Convex Optimization Methods for Bounding Lyapunov Exponents

by

Hans Emanuel Oeri
M.Sc., Simon Fraser University, 2018
B.Sc., Simon Fraser University, 2016

A Dissertation Submitted in Partial Fulfillment of the
Requirements for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Mathematics and Statistics

Convex Optimization Methods for Bounding Lyapunov Exponents

by

Hans Emanuel Oeri
M.Sc., Simon Fraser University, 2018
B.Sc., Simon Fraser University, 2016

Supervisory Committee

Dr. David Goluskin, Supervisor
(Department of Mathematics and Statistics)

Dr. Anthony Quas, Committee Member
(Department of Mathematics and Statistics)

Dr. Nishant Mehta, Committee Member
(Department of Computer Science)

ABSTRACT

In dynamical systems, the stability of orbits is quantified by Lyapunov exponents (LEs), which are computed from the average rate of divergence of trajectories. We develop techniques for computing sharp upper bounds on the largest LE over trajectories using methods from convex optimization, which have previously been used to compute sharp bounds on the time averages of scalar quantities on bounded orbits of dynamical systems. For discrete-time dynamics we develop an optimization-based approach for computing sharp bounds on the geometric mean of scalar quantities. We therefore express LEs as infinite-time averages and as geometric means in continuous-time systems and discrete-time systems, respectively, and then derive optimization problems whose solutions give sharp bounds on LEs. When the system's dynamics is governed by a polynomial vector field, the problems can be relaxed to computationally tractable sum-of-squares (SOS) whose solutions also give sharp bounds on LEs. An approach for the practical implementation of a sequence of SOS feasibility problems whose solutions converge to the maximal LE of discrete systems is provided. We explain how symmetries can be used to simplify and generalize the optimization problems in both continuous-time and discrete-time systems. We conclude by discussing the extension of the techniques developed here to the problem of bounding the sum of the leading LEs. Tractable SOS programs are derived for some special cases of this problem.

The applicability of all the techniques developed here is shown by applying them to various explicit examples. For some systems we numerically compute sharp bounds that agree with the the maximal LEs, and for some we prove analytic bounds on maximal LEs by solving the optimization problems by hand.

Table of Contents

List of Tables

List of Figures

Acknowledgements

First of all, I want to thank David Goluskin for being my supervisor and for guiding and supporting me through this PhD. I thank him for letting me work with him and for allowing me to grow and to explore and learn freely. I am especially grateful for his great patience and guidance which have helped me through many difficult times on this journey and which allowed me to see it through to its completion. He taught me much both in and out of the academic sphere and I learned from him the technical knowledge which was needed for the current work. He was always there for his student and willing to discuss the research and share important insights that helped to make the results here possible. I consider myself lucky to have David as my mentor.

I would next like to thank my committee for their feedback and for all the time that they put into helping me complete this thesis. I thank Anthony Quas for many useful discussion, which directly lead to several of the results here including the equivalence of various notions of LE and generalizations of the results on symmetries.

I would also like to thank Giovanni Fantuzzi for his insights about symmetries of SOS programs and I thank Charles Doering for a productive discussion which clarified the formulation of SOS programs for bounding LEs.

I am also grateful to my mother for allowing me to study and for supporting me during the PhD. Finally, I want to thank my partner Yi Sui for her love and support and for always being by my side during this time.

Dedication

To my mother and to the memory of my father.

Chapter 1

Introduction

In many areas of science, processes that are time-dependent are modelled by dynamical systems whose trajectories represent the evolution of the system's state through either continuous or discrete time. If the initial state of a system is known exactly, then a deterministic dynamical system can predict its future state at any later time—in theory. In reality, however, perfect knowledge about a system's initial state is rarely available, and so an important question is to what extent imperfect knowledge of the initial state can be used to make meaningful predictions about the future state. This question comes down to knowing the intrinsic stability of orbits, which may be understood as the rate at which nearby trajectories converge or diverge from each other. The rate of convergence or divergence of nearby trajectories is quantified by Lyapunov exponents (LEs), which explicitly codify the asymptotic stability of a given trajectory [Lya92a; PP16; Din06]. LEs were first introduced in 1892 by Alexandr Lyapunov [Lya92b], who used these exponents as a means of quantifying the stability of solutions of a dynamical system even when analytical solutions are not known. Today, LEs find extensive use in the study of dynamical systems and in applications to physics [FS95], engineering [AX06] artificial intelligence [Vog+20], economics [Pin88], medical sciences [Fra+90] and many more.

Given their widespread use, the need for practical methods for finding LEs is evident. This is not always an easy task however, as techniques based on computing LEs by numerically integrating the dynamical system suffer from low accuracy and high computational cost [DRVV97; ER85; GPL90]. This is especially true for systems that are chaotic, which are inherently unstable and whose trajectories possess a positive LE which quantifies how unstable solutions are to small perturbations. Furthermore, the trajectories of a system often need to be known for a very long time to accurately compute their LEs [PC12; Ben+80; HUP97; TSR01; Mei17]. In particular, determining a system's maximal Lyapunov exponent over all trajectories does not readily lend itself to integration-based approaches, as it involves the numerical integration of the most unstable orbits of a system. Nonetheless, this exponent is of importance for understanding a system's dynamics as it gives an upper bound on the maximal degree of instability and unpredictability in a system's behaviour. We therefore develop techniques for computing maximal LEs of dynamical systems, which are based on convex optimization methods that do not require numerical integration of the system's trajectories.

Consider continuous-time dynamics governed by an autonomous ODE system,

$$\frac{\mathrm{d}}{\mathrm{d}t}x(t) = f(x(t)), \qquad x(0) = x_0 \in \mathbb{R}^n, \tag{1.0.1}$$

where $f : \mathbb{R}^n \to \mathbb{R}^n$. We assume the system is well-posed in the sense that trajectories $x(t)$ are unique and have continuously differentiable dependence on the initial conditions $x_0 \in \mathbb{R}^n$. At each point along $x(t)$, the linearization of (1.0.1) defines a tangent space in which tangent vectors $y(t)$ evolve according to

$$\frac{\mathrm{d}}{\mathrm{d}t}y(t) = Df(x(t))\, y(t), \qquad y(0) = y_0 \in \mathbb{R}^n, \tag{1.0.2}$$

where $Df(x)$ is the $n \times n$ Jacobian matrix for $f(x)$. The quantity $x(t) + y(t)$ is asymptotic to a trajectory of (1.0.1) with initial condition $x_0 + y_0$ when $y_0 \ll 1$ and $|y(t)| \ll 1$. In a nonlinear system where trajectories diverge from $x(t)$ on average, any nonzero perturbation will eventually grow until the linearization (1.0.2) is no longer valid. Thus (1.0.2) may not approximate nearby trajectories uniformly in time, but it captures, at each instant t, the dynamics of trajectories that are asymptotically near $x(t)$. Since (1.0.2) is linear in y by definition, the dynamics of y are independent of its magnitude. We thus fix $|y_0| = 1$ without loss of generality, meaning that y_0 lies on the sphere $\mathbb{S}^{n-1} \subset \mathbb{R}^n$, but $y(t)$ can grow or shrink in time.

For each initial condition $x_0 \in \mathbb{R}^n$ and initial direction $y_0 \in \mathbb{S}^{n-1}$ in the tangent space, we define the corresponding LE as

$$\mu(x_0, y_0) = \limsup_{t\to\infty} \frac{1}{t} \log |y(t)|. \tag{1.0.3}$$

This means that $|y(t)|$ scales like $e^{\mu t}$ on average as $t \to \infty$. In other words, tangent vectors to the trajectory $x(t)$ in the direction $y(t)$ shrink or grow exponentially in time if and only if $\mu < 0$ or $\mu > 0$, respectively. Assuming that f is continuously differentiable and x_0 leads to a bounded trajectory, μ cannot be $\pm\infty$.

Definition (1.0.3) has an analogue for discrete-time systems where the state of a system is updated by iterating it under a map $f(x)$: Consider the discrete dynamical system on $\mathbb{R}^n$

$$x_{k+1} = f(x_k), \quad x_0 \in \mathbb{R}^n \tag{1.0.4}$$

where the discrete time is indexed by the natural number k. Then the linearlization of this system is governed by the Jacobian of $f(x)$ and is given by

$$y_{k+1} = Df(x_k)y_k, \quad y_0 \in \mathbb{R}^n. \tag{1.0.5}$$

In analogy with (1.0.3), the Lyapunov exponent of the discrete map is defined as [PP16],

$$\mu(x_0, y_0) = \limsup_{k\to\infty} \frac{1}{k} \log\left(\frac{|y_k|}{|y_0|}\right). \tag{1.0.6}$$

The value of μ does not depend on the size of $|y_0|$ so we may take $|y_0| = 1$ without loss of generality, but there are some advantages of leaving $|y_0|$ in the expression for μ for manipulating it later. The LE μ in (1.0.6) has the same properties as μ in (1.0.3) except that its value may be $-\infty$ as seen for instance from the system $x_{k+1} = x_k^2$ on the trajectory at $x_0 = 0$. Trajectories of both (1.0.1) and (1.0.4) have n LEs associated with them, although the LEs may not necessarily be unique [Mei17]. In this text we order them from largest to smallest as $\mu_1(x_0) \geq \cdots \geq \mu_n(x_0)$. The Oseledets ergodic theorem [Ose68] implies that the leading exponent $\mu_1(x_0)$ is attained for almost every initial direction y_0, which we will simply refer to as the LE of the orbit. In many applications this is also the exponent which is of greatest interest, because it is the one that is most readily computed and because it determines the observed stability of an orbit.

Given a collection of orbits of (1.0.1) or (1.0.4) which are bounded forward in time and all eventually enter a trapping set $\Omega \subset \mathbb{R}^n$ we want to bound the LE μ_1 over these trajectories. This means that we want to bound the *maximal Lyapunov* *exponent* μ^*_Ω among all trajectories that remain in Ω for positive time,

$$\mu^*_\Omega = \max_{\substack{x_0 \in \Omega \\ y_0 \in \mathbb{S}^{n-1}}} \mu(x_0, y_0). \tag{1.0.7}$$

This definition applies to continuous-time as well as discrete-time systems. We will assume that all x_0 in (1.0.7) are bounded forward in time, otherwise we will exclude these orbits from the maximization.

Computing the maximal LE using numerical integration is challenging because, as explained in the preceding paragraphs, numerically computing μ_1 for even a single trajectory often only produces results with relatively low accuracy and requires explicit knowledge of the system's trajectories for a very long time. The work involved in finding the LEs of a trajectory is therefore usually greater than simply computing the trajectory itself. This is especially restrictive for systems which contain chaotic trajectories with a positive LE, which due to their inherently unstable nature, make knowing the orbit and its linearization for a long time impossible using numerical integration of finite precision. Bounding the largest LE out of a collection of orbits faces the additional difficulty of not being a convex maximization problem, as the LE of each trajectory needs to be found and compared to those of other trajectories to see which one may be the maximal LE. If the set Ω in (1.0.7) contains infinitely many orbits, then such a search over orbits will never be exhaustive, meaning that no matter how many LEs for different orbits one computes, one cannot conclude with certainty that orbits with a larger LE may not be present in the set Ω. For instance, the strange attractor of a chaotic system often has infinitely many periodic orbits embedded in it [Don18; SN12], each of which may have a distinct LE.

We sidestep these issues entirely by formulating the problem of bounding the maximal LE as a convex optimization problem in the case of continuous-time dynamics, and a sequence of convex feasibility problems in the case of discrete-time dynamics. This allows us to obtain bounds on the maximal LE which do not require any a priori knowledge about solutions and which automatically take into account all orbits in the set Ω. We first express the LE in (1.0.3) and (1.0.6) as an infinite-time average in the continuous-time case and a geometric mean in the discrete-time case. In the continuous time case this formula has appeared in the literature as the Furstenberg–Khasminskii formula for stochastic dynamics [AI95; BBPS22]. We then formulate non-negativity constraints whose enforcement implies an upper bound on μ_Ω^*. When the underlying dynamical system in (1.0.1) and (1.0.4) is polynomial in the state space variables, the non-negativity constraints can be enforced via sum-of-squares constraints.

For continuous-time systems a robust framework for bounding infinite time averages along bounded trajectories of a system exists and has been used to compute accurate upper bounds on the time averages of polynomial quantities in such systems [Che+14; Gol18; Lak+20; Ols+20; KHM21; DM22]. The optimization problems that are solved by these authors and the ones which we will use for computing upper bounds on LEs sum-of-squares (SOS) programs, which can be reformulated as semidefinite programs (SDPs). Besides bounding time averages, SOS programs have been used to verify nonlinear stability [Par00; AP15; Haf+18; FGC22; Kun+16], show that trajectories do or do not enter specified sets [HK14; BG20; PGV21; CKH22], bound expectations in stochastic dynamical systems [Kun+16; Fan+16], and bound extrema over certain trajectories [FG20] or global attractors [Gol20]. SOS programming has also been used to approximate maximally invariant sets both of continuous-time [Wan+21] and discrete-time systems [KHJ13; KHJ14].

In discrete-time systems, convex optimization techniques have been applied to show the stability of these systems using both linear programming methods [GH14] and SOS programming [LA08; Kor22] to find Lyapunov functions. Analogous results to those in [Gol18] and [TGD18] exist for maximizing Birkhoff averages [Boc17; BG19; Jen19], using Mañé lemmas. These works also include ergodic theoretic formulations for extremizing LEs. While these formulations provide a useful theoretical foundation for SOS programming in discrete systems, we find that in practice working with geometric means, rather than Birkhoff averages and arithmetic means is sometimes more useful, especially when trying to bound LEs. To our knowledge, explicit formulations for maximizing geometric means and LEs over orbits of discrete systems have not been given thus far and we therefore develop a method for doing so here.

While the aim of the present work is to bound LEs and not to develop the theory of convex optimization itself, SOS programming nonetheless is at the heart of all of the techniques which we present and develop throughout the rest of this text. In order to make use of the optimization methods for bounding LEs of dynamical systems it is therefore essential that the reader is familiar with the basics of SOS optimization. This is also needed to make use of existing software packages for solving SOS programs. We therefore briefly review the basics of semidefinite and SOS programming in the section below.

1.1 Sum-of-squares polynomials and semidefinite programming

Sum-of-squares (SOS) programs are a class of convex optimization problems which minimize an affine objective function subject to polynomial SOS constraints [Nes00; Par00; Las01]. Semidefinite programs (SDPs) on the other hand are a class of convex optimization problems which minimize an affine objective function subject to matrix inequality constraints. As will be explained in this section, SOS programs can equivalently be formulated as SDPs.

A degree-d, n-variate polynomial is said to be sum-of-squares (SOS) if it can be written as the sum-of-squares of other polynomials. We denote the space of degree-d, n-variate SOS polynomials as $\Sigma_{d,n}$. $S(x) \in \Sigma_{d,n}$ if

$$S(x) = \sum_{j=1}^{m} \Big(h_j(x)\Big)^2, \tag{1.1.1}$$

where $h_j(x) \in \mathbb{R}[x]$ are polynomials and m is finite. The space $\Sigma_{d,n}$ is a cone [BV04]. Note that d must necessarily be even. Furthermore, it can be mapped onto the cone $\mathcal{S}_+^q$ of positive semidefinite $q \times q$ matrices.

Let $\mathbf{b}(x) = \{b_1(x), ..., b_q(x)\}$ be a polynomial basis for the space of all n-variate, degree-$d/2$ polynomials. Then $h_j = a_j^\mathsf{T}\mathbf{b}(x)$ for some $a_j \in \mathbb{R}^q$. More concisely, the vector $\mathbf{h}(x) = [h_1(x), ..., h_m(x)]^\mathsf{T}$ can be written as

$$\mathbf{h}(x) = \begin{bmatrix} a_1^\mathsf{T} \\ \vdots \\ a_m^\mathsf{T} \end{bmatrix} \mathbf{b}(x). \tag{1.1.2}$$

So we can write $S(x)$ as

$$S(x) = \mathbf{h}^\mathsf{T}(x)\mathbf{h}(x) = \mathbf{b}^\mathsf{T}(x) \underbrace{\begin{bmatrix} a_1 & \dots & a_m \end{bmatrix} \begin{bmatrix} a_1^\mathsf{T} \\ \vdots \\ a_m^\mathsf{T} \end{bmatrix}}_{\mathsf{M}} \mathbf{b}(x). \tag{1.1.3}$$

The $q \times q$ matrix $\mathsf{M} \succeq 0$ is in $\mathcal{S}_+^q$, and so M is a representation of $S(x)$ in $\mathcal{S}_+^q$. Note however that M depends on the choice of basis $\mathbf{b}(x)$.

Definition 1 (SOS program)**.** *The following convex optimization problem is called a*

SOS program*:*

$$\begin{aligned} \min_{c\in\mathbb{R}^m} \quad & u^\mathsf{T} c \\ s.t. \quad & S(x) = p_0(x) + \sum_{j=1}^m c_j p_j(x) \in \Sigma_{d,n}. \end{aligned}$$

Here $u \in \mathbb{R}^m$ *is given and the polynomials* $p_0(x), ..., p_m(x)$ *are also given.* $S(x) \in \Sigma_{d,n}$ *is called the polynomial constraint.*

In fact, one is not restricted to only a single polynomial constraint and can impose as many as desired, but for clarity only one is written above. As argued above, any SOS polynomial can be represented by a positive semidefinite matrix. Given a basis $\mathbf{b}(x)$, we can therefore write the polynomial in the SOS constraint from definition 1 as

$$\begin{aligned} S(x,c) &= p_0(x) + \sum_{j=1}^m c_j p_j(x) \\ &= \mathbf{b}^\mathsf{T}(x)\Big[\mathsf{M}_0 + \sum_{j=1}^m c_j \mathsf{M}_j\Big]\mathbf{b}(x), \end{aligned} \tag{1.1.4}$$

where the $\mathsf{M}_0, ..., \mathsf{M}_m$ are some matrices representing $p_0(x), ..., p_m(x)$ in the basis $\mathbf{b}(x)$. Matrices $\mathsf{M}_0, ..., \mathsf{M}_m$ can be constructed by solving a linear system for their entries in terms of the coefficients of the $p_0(x), ..., p_m(x)$. Note that these matrices don't necessarily have to be positive semidefinite since the p_j are not necessarily SOS. In summary, this means that there exist $\mathsf{M}_0, ..., \mathsf{M}_m$ such that

$$\mathsf{M}_0 + \sum_{j=1}^m c_j \mathsf{M}_j \in \mathcal{S}_+^q, \tag{1.1.5}$$

and therefore

$$S(x,c) = p_0(x) + \sum_{j=1}^{m} c_j p_j(x) \in \Sigma_{d,n}. \tag{1.1.6}$$

The SOS program in definition 1 is therefore equivalent to the following semidefinite program, which we state in its dual form for clarity.

Definition 2 (Dual SDP)**.** *The dual form of a semidefinite program (See [BV04]) is the following convex optimization problem:*

$$\min \qquad u^{\mathsf{T}} c \tag{1.1.7}$$

$$s.t. \qquad \mathsf{M}_0 + \sum_{j=1}^{m} c_j \mathsf{M}_j \in \mathcal{S}_+^q. \tag{1.1.8}$$

Several numerical techniques as well as commercial software packages are available to solve this type of SDP. We have used the parser YALMIP [Löf09; Löf16] to formulate the SOS programs with the solver Mosek [AA00] to solve the SOS programs which arose in the course of this work.

The currently used numerical technique chooses the basis (1.1.2) to be a subset of the set of monomials in the variables of the polynomial quantities. Once such a basis is fixed, the dual SDP program in definition 2 is determined and an SDP solver like the one mentioned above is used to compute a solution. An unavoidable consequence of using a monomial basis for $\mathbf{b}(x)$ is that this results in a very ill-conditioned SDP to solve. Additionally, as the iterate in the interior point solver approaches the optimal value, the matrices in the SDP become singular, since at least one of their eigenvalues becomes zero. This is because the problem, being convex, attains its optimum on the boundary of $\mathcal{S}_+^q$, and this boundary is the set of positive semidefinite matrices with at least one zero eigenvalue.

While there do not exist any ways to avoid this issue, an ongoing area of research is whether more stable dual SDP programs can be obtained by representing the polynomials in the SOS program by a different basis than a polynomial basis. Some progress has been made in this area by using for instance Chebyshev polynomials for one-dimensional SOS programs or an interpolant basis in general [PY19].

1.1.1 Organization

This thesis is organized as follows: First, the existing methods for bounding time averages of scalar functions along trajectories of continuous-time systems are explained in chapter 2. We review how optimization problems whose solutions are bounds on the time average can be derived and under what conditions the solutions are guaranteed to give sharp bounds on the maximal time average. We then recall how these optimization problems can be turned into tractable SOS programs and under what conditions these SOS programs must give sharp bounds on time-averages.

We then review the concept of LEs and their generalizations in section 3.1 and attempt to remove any ambiguity about the notion of LE which we treat with our optimization techniques by showing that the two most frequently encountered definition of LEs are equal. The rest of chapter 3 contains the theory of bounding LEs of continuous-time systems using convex optimization. Convex optimization problems, including SOS programs, whose solutions give bounds on the maximal LE, are derived and conditions for sharpness guarantees are given. Section 3.4 discusses how the SOS programs for bounding the maximal LE can be implemented in practice.

In chapter 4 we show how symmetries can be leveraged to simplify the computations and derive more general optimization problems for bounding LEs. We then apply our approach to several explicit examples in chapter 5, including to the Hénon–Heiles system in section 5.4 where we use SOS programming in conjunction with classical methods to compute the system's maximal LE and the periodic orbits on which it is attained.

In chapter 6 we derive optimization problems for bounding arithmetic means as well as geometric means. We then explain why geometric means are a more natural notion of time-average when trying to obtain computationally tractable SOS formulations for the problem of bounding the maximal LE. We therefore develop a theory for bounding geometric means of scalar functions along trajectories of discrete maps by posing the question as a convex optimization problem and prove that under analogous conditions to those in the continuous-time case, these optimization problems will have solutions that equal the maximal geometric mean. We then show how the optimization problems can be solved in practice by turning them into a sequence of SOS feasibility problems, whose solutions are shown to converge to the maximal geometric mean under certain easily-realized conditions.

Chapter 7 applies the methods of the section before to the problem of bounding the maximal LE of a discrete map. We show how the LE of a discrete-time system can be expressed as the geometric mean of a suitable scalar quantity on an enlarged phase space and how the techniques of the previous section can be applied to derive optimization problems and SOS programs whose solutions give sharp bounds on the maximal LE under the conditions explained in those sections. Since some care must be taken in how one formulates the optimization problems for bounding LEs, so that the constraints can be turned into polynomial SOS constraints, we also explain how the constraints and the SOS programs can be implemented in practice.

We then derive symmetry results in chapter 8, which are analogous to those we derived for continuous-time systems and which help to improve the computations that need to be performed to bound the maximal LE with the SOS methods from chapter 7. Several examples are then shown which use our method for bounding LEs, some of which we solve analytically to prove exact bounds on the maximal LE of these systems.

In chapter 10 we outline an extension of the problem of bounding the maximal LE of a continuous-time system: the problem of bounding the sum of the largest k LEs. We argue why this problem is of interest and show how it too can be turned into a convex optimization problem—at least in some special cases. For these special cases we derive explicit SOS programs whose solutions give bounds on the maximal sum of LEs. We also explain why the results in the earlier sections imply that these SOS programs also enjoy the same sharpness guarantees, and the possibility to exploit symmetries, as the optimization problems for bounding the maximal LE. Finally, we apply the methods for bounding the leading sum of LEs to some examples in chapter 11.

1.1.2 Contributions

The contributions in this thesis are based on previous work for bounding time averages using convex optimization, which is explained in chapter 2. The contributions include the formulation of convex optimization problems and SOS programs for bounding the maximal LE of a continuous-time system. We give sharpness guarantees for the optimization problems for bounding LEs, including Propositions 3 and 4. We extend previous results on symmetries to more general groups of symmetries in Propositions 18 and 19 and apply these in chapter 4 where we show in Propositions 5 to 8, and in the discussions following these results, how symmetries of the ODE may be used to simplify the optimization problems for bounding LEs. We compute sharp bounds on the LEs of various systems. In some we confirm previous bounds with our new method and in some we compute these for the first time. In particular, for the Hénon–Heiles system in section 5.4 we compute its maximal LE on a compact domain with our method and with classical methods, and we find and compute the unstable periodic orbits on which the maximal LE occurs.

For discrete-time systems we develop an optimization and SOS-programming approach in chapter 6 for bounding the geometric mean of scalar quantities over bounded orbits of the system and prove various theoretical results, based on Mañé lemmas, which ensure that the optimization problems give sharp bounds on the maximal geometric mean in Propositions 9 to 11. We use these results to derive optimization problems for bounding the maximal LE of discrete-time maps by formulating the problem as a maximization of a geometric mean in chapter 7 and we provide two results in Propositions 12 and 13 which guarantee that the optimization formulations give sharp bounds on the maximal LE. Analogous results about the use of symmetries like those for continuous systems are proved for discrete maps in Propositions 20 and 21 and in Propositions 14 to 17. We work out various examples using these methods and compute bounds on the maximal LE of several maps in chapter 9, including analytic bounds for the Hénon map.

Finally, we explore how the maximal sum of the leading LEs of a system may be bounded by using methods that are based on those which we develop in the previous sections. We derive some optimization problems that bound this sum in a few special cases, including when one wants to bound the sum of all LEs or the sum of $(n-1)$ LEs and the dimension of the phase space is three or four. In sections 10.1.1 and 10.1.4 we also show that these formulations give sharp bounds and can be simplified with symmetries, by explaining how they can be treated with methods that have either been developed elsewhere or which were developed by us in the earlier parts of the thesis. We then compute the sum of the leading LEs for various examples which often possess features which would make computing this sum, and concluding that it is indeed maximizing, more difficult using classical methods.

Chapter 2

Bounding time averages in ODEs with convex optimization

In this section we explain how time averages of scalar quantities $\Phi(x)$ can be extremized over bounded trajectories of (1.0.1) by posing the maximization of the time average as a convex optimization problem. Section 2.1 explains how to derive optimization problems whose solution is guaranteed to be an upper bound on the time average of $\Phi(x)$ and section 2.2 proves a duality result under whose conditions the solution to the optimization problem is guaranteed to be equal to the maximum of the time average of $\Phi(x)$.

In section 2.3 we show how the convex optimization problems whose solutions give a bound on the time average of $\Phi(x)$ can be relaxed to convex SOS programs and how these SOS programs may be formulated to give bounds on the time average on semialgebraic domains by making use of the s-procedure. In section 2.4 we explain a result that guarantees that the solution of the SOS program also equals the maximum of the time average of $\Phi(x)$.

The results in sections 2.1 and 2.2 have been proved in [TGD18] and the duality results for the SOS programs have been proved in [Lak+20]. We show them here because they lay the foundation for the majority of the results in this thesis and to provide some of the details implicit in these sources.

2.1 Derivation of an explicit optimization formulation for bounding time averages

Consider any continuous function $\Phi : \mathbb{R}^n \to \mathbb{R}$ and its infinite-time average along a trajectory $x(t)$ of the ODE (1.0.1),

$$\overline{\Phi}(x_0) = \limsup_{T\to\infty} \frac{1}{T}\int_0^T \Phi(x(t))\,\mathrm{d}t. \tag{2.1.1}$$

We define time averages with lim sup to ensure they exist. If the limit exists, the resulting bounds are unchanged if $\overline{\Phi}$ is defined using lim inf instead. If the trajectory emanating from x_0 is bounded uniformly for $t \geq 0$, then so is $\Phi(x(t))$ for continuous $\Phi(x)$, and thus $\overline{\Phi}$ is finite.

The strategy for bounding $\overline{\Phi}$ without knowing any particular trajectories is to introduce a differentiable *auxiliary function* $V : \mathbb{R}^n \to \mathbb{R}$. For any such V, the continuous function $f \cdot \nabla V : \mathbb{R}^n \to \mathbb{R}$ time-averages to zero along any bounded trajectory. This is true because along trajectories $\frac{\mathrm{d}}{\mathrm{d}t}V(x(t)) = f(x(t)) \cdot \nabla V(x(t))$, in other words $f \cdot \nabla V$ is the Lie derivative of V along ODE solutions. Since $V(x(t))$ is uniformly bounded forward in time along each bounded trajectory, $\overline{f \cdot \nabla V} = \limsup_{T\to\infty} \frac{1}{T}[V(x(T)) - V(x_0)] = 0$. For each V,

$$\overline{\Phi} = \overline{\Phi + f \cdot \nabla V} \leq \sup_{x\in\mathcal{B}} \left[\Phi(x) + f(x) \cdot \nabla V(x)\right], \tag{2.1.2}$$

where $\mathcal{B}$ can be any set (such as $\mathbb{R}^n$) in which the trajectory starting at x_0 eventually remains. For a poor choice of V the right-hand side of (2.1.2) is infinite, which for many Φ is true when $V = 0$, but one can seek V that makes the right-hand side not only finite but as small as possible.

The upper bound (2.1.2) holds for all bounded trajectories and all $V \in \mathcal{C}^1$, so one can maximize the left-hand side over initial conditions x_0 and minimize the right-hand side over V. Let $\mathcal{B}_0$ be a set of initial conditions that each lead to bounded trajectories, and suppose all of these trajectories eventually remain in some set $\mathcal{B}$. The corresponding maximization and minimization of (2.1.2) gives

$$\sup_{x_0 \in \mathcal{B}_0} \overline{\Phi} \le \inf_{V \in \mathcal{C}^1(\mathcal{B})} \sup_{x \in \mathcal{B}} \left[\Phi(x) + f(x) \cdot \nabla V(x)\right]. \tag{2.1.3}$$

One can always choose $\mathcal{B} = \mathbb{R}^n$, but often it is possible to choose smaller $\mathcal{B}$, which may give a better upper bound in (2.1.3). If $\mathcal{B}_0$ lies in the basin of a local attractor, $\mathcal{B}$ can be any set that contains the attractor. If $\mathcal{B}_0$ is a trapping set—i.e., $\mathcal{B}_0$ is forward invariant—one can choose $\mathcal{B} = \mathcal{B}_0$.

2.2 Sharpness guarantees for optimization formulation for bounding time averages

We say that the bound (2.1.3) is *sharp*, if the inequality becomes an equality. Ideally, we would like to find a specific $V \in \mathcal{C}^1(\mathcal{B})$ for which the equality can be realized. However, this may not be possible for a specific $V \in \mathcal{C}^1(\mathcal{B})$.

While it is not known if the inequality (2.1.3) can be strengthened to an equality for any domain $\mathcal{B}$, for certain common domains there exists a positive result. Suppose that $\mathcal{B}$ is a trapping set, so that (2.1.3) holds with $\mathcal{B} = \mathcal{B}_0$, and also that $\mathcal{B}$ is compact. By the well-posedness of (1.0.1) we assume throughout, the mapping $x_0 \mapsto x(t)$ is continuously differentiable for all $x_0 \in \mathcal{B}$. Under these conditions, it has been proved that [TGD18]

$$\sup_{x_0 \in \mathcal{B}} \overline{\Phi} = \inf_{V \in \mathcal{C}^1(\mathcal{B})} \sup_{x \in \mathcal{B}} \left[\Phi(x) + f(x) \cdot \nabla V(x) \right]. \tag{2.2.1}$$

This equality can be seen as a statement of strong convex duality: the left-hand side can be restated with the inner maximization over invariant measures, and the right-hand side can be stated in the same way with the order of minimization and maximization reversed [TGD18]. The result in [TGD18] has similar precedents in the context of discrete-time dynamics—see discussion and references in [Boc17; Jen19]. It is assumed in [TGD18] that $\mathcal{B}$ is a full-dimensional region in $\mathbb{R}^n$, but the proof there applies also if $\mathcal{B}$ is a lower-dimensional manifold embedded in $\mathbb{R}^n$, provided that $V \in \mathcal{C}^1(\mathcal{B})$ means that V has a continuously differentiable extension to an open set containing $\mathcal{B}$. In what follows we will use this fact, applying (2.2.1) for lower-dimensional domains. On unbounded domains, including the important case where $\mathcal{B}_0 = \mathcal{B} = \mathbb{R}^n$, it remains to be characterized when (2.1.3) is an equality or a strict inequality.

Although the right-hand infimum in (2.2.1) may not be attained, there exist V giving arbitrarily sharp upper bounds on the left-hand supremum (which is attained [TGD18]). However, it can be very difficult to evaluate the inner supremum on the right-hand side for a given V, and optimizing over V is still harder. An important exception, in which computationally tractable relaxations of the right-hand minimax problem are possible, is the case where the ODE right-hand side $f(x)$ and the quantity $\Phi(x)$ are polynomial.

2.2.1 Outline of the proof of strong duality

We will now briefly outline the proof of (2.2.1) from [TGD18]. Let $Pr(\mathcal{B})$ be the space of Borel probability measures on $\mathcal{B}$. A measure $m \in Pr(\mathcal{B})$ is said to be invariant with respect to the flow $\varphi_t(x)$ of (1.0.1), if $m\Big(\varphi_t^{-1}(A)\Big) = m(A)$ for all Borel sets $A \subset \mathcal{B}$ and $t \in \mathbb{R}$.

A probability measure m^* is said to be *ergodic* with respect to the flow $\varphi_t(x)$ if the only Borel sets A for which $\varphi_t^{-1}(A) = A$ either satisfy $m^*(A) = 1$ or $m^*(A) = 0$.

The result (2.2.1) holds if the following equalities can be shown to hold:

$$\sup_{x_0 \in \mathcal{B}} \overline{\Phi} = \max_{\substack{m \in Pr(\mathcal{B}) \\ m \text{ is invariant}}} \int \Phi \, dm \tag{2.2.2a}$$

$$= \sup_{m \in Pr(\mathcal{B})} \inf_{V \in \mathcal{C}^1(\mathcal{B})} \int \Phi + f \cdot \nabla V \, dm \tag{2.2.2b}$$

$$= \inf_{V \in \mathcal{C}^1(\mathcal{B})} \sup_{m \in Pr(\mathcal{B})} \int \Phi + f \cdot \nabla V \, dm \tag{2.2.2c}$$

$$= \inf_{V \in \mathcal{C}^1(\mathcal{B})} \max_{x \in \mathcal{B}} \{\Phi + f \cdot \nabla V\}. \tag{2.2.2d}$$

Lemma 1. *Let $\Phi(x) \in \mathcal{C}(\mathcal{B})$ with $\mathcal{B} \subset \mathbb{R}^n$ compact and trapping. For each $x \in \mathcal{B}$, there exists an invariant probability measure m such that*

$$\overline{\Phi}(x) := \limsup_{T \to \infty} \frac{1}{T} \int_0^T \Phi(\varphi_t(x)) \, dt = \int \Phi(x) \, dm \tag{2.2.3}$$

Proof. We will first construct a sequence of probability measures that contains a subsequence which converges weakly to an invariant measure with the above property.

By definition of the lim sup there exists a sequence T_k such that

$$\overline{\Phi}(x) := \lim_{k\to\infty} \frac{1}{T_k} \int_0^{T_k} \Phi(\varphi_t(x))\, dt \tag{2.2.4}$$

Let m_0 be the Dirac distribution on $\mathcal{B}$ centred at $x \in \mathcal{B}$ then define the measures m_k as

$$\begin{aligned} \int \Phi(x)\, dm_k :=& \frac{1}{T_k} \int_0^{T_k} \int \Phi(\varphi_t(x))\, dm_0\, dt, \\ =& \frac{1}{T_k} \int_0^{T_k} \Phi(\varphi_t(x))\, dt. \end{aligned} \tag{2.2.5}$$

Furthermore, each m_k is a probability measure, since taking $\Phi = 1$ gives $m_k(\mathcal{B}) = \int 1\, dm_k = 1$. Since $\mathcal{B}$ is compact, the Prokhorov theorem (Theorem 8.6.2 in [BR07]) guarantees the existence of a subsequence m_{k_j} of m_k which converges weakly to a measure m. Therefore, upon relabelling the indices k_j as k we get

$$\int \Phi(x)\, dm = \lim_{k\to\infty} \frac{1}{T_k} \int_0^{T_k} \Phi(\varphi_t(x))\, dt = \overline{\Phi}(x). \tag{2.2.6}$$

To see that m is invariant, note that if $w = t + s$ then

$$\begin{aligned} \int \Phi(\varphi_s(x))\, dm &= \lim_{k\to\infty} \frac{1}{T_k} \int_0^{T_k} \Phi(\varphi_{t+s}(x))\, dt \\ &= \lim_{k\to\infty} \frac{1}{T_k} \int_s^{T_k+s} \Phi(\varphi_w(x))\, dw \\ &= \lim_{k\to\infty} \frac{1}{T_k} \left(\int_0^{T_k} \Phi(\varphi_w(x))\, dw - \int_0^{s} \Phi(\varphi_w(x))\, dw + \int_{T_k}^{T_k+s} \Phi(\varphi_w(x))\, dw \right) \\ &= \lim_{k\to\infty} \frac{1}{T_k} \int_0^{T_k} \Phi(\varphi_w(x))\, dw \\ &= \int \Phi(x)\, dm. \end{aligned} \tag{2.2.7}$$

Finally, we use the fact that m is invariant with respect to the flow if and only if

$\int \Phi(\varphi_s(x))\, dm = \int \Phi(x)\, dm$ for all $\Phi \in \mathcal{C}(\mathcal{B})$ and for all $s \in \mathbb{R}$. $\square$

A similar result to Lemma 1 is known as the Kryloff–Bogoliouboff Theorem [KB37] and Theorem 4.1 in [EW13]. From Lemma 1 it follows that

$$\sup_{x_0 \in \mathcal{B}} \overline{\Phi} \leq \sup_{\substack{m \,\in\, Pr(\mathcal{B}) \\ m \text{ is invariant},}} \int \Phi\, dm. \tag{2.2.8}$$

Recall that the set of invariant probability measures on $\mathcal{B}$ is non-empty, convex, weak* compact and its extreme points are ergodic measures (Theorem 4.4 in [EW13]). The right hand side of (2.2.8) is the maximization of a continuous linear functional over the set of invariant probability measures. Therefore, its maximum is achieved on an extreme point of the set [Lax02]. In other words, its maximum is achieved for an ergodic measure m^* so that

$$\sup_{x_0 \in \mathcal{B}} \overline{\Phi} \leq \sup_{\substack{m \,\in\, Pr(\mathcal{B}) \\ m \text{ is invariant},}} \int \Phi\, dm = \int \Phi\, dm^*. \tag{2.2.9}$$

We therefore recall the famous Birkhoff ergodic theorem for continuous time (see e.g. [BLM12] or [Wal00]).

Theorem 1. *(Birkhoff) Suppose that $\mathcal{B} \subset \mathbb{R}^n$ is compact and suppose that the probability measure m^* is ergodic with respect to the flow $\varphi_t(x)$ and $\Phi(x) \in \mathbf{L}^1(\mathcal{B}, m^*)$. Then for m^* a.e. $x \in \mathcal{B}$ the following limit exists and satisfies*

$$\overline{\Phi}(x) = \limsup_{T \to \infty} \frac{1}{T} \int_0^T \Phi\big(\varphi_t(x)\big)\, dx = \int_{\mathcal{B}} \Phi\, dm^*. \tag{2.2.10}$$

Applying the Birkhoff ergodic theorem to (2.2.8) implies that the inequality is an equality for m^* a.e. $x_0 \in \mathcal{B}$. Therefore, the supremum over all $x_0 \in \mathcal{B}$ is also equal to the right hand side of (2.2.8). This proves line (2.2.2a).

Lines (2.2.2b) and (2.2.2c) are not trivial to prove. However they are proved in great detail in [TGD18] and we will refer the reader to this source for the proof. Essentially, (2.2.2b) follows from the Lagrangian notion of invariance of the measure m which holds if and only if $\int f \cdot \nabla V \; dm = 0$ and (2.2.2c) can be proved using general minimax theorems like the one in [Sio58].

The last line (2.2.2d) is proved by noting that

$$\sup_{m \in Pr(\mathcal{B})} \int \Phi + f \cdot \nabla V \; dm \leq \max_{x \in \mathcal{B}} \{\Phi + f \cdot \nabla V\}. \tag{2.2.11}$$

Furthermore, maximizing only over Dirac distributions gives

$$\sup_{m \in Pr(\mathcal{B})} \int \Phi + f \cdot \nabla V \; dm \geq \sup_{m = \text{ Dirac}} \int \Phi + f \cdot \nabla V \; dm = \max_{x \in \mathcal{B}} \{\Phi + f \cdot \nabla V\}. \tag{2.2.12}$$

Combining these two inequalities and minimizing over $V \in \mathcal{C}^1(\mathcal{B})$ gives (2.2.2d).

2.3 SOS relaxation of polynomial inequalities

The minimax problem on the right-hand side of (2.1.3) and (2.2.1) can be written equivalently as a minimization subject to a pointwise inequality constraint, in which case (2.1.3) becomes

$$\sup_{x_0 \in \mathcal{B}_0} \overline{\Phi} \leq \inf_{V \in \mathcal{C}^1(\mathcal{B})} B \quad \text{s.t.} \quad \underbrace{B - \Phi(x) - f(x) \cdot \nabla V(x)}_{S(x)} \geq 0 \; \forall x \in \mathcal{B}. \tag{2.3.1}$$

In the case where the given $f(x)$ and $\Phi(x)$ are polynomial in the components of x, and $V(x)$ is also chosen to be polynomial, $S(x)$ defined as in (2.3.1) is a polynomial. Even for polynomial $S(x)$ verifying the right-hand side of (2.3.1) is often intractable; in general, deciding non-negativity of a multivariate polynomial on a given set has NP-hard computational complexity [MK87]. However, there is a standard way to relax the right-hand problem by strengthening the $S(x) \geq 0$ constraint to a SOS constraint, meaning that $S(x)$ is required to be representable as the sum of polynomials that are multiplied by themselves as is explained in section 1.1. We first do the SOS relaxation for the simple case where $\mathcal{B} = \mathbb{R}^n$, followed by more general $\mathcal{B}$ in section 2.3.3.

2.3.1 SOS conditions for $\mathbb{R}^n$

When $\mathcal{B} = \mathbb{R}^n$ on the right-hand side of (2.3.1), $S(x)$ must be non-negative for all $x \in \mathbb{R}^n$. Non-negativity can be enforced by requiring $S(x)$ to admit an SOS representation, as in (1.1.1). In other words, S belongs to the set Σ_n of n-variate SOS polynomials, which is strictly smaller than the set of non-negative polynomials when $n \geq 2$ (for instance, Motzkin's polynomial $x^4y^2 + x^2y^4 - 3x^2y^2 + 1$ is nonnegative on $\mathbb{R}^2$ but it does not possess a SOS decomposition [Mot67]). Then the SOS relaxation of (2.3.1) when $\mathcal{B} = \mathbb{R}^n$ is

$$\sup_{x_0 \in \mathcal{B}_0} \overline{\Phi} \leq \inf_{V \in \mathbb{R}[x]} B \quad \text{s.t.} \quad S \in \Sigma_n, \tag{2.3.2}$$

where $\mathbb{R}[x]$ denotes the set of real polynomials in x. Restricting V to any finite-dimensional subspace of $\mathbb{R}[x]$ gives a finite-dimensional convex optimization problem. In particular, the right-hand side of (2.3.2) will be an SOS program [Par13], as described in detail in section 1.1. In (2.3.2) the tunable parameters, which are denoted by c_j in section 1.1, are B and the coefficients of V (in any chosen basis). S as defined in (2.3.1) is indeed affine in these parameters.

The right-hand problem (2.3.2), with V restricted to a finite subspace, is computationally tractable when the dimension of this subspace and the ODE dimension n are not too large. Solving such SOS programs with V in increasingly large subspaces gives a nonincreasing sequence of infima that are all upper bounds on the left-hand side of (2.3.2). The inequality (2.3.2) is sometimes strict since it holds for any initial set $\mathcal{B}_0$ whose trajectories are bounded forward in time, and different $\mathcal{B}_0$ can give different values. However, for the initial set $\mathcal{B}_0 = \mathbb{R}^n$ in various particular ODEs whose trajectories are bounded forward in time, SOS computations suggest that (2.3.2) is an equality in those cases [Gol18; GF19; Lak+20]. The sharpness of (2.3.2) is an open question. It is not known when the original inequality (2.1.3) is sharp for $\mathcal{B}_0 = \mathcal{B} = \mathbb{R}^n$, nor when the right-hand side of (2.1.3) is equal to its SOS relaxation in (2.3.2). For various particular ODEs, however, solving the SOS program in (2.3.2) with V restricted to increasingly large subspaces has given numerical upper bounds that indeed seem to converge to the maximum of $\overline{\Phi}$ over all trajectories, and we are not aware of counterexamples to such convergence.

2.3.2 Archimedean sets and the s-procedure

The inequality in (2.3.2) can be strengthened to an equality if the domain of the optimization is restricted to a suitable set. This is similar to strong duality (2.2.1), with some practical differences: the domain $\mathcal{B}$ is required to be semialgebraic, Archimedean as well as being compact, and a procedure for constraining the domain needs to be employed. Since SOS constraints are necessarily global, a special technique, called the s-procedure [TP08] is needed to enforce local non-negativity constraints via a global SOS constraint. We will explain both the concept of Archimedean semialgebraic sets and the s-procedure in the following paragraphs.

A set $\mathcal{B}$ is said to be semialgebraic if it can be specified by a finite number of polynomial inequalities and/or equalities:

$$\mathcal{B} = \{x \in \mathbb{R}^n : h_i(x) = 0 \text{ for } i = 1, \ldots, I,\ g_j(x) \geq 0 \text{ for } j = 1, \ldots, J\}\,. \tag{2.3.3}$$

It is clear from the above definition that $\mathcal{B}$ is always closed, so compactness can be enforced by including an inequality $g_0(x) = R^2 - |x|^2 \geq 0$ in the specification of $\mathcal{B}$ for a sufficiently large scalar R. Furthermore, we could also define $\mathcal{B}$ only via inequalities and omit the equalities, since requiring that $h_i = 0$ is equivalent to requiring that both $h_i \geq 0$ and $h_i \leq 0$ simultaneously. The Archimedean property is a technical condition on $\mathcal{B}$ which is needed to prove the strong duality result for SOS programs.

Consider the set

$$\mathbb{M} = \left\{ \sigma_0 + \sum_{j=1}^{J} \sigma_j g_j + \sum_{i=1}^{I} \tilde{\sigma}_i h_i - \sum_{i=1}^{I} \tilde{\sigma}_i' h_i \mid \sigma_j, \tilde{\sigma}_i, \tilde{\sigma}_i' \in \Sigma_n \right\} \tag{2.3.4}$$

of weighted SOS polynomials, where the weights are exactly the polynomials comprising the semialgebraic definition of the domain $\mathcal{B}$. The set $\mathbb{M}$ is called the *quadratic module* of the polynomials defining $\mathcal{B}$. Note that for any real polynomial $\rho \in \mathbb{R}[x]$, we may write $\rho = \tilde{\sigma} - \tilde{\sigma}'$ where $\tilde{\sigma} = (\rho + 1/4)^2 \in \Sigma_n$ and $\tilde{\sigma}' = (\rho - 1/4)^2 \in \Sigma_n$, so any real polynomial may be written as the difference of two SOS polynomials. Therefore, we may equivalently express the quadratic module of $\mathcal{B}$ as

$$\mathbb{M} := \left\{ \sigma_0 + \sum_{j=1}^{J} \sigma_j g_j + \sum_{i=1}^{I} \rho_i h_i \mid \sigma_j \in \Sigma_n,\ \rho_i \in \mathbb{R}[x] \right\}. \tag{2.3.5}$$

The set $\mathcal{B}$ is said to satisfy the Archimedean property (Definition 3.4 in [Las+08]) if

$$N - |x|^2 \in \mathbb{M}, \tag{2.3.6}$$

for some $N \in \mathbb{N}$ [NS07]. Note that this is not at all difficult to satisfy for our purposes, because we may simply add $g_0 := N - |x|^2$ to any compact semialgebraic set $\mathcal{B}$, for large enough $N \in \mathbb{N}$ without changing $\mathcal{B}$, after which satisfying the Archimedean property (2.3.6) becomes trivial.

The second, related, concept that we need is the *s-procedure*, which is a way of using SOS constraints to enforce non-negativity constraints on semialgebraic sets like $\mathcal{B}$. In other words, it is a way of formulating a local pointwise constraint as a global SOS constraint: suppose that we want $S(x) \geq 0$ on $\mathcal{B}$, but we don't care about the sign of S outside of $\mathcal{B}$. One way to say this is to require that

$$\left.\begin{aligned} S(x) - \textstyle\sum_{i=1}^{I} \rho_i(x) h_i(x) - \sum_{j=1}^{J} \sigma_j(x) g_j(x) &\geq 0 \\ \sigma_j(x) &\geq 0 \end{aligned}\right\} \quad \text{for all } x \in \mathbb{R}^n, \tag{2.3.7}$$

because then

$$S(x) \geq \textstyle\sum_{i=1}^{I} \rho_i(x) h_i(x) + \sum_{j=1}^{J} \sigma_j(x) g_j(x) \geq 0 \quad \text{for all } x \in \mathcal{B}, \tag{2.3.8}$$

by the definition of $\mathcal{B}$. In order to make the constraint (2.3.7) computationally tractable we relax it to an SOS constraint:

$$\begin{aligned} S - \textstyle\sum_{i=1}^{I} \rho_i h_i - \sum_{j=1}^{J} \sigma_j g_j &\in \Sigma_n \\ \sigma_j &\in \Sigma_n \end{aligned} \tag{2.3.9}$$

Since $S(x) - \sum_{i=1}^{I} \rho_i(x) h_i(x) - \sum_{j=1}^{J} \sigma_j(x) g_j(x)$ is implied to be non-negative for all $x \in \mathbb{R}^n$ by this condition, it also implies the pointwise non-negativity of $S(x)$ on $\mathcal{B}$. Transforming a local non-negativity constraint into a global SOS constraint by adding the semialgebraic specification of $\mathcal{B}$ into the constraints in the above fashion is called the s-procedure. Note however, that condition (2.3.7) is equivalent to requiring that

S belongs to the quadratic module $\mathbb{M}$ of the polynomials specifying $\mathcal{B}$. Therefore the s-procedure may be summarized as:

$$S \in \mathbb{M} \implies S(x) \geq 0, \text{ for all } x \in \mathcal{B}. \tag{2.3.10}$$

2.3.3 SOS conditions for $\mathcal{B}$

Relaxing the right-hand minimization in (2.3.1) by strengthening its constraint to (2.3.9) gives an immediate upper bound on $\sup_{x_0 \in \mathcal{B}_0} \overline{\Phi}$. In fact this relaxation does not change the value of the infimum in (2.3.1), provided that the semialgebraic set $\mathcal{B}$ is compact and satisfies the Archimedean property (2.3.6). The fact that the right-hand infimum in (2.3.1) is unchanged by the SOS relaxation (2.3.9) when the semialgebraic set $\mathcal{B}$ is compact and Archimedean, follows directly from a theorem of real algebraic geometry called Putinar's Positivstellensatz (Lemma 2).

When $\mathcal{B}$ is a trapping set, in addition to being an Archimedean (and thus compact and semialgebraic set), convergence of SOS computations to $\max_{\mathcal{B}} \overline{\Phi}$ is guaranteed. As observed in [Lak+20], the strong duality result (2.2.1) remains an equality when the right-hand minimization is relaxed by replacing its constraint by the SOS conditions (2.3.9), as guaranteed by Putinar's Positivstellensatz. That is,

$$\begin{aligned} \sup_{x_0 \in \mathcal{B}} \overline{\Phi} = \inf_{V, \rho_i, \sigma_j \in \mathbb{R}[x]} B \quad \text{s.t.} \quad & S - \textstyle\sum_{i=1}^{I} \rho_i h_i - \sum_{j=1}^{J} \sigma_j g_j \in \Sigma_n \\ & \sigma_j \in \Sigma_n, \quad j = 1, \ldots, J, \end{aligned} \tag{2.3.11}$$

where again S depends affinely on B and the coefficients of V as defined in (2.3.1). Choosing finite polynomial spaces in which to seek V, ρ_i, and σ_j gives an SOS program with a computable infimum. As the polynomial spaces are enlarged, the infima of these increasingly expensive computations converge from above to the left-hand side of (2.3.11).

2.4 Sharpness guarantees for SOS programs for bounding time averages

For completeness, we briefly outline the proof of the strong duality result that ensures the equality in (2.3.11) between the supremum of the time average of Φ and the solution to the SOS program on the right hand side. This result is of central importance to the SOS programs that we derive and use in the present work, as it ensures that their solutions do indeed converge to the maximal time average of the scalar quantity Φ. The proof follows the approach in [Lak+20], and uses two key ingredients: Firstly, it uses the strong duality result for the pointwise constrained optimization problem (2.2.1) and the Positivstellensatz (Lemma 4.1 in [Put93]).

Lemma 2. *(Putinar [Put93], Lemma 4.1) Let $\mathcal{B}$ be a semialgebraic set as in* (2.3.3)*. Assume that $\mathcal{B}$ is Archimedean (it satisfies* (2.3.6)*). For all $S \in \mathbb{R}[x]$ one has*

$$S(x) > 0, \text{ for every } x \in \mathcal{B} \implies S \in \mathbb{M}. \tag{2.4.1}$$

We will continue to assume that f and Φ are polynomial in x. First note that since $\mathbb{R}[x]$ is dense in $\mathcal{C}^1(\mathcal{B})$ (Theorem 1.1.2 in [Lla86]), we may restrict V in (2.2.1) to $\mathbb{R}[x]$ while preserving the equality:

$$\sup_{x_0 \in \mathcal{B}} \overline{\Phi} = \inf_{V \in \mathbb{R}[x]} \sup_{x \in \mathcal{B}} \left[\Phi(x) + f(x) \cdot \nabla V(x)\right]. \tag{2.4.2}$$

In particular, by strong duality (2.2.1), if we let $k \in \mathbb{N}$ there exists a sequence B_k and a sequence $V_k \in \mathbb{R}[x]$ satisfying $\lim_{k\to\infty} B_k = \sup_{x_0\in\mathcal{B}} \overline{\Phi}$ and

$$B_k - \Phi(x) - f(x)\cdot\nabla V_k(x) \geq 0, \qquad \text{for all } x \in \mathcal{B}, \tag{2.4.3}$$

for every $k \in \mathbb{N}$. Furthermore let $\epsilon_k > 0$ be any sequence with $\epsilon_k \downarrow 0$ as $k \to \infty$ then

$$\underbrace{B_k + \epsilon_k - \Phi(x) - f(x)\cdot\nabla V_k(x)}_{=:S_k(x)} > 0, \qquad \text{for all } x \in \mathcal{B}. \tag{2.4.4}$$

Therefore, since $S_k(x) > 0$ on $\mathcal{B}$, by Putinar's Positivstellensatz in Lemma 2, we have $S_k \in \mathbb{M}$. But for any k this is just the constraint in (2.3.11) for the particular choice of $B = B_k + \epsilon_k$ and $V = V_k$. Therefore, since $B_k + \epsilon_k \to \sup_{x_0\in\mathcal{B}} \overline{\Phi}$ as $k \to \infty$, the equality in (2.3.11) is proved.

Chapter 3

Bounding LEs of ODEs

We now turn to the problem of bounding LEs of systems of the form (1.0.1). In section 3.1 we expand on the definition of the LE in (1.0.3) by explaining LEs in more detail. We compare two commonly encountered definitions of LE and show that they, in fact, lead to identical values of the LE. This is done to remove any ambiguity about the generality of the optimization-derived bounds we later compute. The concept of a Floquet exponent is also explained in detail and we review how this idea is related to that of a Poincaré map. In section 3.2 we express the LE of an orbit of (1.0.1) as a time average of the form (2.1.1), so that in section 3.3 we can then derive an optimization problem whose solution is guaranteed to converge to the maximal LE (1.0.7) under some assumptions which were explained in the previous section. Finally, we show in section 3.4 how to relax the optimization problem for bounding LEs to computationally tractable SOS programs and explain when its solution is guaranteed to converge to the maximal LE.

3.1 Background on LEs of continuous systems

In the literature on Lyapunov exponents one often encounters slight variations on the definition of LEs [BV17; PP16; Bar05; BF06]. While all definitions of LEs must give the same qualitative information about the stability of an orbit—a negative maximal LE indicates stability, while a positive LE indicates chaos—it is not a priori clear whether they are identical. A crucial question therefore is whether these differing definitions lead to the same numerical values of the LEs. In this section we will attempt to answer this question for the two most commonly encountered definitions of LEs, by showing that they do in fact lead to identical values of LEs for any given orbit. The following sections attempt to provide a partial answer to this question by giving the two most commonly encountered definitions of LEs and showing that they lead to identical numerical values.

In section 3.1.3 the important special case of an LE called a Floquet exponent is described—the LE on an orbit of finite period. Floquet exponents are needed in many practical applications of our optimization procedures for bounding LEs. This is because the maximizing orbits of time averages—and hence of LEs—tend to be periodic orbits, rather than for instance chaotic trajectories of infinite periods. While the results of the previous section apply to bounding the maximal LE, no matter what type of bounded orbit it is maximized on, the examples that we show all attain their maximal LE on finite period orbits. We therefore find ourselves computing upper bounds on the maximal Floquet exponent, specifically.

3.1.1 Two definitions of LEs

Both definitions make use of the concept of a *fundamental matrix* solution to the linearization of (1.0.1). This is the time-dependent matrix $\mathsf{H}(t)$, which evolves initial conditions y_0 to solutions of (1.0.2). In other words, it satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathsf{H} = Df(x)\mathsf{H}, \qquad \mathsf{H}(0) = \mathbf{I}, \tag{3.1.1}$$

and $y(t) = \mathsf{H}(t)y_0$.

The growth rate of small perturbations away from $x(t)$ is therefore related to the eigenvalues and to the singular values of the matrix H. In light of definition (1.0.3), we can therefore define LEs by averaging over the singular values or eigenvalues of H. The singular exponents of (1.0.1) are defined as

$$\nu_i(x_0) := \limsup_{t\to\infty} \frac{1}{t} \log |\mathrm{sv}_i(\mathsf{H}(t))|, \tag{3.1.2}$$

and the Lyapunov exponents as

$$\mu_i(x_0) := \limsup_{t\to\infty} \frac{1}{t} \log |\mathsf{H}(t)y_0|, \tag{3.1.3}$$

where $\mathrm{sv}_i(\mathsf{H}(t))$ denotes the i'th singular value of $\mathsf{H}(t)$. Another way to express the singular exponent is directly via the norm of $\mathsf{H}(t)$.

$$\nu_i(x_0) = \limsup_{t\to\infty} \frac{1}{t} \log |\mathrm{sv}_i(\mathsf{H}(t))| = \limsup_{t\to\infty} \frac{1}{2t} \log |\mathsf{H}(t)^{\mathsf{T}}\mathsf{H}(t)y_0|, \tag{3.1.4}$$

where y_0 is now chosen to be aligned with the i'th Lyapunov vector at $t = 0$. Note that for $i = 1$, (3.1.3) is the same as (1.0.3), because $\mathsf{H}(t)y_0 = y(t)$.

3.1.2 Equivalence of leading singular and Lyapunov exponent

In this section we will show that the definition of singular exponents (3.1.3) and the definition of Lyapunov exponents (3.1.4) lead to the same value of the largest LE. A similar result is proved in [Bar05; BF06], where it is shown that $\nu_i \le \mu_i$ for $i = 2, ..., n$ and equality holds for $i = 1$. Proposition 1 gives an alternative proof for that in [BF06], affirming that indeed $\nu_1 = \mu_1$. The proof in Proposition 1 relies on the following result from measure theory:

Lemma 3 (First Borel-Cantelli Lemma [Kle13])**.** *Let m be a probability measure (for instance the normalized Lebesgue measure on $\mathbb{S}^{n-1}$) on a set Ω. Let $\Upsilon_k \subset \Omega$ for $k = 1, 2, 3, \ldots$ and define $\Upsilon := \limsup_{k\to\infty} \Upsilon_k = \cap_{k=1}^{\infty} \cup_{l=k}^{\infty} \Upsilon_l$.*

If $\sum_{k=1}^{\infty} m(\Upsilon_k) < \infty$, then $m(\Upsilon) = 0$.

This lemma is needed to show that the leading LE is attained for almost any initial perturbation y_0 away from the trajectory $x(t)$.

Proposition 1. *Definitions* (3.1.3) *and* (3.1.4) *are equal when $i = 1$ and we have*

$$\mu(x_0, y_0) = \limsup_{t\to\infty} \frac{1}{t} \log |\mathsf{H}(t)y_0| = \limsup_{t\to\infty} \frac{1}{2t} \log |\mathsf{H}(t)^{\mathsf{T}}\mathsf{H}(t)y_0| = \nu(x_0, y_0), \quad (3.1.5)$$

for almost every $y_0 \in \mathbb{S}^{n-1}$ and initial condition $x_0 \in \mathbb{R}^n$ to (1.0.1)*.*

Proof. Let $\mathbb{S}^{n-1} = \{y \in \mathbb{R}^n | \; |y| = 1\}$. Since $\mathrm{sv}_1(\mathsf{H}(t)) := \sup_{y\in\mathbb{S}^{n-1}} |\mathsf{H}(t)y|$ for every t, we have that $|\mathsf{H}(t)y| \leq \mathrm{sv}_1(\mathsf{H}(t))$ for every t and any vector $y \in \mathbb{S}^{n-1}$. This immediately gives

$$\limsup_{t\to\infty} \frac{1}{t} \log |\mathsf{H}(t)y| \leq \limsup_{t\to\infty} \frac{1}{2t} \log |\mathsf{H}(t)^{\mathsf{T}}\mathsf{H}(t)y| \quad (3.1.6)$$

For the converse, let

$$\nu = \limsup_{t\to\infty} \sup_{y_0\in\mathbb{S}^{n-1}} \frac{1}{2t} \log |\mathsf{H}(t)^{\mathsf{T}}\mathsf{H}(t)y_0| = \limsup_{t\to\infty} \frac{1}{t} \log |\mathsf{H}(t)|, \quad (3.1.7)$$

be the largest LE obtained from (3.1.4) and let $t_k > k^2$ be a fast increasing sequence so that $|\mathsf{H}(t_k)| \geq \exp((\nu - 1/k)t_k)$. This exists since the $\limsup_{t\to\infty} \frac{1}{t} \log |\mathsf{H}(t)| = \nu$ exists, and so there exists a sequence $t_k \to \infty$ as $k \to \infty$ such that $\lim_{k\to\infty} \frac{1}{t_k} \log |\mathsf{H}(t_k)| = \nu$. Therefore choose a subsequence of t_k, which we call t_{k_p}, such that $\frac{1}{t_{k_p}} \log |\mathsf{H}(t_{k_p})| \geq$

$\nu - \frac{1}{p}$ and hence $|\mathsf{H}(t_{k_p})| \geq \exp((\nu - 1/p)t_{k_p})$ for all $p \in \mathbb{N}$. Relabel t_{k_p} as t_k and choose a subsequence from *it* such that $t_{k_l} \geq l^2$ for all $l \in \mathbb{N}$. Relabel t_{k_l} as t_k.

Let y_k be a singular vector for $\mathsf{H}(t_k)$ with $|\mathsf{H}(t_k)y_k| = |\mathsf{H}(t_k)| = \mathrm{sv}_1(\mathsf{H}(t_k))$. In particular, we have $\mathsf{H}(t_k)^T\mathsf{H}(t_k)y_k = |\mathsf{H}(t_k)|^2 y_k$.

Note that for any $y, y_k \in \mathbb{S}^{n-1}$ one has $-1 \leq y^\mathsf{T} y_k \leq 1$. Therefore, let

$$\Upsilon_k = \left\{ y \in \mathbb{S}^{n-1} \ : \ |y^\mathsf{T} y_k| \geq e^{-t_k/k} \right\}. \tag{3.1.8}$$

The quantity $\pi/2 - e^{-t_k/k}$ can be thought of as the maximum angle between y and y_k, up to a reflection. The sets Υ_k are introduced because we want to show that for any y belonging to infinitely many Υ_k, $\limsup_{t\to\infty} \frac{1}{t} \log |\mathsf{H}(t)y| \geq \limsup_{t\to\infty} \frac{1}{2t} \log |\mathsf{H}(t)^\mathsf{T}\mathsf{H}(t)y|$.

First note that for any $b \in \mathbb{R}^n$ and $\bar{y} \in \mathbb{S}^{n-1}$ one has $|b| \geq |b^\mathsf{T}\bar{y}|$. In particular, this is true for y_k as defined above and $b = \mathsf{H}(t_k)^T\mathsf{H}(t_k)y$.

Therefore one gets

$$\begin{aligned} |\mathsf{H}(t_k)||\mathsf{H}(t_k)y| &\geq |\mathsf{H}(t_k)^T\mathsf{H}(t_k)y| \\ &\geq |(\mathsf{H}(t_k)^\mathsf{T}\mathsf{H}(t_k)y)^\mathsf{T} y_k| \\ &= |y^\mathsf{T}\mathsf{H}(t_k)^\mathsf{T}\mathsf{H}(t_k)y_k| \\ &= |\mathsf{H}(t_k)|^2|y^\mathsf{T} y_k|. \end{aligned} \tag{3.1.9}$$

Therefore $|\mathsf{H}(t_k)y| \geq |\mathsf{H}(t_k)||y^\mathsf{T} y_k|$, which, by assumption, satisfies the inequality

$$\begin{aligned} |\mathsf{H}(t_k)y| &\geq |\mathsf{H}(t_k)||y^\mathsf{T} y_k| \\ &\geq e^{((\nu-1/k)t_k)} \left(e^{-t_k/k}\right) \\ &= e^{(\nu-2/k)t_k}. \end{aligned} \tag{3.1.10}$$

Now, let m be the normalized Lebesgue measure on $\mathbb{S}^{n-1}$. Then $m(\Upsilon_k^c) \leq K_n e^{-t_k/k}$

as $k \to \infty$, where K_n is a constant depending on the surface area of the n-dimensional unit sphere. For instance, in 2D we have

$$\begin{aligned} m(\Upsilon_k^c) &= 2\pi - 4\cos^{-1}\left(e^{-t_k/k}\right) \\ &= 2\pi - 4\left(\pi/2 - e^{-t_k/k} + e^{-3t_k/k}/6 - ...\right) \qquad \text{(Taylor)} \\ &= 4e^{-t_k/k} - 4e^{3t_k/k}/6 + ... \\ &\leq 4e^{-t_k/k} \\ &\leq 4e^{-k}. \end{aligned} \tag{3.1.11}$$

for large enough k.

This implies that $\sum_{k=1}^{\infty} m\left(\Upsilon_k^c\right)) = K_n \sum_{k=1}^{\infty} e^{-k} < \infty$.

Therefore, by the first Borel-Cantelli lemma (Lemma 3), one has that a.e. $y \in \mathbb{S}^{n-1}$ lies in all but finitely many Υ_k's. Therefore, for a.e. $y \in \mathbb{S}^{n-1}$ one has

$$\begin{aligned} \limsup_{t\to\infty} \frac{1}{t} \log |\mathsf{H}(t)y| &= \limsup_{k\to\infty} \frac{1}{t_k} \log |\mathsf{H}(t_k)y| \\ &\geq \nu - \frac{2}{k} \\ &= \limsup_{k\to\infty} \frac{1}{2t} \log |\mathsf{H}(t)^{\mathsf{T}}\mathsf{H}(t)y| - \frac{2}{k} \\ &= \limsup_{t\to\infty} \frac{1}{2t} \log |\mathsf{H}(t)^{\mathsf{T}}\mathsf{H}(t)y|, \end{aligned} \tag{3.1.12}$$

which shows the converse.

This, together with the initial direction, then gives the result that

$$\limsup_{t\to\infty} \frac{1}{t} \log |\mathsf{H}(t)y| = \limsup_{t\to\infty} \frac{1}{2t} \log |\mathsf{H}(t)^{\mathsf{T}}\mathsf{H}(t)y| \tag{3.1.13}$$

for a.e. $y \in \mathbb{S}^{n-1}$. □

3.1.3 Floquet exponents

In this section we discuss the theory of Floquet exponents and how they relate to Poincaré maps. While Floquet exponents are only defined on periodic orbits of a system (1.0.1), they are in fact exactly equal to the Lyapunov exponents (3.1.3) on such orbits. In the examples of our optimization formulations for bounding time averages, the maximal time averages are realized on orbits of finite period, which means that we are often bounding Floquet exponents when bounding the maximal LE of a dynamical system.

Definition 3. *Suppose* (1.0.1) *has a periodic solution* $x(t)$ *with period* T*. Let* $\mathsf{H}(t)$ *be the solution to the system* (3.1.1)*, which is the linearization of* (1.0.1) *linearized about* $x(t)$*. Then*

$$\mathsf{M} := \mathsf{H}(T) \tag{3.1.14}$$

is called the monodromy matrix*. Note that the monodromy matrix is independent of the initial condition on the orbit, since the orbit is periodic.*

The monodromy matrix is used in the definition of Floquet exponents. Floquet exponents are characteristic numbers associated with periodic orbits, which quantify the stability of the orbits.

Definition 4. *Let* M *be the monodromy matrix and let* $ev_i(\mathsf{M})$ *be the i'th eigenvalue of* M*. Then in analogy with the definition of LEs, the i'th Floquet exponent is defined as*

$$\tilde{\mu}_i := \frac{1}{T}\log\Big|\mathrm{Re}(\mathrm{ev}_i(\mathsf{M}))\Big|. \tag{3.1.15}$$

Note that since $x(t)$ is a periodic orbit, one will always have a Floquet exponent of value zero because $f(x(t))$ solves (1.0.2) and $\limsup_{t\to\infty} \frac{1}{t}\log(f(x(t))) = 0$. If one has a Poincaré map $P(\tilde{x}_0)$ (see e.g. [Str15] for a definition) for the periodic orbit, then the eigenvalues of $D_{\tilde{x}_0}P(\tilde{x}_0)$ will also correspond to the Floquet exponents. Here $\tilde{x}_0$ is the $(n-1)$-dimensional restriction of x_0 to a periodic orbit. The only difference is that one will not get the value zero Floquet exponent from the eigenvalues of $DP(x_0)$, it's dimension is $(n-1)\times(n-1)$, whereas the dimension of M is $n \times n$.

We want to find a relation between LEs and Floquet exponents and Poincaré maps. Note that by the definition of the monodromy matrix M, $\mathsf{H}(kT) = \mathsf{M}^k$, for any $k \in \mathbb{N}$. But then

$$\begin{aligned} \mu_i &= \limsup_{t\to\infty} \frac{1}{t}\log|\mathrm{ev}_i(\mathsf{H}(t))| \\ &= \lim_{k\to\infty} \frac{1}{kT}\log|\mathrm{ev}_i(\mathsf{H}(kT))| \\ &= \lim_{k\to\infty} \frac{1}{kT}\log|\mathrm{ev}_i(\mathsf{M}^k)| \\ &= \frac{1}{T}\log|\mathrm{Re}(\mathrm{ev}_i(\mathsf{M}))|. \end{aligned} \tag{3.1.16}$$

We first recall a basic result: Let $x(x_0, t)$ be the solution to (1.0.1) with initial condition x_0. Then the fundamental solution of the linearization of (1.0.1) around $x(x_0, t)$ is

$$\mathsf{H}(x(x_0,t),t) = D_{x_0}x(x_0,t). \tag{3.1.17}$$

See e.g. Theorem 8.43 in [KP10] for a proof. Now suppose x_0 lies on a periodic orbit of (1.0.1) so that $\tilde{x}_0$ is the $(n-1)$-dimensional restriction of x_0 to that orbit. Let $P(\tilde{x}_0)$ be a Poincaré map passing through x_0, we will require that this map is transverse to the flow at x_0. Then since $\tilde{x}_0 = P(\tilde{x}_0)$. We have

$$\begin{bmatrix} D_{\tilde{x}_0} P(\tilde{x}_0) & 0 \\ 0 & 1 \end{bmatrix} = D_{x_0} x(x_0, t) = \mathsf{H}(x_0, T), \tag{3.1.18}$$

where we have assumed for simplicity that $\mathrm{ev}_n(\mathsf{H}) = 1$. Now recall Floquet's theorem:

Proposition 2. *Let* $\mathsf{H}(t)$ *be the fundamental solution to the linearization of* (1.0.1) *about* $x(x_0, t)$*, where* $x(x_0, t)$ *is a* T*-periodic solution of* (1.0.1)*. Then there exists a* T*-periodic matrix* $\mathsf{Q}(t)$ *and a matrix* B *such that*

$$\mathsf{H}(t) = \mathsf{Q}(t) e^{t\mathsf{B}}. \tag{3.1.19}$$

For a proof of this theorem see e.g. [Mei17, p. 62]. Let M be the monodromy matrix, i.e. let $\mathsf{M} := \mathsf{H}(T)$. Then since $\mathsf{H}(0) = \mathbf{I}$, $\mathsf{Q}(0) = \mathbf{I}$ and so

$$\mathsf{M} = e^{T\mathsf{B}}. \tag{3.1.20}$$

Let $\mathrm{ev}_i(\mathsf{M})$ be the i'th eigenvalue of M. Most sources (see e.g. [Koo06] or [Mei17]) then define the Floquet exponents by definition 4. Finally, from (3.1.18) the Floquet exponents of a periodic orbit are just the eigenvalues of the Jacobian of the Poincaré map. So the LE of a periodic orbit can be computed via 1.0.3, the eigenvalues of M or the eigenvalues of $D_{\tilde{x}_0} P(x_0)$ if such a map is available. Section 5.3 contains an explicit example showing this relation.

3.2 Integral form of LE

The framework of chapter 2 gives bounds on infinite-time-averaged integrals $\overline{\Phi}$ over bounded trajectories, where Φ is an explicit function of the state space variable. We can apply this framework to the bounding of LEs (1.0.3) by expressing the LE μ as the time integral of some Φ. This is not possible with Φ depending only on the state space variable x of the ODE system (1.0.1), but it is straightforward if Φ is a function on the enlarged state space (x, y),

$$\begin{aligned}
\mu(x_0, y_0) &= \limsup_{T\to\infty} \frac{1}{T}\int_0^T \frac{\mathrm{d}}{\mathrm{d}s}\log|y|\,\mathrm{d}s \\
&= \limsup_{T\to\infty} \frac{1}{T}\int_0^T \frac{y\cdot\frac{\mathrm{d}y}{\mathrm{d}s}}{|y|^2}\mathrm{d}s \\
&= \limsup_{T\to\infty} \frac{1}{T}\int_0^T \frac{y^{\mathsf{T}} Df\big(x\big)y}{|y|^2}\mathrm{d}s.
\end{aligned} \tag{3.2.1}$$

The first equality follows from the definition (1.0.3) by the fundamental theorem of calculus, relying on the absolute continuity of $\log|y(t)|$. The third equality follows from the ODE (1.0.2) for $y(t)$.

In cases where $\mu(x_0, y_0)$ is positive, $y(t)$ is unbounded despite $x(t)$ being bounded, and so the framework of chapter 2 does not apply to trajectories in (x, y) state space. We therefore project the tangent dynamics (1.0.2) onto $\mathbb{S}^{n-1}$ by letting $z(t) = y(t)/|y(t)|$. Then μ can be written as a time average over a trajectory in (x, z) state space as

$$\mu(x_0, z_0) = \overline{z^{\mathsf{T}} Df(x) z}. \tag{3.2.2}$$

Versions of the above formula have appeared in the literature for decades, including with modifications for stochastic dynamics as the Furstenberg–Khasminskii formula [AI95; BBPS22].

3.3 Bounds from minimization over auxiliary functions

To bound the time average on the right-hand side of (3.2.2) using the framework of chapter 2, we need evolution equations for the (x, z) state space. The ODE for z is found by taking the time derivative of the definition $z(t) = y(t)/|y(t)|$ and applying the y ODE (1.0.2). Coupling this with the x ODE (1.0.1) gives

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix} x \\ z \end{bmatrix} = \begin{bmatrix} f(x) \\ \ell(x,z) \end{bmatrix}, \qquad (x,z) \in \mathbb{R}^n \times \mathbb{S}^{n-1}, \tag{3.3.1}$$

where

$$\ell(x,z) = Df(x)z - [z^{\mathsf{T}} Df(x) z] z. \tag{3.3.2}$$

Recall that the framework of chapter 2 was described for an ODE with state space $\mathbb{R}^n$, right-hand side $f(x)$, and initial set $\mathcal{B}_0$ whose trajectories eventually remain in $\mathcal{B}$. To bound the maximal LE of the x ODE, we consider the ODE with enlarged state space $\mathbb{R}^n \times \mathbb{S}^{n-1}$, right-hand side as in (3.3.1), and initial set $\mathcal{B}_0 \times \mathbb{S}^{n-1}$ whose trajectories eventually remain in $\mathcal{B} \times \mathbb{S}^{n-1}$. We choose $\Phi(x,z) = z^{\mathsf{T}} Df(x) z$, so that $\mu = \overline{\Phi}$, and thus the left-hand side of (2.1.3) coincides with the maximal LE $\mu^*_{\mathcal{B}_0}$ defined in (1.0.7). Applying (2.1.3) to our present context immediately gives an upper bound on $\mu^*_{\mathcal{B}_0}$, as stated in the first part of Proposition 3. When $\mathcal{B}$ is compact and trapping (i.e., forward invariant) for the x ODE, so that $\mathcal{B} \times \mathbb{S}^{n-1}$ is compact and trapping for (3.3.1), the equality (2.2.1) proved in [TGD18] can be applied also to give the second part of the proposition. (The fact that (2.2.1) holds for lower-dimensional domains such as $\mathcal{B} \times \mathbb{S}^{n-1}$ is discussed following (2.2.1) above.) In the proposition and elsewhere, the notation $f : \mathcal{C}^k(A, B)$ means that $f : A \to B$ is k times continuously differentiable.

Proposition 3. *Let* $\frac{\mathrm{d}}{\mathrm{d}t}x(t) = f(x(t))$ *with* $f \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R}^n)$.

1. *Let* $\mathcal{B}_0 \subset \mathbb{R}^n$ *and* $\mathcal{B} \subset \mathbb{R}^n$ *be such that all trajectories* $x(t)$ *with* $x(0) \in \mathcal{B}_0$ *are bounded uniformly in* $t \geq 0$ *and eventually remain in* $\mathcal{B}$. *Then the maximal LE* $\mu^*_{\mathcal{B}_0}$ *among trajectories with initial conditions in* $\mathcal{B}_0$ *is bounded above by*

$$\mu^*_{\mathcal{B}_0} \leq \inf_{V \in \mathcal{C}^1(\mathcal{B} \times \mathbb{S}^{n-1})} \sup_{\substack{x \in \mathcal{B} \\ z \in \mathbb{S}^{n-1}}} \left(z^\mathsf{T} Df z + f \cdot \nabla_x V + [Df z - (z^\mathsf{T} Df z) z] \cdot \nabla_z V \right). \tag{3.3.3}$$

2. *Let* $\mathcal{B}$ *be a compact set that is trapping (i.e., forward invariant). Then,*

$$\mu^*_{\mathcal{B}} = \inf_{V \in \mathcal{C}^1(\mathcal{B} \times \mathbb{S}^{n-1})} \sup_{\substack{x \in \mathcal{B} \\ z \in \mathbb{S}^{n-1}}} \left(z^\mathsf{T} Df z + f \cdot \nabla_x V + [Df z - (z^\mathsf{T} Df z) z] \cdot \nabla_z V \right). \tag{3.3.4}$$

In Proposition 3, $\nabla_x V$ and $\nabla_z V$ denote gradients of the auxiliary function $V(x, z)$ with respect to x and z, respectively. In the first part of the proposition, neither $\mathcal{B}_0$ nor $\mathcal{B}$ is necessarily bounded, for instance (3.3.3) holds with $\mathcal{B}_0 = \mathcal{B} = \mathbb{R}^n$ if every trajectory $x(t)$ is bounded forward in time. A similar result applying to linear systems with stochastic forcing is proved in [Kun+16].

3.4 Bounds from SOS relaxations

The right-hand minimization in (3.3.3) and (3.3.4) is intractable in general, but in the polynomial case it can be relaxed into computationally tractable SOS programs by the approach described in section 2.3.1. If the x ODE (1.0.1) is polynomial, then so are $\Phi = z^\mathsf{T} Df(x) z$ and the (x, z) ODE (3.3.1), and we let $V(x, z)$ be polynomial also. The inner supremum on the right-hand side of (3.3.3) and (3.3.4) is bounded above by a constant B if and only if the polynomial

$$P(x,z) = B - z^{\mathsf{T}} Df(x) z - f(x)\cdot\nabla_x V(x,z) - \Big[Df(x)z - (z^{\mathsf{T}} Df(x)z)z\Big]\cdot\nabla_z V(x,z) \tag{3.4.1}$$

is non-negative for all $(x,z)\in\mathcal{B}\times\mathbb{S}^{n-1}$. Note that $P(x,z)$ for the (x,z) system (3.3.1) is the analogue of $S(x)$ defined in (2.3.1) for the x ODE alone. The condition that $P\geq 0$ on $\mathcal{B}\times\mathbb{S}^{n-1}$ can be strengthened into a SOS condition. We do not require that P is SOS since this would imply non-negativity on all of $\mathbb{R}^n\times\mathbb{R}^n$. Following the standard approach described above in section 2.3.3 we formulate SOS conditions that imply non-negativity on $\mathcal{B}\times\mathbb{S}^{n-1}$ but not necessary on $\mathbb{R}^n\times\mathbb{R}^n$.

Assume that $\mathcal{B}$ is a semialgebraic set specified by polynomial inequalities $g_j(x)\geq 0$ and/or equalities $h_i(x)=0$ as defined in (2.3.3). Note that non-compact $\mathcal{B}$ are allowed, including $\mathcal{B}=\mathbb{R}^n$ when there are no constraints on x. The set $\mathbb{S}^{n-1}$ of admissible $z\in\mathbb{R}^n$ has the semialgebraic specification $1-|z|^2=0$. Thus the right-hand minimization in (3.3.3) and (3.3.4) can be relaxed into an SOS program by the approach of section 2.3.3, the result of which is stated in Proposition 4. Provided that $\mathcal{B}$ is a compact set whose semialgebraic specification is Archimedean, the SOS relaxation does not change the value of the right-hand infimum in (3.3.3) and (3.3.4). This follows from the reasoning leading to the equality (2.3.11), with P in place of S and with $1-|z|^2=0$ accompanying the inequalities and/or equalities that specify $\mathcal{B}$. The second part of the following proposition thus follows from the second part of Proposition 3.

Proposition 4. *Let* $\frac{\mathrm{d}}{\mathrm{d}t}x = f(x)$*, where* $x(t)\in\mathbb{R}^n$ *and* $f\in\mathbb{R}^n[x]$*. Let* $\mathcal{B}\subset\mathbb{R}^n$ *be either all of* $\mathbb{R}^n$ *or specified by a finite number of polynomial inequalities,* $g_j(x)\geq 0$ *for* $j\in\{1,\dots,J\}$*, and/or equalities,* $h_i(x)=0$ *for* $i\in\{1,\dots,I\}$ *as in* (2.3.3)*. Define*

$h_0(z) = 1 - |z|^2$ (i.e., $h_0(z) = 0$ coincides with $z \in \mathbb{S}^{n-1}$) and define $P(x,z)$ as in (3.4.1).

1. *Let $\mathcal{B}_0 \subset \mathbb{R}^n$ be such that all trajectories $x(t)$ with $x(0) \in \mathcal{B}_0$ are bounded uniformly in $t \geq 0$ and eventually remain in $\mathcal{B}$. Then the maximal LE $\mu^*_{\mathcal{B}_0}$ among trajectories with initial conditions in $\mathcal{B}_0$ is bounded above by*

$$\mu^*_{\mathcal{B}_0} \leq \inf_{V,\sigma_j,\rho_i \in \mathbb{R}[x,z]} B \quad s.t. \quad P - \textstyle\sum_{i=0}^{I} \rho_i h_i - \sum_{j=1}^{J} \sigma_j g_j \in \Sigma_{2n}$$
$$\sigma_j \in \Sigma_{2n}, \quad j = 1, \ldots, J, \tag{3.4.2}$$

2. *Let $\mathcal{B}$ be a compact set that is trapping (i.e., forward invariant) and whose semialgebraic specification is Archimedean. Then,*

$$\mu^*_{\mathcal{B}} = \inf_{V,\sigma_j,\rho_i \in \mathbb{R}[x,z]} B \quad s.t. \quad P - \textstyle\sum_{i=0}^{I} \rho_i h_i - \sum_{j=1}^{J} \sigma_j g_j \in \Sigma_{2n}$$
$$\sigma_j \in \Sigma_{2n}, \quad j = 1, \ldots, J, \tag{3.4.3}$$

The right-hand minimization in (3.4.2) and (3.4.3) is over $V(x,z)$, $\rho_i(x,z)$, and $\sigma_j(x,z)$ in the infinite-dimensional space $\mathbb{R}[x,z]$ of $2n$-variate polynomials. Restricting each of these tunable polynomials to a finite-dimensional subspace—i.e., choosing a finite-degree ansatz with tunable coefficients—gives an SOS program that can be solved computationally when n, I, J, and the polynomial subspace dimensions are not too large. Note that $P(x,z)$ defined in (3.4.1) is affine in B and V, so the expressions that must be SOS in (3.4.2) and (3.4.3) are affine in all tunable coefficients, giving an SOS program. Restricting the tunable polynomials to finite-degree subspaces may increase the value of the infimum on the right-hand side of (3.4.2) and (3.4.3), but one can always enlarge these subspaces and solve another SOS program with greater computational cost and the possibility of decreasing the value of the infinum. For instance, one can seek each tunable polynomial from the set of $2n$-variate polynomials of degree no larger than d. As $d \to \infty$, the sequence of SOS programs has increasing computational cost but gives a sequence of infima that are nonincreasing and converge to the infimum on the right-hand side of (3.4.2) and (3.4.3). Thus, when the equality (3.4.3) holds, SOS computations can give arbitrarily sharp upper bounds on the maximal LE $\mu^*_{\mathcal{B}}$, at least in theory. The computational examples presented below in chapter 5 show that convergence of upper bounds to the maximal LE indeed can be achieved in practice.

Chapter 4

Symmetries

The exploitation of symmetries in SOS programming significantly reduces computational cost and numerical imprecision [GP04]. In particular, when a polynomial that is constrained to be SOS has a known symmetry, block diagonal structure can be imposed on the matrix representation of that polynomial in the corresponding semidefinite program. For the SOS programs that arise when bounding time averages in ODEs, as described in chapter 2 above, the SOS programs can be formulated to have symmetries if the ODEs and the quantities whose time averages being bounded have corresponding symmetries [GF19; Lak+20]. Since our framework for bounding the maximal LE is a particular application of chapter 2, we can take advantage of symmetries that are shared by the (x, z) ODE (3.3.1), its state space, and the function $\Phi(x, z)$ whose time average gives a Lyapunov exponent. Section 4.1 describes how orthogonal symmetries of the x dynamics induce symmetries of the (x, z) dynamics, and section 4.2 describes how the latter lead to symmetries in the SOS programs that give upper bounds on LEs.

Section 4.3 describes how non-orthogonal symmetries of the x dynamics can be exploited by generalizing the framework described above in section 3.2. In particular, the definition of z as the normalization of the tangent vector y must be generalized using a norm that depends on the symmetry group of the x ODE. We then discuss the specific types of symmetries that the result in Proposition 8 applies to and show how one may make use of such symmetries by applying a suitable change of variables to the original ODE.

4.1 Symmetries of the (x, z) dynamics

An ODE is said to be symmetric under a transformation if solutions are mapped to solutions. Let $\Lambda \in GL(n)$, where $GL(n)$ denotes the group of invertible linear transformations on $\mathbb{R}^n$. An ODE $\frac{\mathrm{d}}{\mathrm{d}t}x = f(x)$ on $\mathbb{R}^n$ is symmetric under Λ if and only if f is *equivariant* under Λ, meaning $f(\Lambda x) = \Lambda f(x)$ for all $x \in \mathbb{R}^n$. A function Φ is *invariant* under Λ if $\Phi(\Lambda x) = \Phi(x)$ for all $x \in \mathbb{R}^n$. The invariant transformations of any function form a group, as do its equivariant transformations. A function is said to be $\mathcal{G}$-invariant (resp. $\mathcal{G}$-equivariant) for a group $\mathcal{G}$ if it is invariant (resp. equivariant) under all $\Lambda \in \mathcal{G}$. Turning particularly to the (x, z) dynamics, we seek a symmetry group for which the right-hand side of the ODE (3.3.1) is equivariant, and the function $\Phi(x, z) = z^\mathsf{T} Df(x) z$ is invariant, under all transformations in the group.

The transformation $(x, z) \mapsto (x, -z)$ is a symmetry both of the (x, z) ODE (3.3.1) and of the function $\Phi(x, z) = z^\mathsf{T} Df(x) z$, reflecting the fact that the growth or decay rate of a tangent vector is unaffected by reversing its direction. The equivariance of the ODE right-hand side under this transformation amounts to the oddness $\ell(x, -z) = -\ell(x, z)$, where $\ell(x, z)$ is defined as in (3.3.2), and the invariance $\Phi(x, -z) = \Phi(x, z)$ is clear.

Further symmetries of the (x, z) dynamics are guaranteed if the x ODE is symmetric under orthogonal transformations, common examples of which include sign changes or rotations. In particular, if the x dynamics are symmetric under $\Lambda \in O(n)$, where $O(n)$ denotes the group of orthogonal linear transformations on $\mathbb{R}^n$, then the (x, z) dynamics are symmetric under $(x, z) \mapsto (\Lambda x, \Lambda z)$, and so is the function $\Phi(x, z)$. To state this precisely in the following proposition, for each group $\mathcal{G} \subset O(n)$ we define the corresponding group $\mathcal{G}' \in O(2n)$ as

$$\mathcal{G}' = \left\{ \begin{bmatrix} \Lambda & 0 \\ 0 & \pm\Lambda \end{bmatrix} \in O(2n) \ : \ \Lambda \in \mathcal{G} \subset O(n) \right\}. \tag{4.1.1}$$

Proposition 5. *Let $f \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R}^n)$. Define $\ell : \mathbb{R}^n \times \mathbb{S}^{n-1} \to \mathbb{R}^n$ and $\Phi : \mathbb{R}^n \times \mathbb{S}^{n-1} \to \mathbb{R}$ by $\ell(x, z) = Df(x)z - [z^\mathsf{T} Df(x)z]z$ and $\Phi(x, z) = z^\mathsf{T} Df(x)z$. If f is $\mathcal{G}$-equivariant for some $\mathcal{G} \subset O(n)$, then (f, ℓ) is $\mathcal{G}'$-equivariant for $\mathcal{G}'$ defined by* (4.1.1)*, and Φ is $\mathcal{G}'$-invariant.*

Proof. Let $\Lambda \in \mathcal{G}$. Invariance under $\mathcal{G}'$ is equivalent to invariance under $(x, z) \mapsto (x, -z)$ and under $(x, z) \mapsto (\Lambda x, \Lambda z)$, and likewise for equivariance. For $(x, z) \mapsto (x, -z)$, the equivariance of (f, ℓ) and invariance of Φ are immediate, as described above in the text. For $(x, z) \mapsto (\Lambda x, \Lambda z)$, the invariance of Φ is shown by

$$\Phi(\Lambda x, \Lambda z) = (\Lambda z)^\mathsf{T} Df(\Lambda x)\Lambda z = z^\mathsf{T} \Lambda^\mathsf{T} \Lambda Df(x)\Lambda^{-1}\Lambda z = z^\mathsf{T} Df(x)z = \Phi(x, z), \tag{4.1.2}$$

where the second equality uses the relation $Df(\Lambda x)\Lambda = \Lambda Df(x)$ that is found by differentiating the equivariance definition $f(\Lambda x) = \Lambda f(x)$. The claimed equivariance of (f, ℓ) holds if $f(\Lambda x) = \Lambda f(x)$, which is true by assumption, and also $\ell(\Lambda x, \Lambda z) = \Lambda \ell(x, z)$. To show the latter, we note that ℓ is related to Φ by $\ell(x, z) = Df(x)z - \Phi(x, z)z$, and we use the relation $Df(\Lambda x)\Lambda = \Lambda Df(x)$ again along with the invariance of Φ,

$$\ell(\Lambda x, \Lambda z) = Df(\Lambda x)\Lambda z - \Phi(\Lambda x, \Lambda z)\Lambda z = \Lambda[Df(x)z - \Phi(x,z)z] = \Lambda\ell(x,z). \quad (4.1.3)$$

This completes the proof. □

4.2 Symmetries of the optimization problems

If an ODE, its domain, and a time-averaged quantity share an orthogonal group of symmetries, then the upper bound on this time average given by the right-hand infimum in (2.1.3) and (2.2.1) is unchanged if the auxiliary function V is further constrained to be invariant under the same symmetries. In particular, any V that yields an upper bound on a time average can be used to construct a symmetrized V yielding the same bound. This is proved in [GF19, Proposition A.1] in the context of a finite cyclic symmetry group. Proposition 18 in the appendix generalizes the result to all groups of linear orthogonal transformations.

We apply these symmetry results in the context of bounding LEs by the infimum on the right-hand side of (3.3.3) and (3.3.4), in which case the pertinent symmetry group $\mathcal{G}'$ of the (x, z) dynamics is given by Proposition 5 above. Provided that the domain $\mathcal{B}$ is $\mathcal{G}$-invariant, the following proposition guarantees that corresponding symmetries can be imposed on $V(x, z)$ in the right-hand minimization of (3.3.3) and (3.3.4) without changing the resulting upper bound.

Proposition 6. *Let $\mathcal{G} \subset O(n)$ and $f \in \mathcal{C}^1(\mathcal{B}, \mathbb{R}^n)$. If f is $\mathcal{G}$-equivariant and $\mathcal{B}$ is $\mathcal{G}$-invariant, then the infimum over $V \in \mathcal{C}^1(\mathcal{B} \times \mathbb{S}^{n-1})$ on the right-hand side of* (3.3.3) *and* (3.3.4) *is unchanged if V is constrained to be $\mathcal{G}'$-invariant for $\mathcal{G}'$ defined by* (4.1.1)*.*

Proof. Under the assumptions on $\mathcal{B}$ and f, Proposition 5 guarantees that (f, ℓ) is $\mathcal{G}'$-equivariant and $\Phi = z^{\mathsf{T}} Df(x) z$ is $\mathcal{G}'$-invariant. The domain $\mathcal{B} \times \mathbb{S}^{n-1}$ is also $\mathcal{G}'$-invariant because $\mathcal{B}$ is $\mathcal{G}$-invariant by assumption and $\mathbb{S}^{n-1}$ is $\mathcal{G}$-invariant for any $\mathcal{G} \subset O(n)$. We can therefore apply Proposition 18 from the appendix in the context of the symmetry group $\mathcal{G}'$, domain $\mathcal{B} \times \mathbb{S}^{n-1}$, ODE right-hand side (f, ℓ), and $\Phi = z^{\mathsf{T}} Df(x) z$. In this context the proposition guarantees that if there exist $V \in \mathcal{C}^1(\mathcal{B} \times \mathbb{S}^{n-1})$ and $B \in \mathbb{R}$ such that $P(x, z) \geq 0$ for all $(x, z) \in \mathcal{B} \times \mathbb{S}^{n-1}$, where P is as defined by (3.4.1), then there exists $\widehat{V} \in \mathcal{C}^1(\mathcal{B} \times \mathbb{S}^{n-1})$ that satisfies the same inequality and is $\mathcal{G}'$-invariant. Therefore, the infimum of

$$\inf_{V \in \mathcal{C}^1(\mathcal{B} \times \mathbb{S}^{n-1})} B \quad \text{s.t.} \quad P(x, z) \geq 0 \ \ \forall (x, z) \in \mathcal{B} \times \mathbb{S}^{n-1} \tag{4.2.1}$$

is unchanged if V is further constrained to be $\mathcal{G}'$-invariant. The constrained minimization problem (4.2.1) is equivalent to the minimax problem on the right-hand side of (3.3.3) and (3.3.4), so the claim is proved. □

Proposition 6, which allows invariance to be imposed on V in the minimization that appears on the right-hand side of (3.3.3) and (3.3.4), has an analogue for the minimization's SOS relaxation that appears on the right-hand side of (3.4.2) and (3.4.3). This is stated by Proposition 7 below, whose assumptions require that the $\mathcal{G}$-invariant set $\mathcal{B}$ is specified in terms of $\mathcal{G}$-invariant polynomials, and whose conclusions allow invariance to be imposed not only on V but also on any other tunable polynomials σ_j and ρ_i. We do not include a proof because the result is a direct application of Proposition 19 in the same way that Proposition 6 is an application of Proposition 18.

Proposition 7. *Let* $\mathcal{G} \subset O(n)$ *and* $f \in \mathbb{R}^n[x]$. *Let* $\mathcal{B} \subset \mathbb{R}^n$ *be either all of* $\mathbb{R}^n$ *or specified by a finite number of polynomial inequalities* $g_j(x) \geq 0$ *and equalities*

$h_i(x) = 0$ as in (2.3.3). If f is $\mathcal{G}$-equivariant and all g_j and h_i are $\mathcal{G}$-invariant, then the infimum on the right-hand side of (3.4.2) and (3.4.3) over $V, \sigma_j, \rho_i \in \mathbb{R}[x, z]$ (or over subspaces of $\mathbb{R}[x, z]$) is unchanged if V and all σ_j and ρ_i are constrained to be $\mathcal{G}'$-invariant for $\mathcal{G}'$ defined by (4.1.1). When V and all σ_j and ρ_i are $\mathcal{G}'$-invariant, the first expression constrained to be SOS in (3.4.2) and (3.4.3) is $\mathcal{G}'$-invariant also.

Proposition 7 has practical implications for the SOS programs whose solutions give upper bounds on the maximal LE—i.e., the SOS programs obtained by restricting the minimization on the right-hand side of (3.4.2) and (3.4.3) to finite spaces of polynomials. Imposing $\mathcal{G}'$-invariance on the tunable polynomials does not change the resulting upper bounds on LEs, but it ensures that all expressions constrained to be SOS are $\mathcal{G}'$-invariant. This latter invariance can be exploited in the numerical solution of SOS programs [GP04].

Remark 1. *If the x ODE has no symmetries, then $\mathcal{G}$ contains only the identity, and $\mathcal{G}'$ contains the identity and the transformation $(x, z) \mapsto (x, -z)$. In this case Proposition 6 reduces to the statement that making $V(x, z)$ even in z does not change the resulting upper bounds on LEs, and Proposition 7 says the same about making all tunable polynomials even in z.*

4.3 Modified formulation preserving non-orthogonal symmetries

This subsection describes how the z dynamics and $\Phi(x, z)$ can be modified so that the propagation of symmetries—from the x dynamics, to the (x, z) dynamics, to the minimization problems giving upper bounds on maximal LEs—is not restricted to orthogonal symmetries. The results of sections 4.1 and 4.2 above apply only to symmetry groups $\mathcal{G}$ of the x dynamics that are subgroups of $O(n)$. To derive analogous results for a chosen $\mathcal{G} \not\subset O(n)$, the framework presented in section 3.2 for bounding maximal LEs must be generalized using $\mathcal{G}$. This generalization is not used in the computational examples of chapter 5 below, where all symmetries are orthogonal, nonetheless we present the theory here for future applications.

We will also discuss the non-orthogonal symmetries which may be used with Proposition 8 after a change of variables: symmetries which are related to orthogonal symmetries up to a similarity transformation. We also explain how such symmetries can be exploited using the results of section sections 4.1 and 4.2 by a change of variables, which turns the original ODE into one with orthogonal symmetries. This means that there are at least two valid approaches for making use of non-orthogonal symmetries, either of which may have advantages over the other in practical applications.

In the definition (1.0.3) of the LE, the Euclidean norm can be equivalently replaced by any other norm on $\mathbb{R}^n$ [Mei17]. Here we replace it by a weighted norm,

$$|y|_{\mathsf{M}} := \left(y^{\mathsf{T}}\mathsf{M}y\right)^{1/2}, \tag{4.3.1}$$

where $\mathsf{M} \in \mathbb{R}^{n\times n}$ is a positive definite symmetric matrix to be chosen based on the symmetry group $\mathcal{G}$. Starting from the expression $\mu(x_0, y_0) = \limsup_{t\to\infty} \frac{1}{t}\log|y(t)|_{\mathsf{M}}$ and defining $z = y/|y|_{\mathsf{M}}$, the arguments used in the beginning of section 3.2 give $\mu(x_0, z_0) = \overline{\Phi}_{\mathsf{M}}$, where

$$\Phi_{\mathsf{M}}(x, z) = z^{\mathsf{T}}\mathsf{M}Df(x)z, \tag{4.3.2}$$

where the phase space of z is now the set $\mathbb{S}^{n-1}_{\mathsf{M}}$ of vectors in $\mathbb{R}^n$ with $|z|_{\mathsf{M}} = 1$. Differentiating $z = y/|y|_{\mathsf{M}}$ gives $\frac{\mathrm{d}}{\mathrm{d}t}z = \ell_{\mathsf{M}}(x, z)$, where

$$\ell_{\mathsf{M}}(x, z) = Df(x)z - [z^{\mathsf{T}}\mathsf{M}Df(x)z]z, \tag{4.3.3}$$

Thus the upper bounds on maximal LEs described in sections 3.3 and 3.4 still hold when Φ and ℓ are replaced by Φ_{M} and ℓ_{M} for any positive definite symmetric M. Choosing M as the identity recovers the original formulas, but there may be advantages to other choices. Other M may give better bounds on LEs from SOS programs with polynomials of fixed degree, for instance, but we do not pursue this idea here. Instead we define M using the symmetry group $\mathcal{G}$ of the x dynamics.

Let the right-hand side f of the x ODE be equivariant under some compact group $\mathcal{G} \subset GL(n)$. Because $\mathcal{G}$ is compact it has a corresponding invariant Haar probability measure m, which by definition is both left invariant and right invariant and therefore satisfies $m(A\Lambda) = m(A)$ for any m-measurable $A \subset \mathcal{G}$ and any $\Lambda \in \mathcal{G}$. This follows from the fact that any compact topological group admits an invariant Haar probability measure (Theorem 8.12 in [Sal16]). From such a measure one may obtain a right-invariant Haar integral (Theorem 6.1 in [Eat07]).

We choose the M matrix as an integral over $\mathcal{G}$

$$\mathsf{M} = \int_{\mathcal{G}} \Lambda'^{\mathsf{T}} \Lambda' \, \mathrm{d}m(\Lambda'), \tag{4.3.4}$$

because this M has the property that $\Lambda^{\mathsf{T}}\mathsf{M}\Lambda = \mathsf{M}$ for all $\Lambda \in \mathcal{G}$ (cf. the proof of Proposition 8 below). Note that M is the identity when $\mathcal{G} \subset O(n)$, reducing to the case of sections 4.1 and 4.2. With this choice of M, the $\mathcal{G}$-equivariance of f induces corresponding symmetry in Φ_{M} and ℓ_{M}. This is stated precisely by the following proposition, which reduces to Proposition 5 in the particular case where $\mathcal{G} \subset O(n)$.

Proposition 8. *Let $f \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R}^n)$. Define $\ell_{\mathsf{M}} : \mathbb{R}^n \times \mathbb{S}^{n-1}_{\mathsf{M}} \to \mathbb{R}^n$ by $\ell_{\mathsf{M}}(x,z) = Df(x)z - [z^{\mathsf{T}}\mathsf{M}Df(x)z]z$, and define $\Phi_{\mathsf{M}} : \mathbb{R}^n \times \mathbb{S}^{n-1}_{\mathsf{M}} \to \mathbb{R}$ by $\Phi_{\mathsf{M}}(x,z) = z^{\mathsf{T}}\mathsf{M}Df(x)z$. Let $\Lambda \in \mathcal{G} \subset GL(n)$ with $\mathcal{G}$ compact. If $f(\Lambda x) = \Lambda f(x)$ for all $x \in \mathbb{R}^n$, then $\ell_{\mathsf{M}}(\Lambda x, \Lambda z) = \Lambda \ell_{\mathsf{M}}(x,z)$ and $\Phi_{\mathsf{M}}(\Lambda x, \Lambda z) = \Phi_{\mathsf{M}}(x,z)$ for all $(x,z) \in \mathbb{R}^n \times \mathbb{S}^{n-1}_{\mathsf{M}}$.*

Proof. Let $\Lambda \in \mathcal{G}$. Since $m(\Lambda)$ is a right-invariant Haar probability measure,

$$\int_{\mathcal{G}} \iota(\Lambda'\Lambda)\, dm(\Lambda') = \int_{\mathcal{G}} \iota(\Lambda')\, dm(\Lambda'), \tag{4.3.5}$$

for all $\iota \in \mathcal{C}(\mathcal{G})$. Therefore,

$$\Lambda^{\mathsf{T}}\mathsf{M}\Lambda = \int_{\mathcal{G}} \Lambda^{\mathsf{T}}\Lambda'^{\mathsf{T}}\Lambda'\Lambda\, dm(\Lambda') = \int_{\mathcal{G}} (\Lambda'\Lambda)^{\mathsf{T}}(\Lambda'\Lambda)\, dm(\Lambda') = \int_{\mathcal{G}} \Lambda'^{\mathsf{T}}\Lambda'\, dm(\Lambda') = \mathsf{M}. \tag{4.3.6}$$

This relation and the relation $Df(\Lambda x)\Lambda = \Lambda Df(x)$, which follows from the $\mathcal{G}$-equivariance of f, together give the invariance

$$\Phi_{\mathsf{M}}(\Lambda x, \Lambda z) = (\Lambda z)^{\mathsf{T}}\mathsf{M}Df(\Lambda x)\Lambda z = z^{\mathsf{T}}\Lambda^{\mathsf{T}}\mathsf{M}\Lambda Df(x)\Lambda^{-1}\Lambda z = z^{\mathsf{T}}\mathsf{M}Df(x)z = \Phi_{\mathsf{M}}(x, z). \tag{4.3.7}$$

For the ℓ_{M} equivariance note that $\ell_{\mathsf{M}}(x, z) = Df(x)z - \Phi_{\mathsf{M}}(x, z)z$, from which it follows that

$$\ell_{\mathsf{M}}(\Lambda x, \Lambda z) = \Lambda[Df(x)z - \Phi_{\mathsf{M}}(x, z)z] = \Lambda\ell_{\mathsf{M}}(x, z). \tag{4.3.8}$$

□

The upper bounds on LEs given by the right-hand minimization in (3.3.3) or its SOS relaxation in (3.4.2) hold also with Φ_{M} and ℓ_{M} in place of Φ and ℓ. In these modified minimization problems, the minima are unchanged if the tunable functions are constrained to be $\mathcal{G}'$-invariant. More precisely, Propositions 6 and 7 generalize to any compact group $\mathcal{G} \subset GL(n)$ upon introduction of Φ_{M} and ℓ_{M}. These generalizations follow from Proposition 8 combined with Proposition 18 or Proposition 19; their proofs are analogous to those of Propositions 6 and 7.

However, there is another way of making use of symmetries that are not orthogonal if these symmetries take the form

$$\mathcal{G}_{\mathsf{A}} = \{\mathsf{A}^{-1}\tilde{\Lambda}\mathsf{A} \ : \ \tilde{\Lambda} \in \mathcal{G} \subset O(n)\}. \tag{4.3.9}$$

In other words $\mathcal{G}_{\mathsf{A}}$ is related to the orthogonal group by a similarity transformation. Note that since $O(n)$ is a compact group, any similarity transformation of it will also be a compact group. Therefore we could make use of this symmetry group by constructing a norm $|\cdot|_{\mathsf{M}}$ from it as in (4.3.1) where M is constructed from the symmetries $\Lambda \in \mathcal{G}_{\mathsf{A}}$ via (4.3.4), and then applying Proposition 8. Rather than using Proposition 8 however, symmetries like in (4.3.9) can also be exploited for the purpose of bounding LEs by transforming the original ODE into a system which is symmetric under the symmetry group $\mathcal{G} \subset O(n)$. Let $\tilde{x}(t) = \mathsf{A}x(t)$, so that tangent vectors for $\tilde{x}(t)$ are $\mathsf{A}y(t)$, we note that the x dynamics and the $\tilde{x}$ dynamics have the same LEs (Lemma 7.19 in [Mei17]). Let $\tilde{f}(\tilde{x}) := \mathsf{A}f(\mathsf{A}^{-1}\tilde{x})$, then the framework of section 3.2 for bounding LEs can be applied to the $\tilde{x}$ ODE $\frac{\mathrm{d}}{\mathrm{d}t}\tilde{x} = \tilde{f}(\tilde{x})$, rather than to the x ODE. The $\tilde{x}$ ODE has the orthogonal symmetry group $\mathcal{G}$, so according to Proposition 5 the $(\tilde{x}, z)$ dynamics is symmetric under $\mathcal{G}'$ in (4.1.1).

To see this, let $\tilde{\Phi}(\tilde{x}, z) := z^{\mathsf{T}} D\tilde{f}(\tilde{x})z$ and let $\tilde{\ell}(\tilde{x}, z) := D\tilde{f}(\tilde{x})z - \tilde{\Phi}(\tilde{x}, z)z$. Then if $\tilde{f}(\tilde{x})$ is equivariant for $\tilde{\Lambda} \in \mathcal{G}$, then the results in sections 4.1 and 4.2 apply to the system $\frac{\mathrm{d}}{\mathrm{d}t}\tilde{x} = \tilde{f}(\tilde{x})$, which has the same LEs as the system $\frac{\mathrm{d}}{\mathrm{d}t}x = f(x)$. By assumption, $f(x)$ is equivariant under $\Lambda = \mathsf{A}^{-1}\tilde{\Lambda}\mathsf{A}$, therefore using this equivariance and making the change of variables $x = \mathsf{A}^{-1}\tilde{x}$ we get

$$\begin{aligned} f(\Lambda x) &= \Lambda f(x) \\ \implies f(\mathsf{A}^{-1}\tilde{\Lambda}\mathsf{A}x) &= \mathsf{A}^{-1}\tilde{\Lambda}\mathsf{A}f(x) \\ \implies f(\mathsf{A}^{-1}\tilde{\Lambda}\tilde{x}) &= \mathsf{A}^{-1}\tilde{\Lambda}\mathsf{A}f(\mathsf{A}^{-1}\tilde{x}) \\ \implies \mathsf{A}^{-1}\tilde{f}(\tilde{\Lambda}\tilde{x}) &= \mathsf{A}^{-1}\tilde{\Lambda}\tilde{f}(\tilde{x}), \end{aligned} \tag{4.3.10}$$

which shows that $\tilde{f}(\tilde{\Lambda}\tilde{x}) = \tilde{\Lambda}\tilde{f}(\tilde{x})$. Also note that in the definitions of $\tilde{\Phi}$ and $\tilde{\ell}$, the Jacobian of $\tilde{f}$ is taken with respect to $\tilde{x}$, not with respect to x. Therefore, the equivariance of $\tilde{f}$ under $\tilde{\Lambda}$ implies that $D\tilde{f}(\tilde{\Lambda}\tilde{x}) = \tilde{\Lambda}D\tilde{f}(\tilde{x})\tilde{\Lambda}^{-1}$, as desired.

Chapter 5

Examples: bounding LEs of continuous systems

We will illustrate the methods in chapters 3 and 4 by applying them to specific examples of ODEs. We first bound the maximal LE of the Lorenz system and show that it is attained at the origin. We then bound the maximal LE of the Arneodo system in section 5.2, a chaotic system of relatively low degree, which is used to demonstrate the effect, which scaling the systems coordinates, can have on the stability of the SOS programs that need to be numerically solved to bound the maximal LE. Example section 5.3 is included because it allows the explicit derivation of a Poincaré map for its period orbit. This allows us to illustrate this concept and analytically confirm that the Floquet exponent on the unstable cycle of the system agrees with the upper bound computed using SOS programming. Finally, we bound a maximal LE of the Hénon–Heiles system in section 5.4 by restricting the dynamics to an energy volume that does not contain any fixed points. The maximal LE in this example is therefore attained on a non-equilibrium orbit. Finding this orbit and computing this LE using classical methods is considerably more challenging than doing so using our convex optimization approach and does not allow one to conclude that there are no orbits with a larger LE. Using the SOS formulation in Proposition 4 we therefore show that the Hénon–Heiles system at the given energy values cannot have any LEs that are greater than the one we report and that it is attained on the orbits which we compute.

5.1 Lorenz system

Consider the Lorenz system [Lor63],

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} \sigma(x_2 - x_1) \\ rx_1 - x_2 - x_1x_3 \\ x_1x_2 - \beta x_3, \end{bmatrix}, \tag{5.1.1}$$

with the standard chaotic parameter values $(\beta, \sigma, r) = (8/3, 10, 28)$. For the unstable fixed point at the origin, or any trajectory approaching it along its stable manifold, the LEs are the real parts of the eigenvalues of $Df(\mathbf{0})$. Thus the leading LE at the origin is $\mu_1 = \left(-1 - \sigma + \sqrt{1 - 2\sigma + 4r\sigma + \sigma^2}\right)/2 \approx 11.82772$. In fact, this is the maximal LE among all trajectories of the Lorenz system, as confirmed by the upper bounds we report below.

We use the SOS approach of section 3.4 to compute upper bounds on the maximal LE $\mu^*_{\mathbb{R}^3}$ of the Lorenz system. For any choice of $\mathcal{B}$ containing the global attractor, (3.4.2) provides a valid upper bound. We choose $\mathcal{B} = \mathbb{R}^3$, in which case the constraints on the right-hand side of (3.4.2) reduce to simply $P - \rho_0 h_0 \in \Sigma_n$. For this non-compact $\mathcal{B}$, Part 2 of Proposition 4 does not apply to guarantee sharpness of the upper bound, nonetheless the computations reported below give the sharp result. The SOS programs we solve are obtained by restricting V and ρ_0 to finite-dimensional spaces, in particular by specifying a maximum degree d and imposing invariances that, according to Proposition 7, do not change the optimum of the SOS program. In addition to the invariance under negation of z that can always be imposed, the symmetry of the Lorenz equations under $(x_1, x_2, x_3) \mapsto (-x_1, -x_2, x_3)$ induces another, so V and ρ_0 are optimized within the spaces

$$\mathcal{V}_d = \left\{ p \in \mathbb{R}[x, z]_d \ : \ p(x, z) = p(x, -z) = p(\Lambda x, \Lambda x) \ \ \forall (x, z) \in \mathbb{R}^3 \times \mathbb{S}^2 \right\}, \quad (5.1.2)$$

where Λ negates the first two coordinates, and $\mathbb{R}[x, z]_d$ denotes the set of polynomials in (x, z) with total degree no larger than d. In our computations we fix $\deg(\rho_0) \leq 2$ and let $\deg(V) \leq d$ with various d. Each fixed d gives an SOS program that we solve to obtain upper bounds on the maximal LE of the Lorenz system:

$$\mu^*_{\mathbb{R}^3} \leq \inf_{\substack{V \in \mathcal{V}_d \\ \rho_0 \in \mathcal{V}_2}} B \quad \text{s.t.} \quad P - \rho_0(1 - |z|^2) \in \Sigma_6, \quad (5.1.3)$$

where $P(x, z; B)$ is defined as in (3.4.1).

Table 5.1 reports the upper bounds (5.1.3) found by numerically solving the SOS program with the maximum degree of ρ_0 fixed to 2, and the maximum degree of V fixed to $d = 2$, 3, and 4. The upper bounds improve as d is raised, and when $d = 4$ the numerical approximation of the upper bound is sharp to 7 digits. We expect that the exact solution to the SOS program with $d = 4$ would be equal to the maximal LE of the Lorenz system, meaning (5.1.3) is an equality in this case. These degrees of ρ_0 and V needed for a sharp bound are particular to the Lorenz system; in other cases one may need larger degrees of all tunable polynomials.

The results of table 5.1 constitute an example of our method giving sharp upper bounds on the maximal LE of a chaotic system using polynomials of modest degree—and, therefore, using SOS computations of modest cost. The example is relatively easy, even though the Lorenz system is chaotic, in the sense that the maximal LE is attained by the fixed point at the origin rather than a more complicated trajectory.

Table 5.1: Upper bounds on the global maximal LE ($\mu^*_{\mathbb{R}^3}$) of the Lorenz system (5.1.1) at the standard parameters, found by numerically solving the right-hand SOS program in (5.1.3) with the maximum degree d of V fixed to various values. Tabulated values are rounded to the precision shown.

Degree of V	Upper bound on maximal LE
2	14.02834
3	13.01444
4	11.82772

In the next subsection we apply our method to a more challenging example.

5.2 Arneodo system

As an example of a system, which despite only containing at most cubic terms in its vector field still requires a high polynomial degree in the auxiliary function in the SOS program for obtaining sharp bounds, consider the Arneodo system [ACT80]

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} x_2 \\ x_3 \\ cx_1^3 + ax_1 + bx_2 - x_3 \end{bmatrix}, \tag{5.2.1}$$

where the parameters are fixed at $a = 5.5$, $b = -3.5$ and $c = -1$, putting the system in the chaotic regime. At these parameter values the system has a fixed point at the origin as well as two additional symmetry-related fixed points on the x_1 axis. The system is also symmetric under negation of all its arguments $(x_1, x_2, x_3) \mapsto (-x_1, -x_2, -x_3)$.

The values in table 5.2 were computed by scaling the attractor by the change of variables $\tilde{x} = x/10$ to improve numerical stability. [PC89] gives $\mu \approx 0.23436$ for the chaotic LE of the Arneodo system. However, the largest eigenvalue of the Jacobian of

Table 5.2: Upper bounds on the global maximal LE ($\mu^*_{\mathbb{R}^3}$) of the Arneodo system (5.2.1) at the parameters given in the text, found by numerically solving the right-hand SOS program in (5.1.3) with $\Lambda = -\mathbf{I}$ for (5.2.1) with the maximum degree d of $V(x, z)$ fixed to the values in the table. The degree of $\rho_0(x, z)$ is fixed to 2. Tabulated values are rounded to the precision shown.

Degree of V	Upper bound on maximal LE
2	2.894867
4	1.275372
6	1.000000
8	1.000000

(5.2.1) at the origin is $\mu = 1$. Therefore, the bounds in table 5.2 appear to be sharp to the shown precision when the maximum degree of the auxiliary function in the SOS program is 6. When the SOS program is solved without the scaling $\tilde{x} = x/10$, sharp bounds are only obtained at degree-8 $V(x, z)$, and only four digits of accuracy are obtained. While we did not need to take these considerations into account when we bounded the maximal LE of the Lorenz system, for a system like the one here it appears that scaling the system correctly can have a significant effect on the accuracy of the solutions and the ability to obtain sharp bounds at low degree $V(x, z)$. This shows that the numerical stability of the SOS programs can be greatly improved with appropriate scalings of the variables, even for relatively simple systems like the Arneodo system.

5.3 A circular limit system

The system in this section attains its maximal LE on a periodic orbit. The orbit's Floquet exponent can be found analytically via a Poincaré map. Consider the following system in polar coordinates:

$$\begin{aligned} \dot{r} &= r(r^2 - 1) \\ \dot{\theta} &= 1. \end{aligned} \tag{5.3.1}$$

In Cartesian coordinates this system has the form

$$\begin{aligned} \dot{x}_1 &= x_1(x_1^2 + x_2^2 - 1) - x_2 \\ \dot{x}_2 &= x_2(x_1^2 + x_2^2 - 1) + x_1. \end{aligned} \tag{5.3.2}$$

This system has an unstable limit cycle which is given by

$$x_1(t) = \cos(t), \qquad x_2(t) = \sin(t). \tag{5.3.3}$$

If we solve the SOS program (3.4.3) with the above system we obtain a bound on the unstable LE of this limit cycle. The results seem to be sharp for degree-2 $V(x, z)$ and give

$$\mu = 2.0000. \tag{5.3.4}$$

In order to compare these results with the actual values of the LE on the periodic orbit, some additional stability analysis was performed. The analogous system to (5.3.2), albeit with stable limit cycle, is treated in [Str15, p. 282] and [Per13, p. 213]. One can derive a Poincaré map by solving the system (5.3.1) by separation of variables

$$\int_{r_0}^{P(r_0)} \frac{1}{r(r^2 - 1)}\, dr = \int_0^{2\pi} dt = 2\pi. \tag{5.3.5}$$

Evaluating this integral and solving for $P(r_0)$ gives

$$P(r_0) = \left(1 + e^{4\pi}\left(\frac{1}{r_0^2} - 1\right)\right)^{-1/2}. \tag{5.3.6}$$

Its derivative with respect to r_0 is

$$P'(r_0) = \frac{e^{4\pi}}{r_0^3}\left(1 + e^{4\pi}\left(\frac{1}{r_0^2} - 1\right)\right)^{-3/2}. \tag{5.3.7}$$

Evaluating this at $r_0 = 1$ implies that the characteristic multiplier obtained from the Poincaré map is $P'(1) = e^{4\pi}$. Since the period of the orbit at $r = 1$ is $T = 2\pi$, this implies that by (3.1.15), the maximal Floquet exponent is

$$\mu = \frac{\log(e^{4\pi})}{2\pi} = 2. \tag{5.3.8}$$

This is also the value of the bound obtained from the SOS program. Therefore this shows that the limit cycle on the unit circle is in fact the LE-maximizing orbit of the system.

5.4 Hénon–Heiles system

To demonstrate the success of our method when the maximal LE is attained on an orbit that is more complicated than a fixed point and the underlying dynamics are chaotic, we consider the Hénon–Heiles system [HH64; SM03], which is a Hamiltonian system. The Hamiltonian function is

$$H(x_1, x_2, x_3, x_4) = \frac{1}{2}(x_1^2 + x_2^2 + x_3^2 + x_4^2) + x_1^2 x_2 - \frac{1}{3}x_2^3, \tag{5.4.1}$$

where x_1 and x_2 are position variables with corresponding momentum variables x_3 and x_4, respectively. This H gives rise to the ODE system

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} x_3 \\ x_4 \\ -x_1 - 2x_1x_2 \\ -x_2 - x_1^2 + x_2^2 \end{bmatrix}. \tag{5.4.2}$$

Since the Hamiltonian function is conserved along trajectories, any chosen range of H values defines an invariant region in the four-dimensional phase space. For our example, we seek a region for which the maximal LE is not attained by a fixed point. The fixed points of (5.4.2) have energies of $H = 0$ and $H = 1/6$. The three fixed points with $H = 1/6$, which are related by symmetry as described below, have a large leading LE. We omit these points by restricting attention to the set where $0 \leq H \leq 1/7$. This set consists of several disconnected regions in $\mathbb{R}^4$, one of which is bounded; fig. 5.1 shows the intersection of these regions with the (x_1, x_2) plane. To restrict to the bounded region we add the condition $x_1^2 + x_2^2 \leq 1$, so the region of interest is the Archimedean, compact, semialgebraic set

$$\mathcal{B} = \{x \in \mathbb{R}^4 : g_j(x) \geq 0,\ j = 1, 2, 3\}, \tag{5.4.3}$$

where

$$g_1 = 1/7 - H(x_1, x_2, x_3, x_4), \tag{5.4.4a}$$

$$g_2 = H(x_1, x_2, x_3, x_4), \tag{5.4.4b}$$

$$g_3 = 1 - x_1^2 - x_2^2. \tag{5.4.4c}$$

The set $\mathcal{B}$ is invariant under the flow of (5.4.2) since the invariant region where $0 \leq H \leq 1/7$ does not intersect the surface where $g_3 = 0$. Because $\mathcal{B}$ is an Archimedean, compact, semialgebraic set, Part 2 of Proposition 4 guarantees that the equality (3.4.3) holds, and so SOS programs with increasing polynomial degrees can give arbitrarily sharp upper bounds on $\mu^*_{\mathcal{B}}$.

Earlier numerical work by Shevchenko and Mel'nikov [SM03] and by Benettin et al. [Ben+80] suggests that the value of the the chaotic Lyapunov exponent of the HH system increases monotonically as a function of the energy H (5.4.1). We therefore expect the maximal LE to be realized on an orbit with $H = 1/7$ in the region as defined above.

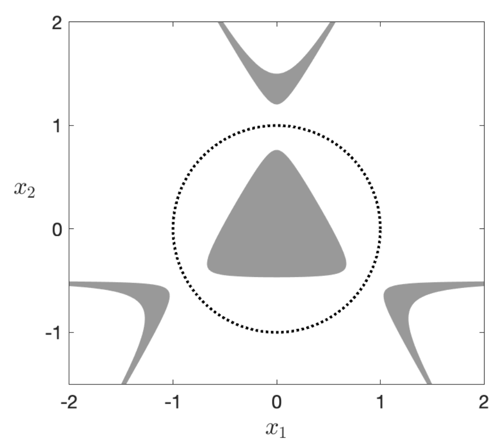

Figure 5.1: The shaded regions show the intersection of the (x_1, x_2) plane with the set where $0 \leq H \leq 1/7$. These disconnected regions are separated by the surface where $x_1^2 + x_2^2 = 1$ $(\cdots\cdots)$. The central shaded region is where $\mathcal{B}$ intersect the (x_1, x_2) plane.

We use the SOS approach of section 3.4 to compute upper bounds on the maximal LE $\mu^*_{\mathcal{B}}$ among trajectories in the invariant set $\mathcal{B}$. In order to impose symmetries on the SOS programs as described in section 4.2, we note that the Hénon–Heiles ODE has a symmetry group $\mathcal{G} \subset O(4)$. Each $\Lambda \in \mathcal{G}$ is a transformation in which some $A \in D_3$ acts on both the position vector (x_1, x_2) and the momentum vector (x_3, x_4), where D_3 is the dihedral group generated by two transformations: Reflection in the first coordinates of (x_1, x_2) and (x_3, x_4), and rotation by $2\pi/3$ in both the (x_1, x_2) plane and the (x_3, x_4) plane. Since the g_j polynomials defining $\mathcal{B}$ are $\mathcal{G}$-invariant also, Proposition 7 guarantees that the upper bound from the SOS program is unchanged if $\mathcal{G}'$-invariance is imposed on all tunable polynomials. However, because the software [Fan19; Löf09] with which we implement SOS programs does not automatically exploit rotational symmetries, the only symmetry of (5.4.2) that we make use of is $(x_1, x_2, x_3, x_4) \mapsto (-x_1, x_2, -x_3, x_4)$. The tunable polynomials are therefore optimized within the spaces

$$\mathcal{W}_d = \left\{ p \in \mathbb{R}[x, z]_d \ : \ p(x, z) = p(x, -z) = p(\Lambda x, \Lambda z) \ \ \forall (x, z) \in \mathbb{R}^4 \times \mathbb{S}^3 \right\}, \quad (5.4.5)$$

where Λ negates the first and third coordinates. In our computations we specify the same maximum degree d for all tunable polynomials. Each fixed d gives an SOS program that we solve to obtain an upper bound on the maximal LE of the Hénon–Heiles system in the region $\mathcal{B}$:

$$\mu^*_{\mathcal{B}} \leq \inf_{V, \rho_0, \sigma_1, \sigma_2, \sigma_3 \in \mathcal{W}_d} B \quad \text{s.t.} \quad P - \rho_0(1 - |z|^2) - \textstyle\sum_{j=1}^{3} \sigma_j g_j \in \Sigma_8, \qquad (5.4.6)$$
$$\sigma_j \in \Sigma_8.$$

where $P(x, z; B)$ is defined as in (3.4.1).

Table 5.3 reports the upper bounds (5.4.6) found by numerically solving the SOS program with the maximum degrees of all tunable polynomials fixed to $d = 2$, 4, 6, 8, and 10. The upper bounds improve as d is raised, seeming to converge to 5 digits once polynomials are of degree 8. To confirm the sharpness of this upper bound, we computed a variety of unstable periodic orbits of the Hénon–Heiles system and calculated the leading LE of each, using a procedure described below. On the shortest-period orbit we computed a leading LE of 0.23081, confirming that the best

Table 5.3: Upper bounds on the maximal LE ($\mu^*_{\mathcal{B}}$) among trajectories of the Hénon–Heiles system (5.4.2) in the set $\mathcal{B}$ defined by (5.4.3) and (5.4.4), found by numerically solving the right-hand SOS program in (5.4.6) with the maximum degree d of all tunable polynomials fixed to various values. Tabulated values are rounded to the precision shown. Figure 5.2 shows the periodic orbits on which numerical integration gives $\mu_1 \approx 0.23081$.

Degree of polynomials	Upper bound on maximal LE
2	0.86999
4	0.43992
6	0.26744
8	0.23081
10	0.23081

upper bounds in table 5.3 are sharp to at least five digits.

To compute a number of unstable periodic orbits—and thus to find one with a large leading LE—we first numerically integrated the Hénon–Heiles system from various initial conditions and searched for close returns to the initial conditions. A fourth–order symplectic integrator [FR90; Hai10] was used to conserve energy. Choosing 121 initial conditions along the curve where $x_3, x_4 = 0$ and $H = 1/7$, we integrated each trajectory for 33 time units. For every trajectory that satisfied the close return condition $|x(t) - x(0)| < 10^{-3}$ at some times $t \geq 1$, we selected the smallest such time T and corresponding initial condition $x(0)$ as producing an approximate periodic orbit.

After compiling 19 pairs of initial conditions and periods giving approximate periodic orbits, we used a shooting method to converge each more precisely to a periodic orbit. Very precise convergence was essential because changes in an orbit produced changes in the computed LE which were proportional to, and often worse than, the size of the error in the initial data. We implemented a shooting method which used MATLAB's `fminsearch` function to minimize $|x(T) - x(0)|$ by tuning the period T and the x_1 coordinate of the initial condition. (The initial x_2 coordinate was determined from the conditions that $x_3, x_4 = 0$ and $H = 1/7$). On each iteration of the shooting method, symplectic integration up to time T was carried out using a fixed time step close to 10^{-3} that evenly divides T. All 19 orbits were converged such that $|x(T) - x(0)| < 10^{-13}$, which required modifying the default options of the `fminsearch` function to tighten tolerances and allow more iterations. Among the 19 converged orbits there were only 7 different periods: approximately 6.97, 8.07, 15.6, 23.0, 29.3, 29.7, and 30.0, with orbits of the same period being related to each other by symmetries of the Hénon–Heiles system. As described below, the largest LE was found on the orbit with the shortest period, whose numerically converged initial condition and period are

$$x(0) = (0.562878385826716, -0.053847890920149, 0, 0), \quad T = 6.966517640959103. \tag{5.4.7}$$

(For the above values, symplectic integration gives $|x(T) - x(0)| \approx 2 \times 10^{-14}$. Figure 5.2 shows this shortest periodic orbit, along with the two symmetry-related orbits obtained by rotating both the position vector (x_1, x_2) and momentum vector (x_3, x_4) by $\pm 2\pi/3$.

To compute the leading LE on each periodic orbit we integrated the linearized equations (1.0.2) and (1.0.1) for the unnormalized tangent vector y and the orbit x, using an RK4 scheme. Starting with y of unit length and pointing in a random direction, we integrated (1.0.2) for one period and set $y(0)$ to $y(T)/|y(T)|$, after which we integrated (1.0.2) again for one period to similarly obtain a new $y(0)$. We carried out this procedure 14 times, after which $y(0)$ did not change its orientiation any more to machine precision. This value for $y(0)$ was taken to be the Lyapunov vector corresponding to the most rapidly expanding direction in the tangent space of the orbit. The system (1.0.2) was then integrated for one more period with this converged

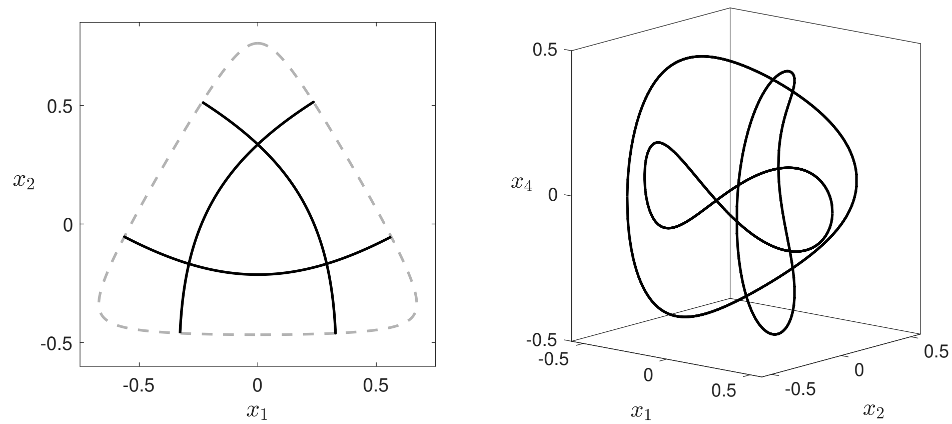

Figure 5.2: Orbits of the Hénon–Heiles system with maximum known LE of $\mu_1 = 0.23081$. The three orbits are related by symmetry (see text). The left panel shows their projection (—) onto the (x_1, x_2) plane along with that of the $H = 1/7$ energy surface (- - -), and the right panel shows their projection onto the (x_1, x_2, x_4) space.

$y(0)$ after which $\frac{1}{T} \log |y(T)|$ gave the LE that we sought. (This way of approximating the LE seemed to give a more accurate leading LE than the QR algorithm [PC12], which approximates the full spectrum of LEs.) For the shortest period orbits that are shown in fig. 5.2, the procedure described above gave a leading LE of $\mu_1 = 0.23081$. This value agrees with our best upper bound in table 5.3 to all five digits, which is the precision to which we trust the SOS computations giving the bounds. These results strongly suggest that the maximal LE among trajectories in $\mathcal{B}$ indeed occurs on the orbits shown in fig. 5.2.

We emphasize that using numerical integration to find periodic orbits and obtain the precise value $\mu_1 = 0.23081$ on one such orbit was much more difficult than using SOS optimization to compute the upper bound $\mu^*_{\mathcal{B}} \leq 0.23081$. Thus there will always be numerical difficulty in finding orbits that attain the maximal LE, and in precisely computing this LE on these orbits. On the other hand, these considerations do not affect the SOS programs that give upper bounds on the maximal LE.

Chapter 6

Bounding time averages of discrete maps

In this section and in chapter 7 the problem of bounding LEs of discrete maps is studied. Many of the results and approaches that were developed for bounding the maximal LEs of continuous-time systems in the previous sections also translate to the case of discrete-time maps. However, the naïve approach of formulating the problem by replacing time-derivatives by finite differences in (3.3.4) does not lead to computationally tractable optimization problems in many cases. We explain the difficulty one faces when trying to bound the LE as arithmetic means via SOS programming and how one can remedy this by expressing the LE as a geometric mean. The remainder of this section develops the framework and theory of bounding geometric means of scalar quantities along bounded orbits of discrete maps. In sections 6.2 and 6.3 we prove two formal results which guarantee sharp bounds on time averages using convex optimization formulations and SOS programming, respectively.

Recall that the LE of a discrete-time system (1.0.6) is computed via the linearization (1.0.5) of the system (1.0.4). Let us see how far we can get when we try to directly follow the convex optimization methods to bound the maximal LE of (1.0.4) like we did for continuous-time systems in the previous sections. We first express (1.0.6) as a time average, which we can bound by enforcing suitable non-negativity—and later SOS—constraints. As we shall see, this approach only works to derive pointwise-constrained optimization problems, but it does not work in general if the underlying dynamics in (1.0.4) are polynomial in the phase-space variables and the constraints are to be relaxed to SOS constraints. Let us begin by expressing (1.0.6) as an infinite-time average:

$$\begin{aligned}\mu(x_0, y_0) &= \limsup_{k\to\infty} \frac{1}{k} \log\left(\frac{|y_k|}{|y_0|}\right) \\ &= \limsup_{k\to\infty} \frac{1}{k}(\log|y_k| - \log|y_0|) \\ &= \limsup_{k\to\infty} \frac{1}{k}\sum_{i=0}^{k-1}(\log|y_{i+1}| - \log|y_i|) \\ &= \limsup_{k\to\infty} \frac{1}{k}\sum_{i=0}^{k-1}(\log|Df(x_i)y_i| - \log|y_i|) \\ &= \limsup_{k\to\infty} \frac{1}{k}\sum_{i=0}^{k-1}\log|Df(x_i)z_i|,\end{aligned} \tag{6.0.1}$$

where $z = \frac{y}{|y|}$ as before. We now have the LE expressed as an arithmetic mean depending on x and z. We will therefore explore how—and if—we can bound such an average using convex optimization and SOS programming techniques.

6.1 Optimization formulations for bounding time averages of discrete maps

In this section we describe two different notions of time-average for discrete maps (1.0.4). In section 6.1.1 we derive optimization problems for bounding the arithmetic mean of a scalar quantity and in section 6.1.2 we derive optimization problems and SOS programs for bounding the geometric mean of a scalar quantity. We explain why the first approach does not usually have an SOS relaxation, while the second approach does, and describe how the resulting non-convex SOS program is solved in practice via a sequence of convex SOS feasibility problems.

6.1.1 Optimization formulations for bounding arithmetic means

In section 2.1 we derived optimization problems to bound time averages by adding the time derivative of an auxiliary function $V(x)$ to the quantity whose average was bounded, noting that this did not change the value of the time average. For discrete time systems the analogue of a time derivative is a finite difference, $V(f(x)) - V(x)$ along orbits, and indeed, we see that the time average of such a finite difference is also zero along bounded orbits of (1.0.4) if $V \in \mathcal{C}(\mathcal{B})$, since

$$\limsup_{k\to\infty} \frac{1}{k} \sum_{i=0}^{k-1} \Big(V\Big(f(x_i)\Big) - V(x_i)\Big) = \limsup_{k\to\infty} \frac{1}{k} \Big(V\Big(f(x_{k-1})\Big) - V(x_0)\Big) = 0. \quad (6.1.1)$$

Therefore, if

$$\overline{\Psi}(x) := \limsup_{k\to\infty} \frac{1}{k} \sum_{i=0}^{k-1} \Psi(x_i), \quad (6.1.2)$$

then like in (2.1.3), we can add the finite difference of an auxiliary function $V \in \mathcal{C}(\mathcal{B})$ to $\Psi(x)$ and bound $\overline{\Psi}$ via a convex minimization of $\Psi(x) + V(f(x)) - V(x)$ over a suitable domain $\mathcal{B} \subset \mathbb{R}^n$ and auxiliary functions $V \in \mathcal{C}(\mathcal{B})$:

$$\sup_{x_0 \in \mathcal{B}_0} \overline{\Psi} \leq \inf_{V \in \mathcal{C}(\mathcal{B})} \sup_{x \in \mathcal{B}} \left[\Psi(x) + V(f(x)) - V(x)\right]. \quad (6.1.3)$$

In fact, sharpness results like those in Lemma 4 below assert that if $\mathcal{B}$ is compact and trapping then the inequality in (6.1.3) can be strengthened to an equality:

$$\sup_{x_0 \in \mathcal{B}} \overline{\Psi} = \inf_{V \in \mathcal{C}(\mathcal{B})} \sup_{x \in \mathcal{B}} \left[\Psi(x) + V(f(x)) - V(x) \right]. \tag{6.1.4}$$

In principle, this result could be applied to the problem of bounding the maximal LE of (1.0.4) by coupling the dynamics of the tangent space of the map, which is encoded by $z_{k+1} = \frac{Df(x_k)z_k}{|Df(x_k)z_k|}$, to the dynamics governing x_k. We would simply need to enlarge the phase space to $\mathcal{B} \times \mathbb{S}^{n-1}$, let $\Psi(x, z) = \log |Df(x)z|$ as in (6.0.1) and apply (6.1.4) to these quantities to obtain a sharp bound on the maximal LE of (1.0.4). While this works in theory, we run into a problem as soon as we attempt to solve it in practice, by turning it into a SOS program via a SOS relaxation. The problem is that for a polynomial map, $\log |Df(x)z|$ and the map governing z_k are not polynomial in the components of x and z. Therefore, it is usually not possible to relax (6.1.4) to a SOS program when we want to bound LEs of polynomial maps. For continuous-time dynamical systems this issue did not arise, because expressing the definition of the LE (1.0.3) as an infinite-time average introduced a derivative of the logarithm, which removed the logarithm and replaced it with something that could be made polynomial in its arguments. Furthermore, the ODE governing the dynamics of $z(t)$ was inherently polynomial for polynomial $f(x)$.

Therefore, we need to amend our approach to expressing time averages of discrete maps so that we can derive an optimization problem for bounding LEs, which can be relaxed to a SOS program that is numerically solvable. It turns out that a more natural notion of time average in this case is a *geometric mean*, rather than the arithmetic mean in this section.

6.1.2 Optimization formulations for bounding geometric means

In this section, the problem of bounding the geometric mean of a scalar quantity along orbits of discrete systems is studied. In order to derive SOS-relaxable optimization problems for bounding LEs of discrete maps (1.0.4), we need to work with a notion of time average, which will allow us to circumvent the inherently non-polynomial expressions arising from the logarithm in the definition of the LE (1.0.6). The key here is that *the arithmetic mean of a logarithm is the logarithm of a geometric mean.* Our strategy then, is to derive optimization problems for bounding the geometric mean of scalar quantities $\Phi(x)$ along bounded trajectories of the system (1.0.4), and to then apply this approach to deriving an optimization procedure for bounding the maximal LE of the system, which can be relaxed to a SOS program.

Definition 5 (Geometric mean). *Let* $\Phi(x) \geq 0$ *on an orbit* x_k *of the system* (1.0.4)*. Then the* geometric mean *of* $\Phi(x)$ *along* x_k *is defined as*

$$\widetilde{\Phi}(x) := \limsup_{k\to\infty} \left(\prod_{i=0}^{k-1} \Phi(x_i) \right)^{1/k}. \tag{6.1.5}$$

The requirement that $\Phi(x)$ *be non-negative for each* x_k *on the orbit is imposed to avoid complex-valued means.*

If we now have a pointwise bound B on $\Phi(x)$ over all $x \in \mathcal{B} \subset \mathbb{R}^n$, then of course B would be a bound on the geometric mean of $\Phi(x)$ along any trajectory. Like in the continuous case however, finding such a bound is generally not possible, and we therefore have to modify $\Phi(x)$ in such a way so that its geometric mean remains unaltered, but that allows one to find a pointwise bound on the modified version of $\Phi(x)$.

To do this, we once again introduce an *auxiliary function* $V(x)$. If $V(x) > 0$ and is continuous, then one has

$$\limsup_{k\to\infty}\left(\prod_{i=0}^{k-1}\frac{V(x_{i+1})}{V(x_i)}\right)^{1/k} = \limsup_{k\to\infty}\left(\frac{V(x_k)}{V(x_0)}\right)^{1/k} = 1, \tag{6.1.6}$$

since $V(x)$ must be uniformly bounded on all bounded trajectories x_k of (1.0.4).

This means that $\Phi(x_i)$ can be replaced by the modified form $\Phi(x_i)V(x_{i+1})/V(x_i)$ in (6.1.5) without changing the mean. Thus

$$\begin{aligned}\widetilde{\Phi}(x) &= \limsup_{k\to\infty}\left(\prod_{i=0}^{k-1}\Phi(x_i)\right)^{1/k} \\ &= \limsup_{k\to\infty}\left(\prod_{i=0}^{k-1}\Phi(x_i)\frac{V(x_{i+1})}{V(x_i)}\right)^{1/k}.\end{aligned} \tag{6.1.7}$$

Clearly, since $x_{k+1} = f(x_k)$, if there exists a number $B \in \mathbb{R}$ such that

$$\Phi(x)\frac{V(f(x))}{V(x)} \leq B \tag{6.1.8}$$

for all $x \in \mathcal{B}$, then B would be a bound on the geometric mean of Φ since $\widetilde{\Phi}(x) \leq B$.

Since $V(x) > 0$ the above inequality can equivalently be written as

$$0 \leq BV(x) - \Phi(x)V(f(x)). \tag{6.1.9}$$

We can therefore find the smallest B satisfying this inequality by solving the following optimization problem:

$$\begin{aligned}\sup_{x_0\in\mathcal{B}_0}\widetilde{\Phi} \leq \inf_{V\in\mathcal{C}(\mathcal{B})} \quad B \qquad s.t. \qquad & BV - \Phi V(f) \geq 0 \\ & V > 0.\end{aligned} \tag{6.1.10}$$

The above bound is always true if $\mathcal{B}_0$ contains only initial conditions which are uniformly bounded in discrete time k and which eventually enter and remain in $\mathcal{B}$. If $\mathcal{B} = \mathbb{R}^n$ and $f \in \mathbb{R}[x]$ and $\Phi \in \mathbb{R}[x]$ then we can make the SOS relaxation from section 2.3.1 to get the following SOS program:

$$\sup_{x_0 \in \mathcal{B}_0} \widetilde{\Phi} \leq \inf_{V \in \mathbb{R}[x]} \quad B \qquad s.t. \qquad \begin{aligned} BV(x) - \Phi(x)V(f(x)) &\in \Sigma_n \\ V(x) - \epsilon &\in \Sigma_n, \end{aligned} \tag{6.1.11}$$

where $\epsilon > 0$ is an arbitrary positive constant which is introduced to enforce the requirement that $V(x)$ be strictly positive. Proposition 10 below explains why it is permissible to replace the constraint $V(x) > 0$ with $V(x) - \epsilon \geq 0$ without affecting the optimum of the optimization problem. While this problem can likely be solved using methods from bilinear semidefinite programming (see e.g. [BEG93]), it has the disadvantage that the term BV contains coefficients of V that are multiplied by the design variable B. This means that we cannot leave B undetermined at the same time as leaving the coefficients of the polynomial V undetermined, since all the design variables have to appear affinely in the polynomial constraint to preserve convexity.

To get around this issue, one can fix B, and solve a series of convex feasibility problems. That is, guess a value of B and fix it. Then solve the convex feasibility problem

$$\text{find} \quad V \in \mathbb{R}[x] \qquad s.t. \qquad \begin{aligned} BV - \Phi V(f) &\in \Sigma_n \\ V - \epsilon &\in \Sigma_n. \end{aligned} \tag{6.1.12}$$

If the problem is feasible, decrease B, if the problem is infeasible, increase B. Decrease the length by which B is increased or decreased, whenever B changes from being feasible to infeasible or infeasible to feasible. Repeat solving this feasibility problem with increasingly accurate guesses for B. The optimal value of B at some fixed degree of $V(x)$ occurs right at the boundary between feasibility and infeasibility of making the polynomial constraint SOS. This can be seen from the fact that if $BV(x) - \Phi(x)V(f(x)) \geq 0$ for some B, then $(B+\delta)V(x) - \Phi(x)V(f(x)) \geq 0$ and if $BV(x) - \Phi(x)V(f(x)) \leq 0$, then $(B-\delta)V(x) - \Phi(x)V(f(x)) \leq 0$ for any $\delta > 0$ because $V(x) > 0$. Therefore, this approach will converge to the optimal value of B. We will refer to this approach as the *bisection method*.

If $\mathcal{B}$ is a semialgebraic subset of $\mathbb{R}^n$, so that it may be expressed via a finite number of equality constraints as in (6.3.1), then we can use the s-procedure to relax (6.1.10) to a global SOS program whose solution gives a bound on the solution of (6.1.10):

$$\sup_{x_0 \in \mathcal{B}_0} \widetilde{\Phi} \leq \inf_{V \in \mathbb{R}[x]} \quad B \qquad s.t. \qquad \begin{aligned} BV - \Phi V(f) - \sum_{j=1}^{J} \sigma_j g_j &\in \Sigma_n \\ V - \epsilon - \sum_{j=1}^{J} \sigma_j g_j &\in \Sigma_n \\ \sigma_j &\in \Sigma_n \end{aligned} \tag{6.1.13}$$

Remark 2. *The argument for the derivation of the above optimization problems works equally well if the $V(f(x))/V(x)$-term in the quotient in* (6.1.8) *is replaced by its reciprocal $V(x)/V(f(x))$. Carrying out the derivation with this quotient gives a non-negativity constraint of the form $BV(f(x)) - \Phi(x)V(x) \geq 0$. While this does not make any difference for theoretical considerations, in practice, using either of these two constraint formulations often has considerable advantages over the other because there may be a significant difference between the required polynomial degree of auxiliary functions that is needed to obtain sharp bounds with the SOS-relaxed feasibility problems.*

6.2 Sharpness guarantees for optimization formulations for bounding time averages of discrete maps

In this section we prove that the solution to the optimization problem (6.1.10) is equal to the maximal geometric mean of $\Phi(x)$ under certain conditions. For this we make use of a duality result which gives the conditions under which (6.1.4) holds. We show how this result can be used to prove a similar strong duality result for (6.1.10). We recall a duality result for ergodic averages:

Lemma 4 ([Boc17])**.** *Suppose that $\mathcal{B}$ is compact, trapping and that $f, \Psi \in \mathcal{C}(\mathcal{B})$ then*

$$\begin{aligned} \sup_{x_0 \in \mathcal{B}} \limsup_{k\to\infty} \frac{1}{k} \sum_{i=0}^{k-1} \Psi(x_i) &= \limsup_{k\to\infty} \frac{1}{k} \sup_{x_0\in\mathcal{B}} \sum_{i=0}^{k-1} \Psi(x_i) \\ &= \inf_{\widehat{V}\in\mathcal{C}(\mathcal{B})} \sup_{x\in\mathcal{B}} \left[\Psi(x) + \widehat{V}\Big(f(x)\Big) - \widehat{V}(x)\right]. \end{aligned} \tag{6.2.1}$$

Proof. For the first line, the supremum and the limsup can be switched according to Proposition 2.2 in [Jen19]. The second line in the lemma comes from [FK83] and from [Boc17]. This can be seen by choosing

$$\widehat{V}_k(x) := \frac{1}{k} \sum_{l=0}^{k-1} \sum_{i=0}^{l-1} \Psi(x_i), \tag{6.2.2}$$

because then

$$\Psi(x) - \widehat{V}_k(f(x)) - \widehat{V}_k(x) = \frac{1}{k} \sum_{i=0}^{k-1} \Psi(x_i), \tag{6.2.3}$$

and so with these $\widehat{V}_k$ the supremum over $x \in \mathcal{B}$ approaches the supremum of $\overline{\Psi}$ as $k \to \infty$. □

This result states that the strong duality in (6.1.4) holds. We will now use this lemma to prove a sharpness result for the optimization problem (6.1.10) whose solution is a bound on the geometric mean of $\Phi(x)$ in definition 5. Specifically, Proposition 9 proves that under some mild assumptions on $\mathcal{B}$, $f(x)$ and $\Phi(x)$, a strong duality result analogous to (2.2.1) holds for the optimization problem of bounding geometric means of $\Phi(x)$ on bounded orbits of (1.0.4).

Proposition 9. *Suppose that $\mathcal{B}$ is compact, trapping and suppose that $f \in \mathcal{C}(\mathcal{B})$ and $\Phi \in \mathcal{C}(\mathcal{B})$ and $\Phi(x) > 0$ for all $x \in \mathcal{B}$ so that $f(x)$ and $\log(\Phi(x))$ satisfy the conditions in Lemma 4. Then the following strong duality holds:*

$$\begin{aligned}\sup_{x_0 \in \mathcal{B}} \widetilde{\Phi} = \inf_{V \in \mathcal{C}(\mathcal{B})} \quad B \quad s.t. \quad & BV(x) - \Phi(x)V(f(x)) \geq 0 \\ & V(x) > 0 \\ = \inf_{\substack{V \in \mathcal{C}(\mathcal{B}) \\ V>0}} \sup_{x \in \mathcal{B}} \left[\Phi(x)\frac{V(f(x))}{V(x)}\right] &\end{aligned} \tag{6.2.4}$$

Furthermore

$$\begin{aligned}\log\left(\sup_{x_0 \in \mathcal{B}} \widetilde{\Phi}\right) &= \sup_{x \in \mathcal{B}} \limsup_{k \to \infty} \frac{1}{k} \sum_{i=0}^{k-1} \log\Big(\Phi(x_i)\Big) \\ &= \log\left[\sup_{x \in \mathcal{B}} \limsup_{k \to \infty} \left(\prod_{i=0}^{k-1} \Phi(x_i)\right)^{1/k}\right]\end{aligned} \tag{6.2.5}$$

Proof. We first prove a weak duality result and then strengthen it to strong duality.

First note that for any $x_0 \in \mathcal{B}$ and $\widehat{V} \in \mathcal{C}(\mathcal{B})$ we have

$$\begin{aligned}\overline{\log \Phi} &= \limsup_{k\to\infty} \frac{1}{k} \sum_{i=0}^{k-1} \log(\Phi(x_i)) + \widehat{V}(f(x_i)) - \widehat{V}(x_i) \\ &= \limsup_{k\to\infty} \frac{1}{k} \sum_{i=0}^{k-1} \left[\frac{\Phi(x_i)V(f(x_i))}{V(x_i)}\right] \\ &\le \sup_{x\in\mathcal{B}} \left[\frac{\Phi(x)V(f(x))}{V(x)}\right]\end{aligned}, \tag{6.2.6}$$

where $V(x) := e^{\widehat{V}(x)} > 0$. But since $\widehat{V} \in \mathcal{C}(\mathcal{B})$ is arbitrary and, the above inequality holds for any $V \in \mathcal{C}(\mathcal{B})$ with $V(x) > 0$. Minimizing the right hand side of the inequality over such $V(x)$ and maximizing the right hand side over $x_0 \in \mathcal{B}$ gives

$$\sup_{x_0\in\mathcal{B}} \overline{\log \Phi} \le \inf_{\substack{V\in\mathcal{C}(\mathcal{B})\\ V>0}} \sup_{x\in\mathcal{B}} \left[\frac{\Phi(x)V(f(x))}{V(x)}\right]. \tag{6.2.7}$$

Furthermore, by Lemma 4

$$\limsup_{k\to\infty} \frac{1}{k} \sup_{x_0\in\mathcal{B}} \sum_{i=0}^{k-1} \log\Big(\Phi(x_i)\Big) = \sup_{x_0\in\mathcal{B}} \limsup_{k\to\infty} \frac{1}{k} \sum_{i=0}^{k-1} \log\Big(\Phi(x_i)\Big). \tag{6.2.8}$$

Since log is monotonically increasing we can move the summation and the maximization over $x_0 \in \mathcal{B}$ into the logarithm to get

$$\begin{aligned}\log(B^*) &= \sup_{x_0\in\mathcal{B}} \limsup_{k\to\infty} \frac{1}{k} \sum_{i=0}^{k-1} \log\Big(\Phi(x_i)\Big) \\ &= \sup_{x_0\in\mathcal{B}} \log \left[\limsup_{k\to\infty} \left(\prod_{i=0}^{k-1} \Phi(x_i)\right)^{1/k}\right] \\ &= \log \left[\sup_{x_0\in\mathcal{B}} \limsup_{k\to\infty} \left(\prod_{i=0}^{k-1} \Phi(x_i)\right)^{1/k}\right],\end{aligned} \tag{6.2.9}$$

and so

$$B^* = \log\left(\sup_{x\in\mathcal{B}} \widetilde{\Phi}\right) \leq \inf_{\substack{V\in\mathcal{C}(\mathcal{B})\\ V>0}} \sup_{x\in\mathcal{B}} \left[\frac{\Phi(x)V(f(x))}{V(x)}\right]. \tag{6.2.10}$$

We now want to prove the above inequality in reverse. For this we note that by Lemma 4 there exists a sequence of $\widehat{B}_k > \log(B^*)$ with $\widehat{B}_k \downarrow \log(B^*)$ as $k \to \infty$ and a corresponding sequence of feasible $\widehat{V}_k \in \mathcal{C}(\mathcal{B})$ such that

$$\widehat{B}_k - \log(\Phi(x)) - \widehat{V}_k(f(x)) + \widehat{V}_k(x) \geq 0 \qquad \forall x \in \mathcal{B}. \tag{6.2.11}$$

Therefore, since the above inequality holds for any $x \in \mathcal{B}$, if $V(x)$ is defined as above we have that for any k

$$\begin{aligned} e^{\widehat{B}_k} &\geq \sup_{x\in\mathcal{B}} \left[\frac{\Phi(x)V_k(f(x))}{V_k(x)}\right] \\ &\geq \inf_{\substack{V\in\mathcal{C}(\mathcal{B})\\ V>0}} \sup_{x\in\mathcal{B}} \left[\frac{\Phi(x)V(f(x))}{V(x)}\right]. \end{aligned} \tag{6.2.12}$$

Taking the limit as $k \to \infty$ gives

$$B^* \geq \inf_{\substack{V\in\mathcal{C}(\mathcal{B})\\ V>0}} \sup_{x\in\mathcal{B}} \left[\frac{\Phi(x)V(f(x))}{V(x)}\right], \tag{6.2.13}$$

from which the result follows. □

In order to relax the optimization problem in Proposition 9 to a SOS program, it cannot contain any positivity constraints, because SOS constraints imply non-negativity—but not positivity. Simply replacing the positivity constraint by a non-negativity constraint also does not work, because then $V(x) = B = 0$ would be optimal and satisfy the constraints. We therefore strengthen the requirement that $V(x)$ be positive by requiring it to be bounded away from zero uniformly by some

$\epsilon > 0$. In other words we constrain $V(x)$ to satisfy $V(x) - \epsilon \geq 0$ for all $x \in \mathcal{B}$. It is not obvious a priori that strengthening the bound in this way will not affect the equality in (6.2.4) because the set of positive $V(x)$ is strictly larger than the set of $V(x)$ that are bounded uniformly away from zero ϵ, for any $\epsilon > 0$. Fortunately, the next result confirms that we can replace $V(x) > 0$ with $V(x) - \epsilon \geq 0$ on the right hand side of (6.2.4) while preserving the equality, and indeed that this can be done for any choice of $\epsilon > 0$, no matter how large it is.

Proposition 10. *Suppose that $\mathcal{B}$ is compact, trapping and suppose that $f \in \mathcal{C}(\mathcal{B})$ and $\Phi \in \mathcal{C}(\mathcal{B})$ and $\Phi(x) > 0$ for all $x \in \mathcal{B}$ so that they satisfy the conditions in Proposition 9. Then for any $\epsilon > 0$ the following strong duality holds:*

$$\sup_{x_0 \in \mathcal{B}_0} \widetilde{\Phi} = \inf_{V \in \mathcal{C}(\mathcal{B})} \quad B \qquad s.t. \qquad \begin{aligned} BV(x) - \Phi(x)V(f(x)) &\geq 0 \\ V(x) - \epsilon &\geq 0 \end{aligned} \tag{6.2.14}$$

Furthermore, if $B^* := \sup_{x_0 \in \mathcal{B}_0} \widetilde{\Phi}$ then

$$\begin{aligned} \log(B^*) &= \sup_{x_0 \in \mathcal{B}} \limsup_{k \to \infty} \frac{1}{k} \sum_{i=0}^{k-1} \log\Big(\Phi(x_i)\Big) \\ &= \log\left[\sup_{x_0 \in \mathcal{B}} \limsup_{k \to \infty} \left(\prod_{i=0}^{k-1} \Phi(x_i)\right)^{1/k}\right] \end{aligned} \tag{6.2.15}$$

Proof. The outline of the proof is as follows: First we show that there exist sequences $B_k > B^*$, $\epsilon_k > 0$ and $W_k \in \mathcal{C}(\mathcal{B})$ which satisfy the constraints in (6.2.14) and $B_k \downarrow B^*$. We then show that for each such B_k there exists a $V_k \in \mathcal{C}(\mathcal{B})$, so that B_k and V_k satisfy the constraints with ϵ_k replaced by ϵ.

First note that since any $V \in \mathcal{C}(\mathcal{B})$ satisfying $V(x) - \epsilon \geq 0$ also satisfies $V(x) > 0$, we immediately get that

$$B^* = \inf_{\substack{V \in \mathcal{C}(\mathcal{B}) \\ V>0}} \sup_{x \in \mathcal{B}} \left[\Phi(x) \frac{V(f(x))}{V(x)} \right] \leq \inf_{\substack{V \in \mathcal{C}(\mathcal{B}) \\ V-\epsilon \geq 0}} \sup_{x \in \mathcal{B}} \left[\Phi(x) \frac{V(f(x))}{V(x)} \right]. \tag{6.2.16}$$

For the other direction, note that by Proposition 9 there exists a sequence $B_k > B^*$ with $B_k \downarrow B^*$ as $k \to \infty$ and a corresponding sequence of feasible $W_k \in \mathcal{C}(\mathcal{B})$ such that

$$B_k = \sup_{x \in \mathcal{B}} \left[\Phi(x) \frac{W_k(f(x))}{W_k(x)} \right]. \tag{6.2.17}$$

Furthermore as in the proof of Proposition 9 $W_k(x) = e^{\widehat{W}_k(x)}$, where $\widehat{W}_k \in \mathcal{C}(\mathcal{B})$. But since $\mathcal{B}$ is compact, $||\widehat{W}_k||_\infty < \infty$ for each k. Therefore, $||W_k||_\infty \geq e^{-||\widehat{W}_k||_\infty}$ on $\mathcal{B}$, and so letting $\epsilon_k := e^{-||\widehat{W}_k||_\infty}$ we get that

$$W_k(x) - \epsilon_k \geq 0. \tag{6.2.18}$$

Therefore,

$$B_k \geq \inf_{\substack{W \in \mathcal{C}(\mathcal{B}) \\ W-\epsilon_k \geq 0}} \sup_{x \in \mathcal{B}} \left[\Phi(x) \frac{W_k(f(x))}{W_k(x)} \right]. \tag{6.2.19}$$

But now note that if we let $V_k := \frac{\epsilon}{\epsilon_k} W_k$ then

$$B_k = \sup_{x \in \mathcal{B}} \left[\Phi(x) \frac{W_k(f(x))}{W_k(x)} \right] = \sup_{x \in \mathcal{B}} \left[\Phi(x) \frac{V_k(f(x))}{V_k(x)} \right], \tag{6.2.20}$$

and $V_k - \epsilon \geq 0$. Therefore

$$B_k \geq \inf_{\substack{V \in \mathcal{C}(\mathcal{B}) \\ V-\epsilon \geq 0}} \sup_{x \in \mathcal{B}} \left[\Phi(x) \frac{V(f(x))}{V(x)} \right], \tag{6.2.21}$$

which is true for every B_k in the sequence and so taking the limit as $k \to \infty$ gives

$$B^* \geq \inf_{\substack{V \in \mathcal{C}(\mathcal{B}) \\ V - \epsilon \geq 0}} \sup_{x \in \mathcal{B}} \left[\Phi(x) \frac{V(f(x))}{V(x)} \right]. \tag{6.2.22}$$

□

6.3 Sharpness guarantees for SOS programs for bounding time averages of discrete maps

We will now prove a sharpness result for the case when the optimization problem (6.1.10) is relaxed to a non-convex SOS-constrained optimization problem. Define the set $\mathcal{B}$ via

$$\mathcal{B} = \{x \in \mathbb{R}^n \ : \ g_1(x) \geq 0, ..., g_J(x) \geq 0, \ j = 1, ..., J\}, \tag{6.3.1}$$

and let $r := \max_j \deg(g_j)$. Let the cone of degree d, n-variate SOS polynomials be defined as in section 1.1. Then define

$$\mathbb{M}_d := \left\{ \sigma_0 + \sum_{j=1}^{J} \sigma_j g_j \ : \ \sigma_0 \in \Sigma_{d,n}, \ \sigma_j \in \Sigma_{d-r,n} \right\}. \tag{6.3.2}$$

As argued in section 2.3.2, writing the set $\mathcal{B}$ using only inequalities also includes the case when we want to include equalities in its definition. Using this, we may write down the following s-procedure-constrained non-convex SOS program:

$$\begin{aligned} \inf_{V \in \mathbb{R}[x]_d} B \ \text{ s.t. } \ & BV - \Phi V(f) \in \mathbb{M}_{d'} \\ & V - \epsilon \in \mathbb{M}_{d'}, \end{aligned} \tag{6.3.3}$$

where $\epsilon > 0$ is defined as in Proposition 10 and $d' = \deg(\Phi) + d \cdot \deg(f)$. We will now use Proposition 9 to prove the following result:

Proposition 11. *Let $\mathcal{B}$ be a compact, trapping, semialgebraic set as in (2.3.3), that satisfies the Archimedean property. Suppose that $f \in \mathbb{R}[x]$ and $\Phi \in \mathbb{R}[x]$ and $\Phi(x) > 0$ for all $x \in \mathcal{B}$. Then for any $\epsilon > 0$*

$$\sup_{x\in\mathcal{B}} \widetilde{\Phi} = \inf_{V\in\mathbb{R}[x]} B \quad s.t. \quad \begin{array}{r} BV - \Phi V(f) \in \mathbb{M}_\infty \ , \\ V - \epsilon \in \mathbb{M}_\infty \end{array} \tag{6.3.4}$$

exists and satisfies

$$\begin{aligned} \log\left(\sup_{x\in\mathcal{B}} \widetilde{\Phi}\right) &= \sup_{x_0\in\mathcal{B}} \limsup_{k\to\infty} \frac{1}{k}\sum_{i=0}^{k-1} \log\Big(\Phi(x_i)\Big) \\ &= \log\left[\sup_{x_0\in\mathcal{B}} \limsup_{k\to\infty} \left(\prod_{i=0}^{k-1} \Phi(x_i)\right)^{1/k}\right]. \end{aligned} \tag{6.3.5}$$

In other words, the SOS-relaxed optimization problem in Proposition 10 is solvable with its optimum equal to the maximum geometric mean of Φ.

Proof. From Proposition 10 we get that

$$\sup_{x\in\mathcal{B}} \widetilde{\Phi} = \inf_{V\in\mathcal{C}(\mathcal{B})} B \quad \text{s.t.} \quad \begin{array}{r} BV(x) - \Phi(x)V\Big(f(x)\Big) \geq 0 \ \cdot \\ V(x) - 2\epsilon \geq 0 \end{array} \tag{6.3.6}$$

Since $\mathcal{B}$ is compact, by the Stone-Weierstrass theorem (see e.g. [Rud+76]) $\mathbb{R}[x]$ is dense in $\mathcal{C}(\mathcal{B})$ in the sense of the $||\cdot||_\infty$ norm. Therefore, we can also maximize over

$V \in \mathbb{R}[x]$, while preserving the equality

$$\sup_{x\in\mathcal{B}} \widetilde{\Phi} = \inf_{V\in\mathbb{R}[x]} B \quad \text{s.t.} \quad \begin{aligned} BV(x) - \Phi(x)V\Big(f(x)\Big) &\geq 0 \\ V(x) - 2\epsilon &\geq 0 \end{aligned} \cdot \tag{6.3.7}$$

By the above infimum, there exists a sequence of $V_k \in \mathbb{R}[x]$ and $B_k \in \mathbb{R}$ such that $B_k \downarrow \sup_{x\in\mathcal{B}} \widetilde{\Phi}$ as $k \to \infty$, and where for each k, we have $B_k V_k(x) - \Phi(x)V_k\Big(f(x)\Big) \geq 0$ and $V_k(x) - \epsilon \geq 0$ for all $x \in \mathcal{B}$. Then let $\delta_k > 0$ be a sequence of real numbers such that $\delta_k \downarrow 0$ as $k \to \infty$. Then

$$\begin{aligned} (B_k + \delta_k)V_k(x) - \Phi(x)V_k\Big(f(x)\Big) &= B_k V_k(x) - \Phi(x)V_k\Big(f(x)\Big) + \delta_k V_k(x) \\ &> \delta_k \epsilon > 0, \end{aligned} \tag{6.3.8}$$

because

$$V_k(x) - 2\epsilon \geq 0 \quad \Longrightarrow \quad V_k(x) - \epsilon > 0. \tag{6.3.9}$$

Now, since $\mathcal{B}$ is compact, semialgebraic and satisfies the Archimedean property, Putinar's Positivestellensatz (Lemma 2) states that any positive polynomial on $\mathcal{B}$ belongs to $\mathbb{M}_\infty$ and so

$$\begin{aligned} (B_k + \delta_k)V_k(x) - \Phi(x)V_k\Big(f(x)\Big) &\in \mathbb{M}_\infty \\ V_k(x) - \epsilon &\in \mathbb{M}_\infty, \end{aligned} \tag{6.3.10}$$

for each k. Since the above holds for each $k \in \mathbb{N}$ and $(B_k + \delta_k) \to \sup_{x\in\mathcal{B}} \widetilde{\Phi}$, the result follows. □

In practice, we solve polynomial optimization problems where the maximum degrees of the tuneable polynomials V and σ_j are fixed, so that $V, \sigma_j \in \mathbb{M}_d$ for some $d < \infty$. When the maximum degrees of V and σ_j are fixed then Proposition 11 does not guarantee that the solution to (6.3.3) may not be strictly larger than the left hand side of (6.3.5). However, in practice we observe that we often do seem to have equality between (6.3.3) and the left hand side of (6.3.5). If we do not have equality, then Proposition 11 tells us that we can solve the polynomial optimization in (6.3.3) with higher-degree V and σ_j for a chance to obtain a better bound on $\sup_{x\in\mathcal{B}} \widetilde{\Phi}$, and that this bound will converge to $\sup_{x\in\mathcal{B}} \widetilde{\Phi}$ as we include polynomials of higher and higher degree in the optimization.

Chapter 7

Bounding LEs of discrete maps

In this section, the problem of bounding the maximal LE of a discrete-time system is studied. While the methods derived here share many features with the continuous case, as argued in section 6.1.2 the fact that one cannot define quantities in terms of integrals or derivatives means that the approach has to be modified somewhat. LEs of discrete systems are defined via (1.0.6). Our strategy then is to apply the methods developed in section 6.1.2 for bounding geometric means to the problem of bounding LEs and to then solve numerically tractable SOS-relaxed feasibility problems to bound the maximal LE over the bounded trajectories of (1.0.4).

7.1 Optimization formulations for bounding LEs of discrete maps

In order to derive an optimization problem for bounding the maximal LE of a discrete map of the form (1.0.4), we begin with the definition of its LE (1.0.6) and build on (6.0.1) to derive an expression for the LE in terms of the logarithm of a geometric mean, rather than the arithmetic mean of a logarithm. This will allow us to derive a SOS-relaxable optimization problem to bound the maximal LE of (1.0.4).

By definition (1.0.6) we have

$$\begin{aligned}\mu &= \limsup_{k\to\infty} \frac{1}{k}\Big(\log|y_k| - \log|y_0|\Big) \\ &= \limsup_{k\to\infty} \frac{1}{k} \sum_{i=1}^{k} (\log|y_i| - \log|y_{i-1}|) \\ &= \limsup_{k\to\infty} \frac{1}{k} \log\left(\prod_{i=1}^{k} \frac{|y_i|}{|y_{i-1}|}\right) \\ &= \limsup_{k\to\infty} \frac{1}{k} \log\left(\prod_{i=0}^{k-1} \frac{|Df(x_i)y_i|}{|y_i|}\right) \\ &= \limsup_{k\to\infty} \frac{1}{k} \log\left(\prod_{i=0}^{k-1} |Df(x_i)z_i|\right),\end{aligned} \tag{7.1.1}$$

where $z := \frac{y}{|y|}$. Therefore, μ depends on $2n$ variables, which evolve according to the map

$$\begin{bmatrix} x_{k+1} \\ z_{k+1} \end{bmatrix} = \begin{bmatrix} f(x_k) \\ \ell(x_k, z_k) \end{bmatrix}, \qquad (x_k, z_k) \in \mathbb{R}^n \times \mathbb{S}^{n-1}, \tag{7.1.2}$$

where

$$z_{k+1} = \ell(x_k, z_k) = \frac{Df(x_k)z_k}{|Df(x_k)z_k|}. \tag{7.1.3}$$

Letting

$$\begin{aligned}\Phi(x, z) &= z^{\mathsf{T}} Df(x)^{\mathsf{T}} Df(x) z \\ &= |Df(x)z|^2,\end{aligned} \tag{7.1.4}$$

we can rewrite the LE as

$$\begin{aligned}\mu &= \limsup_{k\to\infty} \frac{1}{2k} \log\left(\prod_{i=0}^{k-1} |Df(x_i)z_i|^2\right) \\ &= \limsup_{k\to\infty} \frac{1}{2k} \log\left(\prod_{i=0}^{k-1} \Phi(x_i, y_i)\right).\end{aligned} \tag{7.1.5}$$

But because log is monotonically increasing and smooth away from zero we have

$$\mu := \frac{1}{2} \log\left[\limsup_{k\to\infty} \left(\prod_{i=0}^{k-1} \Phi(x_i, z_i)\right)^{1/k}\right]. \tag{7.1.6}$$

The argument of the logarithm is therefore exactly a geometric mean of the type for which we have an optimization framework that can bound it. Letting $\Phi(x, z)$ be defined as in (7.1.4), we just have to bound the geometric mean of $\Phi(x, z)$ along bounded trajectories of (7.1.2) to bound μ. We can do this via (6.1.10), if we expand our phase space to be $\mathcal{B} \times \mathbb{S}^{n-1}$ and our map to include the normalized lienarized dynamics $\ell(x, z)$ (7.1.3). Specifically, we need to solve following optimization problem:

$$\begin{aligned}\inf_{V\in\mathcal{C}(\mathcal{B},\mathbb{S}^{n-1})} \quad B \quad \text{s.t.} \quad BV(x,z) - \Phi(x,z)V\Big(f(x), \ell(x,z)\Big) &\geq 0 \\ V(x,z) &> 0.\end{aligned} \tag{7.1.7}$$

If $\mathcal{B}$ is compact, and if $f(x)$ and $\Phi(x, z)$ satisfy the conditions in Lemma 4 then we can directly apply the sharpness result in Proposition 9 to (7.1.7) and so in that case (7.1.7) is guaranteed to give a sharp bound on the maximal LE of (1.0.4) over trajectories in $\mathcal{B}$. The conditions in Lemma 4 are met if $f \in \mathcal{C}^1(\mathcal{B})$ and $Df(x)$ is invertible. This technical condition on $Df(x)$ ensures that $Df(x)^\mathsf{T} Df(x)$ is positive definite so that $\Phi(x)z$ is strictly positive on $\mathcal{B} \times \mathbb{S}^{n-1}$. We therefore have the following sharpness result for (7.1.7):

Proposition 12. *Let* $x_{k+1} = f(x_k)$ *with* $f \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R}^n)$ *and* $\Phi(x, z) := z^\mathsf{T} Df(x)^\mathsf{T} Df(x) z$.

1. *Let* $\mathcal{B}_0 \subset \mathbb{R}^n$ *and* $\mathcal{B} \subset \mathbb{R}^n$ *be such that all trajectories* x_k *with* $x_0 \in \mathcal{B}_0$ *are bounded uniformly in* $t \geq 0$ *and eventually remain in* $\mathcal{B}$*. Then the maximal LE* $\mu^*_{\mathcal{B}_0}$ *among trajectories with initial conditions in* $\mathcal{B}_0$ *is bounded above by*

$$\exp(2\mu^*_{\mathcal{B}_0}) \leq \inf_{V \in \mathcal{C}(\mathcal{B},\mathbb{S}^{n-1})} \quad B \quad s.t. \quad BV(x,z) - \Phi(x,z)V\Big(f(x), \ell(x,z)\Big) \geq 0$$

$$V(x,z) > 0, \tag{7.1.8}$$

2. *Let* $\mathcal{B}_0 = \mathcal{B}$ *be a compact set that is trapping (i.e., forward invariant) and suppose also that* $Df(x)$ *is invertible for all* $x \in \mathcal{B}$ *. Then,*

$$\exp(2\mu^*_{\mathcal{B}}) = \inf_{V \in \mathcal{C}(\mathcal{B},\mathbb{S}^{n-1})} \quad B \quad s.t. \quad BV(x,z) - \Phi(x,z)V\Big(f(x), \ell(x,z)\Big) \geq 0$$

$$V(x,z) > 0, \tag{7.1.9}$$

7.2 Practical SOS implementation of the optimization problem

We would now like to be able to relax (7.1.7) to a SOS program and enforce a pointwise non negativity constraint with a global SOS constraint by using the s-procedure, so that we may obtain duality results as in Proposition 12. We need to use an s-procedure to enforce the $|z| = 1$ constraint with two barrier functions $g_{J+1} := |z|^2 - 1$ and $g_0 := 1 - |z|^2$, which appear in the domain specification (6.3.1) as the inequalities $|z|^2 - 1 \leq 0$ and $1 - |z|^2 \geq 0$, each of which have SOS-polynomial-constrained weights. Note, however, that we could also do this as we did in section 3.4 with an equality constraint $0 = 1 - |z|^2$ that has a non-SOS polynomial weight.

Relaxing (7.1.7) to a SOS program is still not possible right away, because $\ell(x,z)$ is not a polynomial: the main challenge comes from the $V(f(x),\ell(x,z))$ term in the non-negativity constraint

$$BV(x,z) - \Phi(x,z)V\Big(f(x),\ell(x,z)\Big) \geq 0, \tag{7.2.1}$$

because it is evaluated at $\ell(x,z)$ and will therefore contain rational terms of the form

$$\ell(x,z)^{\alpha} = \frac{(Df(x)z)^{\alpha}}{\Phi(x,z)^{|\alpha|_1/2}}, \tag{7.2.2}$$

where α is a multiindex encoding the degrees of the components of z in each term of $V(x,z)$. The square root in the denominator of $\ell(x,z)$ might pose a serious problem, were it not for the fact that $\ell(x,z)$ and $\Phi(x,z)$ possess the symmetry $z \mapsto -z$, which allows us to consider only $V \in \mathbb{R}[x,z]$ such that $V(x,z)$ contains only even-powered terms in z (i.e. terms z^{α} such that $|\alpha|_1$ is even). This symmetry and some of its implications will be explained in chapter 8 below. This means that $\Phi(x,z)^{|\alpha|_1/2}$ is polynomial in x and z for each term in $V(x,z)$. In order to make the whole SOS-relaxed non-negativity constraint a polynomial, we therefore just have to multiply it through by

$$\Phi(x,z)^{d_z/2-1}, \tag{7.2.3}$$

where d_z is the maximum degree of $V(x,z)$ in z, to give

$$B\Phi(x,z)^{d_z/2-1}V(x,z) - \Phi(x,z)^{d_z/2}V\Big(f(x),\ell(x,z)\Big) \geq 0. \tag{7.2.4}$$

In practice we will also need to make use of the s-procedure to enforce non-negativity on a semialgebraic set via a global SOS constraint. This gives the non-convex SOS program

$$\exp(2\mu^*_{\mathcal{B}_0}) \leq \inf_{V,\sigma_j \in \mathbb{R}[x,z]} \quad B \quad s.t. \quad B\Phi^{d_z/2-1}V - \Phi^{d_z/2}V(f,\ell) - \sum_{j=0}^{J+1} \sigma_j g_j \in \Sigma_{2n}$$
$$V - \epsilon - \sum_{j=0}^{J+1} \sigma_j g_j \in \Sigma_{2n}$$
$$\sigma_j \in \Sigma_{2n}, \tag{7.2.5}$$

with g_0 and g_{J+1} defined as above. Since B multiplies the coefficients of V in the first constraint, this is not a convex SOS program and so we have to use a bisection method, as explained in section 6.1.2, to solve a sequence of convex feasibility problems of the form:

$$\text{find } V \in \mathbb{R}[x,z] \text{ and } B \in \mathbb{R} \quad s.t. \quad B\Phi^{d_z/2-1}V - \Phi^{d_z/2}V(f,\ell) - \sum_{j=0}^{J+1} \sigma_j g_j \in \Sigma_{2n}$$
$$V - \epsilon - \sum_{j=0}^{J+1} \sigma_j g_j \in \Sigma_{2n}$$
$$\sigma_j \in \Sigma_{2n}, \tag{7.2.6}$$

so that for any feasible $V \in \mathbb{R}[x,z]$ and $B \in \mathbb{R}$ the maximal LE of (1.0.4) on $\mathcal{B}$ satisfies $\mu^*_{\mathcal{B}} \leq \frac{1}{2}\log(B)$. Nonetheless, the sharpness result in Proposition 13 below guarantees that (7.2.5) has an optimum and that the bisection method therefore converges to that optimum.

Proposition 13. *Let $\mathcal{B}$ be a compact semialgebraic set as in* (2.3.3)*, that satisfies the Archimedean property and is trapping. Suppose that $f \in \mathbb{R}[x]$ and $\ell(x,z)$ is as defined in* (7.1.3)*. Let $\Phi(x,z) = z^{\mathsf{T}} Df^{\mathsf{T}}(x) Df(x) z$ and suppose that $Df(x)$ is invertible for* *all $x \in \mathcal{B}$. Then for any $\epsilon > 0$*

$$B^* := \inf_{V \in \mathbb{R}[x,z]} \quad B \qquad s.t. \qquad B\Phi^{d_z/2-1}V - \Phi^{d_z/2}V\Big(f, \ell\Big) \in \mathbb{M}_\infty \ , \qquad V - \epsilon \in \mathbb{M}_\infty \tag{7.2.7}$$

exists and satisfies

$$\begin{aligned} \log(B^*) &= \sup_{\substack{x_0 \in \mathcal{B} \\ z_0 \in \mathbb{S}^{n-1}}} \limsup_{k\to\infty} \frac{1}{k} \sum_{i=0}^{k-1} \log\Big(\Phi(x_i, z_i)\Big) \\ &= \log\left(\sup_{x_0 \in \mathcal{B}} \limsup_{k\to\infty} \prod_{i=0}^{k-1} \Phi(x_i, z_i)\right) \\ &= 2\mu^*_{\mathcal{B}}. \end{aligned} \tag{7.2.8}$$

The proof of Proposition 13 is a direct application of Proposition 11, where we note that $B\Phi^{d_z/2-1}V - \Phi^{d_z/2}V\Big(f, \ell\Big) \geq 0$ if and only if $BV - \Phi V\Big(f, \ell\Big) \geq 0$ and so the proof of Proposition 11 directly translates to Proposition 13 by using the strong duality result in Proposition 9.

7.3 Explicit construction of the SOS constraint for bounding LEs

The parser YALMIP, which we use to solve (7.2.6) to obtain bounds on LEs does not allow us to construct the constraints directly, because it does not enable multiplying a rational polynomial by its denominator to get a polynomial. Therefore, we will briefly describe a method for constructing the constraint in (7.2.6) so that it can be passed to YALMIP and solved using Mosek. If we let

$$V(x,z) = \sum_{|\alpha|_1=0, |\beta|_1=0}^{d_x, d_z} c(\alpha,\beta) x^\beta z^\alpha, \tag{7.3.1}$$

where d_x and d_z are the maximal degrees of V in x and z, respectively, and by assumption $|\alpha|_1$ is always even, then

$$\begin{aligned} &B\Phi(x,z)^{d_z/2-1} V(x,z) - \Phi(x,z)^{d_z/2} V\Big(f(x), \ell(x,z)\Big) \\ &= \sum_{|\alpha|_1=0, |\beta|_1=0}^{d_z, d_x} c(\alpha,\beta) \Big[\Phi(x,z)^{d_z/2-1} x^\beta z^\alpha - f(x)^\beta \Big(Df(x)z\Big)^\alpha \Phi(x,z)^{d_z/2-|\alpha|_1/2} \Big]. \end{aligned} \tag{7.3.2}$$

The pseudo code in Algorithm 1 constructs the term $\Phi(x,z)^{d_z/2} V\Big(f(x), \ell(x,z)\Big)$:

Algorithm 1 Constructing $\Phi(x,z)^{d_z/2} V\Big(f(x), \ell(x,z)\Big)$

Require: $V(x,z) = \sum_{|\alpha|_1=0, |\beta|_1=0}^{d_x, d_z} c(\alpha,\beta) x^\beta z^\alpha$
  $\tilde{V} = V\Big(f(x), Df(x)z\Big)$
  $b =$ monomialbasis$(\tilde{V})$
  $c =$ coefficients$(\tilde{V})$
  for $i = 1$: cardinality(b) **do**
    $|\alpha|_1 =$ total degree$\Big(b(i), z\Big)$
    $b(i) = b(i)\Phi(x,z)^{d_z/2-|\alpha|_1/2}$
  end for
  return OUTPUT = dot (c, b)

Chapter 8

Symmetries of discrete maps

Chapter 4 studied how symmetries in the ODE lead to symmetries in the constraints in the optimization problems used to bound LEs. This allowed us to reduce the size and increase the stability of the SOS programs that we had to solve to obtain sharp bounds on LEs. Since the SOS formulations for bounding the LE of discrete maps requires one to solve a sequence of SOS feasibility problems whose solution approaches the maximal LE of a map exactly when the problem becomes infeasible, being able to solve these feasibility problems with increased numerical stability is essential. We therefore would like to be able to prove analogous results about symmetries of discrete maps as we did in chapter 4 about symmetries of continuous ODE systems.

Fortunately, analogous symmetry results do also hold in the discrete-time case. Symmetries in the vector field of the map induce symmetries in the map describing the tangent dynamics of an orbit and in the scalar quantity whose geometric mean gives the LE. This is explained in section 8.1. Section 8.2 argues that Propositions 20 and 21 imply that if the map and the domain are symmetric under a group of orthogonal symmetries $\mathcal{G}$, then constraining the space of auxiliary functions over which the optimization for bounding the maximal LE is carried out to functions which are also symmetric under $\mathcal{G}$ does not alter the value of the solution to the problem. In section 8.3 the symmetry results of the two sections before are extended to more general formulations of the optimization problems for bounding LEs. We discuss how one can make use of non-orthogonal symmetries in the optimization formulation for bounding the maximal LE, by either using a specialized norm to derive non-negativity constraints, or by making a suitable change of variables that allows the results for orthogonal symmetries to be applied.

8.1 Symmetries of the (x, z) dynamics for discrete maps

A discrete dynamical system (1.0.4) is said to be symmetric under an invertible transformation if solutions are mapped to solutions. For a system $x_{k+1} = f(x_k)$ on $\mathbb{R}^n$, a transformation $\Lambda : \mathbb{R}^n \to \mathbb{R}^n$ is a symmetry of the system if and only if f is *equivariant* under Λ, meaning $f(\Lambda x) = \Lambda f(x)$ for all $x \in \mathbb{R}^n$. On the other hand, a function Φ is *invariant* under Λ if $\Phi(\Lambda x) = \Phi(x)$ for all $x \in \mathbb{R}^n$. We seek the symmetry group under which the right-hand side of the (x, z) system (7.1.2) is equivariant, and under which the function $\Phi(x, z) = z^\mathsf{T} Df(x)^\mathsf{T} Df(x) z$ is invariant.

The transformation $z \mapsto -z$ is a symmetry both of the (x, z) map, the normalized linearized map $\ell(x, z)$ (7.1.3) and of the function $\Phi(x, z) = z^\mathsf{T} Df(x)^\mathsf{T} Df(x) z$ whose time average is to be bounded, reflecting the fact that the growth or decay rate of a tangent vector is unaffected by reversing its direction. The equivariance of the right-hand side of (7.1.2) under this transformation follows from the oddness $\ell(x, -z) = -\ell(x, z)$, where $\ell(x, z)$ is defined as in (7.1.3), and the invariance $\Phi(x, -z) = \Phi(x, z)$ is clear.

If the x map alone has symmetries, these can induce additional symmetries in the (x, z) dynamics. Suppose that the x map is symmetric under $x \mapsto \Lambda x$ for some

$\Lambda \in GL(n)$, meaning that $f(\Lambda x) = \Lambda f(x)$ for all $x \in \mathbb{R}^n$. The coupled (x, z) system (7.1.2) does not share this symmetry because the $\ell(x, z)$ function is not equivariant under the transformation of x alone; z must be transformed also. It is natural to consider transforming x and z in the same way by $(x, z) \mapsto (\Lambda x, \Lambda z)$, but this maps $\mathbb{R}^n \times \mathbb{S}^{n-1}$ to itself if and only if Λ is an orthogonal transformation, since otherwise $\Lambda z \notin \mathbb{S}^{n-1}$ for some z. Assuming orthogonality of Λ as needed for $(x, z) \mapsto (\Lambda x, \Lambda z)$ to be well defined on $\mathbb{R}^n \times \mathbb{S}^{n-1}$, this transformation is a symmetry of both the (x, z) map and of the function $\Phi(x, z) = z^{\mathsf{T}} Df(x)^{\mathsf{T}} Df(x) z$. This result is stated by the following proposition, in which $O(n)$ denotes the group of orthogonal linear transformations on $\mathbb{R}^n$.

Proposition 14. *Let $f \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R}^n)$. Define $\ell : \mathbb{R}^n \times \mathbb{S}^{n-1} \to \mathbb{R}^n$ by $\ell(x, z) = \frac{Df(x)z}{|Df(x)z|}$, and define $\Phi : \mathbb{R}^n \times \mathbb{S}^{n-1} \to \mathbb{R}$ by $\Phi(x, z) = z^{\mathsf{T}} Df(x)^{\mathsf{T}} Df(x) z$. Let $\Lambda \in O(n)$. If $f(\Lambda x) = \Lambda f(x)$ for all $x \in \mathbb{R}^n$, then $\ell(\Lambda x, \Lambda z) = \Lambda \ell(x, z)$ and $\Phi(\Lambda x, \Lambda z) = \Phi(x, z)$ for all $(x, z) \in \mathbb{R}^n \times \mathbb{S}^{n-1}$.*

The proof of Proposition 14 is a direct calculation showing the claimed equivariance of ℓ and invariance of Φ. We omit the calculation because it is a special case of the proof of Proposition 17 below. Section 8.2 describes how the symmetries guaranteed by Proposition 14 lead to useful symmetries in the SOS programs that bound LEs from above.

8.2 Symmetries of the optimization problems for discrete maps

If a map, its domain, and a discrete-time-averaged (in the sense of a geometric mean) quantity share a group of symmetries, then the solution to the optimization problem (7.1.7) is unchanged if the auxiliary function V is further constrained to be invariant under the same symmetries. In particular, any V that yields an upper bound on a geometric mean can be used to construct a symmetrized V yielding the same upper bound. Proposition 20 in the Appendix proves the result for any compact symmetry group $\mathcal{G} \subset GL(n)$.

In the context of the upper bounds on LEs provided by the solution of (7.1.7), the pertinent symmetries are described in the preceding section 8.1. Proposition 20 guarantees that corresponding symmetries can be imposed on $V(x, z)$ in the right-hand minimization of (7.1.7) without changing the resulting upper bound. To state these results precisely in the following proposition, we define for each group $\mathcal{G} \subset O(n)$ the group $\mathcal{G}' \subset O(2n)$ via (4.1.1) as in section 4.2.

Proposition 15. *Let $f \in \mathcal{C}^1(\mathcal{B}, \mathbb{R}^n)$. If f is $\mathcal{G}$-equivariant and $\mathcal{B}$ is $\mathcal{G}$-invariant for some $\mathcal{G} \subset O(n)$, then the infimum over $V \in \mathcal{C}(\mathcal{B} \times \mathbb{S}^{n-1})$ in* (7.1.7) *with $\Phi(x, z) = z^{\mathsf{T}} Df^{\mathsf{T}}(x) Df(x) z$ is unchanged if V is constrained to be $\mathcal{G}'$-invariant for $\mathcal{G}'$ defined by* (4.1.1).

Proof. Invariance under $\mathcal{G}'$ is equivalent to invariance under $(x, z) \mapsto (x, -z)$ and under $(x, z) \mapsto (\Lambda x, \Lambda z)$ for all $\Lambda \in \mathcal{G}$. The invariance under negation of z is proved by applying Proposition 20 in the context of the state space $\mathcal{B} \times \mathbb{S}^{n-1}$, to the map (7.1.2), the function $\Phi(x, z) = z^{\mathsf{T}} Df(x)^{\mathsf{T}} Df(x) z$, and the transformation $(x, z) \mapsto (x, -z)$. The assumptions of Proposition 20 are met because, with respect to the transformation, the state space and Φ are invariant, and the right-hand side of the map is equivariant, as described in the second paragraph of section 8.1 above. The invariance under $(x, z) \mapsto (\Lambda x, \Lambda z)$ for each $\Lambda \in \mathcal{G}$ is proved by applying Proposition 20 in the context of this transformation, with the same state space, map, and Φ as in the preceding sentences. To check the assumptions of Proposition 20, first note that $\mathcal{B} \times \mathbb{S}^{n-1}$ is invariant under $(x, z) \mapsto (\Lambda x, \Lambda z)$ since $\mathcal{B}$ is $\mathcal{G}$-invariant by assumption, and the sphere $\mathbb{S}^{n-1}$ is invariant under every orthogonal transformation. Second, note that invariance of Φ and equivariance of the right-hand side of (7.1.2) are given by Proposition 14 since f is $\mathcal{G}$-equivariant by assumption. □

Remark 3. *The group $\mathcal{G}'$ always contains the identity and the transformation $(x, z) \mapsto (x, -z)$. Therefore, even in the absence of any other symmetries, $V(x, z)$ can always be taken to be even in z without worsening the resulting upper bounds on LEs. This fact is essential for being able to construct constraints that are polynomial in their arguments in section 7.3.*

Proposition 15, which allows symmetries of the dynamics to be imposed on V in the right-hand minimization of (7.1.7), has an analogue for the minimization's SOS relaxation (7.2.5). This is stated by Proposition 16 below, whose assumptions require that the semialgebraic specification of the $\mathcal{G}$-invariant set $\mathcal{B}$ uses only $\mathcal{G}$-invariant polynomials, and whose conclusions guarantee invariance not only of V but also of the polynomial weights σ_j.

In short, from any V that satisfies the inequality constraints in (7.1.7), a symmetrized V can be constructed by averaging V over the group orbit of Λ, and this symmetrized V will also satisfy the constraint in (7.1.7). In SOS relaxed minimizations the same symmetry can be imposed on V without changing the infima in (7.2.5), provided that all functions in the semialgebraic specification of the domain $\mathcal{B}$ share the symmetry also.

Proposition 16. *Let $f \in \mathbb{R}^n[x]$. Let $\mathcal{B} \subset \mathbb{R}^n$ be either all of $\mathbb{R}^n$ or specified by a finite number of polynomial inequalities, $g_j(x) \geq 0$ for $j \in \{1, \dots, J\}$. If f is $\mathcal{G}$-equivariant and all g_j are $\mathcal{G}$-invariant for some $\mathcal{G} \subset O(n)$, then the solution to* (7.2.5) *over $V, \sigma_j \in \mathbb{R}[x, z]$ (or in any other vector space of polynomials) is unchanged if V and all σ_j are additionally constrained to be $\mathcal{G}'$-invariant for $\mathcal{G}'$ defined by* (4.1.1).

Proposition 16 has practical implications for the SOS programs whose solutions give upper bounds on the maximal LE for the reasons discussed after Proposition 7.

When a SOS program is translated into a semidefinite program, and each SOS constraint is replaced by semidefiniteness of a matrix, invariances of a polynomial that must be SOS allow block diagonal structure to be imposed on the corresponding matrix. The semidefiniteness constraint thus decomposes into semidefiniteness of each block, allowing for faster and more precise numerical solution.

8.3 Optimization formulation preserving non-orthogonal symmetries for discrete maps

This subsection describes two approaches with which the z dynamics (7.1.3) and the quantity $\Phi(x, z)$ can be modified so that non-orthogonal symmetries can be exploited to bound LEs. These results parallel those in section 4.3 for continuous time systems where the same symmetries may be used to obtain simplified SOS programs. The results of sections 8.1 and 8.2 above apply only to symmetry groups $\mathcal{G}$ of the x dynamics that are subgroups of $O(n)$. However, symmetries which are similarity transformation of orthogonal transformations can still be used in this context, if the original system is transformed using an appropriate change of variables. We will explain how this can be done in detail later in this section. The other technique for using non-orthogonal symmetries is to derive the scalar quantity $\Phi(x, z)$ and the map $\ell(x, z)$ using the norm $|\cdot|_{\mathsf{M}}$ in (4.3.1) and to use these in the optimization problems for bounding LEs. This generalization is not used in the computational examples of chapter 9 below, where all symmetries are orthogonal, but it may be of use for future applications.

In the definition (1.0.6) of the LE, the Euclidean norm can be equivalently replaced by any other norm on $\mathbb{R}^n$ such as the weighted norm $|y|_{\mathsf{M}} = |y^{\mathsf{T}}\mathsf{M}y|$ in (4.3.1). Starting from the expression $\mu(x_0, y_0) = \limsup_{k\to\infty} \frac{1}{k} \log\left(\frac{|y_k|_{\mathsf{M}}}{|y_0|_{\mathsf{M}}}\right)$ and defining $z := y/|y|_{\mathsf{M}}$, the arguments used in the beginning of chapter 7 give $\mu(x_0, z_0) = \frac{1}{2}\log\widetilde{\Phi}_{\mathsf{M}}$, with

$$\Phi_{\mathsf{M}}(x,z) := z^{\mathsf{T}} Df(x)^{\mathsf{T}} \mathsf{M} Df(x) z, \tag{8.3.1}$$

where the phase space of z is now the set $\mathbb{S}^{n-1}_{\mathsf{M}}$ of vectors in $\mathbb{R}^n$ with $|z|_{\mathsf{M}} = 1$. Since $z_k = y_k/|y_k|_{\mathsf{M}}$, (1.0.5) implies that $z_{k+1} = \ell_{\mathsf{M}}(x_k, z_k)$ where

$$\ell_{\mathsf{M}}(x,z) := \frac{Df(x)z}{|Df(x)z|_{\mathsf{M}}}. \tag{8.3.2}$$

The upper bounds on maximal LEs described in chapter 7 still hold when Φ and ℓ are replaced by Φ_{M} and ℓ_{M} for any positive definite symmetric M as in section 4.3. Choosing M as the identity recovers the original formulas.

We choose the M matrix as an integral over compact $\mathcal{G}$ with respect to its Haar probability measure and define it via (4.3.4) this M has the property that $\Lambda^{\mathsf{T}}\mathsf{M}\Lambda = \mathsf{M}$ for all $\Lambda \in \mathcal{G}$ (cf. the proof of Proposition 8 in section 4.3). With this choice of M, the $\mathcal{G}$-equivariance of f induces corresponding symmetries in Φ_{M} and ℓ_{M}:

Proposition 17. *Let $f \in \mathcal{C}^1(\mathbb{R}^n, \mathbb{R}^n)$. Define $\ell_{\mathsf{M}} : \mathbb{R}^n \times \mathbb{S}^{n-1}_{\mathsf{M}} \to \mathbb{R}^n$ by $\ell_{\mathsf{M}}(x,z) = \frac{Df(x)z}{|Df(x)z|_{\mathsf{M}}}$, and define $\Phi_{\mathsf{M}} : \mathbb{R}^n \times \mathbb{S}^{n-1}_{\mathsf{M}} \to \mathbb{R}$ by $\Phi_{\mathsf{M}}(x,z) = z^{\mathsf{T}} Df(x)^{\mathsf{T}} \mathsf{M} Df(x) z$. Let $\Lambda \in \mathcal{G} \subset GL(n)$ with $\mathcal{G}$ compact. If $f(\Lambda x) = \Lambda f(x)$ for all $x \in \mathbb{R}^n$, then $\ell_{\mathsf{M}}(\Lambda x, \Lambda z) = \Lambda \ell_{\mathsf{M}}(x,z)$ and $\Phi_{\mathsf{M}}(\Lambda x, \Lambda z) = \Phi_{\mathsf{M}}(x,z)$ for all $(x,z) \in \mathbb{R}^n \times \mathbb{S}^{n-1}_{\mathsf{M}}$.*

Proof. Let $\Lambda \in \mathcal{G}$. As in the proof of Proposition 8 it follows from the right invariance of the Haar probability measure $m(\Lambda)$ on $\mathcal{G}$ that $\Lambda^{\mathsf{T}}\mathsf{M}\Lambda = \mathsf{M}$. This relation and the relation $Df(\Lambda x) = \Lambda Df(x)\Lambda^{-1}$, which follows from the $\mathcal{G}$-equivariance of f, together give the invariance

$$\begin{aligned} \Phi_{\mathsf{M}}(\Lambda x, \Lambda z) &= (\Lambda z)^{\mathsf{T}} Df(\Lambda x)^{\mathsf{T}} \mathsf{M} Df(\Lambda x) \Lambda z \\ &= z^{\mathsf{T}} \Lambda^{\mathsf{T}} (\Lambda^{-1})^{\mathsf{T}} Df(x) \Lambda^{\mathsf{T}} \mathsf{M} \Lambda Df(x) \Lambda^{-1} \Lambda z \\ &= z^{\mathsf{T}} Df(x)^{\mathsf{T}} \mathsf{M} Df(x) z = \Phi_{\mathsf{M}}(x, z). \end{aligned} \tag{8.3.3}$$

For the ℓ_{M} equivariance note that $\ell_{\mathsf{M}}(x,z) = \frac{Df(x)z}{\sqrt{\Phi(x,z)}}$, from which it follows that

$$\ell_{\mathsf{M}}(\Lambda x, \Lambda z) = \frac{Df(\Lambda x)\Lambda z}{\sqrt{\Phi(\Lambda x, \Lambda z)}} = \frac{\Lambda Df(x) \Lambda^{-1} \Lambda z}{\sqrt{\Phi(x,z)}} = \Lambda \ell_{\mathsf{M}}(x,z). \tag{8.3.4}$$

□

The upper bounds on LEs given by the right-hand minimization in (7.1.7) or its SOS relaxation in (7.2.5) hold also with Φ_{M} and ℓ_{M} in place of Φ and ℓ. In these modified minimization problems, the minima are unchanged if the tunable functions are constrained to be $\mathcal{G}'$-invariant. These generalizations follow from Proposition 17 combined with Proposition 20 or Proposition 21; their proofs are analogous to those of Propositions 15 and 16.

Another way in which we can make use of non-orthogonal symmetries is if they belong to a group $\mathcal{G}_{\mathsf{A}}$ defined in (4.3.9) if we make the change of variables $\tilde{x} = \mathsf{A}x$ and define $\tilde{f}(\tilde{x}) := \mathsf{A} f(\mathsf{A}^{-1} x)$. The same reasoning as in section 4.3 implies that $\tilde{f}(\tilde{x})$ is equivariant under $\tilde{x} \mapsto \tilde{\Lambda}\tilde{x}$ for $\tilde{\Lambda} \in O(n)$, which according to Proposition 14 implies that $\tilde{\Phi}(\tilde{x}, z) := z^{\mathsf{T}} D\tilde{f}(\tilde{x})^{\mathsf{T}} D\tilde{f}(\tilde{x}) z$ is invariant, and

$$\tilde{\ell}(\tilde{x}, z) := \frac{D\tilde{f}(\tilde{x})z}{|D\tilde{f}(\tilde{x})z|} \tag{8.3.5}$$

is equivariant under $(\tilde{x}, z) \mapsto (\tilde{\Lambda}\tilde{x}, \tilde{\Lambda}z)$. Therefore, Propositions 15 and 16 apply to the system defined by $\tilde{f}(\tilde{x})$, which has the same LEs as the untransformed system.

Chapter 9

Examples: bounding LEs of discrete systems

We illustrate the methods for bounding the maximal LE of (1.0.4) with some examples. The first example is the logistic map, which is a classic one-dimensional example of a chaotic map. The SOS program for bounding its maximal LE is simple enough that it can be solved analytically, proving that the maximal LE is attained at the origin of the map. The example in section 9.2 is a map which was chosen because it does not maximize its LE on a fixed point. The example serves to illustrate that the method works for such maps, exactly like the theory predicts. The final example is the Hénon map, which is a chaotic system for which we compute its maximal LE numerically and where we also prove analytic upper bounds on its maximal LE, which are valid for all parameter values.

9.1 Logistic map

Consider the logistic map

$$x_{k+1} = 4x_k(1 - x_k). \tag{9.1.1}$$

We want to bound its maximal LE using an optimization problem of the form:

$$\begin{aligned} \inf_{V \in \mathcal{C}(\mathbb{R}\times\mathbb{S}^0)} \quad B \;\text{ s.t. }\; BV(x,z) - \Phi(x,z)V\Big(f(x), \ell(x,z)\Big) &\geq 0 \\ V(x,z) - \epsilon &\geq 0. \end{aligned} \tag{9.1.2}$$

This SOS program can be solved when B corresponds to the maximal LE, $\mu = \log(4)$ of (9.1.1) if an auxiliary function of the form

$$V = 1 + z^2, \tag{9.1.3}$$

is used so that $V(x,z) > 0$. However, this requires the use of suitable domain constraints, which are furbished by the s-procedure. We enforce the $|z| = 1$ constraint via the barrier function $h_0 = 1 - z^2$. In addition, we also enforce $x \in [0,1]$ via the barrier function $g_1 = x(1-x)$. This constraint on the x-domain is required for the auxiliary function $V = 1 + z^2$. Let us determine the choice of polynomials in the s-procedure. For the auxiliary function (9.1.3) the constraint

$$BV(x,z) - \Phi(x,z)V\Big(f(x), \ell(x,z)) - \rho_0 h_0 - \sigma_1 g_1 \geq 0 \tag{9.1.4}$$

becomes

$$S(x,z) := B(1+z^2) - 32z^2(1-2x)^2 - \rho_0(x,z)(1-z^2) - \sigma_1(x,z)x(1-x) \in \Sigma_2, \tag{9.1.5}$$

where σ_1 is required to be SOS to enforce the constraint $x \in [0,1]$ on x. Choosing $\rho_0 = 16$, $\sigma_1 = 128z^2$ and $B = 16$ gives

$$\begin{aligned} S &= B(1+z^2) - 32z^2(1-2x)^2 - \rho_0(x,z)(1-z^2) - \sigma_1(x,z)x(1-x) \\ &= 16(1+z^2) - 32z^2(1-2x)^2 - 128z^2x(1-x) - 16(1-z^2) \\ &= 0 \in \Sigma_2, \end{aligned} \tag{9.1.6}$$

and of course $\sigma_1 = 128z^2 \in \Sigma_2$ also.

Numerically we also confirm these results when $r = 4$, placing (9.1.1) in the chaotic regime. It can easily be verified that the maximal LE is attained at the fixed point $x^* = 0$ and that its value is $\mu = \log(Df(x^*)) = \log(4-8x^*) = \log(4) \approx 1.38629$. Using the SOS program (6.3.3) with degree-4 $V(x)$ and $\epsilon = 0.1$ also gives $B = 16.00000$, from which we get

$$\mu = 1.38629, \tag{9.1.7}$$

which agrees with the exact value of the LE at the origin to six digits.

9.2 1D map which does not attain its maximal LE on a fixed point

The previous section showed that the logistic map (9.1.1) attains its maximal LE on the orbit corresponding to the fixed point at the origin. The situation that the maximal LE is attained on a fixed point seems to be a common property of many dynamical systems [Boc19; Kal11; Moh19]. For some systems however, the maximal LE is attained on a different orbit, such as a periodic orbit. The following system is included here to show an example of a map which maximizes its LEs on a periodic orbit.

Consider the 1D map

$$x_{k+1} = f(x_k) = -x_k^3 + 10x_k. \tag{9.2.1}$$

Table 9.1 shows the map's fixed points and orbits of period two along with the value of the geometric mean of $\Phi(x,z) = z^\mathsf{T} Df(x)^\mathsf{T} Df(x) z$ and the LE on each orbit. Note that for one-dimensional systems, bounding the maximal LE is equivalent to bounding the geometric mean of $Df(x)^2$ because (7.1.6) reduces to

$$\mu = \frac{1}{2}\log\left[\limsup_{k\to\infty}\left(\prod_{i=0}^{k-1}\Phi(x_i,z_i)\right)^{1/k}\right] = \frac{1}{2}\log\left[\limsup_{k\to\infty}\left(\prod_{i=0}^{k-1}Df(x_i)^2\right)^{1/k}\right], \tag{9.2.2}$$

since $z = 1$ or $z = -1$ only. Since the quantity whose geometric mean we are bounding is z-independent, we do not need to take the z-dynamics (7.1.3) into account. Concretely, this means that the auxiliary function V does not need to depend on z, and we do not need to additionally constrain z via, for instance, an s-procedure. This means that in 1D the problem of bounding the maximal LE entirely reduces to the problem in (6.1.10). However, as the example in the previous section showed, for some systems it may also be possible to obtain a bound on the maximal LE using the approach in (7.2.5) and using a z-dependent V.

Period	Points on orbit		$\widetilde{\Phi}$	μ
1	0		100	2.302595
1	3		289	2.833213
1	-3		289	2.833213
2	$\sqrt{11}$	$-\sqrt{11}$	529	3.135494
2	$-\sqrt{5+2\sqrt{6}}$	$-\sqrt{5-2\sqrt{6}}$	96	2.282174
2	$\sqrt{5-2\sqrt{6}}$	$\sqrt{5+2\sqrt{6}}$	96	2.282174

Table 9.1: Orbits of period two or less of (9.2.1) and the value of the geometric mean of $\Phi(x, z) = Df(x)^2$, together with the LE on each orbit.

From table 9.1 we can see that the largest LE on any of these orbits is on the orbit going from $x_0 = \sqrt{11}$ to $x_1 = -\sqrt{11}$. While there may be higher-degree orbits, such orbits will not have larger LEs than this period-two orbit. We computationally confirmed this assertion by numerically solving (6.1.11) with $\Phi = Df(x)^2$. At degree-2 V and $\epsilon = 300$, the bisection method in (6.1.12) to bound the maximal LE of (9.2.1), gives $B \leq 528.999999$, from which we compute

$$\mu \leq \frac{1}{2}\log(B) \approx 3.135494. \tag{9.2.3}$$

This value approximates the exact value of the maximal LE to seven digits of accuracy. It also suggests that there are no other orbits (such as higher-period orbits) that possess an even greater LE.

9.3 Analytically bounding the maximal LE of the Hénon map

As an example of a higher-dimensional map, where the method in 7.2.5 must be used to bound the maximal LE we consider the Hénon map

$$\begin{aligned} u_{k+1} &= 1 - a u_k^2 + v_k \\ v_{k+1} &= b u_k, \end{aligned} \tag{9.3.1}$$

[Wen14]. For $a = 1.4$ and $b = 0.3$, Hénon claimed that all orbits of the system tend to a chaotic strange attractor [Hén76]. However, this has never been proven. Indeed, Zgliczynski [Zgl97] argued that it is impossible to prove the existence of a strange attractor for the Hénon map using topological methods alone. For instance, at the parameters $(a, b) = (1.39945219, 0.3)$ the system possesses a stable, period-13 limit cycle attractor [Wen14]. In other words if the system is indeed chaotic at the claimed parameter values, then there exists a bifurcation from stability to chaos extremely close to the claimed chaotic values.

The putative attractor is contained in a ball of radius $R = 2$ around the origin of the u, v plane. It contains two fixed points of the form:

$$u^* = \frac{-(1-b) \pm \sqrt{(1-b)^2 + 4a}}{2a} \tag{9.3.2}$$

$$v^* = bu^*. \tag{9.3.3}$$

Furthermore, it has been shown using interval arithmetic methods that the Hénon map contains periodic orbits of all periods inside of the box $[-5, 5]^2$, except for period 3 and period 5 orbits [Gal98]. We note that the system (9.3.1) possesses time reversal symmetry, as can be seen by making the change of variables $u \mapsto -u$ and $t \mapsto -t$, however, the methods in chapter 8 do not enable us to use time symmetry.

We will first attempt to bound the maximal LE of the Hénon map by hand. We make the auxiliary function ansatz

$$V = 1 + z_1^2 + z_2^2, \tag{9.3.4}$$

so that

$$\Phi = b^2 z_1^2 + (2auz_1 - z_2)^2. \tag{9.3.5}$$

Letting $U = [u, v]$ and $Z = [z_1, z_2]$, and $f(u, v)$ be the vector field of (9.3.1), the global SOS constraint becomes

$$\begin{aligned} S &= BV(U,Z) - \Phi(U,Z)V\left(f(U), \frac{Df(U)Z}{\sqrt{\Phi}}\right) \\ &= B + Bz_1^2 + Bz_2^2 - 2b^2z_1^2 - 8a^2u^2z_1^2 - 2z_2^2 + 8auz_1z_2 \\ &= B + Bz_1^2 + Bz_2^2 - 2b^2z_1^2 - 9a^2u^2z_1^2 - 18z_2^2 + (auz_1 + 4z_2)^2, \end{aligned} \tag{9.3.6}$$

where we have completed the square on the last term in the second line. Now define the following barrier functions and s-multipliers:

$$\begin{aligned} h_0 &:= (1 - z_1^2 - z_2^2) \\ \rho_0 &:= (\delta + 8\gamma a^2 u^2) \\ g_1 &:= (r^2 - u^2) \\ \sigma_1 &:= const, \end{aligned} \tag{9.3.7}$$

where δ, γ, r and σ_1 are arbitrary constants to be determined. We need to show that $S - \rho_0 h_0 - \sigma_1 g_1 \in \Sigma_4$ and $\sigma_1 \in \Sigma_4$ for $B \in \mathbb{R}$. Then $\mu_{\mathcal{B}}^* \leq \frac{1}{2}\log(B)$, where

$\mathcal{B} = \{(u, v) \in \mathbb{R}^2 \ : \ r^2 - u^2 \geq 0\}$. Note that this set is not compact and so we do not have guarantees of sharpness, although the computed bounds in the next section nonetheless appear to be sharp. Thus, using this s-procedure, we write

$$\begin{aligned} S &= BV(U,Z) - \Phi(U,Z)V\left(f(U), \frac{Df(U)Z}{\sqrt{\Phi}}\right) - \rho_0 h_0 - \sigma_1 g_1 \\ &= B + Bz_1^2 + Bz_2^2 - 2b^2 z_1^2 - 9a^2u^2z_1^2 - 18z_2^2 + (auz_1 + 4z_2)^2 \\ &\quad - (\delta + 8\gamma a^2u^2)(1 - z_1^2 - z_2^2) - \sigma_0(r^2 - u^2) \\ &= B - \delta - \sigma_0 r^2 + (B + \delta - 2b^2)z_1^2 + (B + \delta - 18)z_2^2 \\ &\quad + (8\gamma - 9)a^2u^2z_1^2 + 8\gamma a^2u^2z_2^2 + (\rho_u - 8\gamma a^2)u^2 + (auz_1 + 4z_2)^2. \end{aligned} \tag{9.3.8}$$

Let $\gamma = 9/8$ and $\sigma_1 = 9a^2$. Clearly $\sigma_1 \in \Sigma_4$ for all $a \in \mathbb{R}$, then

$$\begin{aligned} S = B - \delta - 9a^2r^2 + (B + \delta - 2b^2)z_1^2 + (B + \delta - 18)z_2^2 \\ + 9a^2u^2z_2^2 + (auz_1 + 4z_2)^2. \end{aligned} \tag{9.3.9}$$

But this is SOS whenever

$$\max\{2b^2 - B, 18 - B\} \leq \delta \leq B - 9a^2r^2, \tag{9.3.10}$$

and

$$B = \max\{9a^2r^2,\ 18,\ 2b^2\}. \tag{9.3.11}$$

This proves a bound on the top LE of the Hénon map, which depends on the parameters in the system when these are large enough. If we pick $r = 2$, for instance, then $B = 9a^2r^2 = 70.56$ and so this shows that on the putative strange attractor and on both of the fixed points, the maximal LE can be no larger than

$$\mu \leq \frac{1}{2}\log(B) = \frac{1}{2}\log(70.56) \approx 2.128232. \tag{9.3.12}$$

9.4 Numerically bounding the maximal LE of the Hénon map

Since we obtained finite bounds on the maximal LE in the previous section, it should in principle be possible to use those results as a starting point for setting up a numerical SOS program and attempting to tighten the bounds on $\mu^*_{\mathcal{B}}$ at some fixed parameter values. It turns out that when we allow the s-multipliers of the barrier functions $h_0 = 1 - |z|^2$ and $g_1 = r^2 - u^2$, to be undetermined polynomial ansätze, the bound we obtain for $V(x, z) = 1 + |z|^2$ remains unchanged from the value that we obtain when $r = 2$ and we only allow the variables δ, γ and σ_1 to be tunable with the other auxiliary functions fixed to those in (9.3.7).

$\frac{1}{2}\log(B)$	d
$\frac{1}{2}\log(31.953) \approx 1.732$	2
$\frac{1}{2}\log(10.626) \approx 1.182$	4
$\frac{1}{2}\log(10.624) \approx 1.182$	6

Table 9.2: Bounds on the maximal LE of (9.3.1) using $r = 2$. d is the maximal degree of V, ρ_0 and σ_1, which were all fixed to the same value. We fix $\epsilon = 0.01$ in (7.2.6) to obtain the values in the table. The bounds are accurate to the precision shown.

The bounds in table 9.2 were obtained by considering only polynomials that are even in z, but otherwise allowing all V ρ_0 and σ_1 to be freely varying in the polynomial spaces up to the specified degrees. At the two fixed points the eigenvalues of

$$Df = \begin{bmatrix} -2au & 1 \\ b & 0 \end{bmatrix}, \tag{9.4.1}$$

are

$$\lambda(Df(u_1^*)) = \begin{cases} -1.924 \\ 0.156 \end{cases} \tag{9.4.2}$$

$$\lambda(Df(u_2^*)) = \begin{cases} 3.260 \\ -0.092 \end{cases} . \tag{9.4.3}$$

therefore the largest real part of any of these eigenvalues is $\lambda = 3.260$ and $\mu = \log(\lambda) = 1.182$, so that $\lambda^2 \approx 10.626$, which agrees with the bound obtained at $d = 4$ to three digits. This relatively low accuracy is due to inherent numerical ill-conditioning in the Hénon map at the parameters at which the computations where carried out. However, the fact that one is able to obtain analytic bounds by solving the SOS problem by hand, shows that the issues encountered here are numerical in nature and that the SOS formulation for bounding the LE does work for this example too. In fact, separate numerical investigations have indicated that if the same SOS problem, which we have solved by hand in section 9.3, is solved numerically, then bounds which also seem to be sharp to the same accuracy may be obtained.

Chapter 10

Bounding sums of LEs

In the preceding sections we studied the problem of computing bounds on the maximal LE of continuous and discrete dynamical systems. A related problem which we want to study is bounding the sum of the k largest LEs over bounded orbits of a system. This problem is of interest to us because it is related to the problem of determining the maximal expansion rate of k-dimensional volumes in phase space, and the problem of bounding the dimension of attractors.

An orbit $x(t)$ of the system (1.0.1) has n distinct (but not necessarily different) LEs as given by the definition (1.0.3) [Mei17]. The set of all LEs of an orbit is called the Lyapunov spectrum. Each LE μ_i in the Lyapunov spectrum is computed from the definition (1.0.3) via a solution $y_i(t)$ of the linearization (1.0.2). Each such $y_i(t)$ is called a Lyapunov vector (LV). Note that LVs are necessarily linearly independent.

Consider the paralellotope whose verticies are give by y_i, where each y_i is a distinct LV of the system (1.0.1). Then the length of each vertex evolves as the length of each LV, which is asymptotic to an exponential growth rate dependent on the LE associated with that particular LV:

$$|y_i| \sim e^{\mu_i t}. \tag{10.0.1}$$

Therefore, since the LVs are always linearly independent, this implies that the volume [DG95b]

$$V_k(t)^2 := |y_1 \wedge ... \wedge y_k|^2 = \text{Det} \underbrace{\{y_i^{\mathsf{T}} y_j\}}_{\mathsf{A}}, \tag{10.0.2}$$

where $\wedge$ denotes the n-dimensional outer product and A is the matrix with elements $\mathsf{A}_{ij} = y_i^{\mathsf{T}} y_j$, grows or decays according to [Tem12]

$$V_k(t) \sim \exp(\mu_1 t) \exp(\mu_2 t) ... \exp(\mu_k t) = \exp(\mu_1 t + ... + \mu_k t). \tag{10.0.3}$$

This implies that if $\mu_1 + ... + \mu_k < 0$ then $V_k(t)$ shrinks and if $\mu_1 + ... + \mu_k > 0$ then $V_k(t)$ grows. Knowing the smallest k for which $\mu_1 + ... + \mu_k < 0$ therefore gives a bound on the dimensions of the space to which long-time dynamics of (1.0.1) tends.

If solutions tend to an attractor then the statement that the sum of the largest k LEs of the chaotic orbit on the attractor is negative, implies that the dimension of the attractor is less than k. This is because by the reasoning above, the attractor cannot contain any volumes of dimension k.

If the attractor is a strange attractor, then its dimension will not be an integer in general, as it will be a fractal. Its dimension will be more accurately quantified by a generalization of the usual notion of dimension, such as the Hausdorff dimension or the capacity dimension (also called box-counting dimensions). The capacity dimension of an attractor $\mathcal{A}$ can be computed as follows: let $N(\epsilon)$ be the minimum number of balls of radius ϵ that are needed to completely cover $\mathcal{A}$ (note that if $\mathcal{A}$ is compact, then this number is always finite). Then the capacidty dimension is defined as

$$d_C := \limsup_{\epsilon \to 0} \frac{N(\epsilon)}{\epsilon}. \tag{10.0.4}$$

A famous conjecture by Kaplan and Yorke [Fre+83] states that the sum of LEs and the capacity dimension are directly related, and indeed, that they are equal so that the LEs of a system can be used to exactly compute the capacity dimension. Suppose that the sum of the largest $\tilde{k}$ LEs is negative but the sum of the largest $\tilde{k}-1$ LEs is not negative. Then the Kaplan–Yorke dimension is defined as

$$d_{KY} := \tilde{k} - 1 + \frac{\mu_1 + ... + \mu_{\tilde{k}-1}}{-\mu_{\tilde{k}}}. \tag{10.0.5}$$

It is conjectured that $d_{KY} = d_C$. Note that

$$0 \leq \frac{\mu_1 + ... + \mu_{\tilde{k}-1}}{-\mu_{\tilde{k}}} < 1, \tag{10.0.6}$$

and so this quotient minus one is the correction to the minimal integer-dimensional volume in which the attractor is embedded, so that $\tilde{k} - 1$ plus the correction gives the fractal attractor dimension.

Thus having an upper bound on the sum of the leading k LEs, and finding the smallest k for which the sum is negative, would give a useful tool for bounding the dimension of all the attractors of the orbits which we consider. We therefore try to find a SOS formulation which will allow us to accomplish this.

In Section 4.2 in [DG95b] it is discussed how to compute the growth of volumes along trajectories of the system via:

$$\begin{aligned} \mu_1 + ... + \mu_k &= \limsup_{T\to\infty} \frac{1}{T} \log\Big(V_k(T)\Big) \\ &= \limsup_{T\to\infty} \frac{1}{T} \int_0^T \underbrace{\sum_{i=1}^k \phi_i(s)^\intercal Df\Big(x(s)\Big)\phi_i(s)}_{=\Phi(x,y_1,...,y_k)} \,\mathrm{d}s. \end{aligned} \tag{10.0.7}$$

Here, $\phi_1(t), ..., \phi_k(t)$ are a set of orthonormal basis vectors that span the span of the leading k LVs, $y_1, ..., y_k$. They can, for instance, be constructed using a Gram-Schmidt procedure [PC89].

If $k = 1$, this definition agrees with the LE definition that we have been working with up to this point. Therefore $\phi_1 = z_1 = \frac{y_1}{|y_1|}$ in this case. Hence, the relevant ODE for the z_1 is

$$\frac{\mathrm{d}}{\mathrm{d}t} z_1 = \ell_i(x, z_1) = Df(x)z_1 - \left(z_1^\mathsf{T} Df(x) z_1\right) z_1. \tag{10.0.8}$$

10.1 Optimization problems in special cases

While we do not at this point have an optimization framework for bounding the time average in (10.0.7) in general, in certain special cases we can write down optimization problems which are solvable and which enjoy the same theoretical duality results and symmetry properties as the optimization problems which we have been working with for bounding the maximal LE. In section 10.1.1 we formulate the problem of bounding the maximal sum of all LEs of a system, a problem which also serves to illustrate the validity of the method for bounding time averages in chapter 2. In section 10.1.4 we derive a optimization problem whose solution gives a bound on the maximal sum of the leading $(n-1)$ LEs, when the dimension of the phase space is three or four.

10.1.1 Sum of all LEs

Computing the sum of all LEs of an orbit of (1.0.1) is the simplest problem because in this special case, the dimension of the problem collapses to n for an n-dimensional system. In fact, this problem is more straightforward to treat using convex optimization methods than computing the maximal LE, which was done in the preceding sections and which requires solving optimization problems on a $2n$-dimensional phase space.

Recall the theorem by Lyapunov which states that for an n-dimensional system (1.0.1) one has

$$\sum_{i=1}^{n} \mu_i(x_0) = \limsup_{T\to\infty} \frac{1}{T} \int_0^T \mathrm{tr} Df(x(s))\, ds. \tag{10.1.1}$$

For a proof see e.g. [PP16]. This equality gives an expression for the sum of all the LEs of an orbit in terms of the time average of the trace of the Jacobian, or equivalently the divergence, of the vector field $f(x)$ of (1.0.1) along that orbit.

Therefore, letting $\Phi(x) := \mathrm{tr} Df(x)$, we can immediately use the formulation in (2.2.1) to derive an optimization problem, or the formulation in (2.3.11) to derive a SOS program to bound the sum of all the LEs over orbits of (1.0.1). Note that this does not require any knowledge of the tangent dynamics of $x(t)$, and we do not have to enlarge the nonlinear system to include any ODEs that encode such dynamics as we had to do in chapter 3 when we bounded the maximal LE.

The symmetry results in Propositions 18 and 19 directly apply to this problem if $f(x)$, $\Phi(x)$ and the domain $\mathcal{B}$ satisfy the conditions in those propositions. The reason is that we do not need to worry about symmetries of ODEs governing the tangent dynamics and the given equivariance of $f(x)$ under a group of transformation $\mathcal{G} \subset O(n)$ is enough. Furthermore, if $f \in \mathcal{C}^1(\mathcal{B})$ is equivariant under $\Lambda \in \mathcal{G}$ then $\Phi(x)$ is invariant under Λ, because the trace is invariant under similarity transformations and so

$$\Phi(\Lambda x) = \mathrm{tr}(Df(\Lambda x)) = \mathrm{tr}(\Lambda Df(x)\Lambda^{-1}) = \mathrm{tr}(Df(x)) = \Phi(x). \tag{10.1.2}$$

10.1.2 Theory of cross products and outer products

Besides bounding the sum of all LEs of a system, we also want to study the problem of bounding the sum of $k < n$ LEs—or equivalently—the growth of k-dimensional volumes for $k < n$. We consider volumes defined by (10.0.2), where each y_i represents a vertex of a k-dimensional paralellotope and evoles according to the linearization (1.0.2). Since this is the object which we need to study in order to bound the sum of LEs in general, we will recall some basic notions of the outer product and in particular, when it can be expressed in a more tangible form as a cross product. The outer product or "wedge" product $\wedge$ generlizes the 3D cross product to higher dimensions. In fact, it can be shown that an explicit cross product between two vectors only exists in dimensions 3 and 7 [Mas83]. Nonetheless, if we can explicitly compute a cross product this gives us the advantage of being able to represent an area or a volume as a simple vector. We can then track the evolution of such volumes by tracking the dynamics of the cross-product-defined vector that determines them, much like we tracked the evolution of a vector in the previous sections to track the growth or decay of lengths. While we cannot define a cross product in $\mathbb{R}^n$ for $3 < n \neq 7$, we can however define an $(n-1)$ cross product between $(n-1)$ vectors in an n-dimensional space as follows [Spi18]:

$$y_1 \times ... \times y_{n-1} := \det \begin{bmatrix} \hat{e}_1 & \dots & \hat{e}_k & \dots & \hat{e}_n \\ y_1(1) & \dots & y_1(k) & \dots & y_1(n) \\ \vdots & & \vdots & & \vdots \\ y_k(1) & \dots & y_k(k) & \dots & y_k(n) \\ \vdots & & \vdots & & \vdots \\ y_{n-1}(1) & \dots & y_{n-1}(k) & \dots & y_{n-1}(n) \end{bmatrix}, \tag{10.1.3}$$

where, with a slight abuse of notation, $\hat{e}_j$ denotes the j'th Euclidean basis vector of $\mathbb{R}^n$ and $y_i(j)$ denotes the j'th component of the i'th vector. Crucially, it can be shown that [Spi18]

$$|y_1 \times ... \times y_{n-1}|^2 = \mathrm{Det}\{y_i^\mathsf{T} y_j\}, \tag{10.1.4}$$

which allows us to use this $(n-1)$-cross product to compute the sum of the largest $(n-1)$ LEs. This multi cross product shares the property of the cross product in 3D of being orthogonal to the vectors $y_1, ..., y_{n-1}$. In dimensions $n = 3$ and $n = 4$, it has a closed form evolution equation, which governs its dynamics, much like we have a closed form expression via (1.0.2) that governs the evolution of lengths along an orbit of (1.0.1). Since the $(n-1)$ cross product defines an $(n-1)$-dimensional volume, its rate of growth or decay along orbits is governed by the leading $(n-1)$ LEs of the system. We therefore want to use the $(n-1)$ cross product to derive tractable optimization problems to bound the sum of the leading $(n-1)$ LEs.

10.1.3 Sum of leading n-1 LEs for n=3 and n=4

The problem of bounding the sum of all LEs simplifies to a tractable problem due to it being a special case. It turns out that there exist some other important special cases which make bounding other sums of LEs possible in practice: for instance, the method presented in [DG95b] relies on the fact that in 3D, the area between two vectors which evolve according to the linearization (1.0.2), evolves as

$$\frac{\mathrm{d}}{\mathrm{d}t} v = \left(\mathrm{tr}(Df(x))\mathbf{I} - Df(x)^\mathsf{T}\right) v, \tag{10.1.5}$$

where $v = y_1 \times y_2$. And so

$$\mu_1 + \mu_2 = \limsup_{t\to\infty} \frac{1}{t} \log |v(t)|. \tag{10.1.6}$$

This follows from the fact that $\frac{\mathrm{d}}{\mathrm{d}t} y_i = Df(x) y_i$ and that

$$\Big(\mathsf{M}y_1 \times y_2\Big) + \Big(y_1 \times \mathsf{M}y_2\Big) = \Big(\mathrm{tr}(\mathsf{M})\mathbf{I} - \mathsf{M}^\mathsf{T}\Big)\,(y_1 \times y_2), \tag{10.1.7}$$

where each $y_i \in \mathbb{R}^3$ and M is an arbitrary 3×3 matrix. The above identity can be confirmed by direct calculation. A similar direct calculation shows that in four dimensions we have the identity

$$\Big(\mathsf{M}y_1 \times y_2 \times y_3\Big) + \Big(y_1 \times \mathsf{M}y_2 \times y_3\Big) + \Big(y_1 \times y_2 \times \mathsf{M}y_3\Big) = \Big(tr(\mathsf{M})\mathbf{I} - \mathsf{M}^\mathsf{T}\Big)\,(y_1 \times y_2 \times y_3), \tag{10.1.8}$$

where $y_i \in \mathbb{R}^4$ and M is any 4×4 real matrix and the cross product is defined as in (10.1.3). Therefore, (10.1.5) also holds in four dimensions, with $v = y_1 \times y_2 \times y_3$.

However, it can be shown via a counterexample that the formula (10.1.5) does not generalize to five dimensions. As the counterexample, take M to be the identity matrix and take

$$\begin{aligned} y_1 &= [1, 3, 3, 3, 3] \\ y_2 &= [1, 3, 1, 2, 2] \\ y_3 &= [-1, 2, 2, 1, 1] \\ y_4 &= [-1, 0, 0, 1, 2], \end{aligned} \tag{10.1.9}$$

which are linearly independent. Then

$$\left(\mathsf{M}y_1 \times y_2 \times y_3 \times y_4\right)+\left(y_1 \times \mathsf{M}y_2 \times y_3 \times y_4\right) \\ + \left(y_1 \times y_2 \times \mathsf{M}y_3 \times y_4\right)+\left(y_1 \times y_2 \times y_3 \times \mathsf{M}y_4\right)= [18, 9, 15, -78, 48], \tag{10.1.10}$$

but

$$\left(\mathrm{tr}(\mathsf{M})\mathbf{I} - \mathsf{M}^{\mathsf{T}}\right)(y_1 \times y_2 \times y_3 \times y_4) = [24, 12, 20, -104, 64], \tag{10.1.11}$$

and so (10.1.5) does not hold in dimension five.

10.1.4 Optimization problems to bound the sum of the leading $(n-1)$ LEs

Building on what was done in chapter 3, we will derive optimization problems whose solutions can give sharp bounds on the sum of the leading $(n-1)$ LEs in dimensions three and four. We will then show that the resulting optimization problems share the same symmetry properties as the optimization problems in chapter 3

Using the ODE (10.1.5), we can derive SOS programs to bound the leading two LEs of a three-dimensional dynamical system or the leading three LEs of a four-dimensional system. When working with a single LE, one has to take into account the linearized dynamics (1.0.2) of (1.0.1) and the object that is most prominent in the derived expression is therefore the Jacobian of the vector field $f(x)$. When bounding the sum of the leading $(n-1)$ LEs on the other hand, we do not use the dynamics of a single tangent vector, but rather the dynamics of an $(n-1)$-dimensional volume element, whose dynamics is governed by (10.1.5). The object which takes the place of $Df(x)$ in the derived expressions therefore, is the right hand matrix in (10.1.5):

$$\mathsf{A}(x) := \mathrm{tr}(Df(x))\mathbf{I} - Df(x)^{\mathsf{T}}, \tag{10.1.12}$$

so that (10.1.5) can be written as

$$\frac{\mathrm{d}}{\mathrm{d}t}v = \mathsf{A}(x)v. \tag{10.1.13}$$

Let $z := \frac{v}{|v|}$, then

$$\frac{\mathrm{d}}{\mathrm{d}t}z = \mathsf{A}(x)z - (z^{\mathsf{T}}\mathsf{A}(x)z)z. \tag{10.1.14}$$

Therefore expression (10.1.6) implies that

$$\mu_1 + \mu_2 = \limsup_{T\to\infty} \frac{1}{T}\log(v(T)) = \limsup_{T\to\infty} \frac{1}{T}\int_0^T z^{\mathsf{T}}\mathsf{A}(x)z \,\mathrm{d}s, \tag{10.1.15}$$

and since (10.1.5) also holds for $n = 4$ we similarly have

$$\mu_1 + \mu_2 + \mu_3 = \limsup_{T\to\infty} \frac{1}{T}\log(v(T)) = \limsup_{T\to\infty} \frac{1}{T}\int_0^T z^{\mathsf{T}}\mathsf{A}(x)z \,\mathrm{d}s. \tag{10.1.16}$$

Therefore, if we let $\Phi(x, z) := z^{\mathsf{T}}\mathsf{A}(x)z$ and $\ell(x, z) := \mathsf{A}(x)z - (z^{\mathsf{T}}\mathsf{A}(x)z)z$ we obtain optimization problems analogous to those in Proposition 3 to bound the maximal sum of $(n-1)$ LEs. Define the maximal sum of the leading $(n-1)$ LEs (for $n = 3$ and $n = 4$) as

$$\Sigma\mu^*_\Omega := \max_{\substack{x_0\in\Omega \\ y_i\in\mathbb{S}^{n-1}}} \sum_{i=1}^{I} \mu(x_0, y_i), \tag{10.1.17}$$

for $I = n$, $I = 4$ or $I = 3$. Each y_i is a distinct LV so that we do not simply have

$\Sigma\mu_\Omega^* = I\mu_1$. Then

$$\Sigma\mu^*_{\mathcal{B}_0} \leq \inf_{V\in\mathcal{C}^1(\mathcal{B}\times\mathbb{S}^{n-1})} \sup_{\substack{x\in\mathcal{B}\\ z\in\mathbb{S}^{n-1}}} \left[z^\mathsf{T}\mathsf{A}z + f\cdot\nabla_x V + [\mathsf{A}z - (z^\mathsf{T}\mathsf{A}z)z]\cdot\nabla_z V\right], \quad (10.1.18)$$

when $\mathcal{B}_0$ and $\mathcal{B}$ are defined as in Proposition 3. Furthermore, if the domain $\mathcal{B}$ satisfies the condition of the second part of Proposition 3, then the above inequality is in fact an equality. This convex optimization problem can be relaxed to a SOS program, whose solution also gives an upper bound on $\Sigma\mu^*_{\mathcal{B}_0}$, if the vector field in (1.0.1) is polynomial in x. Let

$$P(x,z) := B - z^\mathsf{T}\mathsf{A}(x)z - f(x)\cdot\nabla_x V(x,z) - \Big[\mathsf{A}(x)z - (z^\mathsf{T}\mathsf{A}(x)z)z\Big]\cdot\nabla_z V(x,z), \quad (10.1.19)$$

then if $\mathcal{B}_0$ and $\mathcal{B}$ are defined as in Proposition 4 and $h_0 = 1-|z|^2$ we have the following upper bound on $\Sigma\mu^*_{\mathcal{B}_0}$:

$$\Sigma\mu^*_{\mathcal{B}_0} \leq \inf_{V,\sigma_j,\rho_i\in\mathbb{R}[x,z]} B \quad \text{s.t.} \quad P - \textstyle\sum_{i=0}^{I}\rho_i h_i - \sum_{j=1}^{J}\sigma_j g_j \in \Sigma_{2n}$$
$$\sigma_j\in\Sigma_{2n}, \quad j=1,\dots,J. \quad (10.1.20)$$

This inequality can be strengthened to an equality if the conditions of the second part of Proposition 4 are also satisfied.

Finally, it would be nice if the symmetry results for the optimization problems for bounding the maximal LE also apply to the optimization problems for bounding the maximal sum of the leading $(n-1)$ LEs, as this would allow us to solve smaller and more stable SOS programs. Fortunately, Propositions 5 and 6 have analogues for the optimization problems in (10.1.18) and (10.1.20). It suffices to show that for

$\Lambda \in O(n)$ the equivariance of $f(x)$ under Λ also implies the equivariance of $\ell(x,z)$ under a transformation of both of its arguments under Λ, and implies the invariance of $\Phi(x,z)$. The invariance of $\Phi(x,z)$ follows from the invariance of the trace under similarity transformations and

$$\begin{aligned}\Phi(\Lambda x, \Lambda z) &= z^{\mathsf{T}}\Lambda^{\mathsf{T}}\mathsf{A}(\Lambda x)\Lambda z \\ &= \operatorname{tr}\Big(\Lambda Df(x)\Lambda^{-1}\Big)z^{\mathsf{T}}\Lambda^{\mathsf{T}}\Lambda z - z^{\mathsf{T}}\Lambda^{\mathsf{T}}\Lambda^{-\mathsf{T}}Df(x)^{\mathsf{T}}\Lambda^{\mathsf{T}}\Lambda z \\ &= \operatorname{tr}\Big(Df(x)\Big)z^{\mathsf{T}}z - z^{\mathsf{T}}Df(x)^{\mathsf{T}}z \\ &= \Phi(x,z).\end{aligned} \tag{10.1.21}$$

Also note that, by definition, if $\Lambda \in O(n)$ then $\Lambda^{-\mathsf{T}} = \Lambda$ and so

$$\mathsf{A}(\Lambda x) = \operatorname{tr}\Big(\Lambda Df(x)\Lambda^{-1}\Big) - \Lambda^{-\mathsf{T}}Df(x)^{\mathsf{T}}\Lambda^{\mathsf{T}} = \operatorname{tr}(Df(x)) - \Lambda Df(x)^{\mathsf{T}}\Lambda^{\mathsf{T}}. \tag{10.1.22}$$

Therefore

$$\begin{aligned}\ell(\Lambda x, \Lambda z) &= \Big(\operatorname{tr}(Df(x)) - \Lambda Df(x)^{\mathsf{T}}\Lambda^{\mathsf{T}}\Big)\Lambda z - \Phi(\Lambda x, \Lambda z)z \\ &= \Lambda\left(\mathsf{A}(x)z - \Phi(x,z)\right)z \\ &= \Lambda\ell(x,z),\end{aligned} \tag{10.1.23}$$

which shoes that $\ell(x,z)$ as defined in this section is equivariant, and $\Phi(x,z)$ as defined in this this section is invariant under $(x,z) \mapsto (\Lambda x, \Lambda z)$ if $f(x)$ is equivariant under Λ. Furthermore $\ell(x,-z) = -\ell(x,z)$ and $\Phi(x,-z) = \Phi(x,z)$.

Chapter 11

Examples: bounding sums of LEs

We apply the methods in chapter 10 to some explicit examples. In section 11.1 we bound the maximal sum of all LEs of the Sprott A system using the method in section 10.1.1, which also serves to illustrate the applicability of the methods in chapter 2 to systems that maximize a quantity of interest on a chaotic trajectory. Section 11.2 shows that the method for bounding the maximal sum of the leading two LEs of a three-dimensional map gives results for the Lorenz system that are consistent with the ones that we obtained in section 5.1. The examples in sections 11.3 and 11.4 show how SOS programming methods can be used to bound the maximal sum of the leading three LEs of four-dimensional maps, and that the approach works even when the system is hyperchaotic or maximizes the sum on a periodic orbit.

11.1 Sprott A system

As mentioned in section 10.1.1 the problem of bounding the maximal sum of all LEs of a dynamical system (1.0.1) reduces to a simpler problem than bounding only a single LE or the maximal sum of the leading $(n-1)$ LEs. The reason for this is that the tangent dynamics need not be know to follow the growth or decay of n-dimensional volumes for an n-dimensional system. This means that this problem falls into the general framework of the optimization problems for bounding time averages discussed in sections 2.1 and 2.3. It therefore serves as a demonstration of these methods in general. Furthermore, we also have the sharpness guarantees of section 2.2 as these results directly apply to the problem of bounding the sum of all LEs as a convex optimization problem.

As a concrete example consider the Sprott A system [Spr94],

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} x_2 \\ -x_1 + x_2 x_3 \\ 1 - x_2^2 \end{bmatrix} \tag{11.1.1}$$

which has the property of not possessing any fixed points. The system possesses the symmetry $(x_1, x_2, x_3) \mapsto (-x_1, -x_2, x_3)$, a property which we make use of in the SOS program that we solve below. The trace of its Jacobean is $\mathrm{tr}(Df(x)) = x_3$. The system (11.1.1) is chaotic. On the chaotic trajectory the Lyapunov spectrum has been computed as $(0.014, 0, -0.014)$ [Spr94], which implies that volumes are preserved on the attractor and that the attractor is three dimensional, as would be predicted by the Kaplan–Yorke formula (10.0.5).

We bounded the maximal sum of all the LEs of the Sprott A system by solving (2.3.2) with $\Phi(x) = \mathrm{tr}(Df(x)) = x_3$ and using an auxiliary function of polynomial degree two. A bound of $\Sigma\mu^*_{\mathbb{R}^3} = 0.00000$ was obtained with a quadratic $V(x)$, which suggests that the bound is already sharp for degree-2 $V(x)$. In fact, since $\Phi(x) = x_3$ using an auxiliary function of the form

$$V(x) = -\frac{1}{2}(x_1^2 + x_2^2 + x_3^2), \tag{11.1.2}$$

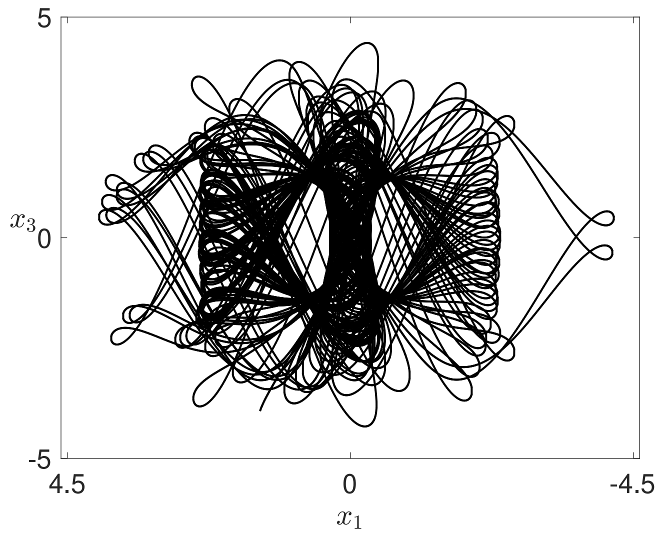

Figure 11.1: The attractor of the Sprott A system. Computed using a RK4 method with a time step of $dt = 0.01$ and with initial condition $x(0) = [0, 5, 0]$.

implies that when $B = 0$ then

$$\begin{aligned} S(x) &= B - \Phi(x) - f(x) \cdot \nabla V(x) \\ &= -x_3 - (x_1 x_2 - x_1 x_2 + x_2^2 x_3 - x_3 - x_2^2 x_3) = 0. \end{aligned} \tag{11.1.3}$$

Therefore, the inequality in (2.3.2) is in fact an equality in this case and $B = 0$ is the maximal sum of all the LEs of (11.1.1). But since the numerics imply that the Lyapunov spectrum of the Sprott A system is symmetric, this shows that a maximizing trajectory is in fact the chaotic trajectory shown in fig. 11.1. That being said, it may well be that the trajectories of (11.1.1) generally have symmetric Lyapunov spectra, in which case any trajectory of the system—including all its periodic orbits—are maximizing orbits for the sum of all LEs.

11.2 Lorenz system

For bounding the leading $(n-1)$ LEs, we first wanted to see if the SOS program in (10.1.20) would give accurate results for bounding the sum of the leading two LEs of a three-dimensional system. Since we know form section 5.1 that the maximal LE of the Lorenz system is attained on the fixed point at the origin, a natural question that we hoped to answer was whether the sum of the leading two LEs is also maximized on this orbit. We therefore solved this SOS program to bound the sum of the leading two LEs of the Lorenz system (5.1.1). We made use of the same symmetries as we did when bounding the maximal LE, when implementing this SOS program. Recall that at the origin, the sum of the two leading LEs is

$$\mu_1 + \mu_2 = \frac{1}{2}\left(-1 - \sigma - 2\beta + \sqrt{1 - 2\sigma + 4r\sigma + \sigma^2}\right) \approx 9.161057, \tag{11.2.1}$$

at the standard parameters. The authors in [DG95a; EFT91] conjectured that the sum of the leading two LEs is maximized on the fixed point at the origin. They did not solve the right-hand maximization analytically but carried out a (non-convex) numerical search over a compact set $\mathcal{B}$ containing the attractor. In general, sharp bounds require a method like ours that accounts for time averaging along trajectories. Table 11.1 shows the bounds on the maximal sum of the leading two LEs, which we computed using (10.1.20) with degree two ρ_0 and the maximum degree of V fixed to the values in the table.

The system (5.1.1) was rescaled with a change of variables $\tilde{x} = x/10$ in order to improve the accuracy of the solution when numerically solving (10.1.20). At degree-4 V and degree-2 ρ_0 this gave $\Sigma^*_{\mathbb{R}^3} \leq 9.161057$, which agrees with the exact value at the origin to seven digits of accuracy. This strongly indicates that the maximal sum

Degree of V	Upper bound on maximal sum of two LEs
2	11.362156
4	9.161057

Table 11.1: Upper bounds on the maximal sum of the two leading LEs of the Lorenz system (5.1.1). The degree of the weight ρ_0 in the s-procedure is fixed to two and the maximum degree of V is fixed to the values in the table.

of the leading two LEs of the Lorenz system is maximized on the fixed point at the origin.

11.3 Hyperchaotic Lorenz-type system

To test (10.1.20) in the case when $n = 4$, we sought an unstable four dimensional system. In four spatial dimensions the phenomenon of hyperchaos becomes a possibility, which means that the system possesses two positive LEs, rather than just the single one required for classical chaos. As an example of a hyperchaotic system, consider the four-dimensional Lorenz-type system [Li+11]:

$$\frac{\mathrm{d}}{\mathrm{d}t}\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix} = \begin{bmatrix} a(x_2 - x_1) \\ bx_1 - cx_2 - x_1x_3 + x_4 \\ -dx_3 + x_1x_2 \\ -kx_2 - rx_4, \end{bmatrix} \tag{11.3.1}$$

when $a = 12$, $b = 23$, $c = 1$ $d = 2.1$ $k = 6$ and $r = 0.2$, this system possesses a single fixed point at the origin and displays hyperchaotic dynamics [CY14]. I.e. the chaotic trajectory of its chaotic attractor possesses two positive LEs. We solve the SOS program

$$\Sigma\mu^*_{\mathbb{R}^4} \leq \inf_{V,\rho_0 \in \mathbb{R}[x,z]} B \quad \text{s.t.} \quad P - \rho_0(1 - |z|^2) \in \Sigma_8 \ , \tag{11.3.2}$$

to obtain a bound on the leading three LEs of the system.

Degree of V	Upper bound on maximal sum of three LEs
2	9.632800
4	8.613661
6	8.613661

Table 11.2: Upper bounds on the maximal sum of the three leading LEs of (11.3.1). All polynomial degrees of auxiliary functions and s-procedure weights are bounded by the value in the left column of the table. The system was rescaled with the change of variables $\tilde{x} = 10x$.

At the origin, the eigenvalues of the Jacobian of (11.3.1) are $-23.913662, 10.638190,$ 0.075472 and -2.100000. Therefore at the fixed point at the origin, the sum of the leading three LEs is $\mu_1 + \mu_2 + \mu_3 = 8.613661$, which agrees with the results in table 11.2 to seven digits of accuracy at degree-4 and degree-6 auxiliary functions. In order to obtain this accuracy from the numerical computations, the system (11.3.1) was rescaled with the change of variables $\tilde{x} = 10x$. Note that even though the system is hyperchaotic, this only means that the leading two LEs of the chaotic trajectory are positive, but not necessarily those of a fixed point, as this example shows.

11.4 Hénon–Heiles system

As we did in section 5.4 we also want to attempt to bound the maximal sum of the leading $(n-1)$ LEs of a system which does not maximize this sum on a fixed point. We note that finding a low-dimensional system whose maximal sum of LEs is not attained on a fixed point and is different from the maximal LE is unlikely, because most systems seem to attain their maximal time averages on fixed points and Hamiltonian systems only have two leading non-zero LEs for dimension six and higher. Notwithstanding, we turn to the Hénon–Heiles system restricted to the domain (5.4.3), because it maximizes its sum of LEs on a periodic orbit. This follows immediately from the work in section 5.4, because of the Hamiltonian nature of the HH system, which implies that for any orbit the sum of the leading LE equals both the sum of the leading two and the leading three LEs. The reason for this is that any non-fixed-point bounded orbit of a dynamical system has at least one LE that is identically zero, and for a Hamiltonian system LEs always come in pairs where one LE is positive and the other is its negation. Therefore the Lyapunov spectrum of each orbit in (5.4.3) has the form $(\mu,\ 0,\ 0,\ -\mu)$. This reasoning implies that the maximal sum of the leading three LEs is equal to the maximal LE of the HH system, and furthermore that the maximizing orbits are the same as those we found in section 5.4 on which the maximal LE is attained.

We therefore solved a sequence of SOS programs of the form

$$\Sigma\mu^*_{\mathbb{R}^4} \leq \inf_{V,\rho_0,\sigma_1,\sigma_2,\sigma_3\in\mathcal{W}_d} B \quad \text{s.t.} \quad \begin{aligned} P - \rho_0(1-|z|^2) - \textstyle\sum_{j=1}^3 \sigma_j g_j &\in \Sigma_8, \\ \sigma_j &\in \Sigma_8, \end{aligned} \tag{11.4.1}$$

with $\mathcal{W}_d$ defined as in (5.4.5) and P defined as in (10.1.19). Table 11.3 shows the upper bounds obtained by numerically solving the right hand side of (11.4.1). The right column in table 11.3 shows the upper bound on the maximal sum of LEs, while the left column shows the degree d of the polynomial space (5.4.5) to which the tunable polynomials in (11.4.1) belong. Table 11.3 shows that at degree-10 auxiliary function the upper bound on the maximal sum of the leading three LEs seems to be sharp and agrees with the maximal LE computed in section 5.4 to the same precision.

Table 11.3: Upper bounds on the maximal sum of the leading three LEs among trajectories of the Hénon–Heiles system (5.4.2) in the set $\mathcal{B}$ defined by (5.4.3) and (5.4.4), found by numerically solving the right-hand SOS program in (11.4.1) with the maximum degree d of all tunable polynomials fixed to various values. Tabulated values are rounded to the precision shown.

Degree of polynomials	Upper bound on maximal sum of three LEs
4	0.80539
6	0.75549
8	0.23852
10	0.23081

Bibliography

[AA00] E. D. Andersen and K. D. Andersen. "The MOSEK interior point optimizer for linear programming: an implementation of the homogeneous algorithm". In: *High performance optimization*. Springer, 2000, pp. 197–232.

[ACT80] A. Arneodo, P. Coullet, and C. Tresser. "Occurrence of strange attractors in three-dimensional Volterra equations". In: *Physics Letters* (1980).

[AI95] L. Arnold and P. Imkeller. "Furstenberg–Khasminskii formulas for Lyapunov exponents via anticipative calculus". In: *Stochastics Stoch. Reports* (1995), pp. 127–168.

[AP15] J. Anderson and A. Papachristodoulou. "Advances in computational Lyapunov analysis using sum-of-squares programming". In: *Discrete & Continuous Dynamical Systems-B* 8 (2015), p. 2361.

[AX06] S. Ariaratnam and W.-C. Xie. "Lyapunov exponents in stochastic structural dynamics". In: Springer. 2006.

[Bar05] E. A. Barabanov. "Singular exponents and properness criteria for linear differential systems". In: *Differential Equations* 2 (2005), pp. 151–162.

[BBPS22] J. Bedrossian, A. Blumenthal, and S. Punshon-Smith. "A regularity method for lower bounds on the Lyapunov exponent for stochastic differential equations". In: *Invent. Math.* (2022), pp. 429–516.

[BEG93] S. Boyd and L. El Ghaoui. "Method of centers for minimizing generalized eigenvalues". In: *Linear algebra and its applications* (1993), pp. 63–111.

[Ben+80] G. Benettin et al. "Lyapunov characteristic exponents for smooth dynamical systems and for Hamiltonian systems; a method for computing all of them. Part 1: Theory". In: *Meccanica* 1 (1980), pp. 9–20.

[BF06] E. A. Barabanov and E. I. Fominykh. "Description of the mutual arrangement of singular exponents of a linear differential systems and exponents of its solutions". In: *Differential Equations* 12 (2006), pp. 1657–1673.

[BG19] J. Bochi and E. Garibaldi. "Extremal norms for fiber-bunched cocycles". In: *Journal de l'École polytechnique et Mathématiques* (2019), pp. 947–1004.

[BG20] J. J. Bramburger and D. Goluskin. "Minimum wave speeds in monostable reaction-diffusion equations: sharp bounds by polynomial optimization". In: *Proc. R. Soc. A* (2020), p. 20200450.

[BLM12] V. Bergelson, A. Leibman, and C. Moreira. "From discrete-to continuous-time ergodic theorems". In: *Ergodic Theory and Dynamical Systems* 2 (2012), pp. 383–426.

[Boc17] J. Bochi. "Ergodic optimization of Birkhoff averages and Lyapunov exponents". In: *Preprint: arXiv:1501.00961* (2017).

[Boc19] J. Bochi. "Genericity of periodic maximization: Proof of Contreras theorem following Huang, Lian, Ma, Xu, and Zhang". In: *ResearchGate* (2019).

[BR07] V. I. Bogachev and M. A. S. Ruas. *Measure theory*. Springer, 2007.

[BV04] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[BV17] L. Barreira and C. Valls. "Lyapunov regularity via singular values". In: *Transactions of the American Mathematical Society* 12 (2017), pp. 8409–8436.

[Che+14] S. I. Chernyshenko et al. "Polynomial sum of squares in fluid dynamics: a review with a look ahead". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 2020 (2014), p. 20130350.

[CKH22] V. Cibulka, M. Korda, and T. Hanis. "Spatio-temporal decomposition of sum-of-squares programs for the region of attraction and reachability". In: *IEEE Control Syst. Lett.* (2022), pp. 812–817.

[CY14] Y. Chen and Q. Yang. "Dynamics of a hyperchaotic Lorenz-type system". In: *Nonlinear Dynamics* 3 (2014), pp. 569–581.

[DG95a] C. R. Doering and J. Gibbon. "On the shape and dimension of the Lorenz attractor". In: *Dynamics and Stability of Systems* 3 (1995), pp. 255–268.

[DG95b] C. R. Doering and J. D. Gibbon. *Applied analysis of the Navier-Stokes equations*. Cambridge University Press, 1995.

[Din06] J. B. Dingwell. "Lyapunov exponents". In: *Wiley encyclopedia of biomedical engineering* (2006).

[DM22] C. R. Doering and A. McMillan. "Optimal time averages in non-autonomous nonlinear dynamical systems". In: *J. Pure Appl. Funct. Anal.* (2022), pp. 231–251.

[Don18] C. Dong. “Topological classification of periodic orbits in Lorenz system”. In: *Chinese Physics B* 8 (2018), p. 080501.

[DRVV97] L. Dieci, R. D. Russell, and E. S. Van Vleck. “On the compuation of Lyapunov exponents for continuous dynamical systems”. In: *SIAM journal on numerical analysis* 1 (1997), pp. 402–423.

[Eat07] M. L. Eaton. “Topological groups and invariant measures”. In: *Multivariate Statistics*. Institute of Mathematical Statistics, 2007, pp. 184–233.

[EFT91] A. Eden, C. Foias, and R. Temam. “Local and global Lyapunov exponents”. In: *Journal of Dynamics and Differential Equations* 1 (1991), pp. 133–177.

[ER85] J.-P. Eckmann and D. Ruelle. “Ergodic theory of chaos and strange attractors”. In: *The theory of chaotic attractors*. Springer, 1985, pp. 273–312.

[EW13] M. Einsiedler and T. Ward. “Ergodic theory”. In: *Springer* (2013).

[Fan+16] G. Fantuzzi et al. “Bounds for deterministic and stochastic dynamical systems using sum-of-squares optimization”. In: *SIAM Journal on Applied Dynamical Systems* 4 (2016), pp. 1962–1988.

[Fan19] G. Fantuzzi. *Aeroimperial-Yalmip*. 2019. URL: https://github.com/aeroimperial-optimization/aeroimperial-yalmip.

[FG20] G. Fantuzzi and D. Goluskin. “Bounding extreme events in nonlinear dynamics using convex optimization”. In: *SIAM journal on applied dynamical systems* 3 (2020), pp. 1823–1864.

[FGC22] F. Fuentes, D. Goluskin, and S. Chernyshenko. “Global stability of fluid flows despite transient growth of energy”. In: *Phys. Rev. Lett.* 128 (2022), p. 204502.

[FK83] H. Furstenberg and Y. Kifer. “Random matrix products and measures on projective spaces”. In: *Israel Journal of Mathematics* 1 (1983), pp. 12–32.

[FR90] E. Forest and R. D. Ruth. “Fourth-order symplectic integration”. In: *Physica D: Nonlinear Phenomena* 1 (1990), pp. 105–117.

[Fra+90] G. Frank et al. “Chaotic time series analyses of epileptic seizures”. In: *Physica D: Nonlinear Phenomena* 3 (1990), pp. 427–438.

[Fre+83] P. Frederickson et al. “The Liapunov dimension of strange attractors”. In: *Journal of Differential Equations* 2 (1983), pp. 185–207.

[FS95] F. Faisal and U. Schwengelbeck. “Unified theory of Lyapunov exponents and a positive example of deterministic quantum chaos”. In: *Physics Letters A* 1-2 (1995), pp. 31–36.

[Gal98] Z. Galias. “Rigorous numerical studies of the existence of periodic orbits for the Hénon map.” In: *J. Univers. Comput. Sci.* 2 (1998), pp. 114–124.

[GF19] D. Goluskin and G. Fantuzzi. “Bounds on mean energy in the Kuramoto–Sivashinsky equation computed using semidefinite programming”. In: *Nonlinearity* 5 (2019), p. 1705.

[GH14] P. Giesel and S. Hafstein. “Computation of Lyapunov functions for nonlinear discrete systems by linear programming”. In: *Journal of Difference Equations and Applications* (2014), pp. 610–640.

[Gol18] D. Goluskin. “Bounding averages rigorously using semidefinite programming: Mean moments of the Lorenz system”. In: *Journal of Nonlinear Science* 2 (2018), pp. 621–651.

[Gol20] D. Goluskin. “Bounding extrema over global attractors using polynomial optimization”. In: *Nonlinearity* (2020), pp. 4878–4899.

[GP04] K. Gatermann and P. A. Parrilo. “Symmetry groups, semidefinite programs, and sums of squares”. In: *Journal of Pure and Applied Algebra* 1-3 (2004), pp. 95–128.

[GPL90] K. Geist, U. Parlitz, and W. Lauterborn. “Comparison of different methods for computing Lyapunov exponents”. In: *Progress of theoretical physics* 5 (1990), pp. 875–893.

[Haf+18] S. Hafstein et al. “Lyapunov function computation for autonomous linear stochastic differential equations using sum-of-squares programming”. In: *Discrete and Continuous Dynamical Systems-Series B* 2 (2018), pp. 939–956.

[Hai10] E. Hairer. “Lecture 2: Symplectic integrators”. In: *Geometric Numerical Integration* (2010).

[Hén76] M. Hénon. “A two-dimensional mapping with a strange attractor”. In: *The Theory of Chaotic Attractors*. Springer, 1976, pp. 94–102.

[HH64] M. Hénon and C. Heiles. “The applicability of the third integral of motion: some numerical experiments”. In: *The Astronomical Journal* (1964), p. 73.

[HK14] D. Henrion and M. Korda. “Convex computation of the region of attraction of polynomial control systems”. In: *IEEE Trans. Automat. Contr.* (2014), pp. 297–312.

[HUP97] F Hubertus, F. E. Udwadia, and W. Proskurowski. “An efficient QR based method for the computation of Lyapunov exponents”. In: *Physica D: Nonlinear Phenomena* 1-2 (1997), pp. 1–16.

[Jen19] O. Jenkinson. “Ergodic optimization in dynamical systems”. In: *Ergodic Theory and Dynamical Systems* 10 (2019), pp. 2593–2618.

[Kal11] B. Kalinin. “Livšic theorem for matrix cocycles”. In: *Annals of mathematics* (2011), pp. 1025–1042.

[KB37] N. Kryloff and N. Bogoliouboff. “La théorie générale de la mesure dans son application à l’étude des systèmes dynamiques de la mécanique non linéaire”. In: *Annals of mathematics* (1937), pp. 65–113.

[KHJ13] M. Korda, D. Henrion, and C. N. Jones. “Convex computation of the maximum controlled invariant set for discrete-time polynomial control systems”. In: *52nd IEEE Conference on Decision and Control*. IEEE. 2013, pp. 7107–7112.

[KHJ14] M. Korda, D. Henrion, and C. N. Jones. “Convex computation of the maximum controlled invariant set for polynomial control systems”. In: *SIAM Journal on Control and Optimization* 5 (2014), pp. 2944–2969.

[KHM21] M. Korda, D. Henrion, and I. Mezić. “Convex computation of extremal invariant measures of nonlinear dynamical systems and Markov processes”. In: *J. Nonlinear Sci.* (2021), p. 14.

[Kle13] A. Klenke. *Probability theory: a comprehensive course*. Springer Science & Business Media, 2013.

[Koo06] W. Koon. *Poincaré map, floquet theory, and stability of periodic orbits*. Tech. rep. Technical report, Control and Dynamical Systems: California Institute of Technology, 2006.

[Kor22] M. Korda. “Stability and performance verification of dynamical systems controlled by neural networks: algorithms and complexity”. In: *IEEE Control Systems Letters* (2022), pp. 3265–3270.

[KP10] W. G. Kelley and A. C. Peterson. *The theory of differential equations: classical and qualitative*. Springer Science & Business Media, 2010.

[Kun+16] J. Kuntz et al. "Bounding stationary averages of polynomial diffusions via semidefinite programming". In: *SIAM Journal on Scientific Computing* 6 (2016), A3891–A3920.

[LA08] J. Lavaei and A. G. Aghdam. "Robust stability of LTI systems over semialgebraic sets using sum-of-squares matrix polynomials". In: *IEEE Transactions on Automatic Control* 1 (2008), pp. 417–423.

[Lak+20] M. V. Lakshmi et al. "Finding extremal periodic orbits with polynomial optimization, with application to a nine-mode model of shear flow". In: *SIAM Journal on Applied Dynamical Systems* 2 (2020), pp. 763–787.

[Las01] J. B. Lasserre. "Global optimization with polynomials and the problem of moments". In: *SIAM J. Optim.* (2001), pp. 796–817.

[Las+08] J. B. Lasserre et al. "Nonlinear optimal control via occupation measures and LMI-relaxations". In: *SIAM journal on control and optimization* 4 (2008), pp. 1643–1666.

[Lax02] P. D. Lax. *Functional analysis*. John Wiley & Sons, 2002.

[Li+11] Y. Li et al. "A new hyperchaotic Lorenz-type system: Generation, analysis, and implementation". In: *International Journal of Circuit Theory and Applications* 8 (2011), pp. 865–879.

[Lla86] J. G. Llavona. *Approximation of continuously differentiable functions*. Elsevier, 1986.

[Löf09] J. Löfberg. "Pre- and post-processing sum-of-squares programs in practice". In: *IEEE Transactions on Automatic Control* 5 (2009), pp. 1007–1011.

[Löf16] J. Löfberg. *Sum-of-squares programming*. https://yalmip.github.io/tutorial/sumofsquaresprogramming/. 2016.

[Lor63] E. N. Lorenz. "Deterministic nonperiodic flow". In: *Journal of the atmospheric sciences* (1963), pp. 130–141.

[Lya92a] A. Lyapunov. *General problem of the stability of motion (translated from Russian to French by E. Davaux and from French to English by AT Fuller)*. 1992.

[Lya92b] A. M. Lyapunov. "The general problem of motion stability". In: *Annals of Mathematics Studies* (1892).

[Mas83] W. Massey. "Cross products of vectors in higher dimensional Euclidean spaces". In: *The American Mathematical Monthly* 10 (1983), pp. 697–701.

[Mei17] J. D. Meiss. *Differential Dynamical Systems*. SIAM, 2017.

[MK87] K. G. Murty and S. N. Kabadi. "Some NP-complete problems in quadratic and nonlinear programming". In: *Mathematical programming* 2 (1987), pp. 117–129.

[Moh19] R. Mohammadpour. "Zero temperature limits of equilibrium states for subadditive potentials and approximation of maximal Lyapunov exponent". In: *arXiv preprint arXiv:1910.07279* (2019).

[Mot67] T. S. Motzkin. "The arithmetic-geometric inequality". In: *Inequalities (Proc. Sympos. Wright-Patterson Air Force Base, Ohio, 1965)* (1967), pp. 205–224.

[Nes00] Y. Nesterov. "Squared functional systems and optimization problems". In: *High Perform. Optim.* Ed. by H. Frenk et al. Springer, 2000, pp. 405–440.

[NS07] J. Nie and M. Schweighofer. "On the complexity of Putinar's Positivstellensatz". In: *Journal of Complexity* 1 (2007), pp. 135–150.

[Ols+20] M. L. Olson et al. "Heat transport bounds for a truncated model of Rayleigh–Bénard convection via polynomial optimization". In: *Phys. D* (2020), p. 132748.

[Ose68] V. I. Oseledets. "A multiplicative ergodic theorem. Characteristic Liapunov, exponents of dynamical systems". In: *Trudy Moskovskogo Matematicheskogo Obshchestva* (1968), pp. 179–210.

[Par00] P. A. Parrilo. "Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization". PhD thesis. California Institute of Technology, 2000.

[Par13] P. A. Parrilo. "Polynomial optimization, sums of squares, and applications". In: *Semidefinite optimization and convex algebraic geometry*. Ed. by G. Blekherman, P. A. Parrilo, and R. R. Thomas. SIAM, 2013. Chap. 3, pp. 47–157.

[PC12] T. S. Parker and L. Chua. *Practical numerical algorithms for chaotic systems*. Springer Science & Business Media, 2012.

[PC89] T. S. Parker and L. O. Chua. *Practical Numerical Algorithms for Chaotic Systems*. New York: Springer, 1989.

[Per13] L. Perko. *Differential equations and dynamical systems*. Springer Science & Business Media, 2013.

[PGV21] J. P. Parker, D. Goluskin, and G. M. Vasil. "A study of the double pendulum using polynomial optimization". In: *Chaos* (2021), p. 103102.

[Pin88] C. Ping. "Empirical and theoretical evidence of economic chaos". In: *System Dynamics Review* 1-2 (1988), pp. 81–108.

[PP16] A. Pikovsky and A. Politi. *Lyapunov exponents*. United Kingdom: Cambridge University Press, 2016.

[Put93] M. Putinar. “Positive polynomials on compact semi-algebraic sets”. In: *Indiana University Mathematics Journal* 3 (1993), pp. 969–984.

[PY19] D. Papp and S. Yildiz. “Sum-of-squares optimization without semidefinite programming”. In: *SIAM Journal on Optimization* 1 (2019), pp. 822–851.

[Rud+76] W. Rudin et al. *Principles of mathematical analysis*. McGraw-hill New York, 1976.

[Sal16] D. Salamon. *Measure and integration*. Citeseer, 2016.

[Sio58] M. Sion. “On general minimax theorems.” In: (1958).

[SM03] I. I. Shevchenko and A. Mel’nikov. “Lyapunov exponents in the Hénon-Heiles problem”. In: *Journal of Experimental and Theoretical Physics Letters* 12 (2003), pp. 642–646.

[SN12] J. Starrett and C. Nicholas. “A suspension of the Hénon map by periodic orbits”. In: *Chaos, Solitons & Fractals* 12 (2012), pp. 1486–1493.

[Spi18] M. Spivak. *Calculus on manifolds: a modern approach to classical theorems of advanced calculus*. CRC press, 2018.

[Spr94] J. C. Sprott. “Some simple chaotic flows”. In: *Physical review E* 2 (1994), R647.

[Str15] S. H. Strogatz. *Nonlinear Dynamics and Chaos*. Boulder, Colorado: Westview Press, 2015.

[Tem12] R. Temam. *Infinite-dimensional dynamical systems in mechanics and physics*. Springer Science, 2012.

[TGD18] I. Tobasco, D. Goluskin, and C. R. Doering. “Optimal bounds and extremal trajectories for time averages in nonlinear dynamical systems”. In: *Physics Letters A* 6 (2018), pp. 382–386.

[TP08] W. Tan and A. Packard. “Stability region analysis using polynomial and composite polynomial Lyapunov functions and sum-of-squares programming”. In: *IEEE Transactions on Automatic Control* 2 (2008), pp. 565–571.

[TSR01] G. Tancredi, A. Sánchez, and F. Roig. “A comparison between methods to compute Lyapunov exponents”. In: *The Astronomical Journal* 2 (2001), p. 1171.

[Vog+20] R. Vogt et al. “On Lyapunov exponents for rnns: Understanding information propagation using dynamical systems tools”. In: *arXiv preprint arXiv:2006.14123* (2020).

[Wal00] P. Walters. *An introduction to ergodic theory*. Springer Science & Business Media, 2000.

[Wan+21] J. Wang et al. “Exploiting term sparsity in moment-sos hierarchy for dynamical systems”. In: *arXiv preprint arXiv:2111.08347* (2021).

[Wen14] H. Wen. “A review of the Hénon map and its physical interpretations”. In: *School of Physics Georgia Institute of Technology, Atlanta, GA* (2014), pp. 30332–0430.

[Zgl97] P. Zgliczynski. “Computer assisted proof of chaos in the Rössler equations and in the Hénon map”. In: *Nonlinearity* 1 (1997), p. 243.

Appendix A

Appendix

B.1 Symmetries in continuous-time systems

Proposition 18 generalizes [GF19, Proposition A.1], which was proved in the context of simple finite symmetry groups, to all subgroups of $O(n)$. Proposition 19 generalizes [Lak+20, Theorem 2] in the same way. The significance of these propositions is that symmetry can be imposed on the function V being optimized over to bound time averages in (2.1.3), or on V and the other tunable polynomials in the SOS relaxation (2.3.11), without changing the resulting optimal bound on $\overline{\Phi}$. Imposing such symmetry ensures the same symmetry in the polynomial expressions that are constrained to be SOS, and the latter can be exploited for faster and more accurate numerical solutions.

Proposition 18. *Let $\mathcal{G}$ be a subgroup of $O(n)$. Suppose the set $\mathcal{B} \subset \mathbb{R}^n$ is $\mathcal{G}$-invariant and the integrable functions $f : \mathcal{B} \to \mathbb{R}^n$ and $\Phi : \mathcal{B} \to \mathbb{R}$ are $\mathcal{G}$-equivariant and $\mathcal{G}$-invariant, respectively. If there exists $V \in \mathcal{C}^1(\mathcal{B})$ and $B \in \mathbb{R}$ such that $B - \Phi(x) - f(x) \cdot \nabla V(x) \geq 0$ for all $x \in \mathcal{B}$, then there exists a $\mathcal{G}$-invariant $\widehat{V} \in \mathcal{C}^1(\mathcal{B})$ satisfying the same inequality.*

Proof. Since $\mathcal{G} \subset O(n)$ is compact, there exists a Haar probability measure $m(\Lambda)$ on

$\mathcal{G}$ that is right-invariant, meaning that $m(A) = m(A\Lambda)$ for all measurable $A \subset \mathcal{G}$ and all $\Lambda \in \mathcal{G}$. Define the symmetrization of V as

$$\widehat{V}(x) = \int_{\mathcal{G}} V(\Lambda' x)\, dm(\Lambda'), \tag{B.1.1}$$

and note that $\widehat{V}$ is $\mathcal{G}$-invariant because it is an integral against a right-invariant measure. Since $B - \Phi(x) - f(x) \cdot \nabla V(x) \geq 0$ holds for all $x \in \mathcal{B}$ and $\mathcal{B}$ is $\mathcal{G}$-invariant, the same inequality holds when x is replaced by Λx for any $\Lambda \in \mathcal{G}$. This inequality evaluated at Λx, after using the invariance of Φ and equivariance of f under Λ, becomes

$$B - \Phi(x) - \Lambda f(x) \cdot \nabla V(\Lambda x) \geq 0 \quad \forall x \in \mathcal{B}. \tag{B.1.2}$$

The gradient in (B.1.2) is with respect to the entire argument Λx, but this alternatively can be expressed using a gradient with respect to x since

$$\Lambda f(x) \cdot \nabla_{x'} V(x')|_{x' = \Lambda x} = f(x) \cdot \nabla_x V(\Lambda x), \tag{B.1.3}$$

because Λ is orthogonal by assumption. Substituting the right-hand side of (B.1.3) into (B.1.2) and integrating over the group $\mathcal{G}$ against $m(\Lambda)$ gives the desired inequality with the symmetrization (B.1.1),

$$B - \Phi(x) - f(x) \cdot \nabla \widehat{V}(x) \geq 0 \quad \forall x \in \mathcal{B}. \tag{B.1.4}$$

□

Proposition 19. *Let $\mathcal{G}$ be a subgroup of $O(n)$. Suppose all $h_i, g_j \in \mathbb{R}[x]$, defining $\mathcal{B}$ via* (2.3.3)*, are $\mathcal{G}$-invariant, $\Phi \in \mathbb{R}[x]$ is $\mathcal{G}$-invariant, and $f \in \mathbb{R}^n[x]$ is $\mathcal{G}$-equivariant.*

If there exist $V, \rho_i, \sigma_j \in \mathbb{R}[x]$ *and* $B \in \mathbb{R}$ *such that* $B - \Phi - f \cdot \nabla V - \sum_{i=1}^{I} \rho_i h_i - \sum_{j=1}^{J} \sigma_j g_j \in \Sigma_n$ *and* $\sigma_j \in \Sigma_n$ *for all* $j \in \{1, \dots, J\}$*, then there exist* $\mathcal{G}$*-invariant* $\widehat{V}, \widehat{\rho}_i, \widehat{\sigma}_j \in \mathbb{R}[x]$ *such that the corresponding expressions belong to* Σ_n*.*

Proof. Define $\widehat{V}$ as the $\mathcal{G}$-invariant symmetrization of V using the right-invariant Haar probability measure $m(\Lambda)$, as in (B.1.1), and define $\widehat{\rho}_i$ and $\widehat{\sigma}_j$ as the analogous symmetrizations of ρ_i and σ_j, respectively. To see that $\sigma_j \in \Sigma_n$ implies $\widehat{\sigma}_j \in \Sigma_n$, let $\mathbf{b}(x)$ be a vector of all monomials with degree up to half the degree of σ_j. There exists a symmetric matrix Q_{σ_j} such that $\sigma_j(x) = \mathbf{b}(x)^\mathsf{T} Q_{\sigma_j} \mathbf{b}(x)$, and because $\sigma_j \in \Sigma_n$ it is possible to choose $Q_{\sigma_j} \succeq 0$. For each $\Lambda \in \mathcal{G}$, every entry in $\mathbf{b}(\Lambda x)$ is a linear combination of the entries of $\mathbf{b}(x)$, so there exists a matrix $\Gamma(\Lambda)$ such that $\mathbf{b}(\Lambda x) = \Gamma(\Lambda)\mathbf{b}(x)$. Therefore,

$$\widehat{\sigma}_j(x) = \mathbf{b}(x)^\mathsf{T} \underbrace{\left[\int_{\mathcal{G}} \Gamma(\Lambda')^\mathsf{T} Q_{\sigma_j} \Gamma(\Lambda')\, dm(\Lambda') \right]}_{\widehat{Q}_{\sigma_j}} \mathbf{b}(x). \tag{B.1.5}$$

This formula makes clear that $\widehat{\sigma}_j$ is not only a polynomial but is SOS: $\Gamma(\Lambda)^\mathsf{T} Q_{\sigma_j} \Gamma(\Lambda) \succeq 0$ for each $\Lambda \in \mathcal{G}$, and so $\widehat{Q}_{\sigma_j} \succeq 0$. Similar arguments confirm that $\widehat{V}$ and $\widehat{\rho}_i$ are polynomials but not that they are SOS since V and ρ_i need not be.

The argument for the first SOS constraint is similar to the argument for $\widehat{\sigma}_j$. There exist a (possibly different) basis vector $\mathbf{b}(x)$ and matrix $Q \succeq 0$ such that

$$B - \Phi - f \cdot \nabla V - \sum_{i=1}^{I} \rho_i h_i - \sum_{j=1}^{J} \sigma_j g_j = \mathbf{b}^\mathsf{T} Q\, \mathbf{b} \tag{B.1.6}$$

since the left-hand expression is assumed to be SOS. The above equality holds for all $x \in \mathbb{R}^n$, including with x replaced by Λx. We replace x by Λx and use the

Λ-invariance of Φ, h_i, g_j, the Λ-equivariance of f, and (B.1.3) to find

$$B - \Phi(x) - f(x) \cdot \nabla_x V(\Lambda x) - \sum_{i=1}^{I} \rho_i(\Lambda x) h_i(x) - \sum_{j=1}^{J} \sigma_j(\Lambda x) g_j(x)$$
$$= \mathbf{b}(x)^{\mathsf{T}} \Gamma(\Lambda)^{\mathsf{T}} Q \, \Gamma(\Lambda) \mathbf{b}(x). \tag{B.1.7}$$

As in (B.1.5), integrating both sides of the above equality over the group $\mathcal{G}$ against the right-invariant Haar probability measure $m(\Lambda)$ gives the desired expression:

$$B - \Phi - f \cdot \nabla \widehat{V} - \sum_{i=1}^{I} \widehat{\rho}_i h_i - \sum_{j=1}^{J} \widehat{\sigma}_j g_j = \mathbf{b}^{\mathsf{T}} \underbrace{\left[\int_{\mathcal{G}} \Gamma(\Lambda')^{\mathsf{T}} Q \Gamma(\Lambda') \, dm(\Lambda') \right]}_{\widehat{Q}} \mathbf{b}, \tag{B.1.8}$$

where $\widehat{Q} \succeq 0$ and thus the left-hand expression is SOS as needed. □

B.2 Symmetries in discrete-time systems

Proposition 20 is the analogue of Proposition 18. Proposition 21 is the analogue of Proposition 19. The significance of these propositions is that symmetry can be imposed on the function V being optimized over to bound geometric means in (6.1.10), or on V and the other tunable polynomials in the SOS program (7.2.5), without changing the resulting optimal bound on $\widetilde{\Phi}$. Imposing such symmetry ensures the same symmetry in the polynomial expressions that are constrained to be SOS, and the latter can be exploited for faster and more accurate numerical solutions. Note that Propositions 20 and 21 are slightly more general than Propositions 18 and 19, because the symmetry groups to which they apply only need to be compact, not orthogonal.

Proposition 20. *Let $\mathcal{G}$ be a compact subgroup of $GL(n)$. Suppose the set $\mathcal{B} \subset \mathbb{R}^n$*

is $\mathcal{G}$-invariant and the integrable functions $f : \mathcal{B} \to \mathbb{R}^n$ and $\Phi : \mathcal{B} \to \mathbb{R}$ are $\mathcal{G}$-equivariant and $\mathcal{G}$-invariant, respectively. If there exists $V \in \mathcal{C}(\mathcal{B})$ and $B \in \mathbb{R}$ such that $BV(x) - \Phi(x)V(f(x)) \geq 0$ for all $x \in \mathcal{B}$, then there exists a $\mathcal{G}$-invariant $\widehat{V} \in \mathcal{C}(\mathcal{B})$ satisfying the same inequality.

Proof. Since $\mathcal{G} \subset GL(n)$ is compact, there exists a Haar probability measure $m(\Lambda)$ on $\mathcal{G}$ that is right-invariant, meaning that $m(A) = m(A\Lambda)$ for all measurable $A \subset \mathcal{G}$ and all $\Lambda \in \mathcal{G}$. Define the symmetrization of V as in (B.1.1) so that $\widehat{V}$ is $\mathcal{G}$-invariant because it is an integral against a right-invariant measure. Since $BV(x) - \Phi(x)V(f(x)) \geq 0$ holds for all $x \in \mathcal{B}$ and $\mathcal{B}$ is $\mathcal{G}$-invariant, the same inequality holds when x is replaced by Λx for any $\Lambda \in \mathcal{G}$. This inequality evaluated at Λx, after using the invariance of Φ and equivariance of f under Λ, becomes

$$BV(\Lambda x) - \Phi(x)V(\Lambda f(x)) \geq 0 \quad \forall x \in \mathcal{B}. \tag{B.2.1}$$

Note that we are multiplying the entire argument of $V(x)$ by Λ, even when $V(x)$ is evaluated at $f(x)$. Therefore, integrating over the group $\mathcal{G}$ against $m(\Lambda)$ gives the desired inequality with the symmetrization (B.1.1),

$$B\widehat{V}(x) - \Phi(x)\widehat{V}(f(x)) \geq 0 \quad \forall x \in \mathcal{B}. \tag{B.2.2}$$

□

Proposition 21. *Let $\mathcal{G}$ be a compact subgroup of $GL(n)$. Suppose all $g_j \in \mathbb{R}[x]$, defining $\mathcal{B}$ via (6.3.1), are $\mathcal{G}$-invariant, $\Phi \in \mathbb{R}[x]$ is $\mathcal{G}$-invariant, and $f \in \mathbb{R}^n[x]$ is $\mathcal{G}$-equivariant. If there exist $V, \sigma_j \in \mathbb{R}[x]$ and $B \in \mathbb{R}$ such that $B\Phi^{d_z/2-1}V - \Phi^{d_z/2}V(f) - \sum_{j=1}^{J} \sigma_j g_j \in \Sigma_n$ and $\sigma_j \in \Sigma_n$ for all $j \in \{1, \ldots, J\}$, then there exist $\mathcal{G}$-invariant $\widehat{V}$ and $\widehat{\sigma}_j \in \mathbb{R}[x]$ such that the corresponding expressions belong to Σ_n also.*

Proof. Define $\widehat{V}$ as the $\mathcal{G}$-invariant symmetrization of V using the right-invariant Haar probability measure $m(\Lambda)$, as in (B.1.1), and define the $\widehat{\sigma}_j$ as the analogous symmetrizations of σ_j, respectively. To see that the $\sigma_j \in \Sigma_n$ implies $\widehat{\sigma}_j \in \Sigma_n$, let $\mathbf{b}(x)$ be a vector of all monomials with degree up to half the degree of σ_j. There exists a symmetric matrix Q_{σ_j} such that $\sigma_j(x) = \mathbf{b}(x)^\mathsf{T} Q_{\sigma_j} \mathbf{b}(x)$, and because $\sigma_j \in \Sigma_n$ it is possible to choose $Q_{\sigma_j} \succeq 0$. For each $\Lambda \in \mathcal{G}$, every entry in $\mathbf{b}(\Lambda x)$ is a linear combination of the entries of $\mathbf{b}(x)$, so there exists a matrix $\Gamma(\Lambda)$ such that $\mathbf{b}(\Lambda x) = \Gamma(\Lambda)\mathbf{b}(x)$. Therefore,

$$\widehat{\sigma}_j(x) = \mathbf{b}(x)^\mathsf{T} \underbrace{\left[\int_{\mathcal{G}} \Gamma(\Lambda')^\mathsf{T} Q_{\sigma_j} \Gamma(\Lambda')\, dm(\Lambda')\right]}_{\widehat{Q}_{\sigma_j}} \mathbf{b}(x). \tag{B.2.3}$$

This formula makes clear that $\widehat{\sigma}_j$ is not only a polynomial but is SOS: $\Gamma(\Lambda)^\mathsf{T} Q_{\sigma_j} \Gamma(\Lambda) \succeq 0$ for each $\Lambda \in \mathcal{G}$, and so $\widehat{Q}_{\sigma_j} \succeq 0$. Similar arguments confirm that $\widehat{V}$ is polynomial but not that it is SOS since V need not be.

The argument for the first SOS constraint is similar to the argument for $\widehat{\sigma}_j$. There exist a (possibly different) basis vector $\mathbf{b}(x)$ and matrix $Q \succeq 0$ such that

$$B\Phi^{d_z/2-1}V - \Phi^{d_z/2}V(f) - \sum_{j=1}^{J} \sigma_j g_j = \mathbf{b}^\mathsf{T} Q\, \mathbf{b} \tag{B.2.4}$$

since the left-hand expression is assumed to be SOS. The above equality holds for all $x \in \mathbb{R}^n$, including with x replaced by Λx. We do this and use the Λ-invariance of Φ, g_j, the Λ-equivariance of f, and (B.1.3) to find

$$B\Phi(x)^{d_z/2-1}V(\Lambda x) - \Phi(x)^{d_z/2}V\Big(\Lambda f(x)\Big) - \sum_{j=1}^{J} \sigma_j(\Lambda x) g_j(x) = \mathbf{b}(x)^\mathsf{T}\Gamma(\Lambda)^\mathsf{T} Q\, \Gamma(\Lambda)\mathbf{b}(x). \tag{B.2.5}$$

As in (B.2.3), integrating both sides of the above equality over the group $\mathcal{G}$ against the right-invariant Haar probability measure $m(\Lambda)$ gives the desired expression:

$$B\Phi^{d_z/2-1}\widehat{V} - \Phi^{d_z/2}\widehat{V}(f) - \sum_{j=1}^{J}\widehat{\sigma}_j g_j = \mathbf{b}^{\mathsf{T}} \underbrace{\left[\int_{\mathcal{G}} \Gamma(\Lambda')^{\mathsf{T}} Q \Gamma(\Lambda')\, dm(\Lambda')\right]}_{\widehat{Q}} \mathbf{b}, \tag{B.2.6}$$

where $\widehat{Q} \succeq 0$ and thus the left-hand expression is SOS as needed. □

Remark 4. *Proposition 21 is stated for the SOS constraint in* (7.2.5) *so that it can more easily be applied to the SOS program for bounding LEs, which uses a specialized SOS constraint whose construction is described in section 7.3. Notwithstanding, the results of Proposition 21 are equally valid when the constraint in it is replaced with* $BV - \Phi V(f) - \sum_{j=1}^{J}\sigma_j g_j \in \Sigma_n$ *and all of the other assumptions in the statement of the proposition hold, in which case it directly applies to the SOS program for bounding geometric means in Proposition 11.*