J Nonlinear Sci (2009) 19: 497–545
DOI 10.1007/s00332-009-9042-5

Journal of
Nonlinear
Science

Reliability of Coupled Oscillators

Kevin K. Lin · Eric Shea-Brown · Lai-Sang Young

Received: 27 August 2007 / Accepted: 6 January 2009 / Published online: 19 February 2009

Abstract We study the *reliability* of phase oscillator networks in response to fluctuating inputs. Reliability means that an input elicits essentially identical responses upon repeated presentations, regardless of the network's initial condition. Single oscillators are well known to be reliable. We show in this paper that unreliable behavior can occur in a network as small as a coupled oscillator pair in which the signal is received by the first oscillator and relayed to the second with feedback. A geometric explanation based on shear-induced chaos at the onset of phase-locking is proposed. We treat larger networks as decomposed into modules connected by acyclic graphs, and give a mathematical analysis of the acyclic parts. Moreover, for networks in this class, we show how the source of unreliability can be localized, and address questions concerning downstream propagation of unreliability once it is produced.

Keywords Coupled oscillators · Random dynamical systems · Neural network dynamics

K.K. Lin (✉)
Department of Mathematics, University of Arizona, Tucson, AZ, USA
e-mail: klin@math.arizona.edu

E. Shea-Brown
Department of Applied Mathematics, University of Washington, Seattle, WA, USA
e-mail: etsb@amath.washington.edu

L.-S. Young
Courant Institute of Mathematical Sciences, New York University, New York, NY, USA
e-mail: lsy@cims.nyu.edu

1 Introduction

This paper contains a mathematical treatment of the question of *reliability* in the context of oscillator networks. Reliability here refers to whether a system produces identical responses when it is repeatedly presented with the same stimulus. Such questions are relevant to signal processing in biological and engineered systems. Consider, for example, a network of interconnected neurons with some background activity. An external stimulus in the form of a time-dependent signal is applied to this neural circuitry, which processes the signal and produces a response in the form of voltage spikes. We say the system is *reliable* if, independent of its state at the time of presentation, the same stimulus elicits essentially identical responses following an initial period of adjustment, that is, the response to a given signal is reproducible (Pikovsky et al. 2001; Bryant and Segundo 1976; Mainen and Sejnowski 1995; Nakao et al. 2005; Zhou and Kurths 2003; Teramae and Tanaka 2004; Ritt 2003; Pakdaman and Mestivier 2001; Kosmidis and Pakdaman 2003; Goldobin and Pikovsky 2005, 2006; Teramae and Fukai 2007).

The present study is carried out in the context of (heterogeneous) networks of interconnected oscillators. We assume the input signal is received by some components of the network and relayed to others, possibly in the presence of feedback connections. Our aim is to understand the relation between a network's reliability properties and its architecture (or "circuit diagram") and the strengths of various connections. This problem is quite different from the simpler and much studied situation of uncoupled oscillators driven by a common input. The latter is equivalent to single, isolated oscillators, the reliability of which has been studied extensively (Pikovsky et al. 2001; Nakao et al. 2005; Zhou and Kurths 2003; Teramae and Tanaka 2004; Ritt 2003; Pakdaman and Mestivier 2001; Kosmidis and Pakdaman 2003; Goldobin and Pikovsky 2005, 2006). To simplify the analysis, we assume the constituent oscillators are *phase oscillators* or circle rotators, and that they are driven by a fluctuating input which, for simplicity, we take to be white noise. Under these conditions, systems consisting of a single, isolated phase oscillator have been shown to be reliable; see, e.g., Teramae and Tanaka (2004), Ritt (2003).

Our results are presented in 3 separate parts. As the reader will see from the brief description below, Parts II and III are largely independent, while they both build on the material in Part I.

Part I: Model and Preliminaries. This part contains a detailed description of the models studied in this paper, a formal definition of reliability, and a review of random dynamical systems theory which justifies the use of Lyapunov exponents as a measure of reliability.

Part II: Two-Oscillator Systems. Part II contains an in-depth study of a small network consisting of two oscillators. It is motivated in part by the following naive (and partly rhetorical) question: *Are networks of coupled phase oscillators reliable, and if not, how large must a network be to exhibit unreliable behavior?* Our answer to this question is that unreliable behavior occurs already in 2-oscillator systems with recurrent connections. Our results demonstrate clearly that such a system can be reliable

or unreliable, and that both types of behaviors are quite prominent depending on the strengths of the feedforward and feedback couplings.

Geometric explanations are given for some of the phenomena observed. Our most striking result, which occupies much of Part II, is to relate reliability properties in response to weak stimuli to shear-induced chaos. Recent advances in dynamical systems theory have identified a mechanism for producing chaos via the interaction of forcing with the underlying shear in a system. The dynamical environment near the onset of phase-locking is particularly susceptible to this mechanism. Applying the cited theory, which we review, we are able to predict the reliability or lack thereof for a range of coupling parameters. At low drive amplitudes, this is the primary cause for unreliability.

Part III. Larger Networks. A clear message from Part II is that networks, even small networks, with recurrent connections can exhibit very complex behavior, and that a complete classification of their reliability properties is not likely to be feasible. On the other hand, *acyclic* networks, or networks whose coupling connections have no cycles, are quite tractable. These two observations suggest that it may be profitable to group nodes which are interconnected with recurrent couplings into *modules*, so that the network as a whole can be described by an acyclic graph connecting these modules. That is to say, we isolate the parts that are too complicated to analyze, treat them as black boxes, and focus on the remaining picture. Such an approach can be effective for certain types of networks, and it is the approach taken in much of Part III.

A rigorous mathematical analysis of the decomposition into modules, including a representation of the Lyapunov exponents of the network in terms of those within modules, is presented. Via this decomposition, we are able to localize the source of any unreliability that may be present, and to study which sites are affected. Among our findings is that once produced, unreliability propagates; no site downstream can be completely reliable. We finish by raising what we believe is an important but seldom addressed question, namely how to deduce reliability of a system from knowledge of its parts. Simple examples are given to demonstrate that this can sometimes be done.

This completes our brief description of what is in each of the 3 parts.

Rigorous mathematical results in this paper are presented in the form of Theorems, Propositions, etc. Simulations and heuristic reasoning are used abundantly where rigorous analysis is not available. Many of our findings are qualitative. Our aim is to identify and explain relevant phenomena, and to make contact with rigorous mathematics whenever we can, hoping in the long run to help bring dynamical systems theory closer to biologically relevant systems.

There is a vast literature on networks of oscillators; it is impossible to give general citations without serious omissions. We have limited our citations to a sample of papers and books with settings closer to ours or which treat specifically the topic of reliability (see below). We mention in particular the preprint (Teramae and Fukai 2007), which discusses the reliability of large, sparsely coupled networks in a way that complements ours.

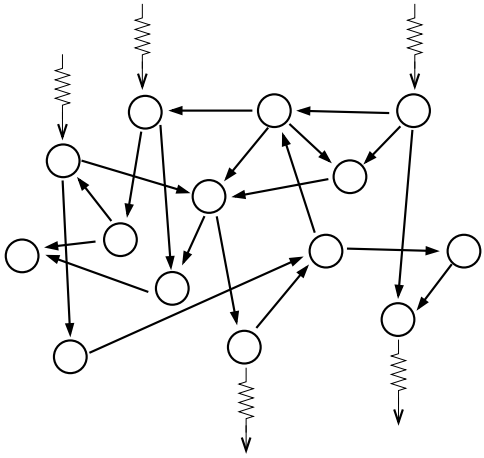

Fig. 1 Schematic of a forced oscillator network

2 Part I. Model and Preliminaries

2.1 Model

2.1.1 Coupled Phase Oscillator Systems

We consider in this paper networks of coupled phase oscillators in the presence of external stimuli. A schematic picture of such a setup is shown in Fig. 1. The networks considered are of the type that arise in many settings; see, e.g., Pikovsky et al. (2001), Strogatz (2000), Kuramoto (1997), Hoppensteadt and Izhikevich (1997), Winfree (2001), Rinzel and Ermentrout (1998), Brown et al. (2003), Ermentrout (1991), Ermentrout and Kopell (1991).

The *unforced* systems, i.e., the systems without external stimuli, are described by

$$\dot{\theta}_i = \omega_i + z(\theta_i) \sum_{j \neq i} a_{ji}\, g(\theta_j), \quad i = 1, \ldots, N. \tag{1}$$

Except in Part II (where we treat two-oscillator systems), N can be any number ≥ 1; in particular, it can be arbitrarily large. The state of oscillator i is described by an angular variable $\theta_i \in \mathbb{S}^1 \equiv \mathbb{R}/\mathbb{Z}$, $i = 1, \ldots, N$. Its intrinsic frequency is given by a constant ω_i. We allow these frequencies to vary from oscillator to oscillator. The second term on the right represents the coupling: $a_{ji} \in \mathbb{R}$, g is a "bump function" vanishing outside of a small interval $(-b, b)$; on $(-b, b)$, it is smooth, ≥ 0, and satisfies $\int_{-b}^{b} g(\theta)\, d\theta = 1$; $z(\theta)$ is a function which in this paper is taken to be $z(\theta) = \frac{1}{2\pi}(1 - \cos(2\pi\theta))$. The meaning of this coupling term is as follows: We say the jth oscillator "spikes" when $\theta_j(t) = 0$. Around the time that an oscillator spikes, it emits a pulse which modifies some of the other oscillators (Peskin 1988; Strogatz and Mirollo 1990; Herz and Hopfield 1995; Nunes and Pereira 1985). The sign and magnitude of a_{ji} describe how oscillator j affects oscillator i: $a_{ji} > 0$ (resp., $a_{ji} < 0$) means oscillator j excites (resp., inhibits) oscillator i when it spikes, and $a_{ji} = 0$ means oscillator i is not directly affected. In this paper, b is taken to be

about $\frac{1}{20}$, and the a_{ji} are taken to be $\mathcal{O}(1)$. Finally, the function $z(\theta)$, often called the phase response curve (Winfree 1974; Guckenheimer 1975; Ermentrout 1996; Kuramoto 1997), measures the variable sensitivity of an oscillator to coupling and stimulus inputs (see below).

The *stimulus-driven* systems are of the form

$$\begin{aligned} d\theta_i &= \omega_i\, dt + z(\theta_i)\Big(\sum_{j\neq i} a_{ji}\, g(\theta_j)\, dt + I_i(t)\Big), \\ I_i(t) &= \varepsilon_i\, dW_i(t), \quad i = 1, \ldots, N. \end{aligned} \tag{2}$$

Here, $I_i(t)$ is the external stimulus received by oscillator i; it is taken to be white noise with amplitude ϵ_i. For simplicity, we assume that for $i \neq j$, $dW_i(t)$ and $dW_j(t)$ are either independent, or they are one and the same, i.e., the same input is received by both oscillators although the amplitudes may differ. As discussed earlier, we are primarily interested in situations where inputs are received by only a subset of the oscillators, with $\epsilon_i \approx 0$ for the rest. Likewise, we are sometimes interested in the response registered at only a subset of the oscillators rather than the whole network.

2.1.2 Neuroscience Interpretations

Coupled phase oscillators arise in many settings (Pikovsky et al. 2001; Strogatz 2000; Kuramoto 1997; Hoppensteadt and Izhikevich 1997; Winfree 2001; Rinzel and Ermentrout 1998). Here, we briefly discuss their use in mathematical neuroscience.

We think of phase oscillators as paradigms for systems with rhythmic behavior. Such models are often derived as limiting cases of oscillator models in two or more dimensions. In particular, the specific form of $z(\cdot)$ chosen here corresponds to the normal form for oscillators near saddle-node bifurcations on their limit cycles (Ermentrout 1996). This situation is typical in neuroscience, where neural models with $z(\theta) \approx 1 - \cos(\theta)$ are referred to as "Type I." The pulse- or spike-based coupling implemented by $g(\cdot)$ may also be motivated by the synaptic impulses sent between neurons after they fire action potentials (although this is not the only setting in which pulsatile coupling arises) (Winfree 2001; Brown et al. 2003; Ermentrout and Kopell 1991; Taylor and Holmes 1998; Strogatz and Mirollo 1990; Herz and Hopfield 1995; Nunes and Pereira 1985; Peskin 1988; Hoppensteadt and Izhikevich 1997).

The general conclusions that we will present do not depend on the specific choices of $z(\cdot)$ and $g(\cdot)$, but rather on their qualitative features. Specifically, we have checked that our main results about reliability and phase locking are essentially unchanged when the $z(\cdot)$ function becomes asymmetric and the location of the $g(\cdot)$ impulse is somewhat shifted, as would correspond more closely to neuroscience (Ermentrout 1996). Therefore, the behavior we find here can be expected to be fairly prototypical for pulse-coupled networks of Type I oscillators.

A standard way to investigate the reliability of a system of spiking oscillators—both in the laboratory and in simulations—is to repeat the experiment using a different initial condition each time but driving the system with *the same* stimulus $\varepsilon I(t)$, and to record spike times in raster plots. Figure 2 shows such a raster plot for an

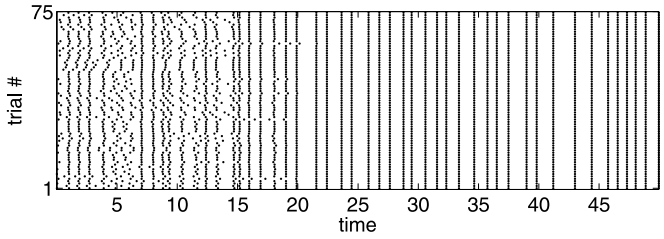

Fig. 2 Raster plot for an isolated oscillator. In each experiment, 75 trials are performed, and *dots* are placed at spike times. Nearly identical spike times are observed after a transient, indicating reliability

isolated oscillator of the present type, as studied by Ritt (2003), Gutkin et al. (2003). Note that, for each repetition, the oscillator produces essentially identical spike times after a transient, demonstrating its reliability.

2.2 Measuring Reliability

2.2.1 A Working Definition of Reliability

We attempt to give a formal definition of reliability in a general setting. Consider a dynamical system on a domain M (such as a manifold or a subset of Euclidean space). A signal in the form of $\iota(t) \in \mathbb{R}^n, t \in [0, \infty)$, is presented to the system. The response $F(t)$ of the system to this signal is given by $F(t) = F(x_0, t, \{\iota(s)\}_{0 \le s < t})$. That is to say, the response at time t may, in principle, depend on $x_0 \in M$, the initial state of the system when the signal is presented, and the values of the signal up to time t.

In the model as described in Sect. 2.2.1, $F(t)$ can be thought of as the N-tuple $(\theta_1(t), \theta_2(t), \ldots, \theta_N(t))$ or $\Psi(\theta_1(t), \theta_2(t), \ldots, \theta_N(t))$, the value of an observable at time t.

We propose now one way to define reliability. Given a dynamical system, a class of signals $\mathcal{I}$, and a response function F, we say the system is *reliable* if for almost all $\iota \in \mathcal{I}$ and $x_0, x_0' \in M$,

$$\left\| F\left(x_0, t, \left\{\iota(s)\right\}_{0 \le s < t}\right) - F\left(x_0', t, \left\{\iota(s)\right\}_{0 \le s < t}\right) \right\| \to 0 \quad \text{as } t \to \infty.$$

Here $\|\cdot\|$ is a norm on the range space of F. We do not claim that this is the only way to capture the idea of reliability, but will use it as our operational definition.

We point out some of the pitfalls of this definition: In practice, signals are never presented for infinite times, and in some situations, responses can be regarded as reliable only if the convergence above is rapid. By the same token, not all initial conditions are equally likely, leaving room for probabilistic interpretations.

Finally, one should not expect unreliable responses to be fully random. On the contrary, as we will show in Sect. 2.2.2, they tend to possess a great deal of structure, forming what are known as *random strange attractors*.

2.2.2 *Reliability, Lyapunov Exponents, and Random Attractors*

We discuss here some mathematical tools that can be used to quantify how reliable or unreliable a driven system is. With $I(t)$ taken to be realizations of white noise, (2) can be put into the framework of a *random dynamical system* (RDS). We begin by reviewing some relevant mathematics (Arnold 2003; Baxendale 1992). Consider a general stochastic differential equation (SDE)

$$dx_t = a(x_t)\,dt + \sum_{i=1}^{k} b_i(x_t) \circ dW_t^i. \tag{3}$$

In this general setting, $x_t \in M$ where M is a compact Riemannian manifold, the W_t^i are independent standard Brownian motions, and the equation is of Stratonovich type. Clearly, (2) is a special case of (3): $x_t = (\theta_1(t), \ldots, \theta_N(t))$, $M = \mathbb{T}^N \equiv \mathbb{S}^1 \times \mathbb{S}^1 \times \cdots \times \mathbb{S}^1$ (and we have the choice between the Itô or Stratonovich integral), and $k = 1$.

In general, one fixes an initial x_0, and looks at the distribution of x_t for $t > 0$. Under fairly general conditions, these distributions converge to the unique stationary measure μ, the density of which is given by the Fokker–Planck equation. Since reliability is about a system's reaction to a single stimulus, i.e., a single realization of $(W_t^1, \ldots, W_t^k)$, at a time, and concerns the simultaneous evolution of all or large ensembles of initial conditions, of relevance to us are not the distributions of x_t but *flow-maps* of the form $F_{t_1,t_2;\omega}$. Here $t_1 < t_2$ are two points in time, ω is a sample Brownian path, and $F_{t_1,t_2;\omega}(x_{t_1}) = x_{t_2}$ where x_t is the solution of (3) corresponding to ω. A well known theorem states that such *stochastic flows of diffeomorphisms* are well defined if the functions $a(x)$ and $b(x)$ in (3) are sufficiently smooth; see, e.g., Kunita (1990). More precisely, the maps $F_{t_1,t_2;\omega}$ are well defined for almost every ω, and they are invertible, smooth transformations with smooth inverses. Moreover, $F_{t_1,t_2;\omega}$ and $F_{t_3,t_4;\omega}$ are independent for $t_1 < t_2 < t_3 < t_4$. These results allow us to treat the evolution of systems described by (3) as compositions of random, *i.i.d.*, smooth maps. Many of the techniques for analyzing smooth deterministic systems have been extended to this random setting. We will refer to the resulting body of work as RDS theory.

Similarly, the stationary measure μ, which gives the steady-state distribution averaged over all realizations ω, does not describe what we see when studying a system's reliability. Of relevance are the *sample measures* $\{\mu_\omega\}$, which are the conditional measures of μ given the past. More precisely, we think of ω as defined for all $t \in (-\infty, \infty)$ and not just for $t > 0$. Then μ_ω describes what one sees at $t = 0$ given that the system has experienced the input defined by ω for all $t < 0$. Two useful facts about these sample measures are

(a) $(F_{-t,0;\omega})_*\mu \to \mu_\omega$ as $t \to \infty$, where $(F_{-t,0;\omega})_*\mu$ is the measure obtained by transporting μ forward by $F_{-t,0;\omega}$, and
(b) the family $\{\mu_\omega\}$ is invariant in the sense that $(F_{0,t;\omega})_*(\mu_\omega) = \mu_{\sigma_t(\omega)}$ where $\sigma_t(\omega)$ is the time-shift of the sample path ω by t.

Thus, if our initial distribution is given by a probability density ρ and we apply the stimulus corresponding to ω, then the distribution at time t is $(F_{0,t;\omega})_*\rho$. For

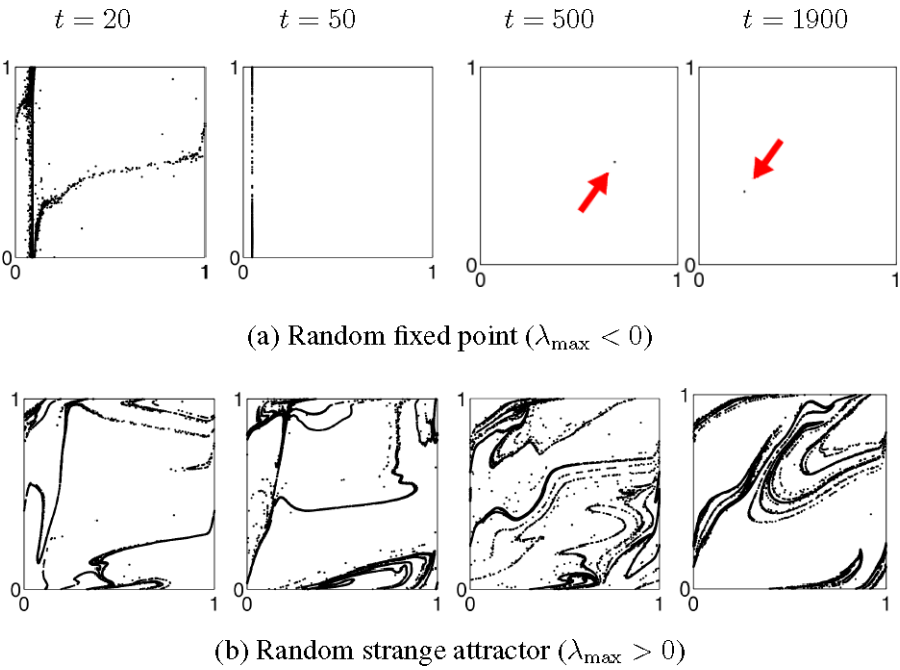

Fig. 3 Snapshots of sample measures for (2) with $N = 2$ oscillators, at various times in response to a single realization of the stimulus. Two different sets of parameters are used in (**a**) and (**b**). In (**a**), the sample measures converge to a random fixed point. In (**b**), the sample measures converge to a random strange attractor. See Theorem 1

t sufficiently large, one expects in most situations that $(F_{0,t;\omega})_*\rho$ is very close to $(F_{0,t;\omega})_*\mu$, which by (a) above is essentially given by $\mu_{\sigma_t(\omega)}$. The time-shift by t of ω is necessary because by definition, μ_ω is the conditional distribution of μ at time 0.

Figure 3 shows some snapshots of $(F_{0,t;\omega})_*\rho$ for a system with $N = 2$ cells, for two different sets of parameters. As noted earlier in the panels corresponding to $t \gg 1$, the distributions approximate $\mu_{\sigma_t(\omega)}$. In these simulations, the initial distribution ρ is the stationary density of (2) with a small-amplitude noise, the interpretation being that the system is intrinsically noisy even in the absence of external stimuli. Observe that these pictures evolve with time, and for large enough t, they have similar qualitative properties depending on the underlying system. This is in agreement with RDS theory, which tells us in fact that the $\mu_{\sigma_t(\omega)}$ obey a statistical law for almost all ω. Observe also the strikingly different behaviors in the top and bottom panels. We will follow up on this observation presently. First, we recall two mathematical results that describe the dichotomy.

In deterministic dynamics, Lyapunov exponents measure the exponential rates of separation of nearby trajectories. Let $\lambda_{\max}(x)$ denote the largest Lyapunov exponent along the trajectory starting from x. Then a positive $\lambda_{\max}$ for a large set of initial conditions is generally thought of as synonymous with chaos, while the presence of stable equilibria is characterized by $\lambda_{\max} < 0$. For smooth random dynamical systems, Lyapunov exponents are also known to be well defined; moreover, they are

nonrandom, i.e., they do not depend on ω. If, in addition, the invariant measure is ergodic, then $\lambda_{\max}$ is constant almost everywhere in the phase space. In Theorem 1 below, we present two results from RDS theory that together suggest that the sign of $\lambda_{\max}$ is a good criterion for distinguishing between reliable and unreliable behavior:

Theorem 1 *In the setting of* (3), *let* μ *be an ergodic stationary measure.*

(1) (*Random sinks*) (Le Jan 1987) *If* $\lambda_{\max} < 0$, *then with probability* 1, μ_ω *is supported on a finite set of points.*
(2) (*Random strange attractors*) (Ledrappier and Young 1988) *If* μ *has a density and* $\lambda_{\max} > 0$, *then with probability* 1, μ_ω *is a random SRB measure.*

In Part (1) above, if in addition to $\lambda_{\max} < 0$, two mild conditions (on the relative motions of two points) are assumed, then almost surely μ_ω is supported on a single point (Baxendale 1992). Observe that this corresponds exactly to reliability for almost every ω as defined in Sect. 2.2.1, namely the collapse of trajectories starting from almost all initial conditions to a point (which continues to evolve with time). See Fig. 3(a). In view of this interpretation, we will take the liberty to equate $\lambda_{\max} < 0$ with reliability in the rest of this paper.

The conclusion of Part (2) requires clarification: In deterministic dynamical systems theory, SRB measures are natural invariant measures that describe the asymptotic dynamics of chaotic dissipative systems (in the same way that Liouville measures are the natural invariant measures for Hamiltonian systems). SRB measures are typically singular. They are concentrated on unstable manifolds, which are families of curves, surfaces, etc. that wind around in a complicated way in the phase space (Eckmann and Ruelle 1985). Part (2) of Theorem 1 generalizes these ideas to random dynamical systems. Here, random (meaning ω-dependent) SRB measures live on random unstable manifolds, which are complicated families of curves, surfaces, etc. that evolve with time. In particular, in a system with random SRB measures, different initial conditions lead to very different outcomes at time t when acted on by the same stimulus; this is true for all $t > 0$, however large. It is, therefore, natural to regard $\lambda_{\max} > 0$ and the distinctive geometry of random SRB measures as a signature of unreliability.

In the special case where the phase space is a circle, such as in the case of a single oscillator, that the Lyapunov exponent λ is nonpositive is an immediate consequence of Jensen's Inequality. In more detail,

$$\lambda(x) = \lim_{t\to\infty} \frac{1}{t} \log F'_{0,t;\omega}(x)$$

for typical ω by definition. Integrating over initial conditions x, we obtain

$$\lambda = \int_{\mathbb{S}^1} \lim_{t\to\infty} \frac{1}{t} \log F'_{0,t;\omega}(x)\, dx = \lim_{t\to\infty} \frac{1}{t} \int_{\mathbb{S}^1} \log F'_{0,t;\omega}(x)\, dx.$$

The exchange of integral and limit is permissible because the required integrability conditions are satisfied in stochastic flows (Kifer 1986). Jensen's Inequality then

gives

$$\int_{\mathbb{S}^1} \log F'_{0,t;\omega}(x)\,dx \le \log \int_{\mathbb{S}^1} F'_{0,t;\omega}(x)\,dx = 0. \tag{4}$$

The equality above follows from the fact that $F_{0,t;\omega}$ is a circle diffeomorphism. Since the gap in the inequality in (4) is larger when $F'_{0,t;\omega}$ is farther from being a constant function, we see that $\lambda < 0$ corresponds to $F'_{0,t;\omega}$ becoming "exponentially uneven" as $t \to \infty$. This is consistent with the formation of random sinks.

The following results from general RDS theory shed light on the situation when the system is multi-dimensional:

Proposition 2.1 (See, e.g., Chap. 5 of Kifer (1986) or Le Jan (1987)) *In the setting of* (3), *assume* μ *has a density, and let* $\{\lambda_1, \ldots, \lambda_d\}$ *be the Lyapunov exponents of the system counted with multiplicity. Then*

(i) $\sum_i \lambda_i \le 0$.
(ii) $\sum_i \lambda_i = 0$ *if and only if* $F_{s,t,\omega}$ *preserves* μ *for almost all* ω *and all* $s < t$.
(iii) *If* $\sum_i \lambda_i < 0$, *and* $\lambda_i \ne 0$ *for all* i, *then* μ_ω *is singular.*

A formula giving the dimension of μ_ω is proved in Ledrappier and Young (1988) under mild additional conditions.

The reliability of a single oscillator, i.e., that $\lambda < 0$, is easily deduced from Proposition 2.1: μ has a density because the transition probabilities have densities, and no measure is preserved by all the $F_{s,t,\omega}$ because different stimuli distort the phase space differently. Proposition 2.1(i) and (ii) together imply that $\lambda < 0$. See also Pakdaman and Mestivier (2001), Nakao et al. (2005), Teramae and Tanaka (2004), Ritt (2003).

For the 2-oscillator system illustrated in Fig. 3, assuming that μ has a density (this is explained in Sect. 3.1), Proposition 2.1(i) and (ii) together imply that $\lambda_1 + \lambda_2 < 0$. Here $\lambda_1 = \lambda_{\max}$ can be positive, zero, or negative. If it is > 0, then it will follow from Proposition 2.1(i) that $\lambda_2 < 0$, and by Proposition 2.1(iii), the μ_ω are singular. From the geometry of random SRB measures, we conclude that different initial conditions are attracted to lower dimensional sets that depend on stimulus history. Thus, even in unreliable dynamics, the responses are highly structured and far from uniformly distributed, as illustrated in Fig. 3(b).

Finally, we observe that since $\lambda_{\max}$ is nonrandom, the reliability of a system is independent of the realization ω once the stimulus amplitude ε is fixed.

Remark In this paper, we compute $\lambda_{\max}$ numerically by solving the variational equation associated with the SDE (3) using the Milstein scheme. For examples of systems where analytical estimates of positive $\lambda_{\max}$ have been obtained, see Baxendale and Goukasian (2002), Arnold et al. (2004), Teramae and Fukai (2007).

3 Part II. Two-Oscillator Systems

We consider in Part II a special case of the general model described in Sect. 2.1.1, namely two coupled phase oscillators forced by an external stimulus as shown:

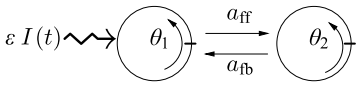

(5)

Equation (2) here simplifies to

$$\begin{aligned}\dot\theta_1 &= \omega_1 + a_{\rm fb} z(\theta_1) g(\theta_2) + \varepsilon z(\theta_1) I(t),\\ \dot\theta_2 &= \omega_2 + a_{\rm ff} z(\theta_2) g(\theta_1),\end{aligned} \tag{6}$$

where we have written $a_{\rm ff}$ and $a_{\rm fb}$ (for "feedforward" and "feedback") instead of a_{12} and a_{21}.

Part II is comprised of three sections. Our main result, which relates reliability properties at low drive amplitudes to the theory of shear induced chaos, is discussed in Sect. 3.2, after some preparatory material in Sect. 3.1 on the undriven system. Other results on the reliability profile as a function of $a_{\rm ff}, a_{\rm fb}$ and ε are discussed in Sect. 3.3.

3.1 Coupling Geometry and Zero-Input Dynamics

How a system responds to weak stimuli is strongly dependent on its geometry at zero input. This section sets the stage for Sect. 3.2 by discussing this geometry, focusing in particular on the set of parameters for which the system is phase-locked (Sect. 3.1.2).

3.1.1 Preliminary Observations

First we describe the flow φ_t on the 2-torus $\mathbb{T}^2$ defined by (6) when the stimulus is absent, i.e., when $\varepsilon = 0$. We begin with the case where the two oscillators are uncoupled, i.e., $a_{\rm ff} = a_{\rm fb} = 0$. In this special case, φ_t is a linear flow; depending on whether ω_1/ω_2 is rational, it is either periodic or quasiperiodic. Adding coupling distorts flow lines inside the two strips $\{|\theta_1| < b\}$ and $\{|\theta_2| < b\}$. These two strips correspond to the regions where one of the oscillators "spikes," transmitting a coupling impulse to the other (see Sect. 2.1.1). For example, if $a_{\rm ff} > 0$, then an orbit entering the vertical strip containing $\theta_1 = 0$ will bend upward. Because of the variable sensitivity of the receiving oscillator, the amount of bending depends on $a_{\rm ff}$ as well as the value of θ_2, with the effect being maximal near $\theta_2 = \frac{1}{2}$ and negligible near $\theta_2 = 0$ due to our choice of the function $z(\theta)$. The flow remains linear outside of these two strips. See Fig. 4.

Because the phase space is a torus and (6) does not have fixed points for the parameters of interest, φ_t has a global section, e.g., $\Sigma_0 = \{\theta_2 = 0\}$, with a return map $T_0 : \Sigma_0 \to \Sigma_0$. The large-time dynamics of φ_t are therefore described by iterating T_0, a circle diffeomorphism. From standard theory, we know that circle maps have nonpositive Lyapunov exponents for almost all initial conditions with respect to Lebesgue measure. This in turn translates into $\lambda_{\max} = 0$ for the flow φ_t.

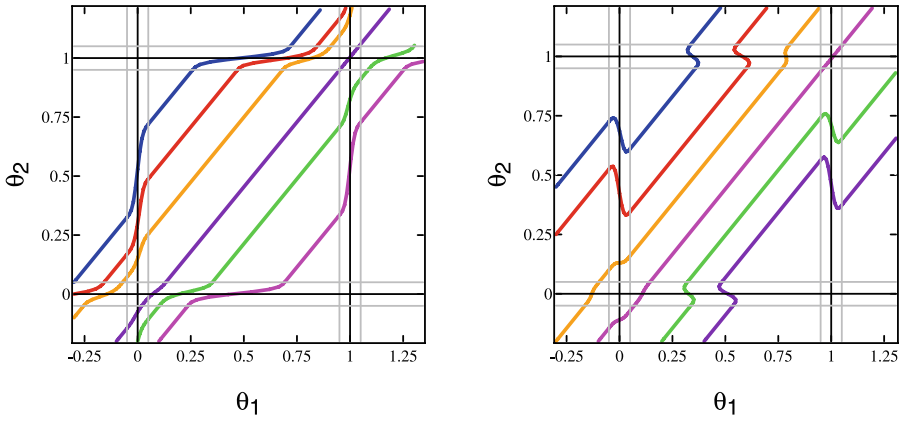

(a) $a_{\mathrm{ff}} = 1$, $a_{\mathrm{fb}} = 1.5$ (b) $a_{\mathrm{ff}} = -1$, $a_{\mathrm{fb}} = -0.5$

Fig. 4 Plots of a few orbits of (6) with $\varepsilon = 0$ showing the strips in which flowlines are distorted. In both sets, $\omega_1 = 1, \omega_2 = 1.1$. Note the directions of bending in relation to the signs of a_{ff} and a_{fb}

We consider next two special but instructive cases of (6), namely the "pure feedforward" case corresponding to $a_{\mathrm{fb}} = 0$ and the "pure feedback" case corresponding to $a_{\mathrm{ff}} = 0$. In these two special cases, the geometry of the system prohibits it from being unreliable:

Proposition 3.1 *For every* $\varepsilon > 0$,

(a) *the system* (6) *has an ergodic stationary measure* μ *with a density*
(b) (i) *if* $a_{\mathrm{fb}} = 0$, *then* $\lambda_{\max} \leq 0$
(ii) *if* $a_{\mathrm{ff}} = 0$, *then* $\lambda_{\max} = 0$.

Proposition 3.1(b) is a special case of a more general theorem on acyclic networks, which we state and prove in Sect. 4.2. Rather than give a formal proof here, we discuss informally the underlying geometric intuition. Proposition 3.1(a) is proved at the end of this subsection.

Consider the case $a_{\mathrm{fb}} = 0$. Notice that when $\varepsilon = 0$, the time-t map of the flow generated by (6) sends vertical circles (meaning sets of the form $\{\theta_1 = c\}$ where c is a constant) to vertical circles. As our stimulus acts purely in the horizontal direction and its magnitude is independent of θ_2, vertical circles continue to be preserved by the flow-maps $F_{s,t,\omega}$ when $\varepsilon > 0$. (One can also see this from the variational equations.) As is well known, $\lambda_{\max} > 0$ usually results from repeated stretching and folding of the phase space. Maps that preserve all vertical circles are far too rigid to allow this to happen. Thus we expect $\lambda_{\max} \leq 0$ when $a_{\mathrm{fb}} = 0$.

Moreover, when $a_{\mathrm{fb}} = 0$ it can be shown that two nearby trajectories that are initially separated in the θ_1 direction will converge exponentially fast (see Theorem 3 in Sect. 4.2.2). The value of $\lambda_{\max}$ therefore hinges on the growth rate of the distance between two trajectories initially separated in the θ_2 direction. However, even though this growth rate also involves compositions of random circle maps, the maps here are

not *i.i.d.*: The kicks received by the second oscillator are in randomly timed pulses that cannot be put into the form of the white noise term in (3). We know of no mathematical result that guarantees a strictly negative Lyapunov exponent in this setting, but believe it is unlikely that (3) will have a robust zero Lyapunov exponent unless $a_{\text{ff}} = 0$.

In the pure feedback case $a_{\text{ff}} = 0$, the second oscillator rotates freely without input from either external stimuli or the first oscillator. Thus, the system always has a zero Lyapunov exponent corresponding to the uniform motion in the direction of θ_2.

Proof of Proposition 3.1(a) To verify that (6) has a stationary measure with a density, first note that because the white noise stimulus $\varepsilon I(t)$ instantaneously spreads trajectories in the horizontal (θ_1) direction, an invariant measure must have a density in this direction. At the same time, the deterministic part of the flow carries this density forward through all of $\mathbb{T}^2$ since flowlines make approximately 45 degree angles with the horizontal axis. Therefore, the two-oscillator system has a 2-D invariant density whenever $\varepsilon > 0$. □

To summarize, *it is impossible for a two-oscillator system to exhibit an unreliable response without recurrent connections*.

3.1.2 Phase Locking in Zero-Input Dynamics

This subsection concerns the dynamics of the coupled oscillator system when $\varepsilon = 0$. Recall that the intrinsic frequencies of the two oscillators, ω_1 and ω_2, are assumed to be ≈ 1. Our main result is that if ω_1 and ω_2 differ by a little, then in regions of the $(a_{\text{ff}}, a_{\text{fb}})$-parameter plane in the form of a square centered at $(0, 0)$, *the two oscillators are 1:1 phase-locked for about half of the parameters*. Two examples are shown in Fig. 5, and more detail will be given as we provide a mathematical explanation for this phenomenon. (Phase locking of pairs of coupled phase oscillators is studied in e.g., Gerstner et al. (1996), Hansel et al. (1995), Ermentrout (1996, 1981), Chow (1998), Taylor and Holmes (1998). A primary difference is that we treat the problem on an open region of the $(a_{\text{ff}}, a_{\text{fb}})$-plane centered at $(0, 0)$.)

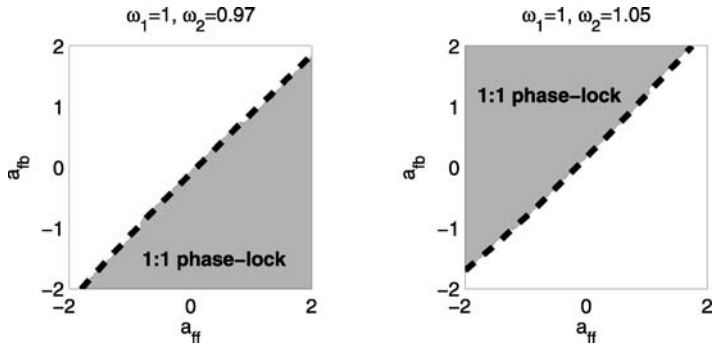

Fig. 5 The critical value a_{fb}^* as functions of a_{ff}. In both plots, the *dashed line* is the a_{fb}^* curve (see text), and the *shaded regions* are the parameters for which 1:1 phase-locking occurs

Let φ_t be the flow on $\mathbb{T}^2$ defined by (6) with $\varepsilon = 0$. To study φ_t, we consider its return map $T : \Sigma_b \to \Sigma_b$ where $\Sigma_b = \{\theta_2 = b\}$. Working with the section Σ_b (as opposed to, e.g., Σ_0) simplifies the analysis as we will see, and substantively the results are not section-dependent. Let $\rho(T)$ be the rotation number of T (Guckenheimer and Holmes 1983). Since $\omega_1 \approx \omega_2$, it is natural to define $\rho(T)$ in such a way that $\rho(T) \approx 1$ when $a_{\text{ff}} = a_{\text{fb}} = 0$, so that $\rho(T)$ may be interpreted as the average number of rotations made by the first oscillator for each rotation made by the second. It is well known that if $\rho(T)$ is irrational, then φ_t is equivalent to a quasi-periodic flow via a continuous change of coordinates, while $\rho(T) \in \mathbb{Q}$ corresponds to the presence of periodic orbits for φ_t. In particular, attracting fixed points of T correspond to attracting periodic orbits of φ_t that are 1:1 phase-locked, "1:1" here referring to the fact that oscillator 2 goes around exactly once for each rotation made by oscillator 1.

We begin with the following elementary facts:

Lemma 3.1

(a) *With ω_1, ω_2, and a_{ff} fixed and letting $T = T_{a_{\text{fb}}}$, the function $a_{\text{fb}} \mapsto T_{a_{\text{fb}}}(\theta_1)$ is strictly increasing for each θ_1; it follows that $\rho(T_{a_{\text{fb}}})$ is a nondecreasing function*[1] *of a_{fb}.*

(b) *With ω_1, a_{ff}, and a_{fb} fixed, $\omega_2 \mapsto T_{\omega_2}(\theta_1)$ is strictly decreasing for each θ_1, and $\rho(T_{\omega_2})$ is a non-increasing function of ω_2.*

Analogous results hold as we vary a_{ff} and ω_1 separately keeping the other three quantities fixed.

Proof This follows from the way each coupling term bends trajectories; see Sect. 3.1.1 and Fig. 4. We show (a); the rest are proved similarly. Consider two trajectories, both starting from the same point on Σ_b but with different a_{fb}. They will remain identical until their θ_2-coordinates reach $1 - b$, as a_{fb} does not affect this part of the flow. Now at each point in the horizontal strip $H = \{1 - b < \theta_2 < 1 + b\}$, the vector field corresponding to the larger a_{fb} has a larger horizontal component while the vertical components of the two vector fields are identical. It follows that the trajectory with the larger a_{fb} will be bent farther to the right as it crosses H. □

Our main result identifies regions in parameter space where T has a fixed point, corresponding to 1:1 phase-locking as discussed above. We begin with the following remarks and notation: (i) Observe that when $\omega_1 = \omega_2$ and $a_{\text{ff}} = a_{\text{fb}}$, we have $T(x) = x$ for all $x \in \Sigma_b$; this is a consequence of the symmetry of $z(\theta)$ about $\theta = \frac{1}{2}$. (ii) We introduce the notation $\Delta\omega = 1 - \frac{\omega_1}{\omega_2}$, so that when $a_{\text{ff}} = a_{\text{fb}} = 0$, $x - T(x) = \Delta\omega$ for all x, i.e., $\Delta\omega$ measures the distance moved by the return map T under the linear flow when the two oscillators are uncoupled. Here, we have used T to denote not only the section map but its lift to $\mathbb{R}$ with $T(1) = 1$ in a harmless abuse of notation.

We will state our results for a bounded range of parameters. For definiteness, we consider $a_{\text{ff}}, a_{\text{fb}} \in [-2, 2]$ and $0.9 \le \omega_i \le 1.1$. The bounds for a_{ff} and a_{fb} are quite arbitrary. Once chosen, however, they will be fixed; in particular, the constants in

[1] This is to allow for phase locking at rational values of $\rho(T_{a_{\text{fb}}})$.

our lemmas below will depend on these bounds. *It is implicitly assumed that all parameters considered in Theorem 2 are in this admissible range*. We do not view b as a parameter in the same way as $\omega_1, \omega_2, a_{\rm ff}$ or $a_{\rm fb}$, but instead of fixing it at $\frac{1}{20}$, we will take b as small as need be, and assume $|g| = \mathcal{O}(\frac{1}{b})$ in the rigorous results to follow.

Theorem 2 *The following hold for all admissible* $(\omega_1, \omega_2, a_{\rm ff}, a_{\rm fb})$ *and all* b *sufficiently small*:

(a) *If* $\omega_2 > \omega_1$, *then there exist* $a^*_{\rm fb} = a^*_{\rm fb}(\omega_1, \omega_2, a_{\rm ff}) > a_{\rm ff}$ *and* $\ell = \ell(\Delta\omega) > 0$ *such that* T *has a fixed point for* $a_{\rm fb} \in [a^*_{\rm fb}, a^*_{\rm fb} + \ell]$ *and no fixed point for* $a_{\rm fb} < a^*_{\rm fb}$.
(b) *If* $\omega_2 < \omega_1$, *then there exist* $a^*_{\rm fb} = a^*_{\rm fb}(\omega_1, \omega_2, a_{\rm ff}) < a_{\rm ff}$ *and* $\ell = \ell(\Delta\omega) > 0$ *such that* T *has a fixed point for* $a_{\rm fb} \in [a^*_{\rm fb} - \ell, a^*_{\rm fb}]$ *and no fixed point for* $a_{\rm fb} > a^*_{\rm fb}$.

Moreover, $|a^*_{\rm fb} - a_{\rm ff}| = \mathcal{O}(\Delta\omega)$; *and for each* $\Delta\omega \neq 0$, ℓ *increases as* b *decreases*.

A proof of Theorem 2 is given in the Appendix.

Our proof shows, in fact, that for as long as $\omega_2 \neq \omega_1$, the lengths of the phase-locked intervals, ℓ, can be made arbitrarily large by taking b small. On the other hand, if we fix b and shrink $\Delta\omega$, then these intervals will shrink. This is consistent with the phenomenon that the phase-locked region lies on opposite sides of the diagonal $a_{\rm fb} = a_{\rm ff}$ when we decrease ω_2 past ω_1, as shown in Fig. 5. Instead of tracking the numerical constants in the proofs, we have checked numerically that for $b = \frac{1}{20}$, the pictures in Fig. 5 are quite typical, meaning about 50% of the parameters are phase locked. Specifically, for $\Delta\omega$ up to about 10% and $|a_{\rm ff}| < 2$, we find $a^*_{\rm fb} - a_{\rm ff} \lesssim 0.2$, so that the $a^*_{\rm fb}$-curve describing the onset of phase-locking is still quite close to the diagonal $a_{\rm ff} = a_{\rm fb}$. Also, for $\Delta\omega$ as small as 0.5%, the phase locked intervals have length $\ell > 4$. These facts together imply that for parameters in the admissible range, the pictures are as shown in Fig. 5.

Figure 6 shows numerically-computed rotation numbers ρ and the rates of contraction to the corresponding limit cycle, i.e., the smaller Lyapunov exponent $\lambda_{\rm min}$ of the flow φ_t. Notice that as $a_{\rm fb}$ increases past $a^*_{\rm fb}$, $\lambda_{\rm min}$ decreases rapidly, so that the fixed point of T becomes very stable, a fact consistent with the large interval on which the system is 1:1 phase-locked. Furthermore, for the full range of $|a_{\rm ff}|, |a_{\rm fb}| \leq 2$ and

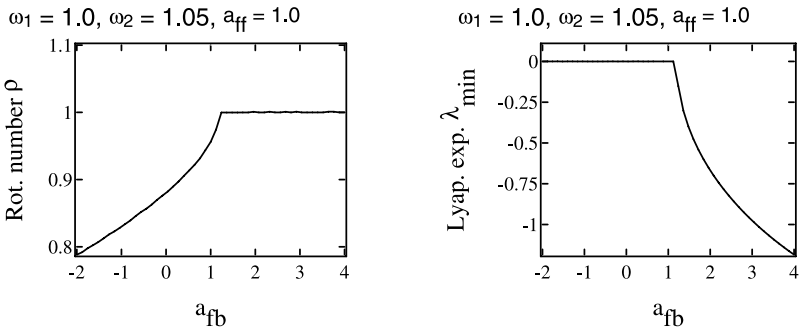

Fig. 6 The rotation number ρ and Lyapunov exponent $\lambda_{\rm min}$ as functions of $a_{\rm fb}$

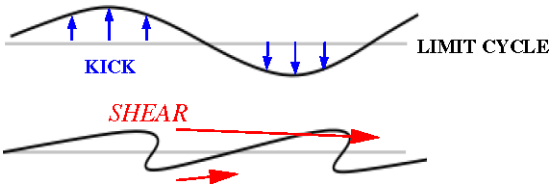

Fig. 7 The stretch-and-fold action of a kick followed by relaxation in the presence of shear

$0.9 \le \omega_i \le 1.1$, we find numerically that $0.53 < \rho(T) < 1.89$. Phase-locking corresponding to rational $\rho(T)$ with small denominators q (e.g., $q = 3, 4, 5$) is detected, but the intervals are very short and their lengths decrease rapidly with q. In other words, when the system is not 1:1 phase-locked—which occurs for about 50% of the parameters of interest—modulo fine details the system appears to be roughly quasi-periodic over not too large timescales. When the white-noise stimulus $\varepsilon I(t)$ is added, these fine details will matter little.

3.2 Responses to Weak Stimuli

Numerical evidence is presented in Sect. 3.2.2 (see Fig. 9) showing that unreliability can occur even when the stimulus amplitude ε is relatively small, and that its occurrence is closely connected with the onset of phase-locking in the zero-input system. A geometric explanation in terms of *shear-induced chaos* is proposed. We begin with a review of the general theory.

3.2.1 A Brief Review of Shear-Induced Chaos

A rough idea of what we mean by "shear-induced chaos" is depicted in Fig. 7: An external force is transiently applied to a limit cycle (horizontal line), causing some phase points to move up and some to move down. Suppose the speeds with which points go around the limit cycle are height dependent. If the velocity gradient, which we refer to as "*shear*", is steep enough, then the bumps created by the forcing are exaggerated as the system relaxes, possibly causing the periodic orbit to fold. Such folding has the potential to lead to the formation of strange attractors. If the attraction to the limit cycle is large relative to the velocity gradient or the perturbation size, however, the bumps in Fig. 7 will dissipate before any significant stretching occurs.

This subsection reviews a number of ideas surrounding the mechanism above. This mechanism is known to occur in many different dynamical settings. We have elected to introduce the ideas in the context of *discrete-time kicking of limit cycles* instead of the continuous-time forcing in (6) because the geometry of discrete-time kicks is more transparent, and many of the results have been shown numerically to carry over with relatively minor modifications. Extensions to relevant settings are discussed later on in this subsection. A part of this body of work is supported by rigorous analysis. Specifically, theorems on shear-induced chaos for periodic kicks of limit cycles are proved in Wang and Young (2001, 2002, 2003, 2009); it is from these articles that many of the ideas reviewed here have originated. Numerical studies extending the ideas in Wang and Young (2002, 2003) to other types of underlying dynamics and forcing are carried out in Lin and Young (2008). For readers who wish to see a more

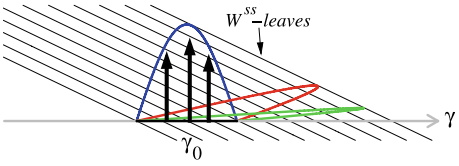

Fig. 8 Geometry of folding in relation to the W^{ss}-foliation. Shown are a segment $\gamma_0 \subset \gamma$, its image after one kick, and two of the subsequent images under the flow

in-depth (but not too technical) account of the material in this subsection, Lin and Young (2008) is a good place to start. In particular, Sect. 1 in Lin and Young (2008) contains a fairly detailed exposition of the rigorous work in Wang and Young (2001, 2002, 2003, 2009).

Discrete-Time Kicks of Limit Cycles

We consider a flow φ_t in any dimension, with a limit cycle γ. Let $T_0 < T_1 < T_2 < \cdots$ be a sequence of kick times, and let $\kappa_n, n = 0, 1, 2, \ldots$, be a sequence of kick maps (for the moment κ_n can be any transformation of the phase space). We consider a system kicked at time T_n by κ_n, with ample time to relax between kicks, i.e., $T_{n+1} - T_n$ should not be too small on average.

Central to the geometry of shear-induced chaos is the following dynamical structure of the unforced system: For each $x \in \gamma$, the *strong stable manifold* $W^{ss}(x)$ of φ_t through x is defined to be

$$W^{ss}(x) = \Big\{ y : \lim_{t\to\infty} d\big(\varphi_t(x), \varphi_t(y)\big) = 0 \Big\}.$$

These codimension 1 submanifolds are invariant under the flow, meaning φ_t carries $W^{ss}(x)$ to $W^{ss}(\varphi_t(x))$. In particular, if τ is the period of the limit cycle, then $\varphi_\tau(W^{ss}(x)) = W^{ss}(x)$ for each x. Together these manifolds partition the basin of attraction of γ into hypersurfaces, forming what is called the *strong stable foliation*.

Figure 8 shows a segment $\gamma_0 \subset \gamma$, its image $\gamma_0^+ = \kappa(\gamma)$ under a kick map κ, and two images of γ_0^+ under $\varphi_{n\tau}$ and $\varphi_{m\tau}$ for $n > m \in \mathbb{Z}^+$. If we consider integer multiples of τ, so that the flow-map carries each W^{ss}-leaf to itself, we may think of it as sliding points in γ_0^+ toward γ along W^{ss}-leaves. (For t that are not integer multiples of τ, the picture is similar but shifted along γ.) The stretching created in this combined kick-and-slide action depends both on the geometry of the W^{ss}-foliation and on the action of the kick. Figure 8 sheds light on the types of kicks that are likely to lead to folding: A forcing that drives points in a direction roughly parallel to the W^{ss}-leaves will not produce folding. Neither will kicks that essentially carry one W^{ss}-leaf to another, because such kicks necessarily preserve the ordering of the leaves. What causes the stretching and folding is the *variation* along γ in how far points $x \in \gamma$ are kicked as measured in the direction transverse to the W^{ss}-leaves; we think of this as the "effective deformation" of the limit cycle γ by the kick.

To develop a quantitative understanding of the factors conducive to the production of chaos, it is illuminating to consider the following linear shear model (Zaslavsky

1978; Wang and Young 2002):

$$\begin{aligned} \dot{\theta} &= 1 + \sigma y, \\ \dot{y} &= -\lambda y + A \sin(2\pi\theta) \cdot \sum_{n=0}^{\infty} \delta(t - nT). \end{aligned} \tag{7}$$

Here, $\theta \in \mathbb{S}^1$ and $y \in \mathbb{R}$ are phase variables, and σ, λ, A are constants.[2] We assume for definiteness that σ and λ are > 0, so that when $A = 0$, the system has a limit cycle γ at $\{y = 0\}$. Its W^{ss}-leaves are easily computed to be straight lines with slope $= -\frac{\lambda}{\sigma}$. When $A \neq 0$, the system is kicked periodically with period T. For this system, it has been shown that the *shear ratio*

$$\frac{\sigma}{\lambda} \cdot A \equiv \frac{\text{shear}}{\text{damping}} \cdot \text{deformation} = \frac{1}{|\text{slope}(W^{ss})|} \cdot \text{deformation} \tag{8}$$

is an excellent predictor of the dynamics of the system. Roughly speaking, if $\lambda_{\max}$ denotes the largest observed Lyapunov exponent, then

(a) if the shear ratio is sufficiently large, $\lambda_{\max}$ is likely to be > 0
(b) if the shear ratio is very small, then $\lambda_{\max}$ is slightly < 0 or equal to 0
(c) as the shear ratio increases from small to large, $\lambda_{\max}$ first becomes negative, then becomes quite irregular (taking on both positive and negative values), and is eventually dominated by positive values.

To get an idea of why this should be the case, consider the composite kick-and-slide action in Fig. 8 in the context of (7). The time-T map of (7) is easily computed to be

$$\begin{aligned} \theta_T &= \theta + T + \frac{\sigma}{\lambda} \cdot \big[A \sin(2\pi\theta) + y\big] \cdot \big(1 - e^{-\lambda T}\big) \pmod 1, \\ y_T &= e^{-\lambda T} \cdot \big(y + A \sin(2\pi\theta)\big). \end{aligned} \tag{9}$$

When the contraction in y is sufficiently strong, the first component of this map gives a good indication of what happens in the full 2-D system. As an approximation, define $f_T(\theta) = \theta_T$ and view f_T as a map of γ to itself. When the shear ratio is large and $(1 - e^{-\lambda T})$ is not too small, $|f_T'|$ is quite large over much of γ, and the associated expansion has the potential to create the positive $\lambda_{\max}$ mentioned in (a). At the opposite extreme, when the shear ratio is very small, f_T is a perturbation of the identity; this is the scenario in (b). Interestingly, it is for intermediate shear ratios that f_T tends to have sinks, resulting in $\lambda_{\max} < 0$ for the 2-D system. The 1-D map $f = \lim_{T\to\infty} f_T$ is known variously as the *phase resetting curve* or the *singular limit*. It is used heavily in Wang and Young (2001, 2002, 2003, 2009) to produce rigorous results for large T.

We now return to the more general setting of φ_t with discrete kicks κ_n, and try to interpret the results above as best we can. To make a meaningful comparison with

[2] In this subsection, λ denotes the damping constant in (7).

the linear shear flow, we propose to first put our unforced system in "canonical coordinates," i.e., to reparametrize the periodic orbit γ so it has unit speed, and to make the kick directions perpendicular to γ—assuming the kicks have well defined directions. In these new coordinates, sizes of vertical deformations make sense, as do the idea of damping and shear, even though these quantities and the angles between W^{ss}-manifolds and γ all vary along γ. Time intervals between kicks may vary as well. The general geometric picture is thus a distorted version of the linear shear flow. We do not believe there is a simple formula to take the place of the shear ratio in this general setting; replacing the quantities σ, λ and A by their averages is not quite the right thing to do. We emphasize, however, that while system details affect the numerical values of $\lambda_{\max}$ and the amount of shear needed to produce $\lambda_{\max} > 0$, the fact that *the overall trends as described in* (a)–(c) *above are valid* has been repeatedly demonstrated in simulation; see, e.g., Lin and Young (2008).

Generalization to Stochastic Forcing. We now replace the discrete-time kicks in the discussion above by a directed (degenerate) continuous stochastic forcing, i.e., by a term of the form $V(x)\,dW_t$ where V is a vector field and dW_t is white noise. By Trotter's product formula, the dynamics of the resulting stochastic flow can be approximated by a sequence of composite maps of the form $\varphi_{\Delta t} \circ \kappa_{\Delta t,\omega}$ where Δt is a small time step, $\kappa_{\Delta t,\omega}$ are kick maps (of random sizes) in the direction of V, and $\varphi_{\Delta t}$ is the unforced flow. Most of the time, the maps $\kappa_{\Delta t,\omega}$ have negligible effects. This is especially the case if the size of V is not too large and damping is present. Once in a while, however, a large deviation occurs, producing an effect similar to that of the discrete-time kicks at times T_n described above. Cast in this light, we expect that the ideas above will continue to apply—albeit without the factors in the shear ratio being precisely defined.

We mention two of the differences between stochastic forcing and periodic, discrete kicks. Not surprisingly, stochastic forcing gives simpler dependence on parameters: $\lambda_{\max}$ varies smoothly, irregularities of the type in (c) above having been "averaged out". Overall trends such as those in (a)–(c) tend to be unambiguous and more easily detected than for deterministic kicks. Second, unlike periodic kicks, very small forcing amplitudes can elicit chaotic behavior without $\frac{\sigma}{\lambda}$ being very large; this is attributed to the effects of large deviations.

Generalization to Quasi-Periodic Flows. We have chosen to first introduce the ideas above in the context of limit cycles where the relevant geometric objects or quantities (such as σ, λ and W^{ss}) are more easily extracted. These ideas apply in fact to flows on a torus that are roughly "quasi-periodic"—meaning that orbits may or may not be periodic but if they are, the periods are large—provided the forcing, stochastic or discrete, has a well-defined direction as discussed earlier. The main difference between the quasi-periodic setting and that of a limit cycle is that W^{ss}-leaves are generally not defined. A crucial observation made in Lin and Young (2008) is that since folding occurs in finite time, what is relevant to the geometry of folding is not the usual W^{ss}-foliation (which takes into consideration the dynamics as $t \to \infty$) but *finite-time* strong stable foliations. Roughly speaking, a time-t W^{ss}-leaf is a curve segment (or submanifold) that contracts in the first t units of time. For the ideas above

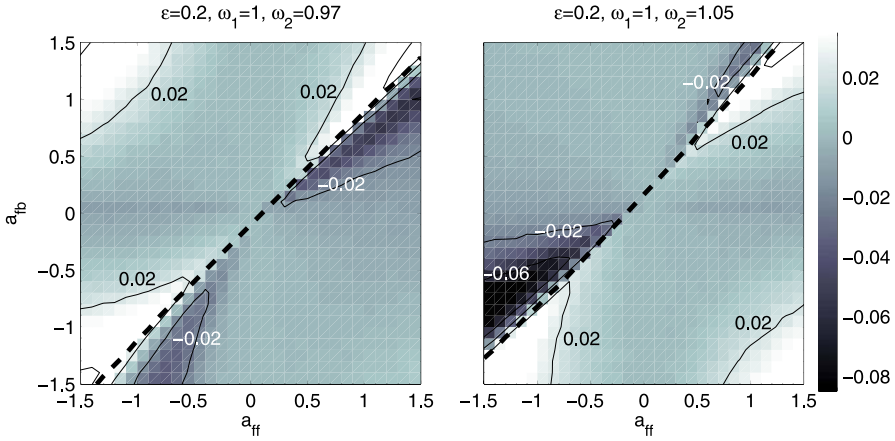

(a) $\varepsilon = 0.2$

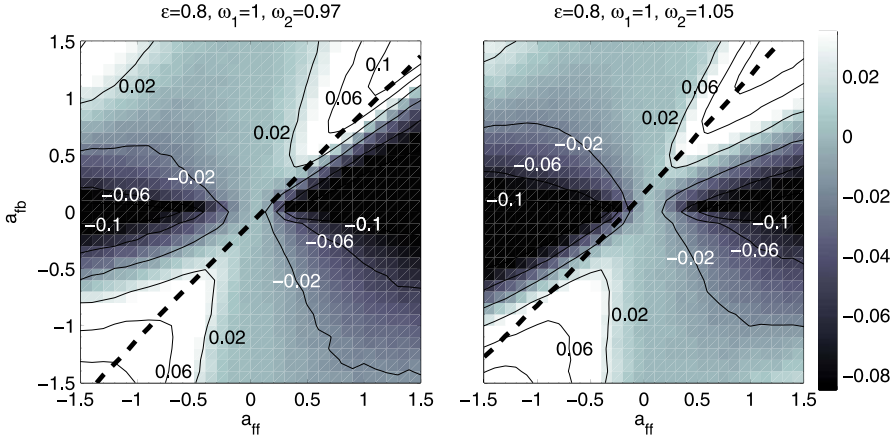

(b) $\varepsilon = 0.8$

Fig. 9 Maximum Lyapunov exponent $\lambda_{\max}$ versus coupling strengths in the two-cell network. In all plots, we use $\omega_1 = 1$. The a_{fb}^*-curve from Fig. 5 is overlaid. All computed $\lambda_{\max}$ shown here have standard errors of ≤ 0.002 as estimated by the method of batched means. By the Central Limit Theorem, this means the actual $\lambda_{\max}$ should lie within $\approx 2.5 \times 0.002 = 0.005$ of the computed value with $\gtrsim 99\%$ probability. *Remark on plots:* We have chosen the dynamic range in shading the figures to allow meaningful comparison of figures; a side effect is that some contour lines may not be visible. We always indicate the actual range of values through explicit labels

to apply, we must verify that time-t W^{ss}-manifolds exist, have the characteristics of a large shear ratio, and that t is large enough for the folding to actually occur. If these conditions are met, then one can expect shear-induced chaos to be present for the same reasons as before.

3.2.2 *Phase-Locking, Shear and Reliability*

We now return to reliability questions for (6). In this subsection, numerical data on $\lambda_{\max}$ as functions of $a_{\rm ff}, a_{\rm fb}$ and ε are discussed. We limit ourselves to relatively small ε, and to phenomena related to the onset of phase-locking, which we have shown in Sect. 3.1.2 to occur at a curve in $(a_{\rm ff}, a_{\rm fb})$-space that runs roughly along the diagonal. We will refer to this curve as the $a^*_{\rm fb}$-curve. Figure 9 shows $\lambda_{\max}$ as a function of $a_{\rm ff}$ and $a_{\rm fb}$ for several choices of parameters, with $a^*_{\rm fb}$-curves from Fig. 5 overlaid for ease of reference. In the top two panels, where ε is very small, evidence of events connected with the onset of phase-locking is undeniable: definitively reliable ($\lambda_{\max} < 0$) and definitively unreliable ($\lambda_{\max} > 0$) regions are both present. Continuing to focus on neighborhoods of the $a^*_{\rm fb}$-curves, we notice by comparing the top and bottom panels that for each $(a_{\rm ff}, a_{\rm fb})$, the tendency is to shift in the direction of unreliability as ε is increased. We will argue in the paragraphs to follow that these observations are entirely consistent with predictions from Sect. 3.2.1.

Shearing Mechanisms

For concreteness, we consider the case $\omega_2 > \omega_1$, and consider first parameters at which the unforced system has a limit cycle, i.e., for each $a_{\rm ff} \in [-1.5, 1.5]$, we consider values of $a_{\rm fb}$ that are $> a^*_{\rm fb}$ and not too far from $a^*_{\rm fb}$. From Sect. 3.2.1, we learn that to determine the propensity of the system for shear-induced chaos, we need information on (i) the geometry of the limit cycle, (ii) the orientation of its W^{ss}-manifolds in relation to the cycle, and (iii) the effective deformation due to the forcing.

The answer to (i) is simple: As with all other trajectories, the limit cycle is linear with slope $\gtrsim 1$ outside of the two corridors $|\theta_1| < b$ and $|\theta_2| < b$, where it is bent; see Fig. 4. As for (ii) and (iii), we already know what happens in two special cases, namely when $a_{\rm ff} = 0$ or $a_{\rm fb} = 0$. As discussed in Sect. 3.1, when $a_{\rm fb} = 0$, vertical circles are invariant under φ_t. Since W^{ss}-leaves are the only manifolds that are invariant, that means the W^{ss}-manifolds are vertical. We noted also that the forcing preserves these manifolds. In the language of Sect. 3.2.1, this means the forcing creates *no variation* transversal to W^{ss}-leaves: The ordering of points in this direction is preserved under the forcing. Hence, shear-induced chaos is not possible here, and not likely for nearby parameters. A similar argument (which we leave to the reader) applies to the case $a_{\rm ff} = 0$. From here on, we assume $a_{\rm ff}, a_{\rm fb}$ are both definitively nonzero.

We now turn to a treatment of (ii) when $a_{\rm ff}, a_{\rm fb}$ are both definitively nonzero, and claim that W^{ss}-leaves generally have a roughly diagonal orientation, i.e., they point in a roughly southwest-northeast (SW–NE) direction. To see this, first recall that the orientation of the W^{ss} leaves can be found in the following way: fix a point p on the limit cycle γ and any nonzero tangent vector v at p (see Fig. 10). We then flow backwards in time, letting $p_{-t} = \varphi_{-t}(p)$ and $v_{-t} = D\varphi_{-t}(p)v$. The strong stable direction at p_{-t} is the limiting direction of $v_{-t-n\tau}$ (τ = period of γ) as $n \to \infty$. (Exact orientations of W^{ss}-leaves depend on $a_{\rm ff}, a_{\rm fb}$ and are easily computed numerically.)

We now demonstrate how one can deduce the general orientation of the W^{ss}-leaves in the two-oscillator system from the signs of its couplings, following the simple

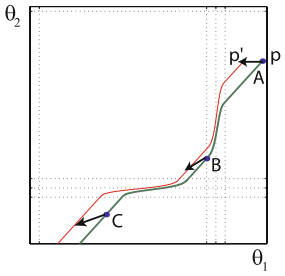

Fig. 10 Rotation of vectors in *backwards* time. Here, $\omega_1 = 1$, $\omega_2 = 1.1$, $a_{\text{ff}} = 1$, $a_{\text{fb}} = 1.5$

strategy outlined above. For definiteness, we consider $a_{\text{ff}} > 0$, so that a_{fb}^* is also positive and slightly larger than a_{ff}. Here, a typical situation is that if we identify the phase space with the square $[0, 1] \times [0, 1]$, then the limit cycle crosses the right edge $\{1\} \times [0, 1]$ in the bottom half, and the bottom edge $[0, 1] \times \{0\}$ in the right half. See Fig. 10, which shows a lift of the limit cycle to $\mathbb{R}^2$. Let A, B and C be as shown, and consider a point p at A. Flowing backwards, suppose it takes time t_B to reach point B, and time t_C to reach point C. We discuss how v_{-t} changes as we go from A to C, starting with the vector v_0 that points from p to a nearby point p'. The rest of the time the flow is linear and v_{-t} is unchanged.

From A to B: Compare the backward orbits of p and p', where $p' = p + kv$ and $k > 0$ is thought of as infinitesimally small. As these orbits reach the vertical strip $\{g > 0\}$, both are bent downwards due to $a_{\text{ff}} > 0$. However, the orbit of p' is bent more because the function $z(\theta)$ peaks at $\theta = 1/2$ (see Fig. 10). Thus, $v_{-t_B} = v + (0, -\delta_1)$ for some $\delta_1 > 0$, as seen in the southwest-pointing arrow by "B" in Fig. 10.

From B to C: Continuing to flow backwards, we see by an analogous argument that as the two orbits cross the horizontal strip $\{g > 0\}$, both are bent to the left, and the orbit of p' is bent more. Therefore, $v_{-t_C} = v_{-t_B} + (-\delta_2, 0)$ for some $\delta_2 > 0$.

Combining these two steps, we see that each time the orbit of p_{-t} goes from A to C, a vector of the form $(-\delta_2, -\delta_1)$ is added to v_{-t}. We conclude that as $t \to \infty$, the direction of v_{-t} is asymptotically SW–NE. Moreover, it is a little more W than S compared to the limit cycle because v_{-t} must remain on the same side of the cycle.

Numerical computations of strong stable directions are shown in Fig. 11. Specifically, the short line segments that intersect the limit cycles are linear approximations to the strong stable manifolds with the corresponding base points. The plot in Fig. 11(a) is an example of the situation just discussed, demonstrating the SW–NE orientation of these manifolds. The plot in (b) is for $a_{\text{ff}}, a_{\text{fb}} < 0$, for which a similar analysis can be carried out. Notice how small the angles are between the limit cycle and its W^{ss}-leaves; this is true for all the parameter sets we have examined where a_{fb} is close to a_{fb}^*. Recall from Sect. 3.2.1 that it is, in fact, the angles in *canonical coordinates* that count. Since the limit cycle is roughly diagonal and the forcing is horizontal, putting the system in canonical coordinates will not change these angles by more than a moderate factor (except in one small region in picture (a)); i.e., the angles will remain small.

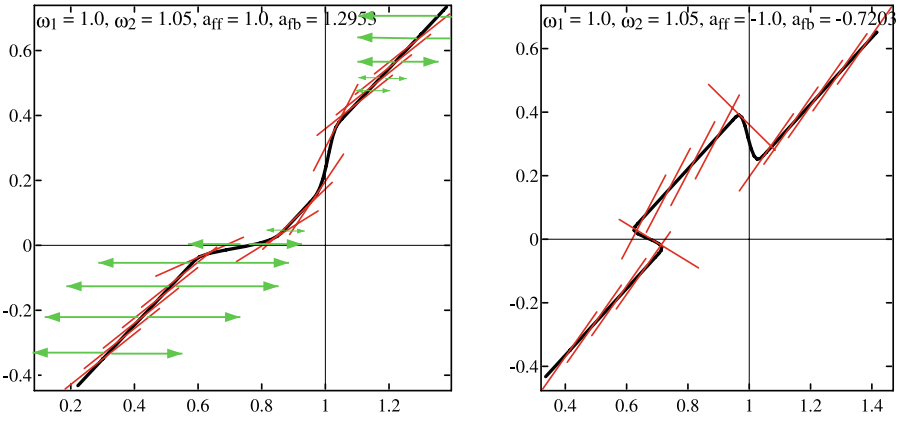

(a) $a_{\mathrm{ff}}, a_{\mathrm{fb}} > 0$ (b) $a_{\mathrm{ff}}, a_{\mathrm{fb}} < 0$

Fig. 11 Strong stable directions along limit cycles. In (**a**), $a_{\mathrm{fb}} = a^*_{\mathrm{fb}} + 0.1$; in (**b**), $a_{\mathrm{fb}} = a^*_{\mathrm{fb}} + 0.2$. Additionally in (**a**), *horizontal arrows* indicate the direction and (via their length) the magnitude of the forcing at various points along the cycle; the variable impact is due to the function $z(\theta_1)$

Finally, we come to (iii), the deformation due to the forcing. Given that the forcing is in the horizontal direction and its amplitude depends on θ_1 (it is negligible when $\theta_1 \approx 0$ and has maximal effect when $\theta_1 \approx \frac{1}{2}$), it causes "bump-like" perturbations transversal to the W^{ss}-manifolds (which are roughly SW–NE) with a geometry similar to that in Fig. 8; see Fig. 11(a).

This completes our discussion of the limit cycle case. We move now to the other side of the a^*_{fb}-curve, where the system is, for practical purposes, quasi-periodic (but not far from periodic). As discussed in the last part of Sect. 3.2.1, the ideas of shear-induced chaos continue to apply, with the role of W^{ss}-leaves now played by finite-time stable manifolds. Since these manifolds change slowly with a_{ff} and a_{fb}, it can be expected—and we have checked—that they continue to make small angles with flowlines. Likewise, the forcing continues to deform flowlines by variable amounts as measured in distances transversal to finite-time stable manifolds.

We conclude that when $a_{\mathrm{ff}}, a_{\mathrm{fb}}$ are both definitively nonzero, the geometry is favorable for shear-induced stretching and folding. Exactly how large a forcing amplitude is needed to produce a positive $\lambda_{\max}$ depends on system details. Such information cannot be deduced from the ideas reviewed in Sect. 3.2.1 alone.

Reliability–unreliability interpretations

We now examine more closely Fig. 9, and attempt to explain the reliability properties of those systems whose couplings lie in a neighborhood of the a^*_{fb}-curve. The discussion below applies to $|a_{\mathrm{ff}}| >$ about 0.3. We have observed earlier that for a_{ff} or a_{fb} too close to 0, phase-space geometry prohibits unreliability.

Consider first Fig. 9(a), where the stimulus amplitude $\varepsilon = 0.2$ is very weak. Regions showing positive and negative Lyapunov exponents are clearly visible in both panels. Which side of the a^*_{fb}-curve corresponds to the phase-locked region is also

readily recognizable to the trained eye (lower triangular region in the picture on the left and upper triangular region on the right; see Fig. 5).

We first explore the phase-locked side of the a_{fb}^*-curve. Moving away from this curve, $\lambda_{\max}$ first becomes definitively negative. This is consistent with the increased damping noted in Sect. 3.1.2; see Fig. 6 (right). As we move farther away from the a_{fb}^*-curve still, $\lambda_{\max}$ increases and remains for a large region close to 0. Intuitively, this is due to the fact that for these parameters the limit cycle is very robust. The damping is so strong that the forcing cannot (usually) deform the limit cycle appreciably before it returns near its original position. That is to say, the perturbations are negligible. With regard to the theory in Sect. 3.2.1, assuming σ remains roughly constant, that $\lambda_{\max}$ should increase from negative to 0 as we continue to move away from a_{fb}^* is consistent with increased damping; see (b) and (c) in the interpretation of the shear ratio.

Moving now to the other side of the a_{fb}^*-curve, which is essentially quasi-periodic, regions of unreliability are clearly visible. These regions, in fact, begin slightly on the phase-locked side of the curve, where a weakly attractive limit cycle is present. We have presented evidence to support our contention that this is due to shear-induced chaos, or folding of the phase space. The fact that $\lambda_{\max}$ is more positive before the limit cycle is born than after can be attributed to the weaker-to-nonexistent damping before its birth. Thus, the general progression of $\lambda_{\max}$ from roughly 0 to definitively negative to positive as we cross the a_{fb}^*-curve from the phase-locked side to the quasi-periodic side is altogether consistent with scenarios (a)–(c) in Sect. 3.2.1 together with the observations in the paragraph on stochastic forcing.

We point out that the unreliability seen in these panels is fairly delicate, perhaps even unexpected *a priori* for the smaller values of a_{ff} and a_{fb}, such as $0.3 < |a_{\text{ff}}|$, $|a_{\text{fb}}| < 0.8$: The bending of the flow lines is rather mild at these smaller coupling parameters. Moreover, we know that no chaotic behavior is possible at $\varepsilon = 0$, and the stimulus amplitude of $\varepsilon = 0.2$ in the top panels is quite small. Recall, however, that the stimulus is a fluctuating white noise, and ε gives only an indication of its *average* amplitude. As noted in Sect. 3.2.1, we believe the unreliability seen is brought about by an interaction between the large fluctuations in the stimulus presented and the shearing in the underlying dynamics.

In Fig. 9(b), the stimulus amplitude is increased to $\varepsilon = 0.8$. A close examination of the plots shows that near to and on both sides of the a_{fb}^*-curve, $\lambda_{\max}$ has increased for each parameter pair $(a_{\text{ff}}, a_{\text{fb}})$, and that the reliable regions are pushed deeper into the phase-locked side compared to the top panels. This is consistent with the shear ratio increasing with forcing amplitude as predicted in Sect. 3.2.1.

This completes our discussion in relation to Fig. 9.

To complement the theoretical description of the geometry of folding given in Sect. 3.2.1, we believe it is instructive to see an actual instance of how such a fold is developed when the system (6) is subjected to an arbitrary realization of white noise. A few snapshots of the time evolution of the limit cycle under such a forcing is shown in Fig. 12. Notice that at the beginning, the combined action of the coupling and forcing causes the curve to wriggle left and right in an uncertain manner, but once a definitive kink is developed (such as at $t = 3$), it is stretched by the shear as predicted.

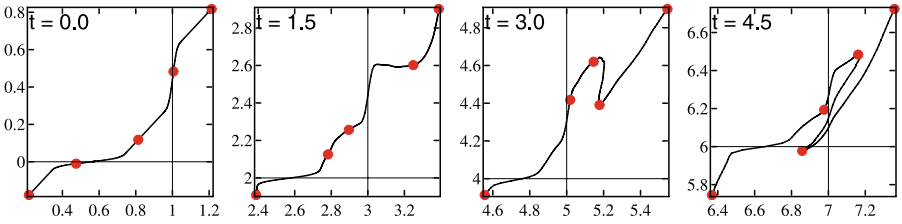

Fig. 12 Folding action caused by white noise forcing and shear near the limit cycle (with $a_{\mathrm{fb}} > a^*_{\mathrm{fb}}$). At $t = 0$, the *curve* shown is the lift of the limit cycle γ to $\mathbb{R}^2$. The remaining panels show lifts of the images $F_{0,t,\omega}(\gamma)$ at increasing times. The parameters are $\omega_1 = 1$, $\omega_2 = 1.05$, $a_{\mathrm{ff}} = 1$, $a_{\mathrm{fb}} = 1.2$, and $\varepsilon = 0.8$. Note that it is not difficult to find such a fold in simulations: very roughly, 1 out of 4 realizations of forcing gives such a sequence for $t \in [0, 5]$

Fig. 13 $\lambda_{\max}$ as function of a_{fb} and ε, with $\omega_1 = 1$, $\omega_2 = 1.1$, and $a_{\mathrm{ff}} = 1$

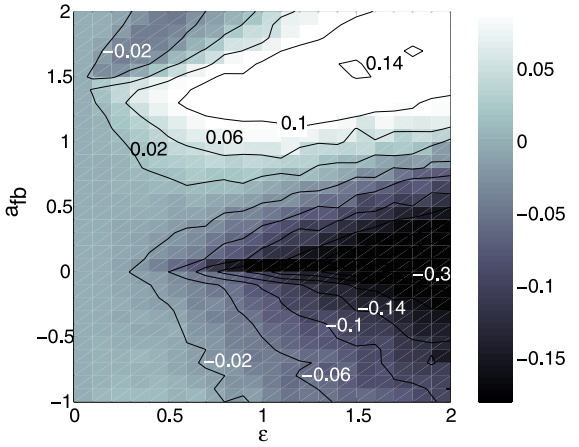

3.3 Dependence of Lyapunov Exponent on Parameters

This section concerns the dependence of $\lambda_{\max}$ on a_{ff}, a_{fb} and ε as these parameters vary over a broad range. Our aim is to identify the salient features in the reliability profile and to attempt to explain the phenomena observed. Our observations are based on plots of the type in Figs. 9 and 13. Some of the explanations we venture below are partial and/or speculative; they will be so indicated.

1. *Triple point:* This phenomenon is the main topic of discussion in Sect. 3.2.2. Figure 13 shows a different view of the parameter dependence of $\lambda_{\max}$. At about $a_{\mathrm{fb}} = 1.4$, which is near a^*_{fb} for the parameters used, both positive and negative Lyapunov exponents are clearly visible for very small ε in a manner that is consistent with Fig. 9 (even though the ω's differ slightly from that figure). We note again that this region, which we refer to as the "triple point," is an area of extreme sensitivity for the system, in the sense that the system may respond in a definitively reliable or definitively unreliable way to stimuli of very small amplitudes, with the reliability of its response depending sensitively on coupling parameters.
2. *Unreliability due to larger couplings:* Figure 9 and other plots not shown point to the occurrence of unreliability when $|a_{\mathrm{ff}}|$ and $|a_{\mathrm{fb}}|$ are both relatively large. We

are referring here specifically to the "off-diagonal" regions of unreliability (far from the $a_{\rm fb}^*$-curve) in Fig. 9. This phenomenon may be partly responsible for the unreliability seen for the larger values of $a_{\rm ff}$ and $a_{\rm fb}$ on the diagonal as well; it is impossible to separate the effects of the different mechanisms.

We do not have an explanation for why one should expect $\lambda_{\max} > 0$ for larger $|a_{\rm ff}|$ and $|a_{\rm fb}|$ aside from the obvious, namely that tangent vectors are more strongly rotated as they cross the strips $\{|\theta_i| < b\}$, making it potentially easier for folding to occur. But folding does *not* occur at $\varepsilon = 0$ in spite of this larger bending. We believe the difference between the two situations is due to the following noise-assisted mechanism: For $p \in \mathbb{T}^2$ and a tangent vector u at p, let us say u is *positively oriented* with respect to flowlines if starting from the direction of the flow and rotating counter-clockwise, one reaches u before reaching the direction of the backward vector field. Without external forcing, if u is positively oriented, $D\varphi_t(p) \cdot u$ will remain positively oriented for all t, because these vectors cannot cross flowlines. Now, in order for folding to occur, as in the formation of a horseshoe (Guckenheimer and Holmes 1983), the flow-maps must reverse the orientations of *some* tangent vectors. Even though larger values of $|a_{\rm ff}|$ and $|a_{\rm fb}|$ mean that tangent vectors are more strongly rotated, a complete reversal in direction cannot be accomplished without crossing flowlines. A small amount of noise makes this crossing feasible, opening the door (suddenly) to positive exponents.

3. *The effects of increasing* ε (up to around $\varepsilon = 2$):
   (a) *Unreliable regions grow larger, and* $\lambda_{\max}$ *increases:* A natural explanation here is that the stronger the stimulus, the greater its capacity to deform and fold the phase space—provided such folding is permitted by the underlying geometry. Because of the form of our stimulus, however, too large an amplitude simply pushes all phase points toward $\{\theta_1 = 0\}$. This will not lead to $\lambda_{\max} > 0$, a fact supported by numerics (not shown).

   (b) *Reliable regions grow larger, and the responses become more reliable:* As ε increases, the reliable region includes all parameters $(a_{\rm ff}, a_{\rm fb})$ in a large wedge containing the $a_{\rm fb} = 0$ axis. Moreover, in this region $\lambda_{\max}$ becomes significantly more negative as ε increases. We propose the following heuristic explanation: In the case of a single oscillator, if we increase ε, $\lambda_{\max}$ becomes more negative. This is because larger distortions of phase space geometry lead to more uneven derivatives of the flow-maps $F_{s,t;\omega}$, which in turn leads to a larger gap in Jensen's Inequality (see the discussion before Proposition 2.1). For two oscillators coupled as in (6), increasing ε has a similar stabilizing effect on oscillator 1. Feedback kicks from oscillator 2 may destabilize it, as happens for certain parameters near the $a_{\rm fb}^*$-curve. However, if $a_{\rm ff}$ is large enough and enhanced by a large ε, it appears that the stabilizing effects will prevail.

4 Part III Larger Networks

We now return to the general model defined in Sect. 2.1.1, where the size of the network, N, is arbitrary and can be $\gg 1$. Our main analytical results in Part III are on

network decomposition and Lyapunov exponents; they are presented in Sect. 4.2. This decomposition is assumed in subsequent sections. Section 4.3 discusses the propagation of unreliability to sites downstream, and Sect. 4.4 discusses a reassembly question, namely how to deduce the reliability of a system from knowledge of its parts. We begin Part III, however, by acknowledging a new issue, one that could not have happened in Part II as it requires at least 3 oscillators, namely intrinsic network chaos.

The mathematical analysis in Sect. 4.2 is rigorous, as are the facts reviewed in Sect. 4.3.1. The rest of the investigation is primarily numerical, supported whenever possible by facts from dynamical systems theory. Findings that have direct bearing on the broader mathematical questions are highlighted as Observations.

4.1 Intrinsic Network Chaos

4.1.1 Chaotic Behavior in Undriven Networks

Our criterion for unreliability, $\lambda_{\max} > 0$, is generally equated with *chaos*. In the 2-oscillator network studied in Part II, this chaotic behavior is triggered by the input: Without this forcing term, the system cannot be chaotic as it is a flow on a two-dimensional surface. The situation is quite different for larger networks.

Observation 1 *Networks of* 3 *or more pulse-coupled phase oscillators can be chaotic in the absence of any external input*.

We explain what we mean by "chaotic" for systems without external input. It is well known that systems of 3 or more coupled oscillators can have homoclinic orbits or horseshoes (see, e.g., Popovych et al. 2005; Nunes and Pereira 1985, and also Guckenheimer and Holmes 1983; Ermentrout 1991 for general references). These orbits by themselves do not necessarily constitute a seriously chaotic environment, however, for they are not always easy to detect, and the chaotic behavior seen can be transient (e.g., when horseshoes coexist with sinks). We claim that networks with 3 or more oscillators can support a stronger form of chaos, a form of chaos that is sustained in time and observable on large parts of the phase space. Mathematically, it is characterized by having positive Lyapunov exponents on large, positive Lebesgue measure sets in the phase space. Under quite general conditions, this is implied by the existence of SRB measures or physical measures (Eckmann and Ruelle 1985; Young 1995, 2002). In the discussion below, we settle for having positive Lyapunov exponents for all initial conditions tested as our operational definition of *intrinsic network chaos*.

One of the simplest configurations that supports intrinsic network chaos is

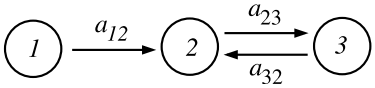

(10)

The setup above is similar to that studied in Part II, except that the (2,3)-subsystem is kicked *periodically* by pulses from oscillator 1 rather than driven by an external stimulus in the form of white noise. Ignoring the presence of oscillator 1 for the

moment, we saw in Sect. 3.2.2 that when a_{23} and a_{32} are roughly comparable and away from 0, the phase space geometry of oscillators 2 and 3 is favorable for *shear-induced chaos*, a geometric mechanism for producing chaotic behavior. Simulations show definitively positive exponents for these values of a_{23} and a_{32} for certain ranges of a_{12}. Specifically, with $(\omega_1, \omega_2, \omega_3) = (0.93, 1, 1.1)$, and $(a_{23}, a_{32}) = (1, 1.45)$, we find that $\lambda_{\max}$ increases from about 0.025 to 0.12 for $a_{12} \in (1, 1.5)$.

While there are certainly large classes of networks that do not exhibit intrinsic network chaos (see, e.g., Sect. 4.2), from the example above one would expect that such chaos is commonplace among larger networks. Moreover, for intrinsically chaotic networks, one would expect $\lambda_{\max}$ to remain positive when relatively weak stimuli are added to various nodes in the network. This raises the following issue of interpretation: In an intrinsically chaotic network that also displays $\lambda_{\max} > 0$ in the presence of inputs, it is unclear how to attribute the source of unreliability, since the effects of inputs and intrinsic chaos are largely inseparable. We adopt here the view that whatever the cause, such a network is unreliable.

4.1.2 Suppression of Network Chaos by Inputs: a Case Study

An interesting question is whether or not networks with intrinsic chaos of the form described above, i.e., with $\lambda_{\max} > 0$ for large sets of initial conditions in the absence of inputs, necessarily respond unreliably to external stimuli. We find that the answer is "no": We have come across a number of instances where weakly chaotic networks are stabilized by sufficiently strong inputs. An example is in the following 3-cell ring:

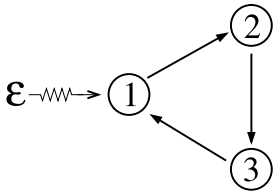

(11)

We find small regions of parameter space near $a_{12} = 1, a_{23} = 2, a_{31} = 0.6$ with intrinsic network chaos, as evidenced by positive values of $\lambda_{\max}$: see Fig. 14(a). As the reader will notice, $\lambda_{\max}$ oscillates quite wildly in this region, with $\lambda_{\max} > 0$ and $\lambda_{\max} = 0$ occurring at parameter values in close proximity.

We now consider the effects of adding a stimulus, that is, taking $\varepsilon > 0$ for system (11). Informally, one may think of the stimulus as sampling nearby parameters, averaging (in a loose sense) the different tendencies, with chaos suppression made possible by the presence of the mixed behavior. A sample of our numerical results is shown in Fig. 14(b). At $\varepsilon = 0$, $\lambda_{\max} > 0$, indicating network chaos. As expected, $\lambda_{\max}$ remains positive for small ε. But, as ε increases, $\lambda_{\max}$ steadily decreases and eventually becomes mildly negative, demonstrating that sufficiently strong forcing can suppress intrinsic network chaos and make a network more reliable. A possible mechanistic explanation is that when ε is large enough, oscillator 1 is affected more by the stimulus than by oscillator 3 due to the relatively weak coupling a_{31}. As in the case of a single oscillator, the stimulus has a stabilizing effect on oscillator 1, making it reliable.

Our understanding from the discussion above may be summarized as follows:

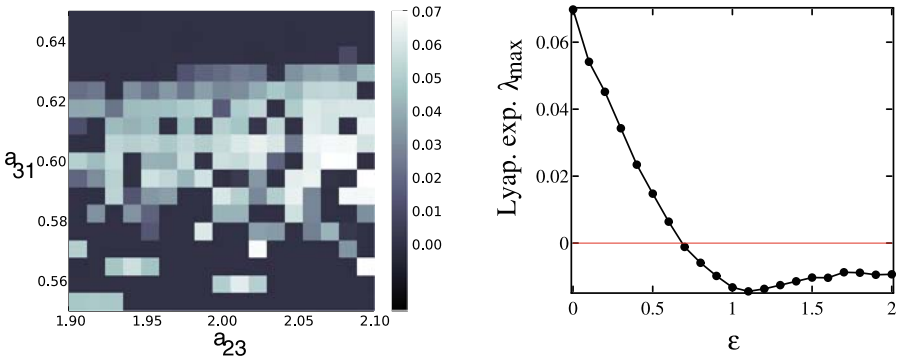

Fig. 14 Suppression of intrinsic chaos in the 3-cell network in Diagram (11). *Left*: the Lyapunov exponent $\lambda_{\max}$ with $\varepsilon = 0$ is shown as a function of a_{23} and a_{31} (a_{12} is fixed at 1). Each exponent is computed using 4 initial conditions. For the vast majority ($\geq$95%) of the points computed, the error in the computed exponent is $\lesssim$0.005. *Right*: we show the response to $\varepsilon > 0$ for $a_{23} = 2.075$ and $a_{31} = 0.6025$. In all plots, we use $\omega_1 = 1$, $\omega_2 = 0.95$, and $\omega_3 = 1.1$

Observation 2 *Some networks with weak intrinsic chaos will respond reliably to moderately strong stimuli.*

4.2 Acyclic Networks and Modular Decompositions

Having seen from Part II that recurrent connections can lead to very complex behavior, we begin with the class of *acyclic* networks. These are networks in which there is a well defined direction of information flow. We will show that acyclic networks are not intrinsically chaotic, and they are never unreliable. Building on the analysis developed for acyclic networks, we identify a broader class that may be accessible to analysis, namely networks that admit a decomposition into modules with acyclic inter-module connections.

For simplicity, we assume throughout that the stimuli are independent; it is trivial to modify the results of this section to accommodate the situation when some of them are identical to each other.

4.2.1 Skew-Product Representation of Acyclic Networks

We first describe the connection graph that corresponds to an oscillator network. Let each node of this graph, $i \in \{1, \ldots, N\}$, correspond to an oscillator. Assign a directed edge from node i to node j, $i \neq j$, if oscillator i provides input to oscillator j, i.e., if $a_{ij} \neq 0$; in this case, we write $i \to j$. (For simplicity, we do not allow nodes to connect to themselves.) A *cycle* in such a directed graph is a sequence of nodes and edges $i_1 \to i_2 \to \cdots \to i_k \to i_1$ for some $k > 1$.

Definition 4.1 We say an oscillator network is *acyclic* if its connection graph has no cycles.

Given any pair of oscillators in an acyclic network, either they are "unrelated", or one is "upstream" from the other. We say oscillator i is "upstream" from oscillator

j if there is a sequence of nodes and edges that goes from i to j. The absence of cycles is precisely what makes this "upstream"–"downstream" notion well defined. We say oscillators i and j are "unrelated" if there is no chain that goes from i to j or vice versa. Unrelated oscillators do not necessarily behave independently: They may, for example, receive input from the same source. Overall, the structures of acyclic networks can still be quite complex, with tree-like branching and recombinations possible.

Our first task is to find a systematic way to treat the dynamics within these networks.

Lemma 4.1 *In an acyclic graph, one can define for each node j a number $m(j)$ representing its maximum number of ancestors, meaning*

(i) *there is a chain of the form $i_1 \to i_2 \to \cdots \to i_{m(j)} \to j$, and*
(ii) *there is no chain of the form $i_1 \to i_2 \to \cdots \to i_{m(j)+1} \to j$.*

The proof is simple: chains cannot be arbitrarily long without a node repeating, and such repeats are impossible in acyclic graphs.

Using the notation introduced in Part I, we now discuss the dynamical structure of acyclic oscillator networks. First, let φ_t denote the flow on $\mathbb{T}^N$ in the absence of inputs, i.e., with $\varepsilon_i \equiv 0$. We say φ_t *factors into a hierarchy of skew-products with* 1*-dimensional fibers* if after a permutation of the names of the N oscillators, we have the following: For each $k = 1, \ldots, N$, there is a vector field $X^{(k)}$ on $\mathbb{T}^k$ such that if $\varphi_t^{(k)}$ is the flow generated by $X^{(k)}$, then (i) $\varphi_t^{(k)}$ describes the dynamics of the network defined by the first k oscillators and the relations among them, and (ii) $\varphi_t^{(k+1)}$ is a *skew-product* over $\varphi_t^{(k)}$, that is, the vector field $X^{(k+1)}$ on $\mathbb{T}^{k+1}$ has the form

$$X^{(k+1)}(\theta_1, \ldots, \theta_{k+1}) = \big(X^{(k)}(\theta_1, \ldots, \theta_k), Y_{(\theta_1,\ldots,\theta_k)}(\theta_{k+1})\big) \tag{12}$$

where $\{Y_{(\theta_1,\ldots,\theta_k)}\}$ is a family of vector fields on S^1 parametrized by $(\theta_1, \ldots, \theta_k) \in \mathbb{T}^k$. In particular, $\varphi_t^{(N)} = \varphi_t$. In the system defined by (12), we refer to $\varphi_t^{(k)}$ on $\mathbb{T}^k$ as the flow on the *base*, and each copy of S^1 over $\mathbb{T}^k$ as a *fiber*.

Proposition 4.1 *The flow of every acyclic network of N oscillators with no inputs can be represented by a hierarchy of skew-products with* 1*-dimensional fibers.*

Proof Let $\mathcal{N}_m$ be the set of oscillators j with $m(j) = m$. Assign an ordering to the oscillators so that all the oscillators in $\mathcal{N}_0$ come first, followed by those in $\mathcal{N}_1$, then $\mathcal{N}_2$, and so on. The order within each $\mathcal{N}_m$ is immaterial. Let us denote this ordering by $i_1 < i_2 < \cdots < i_N$. A skew-product is built inductively as follows: For $k \geq 1$, consider the subnetwork consisting of oscillators $i_1, \ldots, i_k$. None of the oscillators i_j, $j \leq k$, receives input from any of the oscillators i_ℓ, $\ell > k$. (To see this, note that $i_\ell \to i_j$ implies $m(i_j) > m(i_\ell)$ by definition, but this is impossible under the ordering we have chosen.) Therefore, the dynamics of oscillators $i_1, \ldots, i_k$ as a subsystem of the entire network may be described by a vector field of the form $X^{(k)}(\theta_{i_1}, \ldots, \theta_{i_k})$, obtained by projecting $X^{(N)}$ onto these k coordinates. Starting with $k = 1$, the skew-products in the lemma are constructed inductively for increasing k. □

Next we generalize to acyclic networks with stimuli. Such networks can also be represented by a directed graph of the type described above, except that some of the nodes correspond to stimuli and others to oscillators. If i is a stimulus and j an oscillator, then $i \to j$ if oscillator j receives information directly from stimulus i. No arrow can terminate at a stimulus, so that $m(i) = 0$ if i is a stimulus. Clearly, a network driven by stimuli is acyclic if and only if the corresponding network without stimuli is acyclic.

Proceeding to skew-product representations, consider first the case of a single oscillator driven by a single stimulus. Let Ω denote the set of all Brownian paths defined on $[0, \infty)$, and let $\sigma_t : \Omega \to \Omega$ be the time shift. Then the dynamics of the stochastic flow discussed in Sect. 2.2.2 can be represented as the skew-product on $\Omega \times S^1$ with

$$\Phi_t : (\omega, x) \mapsto \big(\sigma_t(\omega), F_{0,t;\omega}(x)\big).$$

Similarly, a network of N oscillators driven by q independent stimuli can be represented as a skew-product with base Ω^q (equipped with product measure) and fibers $\mathbb{T}^N$.

Proposition 4.2 *The dynamics of an acyclic network driven by q stimuli can be represented by a hierarchy of skew-products over Ω^q with 1-dimensional fibers.*

The proof is identical to that of Proposition 4.1, except that when enumerating the nodes of the graph, we first list all of the stimuli (in any order) before listing the oscillators.

4.2.2 Lyapunov Exponents of Acyclic Networks

Consider a network of N oscillators driven by q independent stimuli. We first review some notation. Let $\omega \in \Omega^q$ denote a q-tuple of Brownian paths, and let $F_{0,t;\omega}$ denote the corresponding stochastic flow on $\mathbb{T}^N$. For a fixed $\omega \in \Omega^q$, $x \in \mathbb{T}^N$ and tangent vector v at x, define the *Lyapunov exponent*

$$\lambda_\omega(x, v) = \lim_{t\to\infty} \frac{1}{t} \log\big|DF_{0,t;\omega}(x)v\big| \tag{13}$$

if this limit exists. If μ is a stationary measure of the stochastic flow, then for a.e. ω and μ-a.e. x, $\lambda_\omega(x, v)$ is well defined for all v. The following is the main result of this section:

Theorem 3 *Consider a network of N oscillators driven by q independent stimuli, and let μ be a stationary measure for the stochastic flow. Assume*

(a) *the network is acyclic, and*
(b) *μ has a density on $\mathbb{T}^N$.*

Then $\lambda_\omega(x, v) \le 0$ for a.e. $\omega \in \Omega^q$ and μ-a.e. x.

One way to guarantee that condition (b) is satisfied is to set ε_i to a very small but strictly positive value if oscillator i is not originally thought of as receiving a

stimulus, so that $\varepsilon_i > 0$ for all i. Such tiny values of ε_i have minimal effect on the network dynamics. Condition (b) may also be satisfied in many cases where some $\varepsilon_i = 0$ if suitable hypoellipticity conditions are satisfied, but we do not pursue this here (Nualart 2006).

Before proceeding to a proof, it is useful to recall the following facts about Lyapunov exponents. For a.e. ω and μ-a.e. x, there is an increasing sequence of subspaces $\{0\} = V_0 \subset V_1 \subset \cdots \subset V_r = \mathbb{R}^N$ and numbers $\lambda_1 < \cdots < \lambda_r$ such that $\lambda_\omega(x, v) = \lambda_i$ for every $v \in V_i \setminus V_{i-1}$. The subspaces depend on ω and x, but the exponents λ_j are constant *a.e.* if the flow is ergodic. We call a collection of N vectors $\{v_j\}$ a *Lyapunov basis* if exactly $\dim(V_i) - \dim(V_{i-1})$ of these vectors are in $V_i \setminus V_{i-1}$. If $\{v_j\}$ is a Lyapunov basis, then for any $u, v \in \{v_j\}$, $u \neq v$,

$$\lim_{t\to\infty} \frac{1}{t} \log\left|\sin \angle\left(DF_{0,t;\omega}(x)u, DF_{0,t;\omega}(x)v\right)\right| = 0, \tag{14}$$

that is, angles between vectors in a Lyapunov basis do not decrease exponentially fast; see, e.g., Young (1995) for a more detailed exposition.

Proof of Theorem 3 Since the network is acyclic, it factors into a hierarchy of skew-products. The kth of these is a stochastic flow $F^{(k)}_{0,t;\omega}$ on $\mathbb{T}^k$. It describes the (driven) dynamics of the first k oscillators assuming they have been reordered so that oscillator i is upstream from or unrelated to oscillator j for $i < j$. Let $\mu^{(k)}$ denote the projection of μ onto $\mathbb{T}^k$. Then $\mu^{(k)}$ is an stationary measure for $F^{(k)}_{0,t;\omega}$, and it has a density since μ has a density. We will show inductively in k that the conclusion of Theorem 3 holds for $F^{(k)}_{0,t;\omega}$.

For $k = 1$, $\lambda_\omega(x, v) \leq 0$ for a.e. ω and $\mu^{(1)}$-a.e. x. This is a consequence of Jensen's Inequality; see, e.g., Sect. 2.2.2 for more details.

Assume we have shown the conclusion of Theorem 3 up to $k-1$, and view $F^{(k)}_{0,t;\omega}$ as a skew-product over $\Omega^q \times \mathbb{T}^{k-1}$ with S^1-fibers. Choose a vector v_k in the direction of the S^1-fiber. Note that this direction is invariant under the variational flow $DF^{(k)}_{0,t;\omega}$ due to the skew-product structure. Starting with v_k, we complete a Lyapunov basis $\{v_1, \ldots, v_k\}$ at all typical points. Due to the invariance of the direction of v_k, we may once more use Jensen's Inequality to show that $\lambda_\omega(x, v_k) \leq 0$ for a.e. x and ω. We next consider v_i with $i < k$. First, define the projection $\pi : \mathbb{T}^k \to \mathbb{T}^{k-1}$, and note that

$$\left|DF^{(k)}_{0,t;\omega}(x)v_i\right| = \frac{|\pi(DF^{(k)}_{0,t;\omega}(x)v_i)|}{|\sin\angle(v_k, DF^{(k)}_{0,t;\omega}(x)v_i)|}.$$

Due to (14), we have $\lambda_\omega(x, v_i) = \lim_{t\to\infty} \frac{1}{t} \log|\pi(DF^{(k)}_{0,t;\omega}(x)v_i)|$. But the skew-product structure yields $\pi(DF^{(k)}_{0,t;\omega}(x)v_i) = DF^{(k-1)}_{0,t;\omega}(\pi x)(\pi v_i)$, so by our induction hypothesis, $\lambda_\omega(x, v_i) \leq 0$. □

Remarks on Reliability of Acyclic Networks: Our conclusion of $\lambda_{\max} \leq 0$ falls short of reliability, which corresponds to $\lambda_{\max} < 0$. When there are free-rotating oscillators in a network, i.e., oscillators that are not driven by either a stimulus or another oscillator, then clearly $\lambda_{\max} = 0$. When no free-rotating oscillators are present, typically

one would expect $\lambda_{\max} < 0$. We do not have a rigorous proof, but this intuition is supported by numerical simulations.

Arguments similar to those in the proof of Theorem 3 (but in the absence of stimuli and without the use of an invariant measure) give the following:

Proposition 4.3 *Acyclic networks are never intrinsically chaotic, in the sense that at Lebesgue*-a.e. $x \in \mathbb{T}^N$, *all Lyapunov exponents* (*with* lim sup *instead of* limit *in* (13)) *are* ≤ 0.

4.2.3 Modular Decompositions and Quotient Systems

We next describe how the reliability of more general networks may be analyzed by decomposition into subunits. Consider a graph with nodes $\{1, \ldots, N\}$ as in the beginning of Sect. 4.2.1, and let $\sim$ be an equivalence relation on the set $\{1, \ldots, N\}$. The quotient graph defined by $\sim$ has as its nodes the equivalence classes $[i]$ of $\sim$, and we write $[i] \to [j]$ if there exists $i' \in [i]$ and $j' \in [j]$ such that $i' \to j'$. The following is a straightforward generalization of Proposition 4.2:

Proposition 4.4 *In a network of oscillators driven by q independent stimuli, if an equivalence relation leads to an acyclic quotient graph, then the dynamics of the network can be represented by a hierarchy of skew-products over Ω^q, with the dimensions of the fibers equal to the sizes of the corresponding equivalence classes.*

We pause to discuss in more detail the structure in Proposition 4.4, as it is important in the rest of this paper. An equivalence relation on the nodes of a network describes a decomposition of the network into smaller subunits called *modules*. Introducing directed edges between modules as above, we obtain what we call a *quotient network* in this paper. Assume this quotient network is acyclic, and let $M_1, M_2, \ldots, M_p$ be the names of the modules ordered in such a way that M_i is upstream from or unrelated to M_j for all $j > i$. Let k_j be the number of nodes in module M_j. For $s_i = k_1 + k_2 + \cdots + k_i$, let $F^{(s_i)}_{0,t,\omega}$ denote, as before, the stochastic flow describing the dynamics within the union of the first i modules; we do not consider $F^{(s)}_{0,t,\omega}$ when $s \neq s_i$ for some i. The dynamics of the entire network can then be built up layer by layer as follows: We begin with the stochastic flow $F^{(k_1)}_{0,t,\omega}$. Then proceed to $F^{(k_1+k_2)}_{0,t,\omega}$, which we view as a skew-product over $F^{(k_1)}_{0,t,\omega}$. This is followed by $F^{(k_1+k_2+k_3)}_{0,t,\omega}$, which we view as a skew-product over $F^{(k_1+k_2)}_{0,t,\omega}$, and so on.

Let $\lambda^{(1)}_1, \ldots, \lambda^{(1)}_{k_1}$ denote the Lyapunov exponents of $F^{(k_1)}_{0,t,\omega}$. Clearly, these are the Lyapunov exponents of a network that consists solely of module M_1 and the stimuli that feed into it. We wish, however, to view M_1 as part of the larger network. If $\lambda^{(1)}_{\max} \equiv \max_j \lambda^{(1)}_j > 0$, we say *unreliability is produced within* M_1. An interesting question is the effect of this unreliability on sites downstream. Leaving this to Sect. 4.3, we continue with the present discussion: For $i > 1$, let $\lambda^{(i)}_1, \ldots, \lambda^{(i)}_{k_i}$ denote the *fiber Lyapunov exponents* in the skew-product representation of $F^{(s_i)}_{0,t,\omega}$ over $F^{(s_{i-1})}_{0,t,\omega}$, and let $\lambda^{(i)}_{\max} = \max_j \lambda^{(i)}_j$. Then $\lambda^{(i)}_{\max} > 0$ has the interpretation that *unreliability is produced*

within module M_i as it operates in the larger network. It is important not to confuse this with the Lyapunov exponents of module M_i in isolation, an issue we will follow up in Sect. 4.4.1.

Analogous interpretations for the zero-input systems are obvious: for $i > 1$, $\lambda_{\max}^{(i)} > 0$ at $\varepsilon = 0$ means there is intrinsic network chaos within the module M_i as a subsystem of the larger system, and so on.

The proof of the following result is virtually identical to that of Theorem 3:

Proposition 4.5 *Suppose for a driven network there is an equivalence relation leading to an acyclic quotient graph. Then, with respect to any ergodic stationary measure μ, the numbers $\lambda_j^{(i)}$, $1 \le i \le p$, $1 \le j \le k_i$, are precisely the Lyapunov exponents of the network.*

Proposition 4.5 says in particular that if, in each of the p skew-products in the hierarchy, the fiber Lyapunov exponents are ≤ 0, i.e., if no unreliability is produced within any of the modules, then $\lambda_{\max}$ for the entire network is ≤ 0. Conversely, if unreliability is produced within any one of the modules as it operates within this network, then $\lambda_{\max} > 0$ for the entire network. On the practical level, the skew product structure (which implies that $DF_{0,t;\omega}$ is block-lower-triangular) and the proposition together give a more efficient way to numerically compute Lyapunov exponents of networks with acyclic quotients.

Important Remarks. *In the rest of this paper, we will often take the view that the network in question is equipped with a modular decomposition connected by an acyclic graph.* Such a decomposition always exists for any network, but it can be trivial (e.g., when the entire network is a single module).[3] If the decomposition is nontrivial and $\lambda_{\max} > 0$ for the network, Proposition 4.5 enables us to localize the *source* of the unreliability, i.e., to determine in which module unreliability is produced via their fiber Lyapunov exponents. In particular, modules that are themselves acyclic cannot produce unreliability.

As noted earlier, the idea of "upstream"–"downstream" for acyclic networks extends to modules connected by acyclic graphs, so that it makes sense to speak of a module as being downstream from another module, or a site as being downstream from another site, meaning the modules in which they reside are so related.

4.3 Propagation of Unreliability

In this section, we address the following basic question: *Under what conditions will unreliability generated in one part of a network propagate downstream?* In Sect. 4.3.1, we discuss how to measure unreliability at specific network sites, and in Sect. 4.3.2, we address the question posed.

[3] It is straightforward to show that there is always a *unique* modular decomposition connected by an acyclic graph that is "maximal" in the sense that it cannot be refined any further without introducing cycles into the quotient graph.

4.3.1 Measuring Reliability at Individual Network Sites

Often, it is the response of a network measured at specific oscillators (or *sites*) that is of relevance, rather than the response of the network as a whole. While $\lambda_{\max} > 0$ tells us that there is unreliability somewhere in the system, it does not tell us which oscillators are affected. To capture the idea of reliability at individual sites, recall that the reliability of the entire system is reflected in the sample measures μ_ω (see Sect. 2.2.2). By the same reasoning, the reliability at site i is reflected in the marginals of μ_ω in the variable θ_i; we denote this marginal distribution by $\mu_{\omega,i}$. We say a network is *reliable at site i* if $\mu_{\omega,i}$ is concentrated at a single point; the more uniformly distributed these projected sample measures are, the greater the *site unreliability*. These ideas are easily generalized to groups of more than one site, but we will treat only single site reliability.

The following are three standard ways in which the distribution of $\mu_{\omega,i}$ can be described:

A. Site Entropy. For each i and ω, we let $H(i,\omega)$ denote the entropy of the distribution $\mu_{\omega,i}$, i.e., if $\rho_{\omega,i}$ is the density of $\mu_{\omega,i}$ with respect to Lebesgue measure on S^1, then

$$H(i,\omega) = -\int_{S^1} \log \rho_{\omega,i}\, d\rho_{\omega,i},$$

and we set $H(i,\omega) = -\infty$ if $\mu_{\omega,i}$ is singular. The *site entropy* $H(i)$ is defined to be the expected value of this random variable, i.e., $H(i) = \int H(i,\omega) P(d\omega)$. In practice, $H(i)$ is computed as $\lim_{T\to\infty} \frac{1}{T}\int_0^T H(i,\sigma_t(\omega))\, dt$ via the Ergodic Theorem. This number can range from $-\infty$ to 0, with $H(i) = 0$ corresponding to uniform distribution and $H(i) = -\infty$ corresponding to the distribution being singular with respect to Lebesgue measure. A drawback of site entropy is that it does not distinguish among singular distributions.

B. Information Dimension. We define

$$D(i,\omega) = \lim_{k\to\infty} D_k(i,\omega) \quad \text{where } D_k(i,\omega) = \frac{-\sum_{j=1}^k p_j \log p_j}{\log k}.$$

In the quantity on the right, S^1 is divided into k equal intervals, and p_j is the probability with respect to $\mu_{\omega,i}$ of finding the phase of oscillator i in the jth subinterval. The relevant quantity is then $D(i) = \int D(i,\omega) P(d\omega)$. This takes values on $[0,1]$, with $D(i) = 1$ corresponding to any distribution having a density. Information dimension does not, for example, distinguish between the uniform distribution on S^1 and a distribution supported uniformly on a tiny subinterval of S^1.

C. Cumulative Distribution Functions (CDFs). The most direct way to assess site distributions is to compute the CDF of $\mu_{\sigma_t\omega,i}$, defined as $\mathrm{CDF}(\theta) = \int_0^\theta d\mu_{\sigma_t\omega,i}$, i.e., the probability at time t of finding the phase of oscillator i at a value less than θ. The CDF is very simple to compute numerically, and is especially effective in establishing

whether the measure is concentrated at a single point. A drawback of using CDFs is that it is not a number. Moreover, to be certain that one is seeing "typical" CDFs, one needs to compute them for $\mu_{\sigma_t\omega,i}$ for many values of ω, as CDFs cannot be averaged.

Above, we discussed various ways to assess $\mu_{\omega,i}$ from a purely theoretical standpoint. Returning to the situation at hand, recall from Theorem 1 that when $\lambda_{\max} > 0$, i.e., when the system is unreliable, μ_ω is a random SRB measure. These measures have densities on complicated families of smooth manifolds. In particular, they have dimensions >1; e.g., in the case of 2 oscillators, the dimension of the random SRB measures is $1+\alpha$ for some $\alpha \in (0, 1)$ where α describes the dimension of the fractal part. A well known result in analysis of Mattila (1995) tells us that when measures of dimensions >1 are projected onto 1-dimensional subspaces, the projected measures have a density for projections along *almost all* directions. Now not all projections are "good" in this sense, and we cannot be certain that the projection onto any particular site is "good" (projections onto sites upstream from where unreliability is produced are obviously not). Still, *if* the projection to site i is good, then $\lambda_{\max} > 0$ would imply that $\mu_{\omega,i}$ has a density. This being the case, it is more important to be able to compare different distributions with densities than to distinguish different singular distributions. We therefore favor site entropy over information dimension as a measure of site reliability.

In this paper, we will use a combination of site entropy and CDFs: If the computed values of $H(i)$ appear bounded below, then it is safe to conclude site unreliability, and the closer $H(i)$ is to 0, the more unreliable. If, on the other hand, the computed values of $H(i)$ appear unbounded, then CDFs are used to confirm site reliability.

4.3.2 Sites Downstream from Unreliable Modules

We now return to the question of propagation of unreliability to sites downstream from an unreliable module. The simplest network with which to investigate this is the N-chain system

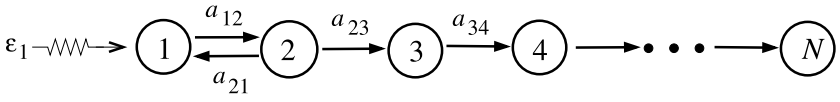

(15)

Suppose we choose a_{12}, a_{21}, and ε_1 so that the (1,2)-subsystem is unreliable. Notice that no unreliability is produced elsewhere. The question is: will this unreliability be observable downstream, at sites $i = 3, 4, \ldots, N$, or does it somehow "dissipate?" We run the following numerical experiment: we set $\omega_1 = 1$, $\omega_2 = 1.05$, $a_{12} = 1$, $a_{21} = 1.05$, and $\varepsilon_1 = 1$, so that the Lyapunov exponent $\lambda_{\max}$ of the (1,2)-subsystem is ≈ 0.1. For $i \geq 3$, we draw randomly (and fix) ω_i from $[0.9, 1.1]$ and $a_{i,i+1}$ from $\pm[0.6, 1.2]$. Computing the site entropies[4] $H(i)$ for sites $i = 3, 7, 10$, we find:

To help interpret these numbers, recall that identifying S^1 with $[0, 1]$, the entropy of the uniform distribution on an interval $[a, b]$ is $\log(b-a)$, and $\log\frac{1}{2} \approx -0.7$. The

[4]We compute site entropies by simulating the response of 10 000 initial conditions to the same realization of the stimulus, then applying the "binless" estimator of Kozachenko and Leonenko (1987) (see also Victor 2002).

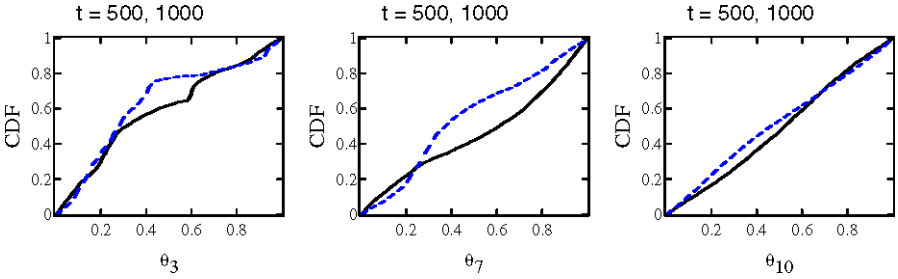

Fig. 15 Site distributions for the system (15) with $N = 10$. Shown are the site-i CDFs for $i = 3, 7, 10$, at times $t = 500$ (*solid*) and $t = 1000$ (*dash*). The parameters are as follows: for the (1,2)-subsystem, we set $\omega_1 = 1$, $\omega_2 = 1.05$, $a_{12} = 1$, $a_{21} = 1.18$, and $\varepsilon = 1$ (so that $\lambda_{\max} \approx 0.13$). For $i \geq 3$, we draw ω_i from [0.9, 1.1] and $a_{i,i+1}$ from $\pm$[0.6, 1.2]

Site	$i = 3$	$i = 7$	$i = 10$
Entropy $H(i)$	−0.4	−0.2	−0.01

data above thus indicate that the site distributions are fairly uniform, and, in fact, seem to become more uniform for sites farther downstream. Figure 15 shows the corresponding CDFs for some sites at representative times; CDFs at other sites are qualitatively similar. The graphs clearly show that unreliability propagates, as they are close to the diagonal lines that would represent uniform distributions at these sites.

We give two explanations for why one should expect this result of propagating site unreliability. The first is a plausibility argument along the lines of the projection argument in Sect. 4.3.1: it is possible—but highly unlikely—that the SRB measure μ_ω would project to point masses in any of the 8 directions corresponding to the 8 sites downstream. A second, perhaps more intuitive explanation, is as follows. Consider site 3 in the system (15). Fix a realization of the stimulus, and let $(\theta_1, \theta_2, \theta_3)$ denote the 3 phase coordinates. For each choice of $(\theta_1(0), \theta_2(0))$, the third oscillator receives a sequence of coupling impulses from the (1,2)-subsystem. Since the (1,2)-subsystem is unreliable, different choices of $\theta_1(0)$ and $\theta_2(0)$ will produce different sequences of coupling impulses to oscillator 3, in turn leading to different values of $\theta_3(t)$. This is synonymous with site unreliability for oscillator 3.

What happens if an oscillator receives inputs from more than one source with competing effects? The simplest situation is an oscillator driven by both an unreliable module upstream and an input stimulus, as depicted in Diagram (16). Note the presence of competing terms: As we have just seen, the unreliable module leads to unreliability at site 3. However, the stimulus ε_3 has a stabilizing effect on that oscillator.

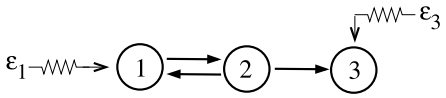

(16)

In the results tabulated below, the parameters used are $\omega_1 = 1$, $\omega_2 = 1.05$, $a_{12} = 1$, $a_{21} = 1.15$, and $\varepsilon_1 = 1$ (so that the (1,2)-subsystem is again unreliable), $a_{23} = 0.5$ and $\omega_3 = 0.93$. The site entropy $H(3)$ is computed for various values of ε_3. We find:

Stimulus amplitude	$\varepsilon_3 = 0.2$	$\varepsilon_3 = 1$	$\varepsilon_3 = 2$
Site entropy $H(3)$	-0.4	-1.2	-2

These numbers tell us that at $\varepsilon_3 = 0.2$, the distribution at site 3 is fairly uniform, and that this distribution becomes more concentrated as ε_3 increases. The CDFs (not shown) confirm this. When $\varepsilon_3 = 2$, for example, about 80% of the distribution $\mu_{\omega,3}$ is concentrated on an interval of length $\leq 1/5$ roughly 70% of the time. These data show that the source of reliability, i.e., the stimulus into oscillator 3, attenuates the propagation of unreliability; we call this phenomenon *interference*. Moreover, we find that when ε_3 is increased further, the support of a large fraction of $\mu_{\omega,3}$ shrinks to smaller and smaller intervals, decreasing the entropy and reflecting a greater tendency to form random sinks in oscillator 3. However, simulations also show that oscillator 3 does not become fully site reliable even at fairly strong forcing. This is also expected: intuitive arguments similar to those above suggest that once created, site unreliability cannot be completely destroyed at downstream sites. To summarize:

Observation 3 *Unreliability, once generated, propagates to all sites downstream*

We finish by clarifying the relation between the material in Sect. 4.2.3 and this section:

Propagation versus Production of Unreliability: The topic of Sect. 4.3 is whether or not unreliability created upstream propagates, i.e., whether it can be observed downstream. This concept complements an idea introduced in Sect. 4.2.3, namely the production of *new* unreliability within a module as measured by the positivity of fiber Lyapunov exponents. Mathematically, site reliability (or unreliability) is reflected in the marginals of μ_ω at the site in question, while the unreliability produced within a module is reflected in the dynamics and conditional measures of μ_ω on fibers. Naturally, when a module produces unreliability in the sense of Sect. 4.2.3, its sites will also show unreliability in the sense of Sect. 4.3.1, and one cannot separate what is newly produced from what is passed down from upstream.

4.4 A Question and a Final Illustrative Example

4.4.1 A Reassembly Problem

We now return to the strategy suggested in Sect. 4.2.3, to study the reliability of networks by analyzing their component modules separately. An obvious benefit of this strategy is that smaller modules are easier to test. However, the strategy will only be successful if the reliability of modules in isolation gives a good indication of their reliability when they are embedded in larger systems.

Pictorially, the problem we face can be represented as follows:

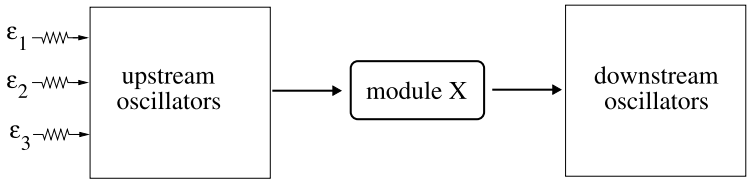

(17)

Suppose our network can be decomposed into three parts: Module X, which is our module of interest, a (possibly large) component upstream from Module X, and a (possibly large) component downstream from Module X. The question is: *Can the reliability properties of Module X as a subsystem of this larger network be predicted by its reliability properties when it is isolated and subjected to a white noise stimulus?*

At the heart of this question is the following issue: A module embedded in a network may receive input from oscillators upstream in the form of coupling impulses; it may also hear directly from the external stimuli. Inputs from other oscillators resemble a point process of impulses with statistics somewhere between Poisson and periodic; moreover, when an oscillator receives kicks from multiple sources, these kicks may be correlated to varying degrees. The question is then whether reliability properties of a module are sufficiently similar under these different classes of inputs that they can be predicted from studies using white noise stimuli.

We do not know the answer to this very difficult yet very important question. To make some initial progress, we have conducted a numerical study in which the which the 2-cell system studied in Part II (see beginning of Part II) is used as Module X in the framework above. Specifically, we consider the following networks with Module X enclosed in dotted lines:

Network A:

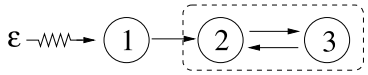

Network B:

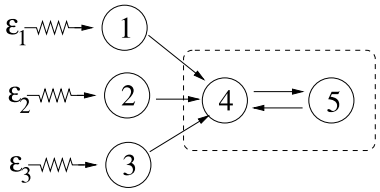

Network C:

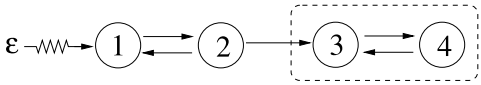

The parameters ω_i and a_{ff} for Module X are exactly as in Fig. 13.[5] Let $\lambda_{\max}^{\text{fib}}$ denote the largest fiber Lyapunov exponent of Module X in each of the networks above (see Sect. 4.2.3 for a discussion of fiber Lyapunov exponents). As in Fig. 13, Lyapunov exponents are plotted in Fig. 16 as functions of ε and a_{fb}—except that here we use $\lambda_{\max}^{\text{fib}}$ and not $\lambda_{\max}$, and the stimulus ε is not fed directly into Module X but to some sites upstream. Since the parts of the networks upstream from Module X in the 4 panels in Fig. 16 are all different (they are reliable in the top two panels and unreliable in the bottom two), we do not expect these panels to be carbon copies of Fig. 13. Yet, the qualitative resemblance i sundeniable. This general pattern persists for many other simulations not shown: As the parameters governing the network within which Module X is embedded are varied, the general tendency is that its reliability when embedded continues to resemble its reliability when it receives white noise inputs, even though the reliable and unreliable regions may shift and the magnitudes of $\lambda_{\max}^{\text{fib}}$ may vary.

We take these results to be limited affirmative answer to the question posed above. In sum:

Observation 4 *The reliability profile of the two-oscillator system driven by white noise is a reasonable guide to its reliability properties as a module embedded in a larger network.*

This observation is not totally unexpected, as we have seen in a number of situations in general dynamical systems theory that the response of a system to external forcing depends considerably more strongly on its underlying geometry than on the type of forcing. One systematic study of this type is carried out in Lin and Young (2008). However, we emphasize that it remains to be seen whether the observation above will carry over to other modules within the present framework.

4.4.2 An Illustrative Example

We finish with the following example, which illustrates many of the points discussed in Part III.

The network depicted in Fig. 17(a) is made up of 9 oscillators. It receives a single input stimulus through oscillator 1 and has "output terminals" at sites 6, 8 and 9—it is here that we will assess the response of the network. We first discuss what to expect based on the ideas above. This is then compared to the results of numerical simulations.

A cursory inspection tells us that this network is acyclic except for the subsystem (4,7). The finest decomposition that yields an acyclic quotient, then, is to regard

[5] That is, for network A: $\omega_2 = 1$, $\omega_3 = 1.1$, $a_{23} = 1$, for network B: $\omega_4 = 1$, $\omega_5 = 1.1$, $a_{45} = 1$, and for network C, $\omega_3 = 1$, $\omega_4 = 1.1$, $a_{34} = 1$.

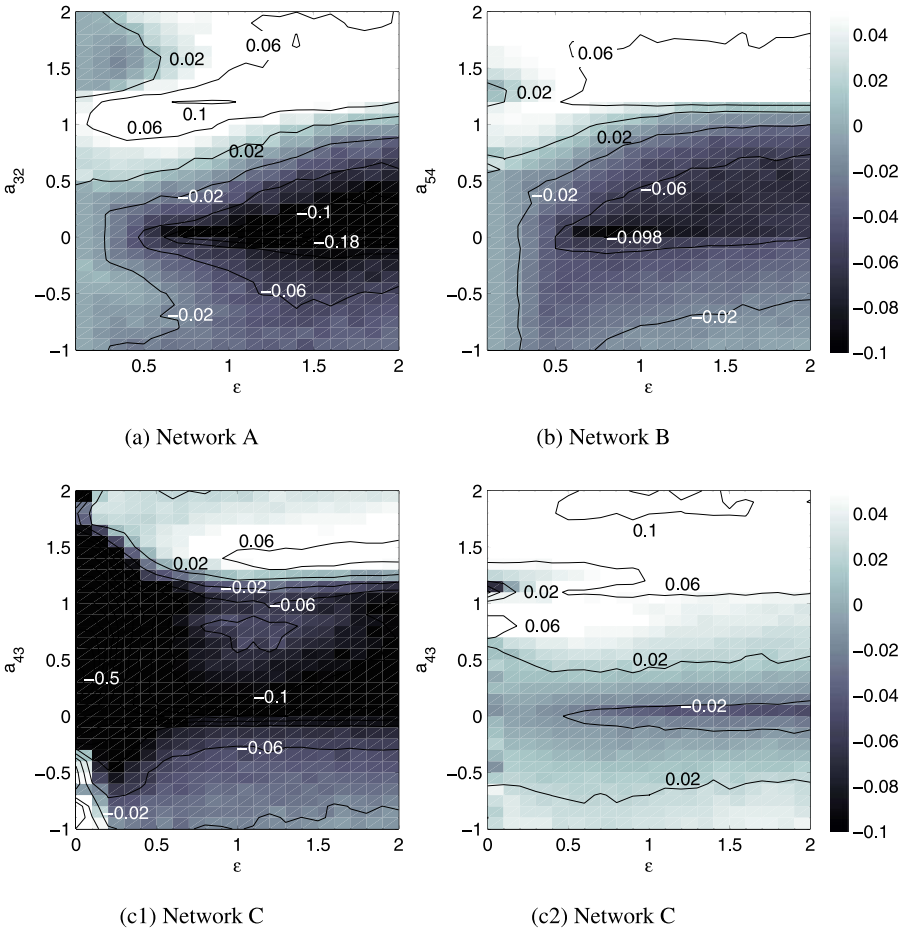

(a) Network A (b) Network B

(c1) Network C (c2) Network C

Fig. 16 The maximum fiber exponent $\lambda_{\max}^{\text{fib}}$ for Module X as an embedded subsystem. The parameters within Module X correspond to the following values for the two-cell network of Part I: $\omega_1 = 1$, $\omega_2 = 1.1$, and $a_{\text{ff}} = 1$ (see Footnote 5). The other parameters are as follows: for Network A, we use $\omega_1 = 0.97$, $a_{12} = 0.7$. For Network B, we draw $\omega_i \in [0.95, 1.05]$ and set $a_{i,4} = \frac{1}{3} \cdot (0.6, 0.8, 0.9)$ and $\varepsilon_1 = \varepsilon_2 = \varepsilon_3 = \varepsilon$. For network C, we explore two different parameter sets. In panel (**c1**), we take $a_{12} = 1.1$, $a_{21} = 0.8$, $a_{23} = 0.5$, and $a_{34} = 1$. In panel (**c2**), we take $a_{12} = -1.1$, $a_{21} = -1.2$, $a_{23} = 0.95$, and $a_{34} = 1$. In both cases, we use $\omega_1 = 1.03$, $\omega_2 = 0.98$, $\omega_3 = 1$, and $\omega_4 = 1.1$

each oscillator as a module except for (4,7), which must be grouped together as one. We choose, however, to work with a coarser decomposition, in which the (1,2,3,5,6)-subsystem is viewed as Module A, the (4,7)-subsystem as Module B, and the (8,9)-subsystem as Module C. Identifying the sites within each of these modules produces an acyclic quotient, as shown in Fig. 17(b).

Since Module B is the only module that is not itself acyclic and hence that is capable of generating unreliability, our results from Sect. 4.2 tell us that $\lambda_{\max}$ for the entire network is ≤ 0 if and only if no unreliability is produced in Module B. In fact, since there are no freely rotating oscillators in the system, we expect $\lambda_{\max}$ to be < 0

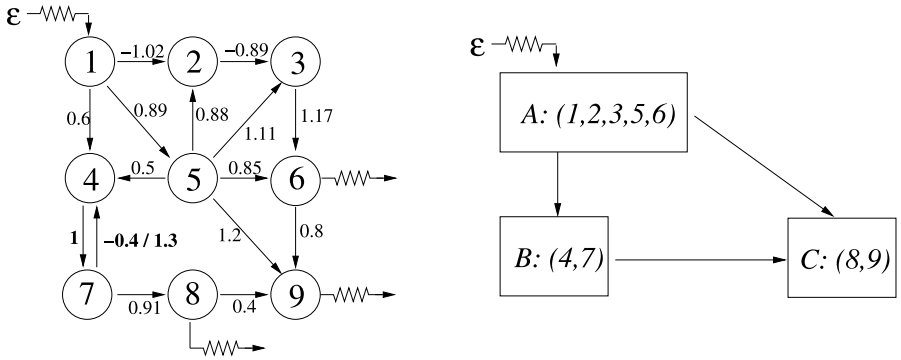

(a) Full network (b) Quotient graph

Fig. 17 Example of a larger network and its quotient graph. In (**a**), we have labeled the edges with a sample of coupling constants. The ω_i are drawn from [0.95, 1.05]

if Module B behaves reliably. The behavior of Module B hinges a great deal on (i) the couplings a_{47} and a_{74}, which determine its reliability in isolation, and (ii) intrinsic network properties in the two Modules A and B together, especially when ε is small; see Sects. 4.1 and 4.4.1. With regard to (i), the reliability profile of the 2-oscillator system compiled in Part II is handy. If unreliability is produced in Module B, then we expect to find sites 8 and 9 to be unreliable, with a lower reading of site entropy at site 9 than 8 due to the stabilizing effects of Module A; see Sect. 4.3. This is what general theory would lead one to predict.

We now present the results of simulations.

First we confirm that Module A alone is reliable as predicted:[6] In addition to the a_{ij} given in Fig. 17, we randomly drew coupling constants from $\pm[0.8, 1.2]$ and frequencies $\omega_i \in [0.95, 1.05]$, and find that $\lambda_{\max}$ for Module A ranges from roughly -0.3 to -0.07 when $\varepsilon = 1$. Site distributions for the sites in Module A are, as predicted, well-localized. For a majority of parameters tested, 90% of an ensemble of 10^4 uniformly-chosen initial conditions has collapsed into a cube of side length $\leq 10^{-2}$ after $t = 60$–110; in all our simulations, the ensemble collapses into a cube of side length $\lesssim 10^{-7}$ after $t \approx 1000$. In particular, the "output terminal" at site 6 is always reliable.

Turning now to the reliability of Module B, we fixed the coupling constants as shown in Fig. 17, with $\omega_i \in [0.95, 1.05]$ and $\varepsilon = 1$, and ran simulations for the following two sets of parameters:

(a) $a_{47} = 1, a_{74} = -0.4$: this case is predicted to be reliable.
(b) $a_{47} = 1, a_{74} = 1.3$: this case is predicted to be unreliable.

The predictions above are based on the behavior of the two-oscillator network "in isolation" (receiving white-noise stimulus of amplitude 1, see Part II), and on the results

[6] All Lyapunov exponents presented in this section have standard errors of ≤ 0.004 as estimated by the method of batched means. By the Central Limit Theorem, this means the actual $\lambda_{\max}$ should lie within $\approx 2.5 \times 0.004 = 0.01$ of the computed value with $\gtrsim 99\%$ probability.

of Sect. 4.2. The following table summarizes the reliability properties of Module B, both in isolation and when embedded within the network:

	In isolation	Embedded in network	
		$\varepsilon = 0$	$\varepsilon = 1$
(a) $a_{47} = 1$, $a_{74} = -0.4$	$\lambda_{\max} = -0.07$	$\lambda^{\text{fib}}_{\max} = 0.014$	$\lambda^{\text{fib}}_{\max} = -0.15$
(b) $a_{47} = 1$, $a_{74} = 1.3$	$\lambda_{\max} = 0.13$	$\lambda^{\text{fib}}_{\max} = 0.076$	$\lambda^{\text{fib}}_{\max} = 0.097$

Note that these values are consistent with the proposal in Sect. 4.4.1: The behavior of the module at $\varepsilon = 1$ when embedded within the network is effectively determined by its behavior in isolation. Furthermore, by Proposition 4.5, we know the Lyapunov exponent $\lambda_{\max}$ of the *entire network* is equal to $\lambda^{\text{fib}}_{\max}$ of Module B in case (b), and is $\geq \lambda^{\text{fib}}_{\max}$ in case (a).

Finally, we study site distributions at sites 8 and 9. In case (a), we find again that computed site distributions are well localized. In case (b), we find the expected evidence of propagated unreliability and interference, with $H(8) = -0.3$ and $H(9) = -0.7$.

In summary, the results of our simulations are entirely consistent with predictions based on the ideas developed in this paper.

5 Summary and Conclusions

In this paper, networks of phase oscillators with pulsatile coupling are considered. External stimuli in the form of white noise are presented to a subset of the oscillators, and reliability properties of network responses are studied. The largest Lyapunov exponent is used (in most places) as a measure of unreliability. Except where indicated otherwise, most of our findings are based on a combination of simulations and heuristic arguments using dynamical systems ideas.

For Two-Cell Systems: Our main finding is that

> they can exhibit both reliable and unreliable responses depending on the signs and strengths of network connections and the stimulus amplitude ε.

Specifically:

1. *Dominantly feedforward networks are always reliable*, and they become more reliable with increasing ε (while dominantly feedback networks are neutral to weakly reliable).
2. When $a_{\text{ff}} \sim a_{\text{fb}}$, i.e., when *feedback and feedforward coupling strengths are comparable*, be they both negative or both positive (mutually inhibitory or excitatory), we have found:

(a) For *smaller ε, reliability depends extremely sensitively on coupling strengths*, with very reliable and quite unreliable configurations occurring in close proximity. This phenomenon is explained by mechanisms of phase locking and shear-induced chaos.
(b) For *larger ε, the system is typically unreliable.*

Results of the type in Item 2 will depend to varying degrees on the nature of the coupling and the phase response function. Our choices here are commonly used in neuroscience.

For Larger Networks:

1. Via a rigorous mathematical analysis, we proved the following:
   (a) *Acyclic networks*, i.e., networks with a well defined direction of information flow, *are never chaotic in the absence of inputs, and never unreliable when inputs are presented.*
   (b) *Networks with a modular decomposition and acyclic quotient are reliable if and only if all individual modules are reliable*, the latter being given by fiber Lyapunov exponents.

The last two points pertain to networks with nontrivial module decompositions:

4. *Once produced, unreliability will propagate*: it can be attenuated but not completely removed; without intervention it actually grows for sites farther downstream.
5. The question "*Can the reliability of modules* in isolation *give information on their behavior when embedded in a larger network?*" is raised and some supporting evidence presented.

Acknowledgements K.L. and E.S-B. held NSF Math. Sci. Postdoctoral Fellowships and E.S-B. a Burroughs-Welcome Fund Career Award at the Scientific Interface; L-S.Y. is supported by a grant from the NSF. We thank Brent Doiron, Bruce Knight, Alex Reyes, John Rinzel, and Anne-Marie Oswald for helpful discussions over the course of this project.

Appendix: Proof of Theorem 2

To prove Theorem 2 (see Sect. 3.1.2), we need two lemmas. Define $\Delta a_{\mathrm{fb}} = a_{\mathrm{fb}} - a_{\mathrm{ff}}$.

Lemma A.1 *There exist b_1 and $K > 0$ such that for all admissible $(\omega_1, \omega_2, a_{\mathrm{ff}}, a_{\mathrm{fb}})$, if $b < b_1$ and $\Delta a_{\mathrm{fb}} < K b^{-2} \Delta\omega$, then $T(1-b) < 1-b$.*

Lemma A.2 *There exist $b_2, C > 0$ and $x_1 \in (0, 1-b)$ such that for all admissible $(\omega_1, \omega_2, a_{\mathrm{ff}}, a_{\mathrm{fb}})$, if $b < b_2$ and $\Delta a_{\mathrm{fb}} > C\Delta\omega$, then $T(x_1) > x_1$.*

We first give the proof of Theorem 2 assuming these two lemmas.

Proof of Theorem 2 We prove (a); the proof of (b) is analogous. Let ω_1, ω_2 and a_{ff} be given. Requirements on the size of b will emerge in the course of the proof.

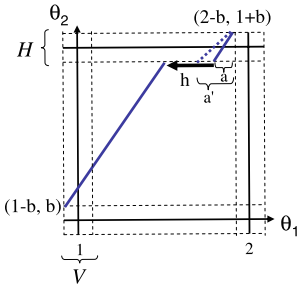

Fig. 18 Values used in proving Lemma A.1

Observe first that with $\omega_2 > \omega_1$, T has no fixed point (and hence there is no 1:1 phase locking) when $a_{\text{fb}} = a_{\text{ff}}$. This is because $T(x) = x$ when $\omega_2 = \omega_1$ and $a_{\text{fb}} = a_{\text{ff}}$ as noted earlier, and using Lemma 3.1(b), we see that for $\omega_2 > \omega_1$ and $a_{\text{fb}} = a_{\text{ff}}$, the graph of T is strictly below the diagonal.

Keeping ω_1, ω_2 and a_{ff} fixed, we now increase a_{fb} starting from $a_{\text{fb}} = a_{\text{ff}}$. By Lemma 3.1(a), this causes the graph of T to shift up pointwise. As a_{fb} is gradually increased, we let a^*_{fb} be the first parameter at which the graph of T intersects the diagonal, i.e., where there exists $x^* \in \Sigma_b$ such that $T(x^*) = x^*$—if such a parameter exists. Appealing once more to Lemma 3.1, we see that $\rho(T) < 1$ for all $a_{\text{fb}} < a^*_{\text{fb}}$, so T can have no fixed point for these values of a_{fb}.

We show now that a^*_{fb} exists, and that the phase-locking persists on an interval of a_{fb} beyond a^*_{fb}. First, if b is small enough, then by Lemma A.1, $T(1-b) < 1-b$ for all $a_{\text{fb}} < a_{\text{ff}} + Kb^{-2}\Delta\omega$. Now if b is small enough that $Kb^{-2} > C$ where C is as in Lemma A.2, then for $a_{\text{fb}} \in [a_{\text{ff}} + C\Delta\omega, a_{\text{ff}} + Kb^{-2}\Delta\omega]$, $T(x_1) > x_1$ for some $x_1 < 1-b$. For a_{fb} in this range, T maps the interval $[x_1, 1-b]$ into itself, guaranteeing a fixed point. It follows that (i) a^*_{fb} exists and is $< a_{\text{ff}} + C\Delta\omega$, and (ii) T has a fixed point for an interval of a_{fb} of length $\ell \geq (Kb^{-2} - C)\Delta\omega$. This completes the proof. □

We now proceed to the proofs of the lemmas. Let $\omega_1, a_{\text{ff}}, \Delta\omega$ and Δa_{fb} be given, with $\Delta\omega, \Delta a_{\text{fb}} > 0$. To study the system where ω_2 is defined by $\Delta\omega = 1 - \frac{\omega_1}{\omega_2}$ and $a_{\text{fb}} = a_{\text{ff}} + \Delta a_{\text{fb}}$, we will seek to compare trajectories for systems with the following parameter sets:

$$\begin{array}{llll} \text{System A}: & a_{\text{ff}}, \quad a_{\text{fb}} = a_{\text{ff}}, & \omega_1, & \omega_2 = \omega_1, \\ \text{System B}: & a_{\text{ff}}, \quad a_{\text{fb}} = a_{\text{ff}} & \omega_1, & \omega_2 = \omega_1 + \Delta\omega \cdot \omega_2, \\ \text{System C}: & a_{\text{ff}}, \quad a_{\text{fb}} = a_{\text{ff}} + \Delta a_{\text{fb}}, & \omega_1, & \omega_2 = \omega_1 + \Delta\omega \cdot \omega_2. \end{array}$$

That is, System C is the system of interest, and Systems A and B are used to help analyze System C. We introduce also the following notation: H and V denote the horizontal and vertical strips of width $2b$ centered at integer values of θ_2 and θ_1. We will work in $\mathbb{R}^2$ instead of $\mathbb{T}^2$.

Proof of Lemma A.1 For each of the 3 parameter sets above, two orbits are considered: Orbit 1 starts from $(\theta_1, \theta_2) = (1-b, b) \in \Sigma_b$ and runs forward in time until it meets $\Sigma_{1-b} = \{\theta_2 = 1-b\}$; orbit 2 starts from $(2-b, 1+b) \in \Sigma_{1+b}$ and runs

backward in time until it meets Σ_{1-b}. We need to prove that for System C, under the conditions in the lemma, the end point of orbit 1 lies to the left of the end point of orbit 2 (as shown in Fig. 18). This is equivalent to $T(1-b) < 1-b$.

For System A, orbits 1 and 2 meet, since for this set of parameters, $T(x) = x$ for all x as noted in Sect. 3.1.2. Comparing Systems A and B, since the vector field for System B has greater slope everywhere, and outside of $H \cup V$ it has slope $\frac{\omega_2}{\omega_1} > 1$, we conclude that for System B the end point of orbit 1 lies to the left of the end point of orbit 2, with a separation $h > \Delta\omega/2$; see Fig. 18.

Next, we compare Systems B and C. Orbit 1 for the two systems is identical, since the equation outside of H does not involve $a_{\rm fb}$. Orbit 2, however, differs for the two systems. To estimate by how much, we compare a, the distance from the end point of orbit 2 to $\theta_1 = 2$ for System B, and a', the corresponding distance for System C as marked in Fig. 18. First, there exist b_1 and $k_1 > 0$ such that for $\theta_1 \in (2-5b_1, 2)$, we have $z(\theta_1) < k_1(2-\theta_1)^2$ and $|z'(\theta_1)| < 2k_1(2-\theta_1)$. Shrinking b_1 further if necessary, we have that for $b < b_1$, orbit 2 has slope $> 1/2$ everywhere and therefore the entire orbit, for both Systems B and C, lies within the region $H \cap \{2-5b < \theta_1 < 2-b\}$. The next step is to apply Gronwall's Lemma to a system that incorporates both Systems B and C. Since $\theta_2(t)$ is identical for the two systems in the relevant region, we may write the equations as

$$\begin{aligned} \dot{\theta_1} &= -\omega_1 - (a_{\rm ff} + \delta) z(\theta_1) \hat{g}(t), \\ \dot{\delta} &= 0, \end{aligned}$$

where $\delta = 0$ corresponds to System B, $\delta = \Delta a_{\rm fb}$ corresponds to System C, and $\hat{g}(t) = g(\theta_2(t))$. Notice that each trajectory reaches Σ_{1-b} after a time $\tau = 2b/\omega_2$.

Motivated by the observation that $z(\theta_1) = \mathcal{O}(b^2)$ in the relevant rectangle, we rescale the variable δ by $\bar{\delta} = b^2\delta$, and define $\bar{z}(\theta_1) = \frac{1}{b^2} z(\theta_1)$. This gives

$$\dot{\theta_1} = -\omega_1 - \left(b^2 a_{\rm ff} + \bar{\delta}\right) \bar{z}(\theta_1) \hat{g}(t), \tag{18}$$

$$\dot{\bar{\delta}} = 0. \tag{19}$$

To find a, we solve (18)–(19) over the time interval $[0, \tau]$, starting from $(\theta_1(0), \bar{\delta}(0)) = (1-b, 0)$; to find a', we do the same, starting from $(1-b, b^2\Delta a_{\rm fb})$. Applying Gronwall's Lemma gives $|a' - a| < b^2 \Delta a_{\rm fb} \exp(L\tau)$, where L is the Lipschitz constant for the allied vector field. To estimate L, note that $|\hat{g}| = \mathcal{O}(\frac{1}{b})$, $|\bar{z}| = \mathcal{O}(1)$, and $|\bar{z}'| = \mathcal{O}(\frac{1}{b})$. This gives $L = \mathcal{O}(1/b)$, and $\exp(L\tau) = \mathcal{O}(1)$. Therefore, $|a' - a| < k_2 b^2 \Delta a_{\rm fb}$ for some constant k_2. Note that this constant can be made independent of $\omega_1, \omega_2, a_{\rm ff}$ or $a_{\rm fb}$.

Recall from the first part of the proof that to obtain the desired result for System C, it suffices to guarantee $|a' - a| < h$. This happens when

$$\Delta a_{\rm fb} < \frac{\Delta\omega}{2k_2 b^2} := K b^{-2} \Delta\omega. \tag{20}$$

□

Proof of Lemma A.2 All orbit segments considered in this proof run from $\Sigma_b \cap \{\theta_1 \in (0,1)\}$ to Σ_{1+b}. We assume $a_{\text{ff}} > 0$; the case $a_{\text{ff}} < 0$ is similar. First, we fix $x_0, x_1 > b$ so that for all admissible $(\omega_1, \omega_2, a_{\text{ff}}, a_{\text{fb}})$, the trajectory starting from x_1 intersects $H = \{1-b < \theta_2 < 1+b\}$ in $H \cap \{\theta_1 \in (1+x_0, \frac{3}{2})\}$. Such x_0, x_1 clearly exist for small enough b. Starting from x_1, we compare the trajectories for Systems A, B and C. We know that the trajectory for System A will end in $(1+x_1, 1+b)$. Thus, to prove the lemma, we need to show the trajectory for System C ends to the right of this point. This comparison is carried out in two steps:

Step 1: Comparing Systems A and B. We claim that the horizontal separation of the end points of these two trajectories is $< c\Delta\omega$ for some constant $c > 0$. It is not a necessary assumption, but the comparison is simpler if we assume b is small: First, a separation $c_1\Delta\omega$ in the θ_2 direction develops between the trajectories as they flow linearly from (x_1, b) to $\theta_1 = (1-b)$. Next, while the trajectories flow through V, Gronwall's Lemma can be used in a manner similar to the above to show they emerge from V with a separation $\le c_2\Delta\omega$ in the θ_2 direction. Third is the region of linear flow, resulting in a separation $\le c_3\Delta\omega$ in the θ_1 direction as the trajectories enter H. Gronwall's Lemma is again used in the final stretch as the trajectories traverse H.

Step 2: Comparing Systems B and C. Notice that up until they reach Σ_{1-b}, the two trajectories are identical. In H, their θ_2 coordinates are equal, and the crossing time is $\tau = \frac{2b}{\omega_2}$. Let $\theta_1^B(t)$ and $\theta_1^C(t)$, $t \in [0, \tau]$, denote their θ_1 coordinates while in H. We write

$$\theta_1^C(\tau) - \theta_1^B(\tau) = \int_0^\tau a_{\text{ff}}\big(z\big(\theta_1^C(t)\big) - z\big(\theta_1^B(t)\big)\big)g\big(\theta_2(t)\big)\,dt + \int_0^\tau \Delta a_{\text{fb}} z\big(\theta_1^C(t)\big)g\big(\theta_2(t)\big)\,dt.$$

The first integral is ≥ 0 by design: Via our choice of x_1, we have arranged to have $1 < \theta_1^B(t) < \theta_1^C(t) < \frac{3}{2}$, and the z-function is monotonically increasing between $\theta_1 = 1$ and $\theta_1 = \frac{3}{2}$. As for the second integral, we know $z(\theta_1^C(t))$ is bounded away from 0 since $1 + x_0 < \theta_1^C(t) < \frac{3}{2}$, so the integral is $> d\Delta a_{\text{fb}}$ for some constant $d > 0$. It follows that $T(x_1) > x_1$ if $d\Delta a_{\text{fb}} > c\Delta\omega$. □

References

Arnold, L.: Random Dynamical Systems. Springer, New York (2003)

Arnold, L., Imkeller, P., Sri Namachchivaya, N.: The asymptotic stability of a noisy non-linear oscillator. J. Sound Vib. **269**, 1003–1029 (2004)

Baxendale, P.H.: Stability and equilibrium properties of stochastic flows of diffeomorphisms. In: Progress Probab., vol. 27. Birkhauser, Boston (1992)

Baxendale, P.H., Goukasian, L.: Lyapunov exponents for small perturbations of Hamiltonian systems. Ann. Probab. **30**(1), 101–134 (2002)

Brown, E., Holmes, P., Moehlis, J.: Globally coupled oscillator networks. In: Kaplan, E., Marsden, J.E., Sreenivasan, K.R. (eds.) Problems and Perspectives in Nonlinear Science: A Celebratory Volume in Honor of Lawrence Sirovich, pp. 183–215. Springer, New York (2003)

Bryant, H.L., Segundo, J.P.: Spike initiation by transmembrane current: a white-noise analysis. J. Physiol. **260**, 279–314 (1976)
Chow, C.C.: Phase-locking in weakly heterogeneous neuronal networks. Physica D **118**, 343–370 (1998)
Eckmann, J.-P., Ruelle, D.: Ergodic theory of chaos and strange attractors. Rev. Mod. Phys. **57**, 617–656 (1985)
Ermentrout, G.B.: $n : m$ phase locking of weakly coupled oscillators. J. Math. Biol. **12**, 327–342 (1981)
Ermentrout, G.B.: Neural networks as spatio-temporal pattern-forming systems. Rep. Prog. Phys. **61**, 353–430 (1991)
Ermentrout, G.B.: Type I membranes, phase resetting curves, and synchrony. Neural Comput. **8**, 979–1001 (1996)
Ermentrout, G.B., Kopell, N.: Multiple pulse interactions and averaging in coupled neural oscillators. J. Math. Biol. **29**, 195–217 (1991)
Gerstner, W., van Hemmen, J.L., Cowan, J.D.: What matters in neuronal locking?. Neural Comput. **8**(8), 1653–1676 (1996)
Goldobin, D., Pikovsky, A.: Synchronization and desynchronization of self-sustained oscillators by common noise. Phys. Rev. E **71**, 045201–045204 (2005)
Goldobin, D., Pikovsky, A.: Antireliability of noise-driven neurons. Phys. Rev. E **73**, 061906-1–061906-4 (2006)
Guckenheimer, J.: Isochronous and phaseless sets. J. Math. Biol. **1**, 259–273 (1975)
Guckenheimer, J., Holmes, P.: Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields. Springer, Berlin (1983)
Gutkin, B., Ermentrout, G.B., Rudolph, M.: Spike generating dynamics and the conditions for spike-time precision in cortical neurons. J. Comput. Neurosci. **15**, 91–103 (2003)
Hansel, D., Mato, G., Meunier, C.: Synchrony in excitatory neural networks. Neural Comput. **7**, 307–337 (1995)
Herz, A.V., Hopfield, J.J.: Earthquake cycles and neural reverberations: collective oscillations in systems with pulse-coupled threshold elements. Phys. Rev. Lett. **75**, 1222–1225 (1995)
Hoppensteadt, F.C., Izhikevich, E.M.: Weakly Connected Neural Networks. Springer, New York (1997)
Kifer, Yu.: Ergodic Theory of Random Transformations. Birkhauser, Boston (1986)
Kosmidis, E., Pakdaman, K.: Analysis of reliability in the Fitzhugh–Nagumo neuron model. J. Comput. Neurosci. **14**, 5–22 (2003)
Kozachenko, L.F., Leonenko, N.N.: Sample estimate of the entropy of a random vector. Probl. Inf. Transm. 23 (1987)
Kunita, H.: Stochastic flows and stochastic differential equations. Cambridge Studies in Advanced Mathematics, vol. 24. Cambridge University Press, Cambridge (1990)
Kuramoto, Y.: Phase- and center-manifold reductions for large populations of coupled oscillators with application to non-locally coupled systems. Int. J. Bifurc. Chaos **7**, 789–805 (1997)
Le Jan, Y.: Équilibre statistique pour les produits de difféomorphismes aléatoires indépendants. Ann. Inst. H. Poincaré Probab. Stat. **23**(1), 111–120 (1987)
Ledrappier, F., Young, L.-S.: Entropy formula for random transformations. Probab. Theory Relat. Fields **80**, 217–240 (1988)
Lin, K.K., Young, L.-S.: Shear-induced chaos. Nonlinearity **21**, 899–922 (2008)
Mainen, Z., Sejnowski, T.: Reliability of spike timing in neocortical neurons. Science **268**, 1503–1506 (1995)
Mattila, P.: Geometry of Sets and Measures in Euclidean Space. Cambridge University Press, Cambridge (1995)
Nakao, H., Arai, K., Nagai, K., Tsubo, Y., Kuramoto, Y.: Synchrony of limit-cycle oscillators induced by random external impulses. Phys. Rev. E **72**, 026220-1–026220-13 (2005)
Nualart, D.: The Malliavin Calculus and Related Topics. Springer, Berlin (2006)
Nunes, A.M., Pereira, J.V.: Phase-locking of two Andronov clocks with a general interaction. Phys. Lett. A **107**, 362–366 (1985)
Pakdaman, K., Mestivier, D.: External noise synchronizes forced oscillators. Phys. Rev. E **64**, 030901–030904 (2001)
Peskin, C.S.: Mathematical Aspects of Heart Physiology. Courant Institute of Mathematical Sciences, New York (1988)
Pikovsky, A., Rosenblum, M., Kurths, J.: Synchronization: A Universal Concept in Nonlinear Sciences. Cambridge University Press, Cambridge (2001)
Popovych, O.V., Maistrenko, Y.L., Tass, P.A.: Phase chaos in coupled oscillators. Phys. Rev. E **71**, 065201-1–065201-4 (2005)

Rinzel, J., Ermentrout, G.B.: Analysis of neural excitability and oscillations. In: Koch, C., Segev, I. (eds.) Methods in Neuronal Modeling, pp. 251–291. MIT Press, Cambridge (1998)
Ritt, J.: Evaluation of entrainment of a nonlinear neural oscillator to white noise. Phys. Rev. E **68**, 041915–041921 (2003)
Strogatz, S.: From Kuramoto to Crawford: Exploring the onset of synchronization in populations of coupled oscillators. Physica D **143**, 1–20 (2000)
Strogatz, S., Mirollo, R.: Synchronization of pulse-coupled biological oscillators. SIAM J. Appl. Math. **50**, 1645–1662 (1990)
Taylor, D., Holmes, P.: Simple models for excitable and oscillatory neural networks. J. Math. Biol. **37**, 419–446 (1998)
Teramae, J., Fukai, T.: Reliability of temporal coding on pulse-coupled networks of oscillators (2007). arXiv:0708.0862v1 [nlin.AO]
Teramae, J., Tanaka, D.: Robustness of the noise-induced phase synchronization in a general class of limit cycle oscillators. Phys. Rev. Lett. **93**, 204103–204106 (2004)
Victor, J.D.: Binless strategies for estimation of information from neural data. Phys. Rev. E 66 (2002)
Wang, Q., Young, L.-S.: Strange attractors with one direction of instability. Commun. Math. Phys. **218**, 1–97 (2001)
Wang, Q., Young, L.-S.: From invariant curves to strange attractors. Commun. Math. Phys. **225**, 275–304 (2002)
Wang, Q., Young, L.-S.: Strange attractors in periodically-kicked limit cycles and Hopf bifurcations. Commun. Math. Phys. **240**, 509–529 (2003)
Wang, Q., Young, L.-S.: Toward a theory of rank one attractors. Ann. Math. (2009, to appear)
Winfree, A.: Patterns of phase compromise in biological cycles. J. Math. Biol. **1**, 73–95 (1974)
Winfree, A.: The Geometry of Biological Time. Springer, New York (2001)
Young, L.-S.: Ergodic theory of differentiable dynamical systems. In: Real and Complex Dynamics. NATO ASI Series, pp. 293–336. Kluwer Academic, Dordrecht (1995)
Young, L.-S.: What are SRB measures, and which dynamical systems have them? J. Stat. Phys. **108**(5), 733–754 (2002)
Zaslavsky, G.: The simplest case of a strange attractor. Phys. Lett. A **69**(3), 145–147 (1978)
Zhou, C., Kurths, J.: Noise-induced synchronization and coherence resonance of a Hodgkin–Huxley model of thermally sensitive neurons. Chaos **13**, 401–409 (2003)