sec. IV.

II. A VARIATIONAL APPROACH

The source of all difficulties of gap-filling comes from the following constraint: the interpolating curve, which shall be as close as possible to a solution of (2), must start at the last valid vector before the gap and reach the first valid vector after the gap in a time T which is prescribed.

To properly satisfy this constraint, we propose to frame the problem of filling gaps as a variational one. We are looking for a differentiable vector function $\xi : [0, T] \rightarrow U$ which minimizes the functional

$$J(\xi) = \int_0^T \left| \dot{\xi}(t) - \mathbf{F}(\xi(t)) \right|^2 dt, \quad (3)$$

with

$$\xi(0) = \mathbf{x}_p; \quad \xi(T) = \mathbf{x}_q.$$

Defining $l = q - p$, we have $T = l \Delta t$. If the curve $\xi(t)$ coincided with the missing curve $\mathbf{x}(t)$ for $t \in [0, T]$, and \mathbf{F} where a perfect reconstruction of the vector field governing the dynamics of the system, then the functional would reach its absolute minimum $J = 0$. The imperfect nature of \mathbf{F} suggests that any curve which makes J small enough can be considered, on the basis of the available information, a surrogate of the true (missing) curve. In section III we shall offer a simple criterion to quantify how small is “small enough”. For the moment we only care to remark that, even for a perfect reconstruction of the vector field, a curve ξ making J arbitrarily small, but not zero, need not to approximate $\mathbf{x}(t)$, in fact, the two curves may be quite different; however, such a curve ξ is consistent with the dynamics prescribed by (2).

Approaching a small value of J from an arbitrary initial curve using standard optimization methods may be computationally very expensive. In addition, discretized forms of (3) may have many relative minima far away from zero, and we expect that downhill algorithms will fall on one of these uninteresting minima for most choices of the initial guess. Thus, our problem really reduces to that of finding an initial guess suitable for easy refinement.

The complexity of the problem is greatly limited if we require that the set of points $\mathcal{L} = \{\mathbf{y}_j\}$ that sample the initial guess, has to be a subset of the set \mathcal{R} of reconstructed vectors. The index $j = 0, 1, \dots, l$ does not necessarily follow the temporal order defined in \mathcal{R} by the index i (cfr. eq. (1)), but we require that $\mathbf{y}_0 \equiv \mathbf{x}_p$ and $\mathbf{y}_l = \mathbf{x}_q$. We shall denote with $S(\mathbf{y}_j)$ the successor of the vector \mathbf{y}_j with respect to the temporal order in \mathcal{R} , and with

$P(\mathbf{y}_j)$ its predecessor. If $P(\mathbf{y}_{j+1}) \neq \mathbf{y}_j$, $S(\mathbf{y}_j) \neq \mathbf{y}_{j+1}$ we say that there is a jump between the position j and $j+1$. We observe that if \mathcal{L} contained all the missing vectors in the correct order ($\mathcal{L} = \{\mathbf{x}_p, \mathbf{x}_{p+1}, \dots, \mathbf{x}_q\}$) there would be no jumps in \mathcal{L} . No subset of \mathcal{R} can be found starting at \mathbf{x}_p , ending at \mathbf{x}_q , and containing no jumps.

To determine properties which characterize good choices for \mathcal{L} , we need to define a discretized form of the functional J that can be evaluated for subsets of \mathcal{R} , rather than for differentiable curves. To this end we may approximate $\dot{\xi} - \mathbf{F}(\xi)$ with $[(\mathbf{y}_{j+1} - \mathbf{y}_j) - (\mathbf{y}_{j+1} - P(\mathbf{y}_{j+1}))]/\Delta t$ or with $[(\mathbf{y}_{j+1} - \mathbf{y}_j) - (S(\mathbf{y}_j) - \mathbf{y}_j)]/\Delta t$. Averaging between the two expressions, and dropping the unimportant Δt , the discretized form of J is

$$J_0(\mathcal{L}) = |P(\mathbf{y}_1) - \mathbf{y}_0|^2 + \sum_{j=1}^{l-2} \frac{|P(\mathbf{y}_{j+1}) - \mathbf{y}_j|^2 + |\mathbf{y}_{j+1} - S(\mathbf{y}_j)|^2}{2} + |\mathbf{y}_l - S(\mathbf{y}_{l-1})|^2. \quad (4)$$

Of course, we shall restrict our choice of vectors to be included in \mathcal{L} only to valid vectors of \mathcal{R} having valid predecessor and successor.

The value of J_0 increases every time that there is a jump in \mathcal{L} . Only the sequence of missing vectors, if it were known, would give $J_0 = 0$. The magnitude of the jump (that is the distance between $P(\mathbf{y}_{j+1})$ and \mathbf{y}_j , and $S(\mathbf{y}_{l-1})$ and \mathbf{y}_l) also affects J_0 .

Although a set \mathcal{L} which performs many very small jumps can conceivably attain a very low value of J_0 , there is an exceedingly small probability to find it within a finite dataset. An illustration of this statement comes from the histogram in figure 1, which shows the distribution of distances between each reconstructed vector and its closest neighbor for the dataset discussed in section III: as expected the frequency of closest neighbors quickly drops to zero for short distances. Then our strategy for finding an initial guess for a downhill minimizer will be that of looking for a set \mathcal{L} which performs as few jumps as possible, giving a somewhat lesser importance to the task of keeping the jumps very small.

Let us call *orbit* any sequence of valid vectors which does not jump. The first vector of an orbit shall have a valid predecessor, and the last a valid successor. Thus we define the *predecessor of the orbit* as the predecessor of its first vector and likewise the *successor of the orbit* as the successor of its last vector. We say that an orbit is *consecutive* to a point if its successor or its predecessor is the closest neighbor of the point. Two orbits are *consecutive* if the successor of one orbit is the closest neighbor of the first vector of the other orbit, or if the predecessor of one orbit is the closest neighbor of the last vector of the other orbit. Let us call *branch* a set made of consecutive orbits. Below we describe a simple algorithm

to construct a set \mathcal{L} by joining together one or more consecutive orbits.

1. We follow forward in time the orbit consecutive to \mathbf{x}_p for l steps, or until it has a valid successor. We store away the set of points made of \mathbf{x}_p followed by the points of this orbit as the *1-jump forward branch*.
2. For each point \mathbf{y}_j of each $(n-1)$ -jumps forward branch (where $j = r, 2r, \dots \leq l_f$, r is an arbitrary stride, $l_f + 1$ is the number of points in the forward branch, and $l_f \leq l$), we follow forward in time the orbit consecutive to \mathbf{y}_j for $l - j$ steps, or until it has a valid successor. We store away all the points up to \mathbf{y}_j of the current forward branch followed by the points of the consecutive orbit as one of the n -jumps forward branches.
3. We repeat step 2 for a fixed number n_f of times.
4. We follow backward in time the orbit consecutive to \mathbf{x}_q for l steps, or until it has a valid predecessor. We store away the points of the consecutive orbit followed by \mathbf{x}_q as the *1-jump backward branch*.
5. For each point \mathbf{z}_j of each $(n-1)$ -jumps backward branch (where $j = 0, r, 2r, \dots < l_b$, r is an arbitrary stride, $l_b + 1$ is the number of points in the backward branch, and $l_b \leq l$), we follow backward in time the orbit consecutive to \mathbf{z}_j for j steps, or until it has a valid predecessor. We store away all the points of this orbit followed by all the points from \mathbf{z}_j to the end of the current backward branch as one of the n -jumps backward branches.
6. We repeat step 5 a fixed number n_b of times.
7. For all possible pairs made by one forward branch and one backward branch we examine *synchronous* pairs of points, that is a point \mathbf{y}_{j_f} in the forward branch, and a point \mathbf{z}_{j_b} in the backward branch such that $j_f + l_b - j_b = l$, where $l_b + 1$ is the number of points in the backward branch. If they coincide, or one is the closest neighbor of the other, then we define $\mathcal{L} = \{\mathbf{y}_0, \dots, \mathbf{y}_{j_f}, \mathbf{z}_{j_b+1}, \dots, \mathbf{z}_{l_b}\}$.

The set \mathcal{L} has two properties which make it unsuitable for filling the gap: it does not approximate well the vector field \mathbf{F} in correspondence of the jumps, and it is piecewise identical to portions of the known signal. Both problems may be fixed by iterating a downhill

minimization method using \mathcal{L} as the initial guess. To this purpose we need a discretized form of (3) which allows as argument any finite sequence of points of U (the discretization (4) is defined only for subsets of \mathcal{R}). The simplest among many possibilities relies on finite differences, leading to the following expression

$$J_1(\mathcal{M}) = \sum_{j=1}^l \left| \frac{\mathbf{w}_j - \mathbf{w}_{j-1}}{\Delta t} - \mathbf{F}(\mathbf{w}_{j-1/2}) \right|^2 \quad (5)$$

where $\mathcal{M} = \mathbf{w}_0, \dots, \mathbf{w}_l$, and the vectors \mathbf{w}_j may or may not belong to \mathcal{R} . Here $\mathbf{F}(\mathbf{w}_{j-1/2})$ is the vector field \mathbf{F} evaluated at the midpoint between \mathbf{w}_{j-1} and \mathbf{w}_j ; J_1 is a function of $m(l-1)$ real variables ($\mathbf{w}_0 = \mathbf{x}_p$ and $\mathbf{w}_l = \mathbf{x}_q$ shall be kept fixed), which can be minimized with standard techniques, using \mathcal{L} as the initial guess.

III. AN EXAMPLE

In this section we show how the algorithm described above performs on a time series generated by a chaotic attractor. We integrate numerically the Lorenz equations [10] with the usual parameters ($\sigma = 10$, $r = 28$, $b = 8/3$). We sample the x -variable of the equations with an interval $\Delta t = 0.02$, collecting 5000 consecutive data points which are our time series. One thousand consecutive data points are then marked as “not-valid”, thus inserting in the time series a gap with a width of 1/5th of the series length, corresponding to a time $T = 20$. For this choice of parameters the Lorenz attractor has a positive Lyapunov exponent $\lambda \approx 0.9$ [6], setting the Lyapunov time scale at $\lambda^{-1} \approx 1.1$. We also find that the autocorrelation function of the time series drops to negligible values in about 3 time units. We conclude that T is well beyond any realistic predictability time for this time series.

In the present example we selected the embedding delay $\tau = 5$ simply by visual inspection of the reconstructed trajectory, and we choose an embedding dimension $m = 3$. However, we checked that results are just as satisfactory at least up to embedding delay $\tau = 15$, and embedding dimension $m = 6$.

We apply the algorithm with $n_f = 2$ and $n_b = 0$. The strides are $r = 1$ for the 2-jumps forward orbits and $r = 100$ for the 3-jumps forward orbits. This leads to 11001 forward orbits to be compared with 1 backward orbit, looking for synchronous pairs points which coincide or are closest neighbor of each other. We find two such pair of points, and the corresponding two initial guesses \mathcal{L}_1 and \mathcal{L}_2 are such that $J_0(\mathcal{L}_1) = 1.21$ and $J_0(\mathcal{L}_2) = 1.04$.

The approximating vector field \mathbf{F} is extremely simple, and its choice is dictated solely by ease of implementation. A comparison between different interpolating techniques is off the scope of this paper, and the interested reader may find further information in [7]. If $\bar{\mathbf{x}}_j$ and $\bar{\bar{\mathbf{x}}}_j$ are the vectors of \mathcal{R} , respectively, closest and second closest to a vector \mathbf{w}_j , then we define

$$\mathbf{F}(\mathbf{w}_{j-1/2}) = \frac{\bar{\mathbf{x}}_j - P(\bar{\mathbf{x}}_j) + \bar{\bar{\mathbf{x}}}_j - P(\bar{\bar{\mathbf{x}}}_j)}{2\Delta t}. \quad (6)$$

Using the definition (6) in (5), we obtain $J_1(\mathcal{L}_1) = 3.46$ and $J_1(\mathcal{L}_2) = 3.17$. In order to smooth-out the jumps in the filling sets, the function J_1 is further decreased by iterating five times a steepest-descent line minimization (see, e.g., [11]) using \mathcal{L}_1 and \mathcal{L}_2 as initial guesses. This procedure yields two sets of $l + 1$ points, \mathcal{M}_1 and \mathcal{M}_2 such that $J_1(\mathcal{M}_1) = 1.54$, and $J_1(\mathcal{M}_2) = 1.34$. The corresponding time series are shown in figure 2. The difference between the smoothed sets \mathcal{M}_1 and \mathcal{M}_2 plotted in figure 2 and the sets with jumps \mathcal{L}_1 and \mathcal{L}_2 would be barely noticeable on the scale of the plot.

The effect of the smoothing may be appreciated by looking at figure 3 which shows the region across the jump between two consecutive orbits of \mathcal{L}_1 . The non-smoothed filling set (dashed line) abruptly jumps from one orbit to the other, but the smoothed trajectory singled out by the points of \mathcal{M}_1 (thick solid line) gently moves between them.

No attempt has been made to approach as closely as possible a minimum of J_1 . In fact, we verified that for orbits in \mathcal{R} having the same length as the interpolating sets \mathcal{M}_1 and \mathcal{M}_2 , J_1 ranges (roughly) between 1 and 9. A general criterion to quantify the accuracy with which the field \mathbf{F} approximates the true dynamics of the observed system, may be described as follows: a) an artificial gap is introduced in the signal, b) the vector field is reconstructed, c) a discretized form of the functional J is evaluated for the missing sequence of vectors, which, in this case, is actually known. For the filling-gap problem there is no point in looking for an interpolating set yielding smaller values of the discretized functional than one obtains with artificial gaps of the same length.

IV. DISCUSSION AND CONCLUSIONS.

In this note we have described an algorithm which fills an arbitrarily wide gap in a time series, provided that the dynamic reconstruction method of Takens is applicable. The goal is to provide a filling signal which is consistent with the observed dynamics, in the sense

that, in the reconstructed phase space, the vector tangent to the filling curve should be close to the vector field modeling the observed dynamics. This request is cast as a variational problem, defined by the functional (3). The acceptable degree of closeness is determined by the level of accuracy of the reconstruction, which is quantified by estimating the functional for known orbits having the same length of the gap.

Obviously, if the time series has more than one gap, our method can be applied to all the gaps, independently of each other.

The key idea that greatly simplifies the problem is that of building a rough filling curve by stitching together pieces of the observed dataset. The actual filling curve, which will not be an exact copy of anything present in the observed dataset, is later obtained by refining the rough one. We have illustrated a basic algorithm that embodies this idea, although no attempt has been made at making it computationally optimal. In particular, with the algorithm in its present form, many of the forward and backward branches will be partial copies of each other, because nothing forbids two distinct branches to jump on the same orbit. This leaves some room for improvement, because the effectiveness of the method relies on a substantial amount of the set \mathcal{R} of reconstructed vectors to be explored by a relatively limited number of branches. In a forthcoming, enhanced version of the algorithm some kind of tagging mechanism shall be incorporated in order to produce non-overlapping hierarchies of forward and backward branches.

We observe that this algorithm does not give a guarantee of success: it is perfectly possible that no point of the forward branches is the closest neighbor of (or coincides with) a synchronous point of the backward branches. In this case the obvious attempt is to deepen the hierarchy of the branches, as much as it is computationally feasible. Or, one may relax the request that branches may jump only between closest neighbors, and accept jumps between second or third neighbors as well. As a last resort, one may stitch any pair of forward and backward branches at their closest synchronous points, hoping that the resulting jump could later be smoothed satisfactorily by minimizing the function (5). However, when facing a failure of the algorithm, we believe that first one should question the goodness and appropriateness of the dynamic reconstruction. The presence of too many gaps, the shortness of the time series, or measurement inaccuracies may make the gap-filling problem an insoluble one. We speculate that the ability of filling gaps with relative ease is a way to test the goodness of a dynamic reconstruction.

The ease with which a gap may be filled, as a function of his width, is a problem deserving further work. For the moment we simply recall that if a set of initial conditions of non-zero measure is evolved in time according to (2), eventually it will spread everywhere on the attractor (here the measure is the physical measure μ of the attractor, cfr. ref. [12]). More rigorously, if ϕ_t is the flow associated to (2), and if it is a mixing transformation, then, for any pair of sets A, B of non-zero measure, $\lim_{t \rightarrow \infty} \mu(\phi_t A \cap B) = \mu(A)\mu(B)$. The dispersion of a set of initial conditions is further discussed in [13], where, for example, it is shown that the essential diameter of a set of initial conditions cannot decrease in time, after an initial transient of finite length.

The notion of mixing transformations leads to the idea that wide gaps should be easier to fill than not-so-wide ones, because forward and backward branches have explored larger portions of the attractor, and so there is a greater chance to find synchronous points where they can be joined together. As a first step toward the verification of this hypothesis, we computed the average minimum distance between synchronous points of the 1-jump forward and backward branches as the gap moves along the dataset, for several gap widths. We used the dataset discussed in sec. III, and a ten times longer extension of it. The results, plotted in figure 4, show that the average separation of the branches initially increases as the gap widens, but then it reaches a well-defined maximum and, from there on, decreases as the gap width is further increased.

We push these speculations even further by formulating the hypothesis that the property of being mixing (or, maybe, the lesser property of being ergodic), rather than determinism, is the crucial factor that allows for filling gaps in time series generated by chaotic dynamical systems. Together with the fact that the dynamic reconstruction technique has been successfully applied to stationary stochastic time series, to generate surrogate data with the same statistics of the observed ones [14], we suggest that some yet-unknown modified version of the method illustrated in this paper should be able to fill gaps in a large class of stochastic time series.

Acknowledgements

This work has been supported by *fondo convezione strana* of the Department of Mathematics of the University of Lecce. We are grateful to Prof. Carlo Sempì and to Dr. Fabio

Paronetto for valuable comments.

-
- [1] F. Takens, in *Dynamical Systems and Turbulence*, R. Mañé, D. Rand, and L. S. Young (eds.), Warwick 1980, *Lecture Notes in Math.* (Springer, Berlin, 1981) , Vol. 898.
- [2] T. Sauer, J. A. Yorke, and M. Casdagli, *J. Stat. Phys.*, **65**, (1991), 579.
- [3] M. Casdagli, S. Eubank, J. D. Farmer, and J. Gibson, *Physica D*, **51**, (1991), 52.
- [4] A. Provenzale, L.A. Smith, R. Vio, and G. Murante, *Physica D*, **58**, (1992), 31.
- [5] J. Theiler, *Phys. Rev. A*, **34**, (1986), 2427.
- [6] J. Froyland and K. H. Alfsen, *Phys. Rev. A*, **29**, (1984), 2928.
- [7] M. Casdagli, *Physica D*, **35**, (1989), 335.
- [8] J. D. Farmer and J. J. Sidorowich, *Phys. Rev. Lett.*, **59**, (1987), 845.
- [9] T. Serre, M. Auvergne, and M. J. Goupil, *Astron. Astrophys.*, **259**, (1992), 404.
- [10] E. N. Lorenz, *J. Atmos. Sci.*, **20**, (1963), 130.
- [11] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing*, Cambridge University Press, Cambridge, (1992).
- [12] J.-P. Eckmann and D. Ruelle, **57**, (1985), 617.
- [13] T. Erber, B. Schweizer, and A. Sklar, *Commun. Math. Phys.*, **29**, (1973), 311.
- [14] F. Paparella, A. Provenzale, L. A. Smith, C. Taricco, R. Vio, *Phys. Lett. A*, **235**, (1997), 233.

Figure Captions

Figure 1 Distribution of distances between each reconstructed vector and its closest neighbor for the dataset discussed in sec. III.

Figure 2 Panel A) shows portion of the time series discussed in sec. III. The bold line was removed and the resulting gap was filled by applying the algorithm described in section II. The bold lines in panels B) and C) are two different fillings.

Figure 3 Plot of the first two components of the reconstructed vectors of: \mathcal{M}_1 (bold solid line with solid circles); \mathcal{L}_1 (bold dashed line with crosses); one orbit of \mathcal{L}_1 and its successors (thin line with open circles); the orbit consecutive to it and its predecessors (thin line with open squares). To illustrate the smoothing effect of minimizing functional (5), we only plot a very small portion of these sets in the vicinity of the jump between the consecutive orbits.

Figure 4 Average minimum distance of synchronous points in the 1-jump forward and backward branch as a function of gap width. The dashed line refers to the dataset discussed in sec. III, the solid line refers to a ten times longer extension of that dataset. The vertical line marks the Lyapunov time $\lambda^{-1} \approx 1.1$.

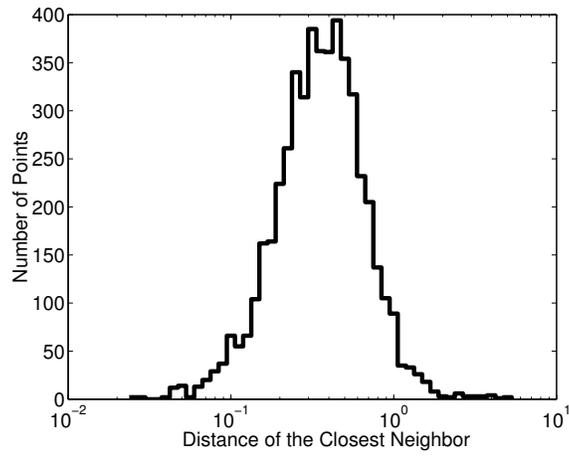


Figure 1: Distribution of distances between each reconstructed vector and its closest neighbor for the dataset discussed in sec. III.

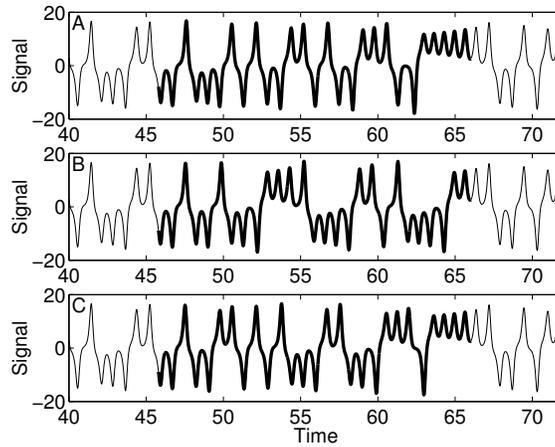


Figure 2: Panel A) shows portion of the time series discussed in sec. III. The bold line was removed and the resulting gap was filled by applying the algorithm described in section II. The bold lines in panels B) and C) are two different fillings.

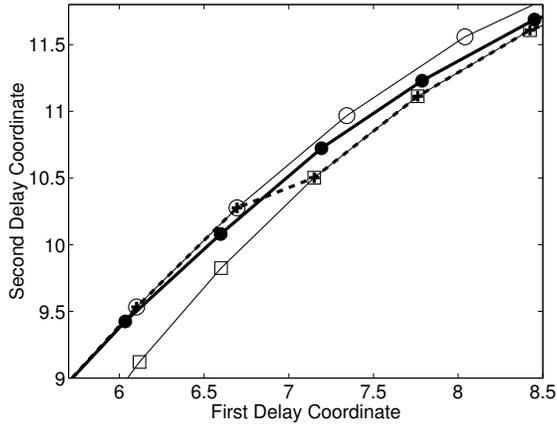


Figure 3: Plot of the first two components of the reconstructed vectors of: \mathcal{M}_1 (bold solid line with solid circles); \mathcal{L}_1 (bold dashed line with crosses); one orbit of \mathcal{L}_1 and its successors (thin line with open circles); the orbit consecutive to it and its predecessors (thin line with open squares). To illustrate the smoothing effect of minimizing functional (5), we only plot a very small portion of these sets in the vicinity of the jump between the consecutive orbits.

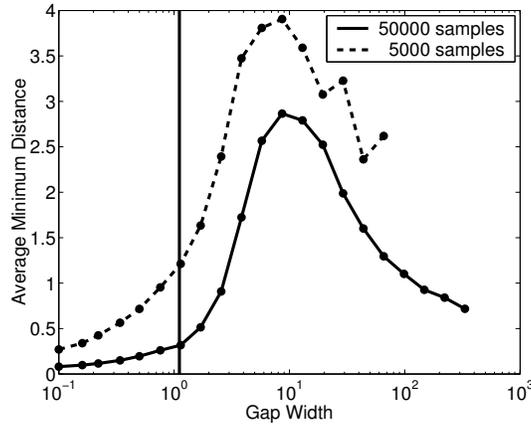


Figure 4: Average minimum distance of synchronous points in the 1-jump forward and backward branch as a function of gap width. The dashed line refers to the dataset discussed in sec. III, the solid line refers to a ten times longer extension of that dataset. The vertical line marks the Lyapunov time $\lambda^{-1} \approx 1.1$.