

SELF-SUSTAINED FIRING IN POPULATIONS OF INTEGRATE-AND-FIRE NEURONS*

CARL VANVREESWIJK[†] AND L. F. ABBOTT[†]

Abstract. A coupled population of simple integrate-and-fire model neurons is analyzed to determine the conditions that lead to stable firing sustained by all-to-all excitatory interactions. The possible temporal firing patterns are determined, and the initial conditions that produce them are discussed. For certain parameter values, two different patterns of firing are possible: one in which the firing is distributed among many groups of neurons that fire sequentially, and another in which the population breaks into two groups that fire alternately. The probabilities of different temporal firing patterns arising from random initial conditions and noise inputs are analyzed by computer simulation.

Key words. firing patterns, integrate and fire, neuronal population

AMS(MOS) subject classifications. 92C20, 58F40, 58F22

1. Introduction. Studies of model neural networks [2], [6] have focused attention on self-sustained firing in neural populations. In network associative memories, spatial patterns of self-sustained firing in populations of simple model neurons are used to represent stored memories and are the outcome of memory retrieval. The model neurons used in these networks are most often binary [7], [12], [13] or are based on a mean-field description [16]. The activity of the resulting network models is described in terms of an average firing rate [1]. The specific timing of action potentials is not computed, so the temporal distribution of firing and issues such as the synchronization or clumping of action potentials cannot be treated. This shortcoming has led to the consideration of associative memories in populations of integrate-and-fire neurons [4]. Integrate-and-fire neuron models [11] go beyond a simple firing rate description by predicting action potential firing times. Therefore, firing patterns, not just average firing rates, in neuronal populations can be studied. Complex firing patterns like those analyzed analytically here have been found in model simulations [3]. Other studies have revealed some of the richness and complexity exhibited by populations of integrate-and-fire neurons as well [8], [9].

The complexity of coupled systems of integrate-and-fire neurons has resulted in a scarcity of analytic results. Our work is inspired by recent results of Mirollo and Strogatz [14], who studied synchronization in populations of integrate-and-fire oscillators with all-to-all coupling. Mirollo and Strogatz proved that, for almost all initial conditions, the model oscillators tend to synchronized steady states. Related work on coupled oscillators has been done by Kuramoto [10]. In this paper, we consider populations of integrate-and-fire neurons that are *not* intrinsic oscillators. We look for steady-state firing that is not the result of any intrinsic oscillatory properties of individual neurons, but instead arises collectively as a cooperative phenomenon involving the entire population. We determine the types of self-sustained temporal firing patterns that can occur, and we indicate the conditions that must be satisfied for them to arise. Using computer simulations, we analyze the probabilities that certain temporal patterns of self-sustained firing appear and study the effect of noise.

* Received by the editors June 17, 1991; accepted for publication (in revised form) December 20, 1991. This research was supported by Department of Energy contract DE-AC0276-ER03230 and by National Institute of Mental Health grant MH46742.

[†] Physics Department and Center for Complex Systems, Brandeis University, Waltham, Massachusetts 02254.

2. The model. The simple integrate-and-fire model we consider uses a variable x to describe the degree of activation of a neuron. The activation variable satisfies the equation (with $\beta \geq 0$)

$$(2.1) \quad \frac{dx}{dt} = \alpha - \beta x$$

for $0 \leq x \leq 1$. If $x > 1$, the neuron fires, and x is reset to zero. The value 1 is thus the activation threshold. We consider populations in which all neurons are coupled to all other neurons [10] with excitatory couplings of equal strength. When a single neuron fires, the x values of all other neurons in the population are increased by a positive amount ϵ . If n neurons fire simultaneously, the x values of all the other neurons in the population are increased by an amount $n\epsilon$. In this way, firing of some neurons can induce other members of the population to fire. Mirollo and Strogatz [14] considered this model with $\alpha/\beta > 1$. In this case, the presence of the α -term in (2.1) causes individual neurons to fire repetitively even in the absence of coupling to other neurons. The role that coupling plays is to synchronize these oscillations [14].

We wish to consider self-sustained firing in a population of neurons that cannot fire without excitatory input from other neurons in the population. Therefore, we set $\alpha = 0$ in (2.1). The resulting model is a leaky integrate-and-fire model with excitatory couplings between neurons. With $\alpha = 0$, there is no intrinsic oscillation period to set the timescale for firing. (The model neuron will not oscillate if $\alpha/\beta < 1$. It is straightforward to extend our results to the case of nonzero α satisfying this bound, but since this adds complication and introduces no new features, we only consider $\alpha = 0$.) Instead, in a real system of neurons, the minimum time between neuronal firings would be determined by the transmission, integration, and response delays in the population. To represent this phenomenon, we let the model evolve in discrete timesteps equal to a typical delay time. In units of this basic timestep, the time variable takes on integer values $t = 0, 1, 2, \dots$. If a neuron is pushed above threshold and fires at time t , the excitation ϵ due to that firing will not be added to the activation of the other neurons until time $t + 1$. The discrete nature of the time evolution in the model plays an essential role in the results that we obtain.

For discrete time evolution, (2.1) can be integrated over one timestep. This has the effect of multiplying x by a factor

$$(2.2) \quad \lambda = \exp(-\beta).$$

A population of N neurons is described by the activation values x_i for $i = 1, 2, \dots, N$. Suppose that at time t there are $n(t)$ neurons with x values greater than 1 that are in the process of firing. At time $t + 1$, this firing raises the level of all the other neurons by an amount $n(t)\epsilon$. At the same time, the $n(t)$ neurons that fired at time t will drop to the $x = 0$ level. Thus, if a given neuron i is at excitation level $x_i(t)$ at time t , then at time $t + 1$

$$(2.3) \quad x_i(t + 1) = \lambda x_i(t) + n(t)\epsilon \quad \text{if } x_i(t) \leq 1$$

and

$$(2.4) \quad x_i(t + 1) = 0 \quad \text{if } x_i(t) > 1.$$

The number of firing neurons $n(t)$ is obtained by counting the total number of neurons

at time t with $x(t) > 1$

$$(2.5) \quad n(t) = \sum_{i=1}^N \Theta(x_i(t) - 1),$$

where Θ is the unit step function. The evolution of the model for discrete timesteps is determined by iterating (2.3) and (2.4) for all $i = 1, 2, \dots, N$. All of the neurons are updated in parallel.

3. Self-sustained firing patterns. We examine a population of N model neurons described by (2.3) and (2.4), focusing in particular on the possibility of self-sustained firing. Since the model neurons have no intrinsic ability to fire, they must continually receive excitation from other neurons in the population to increase their level of activation and ultimately fire themselves. Sustained firing, if it occurs, is a cooperative phenomenon involving the entire population. At every timestep, some firing must take place, otherwise the entire population will stop firing and ultimately sink to the $x = 0$ level.

Suppose that we initialize the population by assigning a distribution of activation levels $x_i(0)$ to the neurons for $i = 1, 2, \dots, N$. If firing is to occur at all, at least some of these neurons must be initialized above threshold, $x_i(0) > 1$ for some values of i . This causes initial firing to occur. If firing is to continue, the initial firing must raise some other neurons above threshold so that they can fire a time $t = 1$. Similarly, if this raises yet another group of neurons above the threshold, firing continues to occur at $t = 2$, and so on. If at any time no new neurons are raised above the threshold, then all firing stops forever.

The firing of $n(t)$ neurons at time t raises all neurons with sufficiently high activation levels above the threshold, causing them to fire at time $t + 1$. All neurons with activation levels $x_i(t)$ at time t satisfying

$$(3.1) \quad \frac{1 - n(t)\epsilon}{\lambda} < x_i(t) \leq 1$$

are raised above the threshold and subsequently fire at time $t + 1$. All these neurons then coalesce into a single group with activation level $x(t + 1) = 0$. Thus, the firing process can take a collection of neurons with different x values and unite them into a single group all having the same activation level. Because all neurons in such a group are treated equivalently, the group cannot subsequently break up. However, it can coalesce with another group if both groups are raised up past the threshold at the same time. Unless firing stops altogether, the population will settle into an equilibrium state through the coalescence of different groups until a situation is reached where the combining of groups stops and the groups fire sequentially without changing their sizes. For a stable periodic temporal pattern of firing to arise, each group must be large enough to raise the next group in the sequence above the threshold when it fires, but not so large that it raises more than one group above the threshold. Once this equilibrium has been attained with M different groups, the activity of the population will be periodic with period M or less. The only form of sustained firing in the model is periodic firing. However, many different M values and many partitionings of the N neurons into M groups are possible.

To begin our analysis, we derive limits on the number of neurons firing at any time t , $n(t)$, if firing is to continue. If firing occurs at time t , then, for it to continue to time $t + 1$, the number of neurons firing must at the very minimum be sufficient to

raise a neuron with the maximal activation $x(t) = 1$ past the threshold. This means that $\lambda x(t) + n(t)\epsilon$ must be greater than 1 for $x(t) = 1$ or equivalently $n(t) > (1 - \lambda)/\epsilon$. It is also essential that the number of neurons that fire at time t be less than an upper limit. After a particular group has fired, there must be enough neurons left over to provide the impulses to reactivate that group back up to the threshold so that it can fire again. If $n(t)$ neurons fire, there are $N - n(t)$ neurons left. For these to be capable of raising the $n(t)$ neurons back up to the threshold, we must have $(N - n(t))\epsilon > 1$. Combining these two limits gives the condition

$$(3.2) \quad \frac{1 - \lambda}{\epsilon} < n(t) < N - \frac{1}{\epsilon},$$

and from this bound we can derive the requirement

$$(3.3) \quad N\epsilon > 2 - \lambda,$$

which must be satisfied if sustained firing is to occur. Of course, this condition is only necessary, not sufficient.

4. Temporal patterns of self-sustained firing. If stable firing does arise in the model, the final periodic firing pattern consists of M groups of neurons firing sequentially at times $t, t + 1, t + 2, \dots$. We now derive bounds on the allowed values of M and constraints on the numbers of neurons in each of these groups.

Suppose that we have achieved sustained, periodic firing in M separate groups. For convenience, we reset the zero of time so that one of these groups, consisting of $n(1)$ neurons, fires at time $t = 1$. Immediately after that, a group of $n(2)$ neurons fires at time $t = 2$, and so on, until the last group containing $n(M)$ neurons fires at time $t = M$. At this point, if the firing is to continue, the first group of $n(1)$ neurons must fire again at time $M + 1$, the second group of $n(2)$ neurons must fire at time $M + 2$, and so on indefinitely. For the first group to fire at time $t = M + 1$, its activation level must be raised above the threshold by the firing of the other neurons. At time $t = 2$, the activation level for the first group is zero, since it has just fired at the previous time $t = 1$. At time $t = 3$, the level of group one is raised to $n(2)\epsilon$ by the firing of the second group. At time $t = 4$ the activation level is $\lambda n(2)\epsilon + n(3)\epsilon$ and so forth. For this group of neurons to fire again at time $M + 1$, we must therefore require that its activation level at that time be greater than the threshold

$$(4.1) \quad [\lambda^{M-2}n(2) + \lambda^{M-3}n(3) + \dots + \lambda n(M-1) + n(M)]\epsilon > 1.$$

By similar reasoning, the second group of neurons will fire at time $M + 2$ if

$$(4.2) \quad [\lambda^{M-2}n(3) + \lambda^{M-3}n(4) + \dots + \lambda n(M) + n(1)]\epsilon > 1.$$

Similar relations hold for the third, fourth, and subsequent groups up to the M th group, which requires that

$$(4.3) \quad [\lambda^{M-2}n(1) + \lambda^{M-3}n(2) + \dots + \lambda n(M-2) + n(M-1)]\epsilon > 1.$$

Adding these M inequalities together and using the fact that

$$(4.4) \quad \sum_{t=1}^M n(t) = N,$$

we find that we must require that

$$(4.5) \quad N\epsilon \sum_{p=0}^{M-2} \lambda^p = \frac{N\epsilon(1 - \lambda^{M-1})}{1 - \lambda} > M.$$

A lower limit on M is obtained by requiring that the groups do not fire prematurely. For example, the group that fired at time $t = 1$ should fire again at time $M + 1$. Therefore, at time M , its activation level must not be above the threshold. This gives the condition

$$(4.6) \quad [\lambda^{M-3}n(2) + \lambda^{M-4}n(3) + \cdots + \lambda n(M-2) + n(M-1)] \epsilon \leq 1.$$

Likewise, for the second group, we require that

$$(4.7) \quad [\lambda^{M-3}n(3) + \lambda^{M-4}n(4) + \cdots + \lambda n(M-1) + n(M)] \epsilon \leq 1,$$

and so on, until for the last group

$$(4.8) \quad [\lambda^{M-3}n(1) + \lambda^{M-4}n(2) + \cdots + \lambda n(M-3) + n(M-2)] \epsilon \leq 1.$$

Adding these M equations together gives the bound

$$(4.9) \quad N\epsilon \sum_{p=0}^{M-3} \lambda^p = \frac{N\epsilon(1 - \lambda^{M-2})}{1 - \lambda} \leq M.$$

Combining the two bounds derived, we have the condition

$$(4.10) \quad \frac{M}{1 - \lambda^{M-1}} < \frac{N\epsilon}{1 - \lambda} \leq \frac{M}{1 - \lambda^{M-2}},$$

which must be satisfied for self-sustained firing of M groups of neurons to be possible. In addition, all of the above inequalities for $n(1), n(2), \dots, n(M)$, as well as the fact that these must all be positive integers, place restrictions on the partitioning of the N neurons into M groups. In particular, some partitionings may be impossible even though the bound (4.10) is satisfied. Of course, we must always have $M \leq N$, but (3.2) sets an even more stringent upper limit on M because the minimum number of neurons in each of the M groups must be an integer greater than $(1 - \lambda)/\epsilon$. The partitioning conditions become less restrictive for large N . Note that, other than these partitioning restrictions, the limits on M depend only on the combination $N\epsilon$ and on the value of λ .

In Fig. 1 we have plotted the values of M allowed by (4.10) as a function of $1/N\epsilon$ for two values of λ . Firing in two groups, $M = 2$, is allowed for any value of $N\epsilon$ greater than two or equivalently for $0 < 1/N\epsilon < 0.5$. Firing in larger numbers of groups, $M > 2$, occurs over an ever-narrowing range of values of $1/N\epsilon$. The behavior of the population is quite different for no leakage, $\lambda = 1$, than with leakage, $\lambda = 0.9$. For example, with no leakage, ever-increasing values of M can be attained as $N\epsilon$ decreases to its minimum value consistent with sustained firing, 1. However, with leakage, $N\epsilon$ near 1 does not permit sustained firing. For $N\epsilon$ around 2, several values of M may be possible. For $\lambda < 1$, two states occur at large values of $N\epsilon$. One state has $M = 2$, and the other has a large M value. The large M value can be estimated, and we find the two possible sustained states for large $N\epsilon$ and $\lambda < 1$,

$$(4.11) \quad M = 2 \quad \text{or} \quad M \approx \frac{N\epsilon}{1 - \lambda}.$$




FIG. 1. The allowed values of M as a function of $1/N\epsilon$ for no leakage, $\lambda = 1.0$, and with leakage, $\lambda = 0.9$. The allowed values of M for sustained firing are the integers lying between the two curves drawn. For a given value of N , the M values may be further restricted by the partitioning requirements. These have not been indicated because they depend on N .

It is interesting to note that this simple model with leakage has the sort of bistable behavior characteristic of epilepsy. Epileptic activity is a complex dynamical behavior involving interplay between excitatory and inhibitory processes [15]. Our model is not intended to provide a complete description of this phenomenon. One stable firing pattern occurs in which the neuronal firing is well distributed in many groups (large M) firing at many different times. However, a second “epileptic” form of firing is also possible for the same value of $N\epsilon$. In this state ($M = 2$), the firing comes as near to synchronizing as possible, and the entire population fires in two large groups. In the absence of leakage ($\lambda = 1$), no similar two-state region exists.

5. The case where $\lambda = 1$. The model we are discussing simplifies considerably if we ignore leakage by setting $\lambda = 1$. This approximation is justified if the amount of leakage in a time of order 1 over the firing rate is small. For $\lambda = 1$, the bounds of the last section simplify, and we find that self-sustained firing can occur in M groups if

$$(5.1) \quad \frac{N\epsilon}{N\epsilon - 1} < M \leq \frac{2N\epsilon}{N\epsilon - 1}.$$

Until now, we have discussed the types of sustained firing that *can* occur, but we have not addressed the question of whether they *will* occur given a set of initial conditions. For the case where $\lambda = 1$, we can answer this second question. As soon as all the neurons have fired at least once, it is possible to determine whether the population will continue firing forever. We initialize the population by assigning activation levels to the N neurons. Starting at the time $t = 0$, $n(0)$ neurons fire, then $n(1)$, $n(2)$, and so on. Suppose that firing continues until some time T when all of the neurons have either fired or are presently firing

$$(5.2) \quad \sum_{t=0}^T n(t) \geq N.$$

Then, if the numbers of neurons that have fired previously or are presently firing satisfy

$$(5.3) \quad 0 < n(t) < N - \frac{1}{\epsilon}$$

for all $t \leq T$, firing continues forever.

To prove this statement, we note that at the time T the neurons are clustered into groups that fired at earlier times $T-1, T-2, \dots$, down to t_0 . Because (5.2) is an inequality, any neurons that fired before time t_0 will have fired more than once by the time T . The key to our analysis is that we know the activation levels of all these groups once they have fired. At time T , the groups are arranged as follows. A group of $n(T)$ neurons is above the threshold and in the process of firing. The first group of neurons below the threshold at this time is a group of $n(t_0)$ neurons that fired at some earlier time $t_0 < T$. At the next lower occupied activation level is a group of $n(t_0+1)$ neurons that fired at time t_0+1 . Next is a group of $n(t_0+2)$ neurons, and so on, until finally we find the group of $n(T-1)$ neurons that just fired at the previous time located at the zero activation level. The time T is arbitrary as long as it is large enough to satisfy condition (5.2). Note that we have used the time at which a group last fired as a label.

We will begin by showing that, if $0 < n(t_0) < N - 1/\epsilon$, then at least $n(t_0)$ neurons will fire at time $T+1$. To do this, we consider the activation level of the highest subthreshold group of $n(t_0)$ neurons at time $T+1$. This can be determined because we know that these neurons fired at time t_0 , so they must have occupied the lowest, $x = 0$ level at time t_0+1 . To determine the activation level at time $T+1$, we must simply count the number of inputs that this group has received from all the other neurons in the intervening time. The analysis is made simpler by the fact that we have taken $\lambda = 1$. Thus, for this group of $n(t_0)$ neurons,

$$(5.4) \quad x(T+1) = \epsilon \sum_{t=t_0+1}^T n(t).$$

Since

$$(5.5) \quad \sum_{t=t_0}^T n(t) = N,$$

we find that

$$(5.6) \quad x(T+1) = \epsilon \left(\sum_{t=t_0}^T n(t) - n(t_0) \right) = \epsilon (N - n(t_0)).$$

Since $n(t_0) < N - 1/\epsilon$, this shows that $x(T+1) > 1$, and the groups of $n(t_0)$ neurons that fired at time t_0 will fire again at time $T+1$.

Thus, we see that, if $n(t_0) < N - 1/\epsilon$, firing will continue until time $T+1$. Using an identical argument, we can show that firing will continue until time $T+2$ if $n(t_0+1) < N - 1/\epsilon$, and so on. Because (5.3) requires that $n(t) < N - 1/\epsilon$ for all $t \leq T$, this proof continues to work until we get to a time when we must use the condition $n(T+1) < N - 1/\epsilon$. Since this requirement for sustained firing is not included in the original conditions, we must prove that it is true.

The number of neurons firing at time $T+1$ is equal to the sum of all neurons that are raised above the threshold by the group of $n(T)$ neurons firing at time T . Using the time at which a group of neurons fired as a label, we can write this sum in the form

$$(5.7) \quad n(T+1) = \sum_{t=t_0}^{t_{\max}} n(t),$$

where t_{\max} is the largest time for which the activation of the corresponding group is above the threshold at time $T + 1$

$$(5.8) \quad x(T + 1) = \epsilon \sum_{t=t_{\max}+1}^T n(t) > 1.$$

In other words, the group of $n(t_{\max})$ neurons is the lowest lying group to be pulled above the threshold. Rearranging the sum in (5.8), we have

$$(5.9) \quad x(T + 1) = \epsilon \sum_{t=t_0}^T n(t) - \epsilon \sum_{t=t_0}^{t_{\max}} n(t) = \epsilon[N - n(T + 1)] > 1.$$

From this inequality, it follows that $n(T + 1) < N - 1/\epsilon$. This means that we can extend the proof of continued firing one more timestep. Furthermore, by repeating the above analysis, we can prove that $n(T + m) < N - 1/\epsilon$ for all m , and thus the proof of continued firing can be extended indefinitely.

6. The probability of sustained firing. Now that we have characterized the possible sustained firing states of the integrate-and-fire population, it is interesting to determine the probabilities that different final states arise from a random starting configuration. We assume that the initial values of the activation variables are chosen from some probability distribution to be specified. We discuss the probability $P(M)$ that the population ends up in a sustained firing pattern consisting of M groups. By definition, $P(0)$ is the probability that the population stops firing.

In principle, the probabilities $P(M)$ can be computed analytically, but the calculation is extremely tedious except when fairly small numbers of neurons are involved. Nevertheless, we outline how it is done. Suppose that initially $n(0)$ neurons are activated above threshold. At subsequent times $n(1), n(2), \dots$, neurons fire. Assume that firing continues until a time $T - 1$, satisfying

$$(6.1) \quad \epsilon \sum_{t=0}^{T-1} \lambda^{T-1-t} n(t) > 1.$$

(If it does not, then, of course, this initial configuration contributes to $P(0)$.) The time T defined by the above equation is the same as the T defined in (5.2). Furthermore, the set of firing numbers from time zero to time T , $\{n(0), n(1), \dots, n(T - 1)\}$, completely determines all subsequent firing. To see this, we note that any neurons that have not fired up until the time $T - 1$ have activation level at time T given by

$$(6.2) \quad x(T) = \lambda^T x(0) + \epsilon \sum_{t=0}^{T-1} \lambda^{T-1-t} n(t).$$

Because of (6.1), this is greater than 1, regardless of what the initial activation $x(0)$ is for this neuron. Thus, by the time T , all neurons will have fired. Using the reasoning of the last section, once all the neurons have fired, their activation levels are completely determined by the set $\{n(0), n(1), \dots, n(T - 1)\}$. Otherwise, all dependence on the initial distribution of activation values is lost by the time T . In particular, a group of neurons that fired a time t_0 have activation level at time T , given by

$$(6.3) \quad x(T) = \epsilon \sum_{t=t_0+1}^{T-1} \lambda^{T-1-t} n(t).$$

To calculate the probabilities $P(M)$, we must compute the probability of a given set of firings $\{n(0), n(1), \dots, n(T-1)\}$ arising from random initial conditions. The probability that the set $\{n(0), n(1), \dots, n(T-1)\}$ arises is just the probability that $n(0)$ neurons were initially above threshold, that $n(1)$ neurons have initial activations in the range

$$(6.4) \quad \frac{[1 - n(0)]\epsilon}{\lambda} < x(0) \leq 1,$$

$n(2)$ neurons have

$$(6.5) \quad \frac{[1 - (n(0)\lambda + n(1))]\epsilon}{\lambda} < x(0) \leq \frac{[1 - n(0)]\epsilon}{\lambda^2},$$

and so on. However, we must also determine what final firing pattern results from the set $\{n(0), n(1), \dots, n(T-1)\}$, and this is extremely tedious if large numbers of neurons are involved. Nevertheless, we have computed the probabilities $P(M)$ in this way for $N = 20$ and $\lambda = \frac{1}{2}$ and obtained excellent agreement with computer simulation results.

For our computer simulations, we have determined the initial state of the population by choosing the initial activation levels for the N neurons randomly and uniformly in the range $0 \leq x(0) \leq x_{\max}$. When this is done, the average number of neurons initialized above threshold is

$$(6.6) \quad \langle n(0) \rangle = \frac{N(x_{\max} - 1)}{x_{\max}}.$$

A natural choice for x_{\max} would be $1 + \epsilon$, giving $\langle n(0) \rangle = N\epsilon/(1 + \epsilon)$. However, from the bound (3.2), we must have $n(0) > (1 - \lambda)/\epsilon$ if any firing is to occur at time $t = 1$. Therefore, to help ensure that at least some firing occurs, we have chosen x_{\max} so that

$$(6.7) \quad \langle n(0) \rangle = \frac{1 - \lambda}{\epsilon} + \frac{N\epsilon}{1 + \epsilon}.$$

The results of the computer simulations are shown in Fig. 2. These were obtained by performing 1,000 repetitions, starting each time with initial activations chosen randomly from the distribution discussed above. We have considered two cases, $N = 20$ and $N = 100$, and taken $N\epsilon$ in the range between 1 and 3, where Fig. 1 shows that the largest range of M values occurs. We have considered no leakage, $\lambda = 1$, and leakage corresponding to $\lambda = 0.9$. As expected, $P(0) = 1$ for M outside the bounds of (4.10). $P(0)$ generally decreases as $N\epsilon$ increases. However, the decrease is not monotonic, most likely due to the difficulties of partitioning the N neurons into M groups for certain values. Figure 2 also shows that $P(M)$ decreases with increasing M . The results for $N = 20$ and $N = 100$ are quite similar, except, of course, that larger N allows for large values of M . In addition, the smaller λ value clearly makes it more difficult to produce states with large numbers of firing groups M .

Using random initial conditions with a uniform probability distribution, it is not very likely that sustained firing with the highest M values will be found. However, if a more uniform distribution of initial activations is used, these final states can be attained. In our simulations, we have seen all of the states allowed by Fig. 1, starting with suitable initial conditions.




FIG. 2. The probability $P(M)$ of sustained firing with M firing groups as a function of $N\epsilon$. $P(0)$ is the probability that firing stops. Initial conditions were chosen randomly over 1,000 trials. The figures at left are for $N = 20$ neurons, and at right for $N = 100$.

7. The effect of noise inputs. In the last section, we initialized the population of integrate-and-fire neurons randomly. Here, we consider a more realistic example in which the random initial state is followed by a period during which external inputs enter the system. We then examine the final state of the system after the inputs are removed. Specifically, during the time that the inputs are active, we modify the evolution equations (2.3) and (2.4) so that

$$(7.1) \quad x_i(t+1) = \lambda x_i(t) + n(t)\epsilon + r_i(t) \quad \text{if } x_i(t) \leq 1$$

and

$$(7.2) \quad x_i(t+1) = 0 \quad \text{if } x_i(t) > 1,$$

where r_i is the external input. In our simulations, we have taken this input to be a random number in the range

$$(7.3) \quad -\sigma < r_i(t) < \sigma$$

with the random choice made independently for each neuron and at each timestep.

The two examples in Fig. 3 show the variety of impacts noise inputs can have. One distinctive feature of the model in the absence of noise is that, once a firing group has formed, it can never break up because all of its members evolve identically.




FIG. 3. *The effect of noise inputs on $P(M)$. Starting from a random initial configuration, the population evolved in the presence of noise at the given strength for 8,000 timesteps. The noise input was then turned off and $P(M)$ was computed using 1,000 repetitions. Both figures correspond to a population of $N = 20$ neurons. Strong noise with $\sigma = 0.2$ greatly reduces the probability of sustained firing especially for large M when $\lambda = 0.9$. Noise with $\sigma = 0.1\epsilon$ had little effect when $\lambda = 0.9$, but, as shown, it increases the probability of sustained firing in large numbers of groups for $\lambda = 1$.*

With noise inputs present, this is no longer true. Thus, one effect of the noise is to separate neurons within a group. On the other hand, the firing process forces neurons that fire together to coalesce into a single group. In the presence of noise, there are two opposing tendencies: noise breaking groups apart, and firing putting them back together again. If the noise is small enough, the coalescing tendencies of the firing win, and noise has little effect. For example, if the noise level $\sigma = 0.1\epsilon$ is applied to $N = 20$ neurons with $\lambda = 0.9$, the results are virtually identical to those shown in Fig. 2 without noise. (Over the range shown, this corresponds to σ in the range 0.005 to 0.015.) For a given noise level, firing patterns with small M values are more stable than those with large M because the groups are more widely separated. This can be seen in Fig. 3, where a noise level of $\sigma = 0.2$ for $\lambda = 0.9$ reduces sustained firing and completely eliminates sustained states with $M > 4$ (compare with Fig. 2). In general, high-noise levels reduce sustained firing and particularly affect large M states.

The results for $\lambda = 1$ in Fig. 3 are considerably different than those of Fig. 2. This is true, despite the fact that the noise level $\sigma = 0.1\epsilon$ had virtually no impact when $\lambda = 0.9$. Thus, we see that a given level of noise has more impact at low leakage. This is due to a second important effect of noise. In the absence of noise inputs, the model we have been discussing has the property that firing will stop forever if it does not occur at every single timestep. If the noise level is large enough for a given λ value, then noise can reinitiate firing. Thus, a population that misses one timestep of firing can, nevertheless, continue to fire. This clearly makes sustained firing more likely, as we can see by comparing the $\lambda = 1$ results in Figs. 3 and 2. For noise to be able to resume firing, it must, at the very least, be strong enough to elevate a neuron sitting right at the threshold at $x = 1$ above the threshold in the next timestep. This requires then that $\lambda + \sigma > 1$, or equivalently $\sigma > 1 - \lambda$. For $\lambda = 1$, clearly any amount of noise satisfies this condition. Our results indicate that noise stronger than this limit can considerably enhance the ability of the system to find sustained firing states, especially those with large M values. However, even higher levels of noise would start destroying the sustained firing states as seen in Fig. 3 for $\lambda = 0.9$.

8. Discussion. The model we have considered is extremely simple, and yet we have found that it not only supports self-sustained firing but that it contains a surprisingly rich spectrum of firing states. The simplicity of the model has enabled us to determine what types of states are allowed and under what conditions they can arise. The fact that the firing process collapses individual neurons into firing groups means that the ultimate fate of any initial state depends only on the numbers of neurons that fire during the first $T - 1$ timesteps, where T is defined by (5.2) or (6.1). This means that the probability of sustained firing resulting from a statistically specified initial state can be computed, although in most cases we have resorted to computer simulation to produce these results due to the length of the calculations involved. Noise inputs of moderate strength, especially at low leakage, result in an increase in the probability of achieving self-sustained firing involving a large number of firing groups. Higher noise levels destroy the sustained firing states with large M values.

We have experimented with some variations of the model, for example, by letting the delay time or the excitation ϵ vary from neuron to neuron. The basic behavior does not change in these cases, and, in particular, we still find periodic sustained firing states. It appears that noise is the only simple way to get nonperiodic firing.

Acknowledgments. The authors thank Steve Strogatz, Charlie Marcus, and Mark Schnitzer for helpful and encouraging comments.

REFERENCES

- [1] L. F. ABBOTT, *Firing rate models for neural populations*, in *Neural Networks: From Biology to High-Energy Physics*, O. Benhar, C. Bosio, P. Del Giudice, and E. Tabet, eds., ETS, Editrice, Pisa, 1991.
- [2] D. J. AMIT, *Modelling Brain Function*, Cambridge University Press, New York, 1989.
- [3] ———, private communication, 1991.
- [4] D. J. AMIT, M. R. EVANS, AND M. ABELES, *Attractor neural networks with biological probe records*, *Network*, 1 (1990), pp. 381–406.
- [5] M. A. DICHTER AND G. F. AYALA, *Cellular mechanisms of epilepsy: A status report*, *Science*, 237 (1987), pp. 157–164.
- [6] J. HERTZ, *Introduction to the Theory of Neural Computation*, Addison-Wesley, Reading, MA, 1990.
- [7] J. J. HOPFIELD, *Neural networks and systems with emergent selective computational abilities*, *Proc. Nat. Acad. Sci. U.S.A.*, 79 (1982), pp. 2554–2258.
- [8] J. P. KEENER, F. C. HOPPENSTEADT, AND J. RINZEL, *Integrate-and-fire models of nerve membrane response to oscillatory input*, *SIAM J. Appl. Math.*, 41 (1981), pp. 503–517.
- [9] B. W. KNIGHT, *Dynamics of encoding in a population of neurons*, *J. Gen. Physiol.*, 59 (1972), pp. 734–766.
- [10] Y. KURAMOTO, *Collective synchronization of pulse-coupled oscillators and excitable units*, *Phys. D*, 50 (1991), pp. 15–30.
- [11] L. LAPICQUE, *Recherches quantitatives sur l'excitation électrique des nerfs traitée comme une polarisation*, *J. Physiol. Pathol. Gen.*, 9 (1907), pp. 620–635.
- [12] W. A. LITTLE, *The existence of persistent states in the brain*, *Math. Biosci.*, 19 (1974), pp. 101–120.
- [13] W. S. MCCULLOCH AND W. PITTS, *A logical calculus of the ideas immanent in nervous activity*, *Bull. Math. Biophys.*, 5 (1943) pp. 115–133.
- [14] R. E. MIROLLO AND S. H. STROGATZ, *Synchronization of pulse-coupled biological oscillators*, *SIAM J. Appl. Math.*, 50 (1990), pp. 1645–1662.
- [15] C. S. PESKIN, *Mathematical Aspects of Heart Physiology*, Courant Institute of Mathematical Sciences, New York University, New York, 1975, pp. 268–278.
- [16] H. R. WILSON AND J. D. COWAN, *Excitatory and inhibitory interactions in localized populations of model neurons*, *Biophys. J.*, 12 (1974), pp. 1–24.