Finding unstable periodic orbits: a hybrid approach with polynomial optimization

Mayur V. Lakshmi[*1], Giovanni Fantuzzi[1], Sergei I. Chernyshenko[1] and Davide Lasagna[2]

[1]Department of Aeronautics, Imperial College London, London, SW7 2AZ, UK
[2]Faculty of Engineering and the Environment, University of Southampton, Highfield, Southampton, SO17 1BJ, UK

January 26, 2021

Abstract. We present a novel method to compute unstable periodic orbits (UPOs) that optimize the infinite-time average of a given quantity for polynomial ODE systems. The UPO search procedure relies on polynomial optimization to construct nonnegative polynomials whose sublevel sets approximately localize parts of the optimal UPO, and that can be used to implement a simple yet effective control strategy to reduce the UPO's instability. Precisely, we construct a family of controlled ODE systems parameterized by a scalar k such that the original ODE system is recovered for $k = 0$, and such that the optimal orbit is less unstable, or even stabilized, for $k > 0$. Periodic orbits for the controlled system can be more easily converged with traditional methods and numerical continuation in k allows one to recover optimal UPOs for the original system. The effectiveness of this approach is illustrated on three low-dimensional ODE systems with chaotic dynamics.

Keywords. Periodic orbits, nonlinear dynamics, polynomial optimization, auxiliary functions, differential equations

1 Introduction

Computing unstable periodic orbits (UPOs) for systems governed by ordinary differential equations (ODEs) is a fundamental problem in the study of nonlinear dynamical systems that exhibit chaotic dynamics. UPOs embedded in a chaotic attractor provide a "skeleton" around which trajectories evolve through a continuous process of attraction and repulsion along the orbits' stable and unstable manifolds [1]. UPOs are also fundamental to the periodic orbit theory introduced by Cvitanović *et al.* [2, 3], which states that the infinite-time average of any observable of interest over a chaotic trajectory can be expressed as a weighted sum of the time averages over individual UPOs.

A widespread and very effective strategy to find dynamically relevant UPOs is to perform recurrence analysis on long chaotic trajectories in order to identify nearly periodic segments that can be used as initial conditions for a variety of UPO-finding algorithms. The simplest family of such algorithms are shooting methods, which apply the Newton–Raphson algorithm to the Poincaré return map of the dynamical system to compute one (single shooting) or more (multiple shooting) points on the periodic orbit as well as its period (see, for instance, [4, 5]

---

*Email address for correspondence: mayur.lakshmi11@imperial.ac.uk

1

and [6, Chapter 12]). A more robust family of methods are variational ones [7, 8], which attempt to deform a (discrete) closed loop in state space into a UPO by minimizing a non-convex cost function that, loosely speaking, measures the deviation between the tangent to the loop and the direction of the ODE's vector field at each point.

Even though this approach to finding UPOs has been applied very successfully to a wide variety of high-dimensional systems, including turbulent fluid flows (see, e.g., [9–11]), it suffers from an inherent drawback: success relies on the availability of a very good initial approximation to a periodic orbit. In the case of shooting methods this is because trajectories starting from inaccurate initial conditions diverge very quickly from the UPO, whereas the basin of attraction for the Newton–Raphson root finder is often very small [4]. For variational methods, instead, poor initial loops may converge to local minima of the nonconvex cost function that are not periodic orbits of the underlying ODE. Since extracting good initial approximations with recurrence analysis relies on chaotic trajectories shadowing a UPO sufficiently closely for one entire period, one generally can find only orbits that are short and/or have only a few unstable directions. UPOs that are very unstable or lie in parts of the state space rarely visited by chaotic trajectories, however, often remain undetected.

In this work, we present a new strategy for constructing good initial approximations to UPOs that may otherwise be difficult to find. This strategy can be implemented computationally on low-dimensional ODE systems with polynomial dynamics and is part of a broader framework to characterize trajectories that maximize or minimize the infinite-time average of a given quantity of interest Φ. Such *extremal* orbits are useful for control purposes, as their knowledge facilitates the design of optimal control actions to stabilize desired dynamics or suppress undesired ones. Of course, focusing on UPOs that optimize time averages is not restrictive because every periodic orbit is extremal for at least one choice of Φ, namely, the indicator function of the orbit itself. Although this is unknown in practice, varying Φ potentially enables one to identify a large number of periodic orbits.

Underpinning our new approach is the recent realization that nearly sharp bounds on extreme values of time averages can be derived by constructing so-called auxiliary functions of the system's state [12–16], which are similar to the Lyapunov functions used in stability analysis but satisfy a different set of constraints. In addition to providing bounds on extremal time averages, Tobasco *et al.* [15] showed that auxiliary functions provide information about the location and shape of the corresponding extremal trajectory—which, very often, is a UPO [17]. These observations have already been exploited by some of the authors [18] to compute near-extremal UPOs for a nine-dimensional ODE model of shear flow with single shooting methods. However, the inherent ill-conditioning of initial value problems for chaotic ODEs limits the applicability of this approach to cases in which very nearly optimal auxiliary functions can be constructed accurately, which is often not the case.

In this work, we go one step further and show that auxiliary functions can also be used to construct an effective open-loop control action that can be expected to reduce the instability of the extremal UPO and, in some cases, provably stabilizes it. Adding this control to the ODE system produces a family of controlled systems parameterized by the control amplitude, to which traditional shooting or variational methods are more easily applied. One can then attempt to numerically continue any orbit computed with control by decreasing the control amplitude, until a UPO for the original ODE is obtained. In particular, combining this approach with the techniques developed in [18] enables one to search for extremal UPOs robustly with auxiliary functions that are not sufficiently close to being optimal for the latter to work in isolation.

The rest of the paper is organized as follows. Section 2 reviews the auxiliary function method and how auxiliary functions can be leveraged to localize extremal trajectories. Section 3 introduces our control strategy to stabilize UPOs and describes how the construction

of a family of controlled systems allows for the computation of extremal UPOs for polynomial ODEs. The practical potential of this approach is demonstrated in section 4 on three low-dimensional ODE systems that display chaotic dynamics. Computational merits and limitations of our method are discussed in section 5 along with possible directions for further improvement. Finally, section 6 offers concluding remarks.

2 Approximating extremal UPOs with auxiliary functions

Consider an autonomous dynamical system governed by the ODE

$$
(1) \qquad \frac{\mathrm{d}\mathbf{a}}{\mathrm{d}t} = \mathbf{f}(\mathbf{a}), \quad \mathbf{a}(0) = \mathbf{a}_0,
$$

where $\mathbf{a} \in \mathbb{R}^n$ and $\mathbf{f} : \mathbb{R}^n \to \mathbb{R}^n$ is smooth. The infinite-time average of a function $\Phi(\mathbf{a})$ along the trajectory starting from $\mathbf{a}_0$ is defined as

$$
(2) \qquad \overline{\Phi}(\mathbf{a}_0) := \lim_{\tau \to \infty} \frac{1}{\tau} \int_0^\tau \Phi\left[\mathbf{a}(t\,;\mathbf{a}_0)\right] \mathrm{d}t,
$$

where $\mathbf{a}(t\,;\mathbf{a}_0)$ denotes the trajectory $\mathbf{a}(t)$ with initial condition $\mathbf{a}_0$, and we assume for simplicity that the limit exists. We are interested in the maximal value of $\overline{\Phi}$ over all bounded trajectories,

$$
(3) \qquad \overline{\Phi}^* := \max_{\substack{\mathbf{a}_0 \in \mathbb{R}^n : \\ \|\mathbf{a}(t\,;\mathbf{a}_0)\| < \infty \ \forall t}} \overline{\Phi}(\mathbf{a}_0),
$$

as well as the initial conditions and corresponding trajectories which achieve it. Note that considering maximal time averages only is not restrictive, as minimizers of $\overline{\Phi}$ coincide with maximizers of $\overline{-\Phi}$.

Upper bounds on $\overline{\Phi}^*$ can be computed in a relatively straightforward way. Suppose that there exist a function $P(\mathbf{a})$ such that $\overline{P} = 0$ and a constant U such that $\Phi(\mathbf{a}) + P(\mathbf{a}) \leq U$ for any $\mathbf{a}$. Averaging this inequality along the trajectory $\mathbf{a}(t\,;\mathbf{a}_0)$ yields $\overline{\Phi}(\mathbf{a}_0) \leq U$ for any $\mathbf{a}_0$, so $\overline{\Phi}^* \leq U$. To construct a function $P(\mathbf{a})$ with zero average, one can take $P(\mathbf{a}) = \mathbf{f}(\mathbf{a}) \cdot \nabla V(\mathbf{a})$ with any $V : \mathbb{R}^n \to \mathbb{R}$ in the class C^1 of continuously differentiable functions. Indeed, along any bounded trajectory of (1) the chain rule gives

$$
(4) \qquad \overline{P[\mathbf{a}(t\,;\mathbf{a}_0)]} = \overline{\mathbf{f}[\mathbf{a}(t\,;\mathbf{a}_0)] \cdot \nabla V[\mathbf{a}(t\,;\mathbf{a}_0)]}
$$

$$
= \overline{\frac{\mathrm{d}}{\mathrm{d}t} V[\mathbf{a}(t\,;\mathbf{a}_0)]}
$$

$$
= \lim_{\tau \to \infty} \frac{V[\mathbf{a}(\tau\,;\mathbf{a}_0)] - V(\mathbf{a}_0)}{\tau}
$$

$$
= 0.
$$

The best upper bound on $\overline{\Phi}^*$ is obtained by optimizing over the choice of V, hereafter called *auxiliary function*:

$$
(5) \qquad \overline{\Phi}^* \leq \inf_{V \in C^1} \left\{ U \mid U - \Phi(\mathbf{a}) - \mathbf{f}(\mathbf{a}) \cdot \nabla V(\mathbf{a}) \geq 0 \quad \forall \mathbf{a} \in \mathbb{R}^n \right\}.
$$

If Φ and $\mathbf{f}$ are polynomial, feasible polynomial auxiliary functions and their corresponding bounds on $\overline{\Phi}^*$ can be constructed computationally upon replacing the inequality constraint in (5) with the stronger requirement that the polynomial $U - \Phi(\mathbf{a}) - \mathbf{f}(\mathbf{a}) \cdot \nabla V(\mathbf{a})$ be a sum of squares (SOS) [12–14, 18–20], which can be handled using efficient algorithms for convex

3

optimization [21–23]. Moreover, if all trajectories of (1) are absorbed by a compact set Ω, then the constraint in (5) can be restricted to Ω and the corresponding SOS computations are guaranteed to return arbitrarily sharp bounds on $\overline{\Phi}^*$ [15, 18]. More precisely, for any $\delta > 0$ one can construct a polynomial auxiliary function V_δ which provides a bound U_δ with $\overline{\Phi}^* \leq U_\delta \leq \overline{\Phi}^* + \delta$. Importantly, this often happens in practice even when no absorbing set Ω exists.

The crucial observation for this work is that any near-optimal auxiliary function constructed with polynomial optimization not only produces an upper bound on $\overline{\Phi}^*$, but can also be used to localize the associated extremal trajectories in state space – which are, very often, UPOs. To understand this localization framework, suppose first that there exists an optimal auxiliary function V_0 proving a sharp upper bound $U_0 = \overline{\Phi}^*$. Then, any extremal trajectory $\mathbf{a}(t)$ must satisfy [13, 15]

$$(6) \qquad \overline{\overline{\Phi}^* - \Phi[\mathbf{a}(t)] - \mathbf{f}[\mathbf{a}(t)] \cdot \nabla V_0[\mathbf{a}(t)]} = 0.$$

Since the quantity being averaged in (6) is nonnegative it must actually vanish pointwise in time, so the extremal trajectory necessarily lies inside the set

$$(7) \qquad \mathcal{S}_0 := \{\mathbf{a} \in \mathbb{R}^n \mid \overline{\Phi}^* - [\mathbf{f}(\mathbf{a}) \cdot \nabla V_0(\mathbf{a}) + \Phi(\mathbf{a})] = 0\}.$$

Although this set may contain also points that are not on the extremal periodic orbit, it provides guidance to locate the extremal trajectory.

In practice, an optimal auxiliary function V_0 is rarely available and one can only construct a δ-suboptimal one, V_δ, with corresponding bound U_δ. Nevertheless, it was shown in [15] that, if the extremal trajectory is a periodic orbit, then for any $\varepsilon \geq \delta$ it must lie inside the set

$$(8) \qquad \mathcal{S}_\varepsilon = \{\mathbf{a} \in \mathbb{R}^n \mid D(\mathbf{a}) \leq \varepsilon\}$$

for a fraction of its time period no smaller than $F := 1 - \delta/\varepsilon$, where

$$(9) \qquad D(\mathbf{a}) := U_\delta - \mathbf{f}(\mathbf{a}) \cdot \nabla V_\delta(\mathbf{a}) - \Phi(\mathbf{a}).$$

Since the polynomial $D(\mathbf{a})$ is nonnegative by construction, when $\delta \ll 1$ (that is, when the auxiliary function is close to being optimal) taking $\delta \ll \varepsilon \ll 1$ often allows one to keep $F \approx 1$ while excluding much of the state space from $\mathcal{S}_\varepsilon$, which therefore localizes large portions of extremal and near-extremal orbits.

While constructing the entire set $\mathcal{S}_\varepsilon$ is computationally intractable except for ODE systems of very low dimension, obtaining points that lie in it is relatively straightforward. For instance, global minimizers for $D(\mathbf{a})$ lie in $\mathcal{S}_\varepsilon$ for all $\varepsilon \geq \delta$ and can sometimes be recovered directly from the solution of the optimization problem resulting from an SOS relaxation of (5) [23, section 6.1.2]. A simpler and more robust procedure, however, is to numerically search for points where $D(\mathbf{a}) \leq \varepsilon$ by minimizing D using any nonlinear minimization algorithm, initialized from a large number of initial conditions [18]. Since all points in $\mathcal{S}_\varepsilon$ along the extremal periodic orbit are close to being global minimizers for D when ε is small, it is reasonable to expect that ∇D will be small along the part of the extremal periodic orbit contained in $\mathcal{S}_\varepsilon$ and large elsewhere. The minimization routine will therefore quickly descend to this flat region and then slowly approach a local minimum, producing a collection of points in $\mathcal{S}_\varepsilon$ on or close to the extremal periodic orbit as part of the process. This typical situation is illustrated in Figure 1, which shows the unstable limit cycle of the reverse-time van der Pol oscillator

$$(10) \qquad \frac{\mathrm{d}a_1}{\mathrm{d}t} = -3a_2, \qquad \frac{\mathrm{d}a_2}{\mathrm{d}t} = -4(1 - 9a_1^2)a_2$$
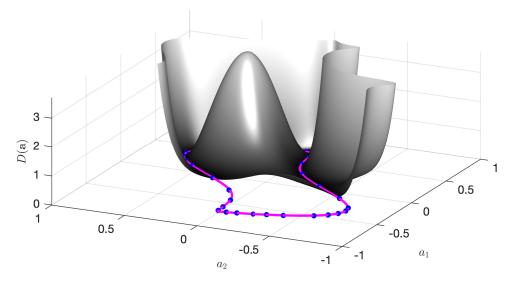
Figure 1: Surface plot of the polynomial $D(\mathbf{a})$ for the rescaled van der Pol oscillator (10) corresponding to the observable $\Phi(\mathbf{a}) = a_1^2 + a_2^2$ and a polynomial auxiliary function V of degree 16. The solid line (—) indicates the limit cycle for (10), which is the extremal trajectory for the chosen Φ. Blue dots (•) show points lying in $\mathcal{S}_{10^{-4}}$ obtained via direct unconstrained minimization of $D(\mathbf{a})$.

along with the polynomial D obtained with $\Phi(\mathbf{a}) = a_1^2 + a_2^2$ (whose infinite-time average on bounded trajectories is maximized on the limit cycle) and a near-optimal polynomial auxiliary function of degree 16. All points on the limit cycle satisfy $D(\mathbf{a}) \leq 10^{-6}$.

As demonstrated in [18], any of the points computed by minimizing $D(\mathbf{a})$ can be used as initial conditions for algorithms that converge to UPOs by evolving the system's dynamics forward in time. However, this basic strategy suffers from two fundamental limitations. The first is that only a finite number of local minima of D may exist when V is suboptimal, to which the minimization routine used will repeatedly converge. The resulting approximation of the extremal periodic orbit can therefore be sparse, making the use of a multiple-shooting or variational method to converge the UPO infeasible in practice. The second fundamental limitation is that, for a given near-optimal auxiliary function and a given ε, the set $\mathcal{S}_\varepsilon$ contains points not on the extremal UPO and possibly far from it [15, section 4]. Thus, the point obtained via minimization of $D(\mathbf{a})$ may not be a good initial condition for single-shooting techniques. In the next section we address the first of these limitations by introducing a novel control methodology to reduce the instability of extremal UPOs.

3   Control methodology

Auxiliary functions can be used not only to localize extremal trajectories, but also be used to formulate an effective control strategy to stabilize them or, at least, reduce their instability. This observation, which is the main contribution of this work, can be used to aid the computation of extremal UPOs with traditional shooting or variational techniques. Subsection 3.1 describes this control strategy in the case of optimal auxiliary functions, for which a rigorous stabilization result can be established without much difficulty. The suboptimal case, which is more relevant in practice but for which we do not have similar theoretical results, is discussed in subsection 3.2. We shall assume throughout that, as is often the case for ODE systems with chaotic dynamics, the trajectory achieving the maximum time average $\overline{\Phi}^*$ is a UPO $\mathcal{O} := \{\mathbf{a}(t) \mid 0 \leq t < T\}$ with period T.

3.1 Optimal auxiliary functions

Suppose that V is an optimal auxiliary function and recall from section 2 that the corresponding nonnegative polynomial $D(\mathbf{a})$ defined in (9) must vanish everywhere along the extremal UPO $\mathcal{O}$, so $\nabla D = 0$ on it. Furthermore, it is not unreasonable to expect that $D(\mathbf{a})$ is steep elsewhere, so $\nabla D \neq \mathbf{0}$ in a neighborhood of $\mathcal{O}$ (cf. Figure 1). Then, consider the ODE

$$\text{(11)} \qquad \frac{\mathrm{d}\mathbf{a}}{\mathrm{d}t} = \mathbf{f}(\mathbf{a}) - k\nabla D(\mathbf{a}), \quad \mathbf{a}(0) = \mathbf{a}_0,$$

where k is an arbitrary nonnegative scalar which takes on the role of a control parameter. Since $\nabla D = 0$ on the extremal UPO $\mathcal{O}$ by construction, $\mathcal{O}$ remains a periodic orbit of (11) for any choice of k, but its stability properties depend crucially on k.

To see this, observe that the two terms on the right-hand side of (11) have a clear dynamical meaning. Given an initial condition $\mathbf{a}_0$ sufficiently close to (but not on) the extremal UPO $\mathcal{O}$, the vector field $\mathbf{f}$ approximately advances the ensuing trajectory along the periodic orbit, but its instability eventually leads to divergence. The $-k\nabla D$ term counteracts this instability by pushing the trajectory back towards the set $\mathcal{S}_0 = \{\mathbf{a} \mid D(\mathbf{a}) = 0\}$, which contains the entire extremal UPO because the auxiliary function V used to construct D is optimal. Thus, the term $-k\nabla D$ effectively acts as a control term that reduces the instability of the UPO and whose authority is proportional to k. In particular, the controlled ODE (11) reduces to the original ODE (1) when $k = 0$, while letting $k \to \infty$ and rescaling time by k^{-1} leads to the ODE governing the steepest-descent minimization of D.

With this intuition in mind, it is not difficult to prove that the orbit $\mathcal{O}$ is locally stable for all sufficiently large k provided that ∇D does not vanish in a neighborhood of $\mathcal{O}$.

Proposition 1. *Assume that the ODE (1) has a periodic orbit $\mathcal{O} = \{\mathbf{a}(t)\}_{t \in (0,T]}$. Suppose also that $D(\mathbf{a})$ is continuously differentiable, that $D(\mathbf{a}) = 0$ on $\mathcal{O}$, and that $D(\mathbf{a}) \geq 0$ in a set $\Omega \subseteq \mathbb{R}^n$ that contains $\mathcal{O}$. Further, assume that there exists a bounded open neighborhood $\mathcal{N}$ of the orbit $\mathcal{O}$ such that $\|\nabla D(\mathbf{a})\| > 0$ for all $\mathbf{a} \in \mathcal{N} \setminus \mathcal{O}$. Then, there exists k_0 such that $\mathcal{O}$ is a locally stable orbit for the controlled system (11) for all $k > k_0$.*

Proof. Since D is continuous, there exists $\gamma_0 > 0$ such that the set $\mathcal{U}_\gamma = \{\mathbf{a} \in \mathcal{N} \mid D(\mathbf{a}) \leq \gamma\}$ is compact for all $\gamma \leq \gamma_0$, so it does not intersect the boundary of $\mathcal{N}$. The same is true for $\mathcal{V}_\gamma = \{\mathbf{a} \in \mathcal{U}_\gamma \mid \frac{1}{2}\gamma \leq D(\mathbf{a}) \leq \gamma\}$ and

$$\text{(12)} \qquad c := \min_{\mathbf{a} \in \mathcal{V}_\gamma} \|\nabla D\|^2 > 0$$

because $\|\nabla D(\mathbf{a})\| > 0$ on $\mathcal{V}_\gamma$. Then,

$$\text{(13)} \qquad \frac{\mathrm{d}}{\mathrm{d}t} D(\mathbf{a}(t)) = \mathbf{f} \cdot \nabla D - k\|\nabla D\|^2 \leq \max_{\mathbf{a} \in \mathcal{V}_\gamma} \|\mathbf{f} \cdot \nabla D\| - kc$$

along any trajectory of (11) starting in $\mathcal{V}_\gamma$. Consequently, D decays along trajectories provided that

$$\text{(14)} \qquad k > \max_{\mathbf{a} \in \mathcal{V}_\gamma} \frac{\|\mathbf{f} \cdot \nabla D\|}{c} =: k_0.$$

In particular, for all such k trajectories cannot escape the set $\mathcal{U}_\gamma$, which is a neighborhood of the orbit $\mathcal{O}$. Moreover, since $\mathcal{O} = \bigcap_{\gamma > 0} \mathcal{U}_\gamma$ we can make this trapping set $\mathcal{U}_\gamma$ arbitrarily small by taking γ arbitrarily small. This proves that $\mathcal{O}$ is locally stable, as claimed. $\qquad \square$

In principle, this stability result enables the computation of the UPO $\mathcal{O}$ for system (1) in a straightforward way: simply find points $\mathbf{a}_0$ with $D(\mathbf{a}_0) \leq \varepsilon$ using nonlinear minimization algorithms, as described in section 2, and use them as initial conditions to simulate the controlled ODE (11). The process can be repeated with increasingly small ε and increasingly large k until $\mathbf{a}_0$ falls within the basin of attraction of $\mathcal{O}$. In practice, however, optimal auxiliary functions required by this approach are rarely available in practice, so one must adjust the procedure to allow for suboptimal ones. We turn to this next.

3.2 Suboptimal auxiliary functions

The controlled ODE (11) can be formulated and solved numerically even when the auxiliary function V used to construct D is suboptimal, but two complications arise. First, the extremal orbit $\mathcal{O}$ is generally not a trajectory of the controlled system, because D and ∇D need not vanish along it when V is suboptimal. Second, and most important, it is possible that taking $k \gg 1$ introduces unwanted stable equilibria in the vicinity of $\mathcal{O}$ where $\mathbf{f} - k\nabla D$ vanishes, preventing trajectories of the controlled system from shadowing $\mathcal{O}$ over an entire period.

To avoid this issue and increase the likelihood that the controlled system possesses a periodic orbit that continuously deforms into $\mathcal{O}$ as k is reduced, we project the control term $-k\nabla D$ onto the subspace perpendicular to $\mathbf{f}$ and replace (11) with

$$(15) \qquad \frac{\mathrm{d}\mathbf{a}}{\mathrm{d}t} = \mathbf{f}(\mathbf{a}) - k\left[I - \frac{\mathbf{f} \otimes \mathbf{f}}{\|\mathbf{f}\|^2}\right]\nabla D(\mathbf{a}), =: \mathbf{h}_k(\mathbf{a}), \qquad \mathbf{a}(0) = \mathbf{a}_0.$$

As proved in Appendix A, the stability result in Proposition 1 for optimal V extends to this modified controlled system under moderate assumptions on the behavior of D and $\mathbf{f}$ near the extremal orbit $\mathcal{O}$. For suboptimal V, however, one cannot guarantee that increasing k will result in the existence of a stable periodic orbit for (15). Nevertheless, provided that V is sufficiently close to optimal, it is not unreasonable to expect that there exists a family of UPOs for the family of systems (15) that connects to the extremal UPO for the original system (1). In addition, if D increases rapidly in the directions normal to $\mathbf{f}$ in the vicinity of the extremal UPO $\mathcal{O}$ as in Figure 1, periodic orbits for large k are likely to be less unstable than the original extremal orbit $\mathcal{O}$ obtained with $k = 0$ because the control term still strongly damps perturbations normal to it.

These heuristic observations suggest that traditional techniques for computing UPOs may be much more effective for large values of k than for $k = 0$. Moreover, since at least part of the original orbit $\mathcal{O}$ must lie in the set $\mathcal{S}_\varepsilon$ defined in (8) when V is near-optimal, points in it are not unlikely to be good initial conditions to find exactly the branch of UPOs that connects to $\mathcal{O}$ as $k \to 0$ (if one exists). Therefore, we propose to search for extremal UPOs using the following 4-step procedure:

(1) Construct a near-optimal polynomial auxiliary function V by solving an SOS relaxation of (5) as described in [12–14, 18, 19].
(2) Construct the polynomial D and, given a tolerance $\varepsilon > 0$, attempt to find points where $D(\mathbf{a}) \leq \varepsilon$ with direct nonlinear minimization of D.
(3) Fix $k > 0$ and search for UPOs of the controlled ODE system (15) using the points obtained at step 2 as initial conditions for shooting methods or for recurrence analysis.
(4) Repeat steps 2 and 3 with increasingly small ε and k until a periodic orbit is found, then perform continuation in k until $k = 0$.

We stress that this procedure is not guaranteed to work because the polynomial D may not behave as illustrated in Figure 1, UPOs for (15) may be hard to find even with $k \gg 1$, and any branch of UPOs one manages to find may not continue up to $k = 0$. Nevertheless, numerical experiments reveal that our strategy is often successful in practice.

7

4 Examples

We now demonstrate the potential of the control methodology described in subsection 3.2 to find extremal UPOs on three low-dimensional ODE systems that display chaotic dynamics. Details of our numerical implementation are discussed in Appendix B.

4.1 A three-dimensional chaotic system

Consider the three-dimensional polynomial ODE system

$$(16) \qquad \frac{\mathrm{d}a_1}{\mathrm{d}t} = a_2 + a_3, \qquad \frac{\mathrm{d}a_2}{\mathrm{d}t} = -a_1 + \frac{1}{2}a_2, \qquad \frac{\mathrm{d}a_3}{\mathrm{d}t} = a_1^2 - a_3,$$

which has two equilibrium points at $(0,0,0)$ and $(-2,-4,4)$ and a chaotic attractor [24]. Trajectories starting outside the basin of the attraction of these invariant structures may become unbounded, but we can still search for extremal UPOs as long as near-optimal auxiliary functions to bound time averages on bounded trajectories can be constructed. This is not guaranteed by the theoretical results in [15, 18] because (16) has no compact absorbing set, but it appears to be true in practice.

We therefore applied our 4-step control strategy to search for extremal UPOs that maximize the infinite-time average of the following observables:

$$(17) \qquad \begin{aligned}
\Phi_1(\mathbf{a}) &= 0.33a_1^2 + 0.27a_1a_2 + 1.28a_1a_3 + 0.88a_2^2 + 0.49a_2a_3 + 0.05a_3^2, \\
\Phi_2(\mathbf{a}) &= 0.71a_1^2 + 0.59a_1a_2 + 0.84a_1a_3 + 0.42a_2^2 + 0.83a_2a_3 + 0.31a_3^2, \\
\Phi_3(\mathbf{a}) &= 0.75a_1^2 + 0.68a_1a_2 + 1.04a_1a_3 + 0.5a_2^2 + 1.52a_2a_3 + 0.38a_3^2, \\
\Phi_4(\mathbf{a}) &= 0.98a_1^2 + 0.3a_1a_2 + 1.42a_1a_3 + 0.6a_2^2 + 1.21a_2a_3 + 0.02a_3^2.
\end{aligned}$$

These were selected from a list of 30 randomly generated quadratic Φ, after removing those whose time average is maximized at one of the equilibria. We deemed this to be the case if the best upper bound on $\overline{\Phi}^*$ obtained with the polynomial optimization techniques of [12–14, 18, 19] differed from the value of Φ at one of the equilibria by less than 0.01 in absolute value.

Results of computations for the observables Φ_1 and Φ_2 are shown in Figures 2 and 3. First, we used polynomial optimization to construct a polynomial auxiliary function V of degree 14, which gives the upper bound $\overline{\Phi_1}^* \leq 0.4217$. Increasing the polynomial degree gives the same bound to within 0.047%, suggesting that our V is very close to being optimal. We then computed local minimizers of the polynomials D corresponding to this near-optimal V. These local minimizers, plotted as red dots in the figure, lie in the set $\mathcal{S}_{10^{-4}}$ and are expected to lie close to the extremal UPO. Recurrence analysis on the trajectory of the controlled ODE system (15) with $k = 0.25$ starting from the best available local minimizer produced a very good initial guess for a UPO, plotted as dot-dashed blue lines in Figure 2, which we converged using the variational algorithm of [8]. Numerical continuation of this UPO in k down to $k = 0$ resulted in a UPO for (16), plotted as a solid black line. The average of Φ_1 over this UPO is within 0.1% of the upper bound on $\overline{\Phi_1}^*$ reported above, strongly suggesting that we have indeed computed the extremal orbit. Results for the observable Φ_2, plotted in Figure 3, lead to a different UPO but are qualitatively analogous, so we do not discuss them for brevity. For both observables, the same results were obtained also when the initial value of the control amplitude k was varied in the range $[0.2, 0.7]$. This suggests that, at least for this particular ODE system, our approach is not very sensitive to the initial choice of k.

Similar calculations for the observables Φ_3 and Φ_4 in (17) led to the discovery of only a third UPO, illustrated in Figure 4, that simultaneously maximizes the average of both quantities. We conclude this because numerical upper bounds on $\overline{\Phi_3}^*$ and $\overline{\Phi_4}^*$ computed
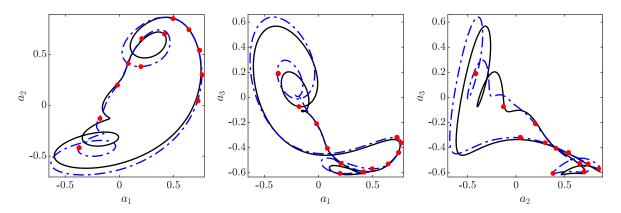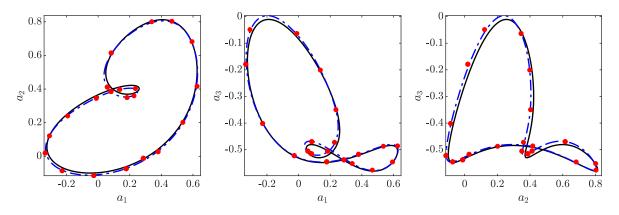
Figure 2: Extremal UPO for system (16) with observable $\Phi_1(\mathbf{a})$ (——). Also plotted are the converged UPO for the corresponding controlled system (15) with $k = 0.25$ (–·–) and local minimizers of the polynomial D (•) obtained with a degree-14 auxiliary function, which lie inside the set $\mathcal{S}_{10^{-4}}$.

Figure 3: Extremal UPO for system (16) with observable $\Phi_2(\mathbf{a})$ (——). Also plotted are the converged UPO for the corresponding controlled system (15) with $k = 0.25$ (–·–) and local minimizers of the polynomial D (•) obtained with a degree-14 auxiliary function, which lie inside the set $\mathcal{S}_{10^{-4}}$.
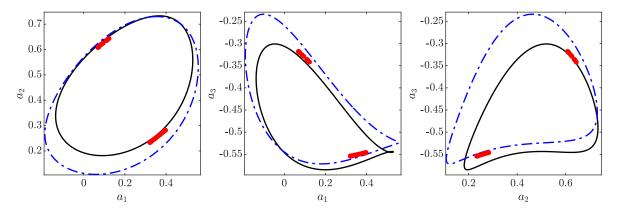

Figure 4: UPO for system (16) that simultaneously maximizes the time average of $\Phi_3(\mathbf{a})$ and $\Phi_4(\mathbf{a})$ (——). Also plotted are the converged UPO for the corresponding controlled system (15) with $k = 0.1$ (–·–) and local minimizers of the polynomial D (•) obtained with a degree-6 auxiliary function, which lie inside the set $\mathcal{S}_{10^{-8}}$.

with degree-10 polynomial auxiliary functions are actually 0.0036% *smaller* than the averages of Φ_3 and Φ_4 on this UPO.[1] However, this orbit could be found using polynomial auxiliary functions of degree as low as six, which yield upper bounds 0.0038% and 0.0041% larger than the respective true averages, and with control amplitude k as small as 0.1. Auxiliary functions of low polynomial degrees are significantly cheaper to construct [12–14, 18, 19] and may be expected to work when the extremal UPO being sought has a simple shape in state space. Indeed, the best control term in (15) is obtained when the function D corresponds to an optimal auxiliary function V. In this ideal case, D must vanish on the extremal UPO (cf. section 2). Therefore, a polynomial V with enough degrees of freedom to approximately satisfy the same constraint may lead to good control terms in (15), even if the corresponding upper bound on $\overline{\Phi}^*$ are far from sharp. We suspect that this is why the extremal UPO for Φ_3 and Φ_4 could be found with degree-6 V while the extremal UPOs for Φ_1 and Φ_2, which have a more complicated shape, required increasing the polynomial degree to 14.

A final interesting observation is that, in Figures 2 and 3, many of the local minimizers of $D(\mathbf{a})$ approximating the extremal UPO actually to lie closer to the converged UPO for the *controlled* system. A possible explanation for this is that the extremal UPO for (16) is not a periodic orbit for the uncontrolled one when V is suboptimal, and need not lie entirely in the set $\mathcal{S}_\varepsilon$ if ε is small. On the other hand, the local minimizers of $D(\mathbf{a})$ do lie in $\mathcal{S}_\varepsilon$ by construction, and the control term in (15) pushes trajectories exactly towards these points. Consequently, local minimizers of D could be better initial conditions for the controlled system, rather than for the uncontrolled one. This, however, may not always be the case (see, for instance, Figure 4) and appears to depend on how suboptimal the auxiliary function is, on the chosen value of the control parameter k, and on how "flat" the polynomial D is in the vicinity of the extremal UPO along directions perpendicular to it.

4.2 Lorenz–96 system

We next study the five-dimensional Lorenz–96 ODE system [25]

$$(18) \qquad \frac{\mathrm{d}a_i}{\mathrm{d}t} = (a_{i+1} - a_{i-2})a_{i-1} - a_i + F, \qquad i = 1, \dots, 5,$$

where we adopt the convention that $a_{-1} = a_4$, $a_0 = a_5$ and $a_6 = a_1$. The scalar F is a constant forcing term and the system has a unique equilibrium point $\mathbf{a}_0 = (F, \dots, F)$. Fixing $F = 8$ we applied the 4-step strategy outlined in subsection 3.2 with a degree-10 polynomial function and $k = 0.5$ to search for the extremal UPO for the observable

$$(19) \qquad \Phi(\mathbf{a}) = (a_1 - F)^2 + (a_4 - F)^2,$$

whose time average is certainly not maximized by the equilibrium. Local minima of the polynomial D, the converged UPO for the controlled system (15), and the final UPO with $k = 0$ are illustrated in Figure 5. The fact that the upper bound on $\overline{\Phi}^*$ obtained in step 1 exceeds the numerical of average of Φ over the converged UPO by only 0.43% strongly suggests that we have indeed found the extremal one.

It is instructive to ask whether our control strategy offers any advantages over the method from [18], which attempts to find the extremal UPO directly by using local minimizers of D (supplemented with a guess for the orbit's period) as initial conditions for a single-shooting Newton–Raphson method. With the set of local minimizers shown in Figure 5, the single-shooting Newton–Raphson algorithm fails to converge to any UPO for the Lorenz–96 system (18) even when the initial guess for the period is taken to be equal to the period T^* of

---

[1]This apparent contradiction is due to unavoidably finite tolerances in the algorithms used to optimize the upper bounds, which may return slightly infeasible answers.
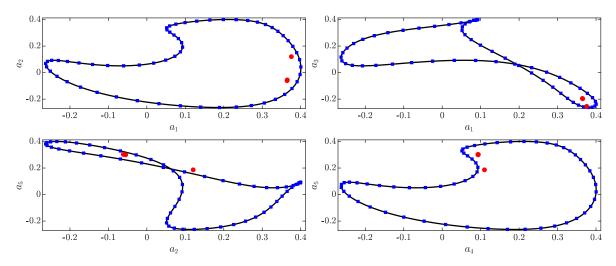
Figure 5: UPO for system (18) that maximizes the time average of $\Phi(\mathbf{a}) = (a_1 - F)^2 + (a_4 - F)^2$ for $F = 8$ (—). Also plotted are the converged UPO for the corresponding controlled system (15) with $k = 0.5$ (■) and local minimizers of the polynomial D (●) obtained with a degree-10 auxiliary function, which lie inside the set $\mathcal{S}_{10^{-4}}$.

the target extremal UPO. To understand this failure in convergence, we integrated both the uncontrolled Lorenz–96 system and the corresponding controlled system with $k = 0.4$ over a time interval of length T^*, taking the best local minimizer of D as the initial condition. The shooting error $\|\mathbf{a}(T^*; \mathbf{a}_0) - \mathbf{a}_0\| / \|\mathbf{a}_0\|$ was 0.256 for the controlled trajectory and 0.396 for the uncontrolled one. This clearly demonstrates the efficacy of the control term in (15) in reducing the instability of the orbit, which is too unstable for the shooting strategy of [18] to work. This difficulty could not be resolved by working with a more accurate polynomial auxiliary function of degree-14, highlighting that the sensitivity of the single-shooting approach to poor initial conditions cannot be easily mitigated by increasing the accuracy of V. In contrast, the control strategy described in this paper works robustly.

4.3 A model of shear flow

As a final example, we consider a nine-dimensional ODE system modeling sinusoidally forced shear flow in a periodic channel [26]. The system takes the form

$$
(20) \qquad \frac{da_i}{dt} = \lambda_1 \delta_{1i} - \frac{1}{Re} \lambda_j a_j + N_{ijk} a_j a_k, \qquad i, j, k = 1, \ldots, 9,
$$

where summation over indices j and k is assumed, δ_{1i} is the usual Kronecker delta, Re is a fixed constant representing the flow's Reynolds number, and λ_j, N_{ijk} are numerical coefficients corresponding to the "NBC" configuration in [26, 27].

At all values of Re, system (20) has a locally stable equilibrium point $\mathbf{a}_l = (1, 0, \ldots, 0)$, which represents the laminar flow state. Trajectories display chaotic behavior as Re is raised and a large number of UPOs have been computed for $Re > 80.54$ [26, 27]. Here, we fix $Re = 120$ and look for the UPO that maximizes (in a time-averaged sense) the energy of perturbations from the laminar flow,

$$
(21) \qquad \Phi(\mathbf{a}) := \|\mathbf{a} - \mathbf{a}_l\|^2.
$$

The result of computations with a degree-10 auxiliary function and $k = 0.4$, are shown in Figure 6. The numerical average of Φ over the converged UPO for system (20) is only 0.63% less than our upper bound on Φ. Once again, this strongly suggests that our method
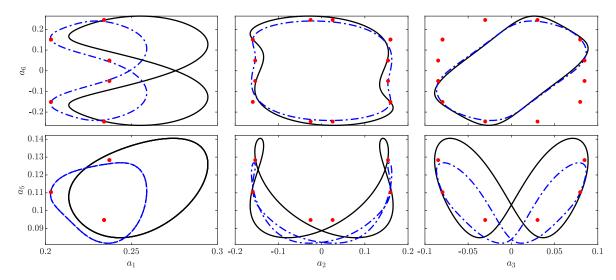
11

Figure 6: UPO for system (20) that maximizes the time average of $\Phi(\mathbf{a}) = \|\mathbf{a} - \mathbf{a}_l\|^2$ at $Re = 120$ (——). Also plotted are the converged UPO for the corresponding controlled system (15) with $k = 0.4$ (–·–) and local minimizers of the polynomial D (•) obtained with a degree-10 auxiliary function, which lie inside the set $\mathcal{S}_{10^{-4}}$.

yields the extremal UPO for Φ. As in subsection 4.2, the uncontrolled single-shooting method of [18] fails to produce any periodic orbit, even when the initial guess for the period is taken to be equal to the period T^* of the target extremal UPO. Integrating both the uncontrolled system (20) and its controlled counterpart (15) with $k = 0.4$ over a time interval of length T^* using the best local minimizer of D as the initial condition, we find that the shooting error $\|\mathbf{a}(T^*; \mathbf{a}_0) - \mathbf{a}_0\|/\|\mathbf{a}_0\|$ is 0.0174 for the controlled trajectory and 0.3751 for the uncontrolled one. This again illustrates the stabilizing effect of the control. Of course, we cannot exclude that the single-shooting method of [18] will work if one uses polynomial auxiliary functions of higher degree, but doing so would require significantly larger computational resources than those available to this study. The control strategy approach presented here, however, could be implemented successfully without difficulties.

5 Discussion

The examples above demonstrate that the four-step strategy introduced in this work enables one to compute extremal UPOs robustly, at least for low-dimensional ODE systems. The ability to focus on extremal orbits is a particular advantage of our method, because performing recurrence analysis on chaotic trajectories might easily miss UPOs with extreme behavior that do not contribute significantly to the statistics of the chaotic attractor.

On the other hand, the need to construct polynomial auxiliary functions that produce sufficiently accurate upper bounds on the extremal time average $\overline{\Phi}^*$ of the quantity of interest currently poses a significant barrier to scalability. This is because the computational complexity of the polynomial optimization techniques described in [12–14, 18, 19] grows combinatorially as the number of ODE states and the degree of the auxiliary function is increased [23]. As a result, the largest polynomial degree that can currently be considered for the nine-dimensional system in subsection 4.3 is approximately 10 on a workstation with 64GB of RAM, and reduces to no more than 4 or 6 for ODEs with a few tens of states. Nevertheless, removing computational bottlenecks in general polynomial optimization and in its applications to dynamical systems are problems that have attracted significant interest

in recent years (see, e.g., [28–37]), so we expect that our UPO search strategy will become practical for ODEs of moderate dimension in the near future.

With computational aspects in mind, a particularly attractive aspect of the control strategy proposed in this paper is that its four steps do not depend on the particular algorithms used to carry them out. This enables one to use more sophisticated numerical techniques not only to construct polynomial auxiliary functions, but also to converge UPOs for the controlled system (15). For example, one could augment multiple-shooting and variational algorithms with recent approaches to identifying near-periodic trajectory segments based on dynamic mode decomposition [38], which are more robust than the traditional recurrence analysis employed in this work.

Finally, all results reported in this paper were obtained by fixing the control amplitude k in (15) to an arbitrary constant value in the interval $(0, 1)$. Methods to optimize k statically or dynamically by taking into account the characteristics of the polynomial $D(\mathbf{a})$ would be a highly valuable addition to our approach, because simply fixing a large k value would mean that the periodic orbit for the controlled system (15), if it exists, lies far away from the extremal UPO for the uncontrolled system (1). This situation is not ideal, because then a continuous branch connecting a periodic orbit for the controlled system to the extremal UPO for the uncontrolled one may not exist. On the other hand, taking k too small could mean that the stabilizing effect of the control term is not sufficient to find any near-periodic trajectory segments. There clearly is an optimal choice for k, but there are also many ways to formulate an optimal control problem for the controlled ODEs (11) or (15) and, without a priori knowledge of an extremal UPO and/or of the polynomial D, it is not immediately clear which one leads to the best UPO approximation. We leave further investigation of this problem to future work.

6 Conclusion

We have presented a novel technique of computing UPOs for ODE systems governed by polynomial dynamics, which overcomes the deficiencies of a related approach presented in [18]. As in that work, the UPO search procedure is initiated by leveraging polynomial optimization techniques to construct an auxiliary function V that proves a near-sharp bound U on the maximal value of the infinite-time average of a prescribed observable Φ. Direct unconstrained minimization of the polynomial $D = U - \Phi - \mathbf{f} \cdot \nabla V$ then yields a set of points which lie close to the extremal UPO for the original system. The novel contribution of this work was to show that this polynomial D can also be used to construct an effective control strategy which reduces the instability of the orbit, and aids its computation with traditional techniques.

More precisely, we have formulated a family of controlled ODE systems (15) parameterized by a control amplitude k, which reduce to the original ODE system when $k = 0$. When the auxiliary function V is optimal and the corresponding D satisfies the assumptions of Proposition 1, a sufficiently large control amplitude k will guarantee the existence of a (locally) stable periodic orbit for the controlled system that can be found simply by time integration. By construction, this periodic orbit coincides with the UPO of the original, uncontrolled system that maximizes the infinite-time average of the given observable Φ. If V is not optimal, which is often the case in practice, stabilization cannot be guaranteed. However, for near-optimal V one expects that periodic orbits for the controlled systems not only still exist, but are also less unstable than the extremal UPOs for the original ODE and, crucially, connect to it continuously as the control amplitude k is reduced to zero. One can therefore first compute an orbit for the controlled system by combining recurrence analysis along trajectories starting from local minimizers of D with shooting or variational algorithms,

and then continue it numerically in k to recover the extremal UPO for the original ODE, which is harder to compute directly.

This process was applied successfully in section 4 to a three-dimensional system with a chaotic attractor [24], a five-dimensional version of the Lorenz–96 system [25], and a nine-dimensional model of shear flow [26]. The control methodology developed in this paper was essential to compute the correct extremal UPOs for a number of different observables Φ. In stark contrast, the uncontrolled approach in [18] failed in all cases. We cannot say if the robust numerical behavior observed in these examples is generic, and further theoretical and computational advances are necessary before complex high-dimensional ODE systems can be tackled using the ideas we have described here. Nevertheless, augmenting traditional techniques for converging UPOs with recent frameworks for ODE analysis via polynomial optimization promises to be a fruitful avenue of research.

A   A stability result for ODE (15)

In the case of optimal V, the stability argument in Proposition 1 can be extended to the controlled system (15) under moderate assumptions on D and $\mathbf{f}$. Proceeding as in the proof of Proposition 1, one must show that the right-hand side of the inequality

$$(22) \qquad \frac{\mathrm{d}D}{\mathrm{d}t} \leq \max_{\mathbf{a} \in \mathcal{V}_\gamma} |\mathbf{f} \cdot \nabla D| - k\|\nabla D\|^2 \left(1 - \frac{|\mathbf{f} \cdot \nabla D|^2}{\|\mathbf{f}\|^2 \, \|\nabla D\|^2}\right)$$

is negative at all points $\mathbf{a} \in \mathcal{V}_\gamma$ for sufficiently large k. This is true if D and $\mathbf{f}$ satisfy

$$(23) \qquad \frac{|\mathbf{f} \cdot \nabla D|}{\|\mathbf{f}\| \, \|\nabla D\|} < \mathrm{const} < 1 \qquad \forall \mathbf{a} \in \mathcal{V}_\gamma.$$

This constraint is weak, because $\nabla(\mathbf{f} \cdot \nabla D) = \nabla \mathbf{f} \, \nabla D + \mathbf{f} \cdot (\nabla \otimes \nabla D)$ vanishes on the orbit: the first term in the sum is zero there because $\nabla D = \mathbf{0}$, while the second term is zero because it is the material derivative of ∇D along the orbit. As a result, one can expect that in the vicinity of the orbit $\mathbf{f} \cdot \nabla D \sim d^2$ while $\|\nabla D\| \sim d$, where d is the distance to the orbit. This, in particular, will be the case at all the points on the orbit where the Hessian matrix $\nabla \otimes \nabla D$ is positive definite in the subspace perpendicular to $\mathbf{f}$. One can therefore ensure (23) by taking γ small enough.

B   Numerical implementation

For all examples presented in section 4, polynomial auxiliary functions were constructed with the polynomial optimization framework described in [12–14, 18, 19] using a customized version[2] of the MATLAB optimization toolbox YALMIP [39] and optimization solver MOSEK v.9.2 [40]. Approximate local minima of the corresponding polynomial $D(\mathbf{a})$ were found using the BFGS quasi-Newton method [41–44] implemented in MATLAB's built-in function `fminunc`, with the step tolerance (relative lower bound on the size of an iteration step) and the first-order optimality tolerance (lower bound on $\|\nabla D\|_\infty$) both set to 10^{-16}.

The best local minimum was then used to integrate the controlled system (15) forward in time for a value of k fixed arbitrarily, and recurrence analysis was employed to identify

---

[2]Available from https://github.com/aeroimperial-optimisation/aeroimperial-yalmip

near-periodic trajectory segments [45, 46]. In our implementation, a near recurrence of period T was deemed to have occurred at time t if the quantity

$$(24) \qquad R(t, T) := \frac{\|\mathbf{a}(t) - \mathbf{a}(t - T)\|}{\|\mathbf{a}(t)\|} \leq 0.025.$$

The portion of trajectory between times $t - T$ and T was then used as an initial guess to converge a periodic orbit using the variational approach described in [8]. This method requires minimizing the cost function

$$(25) \qquad C(\hat{\mathbf{a}}_0, \ldots, \hat{\mathbf{a}}_{N-1}, T) := \frac{N}{2T} \sum_{i=0}^{N-1} \left\| \hat{\mathbf{a}}_{i+1} - \hat{\mathbf{a}}_i - \frac{T}{N} \mathbf{h}_k\left(\frac{\hat{\mathbf{a}}_{i+1} + \hat{\mathbf{a}}_i}{2} \right) \right\|^2$$

over the orbit's period T and over N points $\hat{\mathbf{a}}_0, \hat{\mathbf{a}}_1, \ldots, \hat{\mathbf{a}}_{N-1}$ distributed at equal time intervals T/N along the orbit. (In writing (25), we used the convention that $\hat{\mathbf{a}}_N = \hat{\mathbf{a}}_0$ by periodicity.) We solved this nonlinear least-squares problem using the Levenberg–Marquardt algorithm [47, 48] implemented in the MATLAB built-in function `lsqnonlin`, with the function tolerance on C (relative lower bound on the change in the value of C after an iteration) and the step tolerance on the vector $(\hat{\mathbf{a}}_0, \ldots, \hat{\mathbf{a}}_{N-1}, T)$ both set to 10^{-16}.

Periodic orbits for the controlled ODE system (15) computed in this way were then numerically continued in k down to $k = 0$. This could be done with a variety of sophisticated numerical approaches (see, e.g., [49]); here, however, we simply minimized (25) at increasingly small values of k using the minimizer from the previous computation as the initial condition.

References

[1] P. Cvitanović. Periodic orbits as the skeleton of classical and quantum chaos. *Phys. D*, 51:138–151, 1991.

[2] R. Artuso, E. Aurell, and P. Cvitanovic. Recycling of strange sets: I. Cycle expansions. *Nonlinearity*, 3:325–359, 1990.

[3] P. Cvitanović. Dynamical averaging in terms of periodic orbits. *Phys. D*, 83:109–123, 1995.

[4] C. T. Kelley. *Solving Nonlinear Equations with Newton's Method*. Fundam. Algorithms. SIAM, Philadelphia, 2003.

[5] U. M. Ascher, R. M. M. Mattheij, and R. D. Russell. Initial Value Methods. In *Numerical Solution of Boundary Value Problems for Ordinary Differential Equations*, Classics Appl. Math., chapter 4, pages 132–184. SIAM, Philadelphia, 1995.

[6] P. Cvitanović, R. Artuso, R. Mainieri, G. Tanner, and G. Vattay. *Chaos: Classical and Quantum*. Niels Bohr Institute, Copenhagen, 2017. (version 15.9).

[7] Y. Lan and P. Cvitanović. Variational method for finding periodic orbits in a general flow. *Phys. Rev. E*, 69:016217, 2004.

[8] B. M. Boghosian, L. M. Fazendeiro, J. Lätt, H. Tang, and P. V. Coveney. New variational principles for locating periodic orbits of differential equations. *Philos. Trans. Roy. Soc. A*, 369:2211–2218, 2011.

[9] L. Fazendeiro, B. M. Boghosian, P. V. Coveney, and J. Lätt. Unstable periodic orbits in weak turbulence. *J. Comput. Sci.*, 1:13–23, 2010.

[10] G. J. Chandler and R. R. Kerswell. Invariant recurrent solutions embedded in a turbulent two-dimensional Kolmogorov flow. *J. Fluid Mech.*, 722:554–595, 2013.

[11] D. Lucas and R. Kerswell. Spatiotemporal dynamics in two-dimensional Kolmogorov flow over large domains. *J. Fluid Mech.*, 750:518–554, 2014.

[12] S. I. Chernyshenko, P. Goulart, D. Huang, and A. Papachristodoulou. Polynomial sum of squares in fluid dynamics: A review with a look ahead. *Philos. Trans. Roy. Soc. A*, 372:20130350, 2014.

[13] G. Fantuzzi, D. Goluskin, D. Huang, and S. I. Chernyshenko. Bounds for Deterministic and Stochastic Dynamical Systems using Sum-of-Squares Optimization. *SIAM J. Appl. Dyn. Syst.*, 15:1962–1988, 2016.

[14] D. Goluskin. Bounding Averages Rigorously Using Semidefinite Programming: Mean Moments of the Lorenz System. *J. Nonlinear Sci.*, 28:621–651, 2018.

[15] I. Tobasco, D. Goluskin, and C. R. Doering. Optimal bounds and extremal trajectories for time averages in nonlinear dynamical systems. *Phys. Lett. A*, 382:382–386, 2018.

[16] M. Korda, D. Henrion, and I. Mezić. Convex computation of extremal invariant measures of nonlinear dynamical systems and Markov processes. *J. Nonlinear Sci.*, 31:14(1–26), 2021.

[17] T.-H. Yang, B. R. Hunt, and E. Ott. Optimal periodic orbits of continuous time chaotic systems. *Phys. Rev. E*, 62:1950–1959, 2000.

[18] M. V. Lakshmi, G. Fantuzzi, J. D. Fernández-Caballero, Y. Hwang, and S. I. Chernyshenko. Finding extremal periodic orbits with polynomial optimization, with application to a nine-mode model of shear flow. *SIAM J. Appl. Dyn. Syst.*, 19:763–787, 2020.

[19] D. Goluskin and G. Fantuzzi. Bounds on mean energy in the Kuramoto–Sivashinsky equation computed using semidefinite programming. *Nonlinearity*, 32:1705–1730, 2019.

[20] D. Lasagna, D. Huang, O. R. Tutty, and S. Chernyshenko. Sum-of-squares approach to feedback control of laminar wake flows. *J. Fluid Mech.*, 809:628–663, 2016.

[21] P. A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Math. Program., Ser. B*, 96:293–320, 2003.

[22] P. Parrilo. Polynomial Optimization, Sums of Squares, and Applications. In *Semidefinite Optimization and Convex Algebraic Geometry*, MOS-SIAM Ser. Optim., chapter 3, pages 47–157. SIAM, Philadelphia, 2012.

[23] J. B. Lasserre. *An Introduction to Polynomial and Semi-Algebraic Optimization*. Cambridge Texts Appl. Math. Cambridge University Press, Cambridge, 2015.

[24] J. C. Sprott. Some simple chaotic flows. *Phys. Rev. E*, 50:R647–R650, 1994.

[25] E. N. Lorenz. Predictability: A problem partly solved. In *Proceedings of the Seminar on Predictability*, volume 1 of *ECWF Seminar*, pages 1–18, Reading, UK, 1996. ECMWF.

[26] J. Moehlis, H. Faisst, and B. Eckhardt. A low-dimensional model for turbulent shear flows. *New J. Phys.*, 6:56, 2004.

[27] J. Moehlis, H. Faisst, and B. Eckhardt. Periodic Orbits and Chaotic Sets in a Low-Dimensional Model for Shear Flows. *SIAM J. Appl. Dyn. Syst.*, 4:352–376, 2005.

[28] Y. Zheng, G. Fantuzzi, and A. Papachristodoulou. Exploiting sparsity in the coefficient matching conditions in sum-of-squares programming using ADMM. *IEEE Control Syst. Lett.*, 1(1):80–85, 2017.

[29] J. B. Lasserre, K.-C. Toh, and S. Yang. A bounded degree SOS hierarchy for polynomial optimization. *EURO J. Comput. Optim.*, 5(1-2):87–117, 2017.

[30] T. Weisser, J.-B. Lasserre, and K.-C. Toh. Sparse-BSOS: a bounded degree SOS hierarchy for large scale polynomial optimization with sparsity. *Math. Program. Comput.*, pages 1–32, 2017.

[31] A. A. Ahmadi, G. Hall, A. Papachristodoulou, J. Saunderson, and Y. Zheng. Improving efficiency and scalability of sum of squares optimization: Recent advances and limitations. In *Proceedings of the 56th Annual Conference on Decision and Control*, pages 453–462, 2018.

[32] A. Ahmadi and A. Majumdar. DSOS and SDSOS Optimization: More Tractable Alternatives to Sum of Squares and Semidefinite Optimization. *SIAM J. Appl. Algebra Geom.*, 3:193–230, 2019.

[33] Y. Zheng, G. Fantuzzi, and A. Papachristodoulou. Fast ADMM for sum-of-squares programs using partial orthogonality. *IEEE Trans. Automat. Control*, 64(9):3869–3876, 2019.

[34] M. Tacchi, C. Cardozo, D. Henrion, and J.-B. Lasserre. Approximating regions of attraction of a sparse polynomial differential system. arXiv:1911.09500 [math.OC], 2019.

[35] J. Wang, V. Magron, and J.-B. Lasserre. TSSOS: A moment-SOS hierarchy that exploits term sparsity. *SIAM J. Optim.*, 31(1):30–58, 2021.

[36] J. Wang, V. Magron, and J.-B. Lasserre. Chordal-TSSOS: A moment-SOS hierarchy that exploits term sparsity with chordal extension. *SIAM J. Optim.*, 31(1):114–141, 2021.

[37] C. Schlosser and M. Korda. Sparse moment-sum-of-squares relaxations for nonlinear dynamical systems with guaranteed convergence. arXiv:2012.05572 [math.OC], 2020.

[38] J. Page and R. R. Kerswell. Searching turbulence for periodic orbits with dynamic mode decomposition. *J. Fluid Mech.*, 886, 2020.

[39] J. Löfberg. YALMIP : A toolbox for modeling and optimization in MATLAB. In *2004 IEEE International Conference on Robotics and Automation (IEEE Cat. No.04CH37508)*, pages 284–289, New Orleans, 2004.

[40] MOSEK ApS. MOSEK Optimization Toolbox for MATLAB, 2020. (version 9.2.28).

[41] C. G. Broyden. The Convergence of a Class of Double-rank Minimization Algorithms 1. General Considerations. *IMA J. Appl. Math.*, 6:76–90, 1970.

[42] R. Fletcher. A new approach to variable metric algorithms. *Comput. J.*, 13:317–322, 1970.

[43] D. Goldfarb. A Family of Variable-Metric Methods Derived by Variational Means. *Math. Comp.*, 24:23–26, 1970.

[44] D. F. Shanno. Conditioning of Quasi-Newton Methods for Function Minimization. *Math. Comp.*, 24:647–656, 1970.

[45] D. Viswanath. Recurrent motions within plane Couette turbulence. *J. Fluid Mech.*, 580:339–358, 2007.

[46] P. Cvitanović and J. F. Gibson. Geometry of the turbulence in wall-bounded shear flows: Periodic orbits. *Phys. Scr.*, T142:014007, 2010.

[47] K. Levenberg. A method for the solution of certain non-linear problems in least squares. *Quart. Appl. Math.*, 2:164–168, 1944.

[48] D. W. Marquardt. An Algorithm for Least-Squares Estimation of Nonlinear Parameters. *SIAM J. Appl. Math.*, 11:431–441, 1963.

[49] A. Dhooge, W. Govaerts, Y. A. Kuznetsov, H. G. Meijer, and B. Sautois. New features of the software MatCont for bifurcation analysis of dynamical systems. *Math. Comput. Model. Dyn. Syst.*, 14:147–175, 2008.