To appear in *Nonautonomous and Random Dynamical Systems in Life Sciences*, Springer Lecture Notes in Mathematical Biosciences, edited by P. Kloeden and C. Pötzsche.

Chapter 4
Stimulus-response reliability of biological networks

Kevin K. Lin

Abstract If a network of cells is repeatedly driven by the same sustained, complex signal, will it give the same response each time? A system whose response is reproducible across repeated trials is said to be *reliable*. Reliability is of interest in, e.g., computational neuroscience because the degree to which a neuronal network is reliable constrains its ability to encode information via precise temporal patterns of spikes. This chapter reviews a body of work aimed at discovering network conditions and dynamical mechanisms that can affect the reliability of a network. A number of results are surveyed here, including a general condition for reliability and studies of specific mechanisms for reliable and unreliable behavior in concrete models. This work relies on qualitative arguments using random dynamical systems theory, in combination with systematic numerical simulations.

Key words: Reliability, spike-time precision, coupled oscillators, random dynamical systems, neuronal networks, Lyapunov exponents, SRB measures

4.1 Introduction

If a network of neurons is repeatedly presented with the same complex signal, will its response be the same each time? A network for which the answer is affirmative is said to be *reliable*. This property is of interest in computational neuroscience because neurons communicate information via brief electrical impulses, or *spikes*, and the degree to which a system is reliable constrains its ability to transmit information via precise temporal patterns of spikes.

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

Thus, whether a given system is capable of reliable response can affect the mode and rate with which it transmits and processes information.

The reliability of *single neurons* has been well studied both experimentally and theoretically. In particular, *in vitro* experiments have found that single, synaptically isolated neurons are reliable under a broad range of conditions, *i.e.*, the spike times of an isolated neuron in response to repeated injections of a fixed, fluctuating current signal tend to be repeatable across multiple trials [6, 33, 16]. Theoretical studies have also found that models of isolated neurons tend to be reliable [53, 40, 42, 39, 20, 15, 16]. Less is known at the network or systems level; see, e.g., [4, 35, 7, 2, 32, 17] for some relevant experimental findings, and [41] for a theoretical treatment using a different approach.

This chapter reviews a body of work aimed at discovering network conditions and dynamical mechanisms that can affect the reliability of a network. The ergodic theory of random dynamical systems, i.e., the measure-theoretic analog of the theories surveyed in Chap. **??** and **??** of the present collection, plays a key role in this work: it provides a natural mathematical framework for precisely formulating the notion of reliability and providing tools that, in combination with numerical simulations, enable the analysis of concrete network models.

Much of the material and exposition here follow [26, 28, 27]. The first of these papers is concerned with mathematically-motivated questions, while the latter two concentrate on a more biological class of networks. These papers, as well as the present review, mainly focus on networks of oscillatory (i.e., tonically spiking) neurons. Networks of excitable neurons, which can behave rather differently, are the subject of a recent study [22] (see the discussion).

Relevance outside neuroscience. Reliability is a general dynamical property that is potentially relevant for a wide range of signal processing systems. Since biological systems, on scales ranging from single genes to entire organisms (and even populations), must respond to unpredictable environmental signals, it is possible that some of the mathematical framework and perhaps even the approaches and ideas outlined here may be of use in studying other types of biological information processing. Concepts analogous to reliability have also found use in areas far from biology, e.g., in engineered systems like coupled lasers [43] and in in molecular dynamics simulations [44].

The rest of this chapter is organized as follows: in Sect. 4.2, the concept of reliability is given a precise formulation, and some relevant results from random dynamical systems theory are reviewed. Sect. 4.3 presents a general condition that guarantees reliability, and Sect. 4.4 examines specific mechanisms for reliable and unreliable behavior in some concrete network models.

4.2 Problem statement and conceptual framework

This section describes a class of models which will be used throughout the rest of this chapter. The concept of reliability is given a precise formulation in this context, and relevant ideas from random dynamical systems theory are reviewed.

4.2.1 Model description and a formulation of reliability

To illustrate our ideas, we use networks of so-called "theta neurons" (see, e.g., [12]). These are idealized models of neurons that spike periodically at a fixed frequency ω in the absence of external forcing.[1] Single theta neurons have the form

$$\dot{\theta}(t) = \omega + z\big(\theta(t)\big)I(t) \tag{4.1}$$

Here, the state of the neuron is given by an angle $\theta \in S^1$, which is here mapped onto the interval $[0, 1]$ with endpoints identified; $\omega > 0$ is the intrinsic frequency of the neuron; $I(t)$ represents the sum of all the stimuli driving the neuron; and z is the *phase response curve* (PRC) of the neuron. The angle θ represents the fraction of the cycle that the neuron has completed; the neuron is viewed as generating a spike at $\theta = 0$. If $I(t) \equiv 0$, Eq. (4.1) is just the equation for a phase oscillator. A nonzero input $I(t)$ modulates the firing rate of the neuron, and the phase response z captures the state-dependent response of the neuron to stimuli.

Phase models like Eq. (4.1) are often used in biology to model rhythmic activity (see, e.g., [49, 12], and also Chap. **??**). In the context of neuroscience, the choice of PRC determines the response of the neuron model to stimuli, and a variety of PRCs are commonly used. In this chapter, $z(\theta)$ is taken to be $\frac{1}{2\pi}\big(1 - \cos(2\pi\theta)\big)$, which models so-called "Type I" neurons [11, 5]. This PRC has the property that it is positive when the neuron spikes, and for $\theta \approx 0$, $z(\theta) = O(\theta^2)$. The latter represents a form of *refractory effect:* at the moment when the neuron spikes, it is insensitive to its inputs, and is unable to generate a second spike immediately. This PRC is sometimes justified formally by truncating the normal form of neuron models near a saddle-node-on-invariant-circle bifurcation[2]; for our purposes, it mainly serves as convenient phenomenological model for neuronal response.

The class of models used in this chapter are networks of theta neurons. The network equations have the form

---

[1] Theta neurons can also model neurons operating in an *excitable* regime. The reliability of excitable theta neuron networks is studied in [22].

[2] See, e.g., [5], but note that phase truncations can sometimes miss important dynamical effects [25], and their use in biological modeling should be carefully justified.

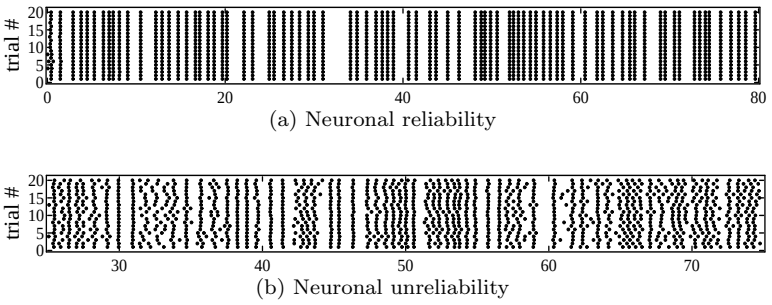

(a) Neuronal reliability

(b) Neuronal unreliability

Fig. 4.1 Raster plots showing the spike times of two neurons across 20 trials in response to a fixed stimulus. The neuron shown in (a) is chosen at random from a network with 100 neurons, with parameters chosen so that the network response is reliable across trials (as are single neuron responses). In (b), the neuron comes from an unreliable network. Figure adapted from [27].

$$\dot{\theta}_i = \omega_i + z(\theta_i) \left(\sum_{j \neq i} a_{ji} \; g(\theta_j) + I_i(t) \right) , \quad i = 1, 2, \cdots, N, \tag{4.2}$$

where the function $g : [0, 1] \to \mathbb{R}$ is an approximate delta function, i.e., it is a smooth function supported in a small interval $[-\delta, \delta]$ (here $\delta \sim 1/20$) satisfying $\int_0^1 g(\theta) \, d\theta = 1$; such pulse couplings are simple models for relatively fast synapses. The coupling matrix $A = (a_{ji})$ encodes the network structure; for simplicity we assume $a_{ii} = 0$ for all i, i.e., no self-loops. A number of different network architectures are considered in this review; these are specified along the way.

The stimuli $I_i(t)$ in Eq. (4.2) are modeled as white noise, i.e., $I_i(t)dt = \varepsilon_i \; dW_t^i$ where W_t^i denotes a standard Wiener process. This is an idealization of sustained, fluctuating signals, and has the convenient mathematical consequence that Eq. (4.2) is a (possibly quite large) system of stochastic differential equations (SDEs). *A priori,* the W_t^i for $i = 1, \cdots, N$ may be independent or correlated; for simplicity let us assume they are either independent or identical, allowing some neurons to receive the same input. Note that in Eq. (4.2), the stochastic forcing terms solely represent external stimuli driving the neuron, and not sources of neuronal or synaptic noise (but see the discussion at the end of the chapter).

These network models, though highly idealized, are broad enough to generate both reliable and unreliable network response without requiring careful tuning of parameters. That is, upon repeated trials with the same realizations of the $I_i(t)$ but different initial conditions, they can generate responses that are essentially the same across trials *(reliable)* or differ substantially across trials *(unreliable).* These behaviors are illustrated in Fig. 4.1.

A notion of neuronal reliability

What would it mean for a system of the form (4.2) to be reliable? Suppose we fix a single realization of the stimulus $(I_i(t))$ and drive Eq. (4.2) with the stimulus realization over a number of repeated "trials," with a new initial condition on each trial. The system (4.2) is said to be *neuronally reliable* (or simply "reliable") if

$$\lim_{t\to\infty} \mathrm{dist}(\Theta(\mathrm{t}), \Theta'(\mathrm{t})) = 0 \ , \tag{4.3}$$

where $\Theta(t) = (\theta_1(t), \cdots, \theta_N(t))$ denotes the state of the entire network at time t, and Θ and Θ' are two trajectories with initial states $\Theta(0) \neq \Theta'(0)$ [27]; it is assumed that $\Theta(0)$ is sampled independently from a fixed probability density ρ, say ρ =Lebesgue, on each trial. In other words, Eq. (4.3) means that given enough time, the entire network state is reproducible across repeated trials with random initial conditions.

As will be seen in Sect. 4.2.2, this notion of reliability can be naturally studied within the framework of random dynamical systems theory, making it a convenient mathematical definition. In biological terms, if a network is neuronally reliable, any network output that is a function of the network state will also be reproducible across repeated trials, so that neuronal reliability is in a sense the strongest form of reliability one might consider. Note, however, that there are other biologically relevant notions of reliability; some of these are mentioned in Sect. 4.2.3.

4.2.2 Relevant mathematical background

I begin by reviewing some relevant mathematical ideas [1, 3]; these can be viewed as ergodic-theoretic analogs of the theories reviewed in Chapter **??** by Kloeden and Pötzsche, and Chapter **??** by de Freitas and Sontag. There is, in particular, some overlap (both in overall goals and specific results) with the latter, though the perspective and emphasis here are different. The setting is a general SDE

$$dx_t = a(x_t)\ dt + \sum_{i=1}^{k} b_i(x_t) \circ dW_t^i \ , \tag{4.4}$$

where $x_t \in M$ with M a compact Riemannian manifold, and the W_t^i are independent standard Brownian motions. Clearly, Eq. (4.2) is a special case of Eq. (4.4): $x_t = (\theta_1(t), ..., \theta_N(t))$, $M = \mathbb{T}^N \equiv S^1 \times S^1 \times \cdots \times S^1$.

(To make sense of the theory outlined below on a general manifold M, Stratonovich calculus is necessary. But for $M = \mathbb{T}^N$ one can use either Itô or Stratonovich, and for simplicity Itô is used in Sect. 4.3 and beyond.)

Stochastic flows. In most physical applications involving SDEs, one fixes an initial x_0, and looks at the distribution of x_t for $t > 0$. These distributions evolve in time according to the Fokker-Planck equation, and under fairly general conditions converge to a unique stationary measure μ as $t \to \infty$. Since reliability is about a system's reaction to a single stimulus, *i.e.*, a single realization of the driving Wiener processes $(W_t^1, \cdots, W_t^N)$, at a time, and concerns the simultaneous evolution of all or large sets of initial conditions, of relevance to us are not the distributions of x_t but *flow-maps* $F_{t_1,t_2;\omega}$, where $t_1 < t_2$ 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 (4.4) 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 Eq. (4.4) are sufficiently smooth; see, *e.g.*, [21]. 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 (4.4) as compositions of random, IID, smooth maps. Many of the techniques for analyzing smooth deterministic systems have been extended to this random setting (see, e.g., Chap. **??** and **??**); the resulting body of results is collectively called "RDS theory" in this review.

The stationary measure μ, which gives the steady-state distribution averaged over all realizations ω of the driving Wiener processes, does not describe what we see when studying a system's reliability. Of relevance are the *sample measures* $\{\mu_\omega\}$, defined by

$$\mu_\omega = \lim_{t\to\infty} (F_{-t,0;\omega})_* \mu \tag{4.5}$$

where $(F_{-t,0;\omega})_*\mu$ denotes the push-forward of the stationary measure μ along the flow $F_{-t,0;\omega}$, and we think of ω as defined for all $t \in (-\infty, \infty)$ and not just for $t > 0$. (That the limit in Eq. (4.5) exists follows from a martingale convergence argument; see, e.g., [18].) One can view μ_ω as the conditional measures of μ given the past history of ω; it describes the distribution of states at $t = 0$ given that the system has experienced the input defined by ω for all $t < 0$. The family of measures $\{\mu_\omega\}$ 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; for this reason they are also sometimes called *random invariant measures.* Sample measures are measure-theoretic analogs of *pullback attractors* (see Chap. **??** and **??**), and are the distributions of *equilibria* (Chap. **??**).

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

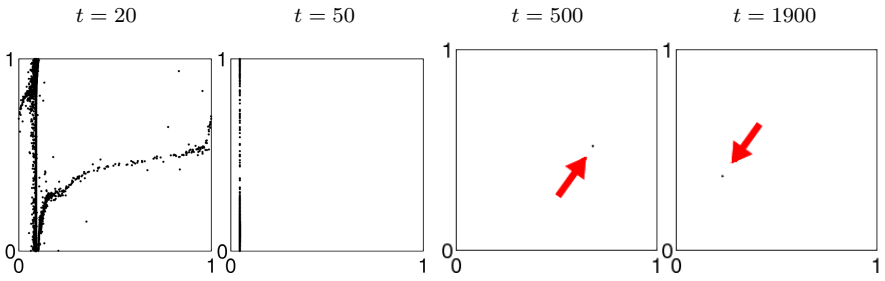

(a) Random fixed point ($\lambda_{\max} < 0$)

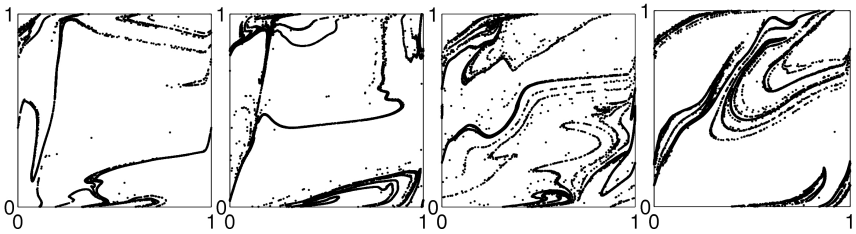

(b) Random strange attractor ($\lambda_{\max} > 0$)

Fig. 4.2 Temporal snapshots of sample measures for Eq. (4.2) with $N = 2$ oscillators driven by a single stimulus realization. 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. These figures (adapted from [26]) illustrate Theorem. 4.2.1.

Fig. 4.2 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, these distributions approximate $\mu_{\sigma_t(\omega)}$ for t sufficiently large. In these simulations, the initial distribution ρ is the stationary density of Eq. (4.2) with a small-amplitude noise, the interpretation being that the system is intrinsically noisy even in the absence of external stimuli; this distribution is then pushed forward in time using a fixed stimulus ω . 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.

The measure μ_ω gives the distribution of all possible states that a system may attain starting in a random state in the past and receiving a given stimulus for a sufficiently long time. Its structure is therefore of natural interest in reliability studies. Below we recall two mathematical results that pertain to the structure of μ_ω, specifically relating Lyapunov exponents to μ_ω .

Lyapunov exponents and sample distributions. For a fixed stimulus realization ω, any $x \in M$, and any nonzero tangent vector $v \in T_xM$, define the *Lyapunov exponent* [50]

$$\lambda_\omega(x,v) = \lim_{t\to\infty} \frac{1}{t} \log |DF_{0,t;\omega}(x)\cdot v| \tag{4.6}$$

when the limit exists. If μ is a stationary measure of the stochastic flow, then for almost every ω and μ-a.e. x, $\lambda_\omega(x,v)$ is well defined for all v. Moreover, if the invariant measure is ergodic, then $\lambda_\omega(x,v)$ is *non-random*, i.e., there exists a set $\{\lambda_1,\cdots,\lambda_r|\lambda_i\in\mathbb{R}\}$, $1\le r\le \dim(M)$, such that for a.e. ω and x and every v, $\lambda_\omega(x,v)=\lambda_i$ for some i . (We can have $r<\dim(M)$ because some of the λ_i may have multiplicity >1.) In what follows, we assume the invariant measure is indeed ergodic, and let $\lambda_{\max}=\max_i \lambda_i$.

As in deterministic dynamics, Lyapunov exponents measure the exponential rates of separation of nearby trajectories. In particular, a positive $\lambda_{\max}$ means the flow is sensitive to small variations in initial conditions, which is generally synonymous with chaotic behavior, while $\lambda_{\max}<0$ means nearby trajectories converge in time; in the deterministic context, this is usually associated with the presence of stable fixed points.

In Theorem. 4.2.1 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 4.2.1. *In the setting of Eq. (4.4), let μ be an ergodic stationary measure.*

(1) (Random sinks) [23] *If $\lambda_{\max}<0$, then with probability 1, μ_ω is supported on a finite set of points.*

(2) (Random strange attractors) [24] *If μ has a density and $\lambda_{\max}>0$, then with probability 1, μ_ω is a random Sinai-Ruelle-Bowen (SRB) measure.*

In Part (1) above, if in addition to $\lambda_{\max}<0$, two non-degeneracy conditions (on the relative motions of two points embedded in the stochastic flow) are assumed, then almost surely μ_ω is supported on a single point, and (as required in Eq. (4.3)) any pair of trajectories will almost surely converge in time [3]. Observe that this corresponds exactly to reliability for almost every ω as defined in Sect. 4.2.1, namely the collapse of trajectories starting from almost all initial conditions to a single, distinguished trajectory. This is the situation in Fig. 4.2(a). In view of this interpretation, we will equate $\lambda_{\max}<0$ with reliability in the rest of this chapter.

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 with respect to Lebesgue, and are concentrated on unstable manifolds, which are families of curves, surfaces etc., that wind around in a complicated way in the phase space [10, 51]. Part (2) of Theorem. 4.2.1 generalizes these ideas to random dynamical systems. Here, random (i.e., 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 acted on by the same stimulus may lead to very different outcomes at time t; this is true for all $t > 0$, however large. Note that in principle, a random strange attractor may still be supported in e.g. a small ball at all times, for the class of oscillator networks at hand we have not observed this to occur. It is, therefore, natural to regard $\lambda_{\max} > 0$ 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, the fact that $\lambda \leq 0$ 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 [18]. Jensen's inequality then gives

$$\int_{S^1} \log F'_{0,t;\omega}(x)\ dx \ \leq \ \log \int_{S^1} F'_{0,t;\omega}(x)\ dx \ = \ 0\ . \tag{4.7}$$

The equality above follows from the fact that $F_{0,t;\omega}$ is a circle diffeomorphism. Since the gap in the inequality in (4.7) 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 some light on the situation when the system is multi-dimensional:

Proposition 4.2.1. (see, *e.g.*, Ch. 5 of [18] or [23]) *In the setting of Eq. (4.4), assume μ has a density, and let $\lambda_1, \cdots, \lambda_d$ be the Lyapunov exponents of the system counted with multiplicity. Then*

(i) $\sum_i \lambda_i \leq 0$;

(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 $\lambda_i \neq 0$ for all i, then μ_ω is singular.

A formula giving the dimension of μ_ω is proved in [24] under mild additional conditions.

The reliability of a single oscillator, *i.e.* that $\lambda < 0$, is also easily deduced from Proposition. 4.2.1: μ has a density because the transition probabilities have densities, and no measure is preserved by all the $F_{s,t,\omega}$ because differ-

ent stimuli distort the phase space differently. Proposition. 4.2.1(i) and (ii) together imply that $\lambda < 0$. See also [39, 36, 42, 40].

For the 2-oscillator system illustrated in Fig. 4.2, assuming that μ has a density (this is straightforward to show; see Part I of [26]), Proposition. 4.2.1(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. 4.2.1(i) that $\lambda_2 < 0$, and by Proposition. 4.2.1(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 the stimulus history. Thus even in unreliable dynamics, the responses are highly structured and far from uniformly distributed, as illustrated in Fig. 4.2(b).

Note on numerical computation of Lyapunov exponents. As is usually the case for concrete models, $\lambda_{\max}$ for Eq. (4.2) can only be computed numerically. As in the deterministic context, the maximum Lyapunov exponent $\lambda_{\max}$ can be computed by solving the variational equation associated with the SDE. In the examples shown here, this is done using the Milstein scheme [19].

4.2.3 Reliability interpretations

RDS theory provides a useful framework for analyzing the reliability properties of specific systems. Before proceeding, however, let us discuss some issues related to the interpretation of the foregoing theory in reliability studies, as these interpretations are useful to keep in mind in what follows.

Lyapunov exponents and neuronal reliability

Advantages of Lyapunov exponents as a theoretical and numerical tool. A natural question is: since reliability measures across-trial variability, why does one not simply perform a number of trials using the same input, and compute e.g. cross-trial variances? Here are some reasons for using $\lambda_{\max}$ — without asserting that it is better for all circumstances:

(1) $\lambda_{\max}$ is a convenient summary statistic. Consider a network of size N. Should one carry out the above procedure for a single neuron, a subset of neurons, or for all N of them? Keeping track of N neurons is potentially computationally expensive, but other possibilities (e.g., using a subset of neurons or other small sets of observables) may involve arbitrary choices and/or auxiliary parameters. A virtue of using $\lambda_{\max}$ is that it is a single non-random quantity, depending only on system parameters. It thus sums up the stability property of a system in a compact way, without requiring any auxiliary, tunable parameters. When plotted as a function of system

parameters it enables us to view at a glance the entire landscape, and to identify emerging trends.

(2) Useful, well-understood mathematical properties. A second reason is that known mathematical properties of Lyapunov exponents can be leveraged. For example, under fairly general conditions, $\lambda_{\max}$ varies continuously, even smoothly, with parameters. This means that if $\lambda_{\max}$ is found to be very negative for a system, then it is likely to remain negative for a set of nearby parameters; the size of this parameter region can sometimes be estimated with knowledge of how fast $\lambda_{\max}$ is changing. Knowing that a system has zero across-trial variance alone will not yield this kind of information.

(3) Computational efficiency. Lyapunov exponents are defined in terms of infinitesimal perturbations. This means they are easily computed by simulating a single long trajectory, rather than requiring evolving an ensemble of trajectories. The latter can be quite expensive computationally.

(4) $\lambda_{\max}$ as a measure of relative reliability. Theorem. 4.2.1 justifies using the *sign* of the Lyapunov exponent as a way to detect reliability. Reliability, however, is often viewed in relative terms, *i.e.*, one might view some systems as being "more reliable" or "less reliable" than others. We claim that, *all else being equal,* the *magnitude* of $\lambda_{\max}$ carries some useful information. There are no theorems to cite here, but ideas underlying the results discussed above tell us that other things being equal, the more negative $\lambda_{\max}$, the stronger the tendencies of trajectories to coalesce, hence the more reliable the system. Conversely, for $\lambda_{\max} > 0$, the larger its magnitude, the greater the instability, which often translates into greater sensitivity to initial conditions. (But there are some caveats; see below.)

In practice, this can be rather useful because for reliable systems, a variance computation will yield 0, while the magnitude of $\lambda_{\max}$ indicates how reliable the system is. If one were to, say, model various sources of noise by adding random terms to Eq. (4.2) that vary from trial to trial, a system with a more negative $\lambda_{\max}$ is likely to be have greater tolerance for such system noise.

Limitations. Having explained some of the advantages, it is important to keep in mind that Lyapunov exponents have a number of serious limitations. The first is that $\lambda_{\max}$ *is a long-time average.* As was pointed out in Chap. **??**, Lyapunov exponents measure asymptotic stability. For one thing, this means $\lambda_{\max}$ may not reflect the initial response of the network upon presentation of the stimulus, which can be important as biological signal processing always occurs on finite timescales. In a bit more detail, for a system with $\lambda_{\max} < 0$, one can view initial transients as an "acquisition" period, during which the system has yet to "lock on" to the signal and after which the system can respond reliably. The magnitude of $\lambda_{\max}$ does not give direct information about this acquisition timescale.

Second, $\lambda_{\max}$ *reflects only net expansion in the fastest-expanding direction.* Because of the pulsatile nature of the interactions in our networks, the ac-

tion of the stochastic flow map on phase space is extremely uneven: at any one time, some degrees of freedom may be undergoing rapid change, while others evolve at a more modest pace. One therefore expects phase space expansion to occur only in certain directions at a time, and that expansion may coexist with strong phase space contraction. A positive $\lambda_{\max}$ only tells us that, on balance, expansion wins over contraction in *some* phase space direction. It does not give information about the relative degrees of expansion and contraction, nor does it tell us the directions in which expansion is taking place. In situations where one is interested in the response of specific "read-out" neurons in the network (see below), it may well be that phase space expansion occurs mainly in directions that do not significantly affect the reliability of these read-out neurons; $\lambda_{\max}$ will not be a useful indicator of relative reliability in that case. (See [26, 22] for examples.)

Moreover, as discussed below, in many applications one may be interested in other notions of reliability other than neuronal reliability. While $\lambda_{\max} < 0$ will typically ensure other forms of reliability, the magnitude of $\lambda_{\max}$ need not necessarily map onto other reliability measures in a one-to-one fashion.

Other notions of reliability

From the point of view of encoding information in the response of networks of biological neurons, neuronal reliability is very stringent: in a large biological network, where individual neurons and synapses may be rather noisy [13], it is clearly quite idealistic to insist that the detailed, microscopic dynamics of each neuron be reproducible across multiple trials. A less stringent way to formulate reliability is to project the dynamics onto a number $k \ll N$ of lower-dimensional signals, e.g., by applying a function Φ to the network state $\Theta(t)$, and to ask whether $\Phi(\Theta(t))$ is reproducible across trials. A neuronally reliable system is clearly guaranteed to be reliable for any choice of Φ, but a neuronally *unreliable* system may still produce reliable responses for some observables Φ . In [27], this idea was carried out in the form of "pooled-response reliability": a "pool," or subset C, of neurons is chosen, and their synaptic outputs are averaged; this average output defines a function Φ_C . A system is said to exhibit *pooled-response reliability* if the signal $\Phi_C(\Theta(t))$ is reproducible. In [27], it is found that (as one might suspect) neuronally unreliable systems often retain some degree of pooled-response reliability. Pooled-response reliability is, however, more difficult to work with, both computationally and mathematically.

There are also a number of other notions of spike-time reliability and precision in use in neuroscience that are distinct from neuronal reliability. Some of these focus on the reproducibility of *spike counts* within narrow time windows, while others focus on the variability of the spike times themselves. (See, e.g., [14] for a discussion.) A recent numerical study [22] has found that while $\lambda_{\max} > 0$ does imply unreliability as indicated by other reliability

measures, the quantitative behavior of these other measures are not always captured by trends in $\lambda_{\max}$.

4.3 Acyclic networks and modular decompositions

We now consider the neuronal reliability of *acyclic networks,* i.e., networks in which there is no feedback, and thus a well defined direction of information flow. We will show that acyclic networks are never unreliable. The proof technique will also suggest a broader class of networks that is more accessible to analysis, namely networks that admit a decomposition into modules with acyclic inter-module connections. The exposition here closely follows Part II of [26], with some details omitted.

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.

4.3.1 Skew-product representations

We first describe the connection graph that correspond to an acyclic network. Each node of this graph, $i \in \{1, \cdots, N\}$, corresponds to an oscillator. If oscillator i provides input to oscillator j, *i.e.*, if $a_{ij} \neq 0$, we assign a directed edge from node i to node j and write "$i \to j$". A *cycle* in such a directed graph is a sequence of nodes such that $i_1 \to i_2 \to \cdots i_k \to i_1$ for some k .

Definition 4.3.1. An oscillator network is *acyclic* if its connection graph has no cycles.

Given any two vertices i and j, let us write "$i \gtrsim j$" if there exists a path from i to j or if $i = j$. It is well known that if a graph is acyclic, then the relation $\gtrsim$ is a partial ordering, so that $i \gtrsim j$ and $j \gtrsim i$ if and only if $i = j$. In terms of information flow in an acyclic network, this means that for any pair of oscillators in an acyclic network, either they are "unrelated" (i.e., not comparable with respect to $\gtrsim$), or one is "upstream" from the other. Unrelated oscillators are not necessarily independent: they may receive input from the same source, for example. Acyclic networks can still be quite complex, with many branchings and recombinations; see Fig. 4.3 for an example.

Now let φ_t denote a flow on $\mathbb{T}^N$ with zero 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 relabeling the N oscillators, the following holds: for each $k = 1, \cdots, N$, 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

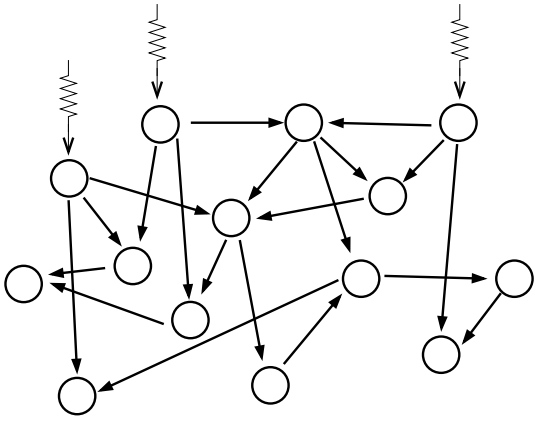

Fig. 4.3 Schematic of an acyclic network. Figure adapted from [26].

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

$$X^{(k+1)}(\theta_1, \cdots, \theta_{k+1}) = (X^{(k)}(\theta_1, \cdots, \theta_k), Y_{(\theta_1, \cdots, \theta_k)}(\theta_{k+1})) \tag{4.8}$$

where $\{Y_{(\theta_1, \cdots, \theta_k)}\}$ is a family of vector fields on S^1 parametrized by $(\theta_1, \cdots, \theta_k) \in \mathbb{T}^k$. In particular, $\varphi_t^{(N)} = \varphi_t$. In the system defined by (4.8), 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 4.3.1. *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.*

The proof is straightforward consequence of the partial ordering $\gtrsim$; see [26].

Next we generalize the notion of skew products 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 and only if oscillator j receives stimulus i. Clearly, since no arrow can terminate at a stimulus, a network driven by stimuli is acyclic if and only if the corresponding network without stimuli is acyclic.

Consider now 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. 4.2.2 can be represented as the skew-product on $\Omega \times S^1$ with

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

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

Proposition 4.3.2. *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 again straightforward; see [26].

4.3.2 Lyapunov exponents of acyclic networks

Consider a network of N oscillators driven by q independent stimuli. As before, let $\omega \in \Omega^q$ denote a q-tuple of Brownian paths, and let $F_{0,t;\omega}$ denote the corresponding stochastic flow on $\mathbb{T}^N$. Let $\lambda_\omega(x, v)$ denote the Lyapunov exponent defined in Eq. (4.6). The following is the main result of this section:

Theorem 4.3.1. *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 $\lambda_\omega(x, v) \leq 0$ 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 [37].

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 vectors $\{v_1, \cdots, v_N\}$ 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 $u, v \in \{v_j\}$, $u \neq v$,

$$\lim_{t\to\infty} \frac{1}{t} \log |\sin \angle(DF_{0,t;\omega}(x)u, DF_{0,t;\omega}(x)v)| = 0 \ . \tag{4.9}$$

That is, angles between vectors in a Lyapunov basis do not decrease exponentially fast; see *e.g.,* [50] for a more detailed exposition.

Proof. Since the network is acyclic, it factors into a hierarchy of skew-products. Supposing the oscillators are labeled so that $i < j$ means oscillator i is upstream from or unrelated to oscillator j, the kth of these is a stochastic flow $F^{(k)}_{0,t;\omega}$ on $\mathbb{T}^k$ describing the (driven) dynamics of the first k oscillators. 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. 4.3.1 holds for $F^{(k)}_{0,t;\omega}$.

First, 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 Eq. (4.7) in Sect. 4.2.2.

Now assume we have shown the conclusion of Theorem. 4.3.1 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 due to the skew-product structure, this direction is invariant under the variational flow $DF^{(k)}_{0,t;\omega}$. Starting with v_k, we complete a Lyapunov basis $\{v_1, \cdots, v_k\}$ at all typical points. Due to the invariance of the direction of v_k, we may once more use Jensen 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 $\pi : \mathbb{T}^k \to \mathbb{T}^{k-1}$ onto the first k coordinates, and note that

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

Due to Eq. (4.9), 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, $\lambda_\omega(x, v_i) \leq 0$. □

Remark 4.3.1. Some remarks concerning Theorem. 4.3.1:

1. Our conclusion of $\lambda_{\max} \leq 0$ falls short of reliability (which corresponds to $\lambda_{\max} < 0$). This is because our hypotheses allow for freely-rotating oscillators, *i.e.*, oscillators that are not driven by either a stimulus or another oscillator, and clearly $\lambda_{\max} = 0$ in that case. When no freely-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.
2. An analogous result in the context of uncertainty propagation was obtained by Varigonda *et. al.*; see [34].

4.3.3 Modular decompositions

Next, we describe how the ideas from the preceding section can be used to analyze the reliability of more general networks, by decomposing large networks into smaller subunits. Consider a graph with nodes $\{1, \cdots, N\}$, and let $\sim$ be an equivalence relation on $\{1, \cdots, N\}$. 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. 4.3.2:

Proposition 4.3.3. *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.*

Proposition. 4.3.3 has a natural interpretation in terms of network structure: observe that an equivalence relation on the nodes of a network partitions the nodes into distinct *modules*. Introducing directed edges between modules as above, we obtain what we call a *quotient network*. Assume this quotient network is acyclic, and let $M_1, M_2, \cdots, M_p$ be the names of the modules, ordered so that M_i is upstream from or unrelated to M_j for all $i < j$. 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 $F^{(s)}_{0,t,\omega}$ except when $s = s_i$ 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 $\lambda^{(1)}_1, \cdots, \lambda^{(1)}_{k_1}$ 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. If $\lambda^{(1)}_{\max} \equiv \max_j \lambda^{(1)}_j > 0$, we say *unreliability is produced within* M_1. We now wish to view M_1 as part of the larger network. To do so, for $i > 1$ let $\lambda^{(i)}_1, \cdots, \lambda^{(i)}_{k_i}$ denote the *fiber Lyapunov exponents*[3] 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 within the larger network* (but see the remark below).

The proof of the following result is virtually identical to that of Theorem. 4.3.1:

Proposition 4.3.4. *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 $\lambda^{(i)}_j, 1 \le i \le p, 1 \le j \le k_i$, are precisely the Lyapunov exponents of the network.*

Proposition. 4.3.4 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

[3] The fiber exponent can be defined exactly as in Eq. (4.6), but with the tangent vector v chosen to lie in the subspace tangent to each fiber; note these subspaces are invariant due to the skew product structure.

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.

Remark 4.3.2. Some comments on Proposition. 4.3.4:

1. The idea of "upstream" and "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 node as being downstream from another node (meaning the modules in which the nodes reside are so related).
2. Note that any network can be decomposed into modules connected by an acyclic graph, but the decomposition may be trivial, *i.e.,* the entire network may be a single module.[4] If the decomposition is nontrivial and $\lambda_{\max} > 0$ for the network, Proposition. 4.3.4 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.
3. It is important to understand that while fiber Lyapunov exponents let us assess the reliability of a module M as it operates within a larger network, i.e., as it responds to inputs from upstream modules and external stimuli, this is not the same as the reliability of M when it operates *in isolation,* i.e., when driven by external stimuli alone. Nevertheless, for many concrete examples, there is reason to think that the two types of reliability may be related. See Part II of [26] for details.
4. On a more practical level, the skew product structure implies that $DF_{0,t;\omega}$ is block-lower-triangular, and this fact together with Proposition. 4.3.4 give us a more efficient way to numerically compute Lyapunov exponents of networks with acyclic quotients.

Example. To illustrate the ideas above, consider the network in Fig. 4.4(a). Here, a single external input drives a network with 9 nodes. The network can be decomposed into modules connected by an acyclic graph, as shown in Fig. 4.4(b). Observe that Module A and Module C are both acyclic, and thus by Proposition. 4.3.4 they cannot generate unreliability. From Proposition. 4.3.4, it follows that whether the overall network is reliable hinges on the behavior of Module B. In [26], the reliability properties of 2-oscillator circuits like Module B are studied using ideas outlined in Sect. 4.4.1, and it is shown that one can indeed give qualitative predictions of the reliability of the network in Fig. 4.4 via modular decomposition.

[4] 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.

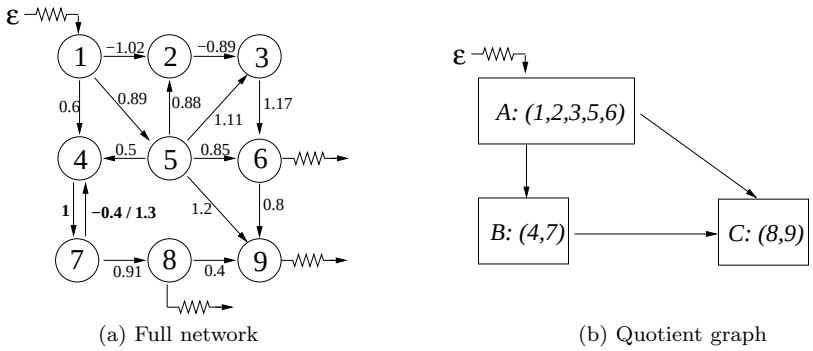

(a) Full network (b) Quotient graph

Fig. 4.4 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 randomly from $[0.95, 1.05]$. Panel (b) shows a modular decomposition. Figure adapted from [26].

4.4 Reliable and unreliable behavior in recurrent networks

Theorem. 4.3.1 highlights the importance of feedback in reliability studies. This section examines some examples of recurrent networks.

4.4.1 Unreliability in a two-oscillator circuit

To better understand what can occur in a system with feedback, we have studied the simplest circuit with recurrent connections, namely a two-oscillator system driven by a single stimulus:

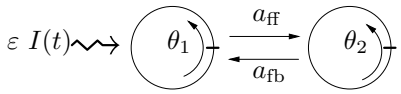

Eq. (4.2) here simplifies to

$$\begin{aligned} \dot{\theta}_1 &= \omega_1 + a_{\rm fb}\ z(\theta_1)\ g(\theta_2) + \varepsilon z(\theta_1)\ I(t)\ , \\ \dot{\theta}_2 &= \omega_2 + a_{\rm ff}\ z(\theta_2)\ g(\theta_1)\ , \end{aligned} \tag{4.10}$$

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

Fig. 4.5 shows the maximal Lyapunov exponent $\lambda_{\max}$ as a function of $a_{\rm ff}$ and $a_{\rm fb}$. In Fig. 4.5(a), the stimulus amplitude is $\varepsilon = 0.2$; in Fig. 4.5(b),

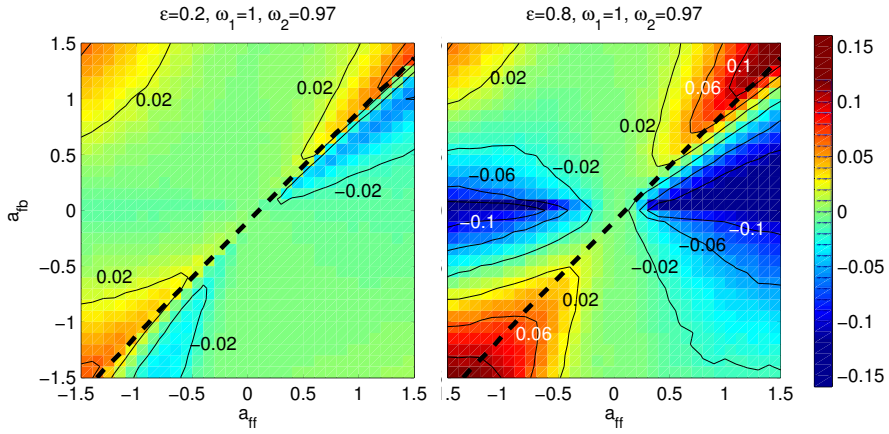

(a) $\varepsilon = 0.2$ (b) $\varepsilon = 0.8$

Fig. 4.5 Lyapunov exponent $\lambda_{\max}$ versus coupling strengths in the two-cell network. In all plots, we use $\omega_1 = 1$. The dashed curve shows the bifurcation curve $a^*_{\text{fb}}(a_{\text{ff}})$. Figure adapted from [26].

it is turned up to $\varepsilon = 0.8$. In both figures, it can be seen that there are large regions of both reliable and unreliable behavior. There is quite a bit of structure in both plots. For example, in Fig. 4.5(b), there is a region of strong reliability along the $a_{\text{fb}} = 0$ axis. This is exactly as expected, since Theorem. 4.3.1 guarantees that $\lambda_{\max} < 0$ when $a_{\text{fb}} = 0$ (there are no freely-rotating oscillators here), and since $\lambda_{\max}$ should depend continuously on a_{ff} and a_{fb}, we would expect it to remain negative for some range of a_{fb} . This "valley" of negative $\lambda_{\max}$ is also present in Fig. 4.5(a), but to a far lesser degree because ε is smaller, and the Jensen inequality argument given in Sect. 4.3.2 suggests that the greater ε is, the more negative $\lambda_{\max}$ should be.

A second, clearly visible structure occurs near the diagonal $\{a_{\text{ff}} = a_{\text{fb}}\}$ in Fig. 4.5(a): one can clearly see $\lambda_{\max} > 0$ on one side and $\lambda_{\max} < 0$ on the other; in Fig. 4.5(b), this structure has expanded and merged with the valley around $\{a_{\text{fb}} = 0\}$. To make sense of what is going on there, it is necessary to first discuss the *unforced dynamics* of the two-cell system. Observe that with $\varepsilon = 0$, Eq. (4.10) is just a deterministic flow on $\mathbb{T}^2$. It is straightforward to show that in the parameter regimes of interest, this 2D flow has no fixed points. One would thus expect essentially two types of behavior: either the flow has one or more limit cycles, or it is essentially quasiperiodic. In [26], an analysis of this ODE and an associated circle diffeomorphism (defined via a Poincaré section) shows that for $\varepsilon = 0$ and $\omega_1 > \omega_2$, the a_{ff}-a_{fb} space is dominated by a large region, roughly equal to $\{a_{\text{fb}} > a_{\text{ff}}\}$, over which the 2D flow is quasiperiodic. As a_{fb} decreases, a bifurcation occurs in which

the system acquires an attracting limit cycle. One can further prove that the corresponding critical value $a^*_{\rm fb}(a_{\rm ff})$ of $a_{\rm fb}$ occurs near $a_{\rm ff}$, and that for $a_{\rm fb} < a^*_{\rm fb}(a_{\rm ff})$, the behavior of the system is dominated by a large region with a single attracting limit cycle. That is to say, when $\varepsilon = 0$ and $a_{\rm fb} < a^*_{\rm fb}(a_{\rm ff})$, the two oscillators are phase-locked in a 1:1 resonance. The critical value $a^*_{\rm fb}(a_{\rm ff})$ can be numerically computed, and is shown as the dashed line in Fig. 4.5.

Since the noise amplitude in Fig. 4.5(a) is fairly small, the structure near the diagonal suggests that the onset of unreliability is connected with the onset of phase-locking in the unforced system. In [26], it was proposed that a dynamical mechanism called *shear-induced chaos* can explain this phenomenon, and a number of its predictions have been checked numerically there. Shear-induced chaos, a version of which was first studied numerically by Zaslavsky [52] and in a general, rigorous developed theory by Wang and Young [45, 46, 47, 48], is a general mechanism for producing chaotic behavior when a dissipative system meeting certain dynamical conditions is subjected to external forcing. We provide a brief summary below; see [25, 30, 29] and references therein for more details.

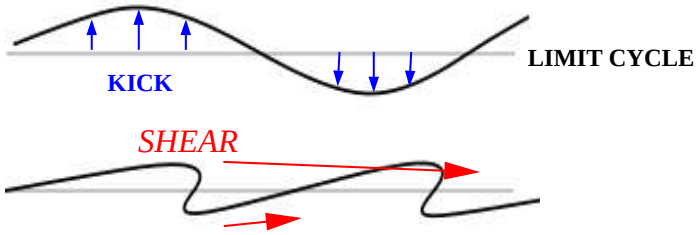

Fig. 4.6 The stretch-and-fold action of a kick followed by relaxation in the presence of shear.

Brief summary of shear-induced chaos. First, recall that in order to generate positive Lyapunov exponents in a dynamical system, it is necessary to have a way of stretching and folding phase space. Shear-induced chaos is a general mechanism for accomplishing this using two main ingredients:

(i) an attracting limit cycle where the surrounding flow exhibits *shear*, and

(ii) a source of external perturbation that forces trajectories off the limit cycle.

By shear, we mean a differential in the velocity as one moves transversally to the limit cycle, as illustrated in the bottom panel of Fig. 4.6. The perturbation in (ii) can take a variety of forms (deterministic or random), the simplest being a sequence of brief "kicks" applied periodically at a fixed time inter-

val; such periodic kicks can be modeled by applying a fixed (deterministic) mapping at a fixed interval.

Fig. 4.6 illustrates the basic geometric ideas: imagine a set of initial conditions along the limit cycle, and that at a certain time a single kick is applied, moving most of the initial conditions off the cycle. Because the limit cycle is attracting, the curve of trajectories will fall back toward the cycle as they evolve. But because of shear, the curve will be stretched and folded as the trajectories fall back toward the cycle. If this process is iterated periodically, it is easy to see that it can lead to the formation of Smale horseshoes, which is well known to be a source of complex dynamical behavior in deterministic dynamical systems. However, horseshoes can coexist with attracting fixed points, so that the associated chaotic behavior may only be transient, i.e., the Lyapunov exponents may still be < 0. In [45, 46, 47, 48], it is proved that for periodically-kicked dissipative oscillators, sustained chaotic behavior characterized by positive Lyapunov exponents, exponential decay of correlations, and the existence of SRB measures (see Sect. 4.2.2) are guaranteed whenever certain conditions are met.[5] In non-technical terms, the conditions are that the limit cycle possesses sufficient shear, that the damping is not too strong, and that the perturbations are sufficiently large and avoids certain "bad" phase space directions (associated with the "strong stable manifolds" of the limit cycle).

The theorems in [45, 46, 47, 48] apply to periodically-kicked oscillators, and the proof techniques do not carry over to the stochastic setting. Nonetheless, the underlying ideas suggest that the shear-induced chaos, as a general dynamical mechanism for producing instabilities, is valid for other types of forcings as well, including stochastic forcing. A systematic numerical study [29] has provided evidence supporting this view, and a recent analysis of a specific SDE with shear has found positive Lyapunov exponents [8].

As explained in detail elsewhere (see, e.g., [29]), shear-induced chaos makes a number of testable qualitative predictions. First, the more shear is present in the vicinity of a limit cycle, the more effective the stretching is, so that all else being equal, increasing shear would lead to a more positive Lyapunov exponent. Similarly, if the limit cycle were strongly attracting, any perturbations would be quickly damped out, reducing the amount of phase space stretching and decreasing the exponent. Finally, as mentioned above, the direction of kicking relative to the geometry of the flow is also important. If these conditions are met, then a system would have the tendency to generate chaotic behavior; the exact nature of the external perturbations (provided they are sufficiently strong) will affect quantitative details, but not gross qualitative features.

Returning now to the two-oscillator system, it has been shown (see Part I of [26]) that as the limit cycle emerges from the bifurcation at $a_{\rm fb} \sim a^*_{\rm fb}(a_{\rm ff})$,

[5] These results have been extended to certain nonlinear parabolic PDEs [31] and periodically-kicked homoclinic loops [38].

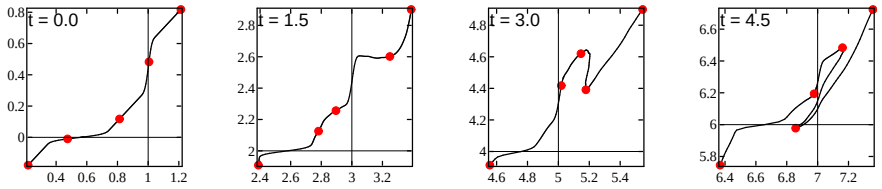

Fig. 4.7 Folding action caused by white noise forcing and shear near the limit cycle (with $a_{\text{fb}} > a^*_{\text{fb}}$). At $t = 0$, the curve shown is the lift of the limit cycle γ to $\mathbb{R}^2$. The remaining panels show lifts of the images $F_{0,t;\omega}(\gamma)$ at increasing times. The parameters are $\omega_1 = 1$, $\omega_2 = 1.05$, $a_{\text{ff}} = 1$, $a_{\text{fb}} = 1.2$, 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 $t \in [0, 5]$. Figure adapted from [26].

there is a great deal of shear in the vicinity of the limit cycle. Moreover, the forcing term in Eq. (4.10) is such that it can take advantage of the shear to produce stretching and folding. Fig. 4.7 illustrates how the folding and stretching occurs. These numerical results show that shear-induced stretching and folding do occur in the system (4.10).

While in this context, there are no theorems linking shear-induced folding with positive exponents, the various features seen around the diagonal in Fig. 4.5(a) can be readily explained using the ideas of shear-induced chaos. First, consider the phase-locked side of the a^*_{fb}-curve, i.e., $a_{\text{fb}} < a^*_{\text{fb}}$. Observe that as a_{fb} decreases, $\lambda_{\max}$ becomes more negative for some range of a_{fb}. This is consistent with increasing damping as the limit cycle (initially weak right after the bifurcation) becomes more strongly attracting. 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, and the value of $\lambda_{\max}$ is close to the value for the unforced flow (which for a limit cycle is always $\lambda_{\max} = 0$). On the other side of the a^*_{fb}-curve, where the system is essentially quasiperiodic, 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. 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 in Fig. 4.5(a) of $\lambda_{\max}$ from roughly 0 to definitively negative to positive as we cross the a^*_{fb}-curve from below consistent with the mechanism of shear-induced chaos.

In Fig. 4.5(b), where the stimulus amplitude is increased to $\varepsilon = 0.8$, the picture one obtains clearly continues some of the trends seen in Fig. 4.5(a): there are still regions of $\lambda_{\max} > 0$ near the diagonal, and a valley of $\lambda_{\max} < 0$ around $\{a_{\text{fb}} = 0\}$. But while the valley around $\{a_{\text{fb}} = 0\}$ can be explained on the basis of Theorem. 4.3.1 (which is valid regardless of the magnitude of ε),

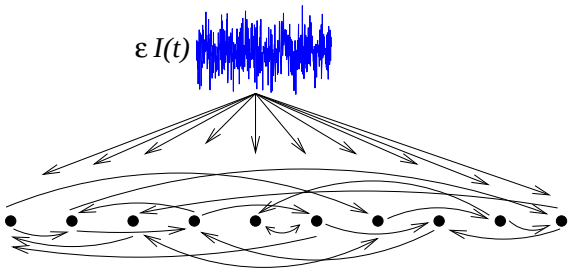

Fig. 4.8 Schematic of a single-layer network.

the behavior around the diagonal now likely involves more global effects as the system takes larger excursions from the limit cycle due to the increased forcing amplitude.

4.4.2 Single-layer networks

Lest the reader think that all recurrent networks are unreliable, for our last example we examine a large recurrent network that is strongly, robustly reliable. The network is shown schematically in Fig. 4.8. It is a sparsely-coupled recurrent network in which each neuron receives exactly κ inputs from other neurons (usually κ is 10 or 20% of the network size N). The intrinsic frequencies ω_i are drawn randomly from $[1-\rho, 1+\rho]$, and the nonzero coupling constants a_{ji} are drawn randomly from $[a(1-\rho), a(1+\rho)]$; the heterogeneity parameter ρ is taken to be 0.1. All oscillators receive the same external drive of amplitude ε . This architecture is motivated by layered models often found in neuroscience.

The Lyapunov exponent $\lambda_{\max}$ of such a single-layer network is plotted against $A := \kappa \cdot a$ in Fig. 4.9(a). Two values of ε are used. As can be seen, as ε increases, $\lambda_{\max}$ decreases. This is not unexpected: just as for the single oscillator in Sect. 4.2.2, we expect the magnitude of $\lambda_{\max}$ to increase with increasing ε (more on this below). Next, observe that $\lambda_{\max}$ is most negative when $A = 0$. This is also expected: with no coupling, the "network" is just a collection of uncoupled oscillators, each of which has $\lambda_{\max} < 0$ by Jensen's inequality; by continuity, this persists for a range of A. Finally, as $|A|$ increases, $\lambda_{\max}$ increases, suggesting that network interactions generally have a destabilizing effect in this network. However, even when A is quite large,[6] $\lambda_{\max}$ remains < 0 for $\varepsilon = 2.5$.

[6] A rough estimate shows that when $A = 2$, each kick should be sufficient to drive the oscillator roughly 1/3 of the way around its cycle.

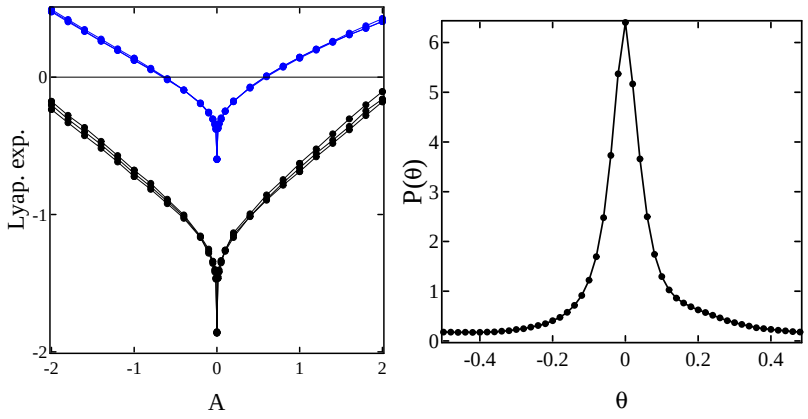

(a) Lyapunov exponent (b) Phase distribution

Fig. 4.9 Behavior of single-layer network. In (a), we plot the maximum Lyapunov exponent $\lambda_{\max}$ versus $A = \kappa \cdot a$; the parameters are $N = 100$, $\kappa = 10$, and $\varepsilon = 1.5$ (upper curve) and 2.5 (bottom). Panel (b) shows the phase distribution of neurons at the arrival of synaptic impulses for a single-layer network with $N = 100$. Figure adapted from [27].

In [27], the following qualitative explanation was proposed: consider first the system with $a = \rho = 0$. In this case, we have a collection of uncoupled, identical theta neurons. Since $\lambda_{\max} < 0$ for single theta neurons (Sect. 4.2.2), the ensemble will become entrained to the common input, and thus synchronize with each other. If we now allow slightly nonzero couplings and heterogeneous frequencies, i.e., $a \approx 0$ and $\rho \approx 0$, then by continuity we expect $\lambda_{\max}$ to still be < 0, and that the oscillators will remain nearly synchronized much of the time. Recall now that in Eq. (4.2), the phase response curve $z(\theta) = O(\theta^2)$ for $\theta \sim 0$. That is, around the time a neuron spikes, it is very insensitive to its inputs. Thus the near-synchrony of the neurons will lead to an attenuation in the *effective* strength of the coupling. This provides an explanation for why A can be made so large in Fig. 4.9 without causing $\lambda_{\max}$ to become positive.

This explanation also leads to a number of testable predictions. First, if we examine the phases of the oscillators when a spike arrives, we should observe a highly-clustered distribution. This has been checked numerically; an example is shown in Fig. 4.9(b). Second, anything that makes it harder for the neurons to synchronize should lead to an *increase* in $\lambda_{\max}$. For example, if we were to increase the amount of heterogeneity ρ in the system, $\lambda_{\max}$ should also increase. This is corroborated by the following results:

Heterogeneity ρ	0	0.01	0.1	0.3
$\lambda_{\max}$	-1.9	-1.7	-0.70	-0.18

As the degree of heterogeneity in the network increases, the oscillators become harder to entrain, and accordingly $\lambda_{\max}$ increases as well.

Concluding discussion

The work surveyed here has shown that the reliability of neuronal networks can be fruitfully formulated and studied within the framework of random dynamical systems. Particularly useful is the maximum Lyapunov exponent of a system, as a summary statistic for detecting reliability in numerical simulations. The results reviewed here show that:

1. *RDS theory, in particular the maximum Lyapunov exponent $\lambda_{\max}$ and results linking $\lambda_{\max}$ to the structure of sample measures, provide a useful framework for studying reliability.*
2. *Acyclic networks of pulse-coupled theta neurons are never unreliable.* This result highlights the importance of feedback in producing unreliability. The underlying ideas also generalize to the setting of modules connected by acyclic graphs, providing a way to analyze larger networks via modular decompositions.
3. *Recurrent networks can be reliable or unreliable.* In particular, a system as small as a two-oscillator circuit can become unreliable; moreover, there is evidence that unreliability in the two-cell circuit can be explained using the ideas of shear-induced chaos. At the same time, large recurrent networks can be robustly reliable as a result of (i) entrainment to a common input, and (ii) the strong refractory effect of the phase response of Type I theta neurons.

There are a number of additional issues relevant to neuroscience that have not been discussed here. We highlight three that are perhaps the most relevant from a biological point of view:

Noise. Neurons and synapses are well known to behave in a noisy fashion, both *in vivo* and *in vitro*. In the context of our model, one can represent the effects of noise by adding stochastic forcing terms that vary from trial to trial. While such a model no longer fits exactly in the framework of standard RDS theory (though the theory of RDS with inputs from Chap. **??** may be relevant), it can be studied numerically via direct measures of reliability such as cross-trial variance. In [27], the theoretical ideas surveyed here (Lyapunov exponents, random attractors) are used to carry out an analysis of the effects of noise on reliability, and to explain the different effects of correlated versus independent noise.

Pooled-response reliability. As mentioned in Sect. 4.2.2, even when $\lambda_{\max} > 0$ there can still be a great deal of structure in the sample distribution μ_ω , suggesting that by suitable projections or pooling of neuronal outputs, one can obtain responses that have some degree of reliability; this has been studied numerically in [27], and has received some attention in the experimental literature [9].

Beyond Lyapunov exponents. As previously mentioned, Lyapunov exponents do not always capture what one wants to know about neuronal response, and

other aspects of sample measures and random attractors may be more relevant in studies of network reliability. Exactly which aspects matter depends, of course, on the application at hand. Some steps in this direction have been taken in [22].

Acknowledgements

The work described in this review were supported in part by the Burroughs-Wellcome Fund (E. S.-B.) and the NSF (L.-S. Y. and K. K. L.).

References

1. L. Arnold. *Random Dynamical Systems.* Springer, New York, 2003.
2. W. Bair, E. Zohary and W.T. Newsome. Correlated Firing in Macaque Visual Area MT: Time Scales and Relationship to Behavior. *J. Neurosci.*, **21**(5):1676–1697, 2001.
3. P.H. Baxendale. Stability and equilibrium properties of stochastic flows of diffeomorphisms. In *Progr. Probab. 27.* Birkhauser, 1992.
4. M. Berry, D. Warland and M. Meister. The structure and precision of retinal spike trains. *PNAS*, **94**:5411–5416, 1997.
5. E. Brown, P. Holmes and J. Moehlis. Globally coupled oscillator networks. In E. Kaplan, J.E. Marsden and K.R. Sreenivasan, editors, *Problems and Perspectives in Nonlinear Science: A celebratory volume in honor of Lawrence Sirovich*, pages 183–215. Springer, New York, 2003.
6. H.L. Bryant and J.P. Segundo. Spike initiation by transmembrane current: a white-noise analysis. *Journal of Physiology*, **260**:279–314, 1976.
7. R. de Reuter van Steveninck, R. Lewen, S. Strong, R. Koberle and W. Bialek. Reproducibility and variability in neuronal spike trains. *Science* **275**:1805–1808, 1997.
8. R.E.L. Deville, N. Sri Namachchivaya and Z. Rapti, "Stability of a stochastic two-dimensional non-Hamiltonian system," *SIAM J. Appl. Math.*, to appear (2011)
9. A. S. Ecker, P. Berens, G. A. Keliris, M. Bethge, N. K. Logothetis and A. S. Tolias. Decorrelated neuronal firing in cortical microcircuits. *Science* **327**:584–587, 2010.
10. J.-P. Eckmann and D. Ruelle. Ergodic theory of chaos and strange attractors. *Rev. Mod. Phys.* **57**:617–656, 1985.
11. G.B. Ermentrout. Type I membranes, phase resetting curves and synchrony. *Neural Comp.* **8**:979–1001, 1996.
12. G.B. Ermentrout and D. Terman. *Foundations of Mathematical Neuroscience.* Springer, 2010.
13. A. Aldo Faisal, Luc P.J. Selen and Daniel M. Wolpert. Noise in the nervous system. *Nat. Rev. Neurosci.* **9**:292–303, 2008.
14. J.-M. Fellous, P.H.E. Tiesinga, P.J. Thomas and T.J. Sejnowski. Discovering Spike Patterns in Neuronal Responses. *J. Neurosci.* **24**:2989–3001, 2004.
15. D. Goldobin and A. Pikovsky. Antireliability of noise-driven neurons. *Phys. Rev. E* **73**:061906–1 –061906–4, 2006.
16. J. Hunter, J. Milton, P. Thomas and J. Cowan. Resonance effect for neural spike time reliability. *J. Neurophysiol.* **80**:1427–1438, 1998.
17. P. Kara, P. Reinagel and R.C. Reid. Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron* **27**:636–646, 2000.

18. Yu. Kifer. *Ergodic Theory of Random Transformations*. Birkhauser, 1986.
19. P.E. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*. Springer-Verlag, New York, 2011.
20. E. Kosmidis and K. Pakdaman. Analysis of reliability in the Fitzhugh–Nagumo neuron model. *J. Comp. Neurosci.* **14**:5–22, 2003.
21. Hiroshi Kunita. *Stochastic flows and stochastic differential equations*, volume 24 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1990.
22. G. Lajoie, K. K. Lin, and E. Shea-Brown. Chaos and reliability in balanced spiking networks with temporal drive. *Phys. Rev. E* **87** (2013)
23. Y. Le Jan. Équilibre statistique pour les produits de difféomorphismes aléatoires indépendants. *Ann. Inst. H. Poincaré Probab. Statist.* **23**(1):111–120, 1987.
24. F. Ledrappier and L.-S. Young. Entropy formula for random transformations. *Probab. Th. and Rel. Fields* **80**:217–240, 1988.
25. K.K. Lin, K.C.A. Wedgwood, S. Coombes and L.-S. Young. Limitations of perturbative techniques in the analysis of rhythms and oscillations. *J. Math. Biol.*, to appear, 2012.
26. K.K. Lin, E. Shea-Brown and L.-S. Young. Reliability of coupled oscillators. *J. Nonlin. Sci.*, **19**:497–545, 2009.
27. —, Spike-time reliability of layered neural oscillator networks. *J. Computat. Neurosci.* **27**:135–160, 2009.
28. —, Reliability of layered neural oscillator networks. *Commun. Math. Sci.* **7**:239–247, 2009.
29. K.K. Lin and L.-S. Young. Shear-induced chaos. *Nonlinearity* **21**:899–922, 2008.
30. —, Dynamics of periodically-kicked oscillators. *J. Fixed Point Theory and Appl.* **7**:291–312, 2010.
31. K. Lu, Q. Wang and L.-S. Young. Strange attractors for periodically forced parabolic equations. *Memoirs of the AMS*, to appear
32. T. Lu, L. Liang and X. Wang. Temporal and rate representations of time-varying signals in the auditory cortex of awake primates. *Nat. Neurosci.* **4**:1131–1138, 2001.
33. Z. Mainen and T. Sejnowski. Reliability of spike timing in neocortical neurons. *Science*, 268:1503–1506, 1995.
34. S. Varigonda, T. Kalmar-Nagy, B. LaBarre, I. Mezić, "Graph decomposition methods for uncertainty propagation in complex, nonlinear interconnected dynamical systems," *43rd IEEE Conference on Decision and Control* (2004)
35. G. Murphy and F. Rieke. Network variability limits stimulus-evoked spike timing precision in retinal ganglion cells. *Neuron* **52**:511–524, 2007.
36. H. Nakao, K. Arai, K. Nagai, Y. Tsubo and Y. Kuramoto. Synchrony of limit-cycle oscillators induced by random external impulses. *Physical Review E* **72**:026220–1 – 026220–13, 2005.
37. D. Nualart. *The Malliavin Calculus and Related Topics*. Springer-Verlag, 2006.
38. W. Ott and Q. Wang. Dissipative homoclinic loops of two-dimensional maps and strange attractors with one direction of instability. *Communications in Pure and Applied Mathematics* **64**:1439–1496, 2011.
39. K. Pakdaman and D. Mestivier. External noise synchronizes forced oscillators. *Phys. Rev. E* **64**:030901–030904, 2001.
40. J. Ritt. Evaluation of entrainment of a nonlinear neural oscillator to white noise. *Phys. Rev. E* **68**:041915–041921, 2003.
41. J. Teramae and T. Fukai. Reliability of temporal coding on pulse-coupled networks of oscillators. *arXiv:0708.0862v1* [nlin.AO], 2007.
42. J. Teramae and D. Tanaka. Robustness of the noise-induced phase synchronization in a general class of limit cycle oscillators. *Phys. Rev. Lett.* **93**:204103–204106, 2004.
43. A. Uchida, R. McAllister and R. Roy. Consistency of nonlinear system response to complex drive signals. *Phys. Rev. Lett.* **93**:244102, 2004.

44. B.P. Uberuaga, M. Anghel and A.F. Voter. Synchronization of trajectories in canonical molecular-dynamics simulations: observation, explanation, and exploitation. *J. Chem. Phys.* **120**:6363–6374, 2004.
45. Q. Wang and L.-S. Young, Strange attractors with one direction of instability. *Comm. Math. Phys.* **218**:1–97, 2001.
46. —, From invariant curves to strange attractors. *Comm. Math. Phys.* **225**:275–304, 2002.
47. —, Strange attractors in periodically-kicked limit cycles and Hopf bifurcations. *Comm. Math. Phys.* **240**:509–529, 2003.
48. —, Toward a theory of rank one attractors. *Annals of Mathematics*, to appear.
49. A. Winfree. *The Geometry of Biological Time*. Springer, New York, 2001.
50. L.-S. Young. Ergodic theory of differentiable dynamical systems. In *Real and Complex Dynamics*, NATO ASI Series, pages 293–336. Kluwer Academic Publishers, 1995.
51. —, What are SRB measures, and which dynamical systems have them? *Journal of Statistical Physics*, **108**(5):733–754, 2002.
52. G. Zaslavsky. The simplest case of a strange attractor. *Phys. Lett.* **69A**(3):145–147, 1978.
53. C. Zhou and J. Kurths. Noise-induced synchronization and coherence resonance of a Hodgkin-Huxley model of thermally sensitive neurons. *Chaos* **13**:401–409, 2003.