Recurrence of Nonlinear Control Systems: Entropy, Bit Rates, and Finite Alphabet Controllers

Hussein Sibai[a], Enrique Mallada[b]

[a]Computer Science and Engineering Department, Washington University in St. Louis,
[b]Electrical and Computer Engineering Department, Johns Hopkins University,

ARTICLE INFO

Keywords:
Entropy
Recurrence
Invariance
Control Systems

ABSTRACT

In this paper, we introduce the notion of recurrence entropy in the context of nonlinear control systems. A set is said to be (τ-)recurrent if every trajectory that starts in the set returns to it (within at most τ units of time). The recurrence entropy of a control system quantifies the complexity of making a set τ-recurrent measured by the average rate of growth, as time increases, of the number of control signals required to achieve this goal. Our analysis reveals that, compared to invariance, recurrence is quantitatively less complex, meaning that the recurrence entropy of a set is no larger than, and often strictly smaller than, the invariance entropy. We provide upper and lower bounds on recurrence entropy and show that they converge to the bounds on invariance entropy as τ decreases to zero. Further, our results show that recurrence entropy lower bounds the minimum data rate between the sensor and controller required for achieving recurrence. We present an algorithm according to which the sensor can send state estimates to the controller over a limited-bandwidth channel to achieve recurrence asymptotically at an exponential rate. Finally, we show that, under mild stricter conditions on the set and dynamics, the control signals that enforce the τ-recurrence of a set can be generated by a finite alphabet of control signals of durations of at most τ units of time, which allows us to store them for quick online execution.

1. Introduction

The topological entropy of a dynamical system is a fundamental property, an invariant (Katok and Hasselblatt, 1995), that describes the rate of the exponential growth of the number of trajectories that are distinguishable with arbitrarily small but finite accuracy. Originally proposed by Adler, Konheim, and McAndrew (Adler, Konheim and McAndrew, 1965), and shortly after reformulated in the form described above by Bowen (Bowen, 1971a,b), it provides a quantitative measure of complexity by capturing how the uncertainty around the system state grows as time evolves. As a result, topological entropy is closely related to information-theoretic notions, such as the average rate of information gathering about the system state above which one can distinguish its trajectories with arbitrary accuracy (Liberzon and Mitra, 2016).

In control theory, wherein one uses the system's state information to perform a task, several notions of entropy have been proposed in the literature, characterizing the complexity of and the minimal data rates necessary for performing a certain control task. Examples of this include estimation entropy (Savkin, 2006; Liberzon and Mitra, 2016; Sibai and Mitra, 2017, 2018, 2023; Kawan and Yüksel, 2018), restoration entropy (Matveev and Pogromsky, 2016, 2019), stabilization entropy (Delchamps, 1990; Nair, Evans, Mareels and Moran, 2004; Colonius, 2012), reachability entropy (Tomar and Zamani, 2022), among others. One notion of entropy particularly instrumental in control is the invariance entropy (Colonius and Kawan, 2009, 2011; Colonius, Kawan and Nair, 2013; Kawan and Delvenne, 2016; Rungger and Zamani, 2017; Tomar, Rungger and Zamani, 2021; Tomar, Kawan and Zamani, 2022), which aims to capture the growth rate of the number of distinct control signals necessary to render a certain set invariant for a period of time, as that period increases to infinity.

Invariance holds a prominent role in control theory. It is, for instance, a core notion in the development of the Lyapunov theory (Khalil, 2002). By trapping trajectories on sub-level sets of a Lyapunov function, one can guarantee boundedness and completeness of trajectories, stability, and even asymptotic or exponential stability via a gradual reduction of the value of the function. Invariant sets can also be used to estimate regions of attractions of an asymptotically stable equilibrium (Genesio, Tartaglia and Vicino, 1985). However, due to intrinsic coupling between the dynamics of

sibai@wustl.edu (H. Sibai); mallada@jhu.edu (E. Mallada)
ORCID(s): 0000-0001-0000-0000 (H. Sibai); 0000-0003-1568-1833 (E. Mallada)

the system and the geometry of the set, finding invariant sets and, by extension, Lyapunov functions is often difficult. Furthermore, in the context of controlled systems, it is not always possible to make a given set (controlled) invariant.

In this work, motivated by recent literature aimed at using the notion of recurrent sets as functional substitutes for invariant sets in control theory (Shen, Bichuch and Mallada, 2022; Siegelmann, Shen, Paganini and Mallada, 2023), we introduce the notion of recurrence entropy for nonlinear control systems. A set is said to be (τ-)recurrent if every trajectory that starts in the set returns to it (within at most τ units of time). Our analysis shows that recurrence, as a control task, is quantitatively less complex than invariance from the point of view that for a given set and dynamical system, the recurrence entropy is no larger than the invariance entropy. Furthermore, we provide upper and lower bounds for the recurrence entropy in terms of the upper box dimension of the set to be made recurrent and a local Lipschitz constant and divergence of the vector field. We also show that recurrence entropy is a lower bound on the minimum bit rate at which the sensor should send state estimates to the controller to render the set recurrent. We further present an algorithm according to which the sensor can choose the information it sends to the controller to achieve this recurrence task asymptotically. Notably, this algorithm results in a bit rate equal to the recurrence entropy upper bound that we derive, plus a linear term equal to the product of the system dimension and the desired rate of exponential convergence toward a recurrent trajectory. We finalize the paper by showing a striking result. By imposing slightly stricter assumptions on the properties of the set and the system capabilities, we show that the control signals necessary to make a set control τ-recurrent can be generated by a finite alphabet of control signals of length at most τ.

Related Work: Our work is closely related to the literature of invariance entropy (Colonius and Kawan, 2009, 2011; Colonius et al., 2013; Kawan and Delvenne, 2016; Rungger and Zamani, 2017; Tomar et al., 2021, 2022). Naturally, since every invariant set is (trivially) τ-recurrent, for all $\tau \geq 0$, the results presented therein apply for $\tau = 0$. Our work also relates to that of Tomar et al. (Tomar and Zamani, 2023). That work relates the minimal bit rates needed to enforce a regular safety property for a discrete-time dynamical system to the invariance entropy of a new system that combines the automaton defining the property and the original system. Particularly, τ-recurrence can be thought of as regular safety property, but as we define it here, it is for continuous-time dynamical systems. Relating our results with (Tomar and Zamani, 2023) would be an interesting future direction. It would also be interesting to design numerical methods such as those proposed in (Tomar et al., 2022) for invariance entropy to estimate recurrence entropy. On the other hand, many works from the formal methods community use discrete abstractions of continuous-time systems to design zero-order-hold controllers to achieve various tasks, mainly reach-avoid ones, that naturally have finite alphabets (Tabuada, 2009; Meyer, Yin, Brodtkorb, Arcak and Sørensen, 2020; Sibai, Huriot, Martin and Arcak, 2024; Rungger and Zamani, 2016). However, these works usually present algorithms that adjust the set being kept invariant when needed while synthesizing the controller. Moreover, they usually predefine the alphabet at the time of the abstraction, mainly by considering constant signals that are equal to the centers of a grid over the control set for a fixed time horizon. In contrast, we fix the set of states of concern while allowing arbitrary control alphabets. We then define fundamental metrics to quantify the minimal size of an alphabet needed to achieve invariance and recurrence as a mean to compare their complexities. We also relate these metrics with the notions of entropy of the system for the same tasks.

Organization of the Paper: The rest of the paper is organized as follows. In Section 2, we provide preliminary definitions regarding the system to be considered, as well as the notion of invariance entropy. In Section 3, we formally introduce the notion of recurrence to be studied in the paper, i.e., τ-recurrence (c.f. Definition 5), as well as the associated notion of entropy. We then introduce, in Section 4, a fundamental result that allows us to bound the distance from a set that recurrent trajectories can travel if they are required to come back to the set within τ units of time. A comparison between recurrence and invariance entropy is performed in Section 5, upper and lower bounds for recurrence entropy are provided in Section 6, and a relationship between entropy and data rates is formally established in Section 7. We introduce an algorithm that can make trajectories (asymptotically) τ-recurrent in Section 8. Finally, Section 9 introduces the notion of finite alphabet controllers and relates its cardinality with entropy notions. Concluding remarks are given in Section 10.

2. Preliminaries

Notation: We denote by $\|\cdot\|$ the supremum norm over $\mathbb{R}^n$, unless otherwise specified. If $N \in \mathbb{N}$, we denote by $[N]$ the set of all non-negative integers less than N. Fix an $\varepsilon > 0$ and a compact set $Q \subset \mathbb{R}^n$. We define $B_\varepsilon(Q) := \{y \in \mathbb{R}^n | \exists x \in Q, \|x - y\| \leq \varepsilon\}$. We also define $\lambda(Q)$ to be the Lebesgue measure of Q. If Q is a singleton set $\{x\}$, we abuse notation and denote $B_\varepsilon(Q)$ by $B_\varepsilon(x)$. We call it a ball, or a hyperrectangle, centered at x with radius ε. Given a set S, $\mathrm{cl}(S)$ denotes its closure. If S is finite, $|S|$ denotes its cardinality. If S is compact subset of $\mathbb{R}^n$, a δ-cover of S is a set of balls of radius δ whose union contains S. We abuse notation and call the set containing the centers of the balls the cover instead of the balls themselves. We denote by $grid(S, \delta)$ the δ-cover of S that is constructed with the centers of the balls 2δ apart and on axis-parallel lines. We also call it a δ-grid of S. We assume that the logarithm function is of base 2 throughout the paper. Consider a function with N arguments. If we replace its i^{th} argument with "$\cdot$" (i.e., dot) and its other arguments with constants, we mean the projection of that function to the one-dimensional domain of the i^{th} argument with the other ones fixed to the specified constants. If we replace an argument with a set in its domain, we mean the function defined only over that set in the domain.

2.1. System description

In this paper, we consider control systems that are defined as follows.

Definition 1. *Consider a nonlinear control system of the form:*

$$\dot{x}(t) = f(x(t), u(t)), \tag{1}$$

where $x(t) \in \mathbb{R}^n$, $u \in \mathcal{U}$, with $\mathcal{U}$ being a set of piece-wise continuous functions mapping $\mathbb{R}_{\geq 0}$ to a compact set $U \subset \mathbb{R}^m$, and the map f is locally Lipschitz. We often abuse notation and interchangeably refer to a function in $\mathcal{U}$ as well as an input vector in U by u.

2.2. Invariance entropy

In this section, we recall the definition of invariance entropy of system (1) from (Colonius and Kawan, 2009). It requires the definitions of controlled invariant sets, invariant trajectories, and invariance spanning sets.

Definition 2 (Controlled invariant sets (Colonius and Kawan, 2009))**.** *A set $Q \subseteq \mathbb{R}^m$ is* controlled invariant *for system (1) if $\forall x \in Q$, $\exists u \in \mathcal{U}$ such that for any $t \geq 0$, $\xi(x, u, t) \in Q$.*

We call controlled invariant sets invariant from hereafter for brevity.

Definition 3 ((T, ε, Q)-invariant trajectories (Colonius and Kawan, 2009))**.** *Fix any $\varepsilon \geq 0$, $T \geq 0$, compact set $Q \subset \mathbb{R}^n$, $x \in Q$, and $u \in \mathcal{U}$. The trajectory $\xi(x, u, \cdot)$ of system (1) is (T, ε, Q)-invariant, if for every $t \in [0, T]$, $\xi(x, u, t) \in B_\varepsilon(Q)$. If the condition is $\xi(x, u, t) \in Q$ instead, we say that ξ is (T, Q)-invariant.*

Fix two non-empty sets $K \subseteq Q \subset \mathbb{R}^n$, an $\varepsilon \geq 0$, and a $T \geq 0$. A set $S \subseteq \mathcal{U}$ is called an *invariance* (T, ε, K, Q)-*spanning* set if for any $x \in K$, there exists a $u \in S$, such that $\xi(x, u, [0, T])$ is (T, ε, Q)-invariant. Let $r_{\text{inv}}(T, \varepsilon, K, Q)$ be the minimal cardinality of such a set if it exists, and be equal to infinity otherwise. The *invariance* entropy of system (1) is defined in (Colonius and Kawan, 2009) as follows:

$$h_{\text{inv}}(K, Q) := \lim_{\varepsilon \searrow 0} \limsup_{T \to \infty} \frac{1}{T} \log r_{\text{inv}}(T, \varepsilon, K, Q). \tag{2}$$

If the trajectories are required to be (T, Q)-invariant in (2), then the minimal cardinality of the corresponding invariance spanning set is denoted in (Colonius and Kawan, 2009) by $r^*_{\text{inv}}(T, K, Q)$. If substituted in (2), the resulting entropy notion $h^*_{\text{inv}}(K, Q)$ is called the *strict invariance* entropy of system (1). When K is equal to Q, we drop the K argument in the definitions above.

3. τ-Recurrence Entropy

In this section, we define the main concept that we introduce in this paper: τ-recurrence entropy. Before being able to define it, we need to define controlled τ-recurrent sets, recurrent trajectories, and recurrence spanning sets, in parallel with the definitions preceding the definition of invariance entropy in the previous section.

3.1. Recurrence spanning sets and entropy

In the following definition, we introduce controlled τ-recurrent sets as compact subsets of the state space of system (1) which satisfy the following condition: for each state in such a set, there exists a control signal that drives the system to have a trajectory that visits the set at least once within each time interval of size τ. This concept generalizes the notion of τ-recurrent sets, first introduced in (Shen et al., 2022), to control systems. We then define the concept of $(T, \varepsilon, \tau, Q)$-recurrent trajectories, which are ones that return to $B_\varepsilon(Q)$ at least once within each time interval of size τ in the interval $[0, T]$.

Definition 4 (Controlled τ-recurrent sets)**.** *A set $Q \subseteq \mathbb{R}^m$ is* controlled τ-recurrent *for system (1), for some finite $\tau \in \mathbb{R}^{\geq 0}$, if for every $x \in Q$, there exists a $u \in \mathcal{U}$ such that for any $t \in \mathbb{R}^{\geq 0}$, there exists a $t' \in [t, t+\tau]$ such that $\xi(x, u, t') \in Q$.*

We call controlled τ-recurrent sets τ-recurrent from hereafter for brevity.

Definition 5 ($(T, \varepsilon, \tau, Q)$-recurrent trajectories)**.** *Fix any $\tau \geq 0$, $\varepsilon \geq 0$, $T \geq \tau$, compact set $Q \subset \mathbb{R}^n$, $x \in Q$, and $u \in \mathcal{U}$. The trajectory $\xi(x, u, \cdot)$ of system (1) is $(T, \varepsilon, \tau, Q)$-recurrent, if for every $t \in [0, T-\tau]$, there exists a $t' \in [t, t+\tau]$ such that $\xi(x, u, t') \in B_\varepsilon(Q)$.*

For simplicity of notation, if $\varepsilon = 0$, we drop the ε argument. Similarly, if $T = \infty$, we drop the T argument. We will also use Definition 5 for piece-wise continuous functions of time that are not necessarily trajectories of system (1).

The final definition before that of τ-recurrence entropy is that of spanning sets. They are sets of control signals which are sufficient to make any trajectory starting from a τ-recurrent set τ-recurrent.

Fix a $\tau \in \mathbb{R}^{\geq 0}$, a compact τ-recurrent set $Q \subset \mathbb{R}^n$, an $\varepsilon \geq 0$, and a $T \geq 0$. A set $S \subseteq \mathcal{U}$ is called a *recurrence $(T, \varepsilon, \tau, Q)$-spanning* set if for any $x \in Q$, there exists a $u \in S$ such that $\xi(x, u, [0, T])$ is $(T, \varepsilon, \tau, Q)$-recurrent. Let $r_{\text{rec}}(T, \varepsilon, \tau, Q)$ be the minimal cardinality of such a set if it exists, and be equal to infinity otherwise. We define the *τ-recurrence* entropy of system (1) as follows:

$$h_{\text{rec}}(\tau, Q) := \lim_{\varepsilon \searrow 0} \limsup_{T \to \infty} \frac{1}{T} \log r_{\text{rec}}(T, \varepsilon, \tau, Q). \tag{3}$$

If we require the trajectories to be (T, τ, Q)-recurrent in (3), then we denote the minimal cardinality of the corresponding spanning set $r^*_{\text{rec}}(T, \tau, Q)$. If substituted in (3), we call the resulting entropy notion $h^*_{\text{rec}}(\tau, Q)$ the *strict τ-recurrence* entropy of system (1).

4. Containment Lemma

In this section, we show how trajectories that are (τ, Q)-recurrent cannot depart arbitrarily from Q. The following assumption is instrumental in achieving this goal.

Assumption 1 (τ-completeness)**.** *For any $x \in Q$ and $u \in \mathcal{U}$, the trajectory $\xi(x, u, \cdot)$ is defined for all $t \in [0, \tau]$ and is continuous in its first argument.*

An immediate consequence of Assumption 1 is that for any $u \in \mathcal{U}$, the closure reachable set $R(Q, u, \tau) := \cup_{t \in [0,\tau], x \in Q} \xi(x, u, t)$ of system (1), i.e., $\text{cl}(R(Q, u, \tau))$, is compact. Moreover, it follows from Proposition 5.2 (Lin, Sontag and Wang, 1996), that under Assumption 1, the set

$$R(Q, \tau) := \bigcup_{u \in \mathcal{U}} R(Q, u, \tau)$$

is bounded. The set $R(Q, \tau)$ contains all states visited by trajectories starting from *some* initial state $x \in Q$ and following *some* control $u \in \mathcal{U}$ for τ seconds. While such a set is, indeed, bounded, it may be quite big, as not all control inputs are meant to make trajectories recurrent. We will therefore consider the subset $\mathcal{U}_r \subseteq \mathcal{U}$ containing all control inputs $u \in \mathcal{U}$ such that for some $x \in Q$, the trajectory $\xi(x, u, \cdot)$ is (τ, Q)-recurrent.

A similar reasoning as before, using the fact that $\mathcal{U}_r \subseteq \mathcal{U}$, leads to fact that

$$R_r(Q, \tau) := \bigcup_{u \in \mathcal{U}_r} R(Q, u, \tau) \subseteq R(Q, \tau), \tag{4}$$

is bounded. For the purpose of estimating how far out (τ, Q)-recurrent trajectories can reach starting from a compact set Q, we define

$$L_\tau = \max_{x_1, x_2 \in \text{cl}(R_r(Q,\tau)), u \in U} \frac{\|f(x_1, u) - f(x_2, u)\|}{\|x_1 - x_2\|} < \infty. \tag{5}$$

Note that L_τ is an upper bound of the Lipschitz constant of the vector field along any (τ, Q)-recurrent trajectory.

The following lemma, which is a generalization of Lemma 2 in (Siegelmann et al., 2023), allows us to obtain an upper bound on how far trajectories can go outside Q.

Lemma 1 (Containment Lemma). *Consider a compact set Q. Then, given any $x \in Q$, and $u \in \mathcal{U}_r$ such that $\xi(x,u,\cdot)$ is (T,τ,Q)-recurrent, the following holds:*

$$\sup_{t\in[0,T]} d(\xi(x,u,t),Q) \le F_Q \tau e^{L_\tau \tau}, \tag{6}$$

where $d(y,Q) := \min_{x\in Q} \|y-x\|$, L_τ is given in (5), and

$$F_Q := \sup_{x\in Q,\, u\in U} \|f(x,u)\| < \infty.$$

Proof. As mentioned before, the proof of this lemma is akin to (Siegelmann et al., 2023), Lemma 2. Given $x \in Q$ and the corresponding $u \in \mathcal{U}_r$ that makes $\xi(x,u,\cdot)$ (T,τ,Q)-recurrent, let $t_1 > 0$ be the first time the trajectory leaves Q, i.e., such that $\xi(x,u,t) \in Q$, for $t \le t_1$, and for all sufficiently small $\delta > 0$, $\xi(x,u,t+\delta) \notin Q$. Without loss of generality, we assume $t_1 < T$. It then follows from the assumption that the trajectory is (T,τ,Q)-recurrent that for all $t \in [0, \min\{t_1+\tau, T\}]$, $\xi(x,u,t)$ can only be outside Q for at most τ seconds. Using now the short notation $x(t) = \xi(x,u,t)$ we have

$$\begin{aligned} a(t) := d(x(t),Q) \le \|x(t)-x\| = \left\| \int_0^t f(x(s),u(s))ds \right\| \le \int_0^t \big(\|f(x(s),u(s)) - f(\Pi_Q[x(s)],u(s))\| \\ + \|f(\Pi_Q[x(s)],u(s))\| \big) ds \le \left(\int_{t_1}^t a(s) L_\tau ds \right)_+ + F_Q(t-t_1)_+, \end{aligned}$$

where $\Pi_Q[y] \in \operatorname{argmin}_{x\in Q} \|y-x\|$ and $(a)_+ := \max\{0,a\}$. It follows then from Grönwall's inequality (c.f Lemma 2.1 in (Khalil, 2002)), with $\lambda = F_Q(t-t_1)_+$, $\mu = L_\tau$, $y(t) = a(t)$) that $\forall t \in [0, \min\{t_1+\tau, T\}]$,

$$a(t) = d(\xi(x,u,t),Q) \le F_Q(t-t_1)_+ e^{L_\tau (t-t_1)_+} \le F_Q \tau e^{L_\tau \tau}.$$

Finally, by repeating the same argument every additional time $\xi(x,u,t)$ leaves Q, the result follows. □

5. Relation between recurrence, τ-recurrence, and invariance entropy

In this section, we show different relations between recurrence and invariance entropy of system (1). In Theorem 1, we show that τ-recurrence entropy is both lower and upper bounded by invariance entropy with different initial and invariant sets. That results in a corollary showing that as τ approaches zero, τ-recurrence entropy approaches invariance entropy, which is in agreement with the intuition that τ-recurrence with $\tau = 0$ is invariance. In Theorem 2, we show that τ'-recurrence entropy is less than τ-recurrence entropy if $\tau' \ge \tau$. That is in agreement with the intuition that faster recurrence to Q requires more information about the state.

Theorem 1. *For any $Q \subseteq \mathbb{R}^n$ that is controlled invariant and $\tau \ge 0$, $h_{\text{inv}}(Q, B_{\delta_\tau}(Q)) \le h_{\text{rec}}(\tau,Q) \le h_{\text{inv}}(Q)$ and $h^*_{\text{inv}}(Q, B_{\delta_\tau}(Q)) \le h^*_{\text{rec}}(\tau,Q) \le h^*_{\text{inv}}(Q)$, where $\delta_\tau = F_Q \tau e^{L_\tau \tau}$, which is the right-hand-side of the containment lemma.*

Proof. The first inequality follows from the containment lemma that shows that any recurrence (T,ε,τ,Q)-spanning (resp. (T,τ,Q)-spanning) set is an invariance $(T,\varepsilon,Q,B_{\delta_\tau}(Q))$-spanning (resp. $(T,Q,\ B_{\delta_\tau}(Q))$-spanning) set. The second inequality follows from the observation that any invariance (T,ε,Q)-spanning (resp. (T,Q)-spanning) set is a recurrence (T,ε,τ,Q)-spanning (resp. (T,τ,Q)-spanning) set as well, for any $\tau \ge 0$ and $\varepsilon \ge 0$. □

Corollary 1. *As $\tau \to 0$, τ-recurrence entropy becomes equal to invariance entropy, i.e., $\lim_{\tau\searrow 0} h_{\text{rec}}(\tau,Q) = h_{\text{inv}}(Q)$.*

Theorem 2. *For any $Q \subseteq \mathbb{R}^n$ that is τ-recurrent for some $\tau > 0$, for any $\tau' \ge \tau$, $h_{\text{rec}}(\tau',Q) \le h_{\text{rec}}(\tau,Q)$ and $h^*_{\text{rec}}(\tau',Q) \le h^*_{\text{rec}}(\tau,Q)$.*

Proof. The result follows from the observation that any (T,ε,τ,Q)-spanning set (resp. (T,τ,Q)-spanning set) is a (T,ε,τ',Q)-spanning one (resp. (T,τ',Q)-spanning) as well. □

Though Theorem 1 only provides a non-strict statement, it is important to notice that it only requires Q to be controlled τ-recurrent. As a result, it is certainly possible to have scenarios wherein

$$h_{\text{rec}}(Q,\tau) < h_{\text{inv}}(Q) = \infty \tag{7}$$

which further emphasizes the fact that achieving τ-recurrence is less demanding than achieving invariance. We will show such an example in the next section.

6. Recurrence entropy bounds

In this section, we present an upper and a lower bound on τ-recurrence entropy. We show that when $\tau = 0$, we recover the upper bound on invariance entropy presented in (Colonius and Kawan, 2009).

Theorem 3 (Upper bound). *For any τ-recurrent set $Q \subseteq \mathbb{R}^n$ and any $\tau' \geq \tau$, $h_{\text{rec}}(\tau', Q) \leq L_\tau \dim_F(Q)/\ln 2 \leq L_\tau n/\ln 2$, where $\dim_F(Q) := \limsup_{\varepsilon\searrow 0} \frac{\ln b(\delta,Q)}{\ln(1/\delta)}$ is the upper box dimension of Q and $b(\delta, Q)$ is the minimal cardinality of a δ-cover of Q.*

Proof. The proof follows that of Theorem 4.2 in (Colonius and Kawan, 2009). Fix any T, ε, and $\tau' \geq \tau$. We define

$$L_{\tau,\varepsilon} = \max_{x_1,x_2\in B_\varepsilon(\text{cl}(R_r(Q,\tau))),u\in U} \frac{\|f(x_1,u) - f(x_2,u)\|}{\|x_1 - x_2\|} < \infty\,. \tag{8}$$

Let C be a minimal $\varepsilon e^{-L_{\tau,\varepsilon}T}$-cover of Q. Since Q is τ-recurrent, then there exists a set $S = \{u_i\}_{i\in[|C|]}$ such that $\xi(x_i, u_i, [0,T])$ is a (T,τ,Q)-recurrent trajectory, where x_i is the i^{th} center in the cover. Using the containment lemma (i.e, Lemma 1), we get that $\sup_{t\in\mathbb{R}^{\geq 0}} d(\xi(x_i,u_i,t),Q) \leq F_Q \tau e^{L_\tau \tau}$.

Using Grönwall's inequality, $\forall t \in [0,T]$ and $\forall x \in B_{\varepsilon e^{-L_{\tau,\varepsilon}T}}(x_i) \cap Q$, $\|\xi(x_i,u_i,t) - \xi(x,u_i,t)\| \leq e^{L_{\tau,\varepsilon}t}\|x_i - x\| \leq e^{L_{\tau,\varepsilon}t}(\varepsilon e^{-L_{\tau,\varepsilon}T}) \leq \varepsilon$. Consequently, $\xi(x,u_i,t)$ is a (T,ε,τ,Q)-recurrent trajectory and S is a recurrence (T,ε,τ',Q)-spanning set, for any $\tau' \geq \tau$. Thus, $r_{\text{rec}}(T,\varepsilon,\tau',Q) \leq b(\varepsilon e^{-L_{\tau,\varepsilon}T}, Q) = |S|$. Now that we have an upper bound on the minimal cardinality of a (T,ε,τ',Q)-spanning set, we can get the upper bound on recurrence entropy by substituting it in equation (3). Formally,

$$\begin{aligned}
h_{\text{rec}}(\tau',Q) &= \lim_{\varepsilon\searrow 0}\limsup_{T\to\infty}\frac{1}{T}\log r_{\text{rec}}(T,\varepsilon,\tau',Q)\\
&\leq \lim_{\varepsilon\searrow 0}\limsup_{T\to\infty}\frac{1}{T}\log r_{\text{rec}}(T,\varepsilon,\tau,Q)\\
&\leq \lim_{\varepsilon\searrow 0}\limsup_{T\to\infty}\frac{1}{T}\log b(\varepsilon e^{-L_{\tau,\varepsilon}T},Q)\\
&\leq \lim_{\varepsilon\searrow 0}\limsup_{T\to\infty}\frac{L_{\tau,\varepsilon}}{\ln(e^{L_{\tau,\varepsilon}T}/\varepsilon)+\ln\varepsilon}\log b(\varepsilon e^{-L_{\tau,\varepsilon}T},Q)\\
&= \lim_{\varepsilon\searrow 0} L_{\tau,\varepsilon}\limsup_{T\to\infty}\frac{1}{\ln(e^{L_{\tau,\varepsilon}T}/\varepsilon)}\log b(\varepsilon e^{-L_{\tau,\varepsilon}T},Q)\\
&= \lim_{\varepsilon\searrow 0} L_{\tau,\varepsilon}\limsup_{\delta\searrow 0}\frac{\ln b(\delta,Q)}{\ln 2\ln(1/\delta)}\\
&= L_\tau \dim_F(Q)/\ln 2.
\end{aligned} \tag{9}$$

The first inequality follows from the fact that any (T,ε,τ,Q)-spanning set is a (T,ε,τ',Q)-spanning one when $\tau \leq \tau'$. The second inequality follows from the set S we constructed earlier being a (T,ε,τ,Q)-spanning set with cardinality $b(\varepsilon e^{-L_{\tau,\varepsilon}T}, Q)$. The third inequality follows from multiplying the numerator and denominator with $L_{\tau,\varepsilon}$ and using the fact that $\ln(e^{L_{\tau,\varepsilon}T}/\varepsilon) + \ln\varepsilon = L_{\tau,\varepsilon}T$. The equality after that follows from the lim sup being unaffected by $\ln\varepsilon$ in the denominator and $L_{\tau,\varepsilon}$ being independent of T. The one before the last equality follows from replacing $\varepsilon e^{-L_{\tau,\varepsilon}T}$ with δ, which transform $\limsup_{T\to\infty}$ to $\limsup_{\delta\searrow 0}$ as well as the fact that $\log c = \ln c/\ln 2$. The last equality follows from substituting the definition of $\dim_F(Q)$ and $\lim_{\varepsilon\searrow 0} L_{\tau,\varepsilon}$ by its value L_τ. □

Remark 1. *Setting τ to zero makes $R_r(Q,\tau)$ as defined in (4) equal to Q and in the definition of L_τ in (5), the domain of the maximum would be $\mathrm{cl}(Q)$. Substituting L_τ in the bound in Theorem 3 results in the same upper-bound as that on invariance entropy in Theorem 4.2 in (Colonius and Kawan, 2009).*

Remark 2. *Theorem 3 shows that if the system is capable of achieving faster recurrence to Q than required, i.e., achieving τ-recurrence while the requirement is τ'-recurrence for some $\tau' > \tau$, then we can obtain a tighter upper bound on recurrence entropy since $L_\tau \le L_{\tau'}$.*

Example 1 (Illustrative Example)**.** *Consider the case following two-dimensional linear system*

$$\begin{bmatrix}\dot{x}_1\\ \dot{x}_2\end{bmatrix} = \begin{bmatrix}0 & 1\\ 0 & 0\end{bmatrix}\begin{bmatrix}x_1\\ x_2\end{bmatrix} + \begin{bmatrix}0\\ 1\end{bmatrix}u \tag{10}$$

We assume $u \in U = [-1,1]$, and consider the set $Q = [-1,1]^2$.

Observe that with simple integration, we can get the closed form solution as follows:

$$\xi(x,u,t) = \begin{bmatrix}x_1(t)\\ x_2(t)\end{bmatrix} = \begin{bmatrix}\int_{s_2=0}^{t}\int_{s_1=0}^{t} u(s_1)ds_1 ds_2 + x_2(0)t + x_1(0)\\ \int_{s=0}^{t} u(s)ds + x_2(0)\end{bmatrix}.$$

Consider the case when $x = [1,1]$. Then, for the trajectory starting at x to not leave Q, the control signal should be chosen so that neither of the two coordinates increase. Both coordinates are monotonically increasing in u. If we choose the control signal to have the minimum value -1 for some interval $[0,T]$ in the effort of preventing the state coordinates from increasing and escaping Q, then $\xi(x,u,t') = [1 + t - \frac{1}{2}t^2, -t+1]$. Thus, for all $t \le 2$, $\xi(x,u,t') \notin Q$. Therefore, there is no piece-wise continuous control signal that can make the trajectory starting from x invariant to Q or even τ-recurrent with $\tau < 2$, and Q is not controlled invariant or τ-recurrent with any $\tau < 2$. Thus, the invariance entropy $h_{\mathrm{inv}}(Q)$ and τ-recurrence entropy $h_{\mathrm{rec}}(Q,\tau)$ of system (10) are infinite for $\tau < 2$.

In contrast, observe that with constant control signals with values in U, any trajectory with an initial state in Q can be driven back to Q within 2 time units. Thus, Q is controlled 2-recurrent and we can use Theorem 3 to compute a finite upper bound on $h_{\mathrm{rec}}(Q,\tau)$. Observe that $L_\tau \le \|\frac{\partial f}{\partial x}\| = \|A\| = 1$, where $\|\cdot\|$ is the induced supremum matrix norm. It therefore follows that

$$h_{\mathrm{inv}}(Q) = +\infty \quad \text{and} \quad h_{\mathrm{rec}}(Q,\tau) = \begin{cases} +\infty & \tau < 2\\ \le 2/\ln 2 & \tau \ge 2\end{cases}.$$

In the following theorem, we present a lower bound on τ-recurrence entropy of system (1). To that end, we will require a stronger version of Assumption 1.

Assumption 2 ((ε,τ)-completeness)**.** *There is some $\varepsilon > 0$, such that, for any $x \in B_\varepsilon(Q)$ and $u \in \mathcal{U}$, the trajectory $\xi(x,u,\cdot)$ is defined for all $t \in [0,\tau]$ and is continuous in its first argument.*

Theorem 4 (Lower bound)**.** *Let Assumption 2 hold. Then, for any τ-recurrent set $Q \subseteq \mathbb{R}^n$,*

$$h_{\mathrm{rec}}(\tau,Q) \ge \frac{1}{\ln 2}\max\left\{0, \min_{(x,u)\in \mathrm{cl}(B_{\delta_\tau}(Q))\times U} \mathrm{div}_x f(x,u)\right\},$$

where $\mathrm{div}_x f(x,u) = \sum_{i=1}^{n} \frac{\partial f_i}{\partial x_i}(x,u) = \mathrm{tr}\frac{\partial f}{\partial x}(x,u)$.

Proof. A small modification of the proof of Theorem 4.1 in (Colonius and Kawan, 2009) would result in the theorem. The modified proof is as follows: first, fix $T \ge 0$, ε that satisfies Assumption 2, and let $S = \{u_j\}_{j\in[M]}$ be a minimal recurrence (T,ε,τ,Q)-spanning set. Let us define the following sets: for any $j \in [M]$,

$$Q_j = \{x \in Q \mid \xi(x,u_j,[0,T]) \text{ is } (T,\varepsilon,\tau,Q)\text{-recurrent}\}. \tag{11}$$

For simplicity of notation, we define $\xi(Q_j,u_j,T) := \cup_{x\in Q_j}\xi(x,u_j,T)$. Note further that the statement $\xi(x,u_j,[0,T])$ is (T,ε,τ,Q)-recurrent is equivalent to the statement that $\xi(x,u_j,[0,T])$ is $(T,\tau,B_\varepsilon(Q))$-recurrent. Then, by Lemma 1 (applied to $B_\varepsilon(Q)$), we have that $\lambda(\xi(Q_j, u_j, T)) \leq \lambda(B_{\delta_{\tau,\varepsilon}}(Q))$, where $\delta_{\tau,\varepsilon} = F_{Q,\varepsilon}\tau e^{L'_{\tau,\varepsilon}\tau}$, which is the right-hand-side of Lemma 1 when Q is replaced by $B_\varepsilon(Q)$ as we define $F_{Q,\varepsilon} = F_{B_\varepsilon(Q)}$ and $L'_{\tau,\varepsilon}$ is defined as follows:

$$L'_{\tau,\varepsilon} = \max_{x_1,x_2\in \mathrm{cl}(R_r(B_\varepsilon(Q),\tau)),u\in U} \frac{\|f(x_1,u) - f(x_2,u)\|}{\|x_1 - x_2\|} < \infty. \tag{12}$$

Note that in the case of invariance (as when $\tau = 0$), we instead have $\lambda(\xi(Q_j, u_j, T)) \leq \lambda(B_\varepsilon(Q))$, as shown in (Colonius and Kawan, 2009).

Now, we can use the transformation theorem and Liouville's trace formula to get:

$$\lambda(\xi(Q_j,u_j,T)) = \int_{Q_j} |\det\frac{\partial \xi}{\partial x}(x,u_j,T)|dx \;\;\geq\;\; \lambda(Q_j) \inf_{\substack{(x,u)\in Q\times\mathcal{U}_r,\\ \xi(x,u,[0,T])\subseteq B_{\delta_{\tau,\varepsilon}}(Q)}} |\det\frac{\partial \xi}{\partial x}(x,u,T)|$$

$$= \lambda(Q_j) \inf_{\substack{(x,u)\in Q\times\mathcal{U}_r,\\ \xi(x,u,[0,T])\subseteq B_{\delta_{\tau,\varepsilon}}(Q)}} \exp\Big(\int_0^T \mathrm{div}_x f(\xi(x,u,s),u(s))ds\Big) \;\;\geq\;\; \lambda(Q_j) \min_{(x,u)\in \mathrm{cl}(B_{\delta_{\tau,\varepsilon}}(Q))\times U} \exp\big(T\mathrm{div}_x f(x,u)\big).$$

Now since $\lambda(Q) \leq M \max_{j\in[M]} \lambda(Q_j)$,

$$\begin{aligned}\lambda(Q) &\leq M\frac{\max_{j\in[M]}\lambda(\xi(Q_j,u_j,T))}{\min_{(x,u)\in \mathrm{cl}(B_{\delta_{\tau,\varepsilon}}(Q))\times U}\exp\big(T\mathrm{div}_x f(x,u)\big)}\\ &\leq M\frac{\lambda(B_{\delta_{\tau,\varepsilon}}(Q))}{\min_{(x,u)\in \mathrm{cl}(B_{\delta_{\tau,\varepsilon}}(Q))\times U}\exp\big(T\mathrm{div}_x f(x,u)\big)}.\end{aligned}$$

Consequently,

$$M \geq \frac{\lambda(Q)}{\lambda(B_{\delta_{\tau,\varepsilon}}(Q))} \min_{(x,u)\in \mathrm{cl}(B_{\delta_{\tau,\varepsilon}}(Q))\times U} \exp\big(T\mathrm{div}_x f(x,u)\big).$$

Recall that M here is equal to $r_{\mathrm{rec}}(T,\varepsilon,\tau,Q)$. Thus, since τ is finite, δ_τ is finite and $h_{\mathrm{rec}}(\tau,Q)$

$$\begin{aligned}&\geq \lim_{\varepsilon\searrow 0}\limsup_{T\to\infty}\frac{1}{T}\min_{(x,u)\in \mathrm{cl}(B_{\delta_{\tau,\varepsilon}}(Q))\times U}\log\exp\big(T\mathrm{div}_x f(x,u)\big)\\ &= \lim_{\varepsilon\searrow 0}\frac{1}{\ln 2}\min_{(x,u)\in \mathrm{cl}(B_{\delta_{\tau,\varepsilon}}(Q))\times U}\mathrm{div}_x f(x,u)\\ &= \frac{1}{\ln 2}\min_{(x,u)\in \mathrm{cl}(B_{\delta_{\tau}}(Q))\times U}\mathrm{div}_x f(x,u).\end{aligned}$$

Note that δ_τ strictly increases with τ. Thus, with a larger τ, the domain over which the minimum is taken in the lower bound becomes larger, and the minimum itself becomes smaller. This is expected since $h(\tau', Q) \leq h(\tau, Q)$ if $\tau' \geq \tau$, according to Theorem 2. Also, as $\tau \to 0$, we get the same lower bound as invariance entropy presented in (Colonius and Kawan, 2009). □

Remark 3. *Theorem 4 does not follow directly from Theorem 4.1 in (Colonius and Kawan, 2009), i.e., from the result that* $h_{\mathrm{inv}}(Q)$ *is greater than or equal to* $\max\{0, \min_{(x,u)\in Q\times U} \mathrm{div}_x f(x,u)\}$, *since* $h_{\mathrm{rec}}(\tau,Q) \leq h_{\mathrm{inv}}(Q)$, *for any* $\tau \geq 0$.

7. Entropy and τ-Recurrence data rates

We assume the setup where there is a sensor that can measure the state of system (1) at any time instant with an arbitrary accuracy. It also has computation capabilities that allows it to simulate the system starting from any initial state and following any control, as long as that trajectory exists. The sensor is connected to a controller over a limited-bandwidth channel. The controller does not have information about the state of the system besides what it receives from the sensor. It does however know the τ-recurrent set $Q \subseteq \mathbb{R}^n$ and the control signals that drive the system when starting from any state in Q to have an (ε, τ, Q)-recurrent trajectory.

An (ε, τ, Q)-*recurrence enforcing* algorithm is a pair of procedures, one for the sensor and the other for the controller. The sensor's procedure determines the bits it sends over the channel to the controller. Based on these bits, the controller's procedure determines how to map these bits to a control signal to drive the system to have an (ε, τ, Q)-recurrent trajectory. The average bit rate of an (ε, τ, Q)-recurrence enforcing algorithm is defined as follows: $\lim_{T\to\infty} \frac{\#\text{bits}(T)}{T}$, where $\#\text{bits}(T)$ is the total number of bits sent by the sensor until time T.

Theorem 5. *For any controlled τ-recurrent set $Q \subseteq \mathbb{R}^n$ and $\varepsilon \geq 0$, there exists no (ε, τ, Q)-recurrence enforcing algorithm with an average bit rate smaller than $h_{\text{rec}}(\tau, Q)$.*

Proof. The proof is by contradiction. If there is such an algorithm with an average data rate smaller than entropy, then there exists a $T > 0$ such that

$$\frac{\#\text{bits}(T)}{T} < \frac{1}{T} \log r_{\text{rec}}(T, \varepsilon, \tau, Q). \tag{13}$$

That implies that $2^{\#\text{bits}(T)} < r_{\text{rec}}(T, \varepsilon, \tau, Q)$. Observe that $2^{\#\text{bits}(T)}$ is the number of control signals that the controller can possibly generate over the interval $[0, T]$. By the assumption that the controller enforces the system to have an $(T, \varepsilon, \tau, Q)$-recurrent trajectory, then for every $x \in Q$, it can generate a control signal that results in a $(T, \varepsilon, \tau, Q)$-recurrent trajectory. Therefore, the set of control signals that the controller can generate is a $(T, \varepsilon, \tau, Q)$-spanning one that has a smaller cardinality than $r_{\text{rec}}(T, \varepsilon, \tau, Q)$, which contradicts the latter's definition being the minimal cardinality of a $(T, \varepsilon, \tau, Q)$-spanning set. □

8. Algorithm for enforcing τ-recurrence over limited-bandwidth channels

In this section, we present Algorithm 1, which when the sensor follows and the controller follows a corresponding symmetric procedure, the resulting control signal drives system (1) to have an exponentially converging trajectory to a (τ, Q)-recurrent one at a user-specified rate $\alpha \geq 0$. When $\alpha = 0$, the trajectory would be an (ε, τ, Q)-recurrent trajectory with a user-specified ε. We define this more formally in Theorem 6 and Corollary 2. After that, we show that the bit rate at which the sensor sends information to the controller when following Algorithm 1 is equal to the upper bound on τ-recurrence entropy presented in Theorem 3 when $\alpha = 0$, and grows linearly with α, otherwise.

In our algorithm, we assume that starting from any state in $B_{\delta_\tau+\varepsilon}(Q)$, for some $\varepsilon > 0$, there exists a control signal that drives system (1) to Q within τ time units. Moreover, we assume that the function that maps the initial states $B_{\delta_\tau+\varepsilon}(Q)$ to the shortest time such a control signal takes to drive system (1) to Q to be Lipschitz continuous. This is formulated as follows.

Assumption 3. $\exists \varepsilon^* > 0$ *such that there exists a control function $h : B_{\delta_\tau+\varepsilon^*}(Q) \times \mathbb{R}^{\geq 0} \to U$ and a corresponding function $t_{\text{rec},Q} : B_{\delta_\tau+\varepsilon^*}(Q) \to [0, \tau]$, such that $\forall x \in B_{\delta_\tau+\varepsilon^*}(Q)$, $\xi(x, h(x, \cdot), t_{\text{rec},Q}(x)) \in Q$ and $\forall t \in [0, t_{\text{rec},Q}(x)]$, $d(\xi(x, h(x, \cdot), t), Q) \leq d(x, Q)$. Moreover, there exists some constant $c^* \geq 0$ such that for any $x_1, x_2 \in B_{\delta_\tau+\varepsilon^*}(Q)$, $|t_{\text{rec},Q}(x_1) - t_{\text{rec},Q}(x_2)| \leq c^* \|x_1 - x_2\|$.*

Next, for any $\tau > 0$ and $\varepsilon \in (0, \varepsilon^*]$, we define a new control function $g : B_{\delta_\tau+\varepsilon}(Q) \times \mathbb{R}^{\geq 0} \to U$ to be used in the algorithm. If the initial state x is in Q, $g(x, \cdot)$ is equal to a control signal that ensures (τ, Q)-recurrence, which exists by the assumption that Q is τ-recurrent. Otherwise, it is equal to the control function h defined in Assumption 3 up until reaching Q, i.e., until $t_{\text{rec},Q}(x)$. After that, it is equal to the control function that ensures the trajectory is τ-recurrent starting from the new initial state in Q.

Formally, let $g' : Q \times \mathbb{R}^{\geq 0} \to U$ be such that for any $x \in Q$, the trajectory $\xi(x, g'(x, \cdot), \cdot)$ is a (τ, Q)-recurrent one. Such a function exists because of the assumption that Q is a τ-recurrent set. We define g as follows: $\forall t \geq 0, g(x, t) := g'(x, t)$ if $x \in Q$. If $x \in B_{\delta_\tau+\varepsilon^*}(Q) \backslash Q$, $\forall t \leq t_{\text{rec},Q}(x), g(x, t) := h(x, t)$ and $\forall t > t_{\text{rec},Q}(x), g(x, t) := g'(\xi(x, h(x, \cdot), t_{\text{rec},Q}(x)), t - t_{\text{rec},Q}(x))$.

8.1. Algorithm description

Algorithm 1 takes as input a τ-recurrent set Q for some $\tau > 0$, an $\varepsilon \in (0, \varepsilon^*]$ (where ε^* is as defined in Assumption 3), and the control function g defined earlier. It also assumes to be given several functions: $sense$, $quantize$, $encode$, $send$, *simulate*, and *sleep*. The function *sense* returns the current state of the system. The function *quantize* returns the closest point in the set given in its second argument to the point given in its first argument. The function *encode* maps the first argument to a bit vector that uniquely identifies it out of the set of states given in the second argument. The function *send* sends the given bit vector over the limited-bandwidth channel to the controller. The function *simulate* simulates the system starting from the state in its first argument following the control signal in its second argument until the time bound specified in its third argument. It returns the last state in the simulated trajectory. If the third argument is an interval, it returns the trajectory segment within that interval. Finally, the function *sleep* makes the sensor wait for the amount of real time passed as argument before continuing the execution of the algorithm. The time of the algorithm execution is assumed to be negligible with respect to τ.

The algorithm starts by initializing S_0 to Q and constructing an $\varepsilon e^{-(L_\tau+\alpha)\tau}$-grid for it, which we denote by C_0. The algorithm then proceeds with an infinite loop. In each iteration, it sends a bit vector that encodes a state estimate to the controller, according to which it can identify the control function u_i the system should follow in the time interval $[i\tau, (i+1)\tau)$. To produce the bit vector, the sensor measures the current state of the system x_i, i.e., $\xi(x_0, \hat{u}, i\tau)$, where $\hat{u}$ is the control signal have been followed so far. Then, it quantizes x_i to one of the centers q_i in the grid C_i. The encoding of q_i with respect of C_i is the bit vector that the sensor sends. The controller, which is running a similar algorithm to Algorithm 1, but without the sensing, can recover q_i as it knows C_i. Using q_i, it can choose the same control function u_i that the sensor intends to use to construct C_{i+1}.

After that, Algorithm 1 computes u_i to be equal to $g(q_i, [0, \tau))$. Then, it simulates the system for τ time units starting from q_i and following u_i. It uses the last state in the simulated trajectory as the center of the ball S_{i+1} which bounds the region where the next sensed state x_{i+1} might be. The radius r_{i+1} of S_{i+1} is an $e^{\alpha\tau}$ factor smaller than that of S_i. After that, it constructs the grid C_{i+1} to be the $r_{i+1}e^{-(L_\tau+\alpha)\tau}$-grid over S_{i+1}, according to which the next state, x_{i+1}, would be quantized. Finally, the sensor waits for the system to evolve for τ time units before sensing it again in the next iteration.

Algorithm 1 Sensor algorithm for achieving recurrence

input: Q, $\varepsilon \in (0, \varepsilon^*]$, $\tau > 0$, $g : B_{\delta_\tau+\varepsilon}(Q) \times \mathbb{R}^{\geq 0} \to U$
$S_0 \leftarrow Q$
$r_0 \leftarrow \varepsilon$
$C_0 \leftarrow grid(S_0, r_0 e^{-(L_\tau+\alpha)\tau})$
$i = 0$
while true **do**
    $x_i \leftarrow sense()$
    $q_i \leftarrow quantize(x_i, C_i)$
    $send(encode(q_i, C_i))$
    $u_i \leftarrow g(q_i, [0, \tau))$
    $r_{i+1} \leftarrow r_i e^{-\alpha\tau}$
    $S_{i+1} \leftarrow B_{r_{i+1}}(simulate(q_i, u_i, \tau))$
    $C_{i+1} \leftarrow grid(S_{i+1}, r_{i+1}e^{-(L_\tau+\alpha)\tau})$
    $i \leftarrow i + 1$
    $sleep(\tau)$

8.2. Algorithm guarantees

Fix the inputs to Algorithm 1, i.e., a controlled τ-recurrent set Q and a corresponding τ-recurrence achieving controller g. Moreover, fix any initial state $x_0 \in Q$. Let $\hat{u} : \mathbb{R}^{\geq 0} \to \mathbb{R}^m$ be the concatenation of the u_is produced by Algorithm 1, i.e., for any $t \geq 0$, $\hat{u}(t) = u_i(t - i\tau)$, where $i = \lfloor t/\tau \rfloor$. Also, let $\hat{\xi} : \mathbb{R}^{\geq 0} \to \mathbb{R}^m$ be the concatenation of the τ-sized fragments of trajectories $\xi(q_i, u_i, [0, \tau))$ produced by the algorithm, i.e., $\forall t \geq 0$ and $t \neq i\tau$ for some $i \in \mathbb{N}$, $\hat{\xi}(t) = simulate(q_i, u_i, t - i\tau)$ and $\hat{\xi}(i\tau) = q_i$, where $i = \lfloor t/\tau \rfloor$. Thus, $\hat{\xi}$ would be right-piece-wise-continuous. Finally, the trajectory that the system would have starting from x_0 following $\hat{u}$ is denoted as usual by $\xi(x_0, \hat{u}, \cdot)$.

Theorem 6. *Algorithm 1 ensures that:*

1. $\forall i \geq 0$, $x_i \in S_i$, *and* $\forall t \geq 0$, $\|\hat{\xi}(t) - \xi(x_0, \hat{u}, t)\| \leq \varepsilon e^{-\alpha t}$,
2. $\forall i \in \mathbb{N}$, $\hat{\xi}[i\tau, \infty)$ *is an* $(\varepsilon e^{-i\alpha\tau}, \tau + c^* \varepsilon e^{-(i\alpha + L_\tau)\tau}, Q)$*-recurrent function, and*

3. $\forall i \in \mathbb{N}$, *$\xi(x_0, \hat{u}, [i\tau, \infty))$ is a $(2\varepsilon e^{-i\alpha\tau}, \tau + c^*\varepsilon e^{-(i\alpha+L_\tau)\tau}, Q)$-recurrent trajectory.*

Proof. First, we will prove part 1) by induction. For the base case: $x_0 \in S_0$ and $\|\hat{\xi}(0) - x_0\| \leq \varepsilon e^{-(L_\tau+\alpha)\tau} \leq \varepsilon$, which hold by the fact that C_0 is a grid over Q with cells of radii $r_0 = \varepsilon e^{-(L_\tau+\alpha)\tau}$ and $\hat{\xi}(0) = q_0$.

Inductive case: fix an $i \in \mathbb{N}$ and assume that $x_i \in S_i$ and $\forall t \in [0, i\tau]$, $\|\hat{\xi}(t) - \xi(x_0, \hat{u}, t)\| \leq \varepsilon e^{-\alpha t}$. By Grönwall's inequality, $\forall t \in [i\tau, (i+1)\tau]$, $\|\xi(x_i, u_i, t) - \xi(q_i, u_i, t)\| \leq e^{L_\tau(t-i\tau)}\|x_i - q_i\| \leq \varepsilon e^{-((i+1)\alpha+L_\tau)\tau}e^{L_\tau(t-i\tau)} \leq \varepsilon e^{-(i+1)\alpha\tau} \leq \varepsilon e^{-\alpha t}$. Recall that $x_{i+1} = \xi(x_i, u_i, \tau)$ and $r_{i+1} = \varepsilon e^{-(i+1)\alpha\tau}$. Thus, $x_{i+1} \in B_{r_{i+1}}(\xi(q_i, u_i, \tau))$, and the latter is S_{i+1}. Then, since C_{i+1} is a grid over S_{i+1} with granularity $r_{i+1}e^{-(L_\tau+\alpha)\tau}$ and $\hat{\xi}((i+1)\tau) = q_{i+1}$, $\|x_{i+1} - \hat{\xi}((i+1)\tau)\| = \|x_{i+1} - q_{i+1}\| \leq r_{i+1}e^{-(L_\tau+\alpha)\tau}$, and thus $\|\xi(x_0, \hat{u}, (i+1)\tau) - \hat{\xi}((i+1)\tau)\| \leq \varepsilon e^{-((i+2)\alpha+L_\tau)\tau} \leq \varepsilon e^{-(i+1)\alpha\tau}$. That proves the inductive argument for part 1).

We prove part 2) also by induction. We will prove the stronger claim that for any $i \in \mathbb{N}$, either $\hat{\xi}(t) \in Q$ for some $t \in [i\tau, (i+1)\tau)$ or $\lim_{t\to((i+1)\tau)^-} \hat{\xi}(t) \in Q$, and $|t_{i+1} - t_i| \leq \tau + c^*\varepsilon e^{-i\alpha\tau}$, where t_i is the last time instant in $[i\tau, (i+1)\tau)$ such that $\hat{\xi}(t_i) \in Q$ or equal to $(i+1)\tau$, otherwise, and t_{i+1} is the first time instant in $[(i+1)\tau, (i+2)\tau)$ such that $\hat{\xi}(t_{i+1}) \in Q$, or $t^*_{i+1} = (i+2)\tau$, otherwise. When it is the case that $\lim_{t\to((i+1)\tau)^-} \hat{\xi}(t) \in Q$, we know from part 1) that $q_{i+1} \in S_{i+1}$, which is centered at the value of that limit and has a radius of r_{i+1}. Thus, $\hat{\xi}((i+1)\tau)$, which is equal to q_{i+1}, would be at most r_{i+1} (i.e., $\varepsilon e^{-(i+1)\alpha\tau}$) from Q.

Base case: by assumption, $x_0 \in Q$ and $\hat{\xi}(0) = q_0$. By part 1), $\|\hat{\xi}(0) - x_0\| \leq \varepsilon e^{-(L_\tau+\alpha)\tau}$ and thus $\hat{\xi}(0) \in B_{\varepsilon e^{-(L_\tau+\alpha)\tau}}(Q) \subseteq B_{\delta_\tau+\varepsilon e^{-(L_\tau+\alpha)\tau}}(Q)$. If $q_0 \in Q$, then u_0 is equal to $g(q_0, [0, \tau))$. That would result in $\hat{\xi}([0, \tau))$ being a prefix of a (τ, Q)-recurrent trajectory starting from q_0, by the definition of g. Thus, either $\hat{\xi}(t) \in Q$ for some $t \in (0, \tau)$ or $\lim_{t\to\tau^-} \hat{\xi}(t) \in Q$. Moreover, by the containment lemma (Lemma 1), $\forall t \in [0, \tau)$, $\hat{\xi}(t) \in B_{\delta_\tau}(Q)$. Thus, $\lim_{t\to\tau^-} \hat{\xi}(t)$, which is the center of S_1, would be in $B_{\delta_\tau}(Q)$. Thus, $\hat{\xi}(\tau)$, which is equal to q_1 and belongs to S_1, would be in $B_{\delta_\tau+\varepsilon e^{-\alpha\tau}}(Q)$.

If $\hat{\xi}(0) \in B_{\delta_\tau+\varepsilon}(Q)\backslash Q$ instead, then, by Assumption 3, applying the control $g(q_0, [0, \tau])$ will result in $\hat{\xi}(t) \in B_{\delta_\tau+\varepsilon}(Q)$, for all $t \in [0, t_{\text{rec},Q}(x)]$. If $t_{\text{rec},Q}(x) = \tau$, then as in the first case, $\lim_{t\to\tau^-} \hat{\xi}(t) \in Q$. Otherwise, if $t_{\text{rec},Q}(x) < \tau$, then $\hat{\xi}(t_{\text{rec},Q}(x)) \in Q$. Moreover, in the interval $(t_{\text{rec},Q}(x), \tau)$, $\hat{\xi}$ would be equal to the trajectory $\xi(\xi(q_0, u_0, t_{\text{rec},Q}(x)), u_0((t_{\text{rec},Q}(x), \tau)), \cdot)$, which is a (τ, Q)-recurrent trajectory by the definition of $u_0((t_{\text{rec},Q}(x), \tau])$ being equal to $g(\xi(q_0, u_0, t_{\text{rec},Q}(x)), [t - t_{\text{rec},Q}(x), \tau - t_{\text{rec},Q}(x)))$. Thus, by the containment lemma (Lemma 1), it is contained in $B_{\delta_\tau}(Q)$. Hence, $\lim_{t\to\tau^-} \hat{\xi}(t)$, the center of S_1, is in $B_{\delta_\tau}(Q)$. By part 1), $x_1 \in S_1$. Also, $q_1 \in S_1$ and $\hat{\xi}(\tau) = q_1$. Thus, $\hat{\xi}(\tau) \in B_{\delta_\tau+r_1}(Q) = B_{\delta_\tau+\varepsilon e^{-\alpha\tau}}(Q)$.

Inductive case: fix an $i \geq 1$ and assume that that part 2) is true until time $i\tau$. Thus, there exists a time instant $t \in [0, \tau)$ such that $\hat{\xi}((i-1)\tau + t) \in Q$ or $\lim_{t\to(i\tau)^-} \hat{\xi}(t) \in Q$. Let t_{i-1} be the largest such instant. If $t_{i-1} < \tau$ and we simulate system (1) following $g(\xi(q_{i-1}, u_{i-1}, i\tau + t_{i-1}), \cdot)$ starting from $\xi(q_{i-1}, u_{i-1}, i\tau + t_{i-1})$, the resulting trajectory will be (τ, Q)-recurrent. Thus, there exists $t' \in (0, \tau]$ such that that trajectory belongs to Q at time $i\tau + t'$. However, $\hat{\xi}$ is equal to that trajectory only in the interval $[(i-1)\tau + t_{i-1}, i\tau)$. In the interval $[i\tau, (i+1)\tau)$, $\hat{\xi}$ will be equal to the trajectory that starts from q_i and follows $g(q_i, \cdot)$. If $q_i \in Q$, then $\hat{\xi}$ would have visited Q within $\tau - t^*_{i-1}$ time units, which is less than τ. If $q_i \notin Q$, then from part 1), we know that $q_i \in S_i$ and thus $\|q_i - \lim_{t\to i\tau^-} \hat{\xi}(t)\| \leq r_i = \varepsilon e^{-i\alpha\tau}$. Then, by Assumption 3, we know that $\hat{\xi}$ would reach Q at or before $\min\{(i+1)\tau, t_{i-1} + \tau + c^*\|q_i - \lim_{t\to i\tau^-} \hat{\xi}(t)\|\}$, which is upper bounded by $\min\{(i+1)\tau, t_{i-1} + \tau + c^*\varepsilon e^{-i\alpha\tau}\}$. We can conclude that the time between two time instants at which $\hat{\xi}$ belongs to Q in the intervals $[(i-1)\tau, i\tau)$ and $[i\tau, (i+1)\tau)$ is less than or equal to $\tau + c^*\varepsilon e^{-i\alpha\tau}$.

Finally, part 3) follows from combining parts 1) and part 2) and using the triangular inequality $d(\xi(x_0, \hat{u}, t), Q) \leq \|\xi(x_0, \hat{u}, t) - \hat{\xi}(t)\| + d(\hat{\xi}(t), Q)$ at the time instants where $\hat{\xi}(t)$ is visiting $B_{\delta_\tau+\varepsilon e^{-i\alpha\tau}}(Q)$. We obtain that $\xi(x_0, \hat{u}, \cdot)$ visits $B_{\delta_\tau+2\varepsilon e^{-i\alpha\tau}}(Q)$ in the $[i\tau, (i+1)\tau)$ interval. Thus, $\xi(x_0, \hat{u}, [i\tau, \infty))$ is a $(2\varepsilon e^{-i\alpha\tau}, \tau+c^*\varepsilon e^{-i\alpha\tau}, Q)$-recurrent trajectory. □

It follows that the trajectory of system (1) when following the controller $\hat{u}$ produced by Algorithm 1 asymptotically approaches a (τ, Q)-recurrent trajectory.

Corollary 2. *As $t \to \infty$, $\xi(x_0, \hat{u}, [t, \infty))$ is a (τ, Q)-recurrent trajectory.*

In the following theorem, we show that the bit rate at which a sensor running Algorithm 1 sends information t the controller matches the upper bound on τ-recurrence entropy in Section 6.

Theorem 7. *The average bit rate at which a sensor running Algorithm 1 will send to the controller is equal to $n(L_\tau + \alpha)/\ln 2$.*

Proof. Fix any $i \in \mathbb{N}$. The number of bits that the sensor running Algorithm 1 sends at the time instant $t = i\tau$ is $\log |C_i|$. Given any time bound $T \geq 0$, the total number of bits sent by the sensor over $[0, T]$ is equal to $\sum_{i=0}^{\lfloor T/\tau \rfloor} \log |C_i|$. Thus, the average bit rate is $\lim_{T\to\infty} \sum_{i=0}^{\lfloor T/\tau \rfloor} \frac{\log |C_i|}{T}$.

We can observe that $\forall i \geq 0$, $C_i = \lceil \frac{\mathrm{diam}(S_i)}{2r_i e^{-(L_\tau+\alpha)\tau}} \rceil^n$. Then, $C_0 = \lceil \frac{\mathrm{diam}(Q)}{2\varepsilon e^{-(L_\tau+\alpha)\tau}} \rceil^n$ and for any $i \geq 1$, $C_i = \lceil \frac{2r_i}{2r_i e^{-(L_\tau+\alpha)\tau}} \rceil^n = \lceil e^{(L_\tau+\alpha)\tau} \rceil^n$.

Thus, the average bit rate is equal to:

$$\begin{aligned} &\lim_{T\to\infty} \frac{1}{T} \sum_{i=0}^{\lfloor T/\tau \rfloor} \frac{\log |C_i|}{T} \\ &= \lim_{T\to\infty} \frac{1}{T} \big(\lceil \frac{\mathrm{diam}(Q)}{2\varepsilon e^{-(L_\tau+\alpha)\tau}} \rceil^n + \sum_{i=1}^{\lfloor T/\tau \rfloor} \log \lceil e^{(L_\tau+\alpha)\tau} \rceil^n \big) \\ &= n(L_\tau + \alpha)/\ln 2. \end{aligned} \quad (14)$$

□

9. Finite alphabet controllers and memory bounds

Having established how recurrence, as a control task, is less complex (entropy) and requires less information (bit rates) than the invariance counterpart, we now center our attention on a striking phenomenon associated with recurrence that is generally unattainable by invariance tasks. That is, though the number of control signals required to enforce recurrence grows exponentially with time, such controls can be generated by a finite alphabet of control signals that can be strategically combined to achieve the desire goal. This allows control signals to be precomputed and stored, thus circumventing one of the key shortcomings of the entropy-based algorithms, which assumes the ability to efficiently compute controls online. We start by introducing the building blocks of our controllers, i.e., the control alphabet.

Definition 6 (Control alphabet). *Fix a $\tau \geq 0$. We call a set $S_\tau := \{v_1, \dots, v_n\}$ consisting of piecewise continuous control signals of the form $v_i : [0, t_i) \to U$, for some $t_i \in (0, \tau]$, a* control alphabet.

Sequentially concatenating signals from the alphabet generate signals over longer time horizons. For example, the signal $u : [0, t_i + t_j) \to U$ such that $u(t) = v_i(t)$ for all $t \in [0, t_i)$ and $u(t_i + t) = v_j(t)$ for all $t \in [0, t_j)$, results from concatenating v_i and v_j, and is denoted by $u = v_i v_j$. This construction allows the generation of control signals of arbitrary duration, by subsequent concatenation of alphabet elements, e.g., concatenating $v_i v_j$ with v_k leads to $(v_i v_j)v_k = v_i v_j v_k$ which has a total duration $t_i + t_j + t_k$. Our goal is then to understand under what condition such alphabets can induce trajectories with the required properties, as in Definition 3 and Definition 5.

More precisely, we call a control alphabet S_τ a $(T, \varepsilon, \tau, Q)$-*invariance alphabet* for system (1) if its members can be sequentially concatenated to generate signals over the interval $[0, T]$ that form an invariance (T, ε, Q)-spanning set. Similarly, we call S_τ a $(T, \varepsilon, \tau, Q)$-*recurrence alphabet* for system (1) if its members can be sequentially concatenated to generate signals that form a recurrence $(T, \varepsilon, \tau, Q)$-spanning set. As before, we drop T from the arguments if equals to infinity and drop ε if equals to zero.

Definition 7 (Nonparametric control). *Given a control alphabet S_τ, a non-parametric controller is defined by a map $\pi : \mathbb{R}^n \to S_\tau$. It* induces *a controller $g : \mathbb{R}^n \times \mathbb{R}^{\geq 0} \to U$ for system (1). The controller g samples the state x_0 of system (1) at time zero, outputs $\pi(x_0)$ for the latter's duration, i.e., if $\pi(x_0) = v_i$, $g(x_0, t) = v_i(t)$ for all $t \in [0, t_i)$, before sampling the state again and repeating the process.*

Our goal is to use non-parametric controllers of the form of Definition 7 as a substitute for having to compute the control signals online via MPC or a similar method. To that end, we define entropy-like metrics that quantify the minimal number of bits necessary to differentiate across elements of an alphabet that achieves invariance and recurrence. We call such metrics *invariance memory* and *recurrence memory*, respectively, as follows:

$$m_{\mathrm{inv}}(\tau, Q) := \lim_{\varepsilon \searrow 0} \limsup_{T\to\infty} \log r_{\mathrm{inv,mem}}(T, \varepsilon, \tau, Q) \quad \text{and} \quad m_{\mathrm{rec}}(\tau, Q) := \lim_{\varepsilon \searrow 0} \limsup_{T\to\infty} \log r_{\mathrm{rec,mem}}(T, \varepsilon, \tau, Q), \quad (15)$$

where $r_{\text{inv,mem}}(T, \varepsilon, \tau, Q)$ and $r_{\text{rec,mem}}(T, \varepsilon, \tau, Q)$ are the minimum cardinalities of $(T, \varepsilon, \tau, Q)$-invariance and recurrence alphabets, respectively. Note that one could consider having the recurrence time and the upper bound on the durations of the signals in the alphabet represented using separate parameters for the recurrence memory definition, instead of considering them both to be equal to τ as in the definition above, but we avoid that here to keep the presentation simple.

Remark 4. *Fix a $\tau \geq 0$ and a compact set $Q \subset \mathbb{R}^n$. Then, $m_{\text{inv}}(\tau, Q)$ is finite if and only if there exists a finite (τ, Q)-invariance alphabet, i.e., a finite set of signals such that for any $T > 0$, they can be concatenated to define signals over $[0, T]$ that keep the trajectories (T, Q)-invariant. This follows from the observation that $r_{\text{inv,mem}}(T, \varepsilon, \tau, Q)$ is monotonically increasing as T increases and ε decreases. Consequently, $m_{\text{inv}}(\tau, Q) = \log r_{\text{inv,mem}}(\tau, Q)$, where $r_{\text{inv,mem}}(\tau, Q)$ is the cardinality of the minimal (τ, Q)-invariance alphabet, since $r_{\text{inv,mem}}(\tau, Q) = \lim_{\varepsilon \searrow 0} \limsup_{T\to\infty} r_{\text{inv,mem}}(T, \varepsilon, \tau, Q)$. Thus, in contrast with invariance entropy, a* strict *version of m_{inv}, where the alphabets should be able to generate strict invariance spanning sets, would not result in a different value. A similar argument can be made for recurrence memory to show that it is finite if and only if there exists a finite (τ, Q)-recurrence alphabet and that $m_{\text{rec}}(\tau, Q) = \log r_{\text{rec,mem}}(\tau, Q)$, where $r_{\text{rec,mem}}(\tau, Q)$ is the cardinality of the minimal such alphabet.*

In the following example, we show that if a system has a zero invariance entropy, that does not imply that it has a finite invariance memory. Moreover, it is possible for a system to have an infinite invariance memory and a zero recurrence memory corresponding to the same controlled invariant set, providing further insight about the complexity difference between recurrence and invariance. Our example follows Example 5.1 in (Colonius and Kawan, 2009).

Example 2. *Consider the system $\dot{x} = -x + u$, where $x \in \mathbb{R}$ and let $Q \subset [-1, 1]$ be an infinite set that is totally disconnected and such that for any $\delta > 0$, $[-\delta, \delta] \cap Q$ is an infinite set. It follows from (Colonius and Kawan, 2009) that $h_{\text{inv}}(Q) = 0$ and $h^*_{\text{inv}}(Q) = \infty$. However, observe that $u = 0$ is sufficient to make any trajectory starting from Q to be (τ, Q)-recurrent, for any $\tau > 0$, as the system is exponentially stable. Thus, $m_{\text{rec}}(\tau, Q) = \log 1 = 0$. On the other hand, by the same argument in (Colonius and Kawan, 2009), there is no finite set of control signals that can keep Q invariant, since starting from any state $x_0 \in Q$, only the signal u, where $u(t) = x_0$ for all $t \geq 0$, results in $\xi(x_0, u, t) \in Q$ for all $t \geq 0$. An infinite alphabet is thus needed to construct an invariance spanning set. Consequently, the limit of $r_{\text{inv,mem}}(\tau, T, \varepsilon, Q)$, as T goes to infinity and as ε goes to zero, is infinity, and $m_{\text{inv}}(\tau, Q) = \infty$, for any $\tau > 0$.*

Remark 5. *Fix $\tau > 0$, a $Q \subseteq \mathbb{R}^n$ that is controlled invariant, and a set $K \subseteq Q$. Let $m_{\text{inv}}(\tau, K, Q)$ and $m_{\text{rec}}(\tau, K, Q)$ be the invariance memories when K is the initial set of states and Q is the set to keep invariant or recurrent. Then, with similar arguments to those in the proof of Theorem 1, we can show that $m_{\text{inv}}(\tau, Q, B_{\delta_\tau}(Q)) \leq m_{\text{rec}}(\tau, Q) \leq m_{\text{inv}}(\tau, Q)$, where $\delta_\tau = F_Q \tau e^{L_\tau \tau}$, which is the right-hand-side of the containment lemma.*

We now show that our notions of memory can be lower-bounded using the strict entropy of the control task.

Theorem 8. *Fix any $\tau > 0$ and a compact set $Q \subset \mathbb{R}^n$. If $m_{\text{inv}}(\tau, Q)$ is finite, then $h^*_{\text{inv}}(Q) \leq \frac{1}{\tau_m} m_{\text{inv}}(\tau, Q)$, where $\tau_m = \min_{i \in [|S_\tau|]} t_i$, and S_τ is the minimal (τ, Q)-invariance alphabet. Similarly, if $m_{\text{rec}}(\tau, Q)$ is finite, then $h^*_{\text{rec}}(\tau, Q) \leq \frac{1}{\tau_m} m_{\text{rec}}(\tau, Q)$, where $\tau_m = \min_{i \in [|S_\tau|]} t_i$ and S_τ is the minimal (τ, Q)-recurrence alphabet.*

Proof. Assume that $m_{\text{inv}}(\tau, Q)$ is finite. Then, there exists a finite minimal (τ, Q)-invariance alphabet S_τ, as discussed in Remark 4. By definition, for any $T \geq 0$, an invariance (T, Q)-spanning set can be generated by concatenating the signals in S_τ. The number of signals over the interval $[0, T]$ that can be generated this way is at most $|S_\tau|^{\lceil T/\tau_m \rceil}$, where $\tau_m = \min_{i \in [|S_\tau|]} t_i$. It follows that $r_{\text{inv}}(T, Q)$, i.e., the cardinality of the minimal invariance (T, Q)-spanning set, is upper-bounded by $\left(r_{\text{inv,mem}}(\tau, Q)\right)^{\lceil T/\tau_m \rceil}$. If we substitute that upper bound in the definition of h^*_{inv}, we get

$$h^*_{\text{inv}}(Q) \leq \limsup_{T\to\infty} \frac{1}{T} \log \left(r_{\text{inv,mem}}(\tau, Q)\right)^{\lceil T/\tau_m \rceil} = \frac{1}{\tau_m} m_{\text{inv}}.$$

The last step follows from adapting Proposition 3.4(ii) in (Colonius and Kawan, 2009) to strict invariance entropy, instead of invariance entropy, to get that

$$h^*_{\text{inv}}(Q) = \limsup_{n\to\infty} \frac{1}{n\tau} \log r^*_{\text{inv}}(n\tau, Q),$$

which implies that assuming that $\lceil T/\tau_m \rceil = T/\tau_m$ in the limit does not change entropy. A similar proof can be followed to obtain the theorem's claim for the recurrence case. □

The next theorem shows that, in the case of recurrence, if we allow for slightly stricter assumptions than Q being control recurrent, one can get a finite upper bound on the recurrence memory.

Theorem 9. *Assume that that there exists an $\epsilon > 0$ s.t. $B_{-\epsilon}(Q) := \{x \in Q \mid \|x - y\| \geq \epsilon,\ \forall y \in \partial Q\}$, where ∂Q is the boundary of Q, is non-empty and that for every $x \in Q$, there exists a control signal u such that $\exists t \in (0, \tau]$, s.t. $\xi(x, u, t) \in B_{-\epsilon}(Q)$. Then, $m_{\text{rec}}(\tau, Q) \leq \log b(\epsilon e^{-L_\tau \tau}, Q)$.*

Proof. We follow a similar construction to the proof of Theorem 3. Let C be a minimal $\epsilon e^{-L_\tau \tau}$-cover of Q[1]. By the theorem's assumption, there exists a set $S_\tau = \{v_i\}_{i \in [|C|]} \subset \mathcal{U}$, where for every center x_i in the cover, there exists a $v_i \in S_\tau$ and $t_i \in (0, \tau]$ such that $\xi(x_i, v_i, t_i) \in B_{-\epsilon}(Q)$. We assume that the domain of v_i is $[0, t_i]$. Using Grönwall's inequality, $\forall t \in [0, t_i]$ and $\forall x \in B_{\epsilon e^{-L_{\tau,\epsilon}\tau}}(x_i) \cap Q$, $\|\xi(x_i, v_i, t) - \xi(x, v_i, t)\| \leq e^{L_{\tau,\epsilon} t}\|x_i - x\| \leq e^{L_{\tau,\epsilon} t}(\epsilon e^{-L_{\tau,\epsilon}\tau}) \leq \epsilon$.

We will prove by induction that S_τ is a (T, τ, Q)-recurrence alphabet. Fix any $x \in Q$ and $T \geq 0$. If $T \leq \tau$, the trajectory resulting from starting from x and following any signal in S_τ is trivially (T, τ, Q)-recurrent. Hence, we assume that $T \geq \tau$. We will concatenate some of the signals in S_τ to get a signal u that is defined over $[0, T]$ that leads the trajectory starting from x to be a (T, τ, Q)-recurrent one. Let x_i be the center in C such that $x \in B_{\epsilon e^{-L_{\tau,\epsilon}\tau}}(x_i)$. Since $\|\xi(x_i, v_i, t_i) - \xi(x, v_i, t_i)\| \leq \epsilon$ and $\xi(x_i, v_i, t_i) \in B_{-\epsilon}(Q)$, then $\xi(x, v_i, t_i) \in Q$. Thus, we let $u(t) = v_i(t)$ for all $t \in [0, t_i]$. Assume that we constructed u until some time $s \leq T$ as a concatenation of signals in S_τ such that $\xi(x, u, s) \in Q$ and ξ is (s, τ, Q)-recurrent. Thus, there exists an $x_j \in C$ such that $\xi(x, u, s) \in B_{\epsilon e^{-L_{\tau,\epsilon}\tau}}(x_j)$. Using the same analysis again, we get that $\|\xi(x_j, v_j, t_j) - \xi(\xi(x, u, s), v_j, t_j)\| \leq \epsilon$. We set $u(s + t) := v_j(t), \forall t \in (0, \min(t_j, T - s)]$. If $t_j \leq T - s$, then $\xi(x, u, s + t_j) = \xi(\xi(x, u, s), v_j, t_j) \in Q$. Finally, notice that the resulting trajectory visits Q every τ seconds at least once. Hence, u drives the trajectory starting from x to be (T, τ, Q)-recurrent and S_τ is a (T, τ, Q)-recurrence alphabet.

Consequently, $r_{\text{rec,mem}}(T, \tau, Q) \leq |S_\tau| = |C| = b(\epsilon e^{-L_{\tau,\epsilon}\tau}, Q)$. Substituting this bound in the definition of $m_{\text{rec}}(\tau, Q)$ results in the theorem. □

Observe that under the same mild conditions of Theorem 9, when combined with Theorem 8, it results in the upper bound on strict recurrence entropy: $h^*_{\text{rec}}(\tau, Q) \leq \frac{1}{\tau_m} \log b(\epsilon e^{-L_\tau \tau}, Q)$, where τ_m and ϵ are as defined in the theorems. We end with an application of Theorem 9 to the dynamics and the set of Example 1, which is another example of finite recurrence memory and an infinite invariance one.

Remark 6. *Consider the system in Example 1. Observe that for any state in Q, we can construct a constant control signal that drives the trajectory to $B_{-\epsilon}(Q)$ in $1 + \sqrt{2}$ time units, where $\epsilon = 0.5$. Thus, by Theorem 9, $m_{\text{rec}}(\tau, Q) \leq \log b(\epsilon e^{-L_\tau \tau}, Q)$, which when substituting $\tau = 1 + \sqrt{2}$, $\epsilon = 0.5$, and $L = 1$, we get $m_{\text{rec}}(\tau, Q) \leq \log b(\epsilon e^{-L_\tau \tau}, Q) \leq \log \lceil \frac{2}{e^{-1-\sqrt{2}}} \rceil^2 = 9.05$. In contrast, by Theorem 8 and the fact that $h^*_{\text{inv}}(Q) \geq h_{\text{inv}}(Q)$ (Proposition 3.1 of (Colonius and Kawan, 2009)), $m_{\text{inv}}(\tau, Q) \geq \tau h^*_{\text{inv}}(Q) \geq \tau h_{\text{inv}}(Q) = \infty$.*

10. Conclusions and Future Work

We present the notion of τ-recurrence entropy for nonlinear control systems as a generalization of the notion of invariance entropy. A system is τ-recurrent to a compact set of states if its trajectories starting from that set can leave it but only for τ seconds at a time. τ-recurrence entropy measures the rate at which the number of control signals that are sufficient to make the system τ-recurrent with respect to a predetermined set increases with time. We show that τ-recurrence entropy is bounded from above and below by the invariance entropy of the system with respect to different compact sets. Moreover, we show that it converges to invariance entropy with respect to the same set as τ decreases to zero, as expected. Then, we derive upper and lower bounds on τ-recurrence entropy as a function of the system dimension, local Lipschitz constant, and the divergence of the vector field. We show that both bounds converge to known corresponding bounds on invariance entropy as $\tau \to 0$, also as expected, and connect the newly defined entropy notion with necessary bit rates for recurrence, providing an algorithm that asymptotically achieves this task. We finalize by showing that, in fact, in many situations, recurrence can be obtained by a finite alphabet of control signals.

When put together, our results suggest that recurrence is a less complex task, in the sense that in many instances, it not only requires smaller spanning sets to perform the task but there are also instances where such spanning sets can be generated via a finite alphabet. This opens up the possibility of storing the control signals and thus avoiding the need to compute controllers online. Future research directions along these lines includes understanding other classes of control

[1] Notice that we use τ to construct the cover, instead of T, which we used in the proof of Theorem 3.

tasks, where a hierarchy of task complexity can be defined, as is the case of invariance and recurrence, as well as other instances where finite alphabets can be generated.

11. Acknowledgements

This work has been supported by the National Science Foundation under the grants CAREER 1752362, CPS 2136324, and Global Centers 2330450.

References

Adler, R.L., Konheim, A.G., McAndrew, M.H., 1965. Topological entropy. Transactions of the American Mathematical Society 114, 309–319.

Bowen, R., 1971a. Entropy for group endomorphisms and homogeneous spaces. Transactions of the American Mathematical Society 153, 401–414.

Bowen, R., 1971b. Periodic points and measures for axiom a diffeomorphisms. Trans. Amer. Math. Soc 154.

Colonius, F., 2012. Minimal bit rates and entropy for exponential stabilization. SIAM Journal on Control and Optimization 50, 2988–3010.

Colonius, F., Kawan, C., 2009. Invariance entropy for control systems. SIAM Journal on Control and Optimization 48, 1701–1721.

Colonius, F., Kawan, C., 2011. Invariance entropy for outputs. Mathematics of Control, Signals, and Systems 22, 203–227.

Colonius, F., Kawan, C., Nair, G., 2013. A note on topological feedback entropy and invariance entropy. Systems & Control Letters 62, 377–381.

Delchamps, D.F., 1990. Stabilizing a linear system with quantized state feedback. IEEE transactions on automatic control 35, 916–924.

Genesio, R., Tartaglia, M., Vicino, A., 1985. On the estimation of asymptotic stability regions: State of the art and new proposals. IEEE Transactions on automatic control 30, 747–755.

Katok, A., Hasselblatt, B., 1995. Introduction to the modern theory of dynamical systems. 54, Cambridge university press.

Kawan, C., Delvenne, J.C., 2016. Network entropy and data rates required for networked control. IEEE Transactions on Control of Network Systems 3, 57–66. doi:path10.1109/TCNS.2015.2440551.

Kawan, C., Yüksel, S., 2018. On optimal coding of non-linear dynamical systems. IEEE Transactions on Information Theory 64, 6816–6829.

Khalil, H.K., 2002. Nonlinear systems; 3rd ed.

Liberzon, D., Mitra, S., 2016. Entropy and minimal data rates for state estimation and model detection, in: Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control, pp. 247–256.

Lin, Y., Sontag, E.D., Wang, Y., 1996. A smooth converse lyapunov theorem for robust stability. SIAM Journal on Control and Optimization 34, 124–160.

Matveev, A., Pogromsky, A., 2016. Observation of nonlinear systems via finite capacity channels: Constructive data rate limits. Automatica 70, 217–229.

Matveev, A.S., Pogromsky, A.Y., 2019. Observation of nonlinear systems via finite capacity channels, part ii: Restoration entropy and its estimates. Automatica 103, 189–199.

Meyer, P.J., Yin, H., Brodtkorb, A.H., Arcak, M., Sørensen, A.J., 2020. Continuous and discrete abstractions for planning, applied to ship docking. ArXiv abs/1911.09773.

Nair, G.N., Evans, R.J., Mareels, I.M., Moran, W., 2004. Topological feedback entropy and nonlinear stabilization. IEEE Transactions on Automatic Control 49, 1585–1597.

Rungger, M., Zamani, M., 2016. Scots: A tool for the synthesis of symbolic controllers, in: Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control, Association for Computing Machinery, New York, NY, USA. p. 99–104. URL: `https://doi.org/10.1145/2883817.2883834`, doi:path10.1145/2883817.2883834.

Rungger, M., Zamani, M., 2017. Invariance feedback entropy of nondeterministic control systems, in: Proceedings of the 20th International Conference on Hybrid Systems: Computation and Control, Association for Computing Machinery, New York, NY, USA. p. 91–100. URL: `https://doi.org/10.1145/3049797.3049801`, doi:path10.1145/3049797.3049801.

Savkin, A.V., 2006. Analysis and synthesis of networked control systems: Topological entropy, observability, robustness and optimal control. Automatica 42, 51–62.

Shen, Y., Bichuch, M., Mallada, E., 2022. Model-free learning of regions of attraction via recurrent sets, in: 61st IEEE Conference on Decision and Control (CDC), pp. 4714–4719. URL: `http://mallada.ece.jhu.edu/pubs/2022-CDC-SBM.pdf`, doi:path10.1109/CDC51059.2022.9993280.

Sibai, H., Huriot, S., Martin, T., Arcak, M., 2024. Symmetry-based abstraction algorithm for accelerating symbolic control synthesis, in: Benz, N., Gopinath, D., Shi, N. (Eds.), NASA Formal Methods, Springer Nature Switzerland, Cham. pp. 166–185.

Sibai, H., Mitra, S., 2017. Optimal data rate for state estimation of switched nonlinear systems, in: Proceedings of the 20th International Conference on Hybrid Systems: Computation and Control, pp. 71–80.

Sibai, H., Mitra, S., 2018. State estimation of dynamical systems with unknown inputs: Entropy and bit rates, in: Proceedings of the 21st International Conference on Hybrid Systems: Computation and Control (Part of CPS Week), Association for Computing Machinery, New York, NY, USA. p. 217–226. URL: `https://doi.org/10.1145/3178126.3178150`, doi:path10.1145/3178126.3178150.

Sibai, H., Mitra, S., 2023. State estimation of continuous-time dynamical systems with uncertain inputs with bounded variation: Entropy, bit rates, and relation with switched systems. IEEE Transactions on Automatic Control , 1–16doi:path10.1109/TAC.2023.3250510.

Siegelmann, R., Shen, Y., Paganini, F., Mallada, E., 2023. A recurrence-based direct method for stability analysis and GPU-based verification of non-monotonic Lyapunov functions, in: 62nd IEEE Conference on Decision and Control (CDC).

Tabuada, P., 2009. Verification and Control of Hybrid Systems: A Symbolic Approach. 1st ed., Springer Publishing Company, Incorporated.

Tomar, M.S., Kawan, C., Zamani, M., 2022. Numerical over-approximation of invariance entropy via finite abstractions. Systems & Control Letters 170, 105395.

Tomar, M.S., Rungger, M., Zamani, M., 2021. Invariance feedback entropy of uncertain control systems. IEEE Transactions on Automatic Control 66, 5680–5695. doi:path10.1109/TAC.2020.3038702.

Tomar, M.S., Zamani, M., 2022. On a notion of entropy for reachability properties, in: 2022 American Control Conference (ACC), pp. 781–786. doi:path10.23919/ACC53348.2022.9867868.

Tomar, M.S., Zamani, M., 2023. Towards minimal data rate enforcing regular safety properties: An invariance entropy approach. IEEE Control Systems Letters .