Optimal Stochastic Control of Linear systems Driven by Neural Dynamics

Levi Burner

December 25, 2020

Contents

1 Introduction **2**

2 Itô versus Stratonovich: A Simple Case Where They are Equal **2**
2.1 Numerical Verification 3

3 Neuron Model **4**

4 Controlling Membrane Potential **6**

5 The LTI Extension **10**

6 Discussion **14**
6.1 Controlling Membrane Potential 14
6.2 Controlling an LTI System 14
6.3 Future Work 14

7 Conclusion **15**

8 Appendix: Codes **17**
8.1 Numerical simulation of SDE $Y = B\,dB$ 17
8.2 Numerical simulation of SDE $Y = B\,dt + t\,dB$ 17

1 Introduction

Neuromorphic engineering offers significant advantages in terms of power-consumption for performing computations related to autonomous systems. The fundamental difference of neuromorphic approaches from traditional digital and analog techniques is the use of spiking neurons as the computational engine. These neurons have internally continuous dynamics except for the "spikes" they receive and send. The literature often models these "spike trains" as Poisson counting processes and for purposes of analysis further approximates the dynamics with diffusion processes.

Since a spiking computation scheme will naturally have spiking outputs, it is of interest to consider efficient control schemes for typical dynamical systems driven by such random impulses.

The literature contains a few works on this subject. In [1], optimal control signals for the diffusion approximations parameter $\lambda(t)$ is derived for controlling a leaky integrate-and-fire (LIF) neuron's membrane potential. Further, it is extended to control a simple LTI system meant to model the human arm. The approach relies on calculus of variations to compute a solution to the minimum variance and fixed expected endpoint problem. This approach is extended in [2] to a large class of LTI systems. The work (and unfortunately many sentences), borrows from [1], but the results are original.

An alternative approach is given in [3]. This work attempts to explicitly control the timing of spike firing. There, closed-loop solutions are derived and solved using dynamic programming and open-loop solutions are found using the Hamilton-Jacobi-Bellman equation to control the probability density function. These approaches are very interesting, and appear to be more flexible than the approaches discussed in this project report. They will be the subject of future work.

This paper will discuss in detail the derivations in [1, 2]. The original goal was to extend the results to the LTV case, but close inspection of [2] showed that was out of scope for the time being. Additionally, a brief discussion of stochastic calculus used in this paper is provided. Surprisingly, the question of Itô versus Stratonovich never actually enters [4, 5, 6]. This is explored up with some basic numerical simulations.

This work is motivatved by current interest in using neuromorphic computation for Micro Air Vehicles (MAV's). Recent work does not consider stochastic control theory [7, 8]. These two works respectively focus on configuring neurons to perform explicit "digital" computations and genetic algorithms to find appropriate neural organization. This paper may lead to an alternative approach based in stochastic control theory that could allow systematic design and rigorous justification of specific neuromorphic control implementations.

2 Itô versus Stratonovich: A Simple Case Where They are Equal

The original intent of this project was to solve some stochastic optimal control problem's using Stratonovich Calculus instead of the Itô calculus. However, upon closer inspection it became apparent that the controversy never enters due to the neuron model used. Consider a Stochastic Difference Equation (SDE) of the following form.

$$dX = f(t, X, B)\, dt + g(t)\, dB \tag{1}$$

Where f and g satisfy appropriate conditions for existence and uniqueness of solutions. By inspection it can be seen that only the Stieltjes integral is required to interpret the SDE.

$$X(t) - X(t_0) = \int_{t_0}^{t} f(t, X, B)\, dt + \int_{t_0}^{t} g(t)\, dB(t) \tag{2}$$

It can be seen that in each integral, only the integrand or the integrator respectively are stochastic. Thus, a slight extension of the Riemann and the Stieltjes integral respectively can be used to interpret the solutions to the SDE's. Further, since $g(t)$ is not stochastic, the Itô and Stratonovich interpretations should be equivalent.

2.1 Numerical Verification

Consider a simple numerical calculation demonstrating the difference in two situations. Consider the following SDE which does not conform to the situation discussed above.

$$dX = B\, dB \tag{3}$$

The Itô solution is trivially verified to be $(B(t)^2 - t)/2$ while the Stratonovich solution is $B(t)^2/2$. We compare the results of numerically evaluating the Itô integral (using the Euler–Maruyama method) with the closed form Stratonovich integral. Because the SDE does not conform with the conditions discussed above, the Itô solutions and the Stratonovich solutions are different, see Figure 1 for results. A listing of the simulation code is given in the appendix.

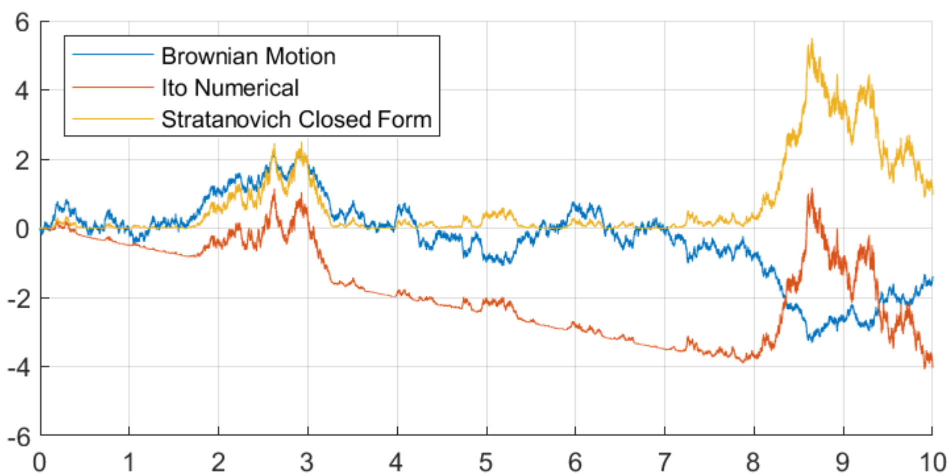

Figure 1: Itô and Stratonovich solutions to the SDE $Y = B\, dB$ for a particular sample path of Brownian motion. As can be seen the solutions are quite different.

Meanwhile, consider a simple case corresponding to the above discussion. The simplest example is $Y = B\, dt + t\, dB$. It can be easily shown that the Itô and Stratonovich solutions are tB for the initial condition $Y(0) = 0$.

This is confirmed by the numeric simulation below where the numeric evaluation of the Itô integral gives the same result as the analytic Stratonovich solution as shown in Figure 2.

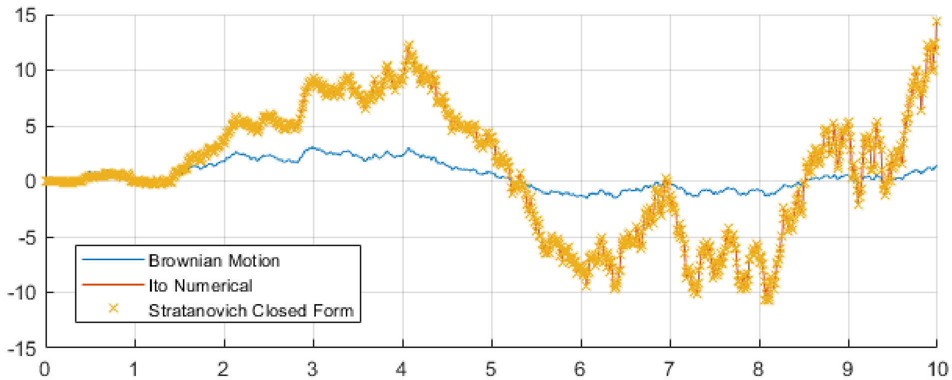

Figure 2: Itô and Stratonovich solutions to the SDE $Y = B\,dt + t\,dB$ for a particular sample path of Brownian motion. As can be seen the solutions are the same!

3 Neuron Model

The following derivations will closely follow the work in [1, 2]. Both use the following leaky integrate-and-fire approximation of neural dynamics.

$$\begin{aligned} dV(t) &= -\frac{1}{\gamma}[V(t) - V_{\text{rest}}]\,dt + dI_{\text{syn}}(t) \\ I_{\text{syn}}(t) &= a\sum_{i=1}^{p} I_i^e(t) - b\sum_{j=1}^{q} I_j^i(t) \end{aligned} \tag{4}$$

Where I_i^e is the i'th excitatory synaptic input and I_j^i is the j'th inhibatory input. For simplicity we assume the constants have the following properties $V_{rest} = 0$, $a = b$, and $p = q$. This corresponds to the case were excitatory and inhibitory inputs have equal strength and there are the same number of excitatory and inhibitory inputs.

As is usual, this simple continuous model holds until $V(t)$ cross a threshold V_{thresh} from below. Then the neuron output's a spike, and $V(t)$ is reset to V_{rest}.

As explained in [1] the above model can be approximated by a diffusion approximation. This relaxes the constraint on $I_{\text{syn}}(t)$ of depending on discontinuous Poisson random variables and instead approximates their effect using Brownian motion. This is reasonable as long as the neurons dynamics are “slower” than the incoming spikes. Of course this assumption is slightly dubious if all neurons are to have similar dynamics, in such a case the inter-spike interval of the incoming spike train would be similar to the out-going inter-spike interval and so the approximation no longer seems reasonable.

Regardless, moving forward the difference equation for the diffusion approximation for the synaptic current is given below.

$$dI_{\text{syn}}(t) = a\lambda(t)\,dt + a\lambda(t)^{\alpha}(t)\,dB(t) \tag{5}$$

The parameter α determines the nature of the spiking neuron. In [1], $\alpha > 0.5$ is dubbed "supra-poisson" meaning the input has greater variance than a Poisson distribution. $\alpha = 0.5$ is the poisson case. And $\alpha < 0.5$ is the sub-Poisson case. For our analysis we will focus on the case of $\alpha > 0.5$ else the control problem is degenerate and the solution is a δ function [1]. Further, we will also assume $\lambda(t) \geq 0\ \forall t$ as this simplifies the notation significantly (and turns out to be valid).

Substitution of (5) into (4) gives the following expression for the membrane potential (recall it was assumed $V_{\text{rest}} = 0$).

$$dV(t) = -\frac{1}{\gamma}V(t)\,dt + a\lambda(t)\,dt + a\lambda(t)^{\alpha}(t)\,dB(t) \tag{6}$$

Of interest is an analytic solution for $V(t)$. Since this is a stochastic differential equation, either an Itô or Stratonovich interpretation is standard. However, recalling the discussion in Section 2 the closed form solution for V requires only the Stieltjes integral and thus the Itô and Stratonovich interpretations are equivalent. More clearly, the partial derivatives with respect to a second parameter x are zero since the solution $V(t)$ is not dependent on any such parameter. Thus, the corresponding term's in Itô's Lemma (shown below) are zero.

$$dV(t) = \left[\frac{\partial V(t)}{\partial t} + \frac{\partial V(t)}{\partial x} + \frac{1}{2}\frac{\partial^2 V(t)}{\partial x^2}\right]dt + \frac{\partial V(t)}{\partial x}\,dB(t) \tag{7}$$

Thus, we are left with $dX(t) = \frac{\partial X(t)}{\partial t}\,dt$, which is the same as the Stratonovich rule and ordinary calculus.

Therefore, we have that the solution for $V(t)$ can be found using ordinary linear system theory.

$$V(t) = A(t)V(t)\,dt + B(t)u(t)\,dB(t) \tag{8}$$

Then

$$V(t) = \Phi_A(t,t_0)V(t_0) + \int_{t_0}^{t}\Phi_A(t,s)u(t)\,dB(t) \tag{9}$$

Assuming without-loss-of-generality that $V_{t_0} = 0$ and recognizing by inspection that $\Phi(t,s) = a\exp(-(t-s)/\gamma)$, $u_1(t) = \lambda(t)$, and $u_2(t) = \lambda^{\alpha}(t)$ we see that only the integral terms are left.

$$V(t) = a\int_{t_0}^{t} e^{-(t-s)/\gamma}\lambda(s)\,ds + a\int_{t_0}^{t} e^{-(t-s)/\gamma}\lambda(s)^{\alpha}\,dB(s) \tag{10}$$

The result is easily verified using the Lebinz Integral Rule.

$$\begin{aligned}
dV(t) &= \frac{\partial V(t)}{\partial t}\,dt \\
&= \frac{\partial}{\partial t}\left[a\int_{t_0}^{t} e^{-(t-s)/\gamma}\lambda(s)\,ds + a\int_{t_0}^{t} e^{-(t-s)/\gamma}\lambda(s)^{\alpha}\,dB(s)\right]dt \\
&= \left[ae^{-(t-t)/\gamma}\lambda(t) - \frac{a}{\gamma}\int_{t_0}^{t} ae^{-(t-s)/\gamma}\lambda(s)ds + ae^{-(t-t)/\gamma}\lambda(t)^{\alpha}\frac{\partial dB(t)}{\partial t} - \frac{a}{\gamma}\int_{t_0}^{t} e^{-(t-s)/\gamma}\lambda(s)\,dB(s)\right]dt \\
&= -\frac{1}{\gamma}V(t)\,dt + a\lambda(t)\,dt + a\lambda(t)^{\alpha}\,dB(t)
\end{aligned} \tag{11}$$

Where sum of the integral terms were recognized to be equal to $-V(t)/\gamma$ and a somewhat dubious heuristic was used to conclude that the Lebinz Integral Rule could be used when the integrator depends on on the term being differentiated by.

Thus, we conclude with this nerual model a solution for $V(t)$ can be conveniently be found. It is independent of the choice of stochastic calculus and simple forms of $\lambda(t)$ could admit closed form solutions of the integrals dependent on the $dB(t)$ integrator.

4 Controlling Membrane Potential

We are interested in driving systems using the output spikes of the leaky integrate-and-fire approximation of the neuron. In particular, is it interesting to attempt control of the firing time. Since the neuron fires when $V(t) = V_{\text{thresh}}$ a reasonable heuristic is to attempt to control the firing time using the following minimum variance, fixed endpoint optimal control problem where the dynamics are subect to (6) with initial condition $V(t_0) = 0$.

$$\begin{aligned} &\min_{\lambda} \operatorname{var}(V(T_f)) \\ &\mathbf{E}\, V(T_f) = V_{\text{thresh}} \end{aligned} \tag{12}$$

The intuition is that since $V(t_0) = 0$, the optimal solution will attempt to raise the membrame potential to the the threshold voltage. However, since the variance is to be minimized, the optimal control will raise the membrane potential in a controlled manner, such that V_{thres} is reached at T_f in a smooth manner.

We follow the derivation of the optimal control signal $\lambda^*(t)$ using the technique in [1] with the details expanded on significantly.

Consider, since by definition, Brownian motion satisfies $B(t) - B(s) \sim N(0, t-s)$ we have that $\mathbf{E}\, dB(t) = \mathbf{E} \lim_{\delta\to 0} [B(t+\delta) - B(t)] = \lim_{\delta\to 0} \mathbf{E}\, [B(t+\delta) - B(t)] = 0$. Thus the mean of $V(t)$ is simple to calculate.

$$\begin{aligned} \mathbf{E}\, V(t) &= \mathbf{E}\left[a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)\, ds + a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)^{\alpha}\, dB(s)\right] \\ &= a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)\, ds + a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)^{\alpha}\, \mathbf{E}\, [dB(s)] \\ &= a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)\, ds \end{aligned} \tag{13}$$

Therefore the endpoint constraint on the control λ is simply in the form of a convolution with a first order low pass filter.

$$a \int_{t_0}^{T_f} e^{-(T_f - s)/\gamma} \lambda(s)\, ds = V_{\text{thresh}} \tag{14}$$

With the first moment understood, we focus on the second moment in order to find the variance.

$$
\begin{aligned}
\mathbf{E}\, V^2(t) &= \mathbf{E}\, \left[a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)\, ds + a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)^\alpha \, dB(s) \right]^2 \\
&= \mathbf{E}\, \left(a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)\, ds \right)^2 \\
&+ \mathbf{E}\, 2 \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)\, ds \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)^\alpha \, dB(s) \\
&+ \mathbf{E}\, \left(a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)^\alpha \, dB(s) \right)^2
\end{aligned} \tag{15}
$$

The first term is deterministic so the expectation can be dropped, the cross terms involves the expectation of an integral with orthogonal increments, so the expectation is zero, and the last term is actually interesting. Next, we combine the moments to get the variance.

$$
\begin{aligned}
&\operatorname{var}(V(t)) = \mathbf{E} V^2(t) - \mathbf{E} V(t)^2 \\
&= \left(a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)\, ds \right)^2 + \mathbf{E}\, \left(a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)^\alpha \, dB(s) \right)^2 - \left(a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)\, ds \right)^2 \\
&= \mathbf{E}\, \left(a \int_{t_0}^{t} e^{-(t-s)/\gamma} \lambda(s)^\alpha \, dB(s) \right)^2 \\
&= \mathbf{E}\, a^2 \int_{t_0}^{t} \int_{t_0}^{t} e^{-(t-s_1)/\gamma} \lambda(s_1)^\alpha e^{-(t-s_2)/\gamma} \lambda(s_2)^\alpha \, dB(s_1)\, dB(s_2)
\end{aligned} \tag{16}
$$

Recall, brownian motion is defined to have independent increments, e.g. the expectation of differences is zero unless the differences start from the same time. Therefore, if $s_1 \neq s_2$ then $\mathbf{E}\left[dB(s_1)dB(s_2)\right] = 0$. However, if $s = s_1 = s_2$ then $\mathbf{E}\left[dB(s)^2\right] = dt$ due to the quadratic variation of brownian motion (see ENEE 762 Lecture 6 materials for proof). Thus, the double integral in (17) collapses to a single integral giving a nice form for the variance.

$$
\begin{aligned}
\operatorname{var}(V(t)) &= a^2 \int_{t_0}^{t} \int_{t_0}^{t} e^{-(t-s_1)/\gamma} \lambda(s_1)^\alpha e^{-(t-s_2)/\gamma} \lambda(s_2)^\alpha \, \mathbf{E}\left[dB(s_1)\, dB(s_2)\right] \\
&= a^2 \int_{t_0}^{t} e^{-2(t-s)/\gamma} \lambda(s)^{2\alpha}\, ds
\end{aligned} \tag{17}
$$

We now have a fully deterministic fixed endpoint minimum variance optimal control problem. As shown in [1] we can solve this with direct application of the calculus of variations. The problem is restated below.

$$
\begin{aligned}
&\min_{\lambda} \left(a^2 \int_{t_0}^{T_f} e^{-2(T_f-s)/\gamma} \lambda(s)^{2\alpha} ds \right) \\
&\text{s.t. } a \int_{t_0}^{T_f} e^{-(T_f-s)/\gamma} \lambda(s)\, ds = V_{\text{thresh}}
\end{aligned} \tag{18}
$$

Consider adding perturbation η such that the constraint is still satisfied. E.g. let $\eta(t)$ satisfy

$$
\begin{aligned}
& a\int_{t_0}^{T_f} e^{-(T_f-s)/\gamma}\left(\lambda(s)+\eta(s)\right)\,ds = a\int_{t_0}^{T_f} e^{-(T_f-s)/\gamma}\lambda(s)\,ds \\
\iff & \int_{t_0}^{T_f} e^{-(T_f-s)/\gamma}\eta(s) = 0
\end{aligned}
\tag{19}
$$

Then by the standard approach in calculus of variations, for all such η and for ϵ in neighborhood of zero, the solution to the minimization problem, $\lambda^*(s)$ satisfies the following.

$$
\left(a^2\int_{t_0}^{T_f} e^{-2(T_f-s)/\gamma}\lambda^*(s)^{2\alpha}ds\right) \leq \left(a^2\int_{t_0}^{T_f} e^{-2(T_f-s)/\gamma}\left(\lambda^*(s)+\epsilon\eta(s)\right)^{2\alpha}ds\right) \tag{20}
$$

Thus using the first order for conditions of optimality, the derivative of the right hand side with respect to ϵ, evaluated at zero, is zero.

$$
\begin{aligned}
& \frac{d}{d\epsilon}\left(a^2\int_{t_0}^{T_f} e^{-2(T_f-s)/\gamma}\left(\lambda^*(s)+\epsilon\eta(s)\right)^{2\alpha}ds\right)\Bigg|_{\epsilon=0} = 0 \\
\iff & \int_{t_0}^{T_f} e^{-2(T_f-s)/\gamma}\,2\alpha\left(\lambda^*(s)+\epsilon\eta(s)\right)^{2\alpha-1}\eta(s)ds\Bigg|_{\epsilon=0} = 0 \\
\iff & \int_{t_0}^{T_f} e^{-2(T_f-s)/\gamma}\lambda^*(s)^{2\alpha-1}\eta(s)ds = 0
\end{aligned}
\tag{21}
$$

By observing (19) and (21) both have conditions equal to zero we can set the equations equal to each other.

$$
\int_{t_0}^{T_f} e^{-(T_f-s)/\gamma}\eta(s)ds = \int_{t_0}^{T_f} e^{-2(T_f-s)/\gamma}\lambda^*(s)^{2\alpha-1}\eta(s)ds \tag{22}
$$

Since this holds for a fairly general class of η we then argue that the integrands components independent of η are equal up to a scaling factor (I am unsure of how to show this rigorously).

$$
\beta e^{-(T_f-s)/\gamma} = e^{-2(T_f-s)/\gamma}\lambda^*(s)^{2\alpha-1} \tag{23}
$$

By applying some algebraic tricks this can be related to the constraint equation to solve for β. At this point the analysis requires the hypothesis that $\alpha > 0.5$.

$$
\begin{aligned}
& \beta e^{-(T_f-s)/\gamma} = e^{-2(T_f-s)/\gamma}\lambda^*(s)^{2\alpha-1} \\
\iff & \beta e^{(T_f-s)/\gamma} = \lambda^*(s)^{2\alpha-1} \\
\iff & \beta^{1/(2\alpha-1)}e^{-((T_f-s)/\gamma)(-1/(2\alpha-1))} = \lambda^*(s) \\
\iff & a\beta^{1/(2\alpha-1)}e^{-((T_f-s)/\gamma)(2\alpha-2)/(2\alpha-1)} = a\lambda^*(s)e^{-(T_f-s)/\gamma}
\end{aligned}
\tag{24}
$$

Where the right hand side is recognized as the integrand in the constraint of the optimization problem. Thus integrating both sides reveals a relation that can be solved for β. Further, considering the second to last simplification above, we see that once β is known, the optimal control λ^* will be known in a closed form.

$$a\int_{t_0}^{T_f} \beta^{1/(2\alpha-1)} e^{-((T_f-s)/\gamma)(2\alpha-2)/(2\alpha-1)} ds = a\int_{t_0}^{T_f} \lambda^*(s) e^{-(T_f-s)/\gamma} ds = V_{\text{thresh}} \tag{25}$$

While, the left hand side at first appears to contain a complicated integral, closer inspection reveals almost all terms are simple multipliers. Finding an expression for β is then a straightforward, but tedious exercise.

$$\begin{aligned} & a\int_{t_0}^{T_f} \beta^{1/(2\alpha-1)} e^{-((T_f-s)/\gamma)(2\alpha-2)/(2\alpha-1)} ds = V_{\text{thresh}} \\ \iff & a\beta^{1/(2\alpha-1)} \exp\left(\frac{T_f}{\gamma}\frac{2\alpha-1}{2\alpha-2}\right)\left[\gamma\frac{2\alpha-1}{2\alpha-2}\exp\left(\frac{s}{\gamma}\frac{2\alpha-2}{2\alpha-1}\right)\right]_{t_0}^{T_f} = V_{\text{thresh}} \\ \iff & \beta^{1/(2\alpha-1)} = \frac{V_{\text{thresh}}}{a\gamma}\frac{2\alpha-2}{2\alpha-1}\left(1-\exp\left(-\frac{T_f-t_0}{\gamma}\frac{2\alpha-2}{2\alpha-1}\right)\right) \end{aligned} \tag{26}$$

Recalling that (24) gave us a form for λ^* we subsitute in the expression for β. Recall that the following holds for $\alpha > 0.5$.

$$\lambda^*(t) = \frac{V_{\text{thresh}}}{a\gamma}\frac{2\alpha-2}{2\alpha-1}\left(1-\exp\left(-\frac{T_f-t_0}{\gamma}\frac{2\alpha-2}{2\alpha-1}\right)\right)^{-1}\exp\left(\frac{T_f-t}{\gamma(2\alpha-1)}\right) \tag{27}$$

Inspection reveals there is a singularity at $\alpha = 1$. In which case the $2\alpha-2$ terms go to zero. The solution can be found using L'Hôpital's rule.

$$\begin{aligned} & \lim_{\alpha\to1}\frac{V_{\text{thresh}}}{a\gamma}\frac{2\alpha-2}{2\alpha-1}\left(1-\exp\left(-\frac{T_f-t_0}{\gamma}\frac{2\alpha-2}{2\alpha-1}\right)\right)^{-1}\exp\left(\frac{T_f-t}{\gamma(2\alpha-1)}\right) \\ =& \left[\lim_{\alpha\to1}\frac{V_{\text{thresh}}}{a\gamma}\frac{2\alpha-2}{2\alpha-1}\left(1-\exp\left(-\frac{T_f-t_0}{\gamma}\frac{2\alpha-2}{2\alpha-1}\right)\right)^{-1}\right]\exp\left(\frac{T_f-t}{\gamma}\right) \\ =& \left[\lim_{\alpha\to1}\frac{V_{\text{thresh}}}{a\gamma}2\left[2\left(1-\exp\left(-\frac{T_f-t_0}{\gamma}\frac{2\alpha-2}{2\alpha-1}\right)\right)\right.\right. \\ & \left.\left.+(2\alpha-2)\left(\exp\left(-\frac{T_f-t_0}{\gamma}\frac{2\alpha-2}{2\alpha-1}\right)\frac{T_f-t_0}{\gamma}\frac{2}{(2\alpha-1)^2}\right)\right]^{-1}\right]\exp\left(\frac{T_f-t}{\gamma}\right) \\ =& \left[\lim_{\alpha\to1}\frac{V_{\text{thresh}}}{a\gamma}2\left[2\left(1-\exp\left(-\frac{T_f-t_0}{\gamma}\frac{2\alpha-2}{2\alpha-1}\right)\right)\right]^{-1}\right]\exp\left(\frac{T_f-t}{\gamma}\right) \\ =& \frac{V_{\text{thresh}}}{a\gamma}\exp\left(\frac{T_f-t}{\gamma}\right) \end{aligned} \tag{28}$$

In summary, it has been shown that if $\alpha > 0.5$ then λ^* is given by (27). In particular, if $\alpha = 1$ then λ^* is given by the final expression in (28). These results are in complete agreeance with [1].

5 The LTI Extension

We now consider the derivation in [2], for the LTI case. The derivation follows closely but expands on the details significantly as in Section 4.

We consider the class of SDE's shown below.

$$dx(t) = A\,xdt + \Lambda(t)\,dt + \Sigma(t)\,dB(t) \tag{29}$$

Where A is a constant matrix with distinct eigenvalues, $\Lambda(t)$ and $\Sigma(t)$ are control signals similar to $\lambda(t)$ in (6). Without, loss of generality let us assume the system is in controllable canonical form.

Note: The requirement that A have distinct eigenvalues may at first glance seem restrictive. However, for real systems, transfer funcitons are frequently recovered from system identification, in which case it is extremely unlikely for two poles to overlap, thus the restriction is not a problem for many practical problems.

$$dx(t) = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ a_1 & a_2 & a_3 & \dots & a_N \end{bmatrix} x(t)\,dt + \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ \lambda(t) \end{bmatrix} dt + \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ \lambda(t)^\alpha \end{bmatrix} dB(t) \tag{30}$$

A reasonable control problem is to attempt to place the highest order integrator $x_1(t)$ at a particular point at a particular time T. Proceeding as in Section 4, we will solve the following control problem for the case $x(t_0) = 0$.

$$\begin{aligned} &\min_{\lambda} \int_{T_1}^{T_2} \operatorname{var}\left(x_1(t)\right)\,dt \\ &\mathbf{E}\,x_1(t) = x_d \ \forall t \in [T_1, T_2] \end{aligned} \tag{31}$$

Using the standard theory, we see that the solution for $x(t)$ given A's transition matrix $\Phi(t,s)$ is easily expressed.

$$x(t) = \Phi_A(t,t_0)x_0 + \int_{t_0}^{t} \Phi(t,s) \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \lambda(s) \end{bmatrix} ds + \int_{t_0}^{t} \Phi(t,s) \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \lambda(s)^\alpha \end{bmatrix} dB(s) \tag{32}$$

Recall, $x(t_0) = 0$, then due to the decomposition into controllable canonical form, a clean expression for $x_1(t)$ is easily found by inspection.

$$x_1(t) = \int_{t_0}^{t} \Phi_{1,N}(t,s)\lambda(s)ds + \int_{t_0}^{t} \Phi_{1,N}(t,s)\lambda(s)^\alpha dB(s) \tag{33}$$

At this point the expected value of $x_1(t)$ given the transition matrix is clear.

$$\mathbf{E}\,x_1(t) = \mathbf{E}\int_{t_0}^{t} \Phi_{1,N}(t,s)\lambda(s)ds + \int_{t_0}^{t} \Phi_{1,N}(t,s)\lambda(s)^\alpha dB(s) = \int_{t_0}^{t} \Phi_{1,N}(t,s)\lambda(s)ds \tag{34}$$

Further, since the system is in controllable canonical we see that the transition matrix entries are trivially related to each other.

$$\begin{aligned}
&\frac{d}{dt}\Phi(t,s) = A\Phi(t) \\
\iff &\frac{d}{dt}\Phi(t,s) = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ a_1 & a_2 & a_3 & \dots & a_N \end{bmatrix} \Phi(t,s)
\end{aligned} \tag{35}$$

From this representation, it can be recognized by inspection that for $n \in 1, 2, \dots, N-1$ and $m \in 1, 2, \dots N$ we have that $\frac{d}{dt}\phi_{n,m}(t) = \phi_{n+1,m}(t)$.

Further, by basic properties of the transition matrix, it is known that $\Phi(0) = I$. This reveals another important property due to the controllable canonical form.

$$\frac{d}{dt}\Phi(t,t) = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ a_1 & a_2 & a_3 & \dots & a_N \end{bmatrix} \tag{36}$$

Thus, for $n \in 1, 2, \dots, N-2$ we have the $\frac{d}{dt}\Phi_{n,N}(t,t) = \Phi_{n+1,N}(t,t) = 0$

At this point it is possible to derive the optimal control $\lambda^*(t)$ for $t \in [T_1, T_2]$ using the endpoint constraint. Recalling the endpoint constraint and evaluating with our new expression for $x_1(t)$ reveals:

$$\mathbf{E}\, x_1(t) = \int_{t_0}^{t} \Phi_{1,N}(t,s)\lambda(s)ds = x_d \ \forall t \in [T_1, T_2] \tag{37}$$

Taking the derivative with respect to t using the Lebinz rule, reveals a recursive relation that holds for $t \in [T_1, T_2]$.

$$\begin{aligned}
&\frac{d}{dt}\int_{t_0}^{t} \Phi_{1,N}(t,s)\lambda(s)ds = \frac{d}{dt}x_d \\
\iff &\Phi_{1,N}(t,t)\lambda(t) + \int_{t_0}^{t} \Phi_{2,N}(t,s)\lambda(s)ds = 0 \\
\iff &\int_{t_0}^{t} \Phi_{2,N}(t,s)\lambda(s)ds = 0 \\
&\vdots \\
\iff &\int_{t_0}^{t} \Phi_{N,N}(t,s)\lambda(s)ds = 0
\end{aligned} \tag{38}$$

Rewritten in matrix form, the following relation holds for $t \in [T_1, T_2]$.

$$\int_{t_0}^{t} \Phi(t,s) \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \lambda(s) \end{bmatrix} ds = \begin{bmatrix} x_d \\ 0 \\ \vdots \\ 0 \end{bmatrix} \tag{39}$$

Differentiating again with respect to t reveals an expression for λ.

$$\begin{gathered} \frac{d}{dt} \int_{t_0}^{t} \Phi(t,s) \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \lambda(s) \end{bmatrix} ds = \frac{d}{dt} \begin{bmatrix} x_d \\ 0 \\ \vdots \\ 0 \end{bmatrix} \\ \iff \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \lambda(t) \end{bmatrix} + A \int_{t_0}^{t} \Phi(t,s) \begin{bmatrix} 0 \\ \vdots \\ 0 \\ \lambda(s) \end{bmatrix} ds = 0 \end{gathered} \tag{40}$$

Inspecting the last row reveals for $t \in [T_1, T_2]$ we have that $\lambda^*(t) = -a_1 x_d$

As an aside, recall that A is assumed to be diagonalizable. It is a basic result that $\exists K$ invertible such that:

$$\Phi(t) = K \begin{bmatrix} e^{\lambda_1 t} & 0 & \dots & 0 \\ 0 & e^{\lambda_2 t} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & e^{\lambda_N t} \end{bmatrix} K^{-1} \tag{41}$$

To extract $\Phi_{1,N}$ we consider by basic matrix multiplication it can be expressed as a linear combination of $e^{\lambda_n t}$.

$$\begin{aligned} \Phi_{1,N}(t) &= \begin{bmatrix} k_{1,1} & k_{1,2} & \dots & k_{1,N} \end{bmatrix} \begin{bmatrix} e^{\lambda_1 t} & 0 & \dots & 0 \\ 0 & e^{\lambda_2 t} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & e^{\lambda_N t} \end{bmatrix} \begin{bmatrix} k^{-1}_{1,N} \\ k^{-1}_{2,N} \\ \vdots \\ k^{-1}_{N,N} \end{bmatrix} \\ &= \sum_{n=1}^{N} k_{1,n} k^{-1}_{n,N} = \sum_{n=1}^{N} \theta_n e^{\lambda_n t} \end{aligned} \tag{42}$$

Therefore, (33) can be expressed using a linear combination of exponentials.

$$x_1(t) = \sum_{n=1}^{N} \left[\int_{t_0}^{t} \theta_n e^{\lambda_n (t-s)} \lambda(s) ds + \int_{t_0}^{t} \theta_n e^{\lambda_n (t-s)} \lambda(s)^{\alpha} dB(s) \right] \tag{43}$$

We notice this form is surprisingly similar to (10). It is just a linear sum of exponentials now instead of a single exponential. The LTI assumption and controllable canonical form were critical to realizing this understandable form.

Regardless, we now turn to understanding how to control the variance. Consider, the variance can be found in an identical manner to the last section. By basic rearrangement of the integrals we find a form not dependent on the already known value of $\lambda^*(t)$ for $t \in [T_1, T_2]$.

$$\begin{aligned}\int_{T_1}^{T_2} \mathrm{var}(x_1(t))dt &= \int_{T_1}^{T_2}\int_{t_0}^{t} \phi_{1,N}^2(t,s)\lambda(s)^{2\alpha}\, ds\, dt \\ &= \int_{t_0}^{T_1}\int_{T_1}^{T_2} \phi_{1,N}^2(t,s)\lambda(s)^{2\alpha}\, dt\, ds + \int_{T_1}^{T_2}\int_{s}^{T_2} \phi_{1,N}^2(t,s)\lambda(s)^{2\alpha}\, dt\, ds\end{aligned} \tag{44}$$

Clearly, the second term is fixed due to the fixed value of λ in the considered time range. Thus, it suffices to minimize the first term using calculus of variations. Consider adding perurbation function η such that the impact on constraint at the endpoint T is satisfied. E.g. η satisfies:

$$\int_{t_0}^{T} \Phi_{1,N}(T,s)\eta(s)ds = 0 \tag{45}$$

Then, as in the previous section, we require that the derivative of the integrated variance with respect to ϵ, evaluated at zero, be zero, so that the optimal λ can be found.

$$\begin{aligned}&\frac{d}{d\epsilon}\left(\int_{t_0}^{T_1}\int_{T_1}^{T_2} \phi_{1,N}^2(t,s)\left(\lambda^*(s)+\epsilon\eta(s)\right)^{2\alpha}\, dt\, ds\right)\Bigg|_{\epsilon=0} = 0 \\ \iff &\int_{t_0}^{T_1}\int_{T_1}^{T_2} \phi_{1,N}^2(t,s)\lambda^*(s)^{2\alpha-1}\, dt\, ds = 0\end{aligned} \tag{46}$$

Setting (45) and (46) equal reveals a relation between the integrands the results in an explict form for λ^*.

$$\begin{aligned}&\int_{t_0}^{T_1}\int_{T_1}^{T_2} \phi_{1,N}^2(t,s)\lambda^*(s)^{2\alpha-1}\, dt\, ds = \int_{t_0}^{T} \Phi_{1,N}(T,s)\eta(s)ds \\ \iff &\int_{T_1}^{T_2} \phi_{1,N}^2(t,s)\lambda^*(s)^{2\alpha-1}\, dt = \Phi_{1,N}(T,s) \\ \iff &\lambda^*(s) = \left[\frac{\Phi_{1,N}(T,s)}{\int_{T_1}^{T_2}\phi_{1,N}^2(t,s)\, dt}\right]^{1/(2\alpha-1)} \qquad \forall t \in [0,T)\end{aligned} \tag{47}$$

So in summary we have the found the optimal control λ^* is the following piecewise function for the case $\alpha > 0.5$.

$$\lambda^*(t) = \begin{cases} -a_1 x_d & t \in [T_1, T_2] \\ \left[\frac{\Phi_{1,N}(T,s)}{\int_{T_1}^{T_2}\phi_{1,N}^2(t,s)\, dt}\right]^{1/(2\alpha-1)} & t \in [0, T_1) \end{cases} \tag{48}$$

Caution: This result is different from [2]. In particular, that paper finds a substantially different function for the numerator for the case $t \in [0,T)$. More work is needed to get to the bottom of this, to determine where a mistake was made, but there is not time right now.

6 Discussion

6.1 Controlling Membrane Potential

Control of the membrane potential using the diffusion approximation for the neural dynamics was a straightforward exercise in stochastic optimal control. While the algebra involved was somewhat tedious, the results were clear and useful. Really, most of the optimal control has a simple exponential form. It is the calculation of the constant multiplier in front of the decaying exponential that is tedious.

While [1] did extensive simulations using the model, it would be of interest to replicate these experiments with simple Matlab simulations. In, particular, it would be interesting to understand the different neuron parameter's impact on the ability to control spike timing. If there were another week before submission this would be the top priority.

6.2 Controlling an LTI System

Extending to mathematics from control of membrane potential to control of an LTI system was not too difficult in theory. However, in practice, [2] was difficult to understand and in the end a result different than that paper was obtained. More time is needed to untangle the results in [2] further in order to understand their approach and confirm if a mistake was made in this project report's derivation.

Regardless, it seems suspect to approximate the effects of driving an LTI system with a spiking neural output by directly inserting the diffusion approximation of the membrane potential as a control (see (30)). While, it is reasonable to assume that the average spiking rate is in some way determined by the membrane potential, the relation is inexact at best. Further, as discussed earlier, the approximation does not seem to be reasonable when many neurons become involved. It seems likely this approximation would significantly limit the utility of the resulting control system and the optimality of λ^* would be of little practical purpose.

6.3 Future Work

The original intention of this project was to extend the results to LTV systems. However, it is clear from Section 5 that the LTI assumption and the resulting controllable canonical form are essential to realizing a reasonable optimal control problem. Generalizing the results to LTV would require significant work that is not in scope at this time.

There is also the issue of diffusion approximation and it's impact on the performance of the system. Further, the solutions found were open-loop, and it was not clear how to proceed in extending them to the closed-loop case. Finally, the degenerate results in the poisson, and sub-poisson case as indicated by [1, 2] are unappealing.

Going forward, it will be of interest to look into the results in [3]. There, the precise timing of spike trains was controlled in both an open and closed loop fashion. While these approaches required solving the HJB or using dynamic programming, the technical advantages of the method are clear and I suspect the methods are far more extensible than those uncovered in this project.

7 Conclusion

Optimal stochastic control of a neuron's membrane potential was explored by deriving an open-loop control. The approach utilized basic stochastic control concepts to result in an interesting deterministic optimal control problem that could be solved with calculus of variations. An extension was made to the LTI case. While the results were interesting, it is suspected that the diffusion approximation used for the leaky integrate-and-fire neuron model is too relaxed for the goal of precise control of MAVs using neurormophic computation. Future work will perform a practical investigation using simulation. Additionally, it is of interest to do a similarly deep investigation into [3] where spike timing was controlled precisely.

References

[1] J. Feng, X. Chen, H. C. Tuckwell, and E. Vasilaki, "Some optimal stochastic control problems in neuroscience - A review," *Modern Physics Letters B*, vol. 18, no. 21-22, pp. 1067–1085, 2004.

[2] Y. Chen, Y. Deng, S. Yue, and C. Deng, "Optimal Stochastic Control Problem for General Linear Dynamical Systems in Neuroscience," *Advances in Mathematical Physics*, vol. 2017, 2017.

[3] A. Iolov, S. Ditlevsen, and A. Longtin, "Stochastic optimal control of single neuron spike trains," *Journal of Neural Engineering*, vol. 11, no. 4, 2014.

[4] N. G. van Kampen, "Itô versus Stratonovich," *Journal of Statistical Physics*, vol. 24, no. 1, pp. 175–187, 1981.

[5] R. Mannella and P. V. McClintock, "ITÔ versus stratonovich: 30 years later," *Fluctuation and Noise Letters*, vol. 11, no. 1, pp. 1–10, 2012.

[6] H. C. Tuckwell, "A study of some diffusion models of population growth," *Theoretical Population Biology*, vol. 5, no. 3, pp. 345–357, 1974.

[7] R. K. Stagsted, A. Vitale, J. Binz, A. Renner, L. B. Larsen, and Y. Sandamirskaya, "Towards neuromorphic control : A spiking neural network based PID controller for UAV," in *Robotics: Science and Systems 2020*, 2020.

[8] J. J. Hagenaars, F. Paredes-Vallés, S. M. Bohté, and G. C. De Croon, "Evolved Neuromorphic Control for High Speed Divergence-Based Landings of MAVs," *IEEE Robotics and Automation Letters*, vol. 5, no. 4, pp. 6239–6246, 2020.

8 Appendix: Codes

8.1 Numerical simulation of SDE $Y = B\,dB$

Results shown in Figure 1.

```
clear
close all
T = 10;
N = 10000000;
t = 0:T/N:T-T/N;
sigma = sqrt(T / N);

dB = normrnd(0, sigma, 1, N-1);
B = zeros(1, N);
for i = 2:N
    B(i) = B(i-1) + dB(i-1);
end

Yi = zeros(1, N);
for i = 2:N
    Yi(i) = Yi(i-1) + B(i-1)*(B(i) - B(i-1));
end

Yia = zeros(1, N);
for i = 1:N
    Yia(i) = (B(i)^2 - (i - 1) * (T / N))/2;
end

Ysa = zeros(1, N);
for i = 1:N
    Ysa(i) = B(i)^2 / 2;
end

hold on
plot(t, B)
plot(t, Yi)
plot(t, Ysa)

legend('Brownian Motion', 'Ito Numerical', 'Stratanovich Closed Form')
grid on
```

8.2 Numerical simulation of SDE $Y = B\,dt + t\,dB$

Results shown in Figure 2.

```
clear
close all
T = 10;
N = 1000;
t = 0:T/N:T-T/N;
sigma = sqrt(T / N);

dB = normrnd(0, sigma, 1, N-1);
B = zeros(1, N);
for i = 2:N
    B(i) = B(i-1) + dB(i-1);
```

```
end

Yi = zeros(1, N);
for i = 2:N
    Yi(i) = Yi(i-1) + B(i-1)*(T/N) + ((i - 1) * (T / N))*(B(i) - B(i-1));
end

Ysa = zeros(1, N);
for i = 1:N
    Ysa(i) = (i * (T / N)) * B(i);
end

hold on
plot(t, B)
plot(t, Yi)
plot(t, Ysa, 'x')

legend('Brownian Motion', 'Ito Numerical', 'Stratanovich Closed Form')
grid on
```