

Reliability of Spike Timing Is a General Property of Spiking Model Neurons

Romain Brette

brette@ccr.jussieu.fr

Centre de Mathématiques et de Leurs Applications, Ecole Normale Supérieure de Cachan, 94230 Cachan, France, and INSERM U483, Université Pierre et Marie Curie, 75005 Paris, France

Emmanuel Guigon

INSERM U483, Université Pierre et Marie Curie, 75005 Paris, France

The responses of neurons to time-varying injected currents are reproducible on a trial-by-trial basis *in vitro*, but when a constant current is injected, small variances in interspike intervals across trials add up, eventually leading to a high variance in spike timing. It is unclear whether this difference is due to the nature of the input currents or the intrinsic properties of the neurons. Neuron responses can fail to be reproducible in two ways: dynamical noise can accumulate over time and lead to a desynchronization over trials, or several stable responses can exist, depending on the initial condition. Here we show, through simulations and theoretical considerations, that for a general class of spiking neuron models, which includes, in particular, the leaky integrate-and-fire model as well as nonlinear spiking models, aperiodic currents, contrary to periodic currents, induce reproducible responses, which are stable under noise, change in initial conditions and deterministic perturbations of the input. We provide a theoretical explanation for aperiodic currents that cross the threshold.

1 Introduction

The responses of neurons to dynamic stimuli have been shown to exhibit high reliability *in vitro* (Mainen & Sejnowski, 1995; Hunter, Milton, Thomas, & Cowan, 1998; Fellous et al., 2001; Beierholm, Nielsen, Ryge, Alstrom, & Kiehn, 2001) and *in vivo* (Berry, Warland, & Meister, 1997; Nowak, Sanchez-Vives, & McCormick, 1997; Reich, Victor, Knight, Ozaki, & Kaplan, 1997; Berry & Meister, 1998; Buracas, Zador, DeWeese, & Albright, 1998; Bair & Koch, 1996). In this case, spike timing is reproducible on a trial-by-trial basis up to a precision of 1 ms or less, even a long time after stimulus onset (1 s in Mainen & Sejnowski, 1995). Consequently, in a variety of neurons and for time-varying stimuli, spike times convey more information about the stimulus than spike count alone does (Victor & Purpura, 1996;




Figure 1: Three cases regarding the reproducibility of neuron responses to a given stimulus. (A) Small variances in ISIs accumulate, so that after some time, spikes are not synchronized between trials. (B) Spike timing is reliable even in the long run but depends on the initial state of the neuron—here, two stable solutions coexist. (C) Spike timing is reliable; noise does not accumulate over time.

de Ruyter van Steveninck, Lewen, Strong, Koberle, & Bialek, 1997; Buračas et al., 1998; Reich, Mechler, Purpura, & Victor, 2000; Reinagel & Reid, 2000). Rate-modulated Poisson processes do not account for these properties (Victor & Purpura, 1996; Baddeley et al., 1997; de Ruyter van Steveninck et al., 1997; Berry et al., 1997; Nowak et al., 1997; Reich, Victor, & Knight, 1998; Kara, Reinagel, & Reid, 2000; Reich et al., 2000). On the other hand, stationary stimuli elicit spike trains with very coarse precision (Mainen & Sejnowski, 1995; Berry et al., 1997; de Ruyter van Steveninck et al., 1997; Buračas et al., 1998). Mainen and Sejnowski (1995) injected a constant current into a neuron and compared the output spike trains over repeated trials. At the beginning, spikes occurred at the same time in all trials, but then small variances in interspike intervals (ISIs) began to accumulate, so that spikes eventually were totally desynchronized over trials. It is not clear why this does not happen with time-varying currents.

What happens when the responses of a neuron to a given stimulus are compared over repeated trials, knowing that each trial is different due to intrinsic noise? Three cases may occur, as illustrated in Figure 1:

Case 1: Small variances in ISIs accumulate, so that eventually spikes tend to desynchronize over trials.

Case 2: Spike times are reproducible in the long run but depend on the initial state of the neuron (e.g., on the membrane potential at the onset of the stimulus); that is, several stable responses exist.

Case 3: Spike times are reproducible in the long run and do not depend on the initial state of the neuron.

Experimentally, case 1 has been shown to occur for constant input currents, whereas time-varying stimuli induce reproducible responses (Mainen & Sejnowski, 1995), although the experimental procedures used in these studies do not allow us to distinguish between cases 2 and 3. Only in the last case can the responses be considered truly reproducible.

In this article, we show that for a general class of spiking neuron models, which includes, in particular, the leaky integrate-and-fire model (Lapicque, 1907; Knight, 1972) as well as nonlinear spiking models, all three cases can occur if the input current is periodic, while aperiodic currents induce reproducible responses. In addition to numerical simulations, we put forth a theoretical explanation of this property for aperiodic currents that oscillate around threshold. The conditions required for our explanation are not fulfilled by the nonleaky integrate-and-fire model—also called perfect integrator—which is never reliable.

2 Materials and Methods

2.1 What Is Reliability? A neuron is said to be reliable when its responses to a given stimulus are reproducible. Thus, reliability is a priori stimulus dependent. The experimental results of Mainen and Sejnowski (1995) show that with a constant current, the responses of a neuron in different trials are initially synchronous, but then small variances in ISIs add up and the spikes become desynchronized. If the intrinsic noise of neurons were lower, it would simply take the spikes longer to desynchronize, but in the end, the result would be the same. However, with nonconstant currents resembling synaptic inputs (low-pass filtered gaussian noise), desynchronization does not happen. In both cases, intrinsic noise is low at the timescale of a single ISI, but it eventually leads to a desynchronization only in the case of a constant current.

Thus, we propose the following definition of reproducibility: responses to a stimulus are considered reproducible if the spike time jitter remains of the order of the intrinsic noise, however long we wait, or, equivalently, that any given precision can be reached if the level of noise is low enough. This definition allows us to distinguish the two cases mentioned above. However, neuron responses can also fail to be reproducible if spike timing depends on the initial condition, that is, the membrane potential at stimulus onset. Therefore, we shall consider in the definition the asymptotic spike time jitter induced not only by intrinsic noise but also by the dispersion of initial conditions.

Our aim is to show that depending on the stimulus, all of the three cases mentioned in section 1 can occur. Furthermore, we shall show that a broad class of models has reproducible responses to aperiodic stimuli. We pay particular attention to leaky models—in a general sense, as defined below. However, we also introduce and study other types of models—the perfect integrator and nonleaky models—to show that the nonreliability of the perfect integrator is the exception rather than the rule.

2.2 Models.

2.2.1 The General Model. We consider spiking neuron models, where the membrane potential V evolves according to the following one-dimensional differential equation,

$$\frac{dV}{dt} = f(V, t), \quad (2.1)$$

with the spike modeled as follows: when the potential reaches a threshold V_t , it is reset to V_r . This general definition includes many different models, such as the leaky integrator (Lapicque, 1907; Knight, 1972) and the perfect integrator (Knight, 1972). The equation needs not be linear.

Such models are not always reliable, as the example of the perfect integrator shows. However we shall show that a very broad class of models of this kind exhibits reproducible responses.

2.2.2 Leaky Models. One example of such a model is the leaky integrator (Lapicque, 1907; Knight, 1972), where the membrane potential V evolves according to the following equation:

$$\tau \frac{dV}{dt} = -V + RI(t), \quad (2.2)$$

where τ is the membrane time constant, R is the resistance, and $I(t)$ the input current.

We generalize the notion of leak by defining a leaky model as a model of type 2.1 satisfying the following condition:

$$\frac{\partial f}{\partial V} < 0.$$

This includes the case when τ and R vary in time and also includes nonlinear equations. Basic mathematical properties of leaky models are summarized in the article appendix.

In our simulations, we used a leaky integrator with $\tau = 33$ ms, $R = 200$ M Ω , $V_t = 15$ mV, and $V_r = -5$ mV, and a nonlinear leaky model defined by

$$\tau \frac{dV}{dt} = -aV^3 + RI(t), \quad (2.3)$$

with the same values for τ , R , V_t , and V_r , and $a = 4444 \text{ V}^{-2}$ (this is for scaling reasons). This corresponds to a neuron model with rectification, that is, with a voltage-dependent leak. We chose these parameter values so as to obtain figures of the same magnitude as in experimental results. However, the actual value of parameters has no influence on the results, since two sets of parameters actually define the same model (change of variables).

2.2.3 Nonleaky Models. We shall show that the leak is not necessary to obtain reliable spike timing. We simulated an example of a nonleaky model, described by the following equation,

$$\tau \frac{dV}{dt} = VI(t) + k, \quad (2.4)$$

where $I(t)$ is the input current, with $V_t = 1$, $V_r = 0$, $\tau = 33 \text{ ms}$, and $k = 1$. This model is leaky only if $I(t) < 0$ for all t . Otherwise, the solutions of equation 2.4 diverge exponentially in all intervals where $I(t) > 0$. However, we shall see that it does not prevent the model from having reliable spike timing.

Equation 2.4 does not correspond to a neuron model; we chose it only to illustrate our explanation of reliability.

2.2.4 Perfect Integrator. As an example of a nonreliable model, we have taken the common nonleaky integrate-and-fire model,

$$\tau \frac{dV}{dt} = f(t), \quad (2.5)$$

where $f(t)$ is an input function. We used $\tau = 12.5 \text{ ms}$ (in order to have comparable timescales), $V_t = 1$, and $V_r = 0$. We shall see in the simulations that this model is never reliable. Mathematically, the unreliability of this model results from the fact that the distance between two solutions is constant modulo 1, and thus dynamical noise always accumulates (Knight, 1972).

2.3 Stimuli. We tested the reproducibility of model responses to various periodic and aperiodic currents. We also assessed the influence of current parameters, such as amplitude and mean, on spike timing. For this purpose, we used the following generic method: we define a basis function $B(t)$ as one with values between -1 and 1 , such as a sinusoidal wave, and then derive a family of currents by a parameterized affine change of variables, for example, $I(t) = I_{th} + 0.5 + pB(t)$ with $p \in [0, 1]$, where I_{th} is the input current threshold (the minimum constant current eliciting a response). In this example, we obtain a family of currents with a fixed DC part and whose amplitude varies from 0 to 1.

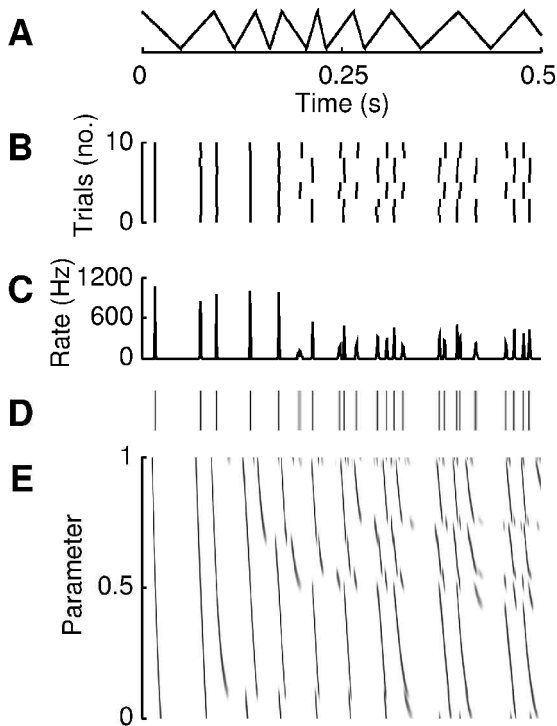


Figure 2: Construction of a 2D spike density figure. (A) Basis function. A family of currents is defined with a parameter controlling the current mean. (B) Trials for a given parameter. Every short vertical line represents a spike. (C) PSTH computed from 2000 trials with the parameter used in B. (D) The PSTH is converted into gray rasters, with the higher the firing rate, the darker the lines. (E) Gray rasters for all parameters are collected and arranged vertically.

In most simulations, we chose the change of variables so that the current is always above threshold for $p \in [0, 0.5]$ and oscillates around it for $p \in [0.5, 1]$.

Subsequently, we want to display the reliability of model responses for all the parameters of a given family of currents on a single figure. How this is done is illustrated by Figure 2. For a given basis function (see Figure 2A), we compute the outputs of a large number of trials for each parameter value (see Figure 2B) and then compute a poststimulus time histogram (PSTH) using small time bins (see Figure 2C), which we convert into a line of gray rasters (see Figure 2D). Arranging these lines vertically, we get a two-dimensional spike density figure (see Figure 2E), where histograms for all parameter values between 0 and 1 are represented simultaneously, with parameter on

the vertical axis and time on the horizontal axis. This allows us to visualize not only the reliability of spike timing for every parameter value, but also the effect of parameters on spike timing and reliability. Thus, we assess the stability of spike timing under noise as well as under deterministic perturbations.

For periodic currents, we used sine waves, and for aperiodic currents, we used two types of basis functions: (1) a triangle-shaped function, with the slope of each line chosen according to a random uniform distribution between two values, and (2) a noisier function: a low-pass filtered gaussian noise that we distorted so that its magnitude distribution over time was uniform between -1 and 1 (see, e.g., Figure 4.2.2). This was done to distinguish between the input current always above threshold from the input current oscillating around it. The specific shape of input currents did not influence the results.

2.4 Measure of Precision. Precision of spike timing was assessed from PSTHs. Most previous measures relied on the identification of “events” (Mainen & Sejnowski, 1995), defined as peaks in the PSTH, whose width is an estimate of the precision. However, assessing the precision only in events might lead to an overestimate, as observed in Tiesinga, Fellous, & Sejnowski (2002a), since the periods when the PSTH is not peaked enough are discarded. Other measures are based on the variance of the PSTH about the mean output frequency (Hunter et al., 1998), but this measure is misleading when the output frequency varies over time.

The measure we propose here is based on the entropy of distributions (Borst & Theunissen, 1999). We cannot directly compute the entropy from a PSTH, because this is not a normalized probability distribution. We choose the successive ISIs of one trial to divide the PSTH in successive windows. The models we use have the property that for every other trial, the neuron spikes exactly once in each of these windows. Equivalently, we can split the PSTH in successive windows containing 2000 spikes (the number of trials). We can evaluate the precision of spike timing in any window from the PSTH using the following formula,

$$H = - \sum_i p(i) \log_2 p(i),$$

where $p(i)$ is the number of spikes in time bin i divided by the number of trials. H is the entropy of the distribution of spike times in the window and represents the number of bits required on average to express in which time bin a neuron spikes. One advantage over event-based measures is that if there are m stable solutions, each PSTH window contains m peaks, which gives a higher entropy. However, this quantity depends on the size of the time bin used for the PSTH. We can translate it into a time measure that does not depend on the size of the time bin in the following way. An entropy of H bits is the entropy of a uniform distribution over 2^H time bins, that is, over

an interval of size $\delta t 2^H$, where δt is the length of the time bin, which gives the following formula:

$$\delta t \exp \left(- \sum_i p(i) \log p(i) \right). \quad (2.6)$$

As δt tends to 0, this quantity tends to

$$\exp \left(- \int p(t) \log p(t) dt \right),$$

where $p(t)$ represents the spike time density. Thus, if δt is small enough, the precision measure given by formula 2.6 does not depend on the length of the time bin. As an example, this measure gives a precision of Δ for a uniform distribution over an interval $[t - \Delta/2, t + \Delta/2]$ and a precision of 4σ for a gaussian distribution with standard deviation σ .

We used formula 2.6 to compute the precision in each window and averaged over all successive windows, discarding the first 2.5 s of the 10 s runs (to take the time of convergence into account).

In the results, we shall say that neuron responses have a high precision if their measure of precision given by formula 2.6 is small.

2.5 Details of Simulations. Input currents were sampled at a rate depending on the required resolution—2000 Hz in most simulations with the leaky integrator. Equations for the linear leaky integrator and the perfect integrator were solved by exact integration. The equations of the two other models—the nonlinear and the nonleaky ones—were integrated by the Euler method (these were simulated mostly for illustration, and only on a short timescale). Spike timing was computed by bisection with a precision of $1 \mu\text{s}$ using the following observation: over a single current sampling step ($1/2000$ s), equation 2.1 reduces to an autonomous scalar equation because the input current is constant, which means that $V(t)$ is monotonous over this duration, and the time when the potential hits threshold can then be computed unambiguously and quickly by bisection.

A small noise was added to the potential at each time step of integration. The type of noise used is irrelevant because it tends to a gaussian noise as the time step tends to 0. Therefore, for computational ease, we used a uniform noise with an amplitude ranging from $2.5 \text{ mV}\cdot\text{s}^{-1/2}$ to $3.5 \text{ mV}\cdot\text{s}^{-1/2}$. Reliability of spike timing was assessed by 2000 repeated trials for every tested model and input current over an extended period of time—10 s for the leaky integrator—and PSTHs were computed with a 0.5 ms time bin. Doing many trials avoided the problem of smoothing PSTHs.

3 Theoretical Predictions

3.1 Periodic Currents. Responses to periodic currents highlight the two difficulties that may prevent the neuron responses from being reproducible: (1) a small noise may accumulate over time, or (2) several stable responses may coexist.

Because periodically driven leaky models are related to the dynamical theory of homeomorphisms of the circle (Denjoy, 1932; Coddington & Levinson, 1955; Arnold, 1961; Herman, 1977; Keener, 1980; Veerman, 1989), their mathematical properties are well known (Knight, 1972; Keener, Hoppensteadt, & Rinzel, 1981). However, the reliability of periodically driven neuron models is a recent issue (Tiesinga, 2002; Tiesinga, Fellous, & Sejnowski, 2002b). We describe the available theoretical results and explain how they relate to our problem.

The dynamics of leaky models depend on the value of the rotation number, defined as the ratio of the input frequency to the output spike rate:

- It is rational if and only if the output pattern of spikes is (asymptotically) periodic, the period being a multiple of the input period. When this pattern is stable under perturbations, this is called phase locking, which has also been studied experimentally (Rescigno, Stein, Purple, & Poppele, 1970; Ascoli, Barbi, Chillemi, & Petracchi, 1977; Guttman, Feldman, & Jakobsson, 1980; Koppl, 1997). There is $p : q$ phase locking when a stable pattern of p spikes is produced every q periods of the input. Thus when the model is phase-locked, a small noise does not accumulate over time. However, several stable responses may coexist. As noted in Tiesinga (2002), if a $p : q$ phase-locked solution is shifted by a multiple of the input period, a new solution is obtained, so that at least q distinct stable solutions exist. Thus, regarding reliability, cases 2 and 3 can appear.
- When it is irrational, under some regularity conditions, the dynamics of the model are topologically equivalent to the dynamics of a model with constant input, which means that a small noise always accumulates over time, however small it may be. Thus, case 1 occurs.

When the input current is parameterized, for example, by its frequency and mean, the rotation number is a continuous function of the parameters, which implies that cases 1, 2, and 3 occur in any area of the parameter space where the rotation number is not constant. Thus, there is no systematic relationship between current parameters and reliability. However, while the set of parameters where case 1 occurs has positive measure if the current is always above threshold (Herman, 1977), it has measure zero, and thus is not observable, if the current oscillates around threshold (Keener, 1980; see Veerman, 1989, for a rigorous mathematical proof). To sum up, cases 2 and 3 occur in any area where the current oscillates around threshold; cases 1,




Figure 3: A current below threshold results in an interval where no spiking is possible. Arrows represent whether the vector field points up or down at threshold.

2, and 3 occur in any area where the current is above threshold (unless the rotation number is constant). The (asymptotic) spike time jitter is an increasing function of the level of intrinsic noise, which tends to 0 as the noise variance tends to 0 when the model is phase-locked and to a positive number when it is not. Quantitatively, this function depends on various detailed aspects of the dynamics, such as the Lyapunov exponent (Tiesinga, 2002).

3.2 Aperiodic Currents. In the case of aperiodic currents, such as filtered gaussian noise, we will see in the numerical results that neither case 1 nor case 2 arises. Other studies suggest that the responses are reliable when the current is above threshold (Pakdaman & Tanabe, 2001). Here, we explain why this should also be so when the input current oscillates around threshold.

Our explanation relies on the construction of a set of possible spike times, which is done with no noise on the dynamics. In the appendix, we prove that this set is stable under noise and deterministic perturbations.

Let t_1 be a time when the input current goes below threshold. A solution arriving at threshold at time t_1 cannot cross it and will hit threshold again only later, at time t_2 , as shown in Figure 3. The area defined by this solution and the line $V = V_t$ cannot be reached by any other solution, since trajectories cannot cross. Therefore, no spike can occur between t_1 and t_2 , which happens every time the current goes below threshold, and so such time intervals can be removed from the possible spike times. The area defined by the solutions starting at reset between t_1 and t_2 cannot be reached by any solution starting before t_1 , since such a solution would spike between t_1 and




Figure 4: Construction of unreachable time intervals. As in Figure 3, time is on the horizontal axis and membrane potential on the vertical axis. The color areas are unreachable by any solution. Thus, the sample solution, shown as a dashed line, always remains in the white area. The simulation was 500 ms long with a leaky integrator.

t_2 . The spike times of these forbidden solutions define a whole sequence of unreachable spike time intervals. Therefore, every time the input current goes below threshold, a whole sequence of intervals is removed from the possible spike times, as shown in Figure 4. We can see that the first time the current goes below threshold, a red area that no solution can reach is created near threshold. The solutions starting from the resulting forbidden time interval are in red. The next time the current goes below threshold occurs outside the red area and creates a new sequence of yellow stripes. The third time occurs inside the red area, and the result is displayed in blue. Thus, a solution starting at time 0 always remains in the white area, as shown by the dashed black line. One can see that the white area gradually gets smaller and smaller.

Based on the above considerations, we predict that when the current oscillates around threshold, the set of possible spike times eventually gets very narrow, so that spike timing should be reproducible in the long run, even if the potential is not fixed at reset at stimulus onset. For the general model given by equation 2.1, “the current goes below threshold” means that the vector field points downward at threshold, that is, $f(V_t, t) < 0$.

In the construction described above, we implicitly assumed that a solution is at reset potential at time t only when it spikes at time t , which is not necessarily so in the general case of equation 2.1. For example, in Figure 4, one might imagine a solution starting at time 0, going below reset, and then going up into the first red stripe. This cannot happen if we assume $f(V_r, t) > 0$ for all t . Another sufficient condition to avoid this problem is to assume that the model is leaky, as discussed in the appendix. None of these two conditions is fulfilled by the perfect integrator. Indeed, in this model, when the vector field points downward at threshold, it also does so at reset. The numerical simulations show that this model is never reliable.

We briefly describe the numerical construction of the set of reachable spike times, as illustrated in Figure 4. First, find the first time t_1 when the input current goes below threshold. This is the starting point of the first

forbidden interval. Then run the solution starting from threshold at time t_1 and stop when it hits threshold again at time t_2 . The interval $I = [t_1, t_2]$ is the first forbidden interval, where no spike can occur. Then run the solution starting from reset at time t_1 (resp. t_2) and stop when it hits threshold at time t_3 (resp. t_4), so as to obtain the next forbidden interval $[t_3, t_4]$. By repeating the process, we get a whole sequence of intervals where no spike can occur. In the remaining set of times, the current goes again below threshold. So find the first time when it does so after t_2 , and repeat the procedure. Thus, when the algorithm ends, we get the remaining set of reachable spike times, and we shall see that this set eventually gets very small.

Since we chose to display the reliability of model responses for all parameters of a given family of currents (see Figure 5A) on a single figure, we computed the set of reachable spike times for each parameter (see Figures 5B and 5C) and displayed it in a two-dimensional white and black figure, with time on the horizontal axis and parameter on the vertical axis (see Figure 5D). Black dots represent reachable spike times.

The explanation above does not apply to the case when the input current is always above threshold. However, in our numerical results, responses seem to be reproducible in both cases, although the required intrinsic noise for reproducibility is significantly lower than for currents crossing the threshold. Our explanation does not rule out the possibility of multiple stable solutions, but this does not seem to happen in our simulations.

4 Numerical Results

4.1 Periodic Currents. The results regarding periodic stimuli are displayed in Figure 6, which illustrates the theoretical results. A leaky integrator is driven by a 20 Hz sinusoidal current (see Figure 6A), whose mean is a decreasing linear function of the parameter. For parameters $p < 0.5$, the current is always above threshold, whereas for $p > 0.5$, the current oscillates around threshold. The spike density figure (see Figure 6B) was computed without noise, with random initial potential for every trial, so that phase locking appears clearly. Adding dynamical noise results only in blurring this figure; it has no qualitative influence on the results. The model was run for 5 s. When phase locking occurs, all the stable solutions can be observed in the last 250 ms of the runs (see Figure 6B). Phase locking appears for a given parameter value when the density line for this value consists of a finite number of black dots only, that is, when the corresponding PSTH consists of a finite number of isolated peaks. Conversely, there is no phase locking when the density line consists of gray shades, that is, when the corresponding PSTH is a smooth distribution. As expected, responses are always phase-locked for $p > 0.5$, whereas there are not always for $p < 0.5$. Mean firing rate and precision were computed for every parameter value with noise-free dynamics, as seen in Figure 6C (solid lines), where phase locking corresponds to plateaus in the firing rate and to drops in precision


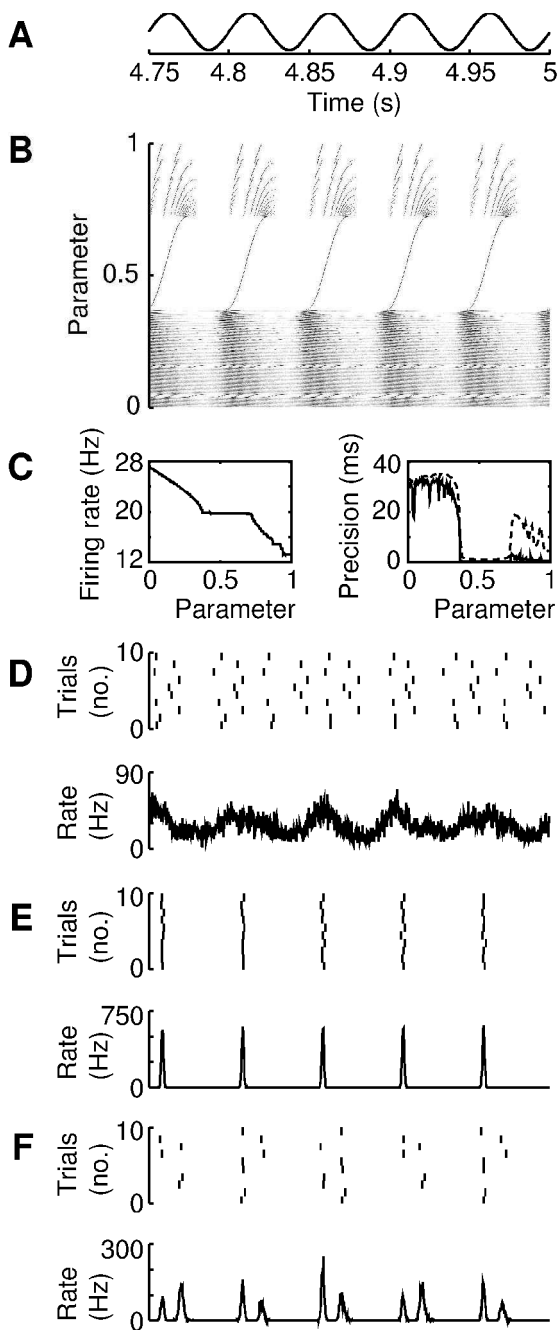


Figure 5: Construction of the set of possible spike times for all parameters of a given family. (A) Basis function. (B) For a given parameter, we compute the times when no spike may occur (dark areas are not reachable). (C) This is converted into a line of white and black rasters, with black ones corresponding to possible spike times. (D) Collecting all these lines for all parameters and arranging them vertically, we get a two-dimensional white and black figure of possible spike times, with time on the horizontal axis and parameter on the vertical axis.

measure. Precision was also computed with a small noise on the dynamics ($1 \text{ mV}\cdot\text{s}^{-1/2}$), which results in a smoothed version of the noise-free precision plot, where only 1:1 phase locking appears clearly (see Figure 6C, dashed line). There is no trivial relationship between precision and parameter. The precision is better for $p > 0.5$, where the model is always phase-locked (the current oscillates around threshold), than for $p < 0.5$. The best precision is obtained for 1:1 phase locking, as observed in Hunter et al. (1998).

We showed the results for three parameter values (see Figures 6D–6F), which illustrate the three possible cases of reliability. The simulations were done with a small dynamical noise.

Case 1: Noise accumulates. For $p = 0$, responses were not reproducible. After 5 s, the PSTH was flat, and individual trials were not synchro-



nized (see Figure 6D). The model behaves in the same way as with a constant current.

Case 2: Multiple stable solutions. For $p = 1$, the responses are phase-locked, as shown in the spike density figure (see Figure 6B). The spike rate is 13.3 Hz (see Figure 6C), which is two-thirds the input frequency. Thus, there is 2:3 phase locking; the model neuron spikes twice every three periods, as shown in the individual trials (see Figure 6F). The density figure (see Figure 6B) and the PSTH (see Figure 6F) indicate six spikes every three periods, which means there are three stable solutions, resulting also in a higher sensitivity to noise (see Figure 6C, dashed line).

Case 3: One stable solution. For $p = 0.5$, the spike rate is 20 Hz (see Figure 6C): the model is 1:1 phase-locked. There is only one stable solution, as seen in the individual trials (see Figure 6E). Thus, responses are truly reproducible, and precision is very high (see Figure 6D).

In these simulations, case 3 occurs in only about 30% of the parameters range (the 1:1 phase-locking area).

4.2 Aperiodic Currents.

4.2.1 Noise Does Not Accumulate over Time. In Figure 7, we assessed the reliability of a noisy leaky integrator in response to a random triangle wave, whose mean is constant and amplitude varies linearly with the parameter (the current was constant for $p = 0$). For all trials, membrane potential was at rest at stimulus onset; thus, we do not distinguish between cases 2 (several stable solutions) and 3 (one stable solution). The input current is always above threshold for $p < 0.5$ and oscillates around threshold for $p > 0.5$.

Figure 6: *Facing page.* Reliability of a leaky integrator driven by a periodic input current. The model is a linear leaky integrator. (A) Input current is a 20 Hz sine wave $I(t) = 85 + 40(1 - p) + 30\sin(40\pi t)$ pA, where p is the parameter. The results are shown from 4.75 s after the onset of the stimulus. (B) The model was run 2000 times for every parameter value (400 values between 0 and 1), with initial potential drawn uniformly between threshold and reset, but no noise on the dynamics, so that phase locking appears clearly. (Adding noise results in a blurred version of this figure.) The results are displayed as a spike density figure. (C) Firing rate was computed as a function of the parameter, with no noise on the dynamics. Precision was computed with no noise on the dynamics (solid line) and with a noise of $1 \text{ mV}\cdot\text{s}^{-1/2}$ (dashed line). Phase locking appears for parameter values below 0.5 as drops in precision. (D–F) For three different parameter values (0, 0.5, 1), we show the results of 10 trials of the model with noise on the dynamics ($2.5 \text{ mV}\cdot\text{s}^{-1/2}$) and random initial potential, and also show the PSTH, computed from 2000 trials. Note the different scales for firing rate in these graphs.

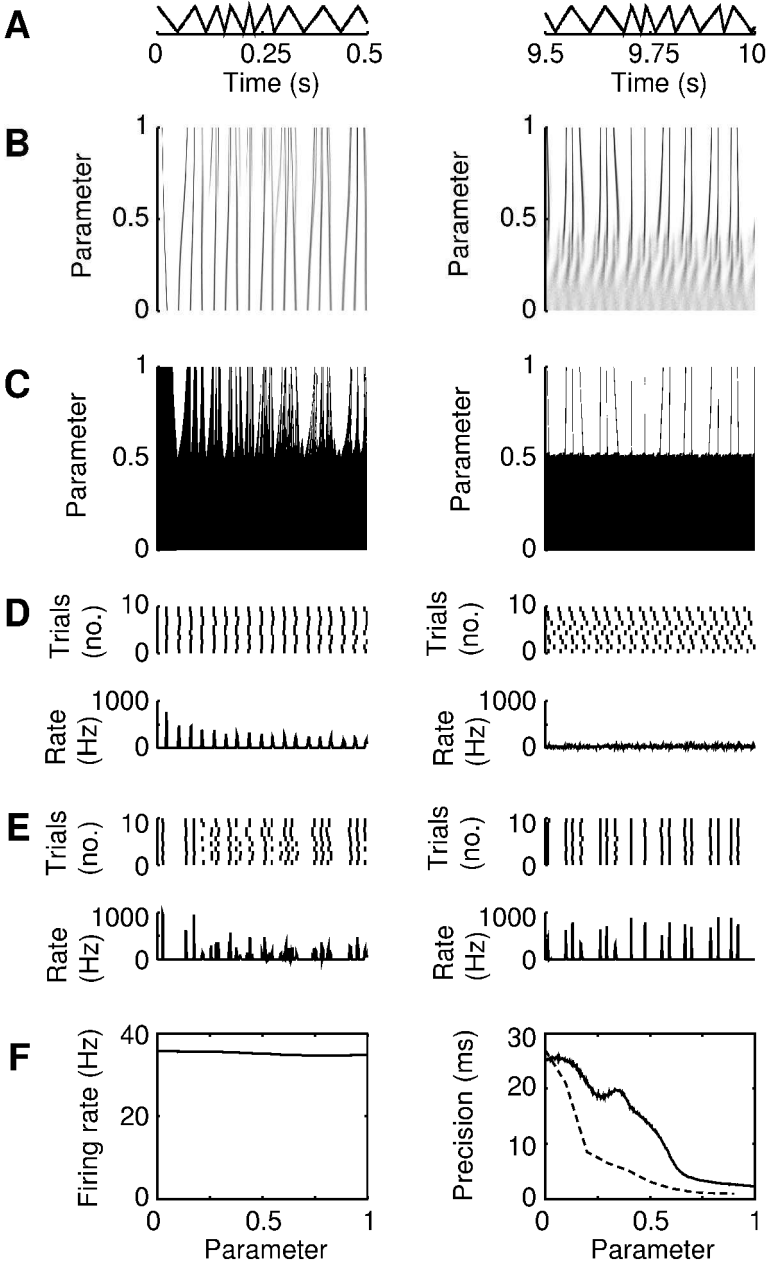
Trials were 10 s long. The spike density figure (see Figure 7B) shows that when the input current oscillated around threshold ($p > 0.5$), spikes were synchronized in the long run, and the actual spike times corresponded with our computational prediction of possible spike times (see Figure 7C). Tight synchronization can be clearly seen in the 10 individual trials computed for parameter $p = 0.8$ in Figure 7E, and the PSTH is very peaked. Thus, noise did not accumulate (case 2 or 3).

Spiking seems less precise when the input current is always above threshold ($p < 0.5$), as shown in the spike density figure (see Figure 7B) and the precision plot (see Figure 7F, solid line). Individual trials for parameter $p = 0.1$ (see Figure 7D) show synchronized spikes at the beginning of the simulation, but not in the last 500 ms of the 10 s runs. This is reflected by the PSTH as a steady decrease in the height of the peaks during the first 500 ms and as a flat distribution during the last 500 ms. However, we also computed the precision of long runs (200 s) with a lower noise on the dynamics ($0.7 \text{ mV}\cdot\text{s}^{-1/2}$) (see Figure 7F, dashed line). This duration was enough for noise to accumulate in the case of a constant current ($p = 0$), but it did not lead to a desynchronization of trials for parameters above 0.2. Thus, it seems that in contrast with periodic currents, responses to aperiodic currents are also reproducible (case 2 or 3) when the current is above threshold, as suggested by other studies (Jensen, 1998; Pakdaman & Tanabe, 2001), though the level of tolerable noise is lower than when the current oscillates around threshold.

The type of current used (e.g., gaussian instead of triangle wave) did not alter the results.

4.2.2 There Is Only One Stable Solution. In Figure 8, we assessed the reliability of a noisy leaky integrator in response to a random triangle wave,

Figure 7: *Facing page.* Reliability of a leaky integrator with initial potential at rest. The model is a linear leaky integrator. (A) Input current is $I(t) = 150 + 150pB(t)$ pA, where p is the parameter and $B(t)$ is a triangular wave taking values between -1 and 1 , with rising and falling time drawn uniformly between 10 ms and 50 ms. We show the results for the first and last 500 ms of the 10 s simulations. (B) The model was run 2000 times for every parameter value (400 values between 0 and 1), with membrane potential initially at rest value, and noise on the dynamics ($3.5 \text{ mV}\cdot\text{s}^{-1/2}$). The results are displayed as a spike density figure. (C) The set of possible spike times was computed for every parameter value. (D, E) For parameter values $p = 0.1$ (E) and $p = 0.8$ (F), 10 individual trials were extracted and displayed, along with the PSTH computed from 2000 trials. (F) Mean firing rate does not vary much with the parameter. The precision was computed as a function of the parameter from the last 7.5 s of 10 s' runs with a noise of $3.5 \text{ mV}\cdot\text{s}^{-1/2}$ on the dynamics (solid line, 400 parameter values), and from the last 10 s of 200 s' runs with a noise $0.7 \text{ mV}\cdot\text{s}^{-1/2}$ (dashed line, 10 parameter values).



whose mean is constant and amplitude varies linearly with the parameter. For every trial, membrane potential was drawn at random uniformly between rest and threshold, so that we could distinguish between cases 2 (several stable solutions) and 3 (one stable solution).


When the input current oscillates around threshold ($p > 0.5$), we can see from the spike density figure (see Figure 8B) that spike timing becomes precise after some time. In the individual trials for parameter $p = 0.95$ (see Figure 8E), spikes initially occur at different times, depending on the membrane potential at onset, and eventually they become synchronized over trials, whatever the initial potential, which corresponds to case 3, and was predicted by the computation of possible spike times (see Figure 8C). The results are shown only for the first 500 ms, but the behavior after 9.5 s is similar to the one shown in Figure 7.

When the input current is always above threshold ($p < 0.5$), responses are not reproducible with this magnitude of intrinsic noise, as the results of individual trials for $p = 0.1$ show (see Figure 8D). However, we found that with a lower intrinsic noise, responses were reproducible in the long run for small parameter values even after 200 s (precision of less than 6 ms for $p > 0.2$).

4.2.3 Spike Times Are Stable Under Deterministic Perturbations. Do spike times in the long run change when the input current is increased? The spike times of a neuron driven by a constant current are not stable under deterministic perturbations: if the (constant) ISI is T and a slight decrease of the current makes it $T + \delta$, then the difference between the times of the n th spike is $n\delta$, leading to a desynchronization of the two patterns of spikes after enough time.

In Figure 9, we simulated a noisy leaky integrator driven by a random triangle wave, whose mean and amplitude increased linearly with the parameter. The current oscillated around threshold for every parameter value. For every trial, membrane potential was drawn at random uniformly between rest and threshold. The mean firing rate, computed over 10 s, increased from 22 to 35 Hz with the parameter (see Figure 9D), but spike times did not vary significantly with the parameter, even after 750 ms (see Figures 9B and 9C).

Figure 8: *Facing page.* Reliability of a leaky integrator with random initial potential. The model is a linear leaky integrator. (A) Input current is $I(t) = 150 + 150pB(t)$ pA, where p is the parameter and $B(t)$ is a normalized filtered gaussian noise with time constant 10 ms. (B) The model was run 2000 times for every parameter value (400 values between 0 and 1), with initial potential drawn uniformly between rest and threshold, and noise on the dynamics ($2.5 \text{ mV}\cdot\text{s}^{-1/2}$). The results are displayed for the first 500 ms as a spike density figure. (C) The set of possible spike times was computed for every parameter value. (D, E) For parameter values $p = 0.1$ (D) and $p = 0.95$ (E), 10 individual trials were extracted and displayed, along with the PSTH computed from 2000 trials.



Thus, an increase in firing rate creates additional spikes rather than a shift of all spikes. Precision was almost constant for all parameter values—around 3 ms (see Figure 9D). Thus, when the input current oscillates around threshold, responses are stable under significantly large deterministic perturbations.

4.2.4 Nonlinear and Nonleaky Models Are Also Reliable. In Figure 10, we assessed the reliability of a noisy leaky nonlinear model driven by equation 2.3 (see Figures 10B and 10C) and a noisy nonleaky model described by equation 2.4 (see Figures 10D and 10E) driven by a gaussian input current whose mean is constant and amplitude varies linearly with the parameter. For every trial, membrane potential was drawn at random uniformly between rest and threshold. Again, responses were reproducible (case 3) when the input current oscillated around threshold ($p > 0.5$).

4.2.5 The Perfect Integrator Is Never Reliable. As a counterexample, we simulated a noisy perfect integrator (Knight, 1972; see Figure 11), driven by a gaussian current whose mean is constant and amplitude varies linearly with the parameter. For all trials, membrane potential was at rest at stimulus onset. The hypotheses required for our theoretical explanation are not fulfilled by this model.

We can see in Figure 11B that whether the input current is above threshold ($p < 0.5$) or oscillates around it ($p > 0.5$), the noise accumulates and leads to a complete desynchronization of spikes over trials (case 1). Figure 11C shows that the set of possible spike times does not get smaller over time, and Figure 11D that the precision is always above 13 ms, which is half the average ISI.

5 Discussion

We identified two ways in which neuron responses can fail to be reproducible: dynamical noise can accumulate over time and lead to a desynchronization over trials, or several stable responses can exist, depending on the initial condition. These two cases can occur when the model neuron is driven by a periodic current, but the former can occur only if the current is above threshold (Keener et al., 1981). In contrast, we showed that responses of a general class of spiking neuron models to aperiodic currents oscillating around threshold are reproducible: spike timing in the long run is stable under noise and deterministic perturbation and does not depend on the initial condition. We provided a theoretical explanation based on a geometrical construction that shows that the set of possible spike times becomes smaller and smaller over time. Our numerical results suggest that responses to aperiodic currents above threshold are also reproducible, which confirms other studies (Pakdaman & Tanabe, 2001), but the level of tolerable noise seems lower in this case.




Figure 9: Stability of spike timing under deterministic perturbations. The model is a linear leaky integrator. (A) Input current is $I(t) = 112.5 + (75 + 37.5p)B(t) + 37.5p$ pA, where p is the parameter and $B(t)$ is a triangular wave taking values between -1 and 1 , with rising and falling time drawn uniformly between 10 ms and 50 ms. Thus, the mean input current increases with the parameter from 112.5 pA to 150 pA. (B) The model was run 2000 times for every parameter value (400 values between 0 and 1), with initial potential drawn uniformly between rest and threshold and noise on the dynamics (2.5 mV.s $^{-1/2}$). The results are displayed for the first 750 ms as a spike density figure. (C) The set of possible spike times was computed for every parameter value. (D) Mean firing rate and precision were computed from 10 s simulations.

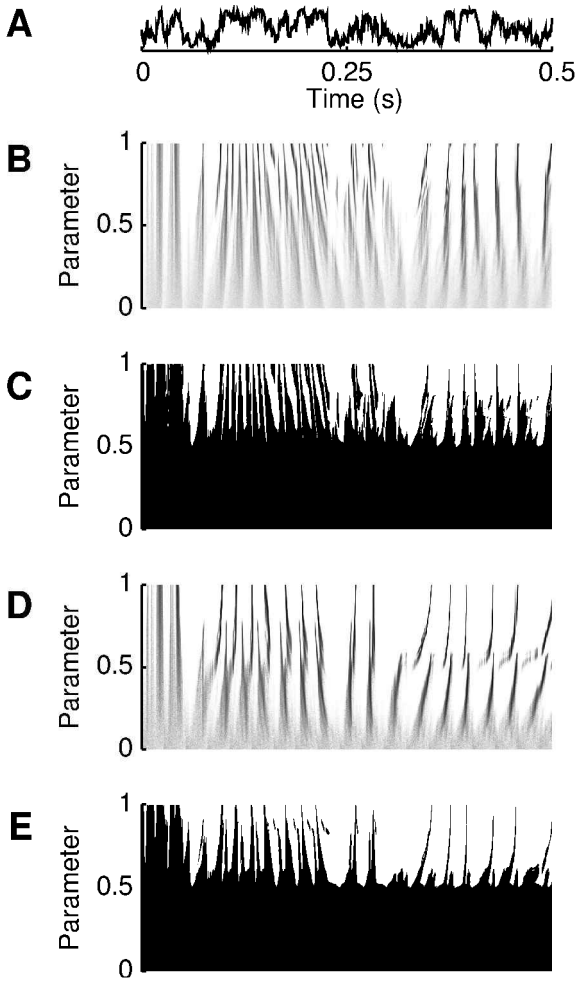


Figure 10: Reliability of a nonlinear and a nonleaky model. We assessed the reliability of two models that differ from the linear leaky integrator: a leaky nonlinear model driven by equation 2.3 (B, C) and a nonleaky model driven by equation 2.4 (D, E). (A) Input current is $I(t) = 150 + 150pB(t)$ pA for the leaky nonlinear model and $I(t) = 0.5 + 3pB(t)$ for the nonleaky model (unit is irrelevant), where p is the parameter and $B(t)$ is a normalized filtered gaussian noise with time constant 10 ms. (B) The nonlinear leaky model was run 2000 times for every parameter value (400 values between 0 and 1), with initial potential drawn uniformly between rest and threshold, and noise on the dynamics ($2.5 \text{ mV}\cdot\text{s}^{-1/2}$ for the leaky model). The results are displayed for the first 500 ms as a spike density figure. (C) The set of possible spike times was computed for every parameter value. (D, E) The same simulations were done for the nonleaky model, with noise on the dynamics $0.1 \text{ V}\cdot\text{s}^{-1/2}$.

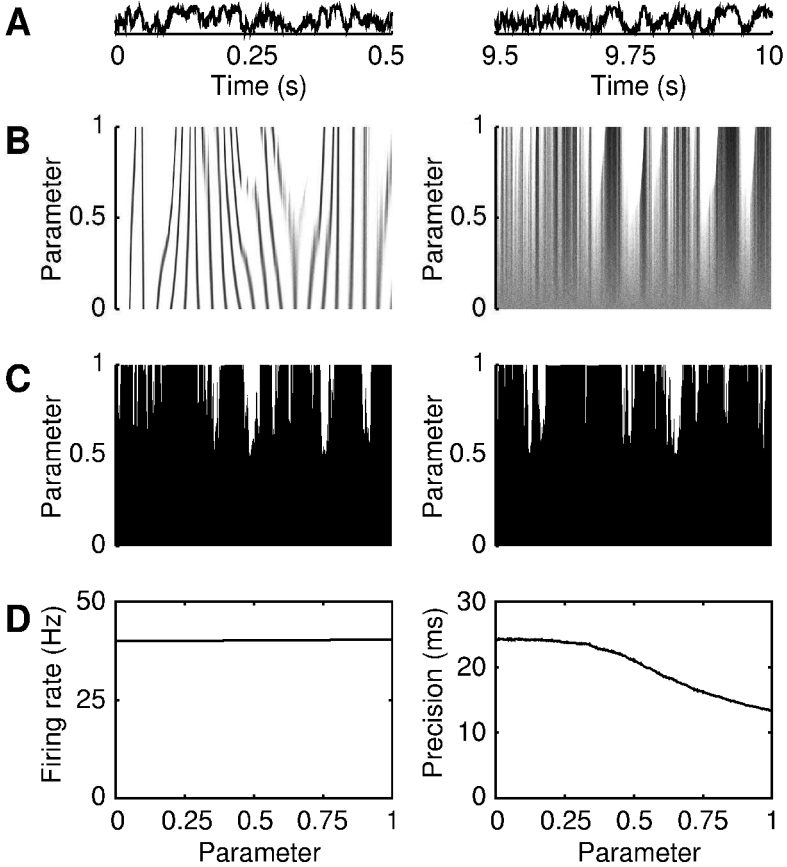


Figure 11: The perfect integrator is not reliable. The model is a perfect integrator driven by equation 2.5. (A) Input current is $I(t) = .5 + pB(t)$ (units are not significant here), where p is the parameter and $B(t)$ is a normalized filtered gaussian noise with time constant 10 ms. We show the results for the first and last 500 ms of the 10 s simulations. (B) The model was run 2000 times for every parameter value (400 values between 0 and 1), with membrane potential initially at rest value and noise on the dynamics ($0.27 \text{ V}\cdot\text{s}^{-1/2}$). The results displayed are displayed as a spike density figure. (C) The set of possible spike times was computed for every parameter value. (D) Mean firing rate is almost constant, around 40 Hz, while precision is always above 13 ms, which is very high, knowing that this is about half the mean ISI.

We showed that for aperiodic currents that cross the threshold, noise does not accumulate over time. In other words, in this case, the asymptotic spike time jitter is a function of the level of intrinsic noise that tends to 0 as the noise variance tends to 0, which is not generally so for periodic currents. This result is not a consequence of particular factors such as the leak, the amount of fluctuations in the input current, or its frequency spectrum. However, quantitatively, the ratio of the asymptotic spike time jitter to the noise variance depends on various characteristics of the input current (and, more generally, of the differential equation; Hunter et al., 1998; Fellous et al., 2001; Beierholm et al., 2001; Tiesinga, 2002; Tiesinga et al., 2002b). Such a relationship is relevant only when responses are reproducible in the way defined above. In particular, it would be illusory to search for a formula describing the precision as a function of a parameter of the input in the case of periodic currents above threshold. Indeed, when varying the parameter, the model alternates between phase locking, where precision is close to 0, and structural instability, where spike timing is imprecise even with an arbitrarily small intrinsic noise (besides being dependent on the initial condition). However, when the model is phase-locked, the (quantitative) precision depends on dynamical properties such as the Lyapunov exponent and the distance to the boundary of the Arnold tongue (Tiesinga, 2002).

Appendix: Stability of the Set of Possible Spike Times

A leaky model is described by a one-dimensional differential equation,

$$\frac{dV}{dt} = f(V, t) \tag{A.1}$$

where V is the membrane potential, with the spike modeled as follows: when the potential reaches a threshold V_t , it is reset to V_r . The leak is described by the following condition on f :

$$\frac{\partial f}{\partial V} < 0.$$

This inequality implies that the distance between any two solutions of equation A.1 (without spike) always decreases and tends to zero—hence, the term *leak*.

Since we are interested in the sequence of spike times, it is natural to introduce the map $\varphi: \mathbb{R} \rightarrow \mathbb{R}$ that gives the time of the first spike following a given spike time, that is, $\varphi(t)$ is the time when a run starting from reset potential at time t hits threshold. We shall call this map the *spike map*. Thus, the sequence of spike times of a run with first spike at time t is just $t, \varphi(t), \varphi^2(t), \dots$. The spike map was first introduced in Rescigno et al. (1970) in the case when the input current is periodic and studied further in Knight (1972) and Keener et al. (1981).

One can easily show that for every t , $f(V_t, \varphi(t)) \geq 0$, which means there can be no spike if the vector field points downward at threshold, that is, in the case of the linear leaky integrator, if the input current is below the current threshold. This result is the first cornerstone of the construction of the set of possible spike times, and it holds also for nonleaky spiking models, including the perfect integrator. It allows us to remove intervals from the set of possible spike times, as described in section 3.2. In the construction of this set, we also remove the image of these intervals by iterates of φ . However, this is correct only if φ is injective. For leaky models, one can show that φ is strictly increasing on its range (see Keener et al., 1981, for the case of periodic currents). This is also true if $f(V_r, t) > 0$ for all t , which is the case of the nonleaky models we simulated.

For periodic currents, that is, in the general case, when f is periodic with respect to t , the fact that φ is strictly increasing on its range means it is a lift of an orientation-preserving circle map. If φ is surjective, that is, if the current is always above threshold, then φ is a lift of a homeomorphism of the circle. The dynamics of both noninvertible orientation-preserving circle maps and homeomorphisms of the circle are well known, as we described in section 3.1.

Our aim is to prove that the set of possible spike times constructed in section 3.2 is stable under noise and deterministic perturbations. We assume that the duration of ISIs is bounded, so that there is an $M > 0$ such that $\varphi(t) - t < M$ for all t .

For mathematical convenience, we consider here that a run may hit threshold without spiking, as long as it does not cross it. Consider runs defined on \mathbb{R} (not only \mathbb{R}^+), with infinitely many spikes on both \mathbb{R}^+ and \mathbb{R}^- . Note $R \subset \mathbb{R}$ the set of spike times of all these runs, that is, the set of possible spike times constructed from $t = -\infty$. Now consider a perturbation of this system defined in the following way: take $\epsilon > 0$, representing the amount of perturbation; the spike following a spike at time t_n occurs at time t_{n+1} such that $|t_{n+1} - \varphi(t_n)| \leq \epsilon$. This accounts for an uncertainty that may be random or deterministic. Note R_ϵ the set of spike times of all possible perturbed runs for a given $\epsilon > 0$. We shall prove the following stability theorem:

Theorem 1. *For any compact set K such that $K \cap R = \emptyset$, there is an $\epsilon > 0$ such that $K \cap R_\epsilon = \emptyset$.*

In other words, any compact set, such as an interval, that is not reachable under the original deterministic dynamical system is not reachable either if a small perturbation is added on the dynamics. Therefore, adding a small noise to the dynamics should have a small influence on the set of spike times. In the same way, a slight change of the vector field should cause a small perturbation of spike times.

First, let us formalize the considerations above. Considering that a run may hit threshold without spiking, as long as it does not cross it, means in the present case that if the spike map is discontinuous at t , then the spike following a spike at time t may occur at either time $\varphi(t^-)$ or $\varphi(t^+)$. Now choose $\epsilon > 0$ as the level of uncertainty; if there is a spike at time t , then the next spike will occur within distance ϵ of either $\varphi(t^-)$ or $\varphi(t^+)$, that is, the set of possible times for next spike is $[\varphi(t^-) - \epsilon, \varphi(t^-) + \epsilon] \cup [\varphi(t^+) - \epsilon, \varphi(t^+) + \epsilon]$. Formally, this framework may be expressed in terms of a set-valued dynamical system. We refer the reader to Aubin and Frankowska (1990) for more details about set-valued analysis. A sequence (t_n) of spike times agreeing with our considerations is one that satisfies

$$t_{n+1} \in F_\epsilon(t_n), \tag{A.2}$$

where $F_\epsilon(t) = [\varphi(t^-) - \epsilon, \varphi(t^-) + \epsilon] \cup [\varphi(t^+) - \epsilon, \varphi(t^+) + \epsilon]$. Thus, F_ϵ is a set-valued map that we write $F_\epsilon: \mathbb{R} \downarrow \mathbb{R}$, as in Aubin and Frankowska (1990).

We introduce a few definitions from set-valued analysis:

Definition 1.

- The graph of a set-valued map $F: X \downarrow Y$ is the set of points $(x, y) \in X \times Y$ such that $y \in F(x)$.
- The inverse of a set-valued map $F: X \downarrow Y$ is the set-valued map $F^{-1}: Y \downarrow X$ defined by $x \in F^{-1}(y)$ if and only if $y \in F(x)$.
- If $A \subset \mathbb{R}$, then

$$F(A) = \bigcup_{t \in A} F(t).$$

- If $F: \mathbb{R} \downarrow \mathbb{R}$ and $G: \mathbb{R} \downarrow \mathbb{R}$ are two set-valued maps, then the composition $G \circ F: \mathbb{R} \downarrow \mathbb{R}$ is the set-valued map defined by $G \circ F(t) = G(F(t))$.
- The notation F^n stands for the n -fold composition $F \circ F \circ \dots \circ F$.
- The intersection of $F: \mathbb{R} \downarrow \mathbb{R}$ and $G: \mathbb{R} \downarrow \mathbb{R}$ is the set-valued map $F \cap G: \mathbb{R} \downarrow \mathbb{R}$ defined by $F \cap G(t) = F(t) \cap G(t)$.

Thus, $F_\epsilon^n(t)$ is the set of possible times for the n th spike following a spike at time t , given an uncertainty ϵ . The set of possible times for the n th spike of any run is $F_\epsilon^n(\mathbb{R})$. Consider a run defined on \mathbb{R} , with infinitely many spikes on \mathbb{R}^- . Every spike of this run must occur within the set $F_\epsilon^n(\mathbb{R})$, for any integer n . Thus, the spike times are contained in the set

$$R(F_\epsilon) = \bigcap_{n>0} F_\epsilon^n(\mathbb{R}).$$

Conversely, for any time $t \in R(F_\epsilon)$, we can find a sequence of spike times satisfying equation A.2. Thus, $R(F_\epsilon)$ is the set of spike times of all such runs.

In order to prove theorem 1, we need another definition:

Definition 2. Let $F: X \downarrow Y$ be a set-valued map. F is said to be compact on compact sets (CC) if for any sequence $x_n \in X$ converging to x , any sequence $y_n \in F(x_n)$ has a cluster point $y \in F(x)$.

Equivalently, F is CC if its graph is closed and it maps bounded sets to bounded sets.

This is a generalization of continuity for set-valued maps. One may note that when F is a single-valued map, it is equivalent to continuity. Let us state a few basic results regarding CC maps:

Proposition 1.

1. If $F: X \downarrow Y$ and $G: Y \downarrow Z$ is CC, then so is $G \circ F$.
2. If $F_\lambda: X \downarrow Y, \lambda \in I$, is a family of CC set-valued maps, then the intersection

$$\bigcap_{\lambda \in I} F_\lambda$$

is also CC.

3. Let

$$I \times X \downarrow Y$$

$$(\lambda, x) \mapsto F_t(x)$$

a CC set-valued map, and let $K \subset X$ be a compact set. Then $\lambda \mapsto F_\lambda(K)$ is CC.

Theorem 1 follows from the following lemma:

Lemma 1. The set-valued map $\epsilon \in \mathbb{R}^+ \mapsto F_\epsilon^n(\mathbb{R}) \cap K$ is CC.

Proof of Theorem 1. It follows from lemma 1 and proposition 1 that the set-valued map $\epsilon \mapsto R(F_\epsilon) \cap K$ is CC; hence, its domain is closed, that is, the set of ϵ such that $R(F_\epsilon) \cap K \neq \emptyset$ is closed. Therefore, if $R(F_0) \cap K = \emptyset$, then for ϵ small enough, $R(F_\epsilon) \cap K = \emptyset$.

Proof of Lemma 1. Note that $F_\epsilon^n(\mathbb{R}) \cap K = F_\epsilon^n(F_\epsilon^{-n}(K))$. It follows from proposition 1 that $(\epsilon, t) \mapsto F_\epsilon^n(t)$ is CC. Therefore, it is sufficient to prove that $(\epsilon, t) \mapsto F_\epsilon^{-n}(t)$ is CC, and since its graph is closed, we need to prove only that it maps bounded sets to bounded sets. That is, we must prove that

for any $A > 0$ and any compact interval I , the set

$$\bigcup_{\epsilon \leq A} F_{\epsilon}^{-n}(I) = F_A^{-n}(I)$$

is bounded. Since the function $t \mapsto \varphi(t) - t$ is bounded, we have

$$\lim_{t \rightarrow -\infty} \varphi^n(t) = -\infty.$$

Therefore, φ^n maps an unbounded set of \mathbb{R}^- to an unbounded set. Hence, $F_A^{-n}(I)$ must be bounded.

Acknowledgments

We thank K. Pakdaman and J.-P. Aubin for many helpful discussions and Christine Lienhart and Marc Maier for revising our English.

References

- Arnold, V. I. (1961). Small denominators. I. Mapping the circle onto itself. *Izv. Akad. Nauk. SSSR Ser. Mat.*, 25, 21–86.
- Ascoli, C., Barbi, M., Chillemi, S., & Petracchi, D. (1977). Phase-locked responses in the *Limulus* lateral eye: Theoretical and experimental investigation. *Biophys. J.*, 19(3), 219–240.
- Aubin, J. P., & Frankowska, H. (1990). *Set-valued analysis*. Boston: Birkhäuser.
- Baddeley, R., Abbott, L., Booth, M., Sengpiel, F., Freeman, T., Wakeman, E., & Rolls, E. (1997). Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proc. R. Soc. Lond. B Biol. Sci.*, 264(1389), 1775–1783.
- Bair, W., & Koch, C. (1996). Temporal precision of spike trains in extrastriate cortex of the behaving macaque monkey. *Neural Comp.*, 8(6), 1185–1202.
- Beierholm, U., Nielsen, C., Ryge, J., Alstrom, P., & Kiehn, O. (2001). Characterization of reliability of spike timing in spinal interneurons during oscillating inputs. *J. Neurophysiol.*, 86(4), 1858–1868.
- Berry, M., & Meister, M. (1998). Refractoriness and neural precision. *J. Neurosci.*, 18(6), 2200–2211.
- Berry, M., Warland, D., & Meister, M. (1997). The structure and precision of retinal spike trains. *Proc. Natl. Acad. Sci. USA*, 94(10), 5411–5416.
- Borst, A., & Theunissen, F. (1999). Information theory and neural coding. *Nat. Neurosci.*, 2, 947–957.
- Buracas, G., Zador, A., DeWeese, M., & Albright, T. (1998). Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron*, 20(5), 959–969.
- Coddington, E., & Levinson, N. (1955). *Theory of ordinary differential equations*. New York: McGraw-Hill.

- de Ruyter van Steveninck, R., Lewen, G., Strong, S., Koberle, R., & Bialek, W. (1997). Reproducibility and variability in neural spike trains. *Science*, *275*, 1805–1808.
- Denjoy, A. (1932). Sur les courbes définies par les équations différentielles à la surface du tore. *J. Math. Pures. Appl.*, *9*(11), 333–375.
- Fellous, J. M., Houweling, A., Modi, R., Pao, R., Tiesinga, P., & Sejnowski, T. (2001). Frequency dependence of spike timing reliability in cortical pyramidal cells and interneurons. *J. Neurophysiol.*, *85*, 1782–1787.
- Guttman, R., Feldman, L., & Jakobsson, E. (1980). Frequency entrainment of squid axon membrane. *J. Membrane Biol.*, *56*, 9–18.
- Herman, M. (1977). *Mesure de Lebesgue et nombre de rotation*. Berlin: Springer-Verlag.
- Hunter, J., Milton, J., Thomas, P., & Cowan, J. (1998). Resonance effect for neural spike time reliability. *J. Neurophysiol.*, *80*(3), 1427–1438.
- Jensen, R. (1998). Synchronization of randomly driven nonlinear oscillators. *Phys. Rev. E*, *58*(6), R6907–R6910.
- Kara, P., Reinagel, P., & Reid, R. (2000). Low response variability in simultaneously recorded retinal, thalamic, and cortical neurons. *Neuron*, *27*(3), 635–646.
- Keener, J. (1980). Chaotic behavior in piecewise continuous difference equations. *Trans. Amer. Math. Soc.*, *261*(2), 589–604.
- Keener, J., Hoppensteadt, F., & Rinzel, J. (1981). Integrate-and-fire models of nerve membrane response to oscillatory input. *SIAM J. Appl. Math.*, *41*, 503–517.
- Knight, B. (1972). Dynamics of encoding in a population of neurons. *J. Gen. Physiol.*, *59*, 734–766.
- Koppl, C. (1997). Phase locking to high frequencies in the auditory nerve and cochlear nucleus magnocellularis of the barn owl, *Tyto alba*. *J. Neurosci.*, *17*(9), 3312–3321.
- Lapicque, L. (1907). Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation. *J. Physiol. Pathol. Gen.*, *9*, 620–635.
- Mainen, Z., & Sejnowski, T. (1995). Reliability of spike timing in neocortical neurons. *Science*, *268*, 1503–1506.
- Nowak, L., Sanchez-Vives, M., & McCormick, D. (1997). Influence of low and high frequency inputs on spike timing in visual cortical neurons. *Cereb. Cortex*, *7*(6), 487–501.
- Pakdaman, K., & Tanabe, S. (2001). Random dynamics of the Hodgkin-Huxley neuron model. *Phys. Rev. E*, *64*(5Pt1), 050902.
- Reich, D., Mechler, F., Purpura, K., & Victor, J. (2000). Interspike intervals, receptive fields, and information encoding in primary visual cortex. *J. Neurosci.*, *20*(5), 1964–1974.
- Reich, D., Victor, J., & Knight, B. (1998). The power ratio and the interval map: Spiking models and extracellular recordings. *J. Neurosci.*, *18*, 10090–10104.
- Reich, D., Victor, J., Knight, B., Ozaki, T., & Kaplan, E. (1997). Response variability and timing precision of neuronal spike trains in vivo. *J. Neurophysiol.*, *77*(5), 2836–2841.
- Reinagel, P., & Reid, R. (2000). Temporal coding of visual information in the thalamus. *J. Neurosci.*, *20*(14), 5392–5400.

- Rescigno, A., Stein, R., Purple, R., & Poppele, R. (1970). A neuronal model for the discharge patterns produced by cyclic inputs. *Bull. Math. Biophys.*, *32*, 337–353.
- Tiesinga, P. (2002). Precision and reliability of periodically and quasiperiodically driven integrate-and-fire neurons. *Phys. Rev. E*, *65*, 041913.
- Tiesinga, P., Fellous, J. M., & Sejnowski, T. (2002a). Attractor reliability reveals deterministic structure in neuronal spike trains. *Neural Comput.*, *14*(7), 1629–1650.
- Tiesinga, P., Fellous, J. M., & Sejnowski, T. (2002b). Spike-time reliability of periodically driven integrate-and-fire neurons. *Neurocomputing*, *44–46*: 195–200.
- Veerman, J. (1989). Irrational rotation numbers. *Nonlinearity*, *2*(3), 419–428.
- Victor, J., & Purpura, K. (1996). Nature and precision of temporal coding in visual cortex: A metric-space analysis. *J. Neurophysiol.*, *76*(2), 1310–1326.

Received March 15, 2002; accepted July 5, 2002.