Washington University in St. Louis
Washington University Open Scholarship

McKelvey School of Engineering Theses & Dissertations | McKelvey School of Engineering

Spring 5-2020

Robust Control of Burst Suppression Amid Physical and Neurological Uncertainty

Stephen Ampleman
Washington University in St. Louis

ShiNung Ching
Washington University in St. Louis

Follow this and additional works at: https://openscholarship.wustl.edu/eng_etds

 Part of the Controls and Control Theory Commons, and the Signal Processing Commons

Recommended Citation

Ampleman, Stephen and Ching, ShiNung, "Robust Control of Burst Suppression Amid Physical and Neurological Uncertainty" (2020). *McKelvey School of Engineering Theses & Dissertations*. 528.
https://openscholarship.wustl.edu/eng_etds/528

This Thesis is brought to you for free and open access by the McKelvey School of Engineering at Washington University Open Scholarship. It has been accepted for inclusion in McKelvey School of Engineering Theses & Dissertations by an authorized administrator of Washington University Open Scholarship. For more information, please contact digital@wumail.wustl.edu.

Washington University in St. Louis

McKelvey School of Engineering

Department of Electrical and Systems Engineering

Thesis Examination Committee:

ShiNung Ching, Chair

James Feher

Sankalp Bhan

Robust Control of Burst Suppression Amid Physical and Neurological Uncertainty

By

Stephen Ampleman

A thesis presented to the McKelvey School of Engineering of Washington University in St. Louis in partial fulfillment of the requirements for the degree of Master of Science

May 2020

St. Louis, Missouri

© 2020 Stephen Ampleman

Acknowledgements

I'd like to thank my advisor, Dr. ShiNung Ching for his patience and willingness to work with me despite unusual constraints due to my part time status. The flexibility he showed to accommodate the large schedule variations was exceptional and I am very grateful for his mentorship throughout this thesis. I'd like to thank my thesis committee members for their time and feedback on this thesis. Dr. Sankalp Bhan has been both a great friend and mentor over the past couple years and I have learned and will continue to learn so much from him. Dr. James Feher has provided excellent advice throughout my graduate experience at Washington University in Saint Louis and appreciate his consideration with all that I wanted to accomplish. I'd also like to thank my parents, James and Laura Ampleman for their high level guidance and support during this process. They raised me to be curious and question everything, as an engineer should, so I owe all of what I am able to accomplish to them. I'd like to thank my siblings, Ryan, Daniel and Rose Ampleman for listening to me talk about this at length at Christmas and Thanksgiving parties for years and not even yawn. Finally, I'd like to thank my girlfriend, Hannah Choi, for her unceasing support during this thesis. As with any large project it is vital to have people you can depend on for help and she was my rock.

Table of Contents

List of Tables

List of Figures

Abstract

Robust Control of Burst Suppression Amid Physical and Neurological Uncertainty

By

Stephen Ampleman

Master of Science in Electrical Engineering

Washington University in St. Louis, 2020

Research Advisor: Professor ShiNung Ching

Burst suppression is a clinical term describing a phenomenon in which the electroencephalogram (EEG) of a sedated patient produces behavior that switches between higher frequency and amplitude bursting to lower frequency and lower amplitude suppression. This phenomenon can be observed during general anesthesia, hypothermia, or in an otherwise induced coma state. In a clinical setting, this phenomenon is typically induced by sedation from a drug such as propofol (2,6-diisopropylphenol). The level of sedation can be quantified by something called the burst suppression ratio (BSR), which is defined as the amount of time that a patient's EEG is in a suppressed state over the amount of time measured. One can vary this ratio by either increasing or decreasing the propofol infusion rate that the patient is given to bring them to a deeper or lighter state of sedation. By measuring the EEG data, one can form a closed loop feedback system where the EEG data is monitored for signs of burst suppression and the propofol is increased or decreased accordingly. Therefore, it becomes desirable to create models of this closed loop system to simulate the kind of behavior that would be expected from a clinical setting such as the one described. Many methods and experimental paradigms have been developed to address this problem including development of pharmacokinetic (PK) models that

describe the dynamics of drug infusion in the body as well as signal processing methods for computing the burst suppression estimation such as the burst suppression probability. Some of these paradigms have been tested in rodent experiments, though human studies remain elusive. In this regard, simulations and detailed physiological modeling and control design can play a key role. This thesis seeks to add on the rich body of work that has been done thus far by incorporating a Schnider PK model with the Wilson-Cowan neural mass model to form a closed loop model which we can use as a basis for more detailed analysis which includes real-time burst suppression estimation as well as uncertainty modeling in both the patient's physical characteristics (such as weight, height, age and gender) in addition to neurological phenomena such as the recovery and consumption rates of neurons during burst suppression behavior. By creating a conversion from the physiological parameters that describe the PK models to the dimensionless and more abstract parameters which guide the Wilson-Cowan equations, and implementing an actuator and burst suppression ratio estimation algorithm, we have effectively modeled the clinical setting with which the BSR is sought to be controlled. Thus, in this study we wished to show that with PID control, one could control this model at a nominal condition (i.e., the patient and neurological parameters which the gains were designed for) as well as at various uncertainty conditions that include both physical and neurological uncertainty, as described above. Using the Zeigler-Nichols tuning method, we were able to design gains to sufficiently control this system at set points of 0.8, 0.5 and 0.2 BSR over a simulation time of roughly 18 hours in both nominal, patient varied with noise added and with reduced performance when including patient variation and noise as well as neurological uncertainty. This time duration was chosen because it was convenient for the model's time constants but also because it is representative of the time a patient may be sedated. The BSR ranges were chosen so as to

show the closed loop system's ability to maintain control at multiple levels of sedation. The reduced performance due to neurological uncertainty was due to the BSR estimation algorithm estimating a lower bound that was too high for the system to be controlled at a BSR of 0.2. The minimum BSR the system with added neurological uncertainty could be controlled to was 0.38, which is where the system held at during the portion of the trajectory that a BSR of 0.2 was commanded. During the achievable parts of the envelope, however, the control scheme worked with similar performance to that of the nominal case. This would suggest that an adaptive estimation algorithm needs to be developed to estimate the neurological deviations from the nominal case. Further, this suggests that if variations in the BSR of a patient due to neurological uncertainty is expected, then accurate estimation of these parameters are vital to reaching a robust solution in a real-time system.

Chapter 1: Introduction

Burst suppression is a phenomenon measured by an electroencephalogram in the neuron dynamics of a patient when the brain is either cooled to a certain temperature (Martin, et al., 1997), is placed under general anesthesia (Brown, Lydic, & Schiff, 2010), or otherwise reaches a coma state (Young, 2000). It is characterized by time intervals containing both high voltage activity (bursts) and low voltage activity (suppression). These alternating states of bursts and suppression can be further characterized by the frequency range in which they occur. When burst suppression occurs, the suppression intervals consist primarily of theta and delta wave signals (<8 Hz) while the burst consist primarily of alpha and beta wave signals (8-30 Hz) (Brown, Lydic, & Schiff, 2010). In addition to frequency content, the amplitudes of each of these signals give evidence for how suppressed the patient really is (Brown, Lydic, & Schiff, 2010). The following figure is Figure 1B and 1C from (Ching, Purdon, Vijayan, Kopell, & Brown, 2012). It shows a comparison of EEG data while the patient is under general anesthesia but has not displayed burst suppression (B) and a patient that has undergone deeper general anesthesia and has indeed achieved a state of burst suppression (C).

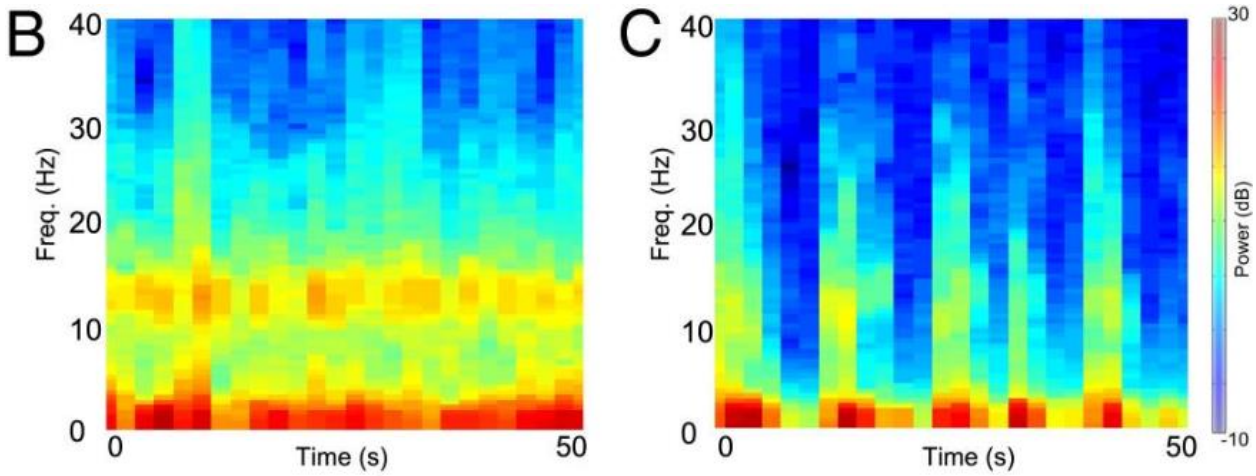

Figure 1: EEG Burst Suppression Absence and Presence

We have chosen the Wilson Cowan equations (Wilson & Cowan 1972, 1973), specifically those augmented to produce a burst suppression phenomenon with a certain parameterization (Liu & Ching, 2017), as a mathematical basis by which we can study how the physiological effects of burst suppression are implemented. Further, an algorithm for calculating the burst suppression given the frequency content and magnitude of the signal is presented and analyzed in comparison to other prevalent burst suppression algorithms.

When pharmacologically inducing a coma state with a drug such as propofol, one can vary the depth of burst-suppression by increasing the dosage applied to the given patient. Therefore, if one desires a certain BSR (measured as the length of time the neurons are suppressed over the total time duration) (Chemali, Wong, Solt, & Brown, 2011) it becomes obvious a feedback control solution is required to target certain ranges of burst suppression. There have been many attempts at formulating a solution to this problem (Ching, et al., 2013), (Chemali, Ching, Purdon, Solt, & Brown, 2013), (Schanechi, Chemali, Liberman, Solt, & Brown, 2013), and (Westover, Kim, Ching, Purdon, & Brown, Robust control of burst suppression for medical coma, 2015). Namely, we strive to further the gains made in the robust

control of burst suppression by introducing actuator dynamics into the closed loop anesthetic delivery (CLAD) system. We further strive to introduce uncertainty not only in the patient's physical characteristics (such as patient height, weight, age and sex), but also in the neuron dynamics of the patient. Once coupled with the presence of an actuator (which would simulate an anesthetic pump), this allows us to study a realistic setting for which CLAD can be applied.

Within the modified Wilson Cowan equations (Liu & Ching, 2017), there are parameters that describe the recovery dynamics of these neurons, which govern how quickly the neurons transfer from a suppressed state back to a bursting state. By varying these parameters, an internal neurological uncertainty can be created and used to increase the robustness of burst suppression feedback control algorithms. Also within these equations, there are various parameters that can be used to quantify the concentration of a pharmacological substance, such as propofol, present in the brain effect site. We can treat the concentration of propofol as the control input to increase the BSR to a desired state. It is crucial to note that clinically, the level of concentration in the effect site cannot be decreased except by a metabolic process. The control input, a pharmacological drug, can only be used to increase the ratio. Therefore, it is vital to both correctly estimate the BSR and use a control method which does not produce high overshoot or steady state error. Therefore, presented in this thesis is a robust control solution which correctly targets certain burst suppression ratios despite uncertainty in the neurological and physical parameters of the patient.

Chapter 2: Wilson Cowan Equations

The Wilson Cowan equations serve to describe neuron dynamics across a large population of neurons. The primary mechanism with which they are modeled is the expression of their dynamics in both an excitatory and inhibitory state. Work has been done by Liu & Ching (2017) to modify these equations to give them a modulating input, $\phi_j(t)$, in addition to original fast dynamics described initially by Wilson and Cowan (1973, 1972). The equations for the excitatory and inhibitory states are below.

$$\dot{e_j} = \omega_e\left(-e_j + \left(k_e - r_e e_j\right)\mathcal{F}\left[c_1 e_j - c_2 i_j + \sum_{k\in\mathbb{N}_j} k_j^{f_e} e_k + P + \phi_j(t)\right]\right) + W_j^e(t) \qquad \textit{Eq. 1}$$

$$\dot{i_j} = \omega_i\left(-i_j + \left(k_i - r_i i_j\right)\mathcal{F}\left[c_3 e_j - c_4 i_j + \sum_{k\in\mathbb{N}_j} k_j^{f_e} e_k + Q + \phi_j(t)\right]\right) + W_j^i(t) \qquad \textit{Eq. 2}$$

The subscript j represents the column of neuron tissue that is being described. Additionally, the subscript k describes a column that would be coupled with column j.This specific study seeks only to analyze a single column of tissue and thus the equations will be hereafter described without the additional subscript or summation for the interconnectivity between columns. Thus, the de-coupled equations for describing neuron activity are below.

$$\dot{e} = \omega_e\left(-e + (k_e - r_e e)\mathcal{F}\left[c_1 e_j - c_2 i + P + \phi(t)\right]\right) + W^e(t) \qquad \textit{Eq. 3}$$

$$i = \omega_i(-i + (k_i - r_i i)\mathcal{F}[c_3 e - c_4 i + Q + \phi(t)]) + W^i(t) \qquad \textit{Eq. 4}$$

The function $\mathcal{F}$ stands for a logistic sigmoid such that is used in the original Wilson-Cowan equations.

$$\mathcal{F}(x) = \frac{1}{1 + \exp[-a(x-\theta)]} - \frac{1}{1 + \exp(a\theta)} \qquad \textit{Eq. 5}$$

The parameters a and θ represent tuning variables that are typically used in such functions to alter the slope and midpoint of the sigmoid. The modulating input mentioned above is described by the following equation.

$$\dot{\phi} = -\mu_1\phi + \left(\frac{\mu_2}{1 + \exp\left(\left[-k_\phi(M-\eta)\right]\right)}\right) \qquad \textit{Eq. 6}$$

Where μ_1 is the time constant for the autonomous part of the dynamics and μ_2 is the time constants of the sigmoidal part of the dynamics. k_ϕ is a tuning variable for the sigmoid slope shape, η represents the midpoint of the sigmoid when the input M is zero. The input M acts as a gating variable for ϕ. The input M is essential in describing the key features of this model as it attempts to accurately model the underlying physiological events occurring. It is a function of two variables, a consumption variable, g_c, and a recovery variable, g_r.

$$\dot{M} = g_r(e) - g_c(e) \quad \textit{Eq. 7}$$

$$g_c = k_c\left(\frac{e^2}{0.01 + e^2}\right) \quad \textit{Eq. 8}$$

$$g_r = k_r\beta \quad \textit{Eq. 9}$$

$$\dot{\beta} = -\nu_1\beta + \left(\frac{\nu_2}{1 + \exp\left(\left[-k_\beta(e - \zeta)\right]\right)}\right) \quad \textit{Eq. 10}$$

The recovery and consumption processes are primarily a function of the excitatory dynamics. Our main focus will be on the evolution of the variable β, and, more specifically, how small changes in the constants that describe its behavior are crucial in how the system as a whole behaves.

Finally, it should be noted that the main variables e and i should not be misunderstood as EEG activity, although there are fundamental similarities between the two. Liu & Ching (2017) made this distinction while also stating that it is a suitable variable for burst-suppression type studies. Therefore, in this paper, we will be analyzing the excitatory firing rate as a surrogate for EEG data.

In an attempt to succinctly summarize the equations, variables, and parameters above, the following tables are given below which names the variable and provides a short description. Units, nominal values, and initial conditions are also given.

Table 1: Variable Definition

Variable	Description	Units	Initial Condition
e	Excitatory firing rate	$\frac{\text{substrate}}{\text{sec}}$	0.3

i	Inhibitory firing rate	$\frac{\text{substrate}}{\text{sec}}$	0.1
ϕ	Slow Process firing rate	$\frac{\text{substrate}}{\text{sec}}$	0.25
M	Consumption/Recovery Gating Variable	$\frac{\text{substrate}}{\text{sec}}$	10
β	Recovery Evolution	$\frac{\text{substrate}}{\text{sec}}$	0.3

Table 2: Parameter Definition

Parameter	Description	Units	Value
k_e, k_i	Maximal value of the excitatory and inhibitory response functions	$\frac{\text{substrate}}{\text{sec}}$	1, 1
r_e, r_i	The absolute refractory period of the excitatory, inhibitory subpopulation	Non-dimensional	1, 1
ω_e, ω_i	Wilson-Cowan Time Constants	$\frac{1}{\text{sec}}$	500, 500
P, Q	Level of background excitation in the excitatory, inhibitory subpopulation	$\frac{\text{substrate}}{\text{sec}}$	900, 0
c_1, c_3	Average number of excitatory synapses per cell	Non-dimensional	16, 15
c_2, c_4	Average number of inhibitory synapses per cell	Non-dimensional	12, 3
$\theta_e, \theta_i, a_e, a_i$	Maximal slope parameters of the logistic curve for the excitatory, inhibitory subpopulation	Non-dimensional	4, 3.7, 1.3, 2
μ_1, μ_2	Modulation time constants	$\frac{1}{\text{sec}}$	2, 2

k_ϕ	Sensitivity to the variations of the metabolic substrate	$\frac{1}{\text{substrate}}$	75
k_β	Sensitivity to the variations of the neuronal activity	$\frac{1}{\text{substrate}}$	4.5
η, ζ	Average number of inhibitory synapses per cell	substrate	0.25, 0.27
k_r, k_c	Metabolic recovery and consumption rates	$\frac{\text{substrate}}{\text{sec}}$	900, 700
ν_1, ν_2	Homeostatic autoregulation time constants	$\frac{1}{\text{sec}}$	.08, .08

In Liu & Ching (2017), a time scale of milliseconds was used. Since we will be studying the model in the span of hours, we have proportionally increased the time constant and metabolic recovery and consumption rates in order to allow the dynamics to scale properly to a timescale of seconds. This time scale allows us to still view the frequency content and the underlying signal mechanisms while also maintaining reasonable simulation times.

The purpose of using this model is to take advantage of its inherent burst suppression qualities as the parameter c_2 is varied. We can use this parameter as a way to model various propofol infusion rates. In a clinical setting, a patient will become more sedated as the propofol infusion rate is increased. Therefore, it is vital to describe bounds with which the model will behave properly as c_2 is varied. In Liu & Ching (2017), the bounds were set by the system's stability in regards to the c_2 parameter. The system has a bifurcation point around the value of 7.7 after which the system develops a stable limit cycle envelope where burst suppression behavior occurs until a c_2 value of 69. At this value a stable steady state solution can be found

and the model becomes completely suppressed. Therefore, we wish to analyze how burst suppression emerges and changes as the value of c_2 increases from 7.7 to 69. The figures below give an example for how the model behaves with multiple c_2 values.

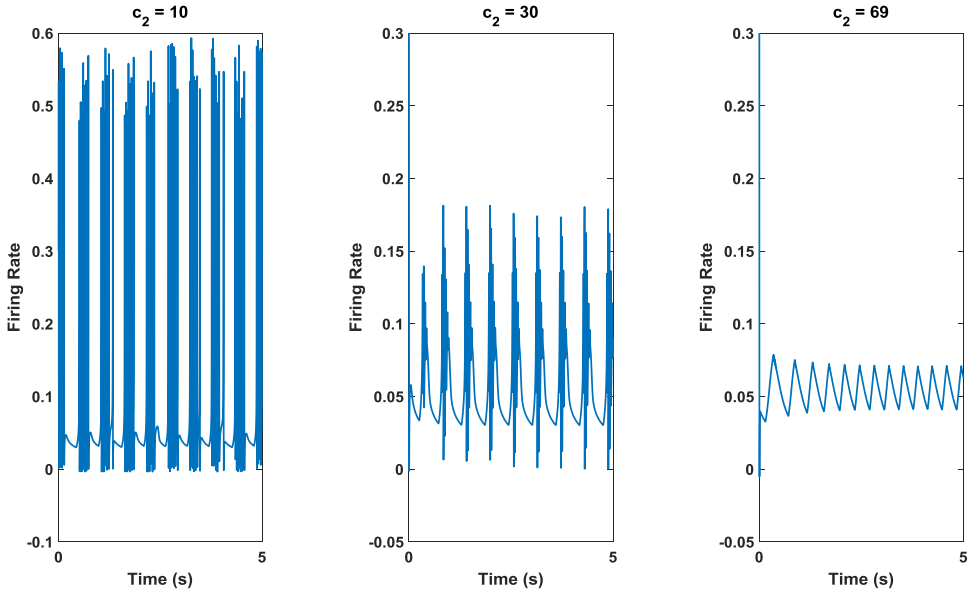

Figure 2: Burst Amplitude Decrease with Increasing c2

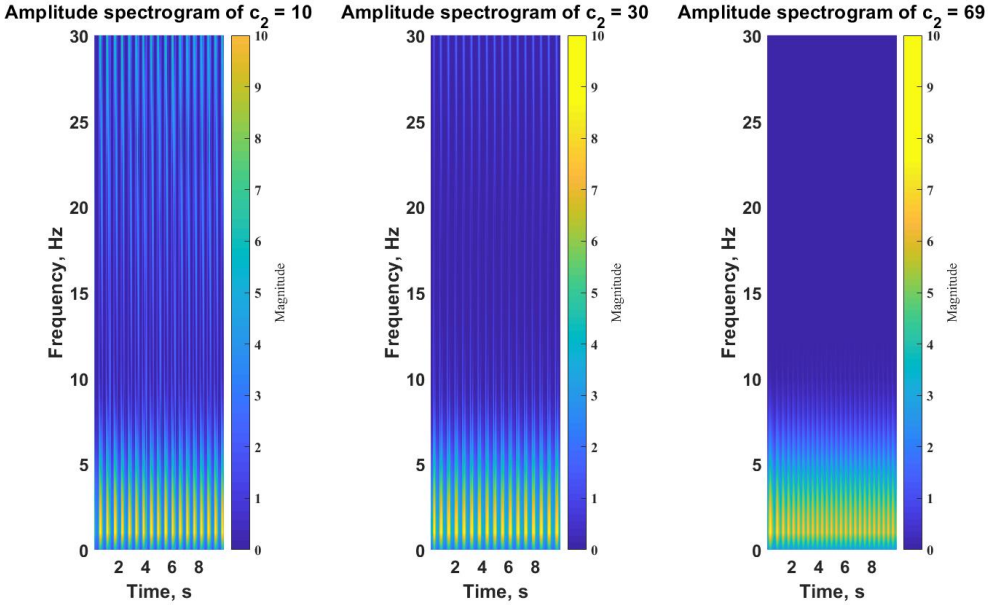

Figure 3: Loss of High Frequency Dynamics as c2 Increases

The values of c_2 were chosen specifically to show how the bursting amplitude changes with time. It is also notable within the spectrogram how when the model reaches near complete suppression, frequencies in the range above 8 Hz are no longer present. These differences in the model expression will be taken into account when the burst suppression algorithms are presented in Chapter 3.

In addition to simply running a propofol (c_2) increase on this model, we wish to generate some uncertainty for different patient types. We focused primarily on the variation of the recovery parameters, k_β, ζ, and ν to simulate how different patients may respond to propofol as it enters their system. The figures below show the same c_2 values but with a change in these recovery parameters. The notable changes are the suppression length and burst amplitude. As these are both taken into account for burst suppression calculation, this will produce a different burst suppression ratio result throughout a range of c_2.

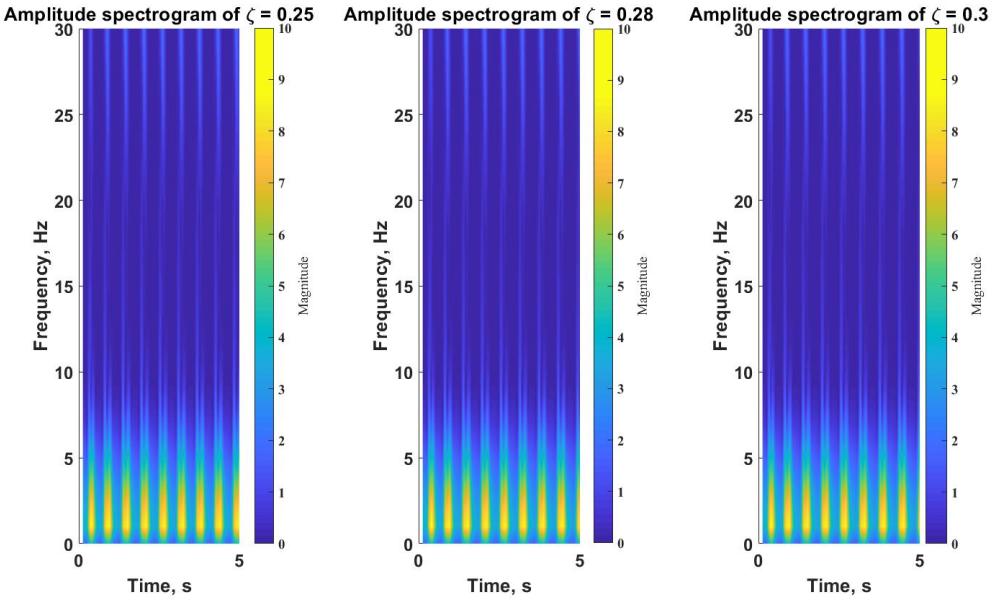

Figure 4: Suppression Length Increase with Increasing zeta

The first set of figures is with a static c_2 value of 30, but a varied ζ. The other recovery parameters were kept the same. There is a subtle difference between the three spectrograms shown above. In the first there are 8 clear bursting events with an additional one trailing off as

the simulation time expired. As ζ increases, the trailing bursting event becomes less visible. This is due to the suppression length increasing as zeta increases.

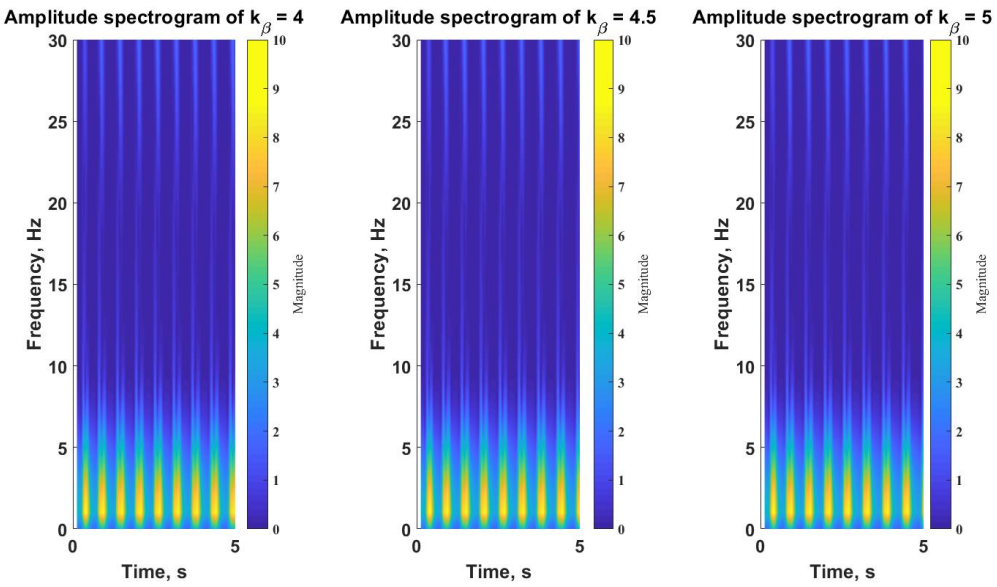

Figure 5: Spectrogram with kB Variation

The second set of figures is with a static c_2 but a varied k_β. The other recovery parameters were kept the same. A similar trend to that seen in the result with a varied ζ is manifested when the parameter k_β is varied. After exploring the effects of varying ν with the remaining parameters remaining constant, we will analyze the effects of coupling these variations.

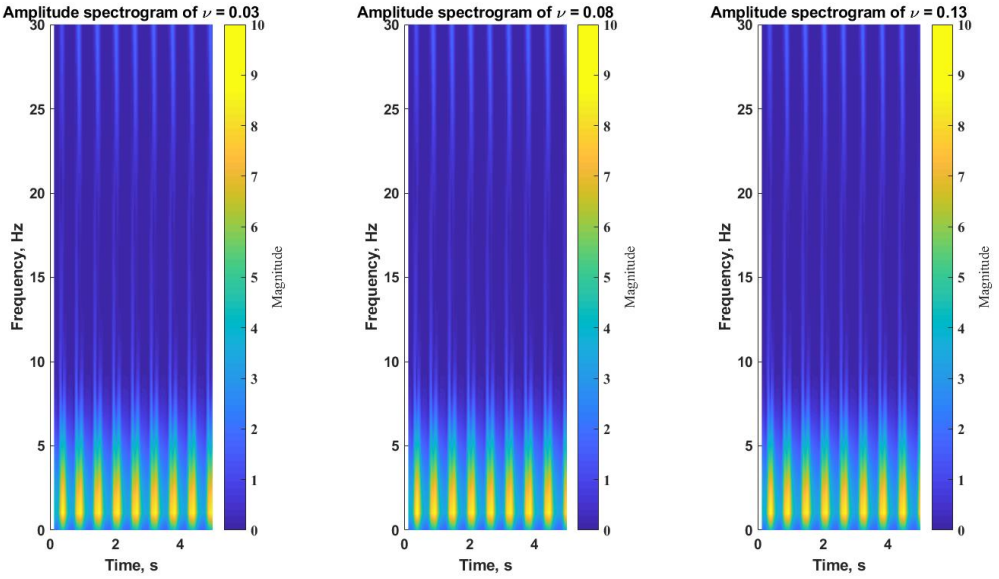

Figure 6: Spectrogram with nu Variation

The third set of figures is with a static c_2 but a varied ν. The other recovery parameters were kept the same. Perhaps unsurprisingly, the same trend that was found with varying the other two recovery parameters showed itself in this variation. To visualize the impact on how changing all three simultaneously impacts the result, the below figure analyzed a configuration with all parameters decreased, a configuration with one parameter increased (this was chosen to be ζ, but as the results above show us, we could have chosen any parameter to use as a point of comparison), and a configuration with all parameters increased.

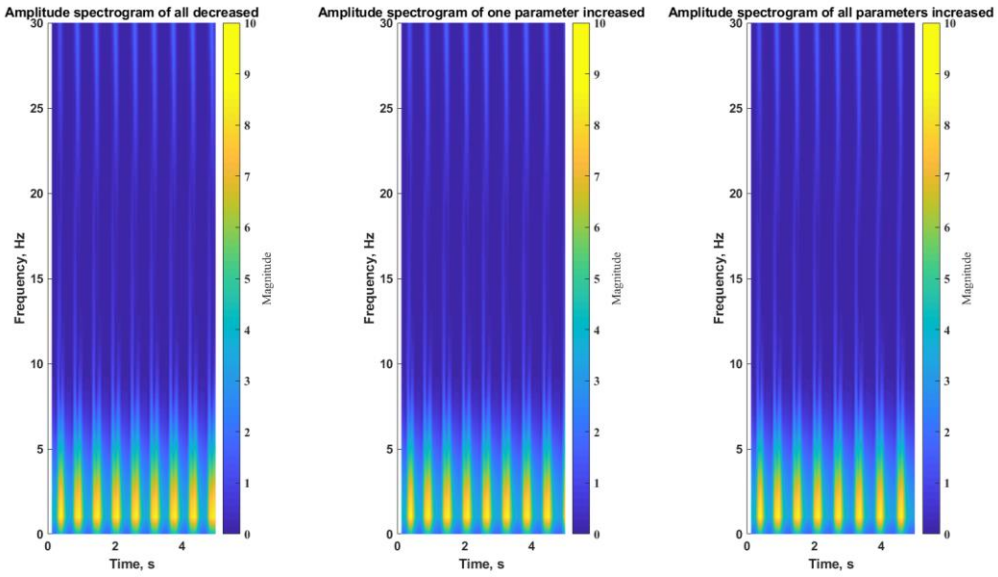

Figure 7: Simultaneous Recovery Parameter Variation

This final analysis confirms our suspicion that simultaneous changes of the recovery parameters results in a more drastic increase of the suppression lengths while a simultaneous decrease in the recovery parameters results in a decrease of the suppression lengths. We see this again by looking at the number of bursting events in each configuration. The first one nearly has a full ninth bursting event take place while the middle one has most of the ninth bursting event cut off due to simulation time. Finally, the case where all configuration parameters are increased shows a complete lack of a ninth bursting event. These recovery parameter variations will again be analyzed in Chapter 3 where we will quantitatively see the result of changing these parameters to confirm the qualitative analysis done here.

Chapter 3: Burst Suppression Ratio Algorithm

As described in the introduction, the burst suppression ratio was used as a method to determine how suppressed the patient currently is. It is measured as the ratio between the time the model is suppressed over the time the data is being measured. The base set of parameters were taken from Liu & Ching (2017) and the algorithm was then tuned to approximate the results shown therein. The algorithm presented attempts to utilize the defining frequency characteristics of burst suppression as its method of determining the correct ratio. The time scale set for this test was in seconds in order to use a sample rate (512Hz) that would emulate a sampling commonly used for EEG processing.

The algorithm utilized a discrete time Fourier transform (DTFT) as its method of determining the frequency content as the signal from the Wilson Cowan equations was sent through. It was desirable to get the frequency bins in integers ranging from 1Hz to 30Hz in order to accurately analyze the expected spectrum. Since our sampling time 512Hz, this meant the number of Fourier points taken needed to be equal to the sampling time. A signal overlay value of 1 was chosen with a window length of 32 samples. After the frequency bins had been created using the DTFT parameters, indices were used to get the magnitudes of the frequencies where suppression takes place (<8Hz) and where bursting takes place (8Hz < w < 30Hz) (Brown, Lydic, & Schiff, 2010), effectively filtering out unwanted frequency ranges. The mean magnitudes for each were taken and a ratio between these values, and pre-chosen magnitudes that would indicate bursting/suppression was found. These values were then put into a sigmoid function to model a percent suppressed based on these ratio values for the particular time step.

Parameter sweeps were done for this sigmoid function and the magnitudes for which the ratios were taken to get a burst-suppression ratio relationship with respect to c_2 that mirrored expectations for how it would evolve as the magnitude of c_2 increased. The equations below provide mathematical insight for the prior description.

$$e(n,k) = \sum_{m=-\infty}^{\infty} e[m] * w[n-m] \exp\left(-\frac{j2\pi k}{N}\right) \quad \textit{Eq. 11}$$

Where n is the current time step that the discrete time Fourier transform (DTFT) is being applied to and k is the frequency "bin" out of N frequencies being analyzed. j represents the imaginary number and w is a Blackman window which is described by the following equation

$$w[k+1] = 0.42 - 0.5\cos\left(\frac{2\pi k}{n-1}\right) + 0.08\cos\left(\frac{4\pi k}{n-1}\right) \quad \textit{Eq. 12}$$

After the transform is taken, the vector is split up into two separate vectors described by the frequencies at which they are calculated. It is desirable to look at the frequency where suppression takes place separate from higher frequency content so they can be analyzed individually. Thus, since the frequency "buckets" were taken in steps of 1 Hz, the first 8 indices of the vector $e(n,k)$ were taken for a certain n to be the suppressed transform while the indices from 9 to 30 were taken to be the bursting transform. Since this operation effectively acts as a filter, choosing only those frequency ranges allows us to filter out unwanted high frequency data that could corrupt the burst suppression calculation. The next operation is a mean calculation of the magnitudes of the frequencies in the suppressed and bursting bins.

$$m_s(n) = \frac{1}{8}\sum_{k=1}^{8} |e(n,k)| \qquad \textit{Eq. 13}$$

$$m_b(n) = \frac{k_b}{21}\sum_{k=9}^{30} |e(n,k)| \qquad \textit{Eq. 14}$$

Where m_s and m_b are the mean magnitudes over that frequency spectrum. It is typical for the bursting magnitudes to be significantly less than the suppressed magnitudes. However, there are large fluctuations between small and large c_2 values in the bursting magnitudes and thus vital to determining the burst suppression ratio. The gain k_b seeks to place the bursting and suppressed magnitudes in a similar value range so as to make both effective inputs. These values are then fed into a sigmoid function and multiplied by the sample time to represent how suppressed the time step is from 0 to 1, where 0 indicates full bursting and 1 indicates full suppression.

$$S(n) = \frac{dt}{1 + \exp\left(-k_{sb}\left(-m_s(n) - m_b(n) - x_{0_{sb}}\right)\right)} \qquad \textit{Eq. 15}$$

The negatives of the magnitudes are taken so as to have a large negative number in the exponential when the bursting and suppressed magnitudes are high, thus giving a value of $S(n)$ closer to zero which indicates a burst event. The following table gave a parameterization that was well suited for a large range of c_2.

Table 3: BSR Algorithm Tuning Parameters

Parameter	Description	Value
k_{sb}	Sigmoid shaping function	1
$x_{0_{sb}}$	Midpoint value for the sigmoid	-30
k_b	Gain on the bursting magnitudes	500

Using this method, the relationship found with burst suppression and c_2 could not be made completely sigmoidal. The lower values of c_2 would express large intervals of suppression before coming into a bursting regime which lasted a similar duration. Using the parameters described in the table, the figure below shows the BSR of each c_2 value run individually. This algorithm calculated a higher value for the lower ranges of c_2, specifically between 8 and 10. After a c_2 value of 10, the BSR increases in a near monotonic fashion until it reaches a near 1 BSR at a c_2 value of 69.

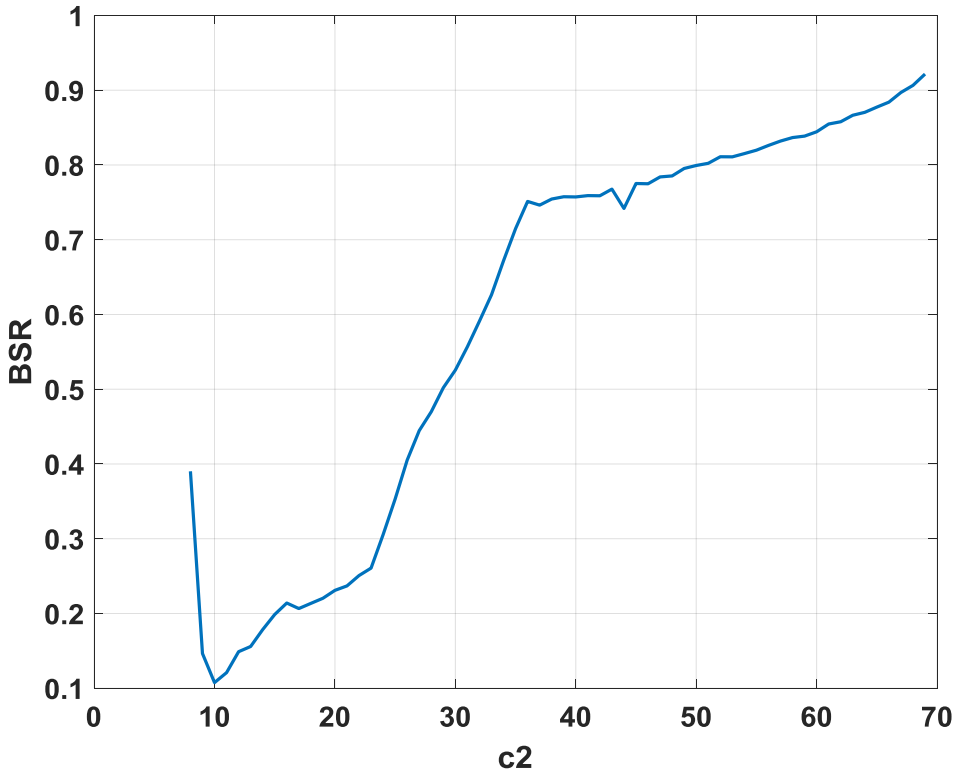

Figure 8: BSR vs c2 Relationship

Despite these limitations on the extreme lower ranges of c_2, the remainder of the values with this parameterization produced results consistent with what would be expected for burst suppression as the propofol concentration is increased. Since this is the type of physiological behavior the model is attempting to reproduce, we can lower bound the effective value of c_2 where the BSR stops decreasing and begins its monotonic increase to 1.

The mean of the suppression is used as a way to calculate the actual suppression over time. This results in a delay before a certain confidence level in the value can be achieved. This is useful information from a control design perspective because these algorithms are very nonlinear in nature and we desire to implement a linear control method. Thus, if we can find the expected delay values for this algorithm to converge, our control design can account for it by

using a linear approximation for delay associated with time to converge to a confidence level. Since we desire tight control on our BSR so as to not fluctuate the patient between multiple BSR levels, a convergence value for the mean signal variance was chosen to be 0.0001.

In order to design the gains for the system, a time for this algorithm to converge had to be selected. Thus, the time for this algorithm to converge with nominal and varied recovery parameters of a limited subset to that of the range run in Chapter 2 were calculated. The results for these convergence times are shown below:

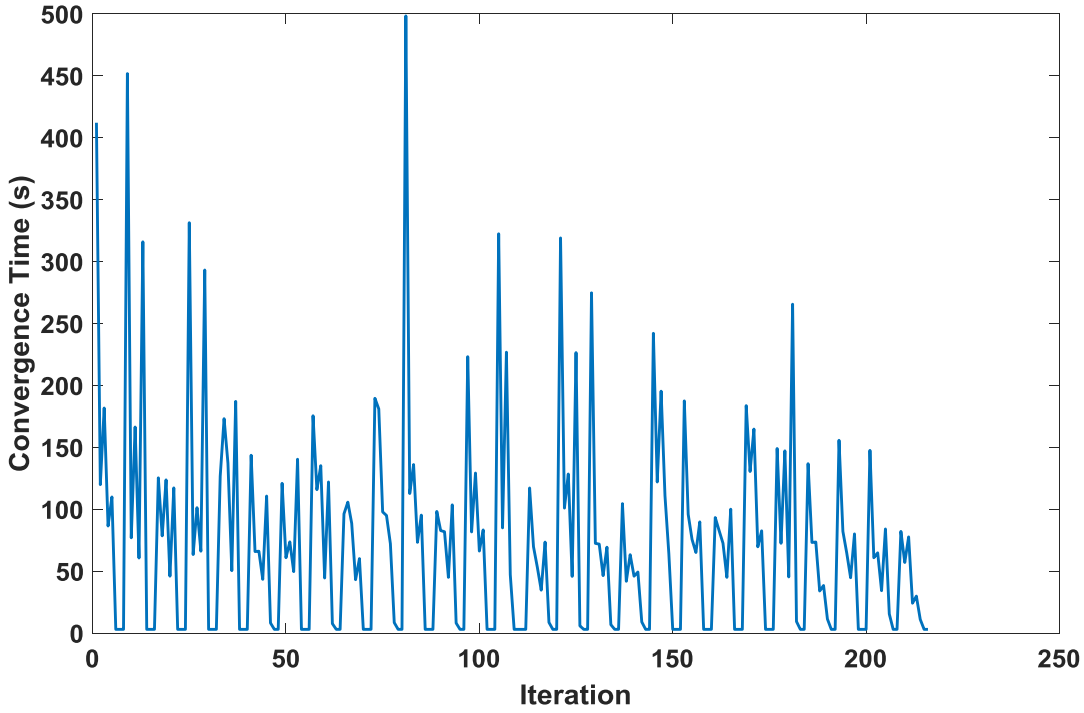

Figure 9: BSR Convergence Times

This variation on convergence times shows some very sharp increases in the lower ranges of c_2 but very quick convergences for large ranges of c_2. This is likely due to the complete absence of higher frequency bursting data at high c_2 values which make it much easier for the

algorithm to converge given the lower frequency data is more consistently defined and is more present at these points. Nonetheless, a delay for this algorithm to be implemented in linear gain design was chosen to be the average of the values above which came out to be approximately 72 seconds. This is consistent for what is expected of convergent BSR values according to (Chemali, Wong, Solt, & Brown, 2011) which gives us further reasoning to select a value of similar magnitude.

Chapter 4: System Architecture

Before looking at the gain design method, we will introduce the nonlinear system architecture which we wish to control. A system architecture was designed around the Wilson Cowan dynamics in order to more accurately simulate a CLAD system. Below in Figure 9 is the Simulink model used to represent the system as a whole. This model was used in Rapid Accelerator mode to be able to run 18 hours of simulation in realistic time.

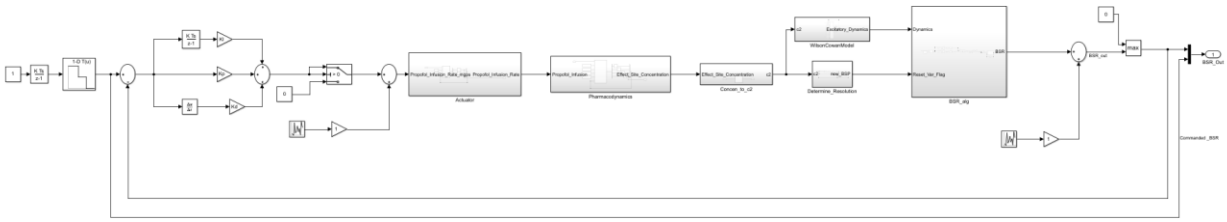

Figure 10: Closed Loop Simulink Model

The plant of this model is comprised of three main pieces. The first is called the pharmacokinetic model (PK model). Pharmacokinetics describes how an infused drug is distributed and discarded from multiple compartments in the body. The model we have selected for this study is the Schnider model for humans (Schnider, et al., 1998), (Absalom, Mani, DeSmet, & Struys, 2009). It is a four state model, each of which describes a different compartment within the model. The main compartment, x_1, describes the concentration of the drug in the central compartment, or the blood, of the patient. Compartments x_2 and x_3 describe the fast and slow compartments, respectively. This can be easily seen by the magnitude of the coefficients that describe the second order relationship between these compartments and the central compartment. Finally, there is x_4, or the effect site compartment concentration. This

represents the concentration of propofol in the brain, which we can then use as an estimate for the variable, c_2, which is an input to the Wilson Cowan equations.

$$\dot{x} = \begin{bmatrix} -(k_{10} + k_{12} + k_{13}) & \frac{V_2 k_{21}}{V_1} & \frac{V_3 k_{31}}{V_1} & 0 \\ \frac{V_1 k_{12}}{V_2} & -k_{21} & 0 & 0 \\ \frac{V_1 k_{13}}{V_3} & 0 & -k_{31} & 0 \\ k_{eo} & 0 & 0 & -k_{eo} \end{bmatrix} x + \begin{bmatrix} \frac{1}{60V_1} \\ 0 \\ 0 \\ 0 \end{bmatrix} y_{act} \qquad \textit{Eq. 16}$$

$$x_{concen} = [0 \quad 0 \quad 0 \quad 1]x \qquad \textit{Eq. 17}$$

The parameters shown in equations 19 and 20 above are functions of the patient's height, weight, age and sex. Typically, these parameters are set in the time scale of $1/min$, however, since we are modeling delay and performing our control design in the timescale of seconds, the parameters were scaled to be in the correct time scale. The equations for non-static parameters are shown below and the static parameters are additionally shown in the following table.

$$LBM_{male} = 1.1m - \frac{128m^2}{h^2} \quad \textit{Eq. 18}$$

$$LBM_{female} = 1.07m - \frac{148m^2}{h^2} \quad \textit{Eq. 19}$$

$$V_2 = 18.9 - 0.391(a - 53) \quad \textit{Eq. 20}$$

$$k_{10} = \left(\frac{1}{60}\right)\left(.443 + 0.0107(m - 77) - 0.0159\,(LBM - 59) + 0.0062(h - 177)\right) \quad \textit{Eq. 21}$$

$$k_{12} = \left(\frac{1}{60}\right)\left(0.302 - 0.0056\,(a - 53)\right) \quad \textit{Eq. 22}$$

$$k_{21} = \left(\frac{1}{60}\right)\left(\frac{1.29 - 0.024(a - 53)}{V_2}\right) \quad \textit{Eq. 23}$$

Table 4: Static PK Coefficients

Parameter	Value
k_{13}	$\left(\frac{1}{60}\right)0.196$
k_{31}	$\left(\frac{1}{60}\right)0.0035$
k_{eo}	$\left(\frac{1}{60}\right)0.456$

Before we send the effect site concentration straight to the Wilson Cowan equations, we must develop a relationship to convert this parameter to the non-dimensional representation that the Wilson Cowan equations use. Based on the relationship established in the BSR algorithm in Figure 7, we can use a similar relationship described in Westover et al (2015) between effect site

concentration and BSP to give an approximation for this transformation. A sigmoid model was made to approximate the BSP vs effect site concentration shown in Westover et al (2015) using equation 26 below.

$$BSP_{approx} = \frac{x_4^{6.6}}{x_4^{6.6} + 5.5^{6.6}} \qquad \textit{Eq. 24}$$

Westover et al (2015) shows that this relationship holds for a certain patient parameterization and is also used for all later experiments run with various patients. The concentration to c_2 values derived from this method are shown in the figure below.

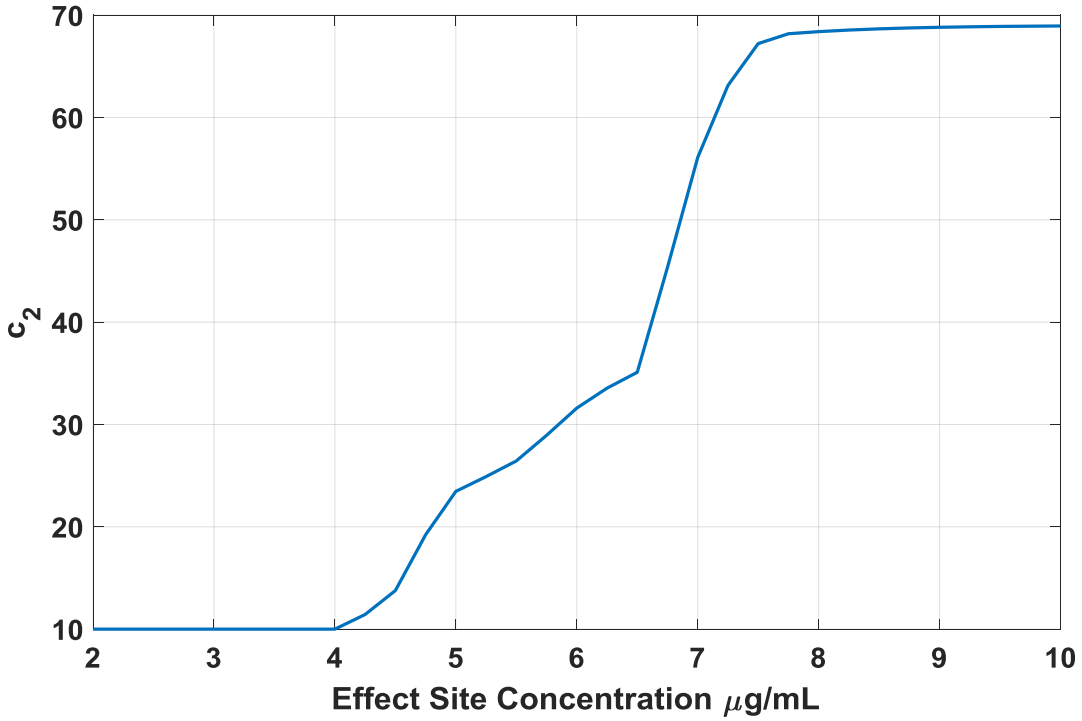

Figure 11: Effect Site Concentration to c2 Parameter

These c_2 values were then sent to the Wilson Cowan equations for the dynamics to be calculated and fed through to the BSR estimation algorithm. The BSR estimation algorithm here

is implemented in the same way that is shown in equations 11-15, with the exception that a reset line has been incorporated for when a new c_2 value is being fed through to the Wilson-Cowan equations. This is necessary since we need to achieve a certain confidence level for each measured BSR. Hence, when a new value of c_2 is fed through the system, the mean and signal variance blocks must be reset. In order to determine the c_2 increments at which the delay is reset, we must look at the performance of the control system. If the system naturally increases c_2 so quickly that the delay caused by the algorithm cannot complete before a new c_2 value is fed through, the system will be caught in a state of constant delay. In this state of constant delay, the BSR will never claim convergence at the current propofol concentration before the concentration changes so drastically that the estimated BSR no longer accurately reflects the state of the system. Hence, a control resolution can be found for BSR as a function of the c_2 delay reset value. To coincide with the confidence level we set of .0001 BSR variance, we set a criteria for each new c_2 value that is sent into the system as described in equation 26 below.

$$\left|c_2 - c_{2_{prev}}\right|^2 \leq c_{2_{thresh}} \qquad \textit{Eq. 25}$$

To choose the threshold value, we can take a linear fit of the sigmoid shown in Chapter 3 where we computed the nominal relationship between c_2 and BSR. Figure (blank) below shows the original BSR curve with the linear approximated plotted over it.

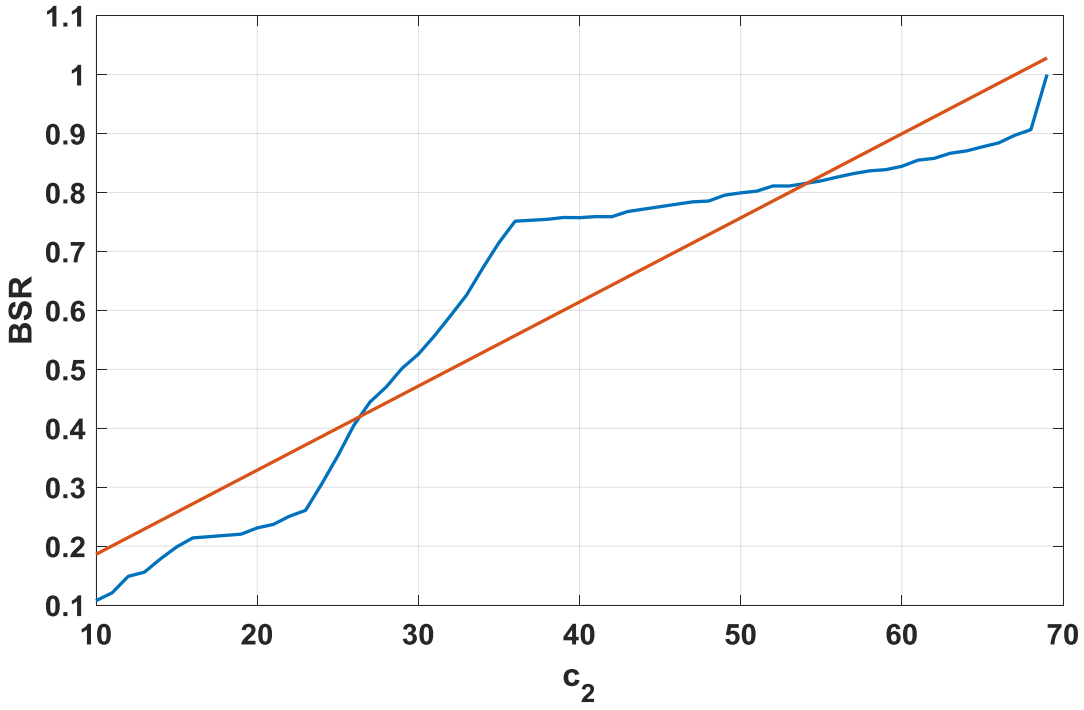

Figure 12: BSR vs c2 with Linear Approximation for BSR growth per c2

This linear fit shows that for every increase of 1 c_2, there is an approximate increase of 0.0143 BSR. Thus we can use a similar relationship that the BSR uses for convergence to calculate what the 2-norm of the difference between the next c_2 value and the previous c_2 should be to reset the delay.

$$|BSR - BSR_{mean}|^2 \leq 0.0001 \qquad \textit{Eq. 26}$$

$$\left|.0143(c_2 - c_{2_{prev}})\right|^2 \leq 0.0001, \qquad \left|c_2 - c_{2_{prev}}\right|^2 \leq 0.489$$

Prior to the plant model we have a second order actuator model and our controller. For this problem, a PID controller architecture was chosen. A specific goal of this study was to

minimize overshoot. Implementing derivative control in addition to a classic PI controller is an excellent way to minimize overshoot in a system as well as minimize steady state error. The equations which describe the controller and actuator are below.

Controller:

$$\begin{bmatrix} \dot{\int} e \\ e \end{bmatrix} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \int e \\ e \end{bmatrix} + \begin{bmatrix} -1 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix} \begin{bmatrix} BSR \\ B\dot{S}R \\ BSR_{ref} \end{bmatrix} \quad \textit{Eq. 27}$$

$$u = \begin{bmatrix} K_I & K_p \end{bmatrix} \begin{bmatrix} \int e \\ e \end{bmatrix} + \begin{bmatrix} 0 & K_d & 0 \end{bmatrix} \begin{bmatrix} BSR \\ B\dot{S}R \\ BSR_{ref} \end{bmatrix} \quad \textit{Eq. 28}$$

Actuator:

$$\dot{x}_{act} = \begin{bmatrix} 0 & 1 \\ -\omega_n^2 & -2\zeta\omega_n \end{bmatrix} x + \begin{bmatrix} 0 \\ \omega_n^2 \end{bmatrix} u \quad \textit{Eq. 29}$$

$$y = \begin{bmatrix} 1 & 0 \end{bmatrix} x \quad \textit{Eq. 30}$$

The method which was used to determine actuator and controller parameters will be shown in further detail in Chapter 5: Gain Selection.

Finally, two switches were added to disallow unrealistic results to come from the controller or the BSR output. First, a switch was put before the PK model that limited the input to the system to be positive (i.e. always providing a propofol infusion rate to the system). This was done because the only inputs we can put into the system are a positive one or zero. However, a useful quality of the PK model is that it is asymptotically stable about a zero concentration

equilibrium point. In this way, it models the system's ability to effectively decay its concentration if no additional propofol is fed into the system. Second, a max block was put at after the noise is added to the output of the BSR estimation algorithm. This was done so that if the BSR estimation algorithm is estimating a near-zero BSR, the noise will not allow a negative BSR to be estimated as that is not physiologically possible.

Chapter 5: Gain Selection

For our PID controller and actuator, there were several goals that needed to be met in order to claim the gains and actuator parameters had been sufficiently designed to meet the system requirements. From a time domain perspective, we desired minimal overshoot, a modest rise time, and sufficient robustness in terms of the noise added at the input and output of the plant, as well as parameter uncertainty. In terms of the time domain requirements, the rise time was a much lower priority for these systems than the steady state error or the overshoot. This is due largely in part to the long periods of time that the system will need to be held at a certain level. Of course, it will have limits on how slow of a system response it can be, but it will not be a driving requirement. Overshoot and steady state error on the other hand are much more important requirements for our application. Overshoot is important because of our control system's lack of ability to remove any propofol from the system. It cannot command a negative infusion rate. Therefore if a target BSR is overshot, our closed loop system must rely on the natural release of the propofol from the effect site compartment to get down to the desired level. A large steady state error is undesirable for obvious reasons that apply to any control system, but for our application a large steady state error could command a much higher propofol dose than the patient requires and thus result in patient overdose.

We relied on frequency domain methods to determine the robustness of our system. A sure-fire way of determining robustness of a system is to look at the sensitivity and co-sensitivity equations. The co-sensitivity and sensitivity functions are uniquely defined by something known as the "waterbed effect". This is due to their algebraic relationship with each other, described in the equation below:

$$S + T = 1 \qquad \textit{Eq. 31}$$

This equation tells us that the closed loop response, T, must be balanced with the sensitivity of the system, S. It is desirable for our system to respond to low frequency inputs, (i.e. set point values or step responses in BSR) and be robust to any high frequency oscillations that could occur due to noise. In this way, the magnitude of the closed loop system, T, should be 1 for low frequency inputs and 0 at high frequency inputs. The opposite should be true for S in order to be robust to noise. An easy tuning method to ensure this will be the case is to look at the loop gain throughout the frequency range. The loop gain should "roll off" at high frequencies which ensures our closed loop system is not responding to high frequency input. In order to maximize gain and phase margin, it is desirable to look at how quickly the loop gain of the system "rolls off" at the loop gain crossover frequency. Making this value near to -20dB/decade ensures that we are getting sufficient gain and phase margin for our closed loop system.

In order to design gains that meet these requirements, we used a Ziegler-Nichols tuning method. This method was created for process control environments and has been refined through numerous studies. It seeks to use the natural characteristics of the process step response as a method to create gains for the system to be controlled (Astrom & Hagglund, 1995). The equations for the tuning method are below.

$$K_p = \frac{1.2T}{L} \quad \textit{Eq. 32}$$

$$K_I = \frac{K_p}{2L} \quad \textit{Eq. 33}$$

$$K_d = \frac{K_p L}{2} \quad \textit{Eq. 34}$$

Where the parameters T and L may be described using the step response of the plant. As an example, the plant model with both the PK model and the selected delay is shown below.

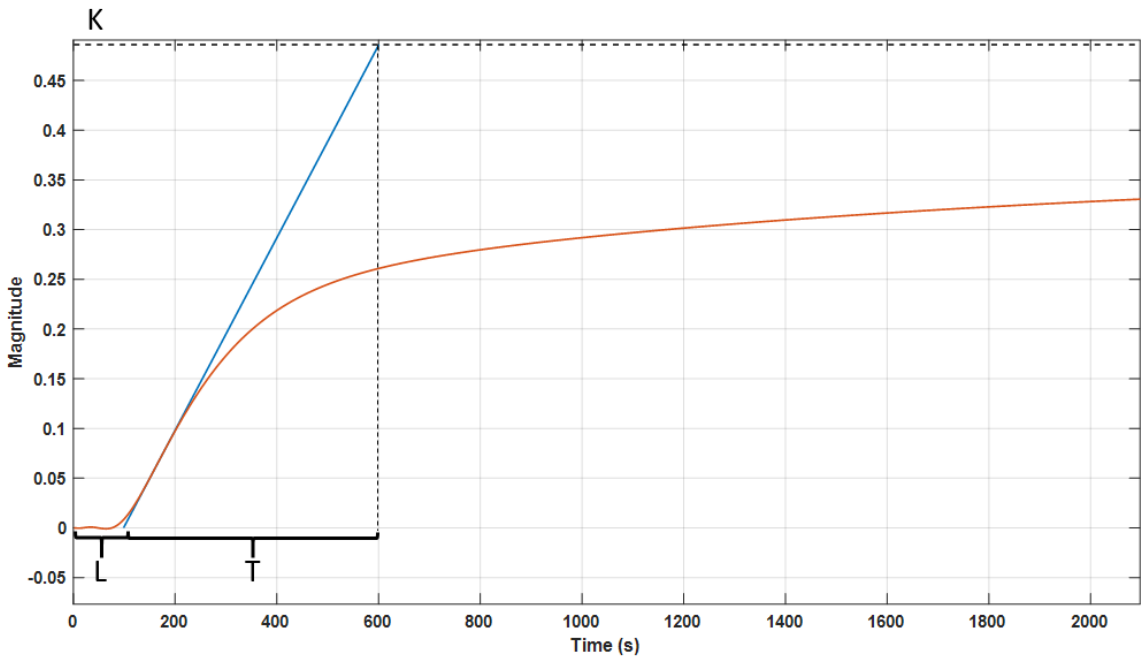

Figure 13: Step Response with Ziegler-Nichols Parameters

The red line is the step response while the blue line is the tangent line at the steepest point in the step response. L is the time from 0 to the x-axis intersection with the tangent line and T is the time from the tangent line intersection with the x-axis and the tangent line intersection with the final value of the step response. It should be noted that the step response for this system

reaches steady state at nearly ten thousand seconds which is why we don't see the step response reach its final value in this figure.

While using these coefficients for design, we also swept the actuator parameters to determine the overall closed loop performance from a linear perspective. A range of damping coefficients from 0 to 1 were run as well as a logarithmic range of time constants.

At each iteration, a linear closed loop system was formed to estimate the nonlinear closed loop system that will ultimately be simulated. The PK model, being linear initially, was coupled with a pade approximation for delay to form the plant model. A pade approximation effectively desires to fit a certain equation order to a delay in order to approximate it in a linear sense. In this way, we can linearly approximate the effects of algorithmic delay in the system. A lower order pade approximation is typically desirable as the higher order approximation may introduce undesirable effects. A third order approximation was found to be sufficient for our application. Since we have a variable delay in our system, the average delay expected for our system throughout a range of recovery parameter uncertainties was used as shown in Chapter 3.

The actuator and controller models were then included to form the loop gain at the input. This is done by breaking the loop at the input and taking the preceding systems in series until a full loop is completed back to the original break point.

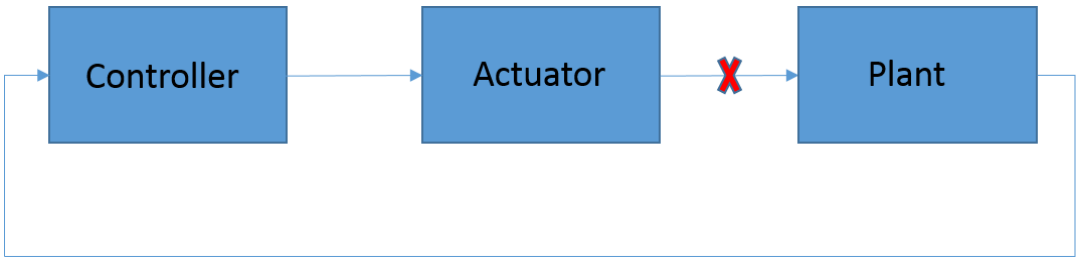

Figure 14: Breaking the Loop

$$Lu_{KG} = -sys_{actuator} * sys_{controller} * sys_{plant} \quad \textit{Eq. 35}$$

$$S = \frac{1}{1 + Lu_{KG}} \quad \textit{Eq. 36}$$

$$T = 1 - S = \frac{Lu_{KG}}{1 + Lu_{KG}} \quad \textit{Eq. 37}$$

This was then used to form the sensitivity and co-sensitivity systems as shown in equations 35 through 37 above. The co-sensitivity, T, also being the closed loop response, was then used to get the time domain performance out of the system as well as gain and phase margin. The loop gain was also taken independently and the loop gain crossover frequency was calculated. The slope of the loop gain (dB/decade) was then determined at this point. Of course, if the chosen actuator and controller destabilized the system then the iteration would be considered a failure and the design would iterate to the next range of values. The stability of the system was simply determined by taking the eigenvalues of the closed loop system and ensuring they were all negative, which indicates all the poles of the system are on the left half plane and the system is stable. The various combinations of actuator time constant and damping coefficient yielded a variety of results. We desired to maintain a gain margin above 6dB and a phase margin greater than 45 degrees in addition to minimal overshoot and a roll off steepness near -20dB. To meet these design requirements and shape a co-sensitivity and sensitivity maximum value that

was relatively small as well as minimize overshoot, the following actuator specifications were chosen.

Table 5: Actuator Design Parameters

Damping Coefficient	Time Constant (s)
1	91.03

These actuator parameters yielded excellent gain and phase margin results. With the Ziegler-Nichols tuning method shown above in equations 32-34, the system yielded a gain margin of 26.4dB and a phase margin of 160.45 degrees as well as a roll off steepness of -21dB at the loop gain crossover frequency. These excellent stability margins should not be altogether surprising due to the fact that our system was already quite stable.

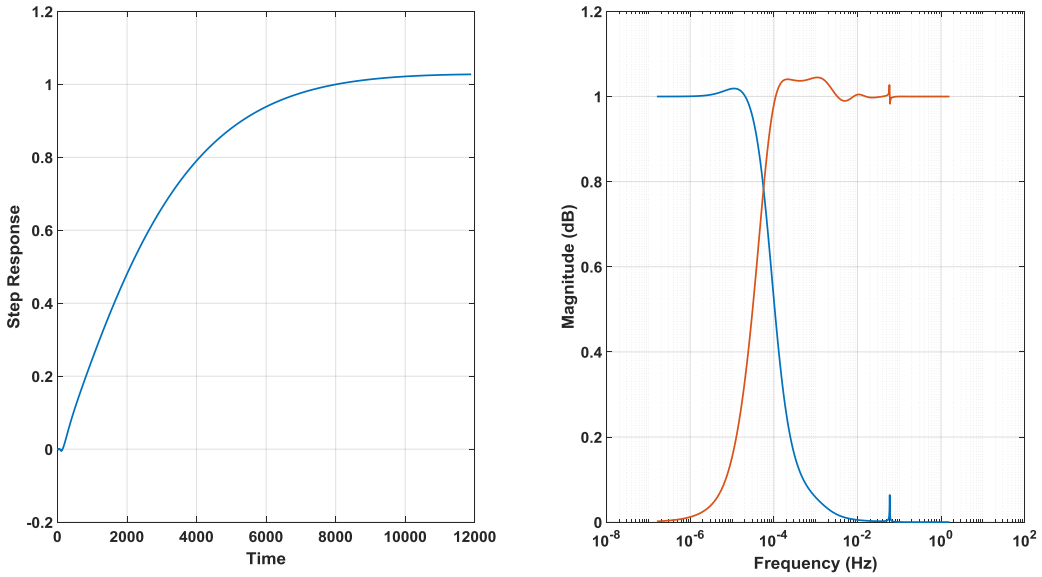

Figure 15: Gain Design Closed Loop System Performance

In addition, the figure above shows the step response and sensitivity/co-sensitivity functions. The figures show a steady-state overshoot with a two percent steady state error. Both functions show magnitudes over 1, however they are relatively small and since our system has excessive margin, these magnitudes should not pose any stability issues.

Chapter 6: Simulation Results

In order to test the model's ability to control the BSR of a certain patient, a BSR trajectory was created that correlates with the values in the table below.

Table 6: CLAD Set Point Reference Commands

BSR	0.8	0.8	0.5	0.5	0.2	0.2
Time (sec)	0	25000	25001	45000	45001	65000

This trajectory was chosen so as to stress a large range of BSR values while also allowing us to view the steady state values at each of these conditions by commanding the system hold at that BSR for an extended period of time. The first simulation done is for the nominal recovery parameters shown in Table 2 with no noise added and with the patient data that the controller was designed to.

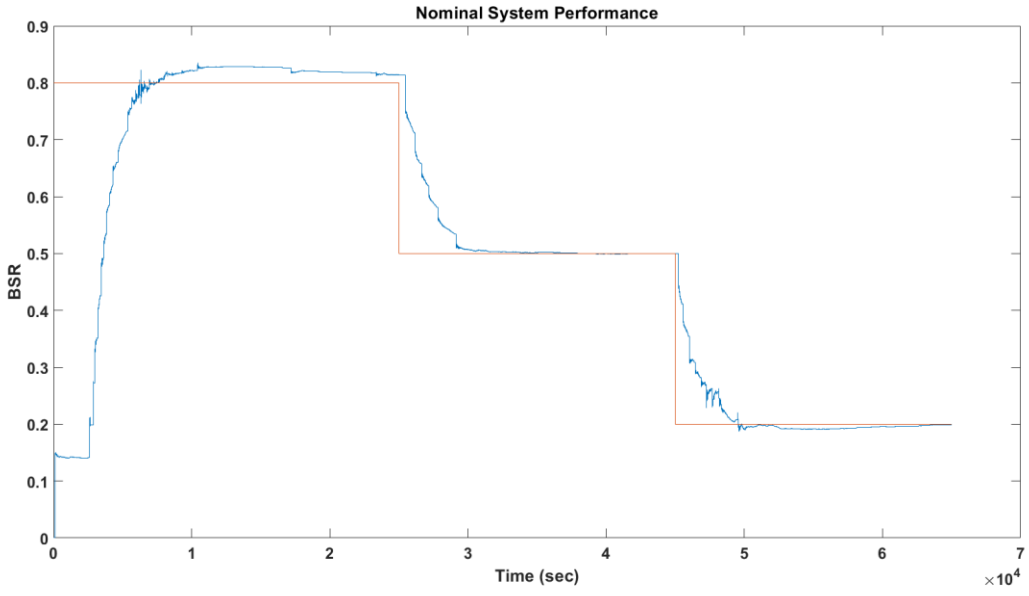

Figure 16: Nominal System Performance

The results show the BSR closely tracking the step responses as they are commanded separately throughout the trajectory. An obvious first takeaway is that the initial BSR comes out to a value of roughly 0.15. This is because the minimum burst suppression ratio our algorithm could estimate while maintaining a relatively monotonic increase throughout the range of c_2 is also 0.15. Due to this, our conversion from concentration in the effect site compartment was limited to the minimum c_2 value that could be attained with near monotonic performance. These results show that there is a steady state error after the system is commanded to a set-point of 0.8 as well as a small initial overshoot. This is not entirely unexpected due to the step response results from our gain selection and actuator design in Chapter 5.

The second simulation run was done with a different patient than what the gains were designed to. In order to model what would result in less propofol infusion than what the original patient needed in an attempt to cause the system to overshoot, a patient type was chosen with a female gender, age of 65, weight of 50 kilograms, and a height of 125 centimeters. Due to the decreased weight and age and increased height in comparison to the original patient, this patient requires less propofol to reach a certain effect site concentration.

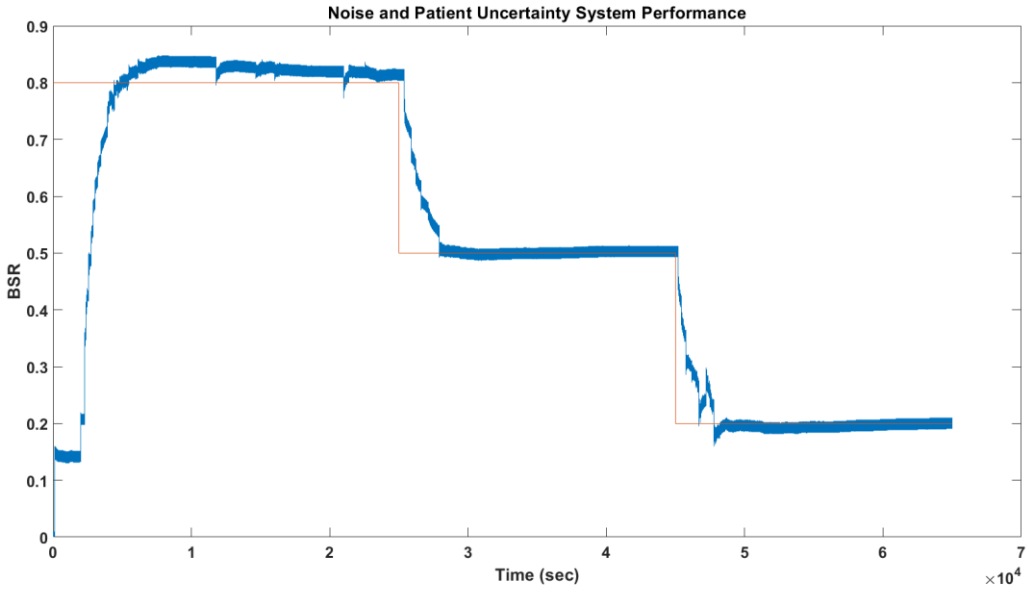

Figure 17: Noise and Patient Uncertainty System Performance

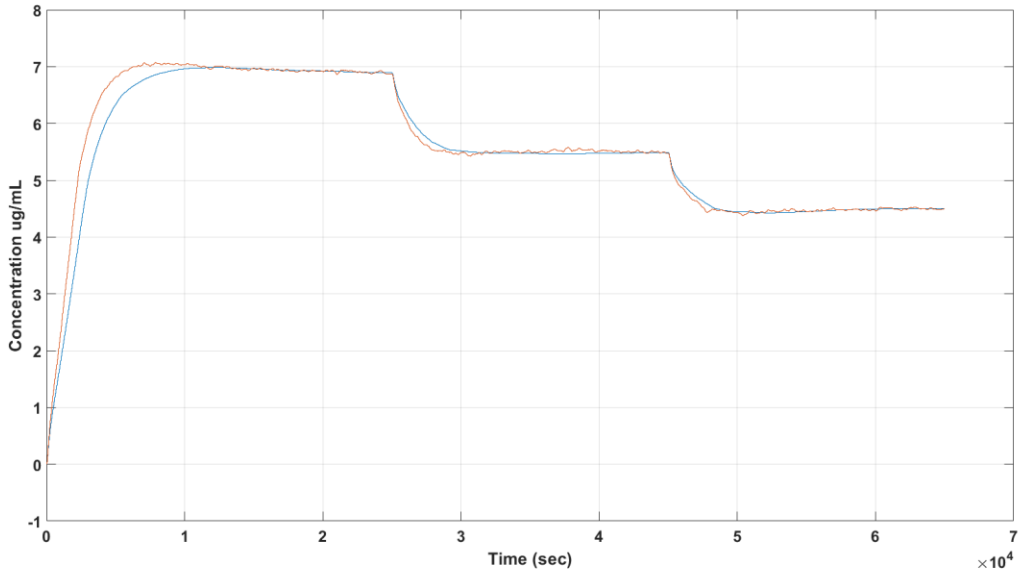

Figure 18: Effect Site Concentration with Patient Uncertainty and Noise

The blue line in the figure above represents the nominal patient concentration while the red line shows the concentration of the patient with added uncertainty. The main difference shown between the two is that the system with patient uncertainty does pose a faster rise time and higher overshoot than the original patient. However, this overshoot is quickly dissipated and

a steady state result similar to that of the nominal case is reached quickly. Aside from this, the closed loop system performance shows the evident noise that was added to the system but maintains a mean value around the commanded BSR.

Finally, a simulation with a drastically different patient and the maximal increase in recovery parameters was run to see how the system deals with maximal uncertainty in the system it was designed for. The maximal increase in recovery parameters are as shown in Figure 6. The immediate impact of this increase is the relationship between BSR and c_2 which will in turn result in a change in the effect site concentration required to achieve a certain BSR. This change also increased the effective lower bound of control due to the increased suppression that we witnessed in Figure 6 when the recovery parameters were all simultaneously increased. This increased suppression, coupled with the static lower bound set on the concentration to c_2 conversion, caused the system to be unable to both attempt a lower c_2 value than 10 for the purpose of searching for a smaller BSR to be controlled and, therefore, unable to reach a BSR at the target level that we desired. The added noise to the system resulted in similar behavior to the model which we added patient uncertainty and noise. In addition, larger spikes than would be expected by noise generation occurred at high BSRs. It is likely that the algorithm had seen a c_2 value or range of c_2 values that caused this jump which could be unique to the recovery parameter uncertainty that was imparted on the model.

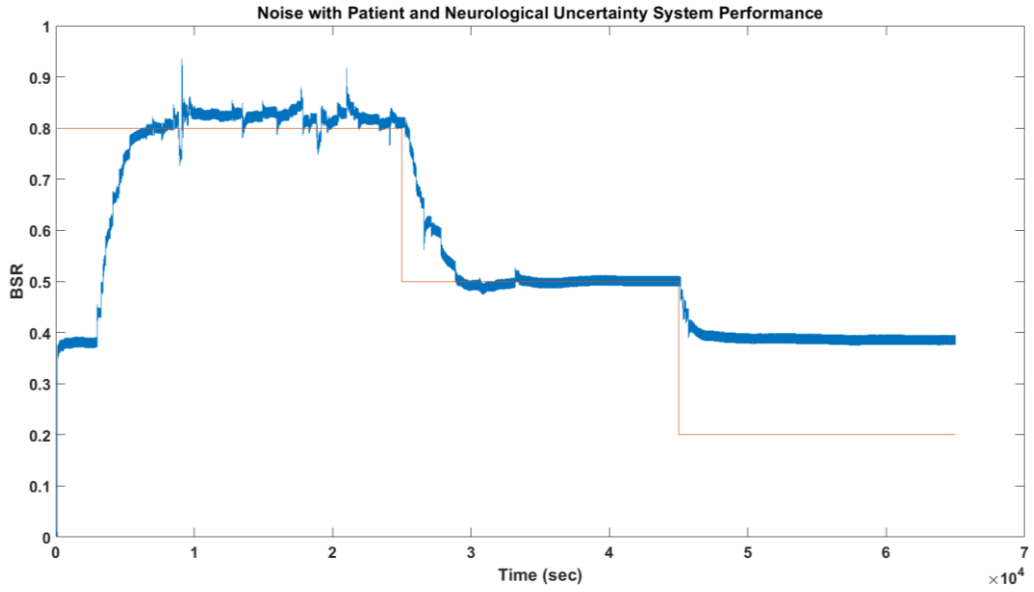

Figure 19: Noise with Patient and Neurological Uncertainty System Performance

Chapter 7: Conclusions and Future Work

This paper presented an algorithm for detecting burst suppression in the modified Wilson-Cowan equations (Liu & Ching, 2017) as well as a closed loop control method for targeting specific BSR's amid patient and neurological uncertainty. Design goals of this system were sufficiently met in the nominal case and the case with added noise and patient uncertainty where the maximum observed overshoot was roughly six percent and the system's steady state error was roughly two percent. In addition to this, a bridge was created to convert the output of the 4-dimensional PK Schnider model (Schnider, et al., 1998 and Absalom, Mani, DeSmet, & Struys, 2009) to c_2 values which are fed as inputs to the modified Wilson-Cowan equations (Liu & Ching, 2017). This allows the two models to be used in conjunction to model the overall system behavior as it relates to propofol impact on neuron firing rates.

Future work may be done to build on these results by improving the burst suppression ratio estimation algorithm to track values at the lower bounds of the equations described by (Liu & Ching, 2017). The reduced performance due to neurological uncertainty that we witnessed in the large uncertainty run was due to the BSR estimation algorithm estimating a lower bound that was too high for the system to be controlled at a BSR of 0.2. The minimum BSR this system could be controlled to was 0.38, which is where the system held during the portion of the trajectory that a BSR of 0.2 was commanded. During the reachable parts of the envelope, however, the control scheme worked with similar performance to that of the nominal case. This would suggest that an adaptive estimation algorithm needs to be developed to estimate the neurological deviations from the nominal case and that the control methodology used herein is sufficient to accomplish our design goals for this application. This adaptive estimation algorithm could be similar to a bolus infusion stage where the algorithm sets itself on initialization before

actively controlling the patient's BSR levels. In addition, a more robust translation between the effect site concentration and the c_2 value could be made to include a wider range of patients and change in the recovery dynamics of the modified Wilson-Cowan equations (Liu & Ching, 2017). Other parameter changes could also be accounted for by varying consumption parameters in addition to the recovery to create a more diverse range of uncertainty which could account for a greater diversity in patient firing rates. The initial tuning done for the adaptive estimation algorithm during bolus could also be done concurrently for the concentration to c_2 conversion to make a more complete robust solution.

References

Absalom, A. R., Mani, V., DeSmet, T., & Struys, M. M. (2009). Pharmacokinetic models for propofol - defining and illuminating the devil int he detail. *British Journal of Anaesthesia*, 26-37.

Astrom, K. J., & Hagglund, T. (1995). *PID Controllers: Theory, Design, and Tuning* (Second Edition ed.). Research Triangle Park, N.C: International Society for Measurement and Control.

Brown, E. N., Lydic, R., & Schiff, N. D. (2010). General Anesthesia, Sleep, and Coma. *The New England Journal of Medicine*, 363:2638-2650.

Chemali, J. J., Wong, K. K., Solt, K., & Brown, E. N. (2011). A state-space model of the burst suppression ratio. *Annual International Conference of the IEEE Engineering in Medicine and Biology* (pp. 1431-1434). Boston: IEEE.

Chemali, J., Ching, S., Purdon, P. L., Solt, K., & Brown, E. N. (2013). Burst Suppression probability algorithms: state-space methods for tracking EEG burst suppression. *Journal of Neural Engineering*, 10(5):056017.

Ching, S., Liberman, M. Y., Chemali, J. J., Westover, B. M., Kenny, J., Solt, K., . . . Brown, E. N. (2013). Real-time Closed-loop Control in a Rodent Model of Medically-induced Coma Using Burst Suppression. *Anesthesiology*, 119(4):848–860.

Ching, S., Purdon, P. L., Vijayan, S., Kopell, N. J., & Brown, E. N. (2012). A neurophysiological-metabolic model for burst suppression. *Proceedings of the National Academy of Sciences of the United States of America*, 109(8): 3095-3100.

Liu, S., & Ching, S. (2017). Homeostatic dynamics, hysteresis and synchronization in a low-dimensional model of burst suppression. *Mathematical Biology*, 74:1011-1035.

Martin, D., Penrod, L., Obrist, W., Kochanek, P., Palmer, A., Wisniewski, S., & DeKosky, S. (1997). Treatment of traumatic brain injury with moderate hypothermia. *The New England Journal of Medicine*, 8:540-546.

Plummer, S. G., Ibala, R., Hahm, E., An, J., Gitlin, J., Deng, H., . . . Akeju, O. (2019). Electroenceaphalogram dynamics during general anaesthesia predict the later incidence and duration of burst-suppression during cardiopulmonary bypass. *Clinical Neurophysiology*, 55-60.

Schanechi, M. M., Chemali, J. J., Liberman, M., Solt, K., & Brown, E. N. (2013). A Brain-Machine Interface for Control of Medically-Induced Coma. *PLOS Computational Biology*, 10(5):056017.

Schnider, T. W., Minto, C. F., Gambus, P. L., Andresen, C., Goodale, D. B., Shafer, S. L., & Youngs, E. J. (1998). The Influence of Method of Administration and Covariates on the Pharmacokinetics of Propofol in Adult Volunteers. *Anesthesiology*, 1170-1182.

Westover, B. M., Kim, S.-E., Ching, S., Purdon, P. L., & Brown, E. N. (2015). Robust control of burst suppression for medical coma. *Journal of Neural Engineering*, 12: 046004.

Westover, B. M., Kim, S.-E., Ching, S., Purdon, P. L., & Brown, E. N. (2015). Robust control of burst suppression for medical coma. *Journal of Neural Engineering*, 12(4):046004.

Wilson, H., & Cowan, J. (1972). Excitatory and inhibitory interations in localized populations of model neruons. *Biophysics*, 12:1-24.

Wilson, H., & Cowan, J. (1973). A mathematical theory of the functional dynamics of cortical and thalmic nervous tissue. *Kybernetik*, 15:55-80.

Young, G. (2000). The EEG in coma. *Journal of Clinical Neurophysiology*, 5:473-485.