

Dynamic Stability of an Axially Loaded Elastically Restrained Column

by

Richard Dromey

A thesis submitted to the Faculty of Graduate Studies
in partial fulfillment of the requirements for the degree of
Master of Science in Civil Engineering

Supervisor

Dr. Jian Deng

Associate Professor - Department of Civil Engineering

Co-Supervisor

Dr. Kefu Liu

Professor - Department of Mechanical Engineering

Lakehead University

Thunder Bay, ON

December 2022

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners. I understand that my thesis may be made electronically available to the public.

Acknowledgements

Words are wholly inadequate to express my gratitude to my wife and son for their infinite patience with my persistent state of distraction while working on this thesis. Many thanks for the countless hours and minutes they respectively allowed me to work without interruption.

On a more directly related note, this thesis would not have been possible without the assistance of Dr. Jian Deng, both for first developing the concept for the numerical method used here and for having the wisdom to know how much is too much to fit in a master's thesis.

On the theoretical side of experimental things, thank you to Dr. Kefu Liu for helping me to avoid a number of particularly inconvenient pitfalls in the design of my apparatus. On the experimental side, many thanks to Cory Hubbard and Morgan Ellis for their assistance with lab equipment and machining, and machining and lab equipment, respectively.

I would also like to thank Dr. Yanglin Gong for introducing me to a variety of non-ideal connections. This thesis would not be what it is without his continued guidance and encouragement in my investigations, and for your time and effort in reviewing this thesis. Thank you also to Dr. Ali Tarokh for taking the time to review my work and provide valuable feedback.

Lastly, thank you to Haining Li for your assistance at various points.

Abstract

The dynamic stability of axially loaded columns is a key problem in structural analysis, and in earthquake engineering particularly. Experience has shown that columns occasionally buckle when subject to dynamic axial loads that are only a small fraction of the load-carrying capacity predicted by static methods. While the mechanism behind these failures has been identified, theoretical studies have focused almost exclusively on pin-pin (and, to a lesser extent, fixed-fixed) connections. In practice, however, most columns have fixities between these two extremes and are described as having semi-rigid or elastically restrained supports. Very few studies have been conducted on this condition, and of the studies performed, none have included experimental verification. Additionally, calculation of the stability based on existing methods is very computationally expensive, as they calculate the response of the column, rather than the stability behavior. As a result, dynamic stability theory is not yet directly applicable to most design situations. To apply this theory to key problems in earthquake and structural engineering, a more practical approach to computing the stability behavior of more complex support conditions is required. This research studies the dynamic stability of axially loaded columns with elastically restrained supports theoretically, numerically, and experimentally. The equation of motion of an elastically restrained column is first derived and converted to a Mathieu equation, from which the stability regions can be obtained using Bolotin's method. Second, a numerical method is developed to compute the state transition matrix for a particular excitation, giving the stability condition and response simultaneously. This numerical method is compared to the analytical results to determine the relative error associated with both the analytical and numerical solutions. Third, a novel experimental apparatus capable of simulating an axially loaded column with elastically restrained supports is designed in the laboratory. The experimental apparatus is capable of providing experimental verification for the matrix-based numerical method and is fully automated to allow the collection of large quantities of data. Due to unforeseen equipment damage, the experimental data could not be collected, but the automation was sufficiently advanced at the time of damage to be proven in principle.

Keywords: Dynamic stability, elastic restraint, semi-rigid, connection, experimental apparatus

List of Symbols

α	fixity coefficient
β_0	generalized damping coefficient of the free response
β_*	numerical damping coefficient for constant excitation
χ	shape factor
Δ_0	initial maximum deflection of the column
Δ_1	peak maximum deflection during one half-cycle of the periodic response
Δ_2	peak maximum deflection during a second half-cycle of the periodic response
δ	deflection of a member
ϵ	relative error
ϵ_M	relative error introduced by the small-angle approximation in a single calculation
ϵ_a	relative discretization error in the amplitude domain
ϵ_c	compounded relative error introduced by the small-angle approximation over a full period
ϵ_f	relative discretization error in the frequency domain
ϵ_{ana}	relative error in the analytical solution
ϵ_{num}	relative error in the numerical matrix solution
η_n	simplified loading coefficient associated with segment n
κ_f	rotational stiffness of the fixed connections
κ_p	rotational stiffness of the pinned connections
λ	root of the characteristic equation in an ODE
μ	loading parameter
$\hat{\nu}$	ratio of the excitation angular velocity to the natural angular velocity of the unloaded column
ν	angular velocity of the axial excitation
ω	natural angular velocity of the column
ω_0	natural angular velocity of the unloaded column
ω_n	natural angular velocity associated during segment n
β	generalized damping coefficient
ϕ	ratio of the excitation angular velocity to the the natural angular velocity of the column

ψ	temporary variable used to avoid consideration in the complex domain
ρ	connection fixity
$\tilde{\rho}$	approximate connection fixity
τ	duration of a single segment
τ	time at the end of the first segment, or duration of a single segment
θ	angle of rotation
θ_r	pinned-end rotation
ζ	generalized Mathieu-Hill function excitation function
A_n	coefficients of the sine-function terms in an infinite series
b	total bits of accuracy
b_Δ	bits of increased accuracy due to oversampling
B_n	coefficients of the cosine-function terms in an infinite series
\mathbf{C}	numerical matrix method constant arising from the hyperbolic cosine function
\mathbf{C}_n	a normalization of $\mathbf{C}_{n\star}$
$\mathbf{C}_{n\star}$	a generalization of \mathbf{C}
\mathcal{C}	an arbitrary matrix
C	an arbitrary constant distinguished by subscript
c	coefficient of internal friction
d_1	distance between opposing bearings in the pin-pin connection
d_2	distance between adjacent bearings in the fixed-fixed connection
\hat{E}_v	ratio of the kinetic energy in the springs to the kinetic energy in the specimen
\hat{E}_b	ratio of the elastic strain energy in the springs to the elastic strain energy in the specimen
E	modulus of elasticity of a member
E_0	energy in the system at $t = 0$
E_Δ	proportion of total energy conserved between cycles
E_∞	energy in the system at $t = \infty$
E_b	elastic strain energy due to bending
E_v	kinetic energy
E_{b_c}	elastic strain energy due to bending in the connection
E_{b_k}	elastic strain energy due to bending in a spring
E_{b_s}	elastic strain energy due to bending in the specimen
E_{v_c}	kinetic energy in the connection
E_{v_s}	kinetic energy in the specimen

F	an arbitrary function distinguished by subscript
f	index of summation for Fourier series
f_0	central frequency of the FFT
F_D	lateral damping forces
F_I	lateral inertial forces
F_P	lateral forces induced by an axial load
I	area moment of inertia of a member
I_k	area moment of inertia of the springs
I_s	moment area of inertia of the specimen
k_f	rotational stiffness of the fixed-fixed connection
k_t	transverse stiffness of the bearings in the pin-pin connection
k_{eq}	equivalent spring stiffness
L	length of a member
L_α	length of the column with elastically restrained supports
L_p	half-length of the square tube
L_f	length of the column with fixed ends
L_p	length of the column with pinned ends
\mathcal{M}_n	a series of unique 2×2 matrices
M	order of the small-angle approximation
m	index of summation when calculating a small-angle approximation
m	mass per unit length of the column
M_k	maximum moment in each spring
m_k	mass of each spring
N	total number of segments considered in the discretization
n	number of the current segment being considered
o	order of the analytical solution
\hat{P}_d	ratio of the dynamic axial excitation to the Euler buckling load of the column
\hat{P}_n	ratio of the axial excitation during segment n with the Euler buckling load
\hat{P}_s	ratio of the static axial excitation to the Euler buckling load of the column
\mathcal{P}	the matrix product in the numerical method
P	applied axial force
P_d	dynamic component of the axial excitation
P_E	euler buckling load for an elastically restrained column
P_f	axial excitation function associated with a specific term in a Fourier series

P_n	dynamic component of the axial excitation with an angular velocity of $n\nu$
P_s	static component of the axial excitation
q	column response function
q_0	initial value of the response function
q_1	value of the response function at the end of the first segment
q_N	position of the column at the end of one full period
q_N	value of the response function at the end of the period
q_n	position of the column at the end of segment n
\mathcal{R}	two-dimensional rotation matrix
\mathcal{R}_N	two-dimensional rotation matrix expressed in terms of N
R	the number of the region being considered
r	transformed response function in undamped coordinates
r_0	initial value of the transformed response function
r_1	value of the transformed response function at the end of the first segment
r_N	value of the transformed response function at the end of the period
\mathbf{S}	numerical matrix method constant arising from the hyperbolic sine function
\mathbf{S}_n	a normalization of \mathbf{S}_{n^*}
\mathbf{S}_{n^*}	a generalization of \mathbf{S}
s	sample rate
s_f	estimate of the slope at the endpoint
s_i	estimate of the initial slope
S_P	spectral density of the axial excitation function
s_{m1}	first estimate of the slope at the midpoint
s_{m2}	second estimate of the slope at the midpoint
T	excitation period
t	time
t_1	time at which Δ_1 was measured
t_2	time at which Δ_2 was measured
t_n	phase-shifted relative time
v_f	estimate of the value at the endpoint
v_{m1}	first estimate of the value at the midpoint
v_{m2}	second estimate of the value at the midpoint
V_{out}	output voltage
V_{sig}	signal voltage

\mathbf{x}	generic variable placeholder
X	normalized deflected shape of an elastically restrained column
x	position along the major axis of a member
X_f	normalized deflected shape of a fixed-fixed column
X_p	normalized deflected shape of a pin-pin column
y	position aligned perpendicular to both the major and bending axes of a member

Contents

Author's Declaration	i
Acknowledgements	iii
Abstract	v
List of Symbols	vii
Table of Contents	xiii
List of Figures	xvii
1 Background	1
1.1 The Dynamic Buckling Mechanism	1
1.2 The Dynamic Equation of Motion for an Axially Loaded Column	3
1.3 Solving Mathieu Equations	6
1.3.1 The Analytical Method	6
1.3.2 Numerical Methods	9
1.4 Semi-Rigid Connections	11
1.5 Experiments on the Dynamic Stability of Columns	13
1.6 Research Objectives and Thesis Outline	14
2 The Analytical Solution	15
2.1 Separating the PDE	15
2.1.1 The Modal Function	16
2.1.2 Converting to Mathieu Equation Form	19
2.2 Finding a Solution	21
2.2.1 Floquet's Theory	21
2.2.2 Transformation to Matrix Form	22
2.2.3 Separating the Matrix	25
2.2.4 Finding the Stability Boundaries	25

2.3	Restating the Stability Boundaries	27
2.3.1	Changing the Variables	28
2.3.2	Re-plotting the Solutions	29
2.4	Summary	29
3	The Numerical Matrix Solution	33
3.1	The Numerical Method	33
3.1.1	Computational Complexity	34
3.2	Discretization	35
3.2.1	Discretization Methods	35
3.2.2	Discretization Errors	37
3.2.3	Approximation Errors	39
3.3	Determining the Column Response	42
3.3.1	Deriving the Response In Regular Coordinates	43
3.3.2	Deriving the Response in Damped Coordinates	47
3.3.3	The Response of an Undamped Column	52
3.4	Stability Boundaries Estimates	52
3.4.1	The Undamped Stability Boundaries	53
3.4.2	The Damped Stability Boundaries	55
3.5	Summary	58
4	Comparison of the Analytical Solution and the Numerical Matrix Solution	59
4.1	Comparing the Solutions Visually	60
4.2	Quantifying The Relative Error in the Undamped Column	64
4.2.1	Accuracy of the Numerical Solution	64
4.2.2	Accuracy of the Analytical Solution	70
4.3	Estimating the Relative Error in the Damped Column	75
4.4	Summary	77
5	An Automated Apparatus for Obtaining the Stability Boundaries Experimentally	81
5.1	The Experimental Apparatus	82
5.1.1	Sensors and Measurement	84
5.2	Designing the Support Connections	85
5.2.1	Determining the Fixity of a Connection	86

5.2.2	Specimen Selection	87
5.2.3	The Classical Connections	88
5.2.4	The Elastically Restrained Support Conditions	92
5.2.5	Rotational Stiffness of the Elastic Connections	98
5.3	Automating the Experiment	101
5.3.1	Exciting the Specimen	101
5.3.2	Detecting Instability and Aborting Tests	104
5.3.3	Retrieving the Data	105
5.4	Data Processing	106
5.4.1	Summarizing the Data	106
5.4.2	Processing the Transfer Functions	108
5.4.3	Processing the Free Response	111
5.4.4	Processing the Experimental Results	115
5.5	Experimental Error	119
5.6	Observed Specimen-Connection Coupling	120
5.6.1	Relative Bending Energy	122
5.6.2	Relative Kinetic Energy	123
5.6.3	Reducing the Influence of the Connection	125
5.7	Summary	127
6	Conclusions and Future Work	129
6.1	Conclusions	129
6.1.1	The Analytical Solution	129
6.1.2	The Numerical Matrix Solution	130
6.1.3	Comparison of the Analytical Solution and the Numerical Matrix So- lution	130
6.1.4	Experimental Observations	131
6.2	Future Work	132
	References	135

List of Figures

1.1	A free-body diagram of an axially excited elastically restrained column	4
1.2	Support conditions for dynamic stability of columns. (1) Pin-pin (2) Fixed-fixed (3) Semi-rigid or elastically restrained	12
2.1	A comparison of ρ , $\tilde{\rho}$, and α	19
2.2	Variation of ω and μ with fixity (α)	21
2.3	Normalized second-order approximation of the stability region boundaries . .	26
2.4	Normalized fifth-order approximation of the stability region boundaries . . .	27
2.5	Transformed fifth-order approximation of the stability region boundaries ($\hat{P}_s = 0$)	30
2.6	Transformed fifth-order approximation of the stability region boundaries, with relative static excitation $\hat{P}_s = \pm 0.2$	31
2.7	Transformed fifth-order approximation of the stability region boundaries, with relative static excitation $\hat{P}_s = \pm 0.5$	32
3.1	Discretizing a continuous function into a step function	35
3.2	Relative errors associated with the use of the small angle approximations by number of segments used	42
3.3	Natural logarithm of stability for an undamped column ($N = 500$) using the first- and second-order approximations	54
3.4	Natural logarithm of stability for an undamped column ($N = 50$) using the second-order and third-order approximations	56
3.5	Effect of varying the order of approximation on the natural logarithm of stability for a damped column ($\beta_0 = 0.1, N = 500$)	57
4.1	Comparison of the undamped numerical solution ($N = 20$) and 2 nd -order through 4 th -order analytical solutions	60
4.2	Zoomed in view of the undamped numerical solution ($N = 20$) and 2 nd -order through 4 th -order analytical solutions	61

4.3	Comparison of the adjusted damped ($\beta_0 = 0.05$) numerical solution ($N = 20$) and 2 nd -order through 4 th -order analytical solutions	62
4.4	Zoomed in view of the damped ($\beta_0 = 0.05$) numerical solution ($N = 20$) and 2 nd -order through 4 th -order analytical solutions	63
4.5	Relative error of the numerical solution when compared to the 5 th -order analytical solution for an undamped column	65
4.6	Relative error as a function of N for the first six stability regions	66
4.7	Minimum relative error as a function of N for the 2 nd through 6 th stability regions expressed as multiple of the minimum relative error of the 1 st stability region	69
4.8	Estimated error in the undamped numerical method	70
4.9	Estimated error in the analytical method ($R = 2$)	71
4.10	Relative error of the numerical solution when compared to the 5 th -order analytical solution for a damped column ($\beta_0 = 0.02$)	76
4.11	Relative error as a function of N for the first five stability regions ($\beta_0 = 0.02$)	78
5.1	Photograph of the experimental apparatus	81
5.2	Schematic of the experimental apparatus	83
5.3	Circuit diagram for the load cell signal amplifier	85
5.4	Pin-type interchangeable connections for the experimental apparatus (attached)	89
5.5	Plan of the pin-type interchangeable connections	90
5.6	Fixed-type interchangeable connections for the experimental apparatus (not attached)	91
5.7	Plan of the fixed-type interchangeable connections	92
5.8	Photograph of the constructed connection (not attached)	93
5.9	Plan of the elastic interchangeable connections	94
5.10	Experimental procedure	102
5.11	Circuit diagram for the deflection limiter	104
5.12	Comparison of the filtered signal and straight-line-connection of signal summary	107
5.13	Comparison of the FFT and DFT in the region of the central frequency measured for the pinned specimen	112
5.14	Effect of filtering on the amplitude of the free response of the pinned specimen	113
5.15	Effect of filtering on the amplitude decay rate over time for the free response of the pinned specimen	114

5.16 Specimen reaction to 0.25 V pre-gain excitation 118

Chapter 1

Background

It has long been observed that members subject to an axial load will, at some load, undergo spontaneous, extreme, and unrecoverable lateral deflection. The best-known formal work on this subject was by Euler in 1757, who derived a convenient formula for predicting the buckling load of straight slender sections of constant cross section that are not subject to any lateral or dynamic loading. While convenient, this work does fail to explain what exactly buckling is, instead only predicting when it occurs. Lagrange later developed the energy approach, which better illustrates what buckling is generally. At its simplest, buckling occurs when the energy required to compress the member axially exceeds the energy required to obtain the same axial deflection through bending. Once this energy level is reached, even an infinitesimal imperfection or deflection will grow exponentially, leading to failure.

In dynamic buckling, the time-sensitive nature of each of the energy components makes the solution less obvious, but fundamental behavior remains the same. The dynamic behavior and stability conditions vary depending on the nature of the problem being considered, but in every case, the solution is based on the governing equation for that problem.

1.1 The Dynamic Buckling Mechanism

The primary difference between dynamic buckling and static buckling is that in the dynamic case the energy level in the column is a function of time. The periodic axial load adds energy to the column proportional to the axial load and the oscillation of the column removes energy through destructive interference based on the relative frequency of the excitation. As the energy in the column increases, so too does the deflection, which causes both the energy added by the axial load and the energy removed through destructive interference to increase. When the axial load adds energy to the system faster than it can be removed, the column buckles.

For an imperfect column subject to an axial load, the equilibrium shape of the column is equal to some multiple of its initial shape. If the axial load is dynamic, the equilibrium shape of the column is always changing, with the column constantly attempting to catch up. When the axial load is compressive, the column accelerates outwards in an attempt to reach stability, and when the axial load is tensile, the column accelerates inwards. As these accelerations coincide with the motion of the column, they cause the energy in the column to increase.

Dynamic instability is most easily observed by considering the case of an imaginary, undamped, vibrating column subject to alternating compressive and tensile loads at twice its natural frequency. Beginning with the column perfectly straight, the vibrations imperceptibly small, and the axial load compressive, each excitation follows the same sequence:

1. As the column moves outwards, the compressive axial load magnifies the deflected shape of the column, causing the column to accelerate outwards.
2. Because the acceleration of the column is only equal to zero once it reaches its equilibrium shape, it reaches this point with some non-zero velocity, and begins to decelerate.
3. When the axial load becomes tensile the column has overshoot its equilibrium shape and is at rest. The deflection of the column is now greater than the maximum deflection of the original vibrations.
4. When the velocity reaches zero, the column is momentarily at rest, and begins to accelerate inwards. Because the excitation is at exactly twice the natural frequency of the column, which has gone through one quarter of a cycle, the excitation has gone through one half of a cycle, and is transitioning from compressive to tensile.
5. The tensile axial load acts to minimize the deflected shape of the column, accelerating the column inwards.
6. When the column is again straight, the axial load is transitioning from tensile to compressive. As the column is oscillating at its natural frequency, all of the elastic strain energy in the column is conserved and becomes kinetic energy, resulting in a higher velocity than in the original vibrations.
7. This process then repeats for the next period of the excitation, but with the direction of motion reversed and a larger initial velocity.

This scenario does not typically occur in practice, as all real columns include some degree of damping due to internal friction and it is unlikely that the excitation will occur at exactly twice the natural frequency of the column, both of which act to decrease the energy in the system and counter this cycle. The internal friction does this by transforming some of the kinetic energy into heat. When the column is excited at some frequency other than exactly twice its natural frequency, energy is lost during the transformation between elastic strain energy and kinetic energy, as the inertial forces and the elastic bending forces are not precisely equal. When the excitation frequency is higher than the natural frequency, the elastic bending forces in the column cannot accelerate the column quickly enough to maintain the oscillations, and the axial excitation must provide the missing energy. When the excitation frequency is lower than the natural frequency, the elastic bending forces in the column attempt to accelerate the column too quickly, and the axial excitation must counter this imbalance to continue the oscillations, again leading to a loss of energy. Both situations lead to a loss of energy in the column, which acts to increase stability.

The overall stability of the column is dependent on the relative influence of each of these effects.

1.2 The Dynamic Equation of Motion for an Axially Loaded Column

Deriving the equation of motion of the column begins with the definition of a Euler column. A Euler column was chosen for two reasons. First, most columns in practical use have length-to-depth ratios well above that required for the assumptions of a Euler column to be appropriate. Second, for short columns, where these assumptions are not reasonable, dynamic stability is not usually a concern.

For a Euler-Bernoulli column, the equation of motion begins by equating the sum of the applied lateral forces to the fourth derivative of the deflection multiplied by the stiffness. For the purposes of studying the dynamic stability of an axially loaded column, the two most important forces that must be considered are the inertial forces and the lateral forces induced by the axial load. In addition to these, damping forces will also be considered as any real column will include some damping. There are numerous other forces that can be considered along the column in special circumstances (elastic foundation forces, external loading, thermal stresses, etc.), but only the general condition where these forces are not

present will be considered here. A free-body diagram of the elastically restrained column being considered is shown in Figure 1.1.

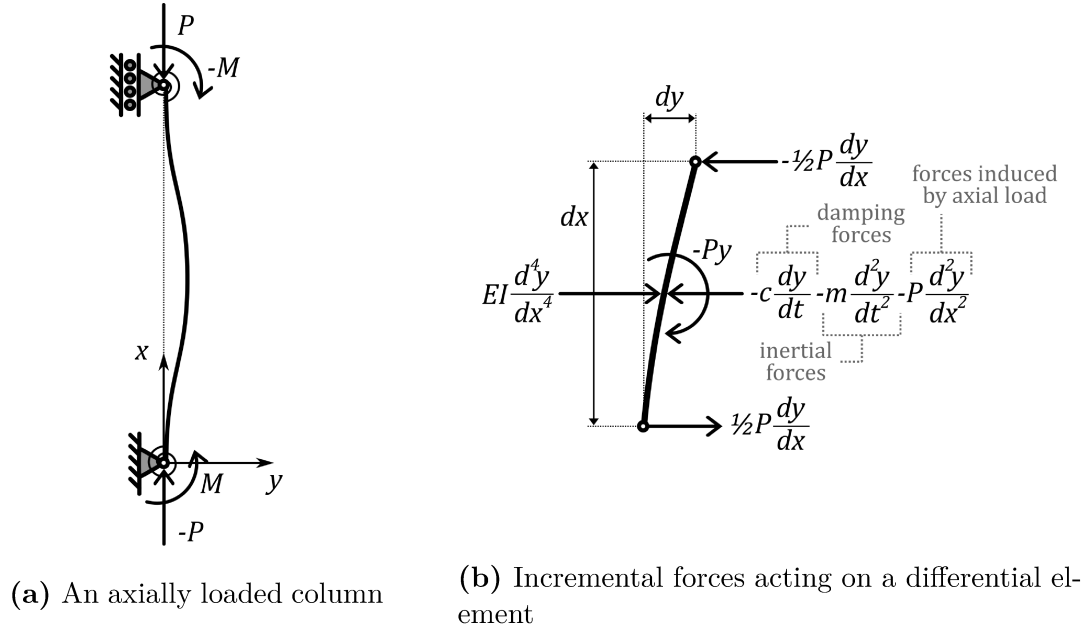


Figure 1.1: A free-body diagram of an axially excited elastically restrained column

The derivation therefore begins with,

$$\frac{\partial^4 y}{\partial x^4} EI = F_I + F_D + F_P \quad (1.1)$$

where E and I are the modulus of elasticity and area moment of inertia of the column, respectively, and y is the coordinate in the transverse direction. F_I , F_D , and F_P are the inertial forces, damping forces, and the lateral forces induced by the applied axial load, respectively.

In the interest of keeping the coordinate system as simple as possible, all measurements upwards and to the right of the column will be considered to be positive. The most important implication of this decision is that compressive forces in the column remain negative.

Lateral Forces Induced by Axial Forces When an axial load is applied to a perfectly straight column, the axial load travels along the length of the column until it is opposed at the opposite end. When any portion of the column is not perfectly straight the axial forces acting at either end do not align, inducing a negative moment proportional to the deflection

of the column at that location. Noting that the moment in a Euler-Bernoulli column is the second derivative of the deflection multiplied by the stiffness of the column, this can be written as,

$$EI \frac{\partial^2 y}{\partial x^2} = -Py \quad (1.2)$$

where P is the axial load applied to the column.

Taking the derivative of this relationship two more times gives the induced lateral forces resulting from the applied axial load:

$$F_P = -P \frac{\partial^2 y}{\partial x^2} \quad (1.3)$$

Lateral Forces From Internal Friction The damping forces from internal friction are assumed to be proportional to the velocity of the column in a direction opposite the velocity. As the coefficient of internal friction (c) is a material property, the magnitude of the internal friction should be constant along the length of the column, giving,

$$F_D = -c \frac{\partial y}{\partial t} \quad (1.4)$$

In practice the behavior of internal friction is somewhat more complex. Several sources[1][2] suggest that the behavior of the internal friction is a function of the internal stress, and therefore some function of y , which would make this simplification a potential source of significant error. As a detailed investigation into the physical mechanics of internal friction is beyond the scope of this work, and given the relatively limited impact of internal friction on the stability boundaries, the assumption of a constant value for internal friction is generally considered to be an acceptable abstraction. This assumption is discussed in greater detail in chapter 5, and the resulting error is found to be acceptably small.

Inertial Forces By considering inertial forces, a static problem can be transformed into a dynamic problem. The inertial forces are simply a restating of Newton's second law, quantifying the acceleration caused by unbalanced forces. Because the inertial forces are the forces required to cancel the forces causing acceleration, they act opposite the direction of acceleration, or in other words,

$$F_I = -m \frac{\partial^2 y}{\partial t^2} \quad (1.5)$$

where m is the mass per unit length of the column.

Assembling the Equation of Motion Substituting equations 1.3 through 1.5 into equation 1.1 gives the equation of motion:

$$EI \frac{\partial^4 y}{\partial x^4} + P \frac{\partial^2 y}{\partial x^2} + c \frac{\partial y}{\partial t} + m \frac{\partial^2 y}{\partial t^2} = 0 \quad (1.6)$$

Chapter 2 goes over the analytical solution to this equation, while Chapter 3 presents a novel numerical solution.

1.3 Solving Mathieu Equations

Because the equation of motion (equation 1.6) is a partial differential equation (PDE), it is not particularly convenient to work with. Beyond the usual inconveniences associated with PDEs, this function belongs to a class of functions called Mathieu-Hill equations, which cannot be solved in the usual sense of the term. A number of methods have been developed to work around this, each with their own particular advantages and limitations.

1.3.1 The Analytical Method

The analytical method used to solve a Mathieu equation is referred to as Bolotin's method, and is covered in greater detail in Chapter 2. This method provides all the usual benefits of an analytical solution, but is problematic in that it is difficult to obtain, limited to a sinusoidal excitation, and gives only an approximate solution. It is important to stress that Bolotin's method is limited in accuracy based on the size of the matrices used to compute the solution, and that it is not therefore a true solution, only an approximation of the true solution.

A number of variations on Bolotin's method exist, but are generally either less flexible or much more complex than the original. For example, the method of averaging can, through a number of assumptions, be used to obtain the variation in amplitude more simply, but only obtains a first-order approximation of the first stability region. On the other end of the spectrum, Bolotin's method can be performed on probability distributions, which is more helpful, but also considerably more complex than the original. The various analytical solutions that can be used as an alternative to Bolotin's method are outside the scope of this thesis, as it focuses primarily on the numerical and experimental solutions.

1.3.1.1 Using Bolotin's Method with Non-Sinusoidal Excitations

In most cases, adaptations of Bolotin's method continue to assume a sinusoidal excitation, which is one of the more limiting assumptions in Bolotin's method. This is particularly limiting in that the stability behavior is directly connected to the excitation frequency, and sinusoids have the distinction of lacking the harmonics present in all other excitations. This disregards the effect of the harmonics, which have the potential to be just as significant as the central excitation frequency.

Bolotin's method assumes an excitation of the form $P(t) = P_s + P_d \cos(\nu t)$, where P_s and P_d are, respectively, the static and dynamic components of the excitation function. If the excitation function is instead assumed to be comprised of a series of sinusoids, Bolotin's method can be extended to the Fourier transform of any arbitrary excitation function,

$$P(t) = P_s + \sum_{n=1}^N P_n \cos(n\nu t) \quad (1.7)$$

where N is the number of sinusoidal terms, and P_n is the dynamic excitation associated with each multiple of the fundamental excitation frequency.

For simplicity, consider an undamped column ($c = 0$). When $N = 1$ this simplifies to a typical application of Bolotin's method, and the matrices shown in section 2.2.3 become tridiagonal. When $N = 2$, however, following Bolotin's method instead produces matrices with five diagonals. Each increment of N adds an additional two diagonals, resulting in matrices with $2N + 1$ diagonals. The 1st-order approximation of the stability region boundaries would therefore be determined from the central $2n \times 2n$ (odd) and $(2n + 1) \times (2n + 1)$ (even) submatrices. In the undamped case, the corresponding half-matrices can be used to reduce the complexity of these matrices, but as they are not truly symmetric this must be done with care.

The primary flaw with this approach is that while this gives an analytical solution to any function, it only does so for a matrix of infinite dimensions. In other words, this approach gives a method for obtaining an approximation of the stability boundaries. As the approximation is only accurate when the errors introduced by truncating the matrix are small relative to the size of the matrix, it is necessary to determine the size of the matrix required to obtain a satisfactory approximation. The accuracy of the analytical method can be determined by comparing successive orders of approximation. The accuracy of the Fourier series cannot be determined as it depends on the specific function being approximated, but

the computational effort required to include a specific number of terms with comparable accuracy can be estimated.

When $N = 1$ (the tridiagonal matrix obtained in Chapter 2), the 2nd- or 3rd-order approximations are usually required to establish the 1st stability region with acceptable accuracy, depending on the application. For each additional stability region to be predicted, the order of the approximation required increases by one, and the dimensions of each matrix increases by two.

As the number of diagonals in the matrix increase, so too does the size of the matrix required to find a solution of acceptable accuracy. Table 1.1 shows the size of the central submatrices required to obtain acceptable levels of accuracy for various stability regions. In general, these sizes should increase with the number of diagonals in the matrix, but the exact nature of the relationship is opaque enough that a much more thorough investigation would be required than is possible here.

Table 1.1: Matrix sizes required to obtain acceptable accuracy ($n = 1$)

Region (R)	Required Dimension
1	4 - 6
2	7 - 9
3	8 - 10
4	11 - 13
...	...
R_{odd}	$2R + 3 \pm 1$
R_{even}	$2R + 4 \pm 1$

If a proportional relationship is assumed, a rough estimate can be made of the effort required to find a solution by scaling these sizes. Because calculating the determinant of a matrix is $O(n^{2.373})$ [3], a computational limit is quickly reached as the dimensions of the matrix increase. Additionally, because the determinant of the matrix must be solved symbolically to determine the boundaries (as opposed to determining the behavior at a specific point, which would allow the determinant to be solved numerically), each operation is relatively time-consuming. For example, some basic testing suggests that it is tedious to calculate the stability boundaries for matrices larger than approximately 9×9 , with the test machine taking just under two minutes to calculate one fifth-order (using 10×10 and 11×11 matrices) stability diagram. The significant majority of this time was attributable to calculating the determinant, and specifically to allocating memory (slightly more than 33GB) for the

symbolic computations. Note that these tests represent only a single machine and a single implementation, and so should be interpreted with a grain of salt, as code optimization and suitable multithreading should easily increase performance by an order of magnitude. Despite this limitation, these results remain instructive because even a $10\times$ improvement in processing capacity would still make anything over perhaps 15×15 extremely inconvenient to calculate, if for no other reason than the time required to allocate the memory. For a proportional relationship between the matrix dimensions and the number of diagonals, this would allow the first stability region to be determined for the first four terms of a Fourier series, and the second to be determined for the first two or three terms of a Fourier series, and the third stability region to be determined for the first two terms of a Fourier series.

Considering the large number of terms required to obtain a reasonable approximation using a Fourier series compared to the limited number of terms available, this approach is expected to have minimal utility.

1.3.2 Numerical Methods

Numerical solutions to the problem of dynamic stability are normally determined by calculating the next state based on the governing equations and the current state. Regardless of the specific method chosen, both the accuracy and cost of the numerical method are primarily dictated by the step size. As the step size decreases, more steps are required to calculate the response, but the error in each step is reduced. This increases the computational cost, but decreases the error, resulting in a higher accuracy.

The simplest of the numerical methods is the constant-slope method, in which the slope is extrapolated to obtain the next value. The upper limit to the complexity of numerical methods is based on balancing the computational cost so that the total error is reduced by the maximum amount. If the method is too simple the error in each step will govern, leading to long computation times on account of the large number of steps required. If the method is too complex, discretization errors will govern, leading to long computation times attributable to the time required to calculate each step. The fourth-order Runge-Kutta method is often a reasonable compromise between these two extremes, and thus represents a good example of a relatively complex numerical method [4].

A more thorough evaluation of existing numerical methods can be obtained elsewhere, but the following examples, summarized from Xie [4], are provided to illustrate the effect of varying the complexity on both the computational effort required and the accuracy obtained.

1.3.2.1 The Constant-Slope Method

The constant-slope method simply assumes that the slope remains constant throughout each step, and increments the value of the function in a linear fashion. For a function F_1 in which the derivative can be solved for, the constant slope method simply extrapolates the current position based on the current slope. If F_1 is a function of time (t) with a step width of t_Δ , this method is defined by the equation,

$$F_1(t + t_\Delta) = F_1(t) + F_1'(t)t_\Delta \quad (1.8)$$

The cumulative error in each iteration is proportional to $1/N[4]$, where N is the number of segments into which the period of the analysis is subdivided. As a consequence of how large this error is, the use of this method is generally not advisable.

1.3.2.2 The Fourth-Order Runge-Kutta Method

The fourth-order Runge-Kutta method is more complex than the constant-slope method in that it has more steps:

1. Calculate the slope at the start of the step (s_i) from the position at the start of the step ($F_1(t)$)
2. Extrapolate the value at the midpoint of the step (v_{m1}) using s_i
3. Calculate the slope at the midpoint of the step (s_{m1}) using v_{m1}
4. Re-estimate the value at the midpoint (v_{m2}) using s_{m1}
5. Recalculate the slope at the midpoint (s_{m2}) of the step using v_{m2}
6. Extrapolate the value at the end of the step (v_f) using s_{m2}
7. Calculate the slope at the end of the step (s_f) using v_f
8. Use a weighted average of the slope estimates to calculate the value at the end of the step using the equation,

$$F_1(t + t_\Delta) = F_1(t) + \frac{s_i + 2s_{m1} + 2s_{m2} + s_f}{6}t_\Delta \quad (1.9)$$

where F_1 and t_Δ are the same as those used in the constant-slope method.

This method produces considerably less error than the constant-slope method (of the order $O(\tau^4)$ or proportional to $1/N^4$)[4], but requires noticeably more calculations to be performed. The order of the error in each method only describes the proportionality of the error, but it is reasonable to assume that the proportionality coefficients are fairly similar when the function being approximated is relatively straight over a single time-step. Assuming that this is true, the fourth-order Runge-Kutta method at $N = 4$ produces results equivalent to those of the constant-slope method when $N = 256$, which require $16\times$ the computational effort. As N increases this disparity becomes even more pronounced, with the constant-slope method requiring $250\times$ the effort when $N = 10$ and $2000\times$ when $N = 20$.

1.3.2.3 Limitations of Current Numerical Methods

These methods, as with many other numerical methods not described in detail here, require considerable computational effort and produce the response associated with the selected initial conditions. While the fact that computational effort is required is common to all numerical methods and cannot be avoided, the fact that these methods determine the response rather than the stability, and the inability to generalize the results are both inconvenient. The fact that the numerical methods do not determine the stability directly is not overwhelmingly problematic, as the stability can be determined from the resulting state by comparing the final state to the initial state after one period. The lack of generalizability is more problematic in that these methods reveal no additional information that can be used to determine the stability outside of the specific conditions calculated. The numerical method developed in Chapter 3 addresses both of these issues by allowing the stability to be calculated directly for multiple initial conditions.

1.4 Semi-Rigid Connections

In all real connections, each degree of freedom in a support has some positive, finite stiffness. This differs from most structural analyses, where connections are commonly idealized by assigning each degree of freedom as being either fixed (meaning that no amount of force is sufficient to cause a deflection) or free (meaning that no amount of deflection is sufficient to cause a reactionary force). The most common idealizations are the pin-pin or fixed-fixed conditions, shown in Figure 1.2.

These conditions are very good approximations in some circumstances, but many connections have intermediate fixity and fall somewhere in between these two extremes. In this




Figure 1.2: Support conditions for dynamic stability of columns. (1) Pin-pin (2) Fixed-fixed (3) Semi-rigid or elastically restrained

case, the connection is referred to as a semi-rigid connection (also shown in Figure 1.2) or is described as being elastically restrained. If such a connection is analyzed with a pinned or fixed support, the analytical solution will be based on an inaccurate assumption. As the quality of any analytical solution is only as good as its assumptions, the results produced by any such analysis will likewise be inaccurate. Relatively little work has been done in the area of dynamic stability relating to semi-rigid connections, as semi-rigid connections have significantly higher complexity than their idealized counterparts.

There seems to be a consensus in favor of using the term semi-rigid over describing a connection as being elastically restrained. While the brevity of the term is no doubt an influencing factor, it is also likely preferred due to its generality. An elastically restrained connection is a specific type of rigid connection in the sense that it exhibits purely elastic behavior. Like with the idealized support conditions, this is an idealization of a real connection in that it ignores the realities of force mobilization and plastic behavior that occur in a real connection. While this is still not perfectly accurate, the degree of approximation

associated with this idealization is much smaller than that used in the pin, roller, and fixed connections.

1.5 Experiments on the Dynamic Stability of Columns

The design of experiments on the dynamic stability of columns can be broadly separated into three categories, depending on the nature of the loads and supports being considered. The first category is when considering impact loading, and is usually conducted by either dropping large weights onto a specimen[5][6][7] or firing specimens at steel plates[8][9]. The second category is when considering a cantilevered column subject to tangential or follower forces. These are conducted either using a shake table[10] or by attaching a rocket engine to the end of the column[11], respectively. The third category is for columns supported at both ends. As the proposed design belongs solidly in the third category, only the third category will be considered in further detail.

The most fundamental experimental setup for assessing the dynamic stability of columns consists of a rigid frame pinning a test specimen in series with a load cell and shaker motor [12]. The shaker motor applies a dynamic load to the specimen, and the response of the specimen is measured to determine its stability. The simplicity of this arrangement makes it easy to adapt to other studies by adding additional features to the experimental design. Iwatsubo[13] uses the same setup but replaces one or both connections with fixed connections. Ahuja[14] includes an elastic foundation and a specimen with varying cross-section. Svensson[15] includes a transverse support preventing buckling in the first mode.

Where only the specimen differs from previous studies, the same experimental setup can be reused with different specimens to study different effects. This can be observed at the National Institute of Technology Rourkela in India, where the same apparatus was used to study viscoelastic damping in sandwich beams[16], the effects of localized damage[17], and slip in bolted connections[18] over three years.

While pinned and fixed connections have been studied on several occasions, there does not appear to be any experimental record including elastic connections of any kind. This is no doubt due to the inconvenience of creating an experimental apparatus capable of including elastic connections of any kind. Beyond simply the inconvenience of developing such an apparatus, there appears to be no attempt made to quantify the fixity of any of the connections used in the experimental record. While the bearings commonly used to represent pinned connections are likely quite adequate for that purpose, the clamped connections are

usually presented as such without any documented effort to confirm that the connections are fixed.

1.6 Research Objectives and Thesis Outline

There are two main objectives of this research.

The first objective was to develop the numerical matrix solution to the problem of an elastically restrained axially excited column and determine its accuracy relative to the analytical solution. Determining the accuracy of the numerical solution also provided an excellent opportunity to estimate the error in the analytical solution, which was previously unquantified.

The second objective was to experimentally verify the analytical and numerical solutions, and to determine the variability associated with the experimental solution. Fulfilling this objective required the development of a suitable experimental apparatus, capable of modeling a range of elastic connections and capable of generating large volumes of data. This apparatus was successfully developed and validated, but due to a combination of equipment failure and supply chain challenges the variability in the experimental results could not be measured.

The organization of this thesis follows from these objectives, with each chapter dedicated to a single step in accomplishing them. First, the analytical solution is obtained using Bolotin's method (Chapter 2). The numerical matrix solution is then developed (Chapter 3). With both the analytical and numerical solutions obtained, the accuracy of each can be determined (Chapter 4), completing the first objective. The second objective is partially fulfilled with the development and validation of the experimental apparatus, which is given its own chapter (Chapter 5).

Chapter 2

The Analytical Solution

The analytical solution to the problem of the dynamic stability of an axially excited column is known as Bolotin's method[19][20]. The process of applying Bolotin's method to the case of an axially excited column with elastic restraints begins with the equation of motion obtained in section 1.2 and follows in sections 2.1 through 2.3.

2.1 Separating the PDE

As the equation of motion, like most PDEs, is not particularly convenient to work with, it is desirable to reduce it to an ordinary differential equation (ODE) by assuming that y (a function of both x and t) can be expressed as the product of a modal or shape function (X , a function of only x) and a response function (q , a function of only t), making the equation separable. After making this substitution and rearranging slightly, the equation of motion becomes,

$$\left(EI \frac{X^{(4)}}{X} + P \frac{X''}{X} \right) q + c\dot{q} + m\ddot{q} = 0 \quad (2.1)$$

where the derivatives of X with respect to x are shown using Lagrange's notation (using primes) and the derivatives of q with respect to t are shown using Newton's notation (using dots) to help differentiate them from each other.

Because P is a periodic function of time, this equation is still not technically separated, but this assumed solution is acceptable provided that X and q can be chosen to satisfy the PDE. As a fourth-order ODE of position (x), the solution to X is based on four boundary conditions, which can be solved for based on the supports at $x = 0$ and $x = L$. As a second-order ODE of time (t), the solution to q is obtained from two initial conditions, based on the position and velocity of the column at $t = 0$. As the boundary conditions of X allow it to be solved generally, unlike the initial conditions of q , the first step to solving the PDE is

to obtain a solution for X and eliminate it from the equation, reducing the problem to an ODE.

2.1.1 The Modal Function

The modal functions traditionally used in the pin-pin and fixed-fixed cases for buckling in the first mode are, respectively,

$$X_p = \Delta_0 \sin\left(\frac{\pi}{L}x\right) \quad (2.2a)$$

$$X_f = \frac{\Delta_0}{2} \left(1 - \cos\left(\frac{2\pi}{L}x\right)\right) \quad (2.2b)$$

where Δ_0 is the maximum initial deflection and L is the length of the column.

The primary limitation of these functions is that they are restricted to integer modal values, which guarantees that the slope at the supports will be either maximum (as in the pin-pin function) or zero (as in the fixed-fixed function). This restriction prevents them from representing any shape with intermediate fixity. Secondly, these modal functions are not capable of representing end conditions with opposite fixity (the pin-fixed and fixed-pin cases).

In the case of elastically restrained connections, the shape of the column should logically fall somewhere between the pin-pin and fixed-fixed cases. A very promising function was proposed (though not experimentally verified) by Giraldo-Londono[21],

$$X = \Delta_0 \frac{\cos\left(\frac{\alpha\pi}{L}x - \frac{\alpha\pi}{2}\right) - \cos\left(\frac{\alpha\pi}{2}\right)}{1 - \cos\left(\frac{\alpha\pi}{2}\right)} \quad (2.3)$$

where α is the fixity constant which satisfies $1 \leq \alpha \leq 2$. When $\alpha = 1$, this reduces to the pin-pin case (equation 2.2a), while when $\alpha = 2$, this reduces to the fixed-fixed case (equation 2.2b). Between these two extremes the fixity assumes some intermediate value.

In essence, this modal function is based on the idea that buckling is not restricted to integer modes, but as the fixity at the supports increases the buckling mode transitions between integer values. The function proposed is primarily limited in that it is not suitable for estimating buckling beyond the first mode, and it assumes that both ends of the column exhibit identical fixity. The first limitation is acceptable because the first buckling mode is nearly always the mode of primary interest. The second is more inconvenient, but as this

limitation is common to the traditional modal functions it is not any more inconvenient than the alternatives.

2.1.1.1 Boundary Conditions

Like the traditional shape functions, the elastically restrained shape function shares the limitation that the deflection at either end must be equal to zero. The remaining two functions are obtained from the rotational stiffness of the elastic restraints (κ). The set of four boundary conditions for the elastically restrained column can therefore be expressed as,

$$X(0) = 0 \quad X(L) = 0 \quad X''(0) = \frac{-\kappa}{EI}X'(0) \quad X''(L) = \frac{\kappa}{EI}X'(L) \quad (2.4)$$

2.1.1.2 Converting the Rotational Stiffness to Fixity

Rather than solving for the fixity (α) as a function of κ , it is easier to first convert the rotational stiffness to the connection fixity, ρ , which is more easily compared with α . The connection fixity, ρ , satisfies $0 \leq \rho \leq 1$ and represents the proportion of the fixed-end moment that develops in the connection or the portion of the pinned-end rotation that is prevented by the connection,

$$\rho = EI \frac{X''}{M_f} = 1 - \frac{X'}{\theta_p} \quad (2.5)$$

where M_f is the fixed-end moment, θ_p is the pinned-end rotation, and the derivatives of X are evaluated at either $x = 0$ or $x = L$.

Since the fixed-end moment and pinned-end rotation are known to be,

$$M_f = 2EI\Delta_0 \left(\frac{\pi}{L}\right)^2 \quad \theta_p = \Delta_0 \left(\frac{\pi}{L}\right) \quad (2.6)$$

the moment-rotation line for the member can be expressed as a straight line:

$$M = 2EI\Delta_0 \left(\frac{\pi}{L}\right)^2 - 2\pi \frac{EI}{L} \theta \quad (2.7)$$

By finding the intersection of this line and that of the connection ($M = \kappa\theta$), the rotation in the connection can be obtained:

$$\theta = \frac{2EI\Delta_0 \left(\frac{\pi}{L}\right)^2}{\kappa + 2EI \left(\frac{\pi}{L}\right)} \quad (2.8)$$

Substituting this value into the equation for ρ converts this rotation to a measure of fixity:

$$\rho = \frac{\kappa}{\kappa + 2\pi \left(\frac{EI}{L}\right)} \quad (2.9)$$

At the same time, α can be converted to ρ by simply substituting the shape function into equation 2.5, giving:

$$\rho = 1 - \frac{\alpha \sin\left(\frac{\pi}{2}\alpha\right)}{1 - \cos\left(\frac{\pi}{2}\alpha\right)} \quad (2.10)$$

Eliminating ρ between these two equations leaves a function of both α and κ , but it is not conveniently solved.

If ρ is instead approximated with,

$$\tilde{\rho} = \frac{\alpha^2 - \alpha}{2} \quad (2.11)$$

then an approximate function for α can be obtained by solving the resulting quadratic equation:

$$\tilde{\rho} = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2\kappa}{\kappa + 2\pi \frac{EI}{L}}} \quad (2.12)$$

While this relationship is not perfect, it does have better than $\pm 1\%$ accuracy.

Figure 2.1 shows the relationship between ρ , $\tilde{\rho}$, and α .

2.1.1.3 Substituting the Modal Function Into the PDE

Substituting this modal function into the equation of motion (equation 1.6) reduces it to an ODE,

$$m\ddot{q} + c\dot{q} + \chi \left(\frac{\alpha\pi}{L}\right)^2 (P_E + P)q = 0 \quad (2.13)$$

where P_E is the Euler buckling load for a column with intermediate fixity α ,

$$P_E = EI \left(\frac{\pi\alpha}{L}\right)^2 \quad (2.14)$$

and χ is the shape factor derived from the modal function.

$$\chi = \frac{1}{1 - \frac{\cos\left(\frac{\alpha\pi}{2}\right)}{\cos\left(\frac{\alpha\pi}{L}x - \frac{\alpha\pi}{2}\right)}} \quad (2.15)$$




Figure 2.1: A comparison of ρ , $\tilde{\rho}$, and α

At midspan, χ simplifies to:

$$\chi = \frac{1}{1 - \cos\left(\frac{\alpha\pi}{2}\right)} \quad (2.16)$$

2.1.2 Converting to Mathieu Equation Form

Assuming that $P(t)$ is a sinusoidal excitation with constant offset, it is easily expressed in terms of the static (P_s) and dynamic (P_d) excitations. Assuming that these excitations are compressive (negative) this becomes,

$$m\ddot{q} + c\dot{q} + \chi\left(\frac{\alpha\pi}{L}\right)^2 (P_E - P_s - P_d \cos(\nu t))q = 0 \quad (2.17)$$

After rearranging, this becomes,

$$\ddot{q} + \frac{c}{m}\dot{q} + \frac{\chi}{m}\left(\frac{\alpha\pi}{L}\right)^2 (P_E - P_s) \left(1 - \frac{P_d \cos(\nu t)}{P_E - P_s}\right)q = 0 \quad (2.18)$$

Once in this form, the standard Mathieu equation form can be easily obtained by solving

for the particular μ and ω associated with the case of an elastically restrained column. Doing so gives the equation of motion as a Mathieu equation in standard form,

$$\ddot{q} + \frac{c}{m}\dot{q} + \omega^2(1 - 2\mu \cos(\nu t))q = 0 \quad (2.19)$$

where μ is the loading parameter,

$$\mu = \frac{P_d}{2(P_E - P_s)} \quad (2.20)$$

and ω is the natural frequency of the column:

$$\omega = \sqrt{\frac{\chi}{m} \left(\frac{\alpha\pi}{L}\right)^2 (P_E - P_s)} \quad (2.21)$$

Noting that the natural frequency of the column is affected by the applied load, it is convenient to express this in terms of the natural frequency of the unloaded column, ω_0 , which can also be expressed as a function of the Euler buckling load. As only the static load should affect this value, the natural frequency of the unloaded column can be obtained factoring out P_E and isolating P_s .

$$\omega = \omega_0 \sqrt{1 - \frac{P_s}{P_E}} \quad \omega_0 = \sqrt{\frac{\chi}{m} \left(\frac{\alpha\pi}{L}\right)^2 P_E} \quad (2.22)$$

2.1.2.1 Comparing the Elastically Restrained Case to the Pinned Case

As only ω and P_E are functions of α , the solution for any α is the same as for a column that is not elastically restrained, and is obtained in the usual method. While the stability diagrams obtained in this fashion are identical to those obtained for the pin-pin case, it is important to note that the frequencies and excitations associated with a particular ω and μ are not the same as for the pin-pin case, as both of these variables are dependent on α . Figure 2.2 shows how the loading and frequency parameters vary with the fixity.

As the fixity increases, a given dynamic load results in a smaller excitation parameter (decreasing by a factor of 4 at full fixity), while the natural frequency of the column increases (by up to a factor of $2\sqrt{2}$ at full fixity).

As this solution reduces to the pin-pin and fixed-fixed cases when $\alpha = 1$ and $\alpha = 2$, respectively, the pin-pin and fixed-fixed solutions represent a subset of the elastically restrained solution. Therefore, while the solution procedure from this point onward is the




Figure 2.2: Variation of ω and μ with fixity (α)

same, it is important to note that it is also more generally applicable due to the consideration of partial fixity in the parameters used.

2.2 Finding a Solution

While the ODE is much more convenient to work with than the PDE, it is worth noting that Mathieu equations can only be solved approximately.

2.2.1 Floquet's Theory

The solution to the ODE depends on Floquet's theory, which states[22] that if F_1 is a function of \mathbf{x} , F_2 is a periodic function of \mathbf{x} , and,

$$F_1'(\mathbf{x}) = F_2(\mathbf{x})F_1(\mathbf{x}) \quad (2.23)$$

then F_1 is periodic and has the form,

$$F_1 = C_1 e^{C_2 \mathbf{x}} F_3(\mathbf{x}) \quad (2.24)$$

where F_3 is a periodic function of \mathbf{x} that has the same period as F_2 , and C_1 and C_2 are constants.

Therefore, the response of the column will be periodic with the same frequency as the excitation. Additionally, because the Mathieu equation is a second-order differential equation, there should be two linearly independent solutions.

2.2.2 Transformation to Matrix Form

Knowing that the response has the same frequency as the excitation and that any function can be expressed as an infinite series of sinusoids, the solution to the ODE can be expressed in the form,

$$q = \sum_{n=0}^{\infty} \left(A_n \sin\left(\frac{n\nu t}{2}\right) + B_n \cos\left(\frac{n\nu t}{2}\right) \right) \quad (2.25)$$

where A_n and B_n are unknown constants unique to each term of the summation.

By substituting the assumed solution into the Mathieu equation (equation 2.19), every term can be moved inside the summation, giving,

$$\sum_{n=0}^{\infty} \left(\ddot{q}_n + \frac{c}{m} \dot{q}_n + (1 - 2\mu \cos(\nu t)) q_n \right) = 0 \quad (2.26)$$

where q_n is the n^{th} term of the assumed solution.

The derivative of each term can now be taken, resulting in a purely algebraic equation. It is also convenient at this point to simplify the equation by dividing by ω^2 :

$$0 = \sum_{n=0}^{\infty} \left(\begin{aligned} & \left[A_n \left(1 - \frac{n^2 \nu^2}{4\omega^2} \right) - B_n \frac{n c \nu}{2m\omega^2} \right] \sin\left(\frac{n\nu t}{2}\right) + \\ & \left[B_n \left(1 - \frac{n^2 \nu^2}{4\omega^2} \right) - A_n \frac{n c \nu}{2m\omega^2} \right] \cos\left(\frac{n\nu t}{2}\right) - \\ & 2\mu \left[A_n \cos(\nu t) \sin\left(\frac{n\nu t}{2}\right) + B_n \cos(\nu t) \cos\left(\frac{n\nu t}{2}\right) \right] \end{aligned} \right) \quad (2.27)$$

Applying the product-to-sum trigonometric identities,

$$2 \cos(\alpha) \cos(\beta) = \cos(\alpha + \beta) + \cos(\alpha - \beta) \quad (2.28a)$$

$$2 \cos(\alpha) \sin(\beta) = \sin(\alpha + \beta) - \sin(\alpha - \beta) \quad (2.28b)$$

to the third term of equation 2.27 eliminates the product terms, reducing each term in the summation to a linear combination of sines and cosines:

$$0 = \sum_{n=0}^{\infty} \left(\begin{array}{l} \left[A_n \left(1 - \frac{n^2 \nu^2}{4\omega^2} \right) - B_n \frac{n c \nu}{2m\omega^2} \right] \sin \left(\frac{n \nu t}{2} \right) + \\ \left[B_n \left(1 - \frac{n^2 \nu^2}{4\omega^2} \right) - A_n \frac{n c \nu}{2m\omega^2} \right] \cos \left(\frac{n \nu t}{2} \right) - \\ A_n \mu \left[\sin \left(\frac{(n+2)\nu t}{2} \right) + \sin \left(\frac{(n-2)\nu t}{2} \right) \right] + \\ B_n \mu \left[\cos \left(\frac{(n+2)\nu t}{2} \right) + \cos \left(\frac{(n-2)\nu t}{2} \right) \right] \end{array} \right) \quad (2.29)$$

Because sine and cosine functions are linearly independent, and therefore must be independently equal to zero to satisfy equation 2.29, equation 2.29 can be split into a system of two equations:

$$0 = \sum_{n=0}^{\infty} \left(\begin{array}{l} \left[B_n \left(1 - \frac{n^2 \nu^2}{4\omega^2} \right) - A_n \frac{n c \nu}{2m\omega^2} \right] \cos \left(\frac{n \nu t}{2} \right) - \\ B_n \mu \left[\cos \left(\frac{(n+2)\nu t}{2} \right) + \cos \left(\frac{(n-2)\nu t}{2} \right) \right] \end{array} \right) \quad (2.30a)$$

$$0 = \sum_{n=0}^{\infty} \left(\begin{array}{l} \left[A_n \left(1 - \frac{n^2 \nu^2}{4\omega^2} \right) - B_n \frac{n c \nu}{2m\omega^2} \right] \sin \left(\frac{n \nu t}{2} \right) - \\ A_n \mu \left[\sin \left(\frac{(n+2)\nu t}{2} \right) + \sin \left(\frac{(n-2)\nu t}{2} \right) \right] \end{array} \right) \quad (2.30b)$$

Because each frequency of sines and cosines is also linearly independent from all others, each of these two equations can be replaced with an infinite system of equations by collecting sines and cosines between terms to isolate each frequency in its own equation. Note that as terms near $n = 0$ include negative frequencies (which are not linearly independent from their positive-frequency counterparts), these frequencies must be converted to their positive-frequency counterparts before separating the frequencies into the system of equations. The sine can cosine terms can then be eliminated, giving a system of algebraic equations.

Equations 2.30a and 2.30b then become,

$$0 = \begin{cases} B_0 + B_2\mu & n = 0 \\ B_1 \left(1 - \mu - \frac{\nu^2}{4\omega^2}\right) - A_1 \frac{c\nu}{2m\omega^2} & n = 1 \\ B_2 \left(1 - \frac{4\nu^2}{4\omega^2}\right) - A_2 \frac{2c\nu}{2m\omega^2} + 2B_0\mu & n = 2 \\ B_n \left(1 - \frac{n^2\nu^2}{4\omega^2}\right) - (B_{n-2} + B_{n+2})\mu - A_n \frac{nc\nu}{2m\omega^2} & n \geq 3 \end{cases} \quad (2.31a)$$

$$0 = \begin{cases} A_1 \left(1 + \mu - \frac{\nu^2}{4\omega^2}\right) - B_1 \frac{c\nu}{2m\omega^2} & n = 1 \\ A_2 \left(1 - \frac{4\nu^2}{4\omega^2}\right) - B_2 \frac{2c\nu}{2m\omega^2} & n = 2 \\ A_n \left(1 - \frac{n^2\nu^2}{4\omega^2}\right) - (A_{n-2} + A_{n+2})\mu - B_n \frac{nc\nu}{2m\omega^2} & n \geq 3 \end{cases} \quad (2.31b)$$

At this point, the systems of equations can be significantly reduced in complexity by performing two substitutions. First, the excitation frequency can be expressed relative to the critical dynamic stability frequency:

$$\phi = \frac{\nu}{2\omega} \quad (2.32)$$

Second, the damping can be expressed in terms of a damping coefficient, (β) , instead of the internal friction coefficient:

$$\beta = \frac{c}{m\omega} \quad (2.33)$$

Equations 2.31a and 2.31b can then be put into matrix form. Because the resulting matrix is infinite, it is impossible to represent in its entirety, but as the outer part of the matrix follows a consistent pattern, only a submatrix must be depicted to allow the construction of

an arbitrarily large matrix. The central 7×7 submatrix is:

$$\begin{bmatrix} 1-9\phi^2 & . & -\mu & . & . & . & -3\beta\phi \\ . & 1-4\phi^2 & . & 0 & . & -2\beta\phi & . \\ -\mu & . & 1+\mu-\phi^2 & . & -\beta\phi & . & . \\ . & 0 & . & 1 & . & -\mu & . \\ . & . & \beta\phi & . & 1-\mu-\phi^2 & . & -\mu \\ . & 2\beta\phi & . & -2\mu & . & 1-4\phi^2 & . \\ 3\beta\phi & . & . & . & -\mu & . & 1-9\phi^2 \end{bmatrix} \begin{bmatrix} A_3 \\ A_2 \\ A_1 \\ B_0 \\ B_1 \\ B_2 \\ B_3 \end{bmatrix} = 0 \quad (2.34)$$

2.2.3 Separating the Matrix

Before using this matrix, it is helpful to note that the odd-numbered rows and columns are completely independent of the even-numbered rows and columns. This is not unexpected, as Floquet's theory suggests that there should be two linearly independent solutions. Separating the matrix into these two independent matrices makes it much more convenient to work with. The 4×4 center of the odd matrix (for A_{odd} and B_{odd}) and the 5×5 center of the even matrix (for A_{even} and B_{even}) are, respectively,

$$\begin{bmatrix} 1-9\phi^2 & -\mu & . & -3\beta\phi \\ -\mu & 1+\mu-\phi^2 & -\beta\phi & . \\ . & \beta\phi & 1-\mu-\phi^2 & -\mu \\ 3\beta\phi & . & -\mu & 1-9\phi^2 \end{bmatrix} \begin{bmatrix} A_3 \\ A_1 \\ B_1 \\ B_3 \end{bmatrix} = 0 \quad (2.35a)$$

$$\begin{bmatrix} 1-16\phi^2 & -\mu & . & . & -4\beta\phi \\ -\mu & 1-4\phi^2 & 0 & -2\beta\phi & . \\ . & 0 & 1 & -\mu & . \\ . & 2\beta\phi & -2\mu & 1-4\phi^2 & -\mu \\ 4\beta\phi & . & . & -\mu & 1-16\phi^2 \end{bmatrix} \begin{bmatrix} A_4 \\ A_2 \\ B_0 \\ B_2 \\ B_4 \end{bmatrix} = 0 \quad (2.35b)$$

2.2.4 Finding the Stability Boundaries

For the matrices 2.35a and 2.35b to have a non-trivial solution, their determinant must be equal to zero. As a general rule, this is not particularly convenient to do except in extreme cases, such as where μ is very small and $\beta = 0$. To illustrate this, consider the matrices shown so far. These matrices give a second-order approximation of the stability boundaries

by solving the following equations:

$$0 = 81\phi^8 + (90\beta^2 - 180)\phi^6 + (9\beta^4 - 36\beta^2 - 99\mu^2 + 118)\phi^4 + ((-3\mu^2 + 10)\beta^2 + 38\mu^2 - 20)\phi^2 + \mu^4 - 3\mu^2 + 1 \quad (2.36a)$$

$$0 = 4096\phi^8 + (1280\beta^2 + 2048\mu^2 - 2560)\phi^6 + (64\beta^4 + (128\mu^2 - 256)\beta^2 - 896\mu^2 + 528)\phi^4 + ((-16\mu^2 + 20)\beta^2 - 32\mu^4 + 112\mu^2 - 40)\phi^2 + 3\mu^4 - 4\mu^2 + 1 \quad (2.36b)$$

Neither of these equations is particularly convenient to work with, and the second-order approximation of the stability boundary conditions isn't even particularly accurate. The plots for equations 2.36a and 2.36b are shown in Figure 2.3. Comparable curves produced by the fifth-order approximation are shown in Figure 2.4.




Figure 2.3: Normalized second-order approximation of the stability region boundaries

As higher-order boundaries are considered the complexity increases rapidly, quickly ruling out any realistic attempt at an algebraic solution. For this reason, it is usually more effective to use numerical methods to find the roots of the equations produced by Bolotin's method.




Figure 2.4: Normalized fifth-order approximation of the stability region boundaries

2.3 Restating the Stability Boundaries

The stability behavior (for a given set of initial and boundary conditions) is dependent on the natural frequency and damping of the column, the applied static and dynamic loads, and the excitation frequency. This gives five possible variables to represent in a stability diagram, which is considerably more than the $2^{1/2}$ axes it is usually practical to present in a two-dimensional figure, necessitating the use of either a transformation of variables (as used in the derivation and its traditional presentation) or a series of diagrams. The transformation of these variables into a loading parameter and relative frequency parameter makes for a convenient presentation (doubly so when the derivation has already done the transformation), but makes it more difficult to use for design purposes. The best example of this is how a change in the static excitation requires both the x-axis and y-axis entry points to be reconsidered.

2.3.1 Changing the Variables

This challenge can be addressed through a change of variables, at the small expense of some algebraic manipulation and the use of additional plots. Because the difficulty of obtaining the stability boundaries is independent of the number of variables considered, it is most convenient to obtain the stability diagrams in the usual manner, and transform the result using the changes of variables detailed here. In this manner, the computationally expensive act of obtaining the stability boundaries need only be completed once. As the static excitation affects both axes, it will be treated as an independent variable in both transformations.

2.3.1.1 Expressing Frequencies Relative to the Natural Frequency of the Unloaded Column

The relative frequency variable (ϕ) is inconvenient in that it expresses the stability relative to twice the natural frequency of the column after adjusting for the effects of preload. This value is, to say the least, not likely to be foremost in the mind of a design engineer in the same way the natural frequency of the unloaded column is, introducing both a greater potential for error and another layer of abstraction separating the physical behavior from the mathematical model. This can be addressed by substituting equation 2.22 into equation 2.32 and rearranging to obtain,

$$\hat{\nu} = \frac{\nu}{\omega_0} = 2\sqrt{1 - \hat{P}_s} \phi \quad (2.37)$$

where $\hat{\nu}$ is the relative excitation frequency, and \hat{P}_s is the relative static excitation:

$$\hat{P}_s = \frac{P_s}{P_E} \quad (2.38)$$

For any relative static excitation, the frequency variables still occupy only a single axis, and the change in the central frequency is not hidden behind a transformation.

2.3.1.2 Expressing the Excitation in Terms of the Buckling Load

Because the loading parameter (μ) changes any time the static or dynamic excitations change, it is inconvenient to use as a measure of the excitation. By stating the static and dynamic load in terms of the buckling load of the column section this issue can be significantly mitigated. This can be done by rearranging equation 2.20 to obtain the relative dynamic

excitation (\hat{P}_d):

$$\hat{P}_d = \frac{P_d}{P_E} = 2\mu \left(1 - \hat{P}_s\right) \quad (2.39)$$

Again, for any given value of relative static excitation, the dynamic excitation requires only a single axis, and the change in the scaling of the dynamic excitation is not hidden behind a transformation.

2.3.2 Re-plotting the Solutions

While the relative frequency and dynamic excitation are the most obvious candidates for being presented on the x- and y-axes, the damping coefficient and static excitation are more flexible. As the readability of a plot decreases significantly when more than one series of lines are presented, one of them will be represented by a series of lines on each plot, and the other will be represented by using a series of plots. As the static excitation causes a horizontal shift in (and therefore overlap between) the stability regions, it is most convenient to plot multiple damping coefficients together, with separate plots for each value of relative static excitation.

Figure 2.5 is therefore the transformed equivalent to Figure 2.4. Note that the relative dynamic excitation is shown only up to 90% of the static buckling load, to allow for easy comparison between these figures and those determined numerically. Figures 2.6 and 2.7 shows the resulting diagrams for various values of relative static excitation between -0.5 and 0.5 .

Two effects are clear from these figures. First, an increase in the static axial load causes the central frequency of each region of dynamic instability to decrease. Second, as the static axial load increases, the dynamic excitation required to cause instability at a given frequency decreases.

2.4 Summary

Beginning with the equation of motion obtained earlier, Bolotin's method was used to obtain analytically the stability boundaries of an elastically restrained column subject to a sinusoidal excitation. This solution was in the form of an algebraic equation obtained by setting the determinant of an infinite matrix to zero, and was solved numerically to obtain the stability boundaries. The fixity term, α , only appears in this solution as part of the frequency term, ω , resulting in the same stability diagram as is obtained for the pin-pin and fixed-fixed columns






Figure 2.5: Transformed fifth-order approximation of the stability region boundaries ($\hat{P}_s = 0$)

except that the frequency associated with any particular point on the diagram varies with α . The solution was then further manipulated so that it could be expressed in terms of the excitation relative to the natural frequency of the unloaded column ($\hat{\nu}$), the peak axial excitation relative to the Euler buckling load (\hat{P}), and the damping coefficient (β). The effect of varying the static load acting on the column was then shown.

(a) $\hat{P}_s = -0.2$ (b) $\hat{P}_s = 0.2$ **Figure 2.6:** Transformed fifth-order approximation of the stability region boundaries, with relative static excitation $\hat{P}_s = \pm 0.2$



(a) $\hat{P}_s = -0.5$ (b) $\hat{P}_s = 0.5$

Figure 2.7: Transformed fifth-order approximation of the stability region boundaries, with relative static excitation $\hat{P}_s = \pm 0.5$

Chapter 3

The Numerical Matrix Solution

Consider the more general form of the Mathieu-Hill equation obtained previously, in which ζ represents an arbitrary periodic excitation:

$$\ddot{q} + 2\beta\dot{q} + \omega^2\zeta(t)q = 0 \tag{3.1}$$

In the case where the excitation function (ζ) is sinusoidal (for example, $\zeta = 1 - 2\mu \cos(\nu t)$), equation 3.1 is a Mathieu equation, and the solution can be obtained analytically through Bolotin's method as outlined previously. Where the excitation function is not sinusoidal, equation 3.1 is a Hill equation, and other means must be used to find the solution. Rather than developing unique methods for determining the solution to every possible type of excitation function, it is convenient to be able to obtain solutions using a single method for all functions.

3.1 The Numerical Method

As there are many cases where analytical solutions are impractical, it is often more convenient to use numerical methods. These have the advantage of being applicable to any function and require very little algebraic manipulation relative to analytical methods. Furthermore, they can provide relatively high accuracy with relatively low computational effort when designed carefully.

Two possible methods by Deng [23] appear in the literature. The more accurate of these two is performed by breaking the excitation into a series of line segments. The tradeoff for this accuracy is that the results of each computation are a series of Bessel functions, which are not particularly convenient to work with. The less accurate of the two is performed by discretizing the excitation function into a step function. This has the advantage of being relatively easy to compute and work with. This should not be understood to mean that the first method is good and the second is bad. Rather, the first uses mathematical complexity

to obtain high accuracy with a coarse discretization, while the second uses mathematical simplicity to obtain high accuracy through a fine discretization.

The numerical simulations presented here are based on the step-function approach. In this approach, the discretization replaces the sinusoidal term ($P(t)$) in the Mathieu equation with a constant, converting the equation to a second-order linear ODE. Each segment can then be considered separately, with the initial conditions at the start of each segment determined by the final conditions at the end of the previous segment. Noting that the response during each step is exponential (real or imaginary) and that each step multiplies the previous step, these responses can be multiplied together to obtain the state transition matrix (STM) which gives the response of the column over the entire period. The stability boundaries for any arbitrary set of initial conditions can then be determined by multiplying the STM by the initial conditions.

Because from Floquet (equation 2.23) the response must have the same period as the excitation, the resulting state must then be in phase with the initial conditions, and the ratio of the two can be used to determine the stability. The boundaries are found when this ratio is equal to one, representing that the net effect of the excitation is exactly stable, neither exponentially growing nor exponentially decaying. When this ratio is less than one the column is said to be stable, as the response decays exponentially with time, and when this ratio is greater than one the column is said to be unstable, as the response grows exponentially with time.

3.1.1 Computational Complexity

The computational complexity of computing the STM for a particular excitation function and excitation frequency, and the computational complexity associated with calculating the stability using a previously calculated STM are, respectively, $O(n)$ and $O(1)$. While calculating the stability associated with multiple initial conditions increases the computation effort required from $O(1)$ to $O(n)$, because the number of segments is expected to be larger than the number of initial conditions under consideration, the complexity of calculating the STM almost certainly governs.

A more thorough analysis follows.

3.2 Discretization

The use of a step function introduces errors due to its deviation from the original excitation function (discretization error). In addition to these errors, the inaccuracy of the small-angle approximation over the duration of each segment introduces additional errors (approximation error). Both of these sources of error decrease as the number of segments (N) increases and the step function more faithfully represents the actual excitation. At the same time, the computational cost of computing the stability increases in direct proportion to the number of segments used. It is therefore important to establish the number of segments required to obtain a desired level of accuracy.

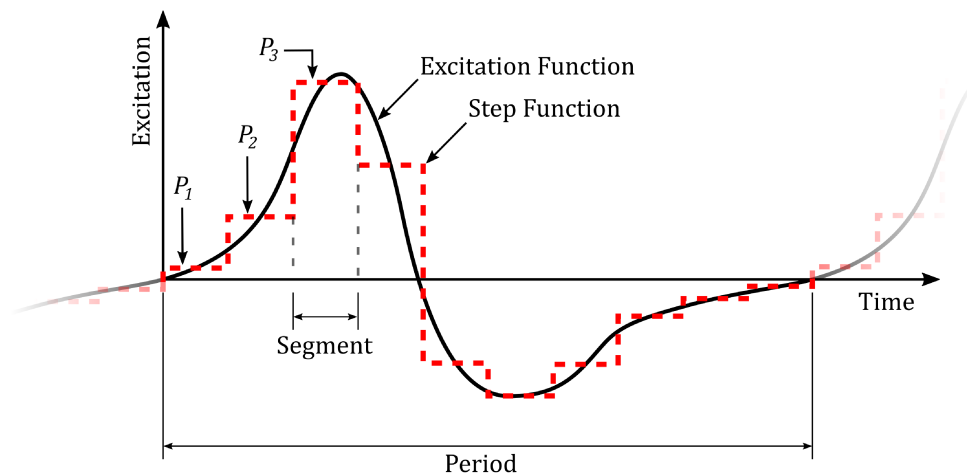


Figure 3.1: Discretizing a continuous function into a step function

3.2.1 Discretization Methods

There are innumerable methods for discretization, each with its own particular set of advantages and disadvantages. The most mathematically rigorous of these is probably the integration method, in which the function is integrated over each step and divided by the step width. The primary disadvantage to this approach is that it requires the calculation of

an integral, which is a very computationally expensive operation, and many other methods have dramatically better performance with negligible loss in accuracy. On the other end of the spectrum, the near-edge evaluation method simply evaluates the function at the start of each step. The primary disadvantage of this method is that it is dramatically less accurate than other methods with comparable computational difficulty.

As a general rule, it is not possible to definitively judge any method as being 'better' or 'worse' than any other method, only better or worse for a particular application. For example, the near-edge evaluation method is one of the worst possible methods to use when other methods are available but is useful in that it does not require any information about future data, and is therefore invaluable in cases where computation power is limited and no future data is available (such as with some PID controls). Similarly, despite its significant drawbacks during computation, the integration method is often ideal for mathematical manipulations, as it does not simplify the function prematurely. The discretized step function (P_n) is most conveniently described using the integration method,

$$P \approx P_n = \frac{1}{\tau} \int_{\tau \lfloor t/\tau - 1/\infty \rfloor}^{\tau \lceil t/\tau \rceil} P dt \quad (3.2)$$

where τ is the period (T) divided by the number of steps in the discretization (N),

$$\tau = \frac{T}{N} \quad (3.3)$$

and the subtraction of $1/\infty$ from the lower bound of the integral is used to prevent the floor ($\lfloor x \rfloor$) and ceiling functions ($\lceil x \rceil$) from returning the same value when t/τ is a natural number, which would result in a zero-width step being generated at each discontinuity.

While this is helpful for derivation, it is too computationally expensive for convenient use during actual computation. Because the excitation function is assumed to reflect a real (and therefore continuous) excitation, and the number of steps (or segments) is assumed to be relatively large, it is often reasonable to assume that each step is roughly linear over its width. As the integral of a line is simply the area of a trapezoid, this assumption suggests that the integration method can be very closely approximated during computation either by evaluating the function at the midpoint of each step (the midpoint method), or by evaluating the function at the beginning and end of each step and averaging the results (the averaging method).

In this thesis, the integration method will be used during derivation. As the excitation

being considered here is sinusoidal, the midpoint method will be very reflective of the integral method and will therefore be used during computation.

3.2.2 Discretization Errors

Because the step function introduces discontinuities in the excitation function, it introduces some degree of discretization error. This error shows up in the amplitude and frequency domains, both of which are of interest and must be considered.

3.2.2.1 Errors in the Amplitude Domain

The discretization error in the amplitude domain arises from the vertical discontinuity that exists between steps. Because the height of each vertical discontinuity is proportional to the slope of the excitation function, the magnitude of the discretization error is also proportional to the slope of the excitation function. If the error is expressed as a relative error, rather than as an absolute error, the scale of the excitation function becomes irrelevant, and the relative discretization error in the amplitude domain depends only on the number of steps in the discretization.

Assuming that the discretization is fine enough that the portion of the excitation function passing through each step is approximately straight, these errors tend to cancel each other out. In this case, the excitation function intersects the step function at the midpoint of each step, and the first half of the step is too high (or low) by the same amount that the latter half of the step is too low (or high). If the excitation function is perfectly straight, these errors should cancel each other out, yielding a net relative error in the amplitude domain of zero. In practice, the excitation function is not perfectly straight during each segment, and the actual error depends on the curvature, which can be estimated from the small-angle approximations. If the behavior is modeled successfully by a first-order approximation, the error in the amplitude domain should be approximately zero, otherwise the error should be proportional to $1/N^2$.

3.2.2.2 Errors In the Frequency Domain

To determine the significance of the discretization error in the frequency domain, the excitation function and the step function must be expressed in terms of their spectral density. This is most easily accomplished by performing a trigonometric Fourier transform on each function to convert them into infinite series of sinusoids. By combining a series of pulse

functions, the step function can be expressed as,

$$P \approx \sum_{f=0}^{\infty} \left[\frac{2 \sin\left(\frac{f\pi}{N}\right)}{f\pi} \sum_{n=1}^N P_n \cos(f\nu t_n) \right] \quad (3.4)$$

where f is the index of summation in the frequency domain, P_n is the height of the step function at step n , and t_n is the phase-shifted relative time:

$$t_n = 2\pi \left(\frac{t}{T} - \frac{n + \frac{1}{2}}{N} \right) \quad (3.5)$$

If the inner summation is combined using phasor addition, the product-to-sum trigonometric identities can be applied to the result to obtain an infinite series of linearly independent equations from each frequency. Each term in this series is given by,

$$P_f = \frac{2 \sin\left(\frac{f\pi}{N}\right)}{f\pi} \left(A_n \sin\left(\frac{2\pi}{T} f\nu t\right) + B_n \cos\left(\frac{2\pi}{T} f\nu t\right) \right) \quad (3.6)$$

where A_n and B_n are unknown constants, and P_f is the spectral density of the step function at $f\nu$. The error in the frequency domain (ϵ_f) can then be compared to the spectral density of the excitation function, S_P , which can be obtained through a Fourier transform:

$$\epsilon_f = \frac{S_P - P_f}{S_P} \quad (3.7)$$

If instead of calculating the value of P_n directly from the excitation function, these values are solved for algebraically such that they minimize the error function, an under-constrained system of equations with N equations and ∞ unknowns is obtained. As a result, values of P_n can be solved for to eliminate errors in the frequency domain up to $N \times$ the fundamental frequency, at the cost of having no control over the errors at higher frequencies. It is not practical to do so, as N is expected to be large and the complexity of finding an algebraic solution increases rapidly as the number of variables increases, but this gives a good estimate of the expected error in the frequency domain. As values of P_n are instead likely to be calculated from the excitation function using some discretization method, and the actual frequency response will be somewhat less accurate than this suggests, but the error in the lower frequencies should be at least as small as the error in the discretization function.

Because only the first few stability regions are generally of interest, the range of frequencies over which this accuracy estimate is reasonable is not typically limiting. For example, the size of the matrices considered in Table 1.1 suggests that only a few frequencies past the fundamental frequency associated with each stability region have a considerable impact on the stability behavior. As such, for N to be sufficiently large that the highest frequencies have essentially no impact on the result, N only needs to be several times larger than the highest frequency considered. This suggests that $N = 5$ is probably sufficient for the first stability region, perhaps $N = 10$ for the second, and so on. Because the calculation of higher-numbered stability regions requires greater precision in general as they are more sensitive to sources of error, the minimum acceptable value of N is expected to increase significantly as the number of the stability region being calculated increases.

3.2.3 Approximation Errors

The small-angle approximation is typically used either to simplify algebra by eliminating trigonometric and hyperbolic functions or to improve performance by avoiding their full computation. As both algebraic derivation and computational performance are of concern here, it is desirable to use the small-angle approximation for both purposes. While the first-order approximation is the one most commonly used, it is not the only one, with higher-order approximations providing greater accuracy (or smaller error) at the expense of greater complexity in both derivation and computation. As the error changes with the order of the approximation, the approximation error can be chosen so that it is smaller than the error appearing elsewhere in the analysis. Doing so is not without its complexities, however, as it requires that the trigonometric and hyperbolic functions be replaced with infinite series, which does not guarantee a simpler algebraic experience. As the computational performance is of greater concern than convenience, it is important to establish how much the order of the small-angle approximation affects the resulting error, so that an appropriate approximation can be selected.

The M^{th} -order approximations of the trigonometric sine and cosine consist of the first M terms of,

$$\cos(\mathbf{x}) \approx \cos_M(\mathbf{x}) = \sum_{m=0}^M (-1)^m \frac{\mathbf{x}^{2m}}{(2m)!} \quad (3.8a)$$

$$\sin(\mathbf{x}) \approx \sin_M(\mathbf{x}) = \sum_{m=0}^M (-1)^m \frac{\mathbf{x}^{2m+1}}{(2m+1)!} \quad (3.8b)$$

The M^{th} -order approximations for the hyperbolic sine and cosine functions are simpler in that every term is positive. They are otherwise identical to the trigonometric approximations:

$$\cosh(\mathbf{x}) \approx \cosh_M(\mathbf{x}) = \sum_{m=0}^M \frac{\mathbf{x}^{2m}}{(2m)!} \quad (3.9a)$$

$$\sinh(\mathbf{x}) \approx \sinh_M(\mathbf{x}) = \sum_{m=0}^M \frac{\mathbf{x}^{2m+1}}{(2m+1)!} \quad (3.9b)$$

The relative error associated with each function can be determined by,

$$\epsilon_M = \frac{f_M(\mathbf{x})}{f(\mathbf{x})} - 1 \quad (3.10)$$

where ϵ_M is the error associated with the M^{th} -order approximation, and f and f_M are a trigonometric or hyperbolic function and its M^{th} -order approximation, respectively.

As the direction of the error is not important when establishing the accuracy (though it is good to know if an approximation over-estimates or under-estimates stability), the absolute error is more meaningful metric. Additionally, because the width of each step is constant, the absolute relative error can be expressed in terms of the number of steps in the discretization, giving,

$$|\epsilon_M| = \left| \frac{f_M(2\pi/N)}{f(2\pi/N)} - 1 \right| \quad (3.11)$$

Because equation 3.11 only gives the relative error occurring in a single calculation, it does not indicate the error resulting from the repeated use of these approximations. Fortunately, the trigonometric and hyperbolic functions have essentially equal errors at small angles, with the primary difference being that the hyperbolic approximations are always less than the true value, while the trigonometric approximations alternate between being higher and lower. While the compound hyperbolic errors are difficult to estimate, the trigonometric errors can be estimated with the standard two-dimensional rotation matrix (\mathcal{R}):

$$\mathcal{R} = \begin{bmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{bmatrix} \quad (3.12)$$

When the period is divided into N segments, the rotation matrix can also be expressed in terms of N rather than θ :

$$\mathcal{R}_N = \begin{bmatrix} \cos(2\pi/N) & -\sin(2\pi/N) \\ \sin(2\pi/N) & \cos(2\pi/N) \end{bmatrix} \quad (3.13)$$

Applying this matrix N times to the initial conditions (q_0 and \dot{q}_0) should, in theory, rotate those initial conditions exactly one full revolution (to q_N and \dot{q}_N) without altering their magnitude. By comparing the magnitude of the resulting q and \dot{q} components, the cumulative error can be estimated. Or, expressed algebraically,

$$\begin{bmatrix} q_N \\ \dot{q}_N \end{bmatrix} = [\mathcal{R}_N]^N \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \approx \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \quad (3.14)$$

The absolute compound relative error is then the ratio of the scale of q_0 and \dot{q}_0 and q_N and \dot{q}_N :

$$|\epsilon_c| = \sqrt{\frac{q_N^2 + \dot{q}_N^2}{q_0^2 + \dot{q}_0^2}} - 1 \quad (3.15)$$

The absolute relative error and absolute compound relative errors resulting from this process are shown in Figure 3.2 for the first- and second-order approximations.

Unsurprisingly, the absolute relative error associated with the first-order small-angle approximations is proportional to $1/N^2$, corresponding to the most significant omitted term in the equivalent power series. The second-order approximation continues this trend, with a relative error proportional to $1/N^4$, for the same reason. The relative error of the multi-step rotations reveals that during each step the error from the small-angle approximations increases additively, resulting in a relative error that is $N \times$ as large as the relative error of a single small-angle approximation. The error associated with the repeated use of the M^{th} -order approximation is, therefore,

$$\epsilon_M \propto \frac{1}{N^{2M-1}} \quad (3.16)$$

Given the low accuracy of the first-order approximation and the high number of segments expected, the second-order (or higher) approximations are likely to be required. For example, if $\pm 1\%$ represents the minimum acceptable accuracy, the first-order approximation requires at least $N \approx 2000$ to be useful, while the second-order approximation requires only $N \approx 20$




Figure 3.2: Relative errors associated with the use of the small angle approximations by number of segments used

to produce useful results. At higher accuracies, this becomes even more pronounced, with the 3rd-order approximations requiring $N \approx 55$ to reach 1ppm accuracy, as compared to the $N \approx 430$ required by the 2nd-order approximations.

3.3 Determining the Column Response

Rather than consider the damped and undamped columns separately, it is clearer to consider the general case of a damped column, as the undamped column is simply a special case of the damped column in which $\beta = 0$.

For an undamped column ($c = 0$) with constant excitation ($\zeta = P_s$), the Mathieu-Hill equation (3.1) reduces to linear second-order ODE with constant coefficients. In the interest of keeping the analysis easily comparable to the analytical method used previously, substituting $P_d = 0$ into equations 2.19 and 2.20 gives the equation of motion of an undamped

column subject to constant excitation as,

$$\ddot{q} + \frac{c}{m}\dot{q} + \omega^2 q = 0 \quad (3.17)$$

where ω is given by equation 2.22.

From this point, the analysis can proceed one of two ways. The first option is to introduce a change in coordinate systems to eliminate the first derivative term. The second is to proceed with the first derivative in place. As coordinate changes have a tendency to muddy the waters, the derivation proceeds without the change in coordinate systems.

In the interests of clarity, the response during a single segment will be considered first, before generalizing the result to the entire period.

3.3.1 Deriving the Response In Regular Coordinates

The response of the column can be determined directly, without resorting to a change of variables. This has the obvious benefit of clarity (as there is no ambiguity about the relationship between variables), and the obvious drawback of complexity (as the first derivative term is not eliminated).

3.3.1.1 The Response Over the First Segment

Before the response can be determined, a solution to the ODE is necessary, which begins with the characteristic equation,

$$\lambda = -\beta_{\star} \pm \psi \quad (3.18)$$

where β_{\star} is the numerical damping term,

$$\beta_{\star} = \frac{c}{2m} \quad (3.19)$$

and ψ is an arbitrary substitution made for convenience:

$$\psi = \sqrt{\beta_{\star}^2 - \omega^2} \quad (3.20)$$

While ψ is almost certainly complex because ω is generally larger than β_{\star} , it is convenient to simply pretend that it is real for the time being. Assuming that the solution can be expressed such that all odd powers of ψ are eliminated, the complex values will be eliminated, and there will be no problems with this assumption. From this assumption, the solution to

the ODE is easily obtained in the usual fashion:

$$q = e^{-\beta_* t} (C_1 e^{\psi t} + C_2 e^{-\psi t}) \quad (3.21)$$

Putting the solution and its derivative into matrix form gives the system of equations:

$$\begin{bmatrix} q(t) \\ \dot{q}(t) \end{bmatrix} = e^{-\beta_* t} \begin{bmatrix} e^{\psi t} & e^{-\psi t} \\ (\psi - \beta_*)e^{\psi t} & -(\psi + \beta_*)e^{-\psi t} \end{bmatrix} \begin{bmatrix} C_1 \\ C_2 \end{bmatrix} \quad (3.22)$$

Setting $t = 0$ allows the integration constants (C_1 and C_2) to be solved for based on the initial conditions, giving,

$$\begin{bmatrix} C_1 \\ C_2 \end{bmatrix} = \frac{1}{2\psi} \begin{bmatrix} (\psi + \beta_*)q_0 + \dot{q}_0 \\ (\psi - \beta_*)q_0 - \dot{q}_0 \end{bmatrix} \quad (3.23)$$

where q_0 and \dot{q}_0 are the values of q and \dot{q} when $t = 0$.

At this point, it is helpful to introduce some additional changes to the notation. First, as the response at the end of the step is more interesting than the response in the middle of the step, call the time at the end of the first segment τ . Substituting the angular velocity into equation 3.3 gives τ without referring to the excitation period (T) explicitly:

$$\tau = \frac{2\pi}{N\nu} \quad (3.24)$$

Second, like with q_0 and \dot{q}_0 it is convenient to refer to q and \dot{q} as q_1 and \dot{q}_1 , respectively, when they are evaluated at $t = \tau$.

Substituting the integration constants back into the matrix and rewriting them in terms of the hyperbolic sine and cosine functions gives the response at the end of the first segment as a function of the initial conditions:

$$\begin{bmatrix} q_1 \\ \dot{q}_1 \end{bmatrix} = e^{-\tau\beta_*} \begin{bmatrix} \cosh(\psi\tau) + \frac{\beta_* \sinh(\psi\tau)}{\psi} & \frac{\sinh(\psi\tau)}{\psi} \\ (\psi^2 - \beta_*^2) \frac{\sinh(\psi\tau)}{\psi} & \cosh(\psi\tau) - \frac{\beta_* \sinh(\psi\tau)}{\psi} \end{bmatrix} \times \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \quad (3.25)$$

As the hyperbolic functions are not particularly convenient to write, they can be removed using the small-angle approximations in equations 3.9a and 3.9b. In addition to simplifying the equation, this has the beneficial effect of eliminating any odd powers of ψ ,

$$\begin{bmatrix} q_1 \\ \dot{q}_1 \end{bmatrix} = e^{-\tau\beta_*} \begin{bmatrix} \mathbf{C} + \mathbf{S}\beta_* & \mathbf{S} \\ -\mathbf{S}\omega^2 & \mathbf{C} - \mathbf{S}\beta_* \end{bmatrix} \times \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \quad (3.26a)$$

where,

$$\mathbf{C} = \sum_{m=0}^M \frac{\tau^{2m} (\beta_*^2 - \omega^2)^m}{(2m)!} \quad \mathbf{S} = \sum_{m=0}^M \frac{\tau^{2m+1} (\beta_*^2 - \omega^2)^m}{(2m+1)!} \quad (3.26b)$$

3.3.1.2 The Response Over the Full Period

The response of the damped column can be generalized to the entire period by noting that the column conditions at the end of the first step give the initial conditions for the analysis of the second step, which, aside from the differing axial load, results in an identical solution. Equation 3.26 is therefore easily generalized by labeling every variable with the segment number (n):

$$\begin{bmatrix} q_n \\ \dot{q}_n \end{bmatrix} = e^{-\tau\beta_*} \begin{bmatrix} \mathbf{C}_{n*} + \mathbf{S}_{n*}\beta_* & \mathbf{S}_{n*} \\ -\mathbf{S}_{n*}\omega_n^2 & \mathbf{C}_{n*} - \mathbf{S}_{n*}\beta_* \end{bmatrix} \times \begin{bmatrix} q_{n-1} \\ \dot{q}_{n-1} \end{bmatrix} \quad (3.27a)$$

where,

$$\mathbf{C}_{n*} = \sum_{m=0}^M \frac{\tau^{2m} (\beta_*^2 - \omega_n^2)^m}{(2m)!} \quad \mathbf{S}_{n*} = \sum_{m=0}^M \frac{\tau^{2m+1} (\beta_*^2 - \omega_n^2)^m}{(2m+1)!} \quad (3.27b)$$

Chaining the segments together by substituting q_{N-1} and \dot{q}_{N-1} , q_{N-2} and \dot{q}_{N-2} , and so on into the equation for q_N and \dot{q}_N expresses the response at the end of the period as a function of the initial conditions:

$$\begin{bmatrix} q_N \\ \dot{q}_N \end{bmatrix} = e^{-N\tau\beta_*} \prod_{n=1}^N \begin{bmatrix} \mathbf{C}_{n*} + \mathbf{S}_{n*}\beta_* & \mathbf{S}_{n*} \\ -\mathbf{S}_{n*}\omega_n^2 & \mathbf{C}_{n*} - \mathbf{S}_{n*}\beta_* \end{bmatrix} \times \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \quad (3.28)$$

3.3.1.3 Normalizing the Solution

Ideally, the stability can be expressed in the same relative terms used previously in section 2.3. This requires both τ and ω_n to be eliminated throughout equation 3.28, which can be accomplished by expanding τ (using equation 3.24) and β_* (using equation 3.19), rearranging equation 2.37 to get $\nu = \hat{\nu}\omega_0$ (and substituting it in), and expanding ω_0 (using equation 2.22). The same β used in the analytical solution (equation 2.33) can then be substituted in, though with the caveat that the ω , in this case, is ω_0 , and the β is, therefore, β_0 , and is specific to ω_0 . As with \hat{P}_d and \hat{P}_s , the axial excitation is readily normalized by introducing $\hat{P}_n = P_n/P_E$:

$$\begin{bmatrix} q_N \\ \dot{q}_N \end{bmatrix} = e^{-\frac{\pi\beta_0}{\nu}} \prod_{n=1}^N \begin{bmatrix} \mathbf{C}_n + \mathbf{S}_n\beta_0 & \frac{2}{\omega_0}\mathbf{S}_n \\ -4\frac{\omega_0}{2}\mathbf{S}_n \left(1 + \hat{P}_n\right) & \mathbf{C}_n - \mathbf{S}_n\beta_0 \end{bmatrix} \times \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \quad (3.29a)$$

where,

$$\mathbf{C}_n = \sum_{m=0}^M \frac{\left(\frac{\pi}{N\hat{\nu}}\right)^{2m} \left(\beta_0^2 - 4 + 4\hat{P}_n\right)^m}{(2m)!} \quad \mathbf{S}_n = \sum_{m=0}^M \frac{\left(\frac{\pi}{N\hat{\nu}}\right)^{2m+1} \left(\beta_0^2 - 4 + 4\hat{P}_n\right)^m}{(2m+1)!} \quad (3.29b)$$

Further simplification requires the initial conditions to be considered. From Floquet (see section 2.2.1) we know that the response of the column must have the same period as the excitation, with the amplitude scaled by some exponential function. Therefore, at the end of a full period of excitation, the column must have a velocity and deflection that are each some multiple of its initial conditions, and only one of the stability criteria in equation 3.30 need be considered to determine the stability,

$$\frac{q_N}{q_0} = \mathcal{P}_{11} + \mathcal{P}_{12} \frac{\dot{q}_0}{q_0} \quad \frac{\dot{q}_N}{\dot{q}_0} = \mathcal{P}_{21} \frac{q_0}{\dot{q}_0} + \mathcal{P}_{22} \quad (3.30)$$

where \mathcal{P} is the 2×2 matrix product series from equation 3.29. As usual, values less than one indicate stability, while values greater than one indicate instability. The stability boundaries occur when the result is precisely equal to one.

The most commonly considered initial condition for a column (at least in civil engineering) is the at-rest condition. In this case, the column is assumed to have a small (often assumed to be both imperceptible and unknown) initial deflection ($q_0 = \delta_0$) and an initial velocity of zero ($\dot{q}_0 = 0$). The second most commonly considered initial condition is, arguably, the in-motion condition, in which the initial deflection is zero ($q_0 = 0$) and the initial velocity has some non-zero value ($\dot{q}_0 = v_0$). Both of these conditions have the benefit of eliminating terms associated with the other, eliminating the need to consider the ratio at which energy in one form transforms into energy in the other. This simplifies stability conditions considerably, to,

$$\frac{q_N}{q_0} = \mathcal{P}_{11} \quad \frac{\dot{q}_N}{\dot{q}_0} = \mathcal{P}_{22} \quad (3.31)$$

For these two common initial conditions, the top-right and bottom-left terms ($[\text{STM}]_{12}$ and $[\text{STM}]_{21}$) have no bearing on the result. Closer investigation reveals that where a term and its reciprocal appear in the top-right and bottom-left corners of two 2×2 matrices, they have no impact on the top-left and bottom-right terms of the product, and appear

unchanged in their original positions in the product. Or, in more precise terms,

$$\left([\mathcal{C}] \odot [\mathcal{M}_1] \right) \times \left([\mathcal{C}] \odot [\mathcal{M}_2] \right) = [\mathcal{C}] \odot [\mathcal{M}_1] [\mathcal{M}_2] \quad (3.32)$$

where \odot is the element-wise multiplication operator, and \mathcal{C} is a 2×2 matrix of the form,

$$[\mathcal{C}] = \begin{bmatrix} 1 & 1/x \\ \mathbf{x} & 1 \end{bmatrix} \quad (3.33)$$

This can be easily extended by noting that because the constant appears in the result, this relationship also holds for any number of matrices multiplied together, provided that the coefficients in \mathcal{C} are common to all of the matrices being multiplied:

$$\prod_n \left([\mathcal{C}] \odot [\mathcal{M}_n] \right) = [\mathcal{C}] \odot \prod_n [\mathcal{M}_n] \quad (3.34)$$

From 3.34 it is clear that the calculated stability from equation 3.31 does not change if \mathcal{P} has its top-right entry multiplied by $\omega_0/2$ and its bottom-left entry by $2/\omega_0$. Having done this, ω_0 has been isolated to β_0 , indicating that the numerical method is truly normalized only in the case of an undamped column. Where the damping is non-zero, the numerical method can be normalized for any particular β_0 , the same as with the analytical method.

The response of the column over a single period in the at-rest or in-motion conditions is therefore,

$$\begin{bmatrix} q_N \\ \dot{q}_N \end{bmatrix} = e^{-\frac{\pi\beta_0}{\hat{\nu}}} \prod_{n=1}^N \begin{bmatrix} \mathbf{C}_n + \mathbf{S}_n\beta_0 & \mathbf{S}_n \\ -4\mathbf{S}_n(1 + \hat{P}_n) & \mathbf{C}_n - \mathbf{S}_n\beta_0 \end{bmatrix} \times \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \quad (3.35a)$$

where,

$$\mathbf{C}_n = \sum_{m=0}^M \frac{\left(\frac{\pi}{N\hat{\nu}}\right)^{2m} \left(\beta_0^2 - 4 + 4\hat{P}_n\right)^m}{(2m)!} \quad \mathbf{S}_n = \sum_{m=0}^M \frac{\left(\frac{\pi}{N\hat{\nu}}\right)^{2m+1} \left(\beta_0^2 - 4 + 4\hat{P}_n\right)^m}{(2m+1)!} \quad (3.35b)$$

3.3.2 Deriving the Response in Damped Coordinates

Alternatively to the previous derivation, the response of the column can be derived using a change in coordinates. This eliminates the first derivative term but also introduces the conceptual complexity associated with a change of variables.

3.3.2.1 Converting to Undamped Coordinates

Before determining the response of the column, the first-derivative term can be eliminated using a variable change, namely,

$$q = e^{-\beta_{\star}t}r \quad (3.36)$$

where r is q in undamped coordinates, and β_{\star} has the same meaning as in section 3.3.1 (equation 3.19):

$$\beta_{\star} = \frac{c}{2m} \quad (3.37)$$

After this substitution, the ODE becomes,

$$\ddot{r} - (\beta_{\star}^2 - \omega^2)r = 0 \quad (3.38)$$

which has the solution,

$$r = C_1 e^{\psi t} + C_2 e^{-\psi t} \quad (3.39)$$

where ψ has the same meaning as in section 3.3.1 (equation 3.20).

The response can then be determined without the interference of the first derivative term. Naturally, because the ODE is not quite the same, the derivation looks slightly different at each step.

3.3.2.2 The Undamped Response Over the First Segment

Putting the solution and its derivative into matrix form gives the system of equations:

$$\begin{bmatrix} r(t) \\ \dot{r}(t) \end{bmatrix} = \begin{bmatrix} e^{\psi t} & e^{-\psi t} \\ \psi e^{\psi t} & -\psi e^{-\psi t} \end{bmatrix} \begin{bmatrix} C_1 \\ C_2 \end{bmatrix} \quad (3.40)$$

Setting $t = 0$ allows the integration constants (C_1 and C_2) to be solved for based on the initial conditions ($r_0 = r(0)$ and $\dot{r}_0 = \dot{r}(0)$):

$$\begin{bmatrix} C_1 \\ C_2 \end{bmatrix} = \frac{1}{2\psi} \begin{bmatrix} \psi r_0 + \dot{r}_0 \\ \psi r_0 - \dot{r}_0 \end{bmatrix} \quad (3.41)$$

Like with q_1 and \dot{q}_1 in section 3.3.1, it is convenient to introduce $r_1 = r(\tau)$ and $\dot{r}_1 = \dot{r}(\tau)$. Substituting the integration constants back into the matrix and rewriting them in terms of the hyperbolic sine and cosine functions gives the response at the end of the first segment

as a function of the initial conditions:

$$\begin{bmatrix} r_1 \\ \dot{r}_1 \end{bmatrix} = \begin{bmatrix} \cosh(\psi\tau) & \frac{\sinh(\psi\tau)}{\psi} \\ \psi \sinh(\psi\tau) & \cosh(\psi\tau) \end{bmatrix} \times \begin{bmatrix} r_0 \\ \dot{r}_0 \end{bmatrix} \quad (3.42)$$

After substituting the M^{th} -order small-angle approximation and using equation 3.20 to eliminate ψ , this becomes,

$$\begin{bmatrix} r_1 \\ \dot{r}_1 \end{bmatrix} = \begin{bmatrix} \mathbf{C} & \mathbf{S} \\ \mathbf{S}(\beta_\star^2 - \omega^2) & \mathbf{C} \end{bmatrix} \times \begin{bmatrix} r_0 \\ \dot{r}_0 \end{bmatrix} \quad (3.43a)$$

where,

$$\mathbf{C} = \sum_{m=0}^M \frac{\tau^{2m} (\beta_\star^2 - \omega^2)^m}{(2m)!} \quad \mathbf{S} = \sum_{m=0}^M \frac{\tau^{2m+1} (\beta_\star^2 - \omega^2)^m}{(2m+1)!} \quad (3.43b)$$

3.3.2.3 The Damped Response Over the First Segment

Because the product term that appears when chaining the segments together restricts the ability to manipulate the equation, the change in variables must be eliminated at this point. This can be done by expressing $r(t)$ and $\dot{r}(t)$ in terms of $q(t)$ and $\dot{q}(t)$ beginning from equation 3.36. In doing so, it is important to note that because the matrix has already been normalized, this also requires that $\beta_\star = \beta_0\omega_0/2$ be substituted into the solutions. By rearranging equation 3.36 and its derivative, the following relationships are obtained:

$$r = e^{\beta_\star t} q \quad \dot{r} = e^{\beta_\star t} \beta_\star q + e^{\beta_\star t} \dot{q} \quad (3.44)$$

At the beginning of the period, these reduce to,

$$r_0 = q_0 \quad \dot{r}_0 = \beta_\star q_0 + \dot{q}_0 \quad (3.45)$$

while at the end of the segment ($t = \tau$) these become,

$$r_1 = e^{\tau\beta_\star} q_1 \quad \dot{r}_1 = e^{\tau\beta_\star} \beta_\star q_1 + e^{\tau\beta_\star} \dot{q}_1 \quad (3.46)$$

Substituting these values back into equation 3.43 gives the response of the column over a single segment in the original coordinate system,

$$\begin{bmatrix} q_1 \\ \dot{q}_1 \end{bmatrix} = e^{-\tau\beta_\star} \begin{bmatrix} \mathbf{C} + \mathbf{S}\beta_\star & \mathbf{S} \\ -\mathbf{S}\omega^2 & \mathbf{C} - \mathbf{S}\beta_\star \end{bmatrix} \times \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \quad (3.47)$$

which is the same as equation 3.26, and leads to the solution obtained in section 3.3.1.

3.3.2.4 The Damped Response Over the Full Period

If the response is generalized to the entire period before reversing the change of variables, which is done in the same manner as in section 3.3.1, equation 3.43 becomes,

$$\begin{bmatrix} r_N \\ \dot{r}_N \end{bmatrix} = \prod_{n=1}^N \begin{bmatrix} \mathbf{C}_{n\star} & \mathbf{S}_{n\star} \\ \mathbf{S}_{n\star}(\beta_\star^2 - \omega_n^2) & \mathbf{C}_{n\star} \end{bmatrix} \times \begin{bmatrix} r_0 \\ \dot{r}_0 \end{bmatrix} \quad (3.48a)$$

where,

$$\mathbf{C}_{n\star} = \sum_{m=0}^M \frac{\tau^{2m} (\beta_\star^2 - \omega_n^2)^m}{(2m)!} \quad \mathbf{S}_{n\star} = \sum_{m=0}^M \frac{\tau^{2m+1} (\beta_\star^2 - \omega_n^2)^m}{(2m+1)!} \quad (3.48b)$$

This can be converted into damped coordinates in the same manner as in section 3.3.2.3 except that the end of the period ($t = N\tau$) is considered instead of the end of the segment ($t = \tau$):

$$r_N = e^{N\tau\beta_\star} q_N \quad \dot{r}_N = e^{N\tau\beta_\star} \beta_\star q_N + e^{N\tau\beta_\star} \dot{q}_N \quad (3.49)$$

The result differs from section 3.3.2.3 significantly because these substitutions occur outside the product, and therefore cannot be rearranged inside the product. The solution in regular coordinates must therefore be expressed in terms of the matrix product,

$$\begin{bmatrix} q_N \\ \dot{q}_N \end{bmatrix} = e^{-N\tau\beta_\star} \begin{bmatrix} \mathcal{P}_{11} + \beta_\star \mathcal{P}_{12} & \mathcal{P}_{12} \\ \mathcal{P}_{21} + \beta_\star (\mathcal{P}_{22} - \mathcal{P}_{11}) - \beta_\star^2 \mathcal{P}_{12} & \mathcal{P}_{22} - \beta_\star \mathcal{P}_{12} \end{bmatrix} \times \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \quad (3.50)$$

where \mathcal{P} is the matrix product in equation 3.48.

3.3.2.5 Normalizing the Solution

Equations 3.48 and 3.50 can be normalized in the same manner as in section 3.3.2.3,

$$\begin{bmatrix} q_N \\ \dot{q}_N \end{bmatrix} = e^{-\frac{\pi\beta_0}{\hat{\nu}}} \begin{bmatrix} \mathcal{P}_{11} + \frac{\beta_0\omega_0}{2}\mathcal{P}_{12} & \mathcal{P}_{12} \\ \mathcal{P}_{21} + \frac{\beta_0\omega_0}{2}(\mathcal{P}_{22} - \mathcal{P}_{11}) - \left(\frac{\beta_0\omega_0}{2}\right)^2\mathcal{P}_{12} & \mathcal{P}_{22} - \frac{\beta_0\omega_0}{2}\mathcal{P}_{12} \end{bmatrix} \times \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \quad (3.51a)$$

where,

$$[\mathcal{P}] = \prod_{n=1}^N \begin{bmatrix} \mathbf{C}_n & \frac{2}{\omega_0}\mathcal{S}_n \\ \frac{\omega_0}{2}\mathcal{S}_n \left(\beta_0^2 - 4 + 4\hat{P}_n\right) & \mathbf{C}_n \end{bmatrix} \quad (3.51b)$$

and,

$$\mathbf{C}_n = \sum_{m=0}^M \frac{\left(\frac{\pi}{N\hat{\nu}}\right)^{2m} \left(\beta_0^2 - 4 + 4\hat{P}_n\right)^m}{(2m)!} \quad \mathcal{S}_n = \sum_{m=0}^M \frac{\left(\frac{\pi}{N\hat{\nu}}\right)^{2m+1} \left(\beta_0^2 - 4 + 4\hat{P}_n\right)^m}{(2m+1)!} \quad (3.51c)$$

Once again ω_0 shows up in the resulting function and must be removed if the function is to be normalized in terms of β_0 . As with the previous derivation it can only be accomplished for the at-rest and in-motion conditions. In these cases either q_0 or \dot{q}_0 is zero, eliminating the top-right and bottom-left terms in the first matrix. With these terms gone, the only remaining references to ω_0 in the matrix product are attached to \mathcal{P}_{12} and \mathcal{P}_{21} , and every reference to \mathcal{P}_{12} in the solution contains $\frac{\omega_0}{2}$ (the solution does not reference \mathcal{P}_{21} once the bottom-left term has been eliminated). Using equation 3.34, ω_0 can then be brought inside the matrix product, eliminating it from the equation.

The solution in the at-rest or in-motion cases can therefore be written as,

$$\begin{bmatrix} q_N \\ \dot{q}_N \end{bmatrix} = e^{-\frac{\pi\beta_0}{\hat{\nu}}} \begin{bmatrix} \mathcal{P}_{11} + \beta_0\mathcal{P}_{12} & 0 \\ 0 & \mathcal{P}_{22} - \beta_0\mathcal{P}_{12} \end{bmatrix} \times \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \quad (3.52a)$$

where,

$$\mathcal{P} = \prod_{n=1}^N \begin{bmatrix} \mathbf{C}_n & \mathbf{S}_n \\ \mathbf{S}_n \left(\beta_0^2 - 4 + 4\hat{P}_n\right) & \mathbf{C}_n \end{bmatrix} \quad (3.52b)$$

and,

$$\mathbf{C}_n = \sum_{m=0}^M \frac{\left(\frac{\pi}{N\hat{\nu}}\right)^{2m} \left(\beta_0^2 - 4 + 4\hat{P}_n\right)^m}{(2m)!} \quad \mathbf{S}_n = \sum_{m=0}^M \frac{\left(\frac{\pi}{N\hat{\nu}}\right)^{2m+1} \left(\beta_0^2 - 4 + 4\hat{P}_n\right)^m}{(2m+1)!} \quad (3.52c)$$

3.3.3 The Response of an Undamped Column

In the undamped case ($\beta_0 = 0$), either solution reduces to,

$$\begin{bmatrix} q_N \\ \dot{q}_N \end{bmatrix} = \prod_{n=1}^N \begin{bmatrix} \mathbf{C}_n & \mathbf{S}_n \\ -4\mathbf{S}_n (1 + \hat{P}_n) & \mathbf{C}_n \end{bmatrix} \times \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \quad (3.53a)$$

where,

$$\mathbf{C}_n = \sum_{m=0}^M \frac{\left(\frac{\pi}{N\hat{\nu}}\right)^{2m} (1 + \hat{P}_n)^m}{(2m)!} \quad \mathbf{S}_n = \sum_{m=0}^M \frac{\left(\frac{\pi}{N\hat{\nu}}\right)^{2m+1} (1 + \hat{P}_n)^m}{(2m+1)!} \quad (3.53b)$$

3.4 Stability Boundaries Estimates

The primary control over both the computational effort required and the accuracy of the numerical method is provided by the number of segments in the discretization. On the one hand, a sufficiently large number of steps must be selected to provide the required accuracy, while on the other, too many points will require excessive computation, delaying the results and increasing their cost. Secondarily, the order of the approximation also affects the accuracy of the solution, along with the associated computational cost. Table 3.1 shows the time required to compute the stability conditions for a quarter of a million points using the first-, second-, and third-order approximations of a damped column. The undamped column had slightly better performance than this (around 10-20% faster), but the general trend was very consistent between the two cases and the difference was not significant enough to warrant a more detailed comparison.

Table 3.1: Computation time by order of approximation (250k pts)

N	Order of approximation		
	1 st	2 nd	3 rd
20	156ms	171ms	194ms
50	306ms	362ms	378ms
100	574ms	748ms	796ms
200	1.37s	1.28s	1.37s
500	2.58s	3.19s	3.38s


From these times it is clear that the computational complexity is indeed $O(n)$, and that the cost of using a higher-order approximation is not particularly large. Unlike N , which can take a wide range of values, the order of the approximation also has relatively few reasonable values. The first task in optimizing the numerical solution is therefore to determine the order of approximation to use. Once an order of approximation is selected, the accuracy can be determined in greater detail.

3.4.1 The Undamped Stability Boundaries


To make the stability behavior of the column more legible, the plots in this section are presented as contour plots with a logarithmic z-axis, rather than just the stability boundaries. As a result, the stability boundaries are located at zero (indicated with a thick black line), rather than at one as presented previously, and the other lines represent stable responses (for negative contours) and unstable responses (positive contours). The colored lines indicate the general nature and scale of the response (with yellow representing unstable responses and purple representing stable responses) as the frequency and excitation move away from the stability boundaries, and are useful primarily for illustrative purposes.

The first-order approximation is faster than the alternatives, but Figure 3.3a deviates significantly from the analytical solution at low relative loadings, suggesting that even at very fine discretizations the accuracy is still fairly low. Interestingly, the correct behavior can be seen in the diagram, with the logarithmic stability contour at 0.01 being a very close match to the analytical stability boundaries. Some rudimentary testing suggests that the logarithmic stability that provides the best match for the actual stability behavior is very slightly lower than 0.01 and that the error exists not only at low relative loadings (though it is most noticeable there), but is consistent throughout the entire plot. While increasing N decreases this issue, it appears to decrease proportionally to N , requiring very fine discretizations for any practical accuracy. As the stability boundary is located at 1, and $e^{0.01} \approx 1.01$, this suggests that the first-order approximation introduces a relative error of $5/N$ in addition to any other sources of error from the analysis. Assuming this relationship is exact suggests that five thousand steps are required for 0.1% accuracy, and 1ppm accuracy requires five million steps.

In contrast, the stability diagrams produced by the second-order approximation (figure 3.3b) are indistinguishable from those produced using higher-order approximations but are obtained more quickly. For example, a sample point tested at $\hat{P} = 0.2$ and $\hat{\nu} = 1.9$ (where



(a) Using the first-order approximations



(b) Using the second-order approximations

Figure 3.3: Natural logarithm of stability for an undamped column ($N = 500$) using the first- and second-order approximations


the error is expected to be most visible) using the second-order analysis differs by only 41ppb from the results computed using the 10th-order approximations. This is very close to $5/N^3$, which agrees well with the accuracy predicted in section 3.2.3. Again assuming this relationship is exact, achieving 0.1% and 1ppm accuracies should require only $N = 18$ and $N = 171$, respectively. Overall, it is difficult to justify using the first-order approximation in any circumstance, as the second-order approximation offers dramatically better computational efficiency. Comparing the first- and second-order analyses at $N = 500$ and $N = 200$, respectively, the second-order analysis achieves ten-thousand times the accuracy in half the time. As the accuracy of every part of the numerical method decreases as N increases, increasing the number of steps has the effect of reducing error throughout the entire analysis, not just in the approximation.

While the plots of the second-order and third-order approximations are indistinguishable at $N = 500$, differences are clear at $N = 50$. This suggests that the order of the approximation is more important as the value of N drops and that any desired accuracy has an associated value of N at which the computation time is minimized. This is not entirely accurate, however, as the accuracy is not only a function of N but also depends on the particular excitation function being used. As a result, there is likely to be some minimum value of N below which the excitation function is not well represented, and the accuracy with which the excitation is represented is likely to govern the accuracy of the solution. Judging by how similarly the stability regions appear to the analytical solution, it does not appear likely that approximations higher than the third-order approximation are required for this to occur, with an estimated accuracy of 50ppm at only $N = 10$.


3.4.2 The Damped Stability Boundaries

Before discussing the damped stability diagrams in greater detail, it is important to note that the stability diagrams shown here are for the adjusted solutions as discussed in section 4.1. The primary intent with the figures shown in this chapter is to demonstrate the similarity between the behavior of the damped and undamped numerical methods, as this discussion focuses primarily on the order of the approximation that is most appropriate.

The first-order approximation (figure 3.5a) shows the same offset behavior in the damped case as it does in the undamped case, with the 0.01 contour being fairly accurate in predicting the actual stability boundaries at $N = 500$. Again, the second-order approximation (figure 3.5b) suffers a slight speed penalty but offers greatly improved accuracy, suggesting that there




(a) Using the second-order approximations




(b) Using the third-order approximations

Figure 3.4: Natural logarithm of stability for an undamped column ($N = 50$) using the second-order and third-order approximations



(a) Using the first-order approximations



(b) Using the second-order approximations

Figure 3.5: Effect of varying the order of approximation on the natural logarithm of stability for a damped column ($\beta_0 = 0.1$, $N = 500$)

is no real justification for using the first-order approximation. As this trend is the same as that in the case of the undamped column, it is reasonable that the third-order approximations represent the most suitable solutions in the damped case as well. As with the undamped solution, higher-order approximations can be used for the damped column, but they are not recommended as it is doubtful that they will result in a significant improvement in accuracy and they take longer to perform.

3.5 Summary

As any numerical solution requires discretization, the effect of discretization was considered and the process was shown to produce errors that are controllable in both frequency and amplitude. The excitation function was approximated by a step function, and the numerical matrix method was developed for a damped column in two different ways. Both of these solutions reduce to the same solution in the case of an undamped column. The solutions produced by the numerical method were then compared visually to estimate the order of the small-angle approximation that should be used in the hyperbolic functions, with the third-order approximations being considered acceptable.

Chapter 4

Comparison of the Analytical Solution and the Numerical Matrix Solution

Like the small-angle approximation error, the overall accuracy of the numerical method can be established by comparing it to the analytical solutions. Unlike the small-angle approximations, however, it is important to note that the error is theoretical, in that it represents the deviation from the theoretical results. In an ideal world, the numerical method could be compared to experimental values to determine a true error.

As the analytical solution is not burdened by troublesome factors like noise or idealizations, it makes sense to first compare the numerical method to the analytical method to establish that it is at least theoretically sound. Once this has been completed, the numerical method can then (using a very large number of points), provide a reasonable indication of the error associated with each order of the analytical solution. In both cases, because the analytical solution only exists at the stability boundary (where it has a value of one) the error can only be determined at the stability boundaries predicted by the analytical function. As a result, the relative error is not a measure of the error in the location of the stability boundary, only the error in the magnitude of the stability.

From the previous analysis in section 3.4, it is clear that the magnitude of the exponential response increases as the distance from the stability boundary increases. Thus, a point not located close to a stability boundary (where the error is known) will have an error smaller than the difference between the measured stability and one. Or, in other words, if a point is far enough from a stability boundary that the error associated with it is unknown, it will be at least small enough that the numerical method will report the correct stability behavior. If the response is to be determined, then this simplification does not necessarily hold, and some other means must be devised to determine the relative error.

4.1 Comparing the Solutions Visually

When the undamped numerical and analytical solutions are compared visually (see Figure 4.1), the differences between them are few.




Figure 4.1: Comparison of the undamped numerical solution ($N = 20$) and 2nd-order through 4th-order analytical solutions

The most obvious difference between the two is that 2nd-order analytical solution does a poor job of predicting the 3rd stability region. This failure should come as no surprise, given that the order of the analytical solution required for reasonable accuracy is usually one higher than the stability region being considered, and this represents the case where the region is instead higher than the order of the analytical solution. As a result, the accuracy in this case is expected to be very poor.

If the third stability region is zoomed in on (Figure 4.2), the differences become more apparent. As the order of the analytical solution increases, so too does the accuracy, with the stability boundaries approaching the true value. The numerical solution also moves but

showing this movement only serves to obscure the plot because it does so very gradually. Interestingly, while the analytical solution only approaches the true value when the order of the analytical solution meets or exceeds the number of the region being considered, the numerical solution is located in approximately the right location even when the discretization is very coarse.

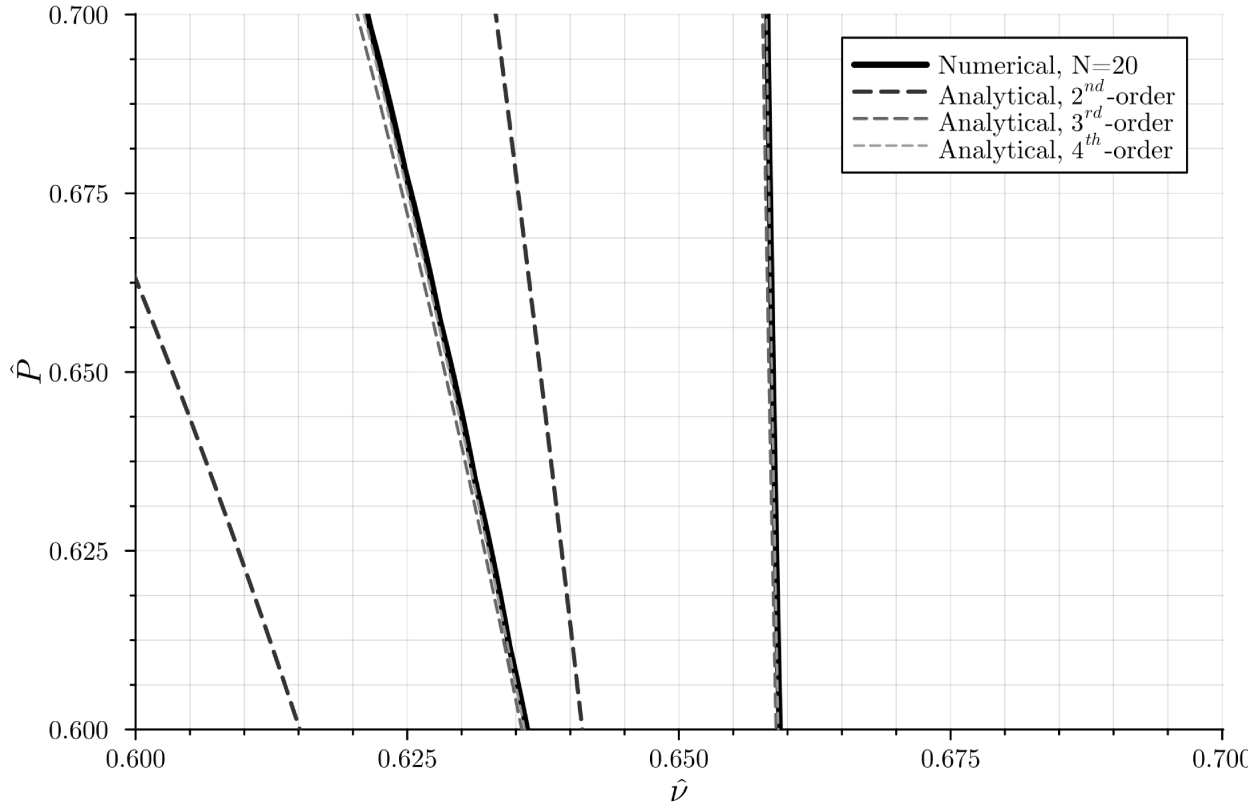


Figure 4.2: Zoomed in view of the undamped numerical solution ($N = 20$) and 2nd-order through 4th-order analytical solutions

The comparison in the damped case is more difficult. As written in Chapter 3, the derived matrix produces results that are very much at odds with the analytical method. They are internally consistent in that both derivations produce identical results, which lends support to the derivation performed, but this isn't overly helpful given that the results are not accurate. While they can be made to fit the analytical solution almost perfectly, as shown in Figure 4.3, there is no discernible justification for the changes that must be made to do so.




Figure 4.3: Comparison of the adjusted damped ($\beta_0 = 0.05$) numerical solution ($N = 20$) and 2nd-order through 4th-order analytical solutions

To obtain these curves, the exponential term in both methods should be $e^{-\beta_0/4\hat{\nu}^2}$ instead of $e^{-\pi\beta_0/\hat{\nu}}$. Additionally, in the matrix solution derived using the change of variables, the stability is given by \mathcal{P}_{11} instead of $\mathcal{P}_{11} + \beta_0\mathcal{P}_{12}$, while in the matrix derived without the change of variables the stability is given by $\mathcal{P}_{11} - \beta_0\mathcal{P}_{12}$ instead of \mathcal{P}_{11} . With these changes made, the response of the column becomes,

$$\begin{bmatrix} q_N \\ \dot{q}_N \end{bmatrix} = e^{-\frac{\beta_0}{4\hat{\nu}^2}} \prod_{n=1}^N \begin{bmatrix} \mathbf{C}_n + \mathbf{S}_n\beta_0 & \mathbf{S}_n \\ -4\mathbf{S}_n(1 + \hat{P}_n) & \mathbf{C}_n - \mathbf{S}_n\beta_0 \end{bmatrix} \times \begin{bmatrix} q_0 \\ \dot{q}_0 \end{bmatrix} \quad (4.1a)$$

where,

$$\mathbf{C}_n = \sum_{m=0}^M \frac{\left(\frac{\pi}{N\hat{\nu}}\right)^{2m} \left(\beta_0^2 - 4 + 4\hat{P}_n\right)^m}{(2m)!} \quad \mathbf{S}_n = \sum_{m=0}^M \frac{\left(\frac{\pi}{N\hat{\nu}}\right)^{2m+1} \left(\beta_0^2 - 4 + 4\hat{P}_n\right)^m}{(2m+1)!} \quad (4.1b)$$

and the stability for the at-rest condition is given by,

$$\frac{q_N}{q_0} = \mathcal{P}_{11} - \beta_0 \mathcal{P}_{12} \tag{4.2a}$$

where,

$$\mathcal{P} = e^{-\frac{\beta_0}{4\hat{v}^2}} \prod_{n=1}^N \begin{bmatrix} \mathbf{C}_n + \mathbf{S}_n \beta_0 & \mathbf{S}_n \\ -4\mathbf{S}_n (1 + \hat{P}_n) & \mathbf{C}_n - \mathbf{S}_n \beta_0 \end{bmatrix} \tag{4.2b}$$

As mentioned earlier, there is no apparent theoretical justification for these changes, but making them causes both methods to produce identical results which closely approximate the analytical solution, as can be seen in Figure 4.4.




Figure 4.4: Zoomed in view of the damped ($\beta_0 = 0.05$) numerical solution ($N = 20$) and 2nd-order through 4th-order analytical solutions

The change in the exponential term primarily corrects the vertical position of the stability boundaries so that it is close to that of the analytical solution. Without this change the stability boundaries are several times higher than the analytical solution, dramatically over-

predicting the stability. The change in the method of calculating the stability primarily addresses issues with the positioning of one side of the stability boundaries or the other and is difficult to define without the correction in the exponential term. For each stability region observed, the addition, subtraction, or elimination of the $\beta_0\mathcal{P}_{12}$ has the effect of shifting one side of each stability region along a line roughly parallel to the other side. The direction of this movement varies depending on the direction of the excitation (initially positive or initially negative), and when the correction is performed this movement disappears entirely.

Once this correction has been performed, the same behavior is observed for the damped solution as for the undamped solution, with the numerical method very closely approximating the stability region boundaries even when N is small.


4.2 Quantifying The Relative Error in the Undamped Column

Figure 4.5 shows the relative error in the numerical solution for various values of N , obtained by calculating the value of the numerical solution at the stability boundaries predicted by the analytical solution. In both cases the results of the numerical method were compared to the 5th-order analytical solution, as that represents the most authoritative solution available. As expected, the error generally decreases with N and increases with both \hat{P} and the number of the stability region. This is a bit too general of an observation to be particularly useful in quantifying the error, but it is a helpful starting point.


One item of particular note is that there are multiple instances where the error reverses suddenly. These are most clearly noticeable in Figure 4.5b, but can also be seen in Figure 4.5a in the line for $N = 20$. These represent places where the stability behavior predicted by the numerical method transitions from stability to instability at the analytical boundary. In essence, these represent points where the relative error in the numerical analysis happens to be zero through sheer luck, and are not representative of the overall error. However the error is quantified, it cannot include these points in detail, as they are not representative of the actual error in the solution.

4.2.1 Accuracy of the Numerical Solution

Using the second stability region in the $N = 20$ curve as an example, it is clear that the actual error near the central frequency is 10-20ppm, not the 80-200ppb advertised by the

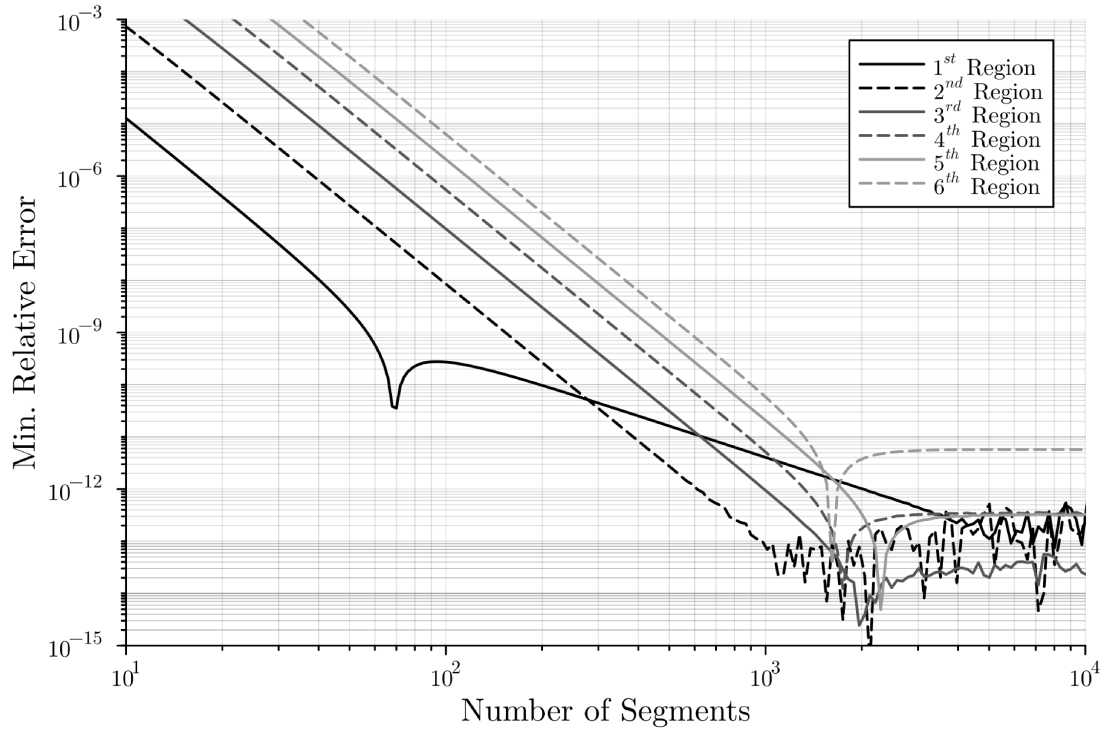


(a) As a function of the excitation frequency




(b) As a function of the peak axial excitation

Figure 4.5: Relative error of the numerical solution when compared to the 5th-order analytical solution for an undamped column



(a) Minimum error



(b) Maximum error

Figure 4.6: Relative error as a function of N for the first six stability regions

smallest error. Considering just this single region, it is clear from Figure 4.5b that aside from the crossings, the error transitions fairly smoothly between some minimum value of $\hat{P} = 0$, to some maximum value near the buckling load. Similar behavior is shown for each of the other buckling regions, and plotting the variation in these minimum and maximum errors for various values of N gives the limits shown in Figure 4.6.

The minimum error in these figures is the error measured near $\hat{P} = 0$. The maximum error is somewhat more difficult to obtain, as the location of the crossings makes the point of maximum error more difficult to locate. This can be clearly seen in the $N = 20$ and $N = 200$ curves in Figure 4.5b, where the crossings are located anywhere from near $\hat{P} = 0$ all the way to $\hat{P} = 0.9$. As a result of this challenge, the maximum error indicated is equal to the greatest error measured at any point in each region.

4.2.1.1 Estimating the Error as a Function of N

Considering only the slope of the lines in Figure 4.6a the minimum error of the numerical method can be estimated to be proportional to $1/N^5$, which is the cumulative error predicted by the third-order approximations. At the same time, Figure 4.6b clearly indicates that the maximum error is proportional to $1/N^2$. Taken together, the influence of N on the relative error of the numerical method appears to be of the form,

$$\epsilon_{num} \propto \frac{F_1(\hat{P}, R)}{N^5} \quad (4.3)$$

where F_1 is some function of the excitation (\hat{P}) and the stability region (R) which dictates the transition between the minimum and maximum errors.

Note that this is not perfectly generalized, as the relative error of the first stability region changes in slope from $1/N^5$ to $1/N^2$ after about $N = 70$. This is considered an acceptable loss for two reasons. First, the relative error at that particular point is already less than 1ppb, which far outstrips the accuracy with which any of the inputs to the numerical method can be determined. Thus, any analysis that depends on the accuracy of the numerical method will be governed by the accuracy of the inputs to the numerical method, rather than by the numerical method itself. Second, because the first stability region is much easier to find than the higher-numbered stability regions, N will most likely be selected based on the accuracy required to estimate the higher-numbered stability regions. As a result, it is more practical to express the accuracy such that the results are as accurate as possible for

the higher-numbered stability regions and include a footnote about the accuracy of the first stability region than it is to do the reverse.

4.2.1.2 Error at a Particular Axial Load

Once the influence of N on the relative error was isolated, the effects of variation in R and \hat{P} were needed. As R did not appear to affect the general curvature in an intuitive manner, the influence \hat{P} was considered, and found to be effectively represented as a power series, giving F_1 with the form,

$$F_1 = F_2(R) \hat{P}^{F_3(R)} \quad (4.4)$$

where F_2 and F_3 were functions of only R , and control the scale and curvature, respectively. The variation in scale was found to be exponential ($F_2 = 2^R$), while the variation in curvature was found to be linear ($F_3 = 2R$). After substituting these findings into equation 4.3, the estimated error became,

$$\epsilon_{num} \propto \frac{2^R \hat{P}^{2R}}{N^2} \quad (4.5)$$

4.2.1.3 Difference in Error Between Regions

While equation 4.5 gives the correct curvatures, it fails to identify the difference in the scale of the relative error between regions. As equation 4.5 already accounts for the difference in the scale of the relative errors when \hat{P} is large, it was only necessary to compare the relative scale of the minimum errors to address this. This was accomplished by normalizing the relative errors of the 2nd-through-6th regions against the relative error of the 1st region, as shown in Figure 4.7. As the behavior of the normalized relative errors varies as N increases, it is necessary to break the behavior into three separate zones before proceeding.

Of the three zones, it is easiest to first discuss the third zone, beginning around $N = 2000$. In the third zone, the normalized relative errors settle near unity, indicating that all regions have approximately the same error. This third zone appears to indicate the limiting accuracy of the matrix method, with Figure 4.6a suggesting a limiting accuracy of approximately 0.1ppt. This limit is very close to the limit of the precision of a 64-bit floating point number and is almost certainly limited by the available computing precision. Values of N greater than approximately 2000 produce no significant increase in accuracy

The first and second zones are delineated by the spike in relative error near $N = 70$ caused by the sudden increases in the relative accuracy of the first stability region. Before

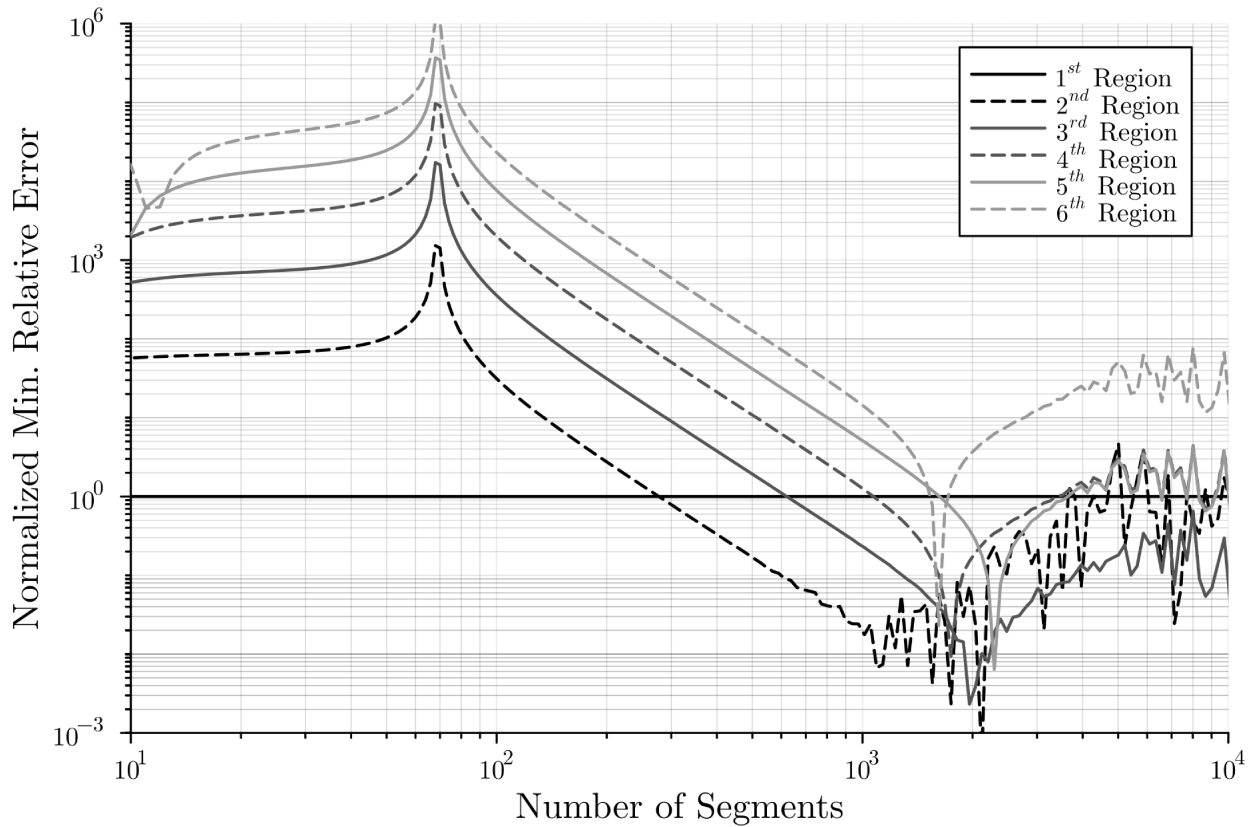


Figure 4.7: Minimum relative error as a function of N for the 2nd through 6th stability regions expressed as multiple of the minimum relative error of the 1st stability region

this point, the normalized relative error for each region is roughly constant, while after this point the normalized relative error decreases proportionally to $1/N^2$. This decrease is a direct result of the change of slope in the first stability region (to which the other regions are compared). If the change in slope in the first zone is neglected (as mentioned earlier), the change in slope in the second zone is also neglected, making the second zone identical to the first zone.

As only the first region is particularly meaningful, the normalized relative error is reasonably well approximated by comparing the values at $N = 30$. Doing so indicated that there is a clear sixth-order relationship between the scale of the error in each region and the number of the stability region (R), giving the minimum relative error in each stability region as R^6/N^5 . After Adding this to equation 4.5, a fairly accurate estimation of the relative error

in any particular stability region was obtained:

$$\epsilon_{num} \approx \frac{2^R \hat{P}^{2R}}{N^2} + \frac{R^6}{N^5} \quad (4.6)$$

Plotting the approximate stability gives Figure 4.8, which is a very reasonable approximation of Figure 4.5b.

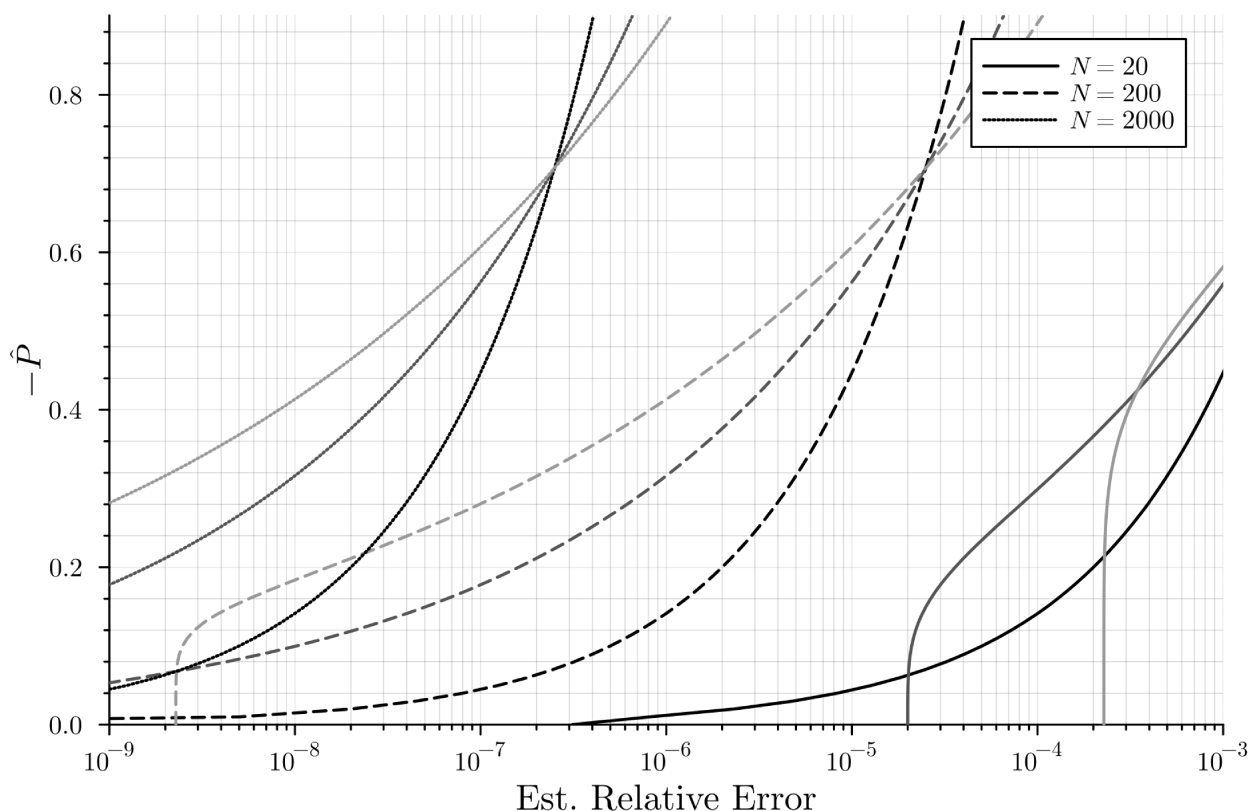


Figure 4.8: Estimated error in the undamped numerical method

4.2.2 Accuracy of the Analytical Solution

The accuracy of the analytical solution is easier to determine than that of the numerical solution in that provided N is sufficiently large, the difference between the numerical and analytical solutions is attributable overwhelmingly to the analytical solution. This can be seen in Figure 4.6b, where the relative error of the 3rd through 6th stability regions approach a particular value despite continued improvements to the accuracy of the numerical method.

To determine the accuracy of the analytical method, the primary challenge is to ensure that N is sufficiently high to capture the settled value of the stability region being considered.

Whether or not the error has converged can be determined by comparing the actual relative error at a point with the theoretical relative error predicted by equation 4.6. This equation represents the maximum error produced by the numerical method, and so if the actual relative error exceeds this value, then the actual error must contain an additional source of error, which is attributed to the analytical method. For simplicity, rather than searching for the error directly, the total error associated with each order of approximation of the first stability region was calculated at $N = 100,000$. The expected error (equation 4.6) was then subtracted from this value, with the residual attributed to the analytical error. Figure 4.9 shows the residual error for the second stability region as a function of the axial load and the order of the analytical solution used.




Figure 4.9: Estimated error in the analytical method ($R = 2$)

This residual error was found to vary with the peak axial excitation and region being determined (\hat{P} and R , as before) and the order of the analytical approximation (o). The

axial load was easily isolated, giving the analytical error in the form

$$\epsilon_{ana} \approx F_1(o, R) \hat{P}^{F_2(o, R)} \quad (4.7)$$

Where F_1 and F_2 are both functions of o and R .

Of these two functions, F_1 was more easily identified. Fitting an exponential function to F_1 isolated o , giving,

$$F_1 = F_{1a}(R) e^{oF_{1b}(R)} \quad (4.8)$$

where F_{1a} and F_{1b} are both functions of only R . It was observed that F_{1a} and F_{1b} were both well approximated by power series:

$$F_{1a} = 10^R \quad F_{1b} = -5.5R^{-1/3} \quad (4.9)$$

Substituting these solutions back into F_1 gave,

$$F_1 = 10^R e^{-5.5R^{-1/3}o} \quad (4.10)$$

The value of F_2 was less convenient to obtain because the stability regions occur as pairs (1st region, 2nd and 3rd regions, and so on) in which both regions are associated with similar coefficients. Intuitively, the accuracy of the analytical solutions should be related to the difference between the order of the solution and the stability region number ($o - R$), with higher accuracies associated with a larger difference. Plotting each difference as a function of the region indicated a linear relationship aside from the effect of the plateaus. This was addressed by replacing R in F_2 with $R + C_1(-1)^R$, which had the effect of pulling the adjacent regions towards each other and eliminating the plateaus. The coefficients of each line obtained in this manner were then found to vary linearly with the difference, and substituting each solution back into the previous one gave F_2 :

$$F_2 = 2.7o - 0.6R + 0.7(-1)^R + 2 \quad (4.11)$$

The solutions for F_1 and F_2 (equations 4.10 and 4.11) were then substituted into equation 4.7 to obtain the approximate error associated with the analytical solution:

$$\epsilon_{ana} \approx 10^R e^{-5.5oR^{-1/3}} \hat{P}^{2.7o-0.6R+0.7(-1)^R+2} \quad (4.12)$$

After some testing, the predicted error was found to be slightly too high in the odd-numbered regions, and slightly too low in the even-numbered regions. This is easily corrected by the addition of an alternating scaling term of magnitude ≈ 3 :

$$\epsilon_{ana} \approx 3^{(-1)^{R+1}} 10^R e^{-5.5oR} \hat{P}^{2.7o-0.6R+0.7(-1)^{R+2}} \quad (4.13)$$

Given that the scaling term is approximate, $3 \approx \pi$, and $10 \approx \pi^2$ (and that 10^R actually overestimates the fitting function prior to $R = 15$) it is convenient to rewrite the function to combine these two terms. After tinkering with the coefficients slightly, this can be rewritten so that it is the product of powers of π and \hat{P} with fewer coefficients:

$$\epsilon_{ana} \approx \pi^{2R-5oR} e^{-3/8oR} (-1)^R \hat{P}^{2.7o-0.6R+0.7(-1)^{R+2}} \quad (4.14)$$

With this adjustment made, and within the conditions that follow, the measured error predicted by equation 4.14 is within a factor of 2 of the error measured using the numerical method, slightly outperforming its predecessor. The first condition is that this function is very conservative when the order of the analytical solution is two or more higher than the region being estimated. This is a very manageable limitation in that the errors predicted for this situation are less than 1ppb up to around $\hat{P} = 0.6$ for the first three stability regions, and errors only peeking above 1ppm around $\hat{P} = 0.9$ after the third stability region. The predicted error is not greater than 0.1% until the 10th-order solution is used to estimate the 8th stability region, and even in this case the error is likely to be significantly smaller than the estimate, as equation 4.14 is conservative in these cases.

The second, and more limiting restriction is that this estimate is only as good as the inputs used in determining it. The numerical method is limited to errors greater than $\epsilon \approx 1 \times 10^{-11}$, with accuracy beginning to suffer somewhat before that point. Because the curve fitting was performed over $1 \leq o \leq 5$, $1 \leq R \leq 5$, and $0.05 \leq \hat{P} \leq 0.90$, the estimate may prove less accurate outside these bounds. Most notably, the curve fitting begins to be quite poor beginning at $R = 5$, where it suggests errors 4 \times larger than those measured. This was taken into account when the scale of the error ($1/2 \leq \epsilon_{est}/\epsilon_{meas} \leq 2$) was estimated, and so does not negate the validity of that estimate, but merely reinforces the bounds within which this estimate is reasonable.

As the error in each region increases with the axial excitation, this can be dramatically simplified if the axial excitation is taken as $\hat{P} = 1$. This eliminates the entire \hat{P} term, allowing for a greatly simplified comparison of the relative error in each stability region and

the effects of increasing R or o .

If each term in equation 4.14 is considered individually, the effect of changing a single variable is more clearly illustrated. The first term indicates that each pair of stability regions is $100\times$ less accurate than the preceding pair. The second term indicates that the rate at which each order of the analytical solution improves the accuracy decreases slightly as the region increases ($245\times$, $194\times$, $170\times$, and $154\times$ for the first four regions). Comparison with the first term suggests that each order of the analytical solution improves the accuracy by more than a two stability region increase decreases the accuracy (up to the 14th-order analytical solution, after which this no longer holds). The third term is interesting in that it primarily describes the curvature of the estimated error, but because \hat{P} is almost always less than one it also usually results in an increase in accuracy. The degree to which the accuracy improves is highly dependent on the magnitude of \hat{P} , with $\hat{P} = 1/10$ giving $500\times$ greater accuracy with every order of the analytical solution and $4\times$ lower accuracy with every region. In contrast, at $\hat{P} = 1/2$, the analytical solution gives only $6.5\times$ more accuracy with each order and $1.5\times$ less accuracy with each region. This effect is much stronger when \hat{P} is small than when it is large, and so this term primarily (but not exclusively) affects the expected error for small axial excitations.

4.2.2.1 Accuracy Thresholds

Table 4.1: Maximum axial excitation at 0.1% accuracy by region and order of the analytical solution

Analytical Solution	Instability Region				
	1	2	3	4	5
1 st -order	0.35	0.33	-	-	-
2 nd -order	1.3	0.94	0.31	0.33	-
3 rd -order	-	1.5	0.76	0.70	0.29
4 th -order	-	-	1.2	1.0	0.59
5 th -order	-	-	-	-	0.86
6 th -order	-	-	-	-	1.1

Table 4.2: Maximum axial excitation at 1ppm accuracy by region and order of the analytical solution

Analytical Solution	Instability Region				
	1	2	3	4	5
1 st -order	0.059	0.079	-	-	-
2 nd -order	0.50	0.39	0.090	0.11	-
3 rd -order	1.1	0.79	0.34	0.34	0.11
4 th -order	-	1.2	0.65	0.60	0.30
5 th -order	-	-	0.94	0.84	0.52
6 th -order	-	-	1.2	1.1	0.72
7 th -order	-	-	-	-	0.91
8 th -order	-	-	-	-	1.1

Equation 4.14 can be easily rearranged to obtain the maximum axial excitation associated with a particular accuracy (for a given region and order of the analytical solution),


$$\hat{P} \leq \ln(\pi) \exp \left(\frac{\log_{\pi}(\epsilon) - 2R + 5oR^{-3/8} + (-1)^R}{2.7o - 0.6R + 0.7(-1)^R + 2} \right) \quad (4.15)$$

which is very convenient in that it allows the maximum value of \hat{P} to be summarized very briefly, such as in tables 4.1 and 4.2.

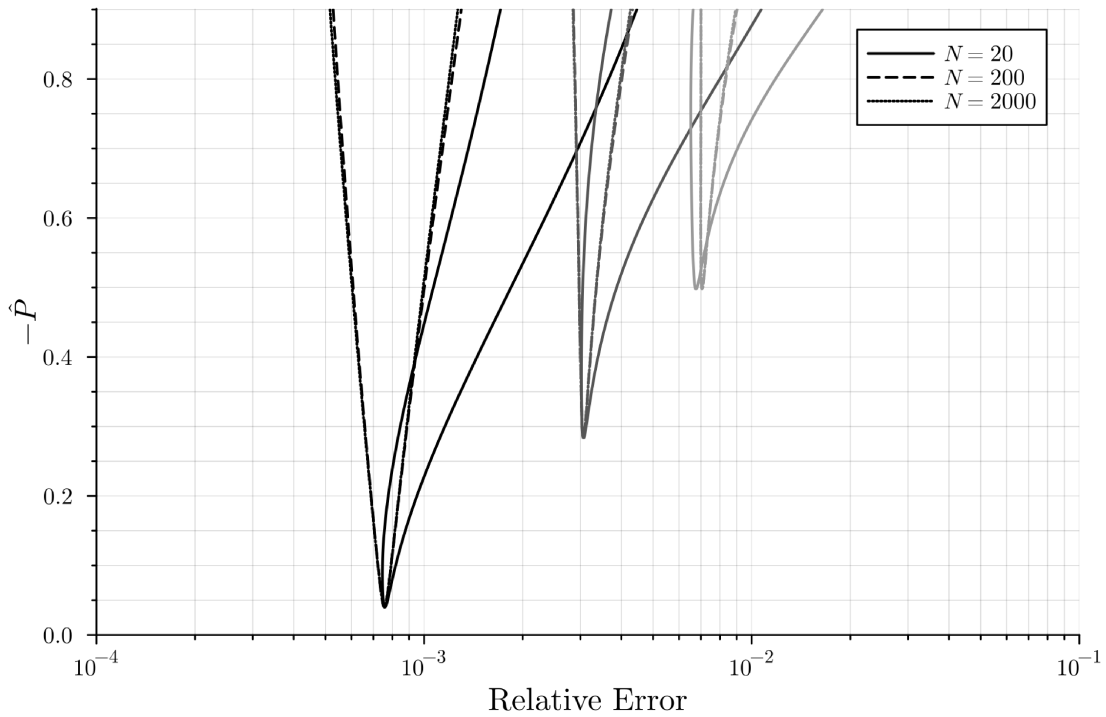
4.3 Estimating the Relative Error in the Damped Column

The relative error in the damped column is found in the same way as for the undamped column, with Figure 4.10 depicting the relative error proportional to the peak axial load and relative excitation frequency. This figure differs most notably from the undamped case in that the relative error approaches some value that depends exclusively on the frequency being considered. This curve appears regardless of β_0 , and defines the maximum achievable accuracy.

After fitting a power series to the relative errors for all regions obtained at $N = 10,000$, it was very clear that this curve is proportional to $1/\nu^2$. With the effect of $\hat{\nu}$ accounted for, the residual was found to be a quadratic function of β_0 . Combining these two values, the



(a) As a function of the excitation frequency



(b) As a function of the peak axial excitation

Figure 4.10: Relative error of the numerical solution when compared to the 5th-order analytical solution for a damped column ($\beta_0 = 0.02$)

minimum error obtainable for a damped column was found to be,

$$\epsilon_{num} \geq \frac{20\beta_0^2 - \beta_0}{4\hat{\nu}^2} \quad (4.16)$$

While this gives the best-case behavior, the worst-case behavior must also be quantified. Fortunately, because the maximum error is quite close to the minimum error, as shown in Figure 4.11, this can be done fairly easily. Noting that the maximum error is fairly well established (less than approximately 50 % higher than its ultimate value) as soon as $N \geq 40$, it is quite reasonable to assume that for any practical purpose, the maximum error will be a fairly good indicator of the actual error. Because the minimum error is at most 40% lower than this value, and the maximum error is at most 50% higher than this number, a very reasonable estimate of the error in the damped numerical analysis is simply,

$$\epsilon_{num} \approx \frac{60\beta_0^2 - 3\beta_0}{8\hat{\nu}^2} \quad (4.17)$$

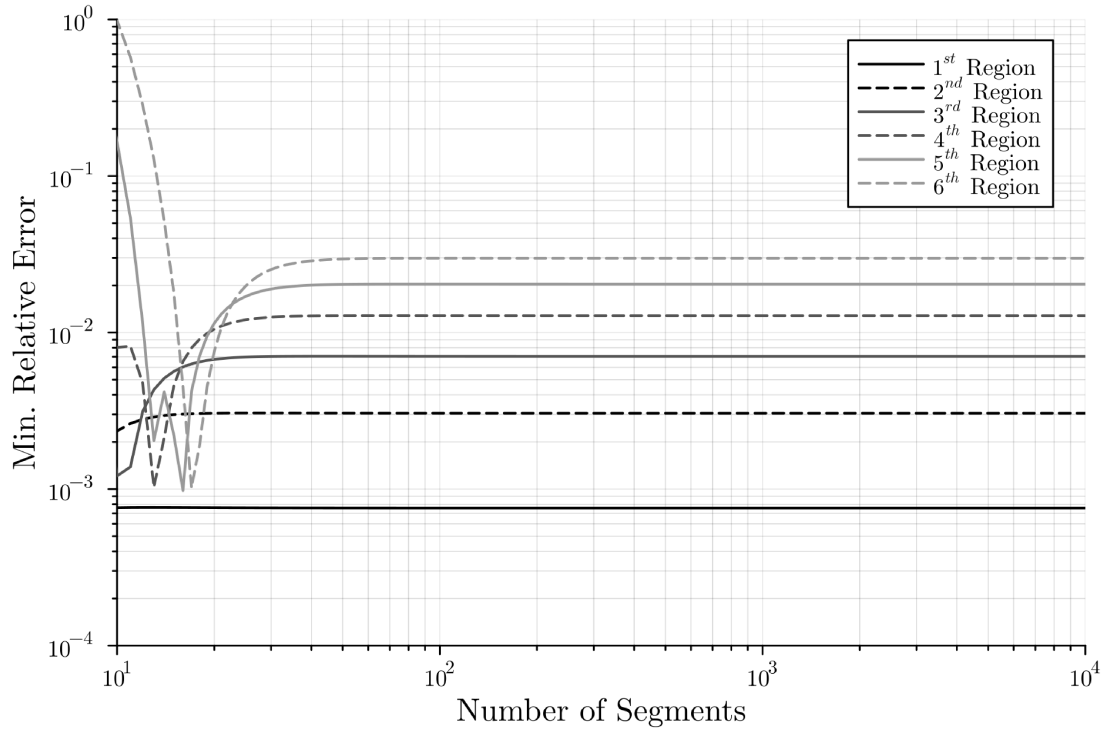
Because the error in the damped case is so strongly associated with β_0 , it is possible to overestimate the resulting accuracy for very small values of β_0 . To avoid this, it is reasonable to simply combine the two additively, as the worse-case error resulting from such a combination is a doubling of the estimated error, which is within the margin of error estimated in both cases. In the undamped case, the damped terms are eliminated, and unless β_0 is very small, the damped case governs, giving a single equation from which the error in both cases can be estimated:

$$\epsilon_{num} \approx \frac{2^R \hat{P}^{2R}}{N^2} + \frac{R^6}{N^5} + \frac{60\beta_0^2 - 3\beta_0}{8\hat{\nu}^2} \quad (4.18)$$

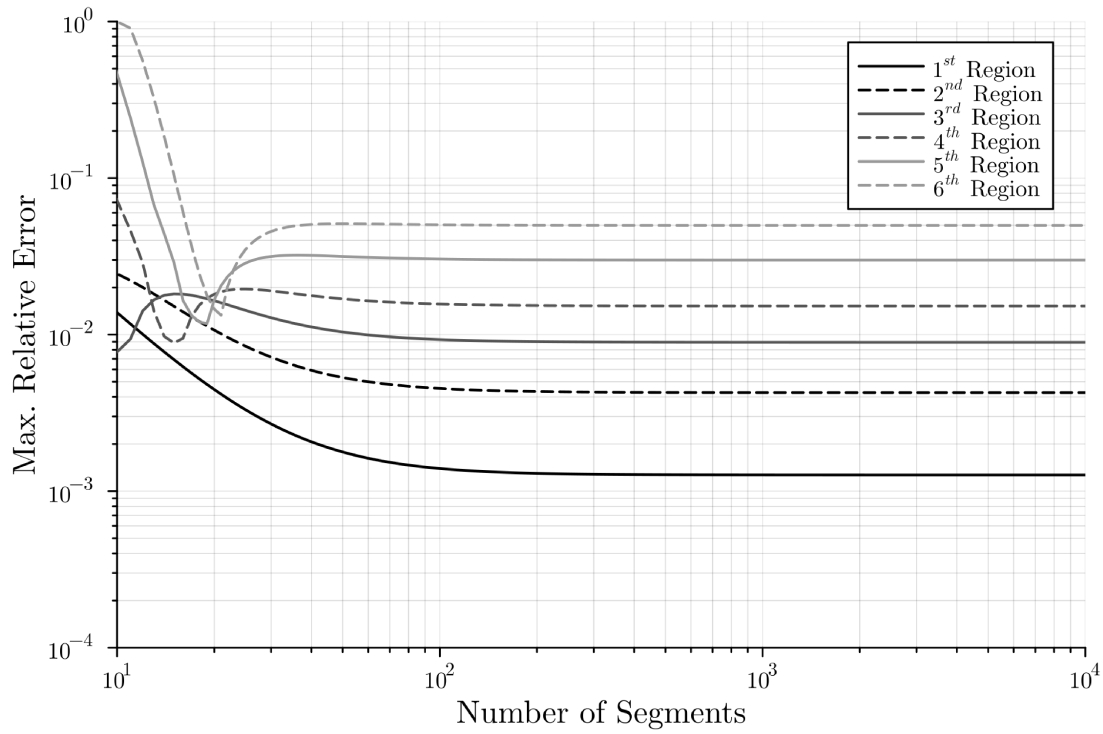
As the accuracies in the numerical method are relatively limited, it is not possible to use these accuracies to determine the error in the analytical analysis.

4.4 Summary

The numerical and analytical solutions were compared visually, and it was observed that the numerical method approaches the final position of the analytical method very quickly. The relative error in the undamped numerical matrix was determined from the fifth-order analytical solutions, with an estimation function being developed to account for the different



(a) Minimum error



(b) Maximum error

Figure 4.11: Relative error as a function of N for the first five stability regions ($\beta_0 = 0.02$)

regions, axial loads, and the number of segments used in the discretization. Once this error could be estimated, the error in the analytical solution was estimated based on the difference between the numerical and analytical solutions when the number of segments used in the numerical matrix method was selected to be excessively high. A fitting function was then developed from which the relative error in the analytical solution could be estimated for any region, axial load, and order of the analytical solution. The relative error in the (adjusted) damped numerical matrix was then determined, and was found to be well estimated with only the damping coefficient and the frequency being considered. The numerical error estimation functions were then combined to give the error in the numerical method as a function of the frequency, axial load, number of segments considered, stability region, and damping coefficient.

In general, the accuracy of the numerical method suffers as the frequency decreases, the axial load increases, or the damping coefficient increases. As with most numerical methods, the accuracy of the solution improves with the fineness of the discretization, though there are limits to the maximum accuracy obtainable through this method. Similar trends were observed for the analytical solutions, though the effect of damping could not be predicted due to the nature of the error in the numerical solution. The maximum excitation levels associated with $\pm 1\%$ and $\pm 1\text{ppm}$ accuracy were then developed. Overall, the analytical solution is more sensitive than the numerical solution to the degree of excitation, making the two methods difficult to equate directly.

Chapter 5

An Automated Apparatus for Obtaining the Stability Boundaries Experimentally

It appears that the dynamic stability of an elastically restrained column has not yet been verified experimentally. This presents something of a problem in that there is no firm practical evidence to support the numerical and analytical claims. The experimental apparatus (shown in Figure 5.1) was designed to address this issue.




Figure 5.1: Photograph of the experimental apparatus

The primary mechanical difference from the traditional experimental apparatus is in the supports, which can be configured over a wide range of fixities. These must be designed with care, as the degree of fixity must be fairly tightly controlled and the many required considerations at the supports present a complex set of requirements to navigate. Additionally, because a given connection can only be designed for a limited range of fixities, the design accommodates a range of interchangeable connections, each with its own range of fixities.

The automation of the experimental apparatus is valuable for two reasons. The first is that the probability of human error is dramatically higher than the probability of computer error (provided the computer is programmed correctly). The second, and more important reason, is that preliminary testing suggests that there exists a not-inconsiderable degree of variability in the results. This, in turn, suggests that it would be highly valuable to determine the variability of real-world behavior, as this cannot be easily measured using analytical and numerical solutions. To obtain the variability, a large volume of measurements need to be made, suggesting that it is worth the effort to automate the experimental procedure.

5.1 The Experimental Apparatus

The experimental apparatus is similar to that used in most experiments on the dynamic stability of columns, differing primarily in two respects. First, the design permits varying degrees of rigidity in the column connections. Second, the design is automated, allowing for the rapid collection of large volumes of data.

Because the automation of the testing procedure is fundamentally dependent on the experimental apparatus, it makes sense to first discuss the apparatus and its various components (as shown in Figure 5.2). Many of these components relate to the automation of the test, and only the physical assembly and the immediately connected items must be considered for now.

The specimen is supported between a load cell and a shaker motor, with the end conditions dictated by a pair of supports placed at each end. The shaker, load cell, and supports are all bolted to the structural steel frame, which rests on a pair of wood 4×4s rails. These are bolted firmly to the concrete bases on which they rest. Each half of the structural steel frame is attached to the wood rails with threaded rods, the length of which can be varied by turning the attached bolts. This allows for the two side channels in the steel frame to be positioned at any desired distance from each other. As a range of fixities is more readily

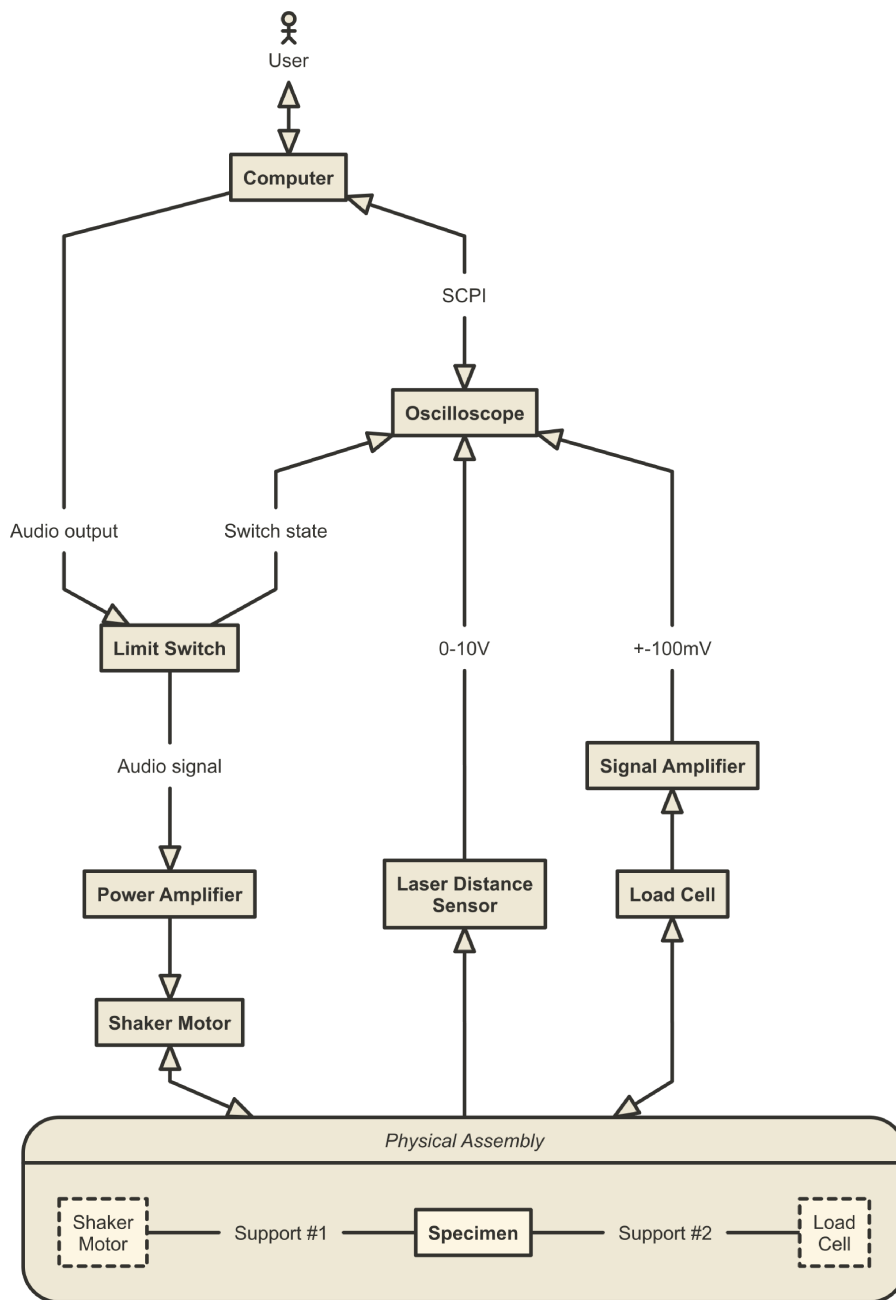


Figure 5.2: Schematic of the experimental apparatus

obtained from a variety of connections than from a single connection, the connections are designed to be interchangeable.

As with most similar experiments, a shaker motor is used to apply an axial load to a specimen, which is connected at either end to a support, with the far support containing a load cell through which the applied force is measured. The deflection of the member is measured at the midspan using a laser distance sensor.

5.1.1 Sensors and Measurement

To measure the behavior of the test specimen under dynamic loading, the load and deflection of the specimen must be measured. To identify the deflected shape of the column, the deflection must be taken in at least two places, from which the relative deflection can be determined. Keeping one deflection sensor in the same place, repeating this process multiple times gives the relative deflection at multiple points, and therefore the deflected shape of the column. In the case of a specimen supported with elastically restrained connections, the rotation of the specimen at the supports must also be measured to determine the actual fixity achieved.

The load is obtained using a 1 kN load cell. The deflection is obtained using a set of laser distance sensors reconfigured for various locations depending on what is being measured.

5.1.1.1 Load Cell

A 1kN load cell was used to measure the applied load. This is considerably higher than the expected peak force of 220 N generated by the shaker motor, but also represents the smallest load cell conveniently available. As the signal produced by a load cell is too small to read practically even if the load cell were more ideally sized, the signal will require amplification, rendering small optimizations in the size of the load cell largely irrelevant.

A schematic of the amplifier circuit is shown in Figure 5.3. One side of the load cell signal (in this case the positive side) uses one of the operational amplifier channels as a voltage buffer. The other side of the load cell signal (in this case the negative side) uses a second operational amplifier channel as a 100× differential amplifier references to the buffer output of the first channel, amplifying the signal without changing the reference point. The DC offset from the load cell dramatically reduces the accuracy of the measurements and is eliminated using a pair of capacitors to AC-couple the output.

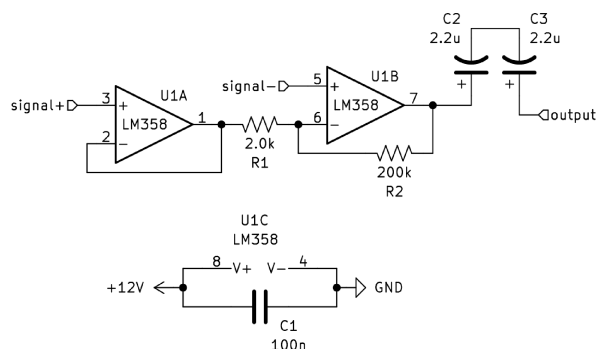


Figure 5.3: Circuit diagram for the load cell signal amplifier

Combined with the input impedance of the oscilloscope ($1\text{ M}\Omega$), the capacitors form a first-order high-pass filter. This affects the amplitude of the output, and so the coupling is designed so that it does not attenuate any of the frequencies of interest, with a corner frequency of approximately 0.145 Hz . Obtaining a corner frequency this low requires a relatively large capacitance ($1.1\text{ }\mu\text{F}$), which is inconvenient to obtain using unpolarized capacitors. This is addressed by placing two polarized capacitors back-to-back, which has the effect of halving their capacitance and eliminating their polarization.

5.1.1.2 Laser Distance Sensors - CP24MHT80

The deflections of the specimen were measured using three CP24MHT80 laser distance sensors. These sensors feature a sensing range of 40 mm to 160 mm , with a resolution of $30\text{ }\mu\text{m}$ and a response time of $\leq 660\text{ }\mu\text{s}$ (1500 Hz). These sensors reflect a laser off of a target to measure the distance and are insensitive to the material, color, and brightness. Additionally, the sensing laser is visible during operation, making them ideal for situations where it is advantageous to know the exact location at which a measurement is obtained.

The maximum deflection of the specimen is obtained by placing one of these sensors at midspan. The others can be placed either at various positions along the specimen to determine the relative deflection, or at either end of the square tube in the connection at the load cell end of the apparatus to obtain the degree of fixity.

5.2 Designing the Support Connections

The first step in designing a connection is to identify the forces to be resisted (or to leave free). Regardless of the fixity of the connection, any axial resistance in the connections will

either reduce the load applied to the specimen (at the shaker end) or cause a difference between the applied load and the measured load (at the load cell end). Either of these options is undesirable, indicating that the connections should offer no axial support. At the same time, any misalignment between the two ends will induce additional moments in the specimen, indicating that they should be laterally fixed. Thus, the specimen should be roller-supported at both ends, with all of the axial force at the load cell end being concentrated in the load cell.

For convenience, the proposed design separates the pin-pin and fixed-fixed connections from the elastically-restrained condition. This allows the relatively extreme cases of the classic support conditions to be created with bearings and allows a single elastically restrained connection to cover as broad a range of intermediate fixities as possible.

5.2.1 Determining the Fixity of a Connection

When a moment is applied to a real connection it causes some degree of rotation, which causes a counteracting moment based on the stiffness of the connection. Assuming that the member attached to the connection is supported at the other end, the fixity (ρ , in this case) is obtained from either the ratio of the actual rotation to the pinned-end rotation or the ratio of the actual moment developed in the connection to the fixed-end moment. Mathematically, this is expressed as,

$$\rho = 1 - \frac{\theta}{\theta_p} = 1 + \frac{M}{M_f} \quad (5.1)$$

The pinned-end rotation and fixed-end moment of a member are obtained assuming a uniformly distributed load of intensity w :

$$\theta_p = \frac{wL^3}{24EI} \quad M_f = \frac{wL^2}{12} \quad (5.2)$$

The actual rotation and moment that develop in the support are given by the point where the moment-rotation relationship of the member and the moment-rotation relationship of the connection intersect. The moment-rotation relationship of the member is simply a straight line connecting the fixed-end moment to the pinned-end rotation, or,

$$M = \frac{wL^2}{12} - \frac{2EI}{L}\theta \quad (5.3)$$

The moment-rotation relationship of the connection usually has an inconvenient shape including a mobilization region, a linear elastic region, and a non-linear plastic region, meaning that the stiffness varies with the applied load. As the connections used in the apparatus were specifically designed to minimize the mobilization region and to avoid entering the plastic region, they are well approximated as by the stiffness of the linear elastic region (κ):

$$M = \kappa\theta \quad (5.4)$$

This simplification reduces the intersection of the moment-curvature relationships to the much simpler problem of finding the intersection of two straight lines. Equating the moments in each side and rearranging gives the rotation, which can be compared to the pinned-end rotation to obtain the fixity:

$$\rho = 1 - \frac{2EI}{2EI + L\kappa} \quad (5.5)$$

As the specimen is the only member being considered in the apparatus, it is helpful to write this such that it is specific to the specimen, rather than in its general form:

$$\rho = 1 - \frac{2E_s I_s}{2E_s I_s + L_s \kappa} \quad (5.6)$$

5.2.2 Specimen Selection

The experimental apparatus was designed with a nominal 3.18 mm \times 25.4 mm strip of mild steel in mind. Aside from obvious reasons like cost and availability, this section is desirable for its (presumably) relatively low damping and for its weldability. Beyond these basic considerations, this section was also selected because it was significantly thicker than most of the existing experimental record.

The first motive for this decision is that the use of larger sections reduces the effect of manufacturing and construction defects on the behavior of the structure, and produces larger deflections that are easier to observe. The second motive is that while the studies previously cited have a wide variety of length-to-depth ratios (100[16], 130[14], 202[12], 285 and 589[17], 324[13], 400[15], and 862[17] [18]), none of them feature specimens with significant dimension in the direction of buckling. The thickest specimen (used with the elastic foundation) had a varying cross-section ranging from 2.5 mm to 7.5 mm thick, but because it (a) was used with an elastic foundation, (b) had a varying cross-section, and (c) had only a single sample tested, it is not comparable to the majority of other experiments. Similarly, the second-

thickest and third-thickest specimens (with thicknesses of 5 mm and 4.5 mm, respectively) were both built-up sections used in the study of viscoelastic damping, which again leaves the majority of the experimental record restricted to extremely thin sections. If these studies are excluded, the smallest length-to-depth ratio used in the previously cited experiments is 202[12], using a specimen with very similar dimensions to those for which this apparatus is designed.

Two specimens were obtained, as the difference in connections between the classical and elastically-restrained connections prevents the same specimen from being used in both cases. Averaging five points spaced roughly equally along the length of each section with a micrometer suggests that the actual dimensions of the specimens are 2.63 mm \times 25.42 mm. Using these measured dimensions gives the area and area moment of inertia of the specimen as $A_s = 66.85 \text{ mm}^2$ and $I_s = 38.54 \text{ mm}^4$. When used with the pinned, fixed, and elastically-restrained connections, the length between supports was measured to be $L_{pp} = 803 \text{ mm}$, $L_{ff} = 753 \text{ mm}$, and $L_\alpha = 871 \text{ mm}$, respectively.

5.2.3 The Classical Connections

By designing specific connections for the extreme cases of the pinned-pinned and fixed-fixed supports, the elastically-restrained connections can be designed to target an intermediate fixity, giving as broad a range of fixities as possible. It is also helpful to quantify the approximate fixity that can be associated with a connection for which the fixity was not calculated explicitly.

5.2.3.1 Pinned-Pinned Connections

The pin-pin condition is easily created by placing a pair of bearings on each side of the specimen. This arrangement gives a literal roller, while still preventing the specimen from rotating about its major axis.

This is not truly a pinned connection, as the rotation of the specimen causes the opposed bearings to misalign and move outwards. The outward movement generates a force, while the misalignment generates a moment arm, resulting in a small bending moment. The rotational stiffness of the pin-pin connection is,

$$\kappa_p = d_1 \left(\frac{1}{\cos(\theta)} - 1 \right) k_t \quad (5.7)$$




Figure 5.4: Pin-type interchangeable connections for the experimental apparatus (attached)

where d_1 is the distance between the center of the bearings, θ is the rotation at the support, and k_t is the transverse stiffness of the bearings.

While k_t is expected to be relatively large, θ and d_1 are both relatively small, resulting in a fairly minimal moment. Additionally, because the slope of $\cos(\theta)$ is zero at $\theta = 0$, both the stiffness and the slope of the stiffness at $\theta = 0$ are zero, suggesting that the quality of this idealization is about as good as that of the small-angle approximation. This stiffness indicates that the lateral restraint at the supports should be as soft as possible, as an excessively stiff lateral restraint causes the moment to increase more quickly. The bearing pairs should therefore be clamped down just tightly enough to prevent lateral deflection, rotation, and bouncing during the experiment, but no tighter.




Figure 5.5: Plan of the pin-type interchangeable connections

For the intended design, for which the maximum permissible lateral deflection is approximately $L/100$, the pinned-end rotation is expected to be approximately 0.032 rad. This results in a moment-arm of 12.90 μm , suggesting that in the absence of a lateral restraint with truly stupendous stiffness, this idealization is almost perfectly ideal. The actual performance of the connection is likely to be somewhat worse than this suggests, as this does not consider the effects of construction tolerances (which are likely to be greater than the moment arm produced as a result of the rotation) or attaching the specimen to the shaker motor and load cell (both of which require connections with some additional stiffness). Of these two sources of error, the connections to the shaker motor and load cell are expected to provide greater error than the construction tolerances, but as the pin itself is very nearly ideal, they can be considered in isolation without introducing substantial error into the calculations (which is helpful, as the lateral stiffness of the support is neither particularly easy to obtain nor consistent between configurations).

5.2.3.2 Fixed-Fixed Connections

The fixed-fixed condition is only slightly more complex and is created by placing three pinned connections in rapid succession, such that a rotation at the first pin necessitates corresponding opposing rotations in each subsequent pin. As the distance between each of



Figure 5.6: Fixed-type interchangeable connections for the experimental apparatus (not attached)

these pins is much shorter than the free span of the specimen they are much more effective in generating moments opposing this rotation, suppressing most of the rotation at the first pin. This is less ideal than the pin-pin connection, but still very close to the ideal behavior. For a given rotation at the first set of bearings, the stiffness provided by the second can be estimated by considering the length of the specimen between the two pins as a beam that is fixed at one end (the first set of bearings) and subject to a point load at the other (the second set of bearings). Noting that the deflection in this case is simple the distance between the two bearings (d_2) multiplied by the rotation at the fixed end (θ), and that the resulting moment (M) is the length of the beam multiplied by the magnitude of the point load,

$$\Delta = \frac{PL^3}{3EI} \rightarrow \theta L = \frac{ML^2}{3EI} \quad (5.8)$$




Figure 5.7: Plan of the fixed-type interchangeable connections

For the fixed-fixed connection shown in Figure 5.6 ($L = d_2$, $E = E_s$, and $I = I_s$), rearranging this to give the moment as a function of θ , from which the rotational stiffness for the fixed connection (κ_f) is easily extracted:

$$\kappa_f = \frac{3E_s I_s}{d_2} \quad (5.9)$$

Substituting the relevant variables ($L_s = 753$ mm, $d_2 = 25$ mm) for the fixed-fixed case into equation 5.6 gives a fixity of,

$$\rho = 1 - \frac{2d_2}{2d_2 + 3L_s} = 0.978 \quad (5.10)$$

which is very nearly equal to full fixity ($\rho = 0.978$).

5.2.4 The Elastically Restrained Support Conditions

Figure 5.8 depicts the designed connection, with the function equivalent overlaid in Figure ???. The end of the specimen is welded to a small steel plate, which is bolted firmly to a pivot constructed from square tube, forming a fixed connection. The pivot bears axially on two cantilevered steel springs, which it must bend to move. Any bending moment applied to the connection will act to rotate the square tube, which will push the ends of the can-



Figure 5.8: Photograph of the constructed connection (not attached)

tilevered springs in opposite directions, generating a moment proportional to the rotation of the connection. By varying the length of the cantilevered steel spring the stiffness can be controlled, allowing a range of stiffnesses to be produced.

The steel springs are made from steel strapping with a cross-section of $0.774 \text{ mm} \times 31.70 \text{ mm}$, giving them $A_k = 24.54 \text{ mm}^2$ and $I_k = 1.225 \text{ mm}^4$. The pivot is only mostly square ($25.49 \text{ mm} \times 25.29 \text{ mm}$), with a wall thickness of 1.177 mm , giving it $A_p = 114.0 \text{ mm}^2$ and $I_p = 11.1 \times 10^3 \text{ mm}^4$.

5.2.4.1 Axial Stiffnesses

The axial load applied by the shaker motor is resisted by a combination of the steel springs, the specimen, and the load cell. By converting each component into its equivalent spring,




Figure 5.9: Plan of the elastic interchangeable connections

the proportion of the applied load carried by each component can be determined, and the degree of support identified.

Steel Springs Each spring is clamped between steel angles and is fitted to the pivot with a set of matching notches. The pivot is then bolted to a plate that has been welded to the specimen. While intended to provide a rotation proportional to the applied moment, this arrangement also provided an axial deflection proportional to the applied axial load. The steel springs, resist the applied axial load through bending, with the applied force transferred from the center of the pivot out to each end of the pivot, and from there back to the base of the spring. This is easily idealized as a pair of end-to-end cantilever beams, with the equivalent stiffness combined in series.

Rearranging the equation for the deflection of a cantilevered beam gives the equivalent stiffness:

$$\delta = \frac{PL^3}{3EI} \longrightarrow k_{eq} = \frac{P}{\delta} = \frac{3EI}{L^3} \quad (5.11)$$

For the pivot, this gives an equivalent stiffness of,

$$k_{eq} = \frac{3(200 \times 10^3 \text{ N mm}^{-2})(11.1 \times 10^3 \text{ mm}^4)}{(40 \text{ mm})^3} = 104.1 \times 10^3 \text{ N mm}^{-1}$$

The stiffness of each spring depends on its length and will be stiffest when the length is shortest (around $L_k = 10 \text{ mm}$):

$$k_{eq} = \frac{3(200 \times 10^3 \text{ N mm}^{-2})(1.225 \text{ mm}^4)}{(10 \text{ mm})^3} = 735.0 \text{ N mm}^{-1}$$

Combining the two in series gives the equivalent spring stiffness of each half of each lateral support:

$$k_{eq} = \left(\frac{1}{104.1 \times 10^3 \text{ N mm}^{-1}} + \frac{1}{91.88 \text{ N mm}^{-1}} \right)^{-1} = 729.8 \text{ N mm}^{-1}$$

Specimen The axial stiffness of the test specimen can be computed directly from Hooke's law, with the caveat that this represents an idealized estimate of the axial stiffness:

$$k_{eq} = \frac{EA}{L} \quad (5.12)$$

The true stiffness will be somewhat lower than suggested by equation 5.12, as the axial force will induce a destabilizing moment which will reduce the stiffness as the deflection increases. As the length of the section also decreases as the deflection increases, this indicates that the axial stiffness of the section is a non-linear function of deflection. In the interest of simplicity, and noting that the change in stiffness is relatively small over the deflections of interest, the stiffness suggested by Hooke's law will be assumed to be reasonably accurate.

This gives the specimen an axial stiffness of,

$$k_{eq} = \frac{(200 \times 10^3 \text{ N mm}^{-2})(66.85 \text{ mm}^2)}{(803 \text{ mm})} = 16.65 \times 10^3 \text{ N mm}^{-1}$$

Load Cell The stiffness of the load cell results from a combination of the axial stiffness of the load cell and the bending stiffness of its support.

The stiffness of the load cell cannot be easily calculated due to its shape, but the approximate magnitude of the stiffness can be estimated. The load cell is approximately one-eighth the length of the member, with (on average) roughly twelve times the cross-sectional area. Knowing that aluminum has a little more than one-third of the stiffness steel has, the load cell should have around $30\times$ the stiffness of the specimen.

The load cell is supported by a short length of steel channel that is simply supported at either end. Rearranging the equation for the deflection of a simply supported beam with a concentrated point load at center-span gives the equivalent spring stiffness:

$$\delta = \frac{PL^3}{48EI} \longrightarrow k_{eq} = \frac{48EI}{L^3} \quad (5.13)$$

For the C100×9 steel channel ($EI = 336 \times 10^9 \text{ N mm}^2$) used to support the load cell, the stiffness will be lowest when the springs are longest. When the springs are approximately 100 mm long, this gives the steel channel an approximate length of $L = 380 \text{ mm}$ in the limiting case. This assumed length is very conservative, as the stiffness of both the springs and the channel are correlated with the length of the springs, but is the most convenient way to illustrate the vast difference in stiffness between the two sections. With this assumed length, the stiffness of the channel is,

$$k_{eq} = \frac{48(336 \times 10^9 \text{ N mm}^2)}{(380 \text{ mm})^3} = 293.9 \times 10^3 \text{ N mm}^{-1}$$

5.2.4.2 Transverse Stiffnesses

While there are no external lateral loads applied to the specimen, transverse internal forces appear throughout the system. For the connection to perform adequately, the transverse support conditions must also match the intended design conditions.

Lateral forces are resisted by some combination of the stinger, steel springs, and the specimen (in the case where the transverse forces differ between supports. The load cell is not expected to provide any significant lateral resistance, as it requires a few millimeters of movement to begin providing resistance.

Stinger The stinger can be thought of as a cylindrical beam with a diameter of 2.38 mm (giving $A = 4.449 \text{ mm}^2$ and $I = 1.575 \text{ mm}^4$) and length $L = 120 \text{ mm}$ that is fixed against rotation at either end. In this case, the stiffness is based on the maximum deflection of a

beam that is fixed at one end, and free but guided at the other:

$$\delta = \frac{PL^3}{12EI} \longrightarrow k_{eq} = \frac{P}{\delta} = \frac{12EI}{L^3} \quad (5.14)$$

This gives the stinger a lateral stiffness of,

$$k_{eq} = \frac{12(200 \times 10^3 \text{ N mm}^{-2})(1.575 \text{ mm}^4)}{(120 \text{ mm})^3} = 2.188 \times 10^3 \text{ N mm}^{-1}$$

Steel Springs For an axial load applied to the specimen (and, by association, the 1-in square tube), the steel springs resist deformation through axial compression on only one side. The stiffness is therefore obtained using Hooke's law (equation 5.12), and is lowest when the length of spring is maximized ($L_k = 100 \text{ mm}$):

$$k_{eq} = \frac{(200 \times 10^3 \text{ N mm}^{-2})(24.54 \text{ mm}^2)}{(100 \text{ mm})} = 49.08 \times 10^3 \text{ N mm}^{-1}$$

Specimen The lateral stiffness of the specimen is found the same way as for the stinger (equation 5.14), and is maximized in the fixed-fixed condition ($L_s = 753 \text{ mm}$):

$$k_{eq} = \frac{12(200 \times 10^3 \text{ N mm}^{-2})(66.85 \text{ mm}^4)}{(753 \text{ mm})^3} = 0.3758 \text{ N mm}^{-1}$$

Note that this stiffness only applies to differential forces, which are expected to be small provided that the end conditions are symmetric.

5.2.4.3 Connection Classification

Having calculated the stiffness associated with each component, the end connections can now be classified. At both ends of the specimen, the lateral resistance is almost entirely dominated by the stiffness of the steel springs (with 99.99% of the total stiffness), which are easily stiff enough to constitute lateral supports. At the loaded end, axial forces transfer primarily into the specimen (98.91%), suggesting no axial restraint. At the unloaded end, the axial forces in the specimen transfer almost exclusively into the load cell (99.94%), suggesting that the specimen is axially restrained.

Setting aside the elastic restraints, the specimen is therefore pinned at the unloaded end, with an axial roller at the loaded end.

5.2.5 Rotational Stiffness of the Elastic Connections

The rotational stiffness (κ) of the connections can be measured directly with relatively little difficulty. As the modal function (equation 2.16) is dependent on α , which is a measure of the degree of fixity or relative rotational stiffness, instead of κ , which is a measure of the absolute rotational stiffness, it is necessary to convert between the two. This conversion is complicated somewhat by the fact that the two are not linearly related.

The mechanism by which the connection resists rotation is depicted in Figure ???. The fixity is based on the point where the moment-rotation lines of the member and the connection intersect. The moment-rotation line of the member is easily determined as it connects the fixed-end moment and the pinned-end rotation in a straight line. The line associated with the connection is more difficult to determine, but assuming that the connection operates exclusively in the linear-elastic region simplifies it to a straight line passing through the origin with a slope equal to the rotational stiffness. Both moment-rotation lines are limited by the maximum moment and rotation each component can resist without failure, and if the intersection falls outside these boundaries the connection will fail.

5.2.5.1 Connection Stiffness

The stiffness of the connection is determined from a combination of the bending stiffness of the springs and the length of the moment arm. Beginning with the deflection at the free end of a spring (which acts as a cantilevered beam) and noting that the deflection must be equal to the product of the connection rotation and the length of the pivot arm, the force in each half of the connection can be converted to an equivalent moment,

$$M = \frac{3E_k I_k L_p^2}{L_k^3} \theta \quad (5.15)$$

from which the stiffness can be extracted:

$$\kappa_k = \frac{3E_k I_k L_p^2}{L_k^3} \quad (5.16)$$

Substituting this stiffness into equation 5.6 gives the fixity associated with this connection. Note that because the rotational stiffness κ_k is associated with only one half of the

connection, it must be doubled before being substituted in:

$$\rho = 1 - \frac{1}{1 + 3 \frac{L_p^2 L_s}{L_k^3} \frac{E_k I_k}{E_s I_s}} \quad (5.17)$$

The connection cannot practically be made shorter than about 10 mm or longer than about 100 mm ($10 \text{ mm} \leq L_k \leq 100 \text{ mm}$) due to space constraints, which suggests that a reasonable range of expected fixities is $0.118 \leq \rho \leq 0.993$ ($1.197 \leq \alpha \leq 1.995$).

5.2.5.2 Strength and Ductility Limits

For the specimen, the limiting conditions are based on the fixed-end moments that develop in the section. These limiting moments will depend on the fixity of the supports but will be greatest when the connection is fully fixed. Assuming that the elastically restrained connections are capable of achieving full fixity, this moment can be obtained from the second derivative of the shape function (equation 2.3). As the shape function is normalized, and the limiting moment is proportional to the maximum deflection required, this moment must then be multiplied by the maximum deflection. This maximum moment occurs at both the ends and at midspan, unlike in the case of a fixed-fixed beam subject to external loading (where the supports experience the larger moment).

ASTM A6-09[24] gives factory tolerances of $L/1000$ for steel sections, with $L/500$ permitted in the out-of-plane direction for sections less than 200 mm wide and stricter tolerances permitted for sections intended to be used as columns. In a similar vein, CSA S16-14[25] requires all sections to have deviations less than $L/500$. At the other end of the spectrum, the maximum deflections permitted by CSA O86-14[26] for wood structures prior to serviceability failure range from $L/180$ to $L/360$, suggesting that it is reasonable to expect the initial straightness to be significantly higher than these values, even in wood.

In the interest of simplicity, values of $L/1000$ for the initial imperfection and $L/100$ for the deflection at failure are convenient. $L/1000$ represents a reasonable guess at a value that meets manufacturing criteria with a comfortable safety margin. $L/100$ represents a reasonable guess at a value that represents a reasonable level of certainty of failure. While these are not likely to be perfectly accurate for any specific application, they are about as representative as it is possible to get when considering the wide range of structural members in use.

For this maximum deflection, the strength and ductility limits of the connection are,

$$\theta \leq \frac{\pi}{100} \quad M \leq \frac{\pi^2 EI}{50L} \quad (5.18)$$

For the selected specimen, this gives a maximum expected fixed-end moment of 1747 N mm. Each spring will need to resist half this amount, or 21.84 N after converting the moment to an equivalent point load.

Given the specimen has a moment resistance of 8718 N mm, the strength of the specimen will not limit the connection, but the strength of the springs is more difficult to determine. When the length of the springs is very short, the moment in the springs will be small because there is no moment arm over which for it to develop. At the same time, when the length is very long, the moment in the springs will also be small because the stiffness is so low. The most limiting length will be somewhere in between these two extremes, where the moment arm is long enough to be significant, but the spring is still stiff enough to develop substantial forces.

Multiplying the relative rotation (the second term in equation 5.17) by the maximum rotation (equation 5.18) gives the maximum rotation associated with a particular spring length. Substituting the result into the moment-rotation relationship (equation 5.15) converts this rotation to the maximum moment in the connection, which can be converted to the maximum moment in the spring (M_k) by dividing half this moment by the length of the pivot arm and multiplying by the length of the spring:

$$M_k = \frac{3\pi}{200} \times \frac{E_k I_k L_p L_k}{L_k^3 + 3L_p^2 L_s \frac{E_k I_k}{E_s I_s}} \quad (5.19)$$

This function reaches its maximum value around $M_k = 93.84$ N mm at $L_k = 40.50$ mm. Like the specimen, this is significantly less than the 1108 N mm moment resistance of each spring, suggesting that the connections will not limit the response unless the deflection is much greater than the intended $L/100$. This provides a comfortable factor of safety (approximately 5 for the specimen and 12 for the connections) against any additional forces that may occur. Additionally, the large factors of safety ensure that the behavior is firmly restricted to the elastic region.

5.3 Automating the Experiment

Automating the experimental procedure (shown in Figure 5.10) can be largely broken into four tasks. The most fundamental of these is to actually run the test (ie. excite the specimen) and to collect the data. To ensure that the results are consistent, the apparatus must also be able to detect instability and stop the test. This functionality has the added benefit of preventing the apparatus from bludgeoning an expensive sensor to death (likely damaging the specimen in the process). Lastly, as the volume of data collected is very large, it is helpful to be able to process the data such that it does not require an excessive quantity of hard drive space.

With each of these tasks, the primary obstacle was to permit a single computer to control each component simultaneously. Owing to the large quantity of data being processed, the automation program was written in Julia.

5.3.1 Exciting the Specimen

Specimen excitation is typically performed using a signal generator or digital-to-analog converter (DAC). As the available equipment in the lab did not permit any sort of programmable control, it was not possible to use any of the existing equipment for this purpose. Fortunately, most computers come with a high-accuracy signal generator in the form of a 3.5 mm audio output. Thus, an audio signal was used to control the motor excitation.

Most audio outputs consist of a 16-bit DAC referenced to a timing crystal with a throughput of 44.1ksp/s. There is remarkably little difference between various audio DACs, as their specifications are largely dictated by the performance of the human ear. For example, the sampling frequency is slightly more than twice the highest frequency the human ear is typically capable of detecting, and thus is the lowest (and therefore least expensive) sample rate that satisfies the Nyquist-Shannon sampling theorem. The primary drawback to this arrangement is that most audio outputs are AC-coupled, leading to a fall-off in output somewhere below the limit of human hearing. This attenuation typically takes the form of a first-order high-pass filter, as the performance is not particularly critical and it is achieved with a single capacitor. Of the three audio outputs tested, two (one desktop and one laptop) were well approximated by a low-pass filter with a central frequency of around 3 Hz to 4 Hz. The third was a USB audio card, for which attenuation was not apparent even at 0.2 Hz. The power amplifier intended for the experiment is capable of applying up to $20\times$ gain to an

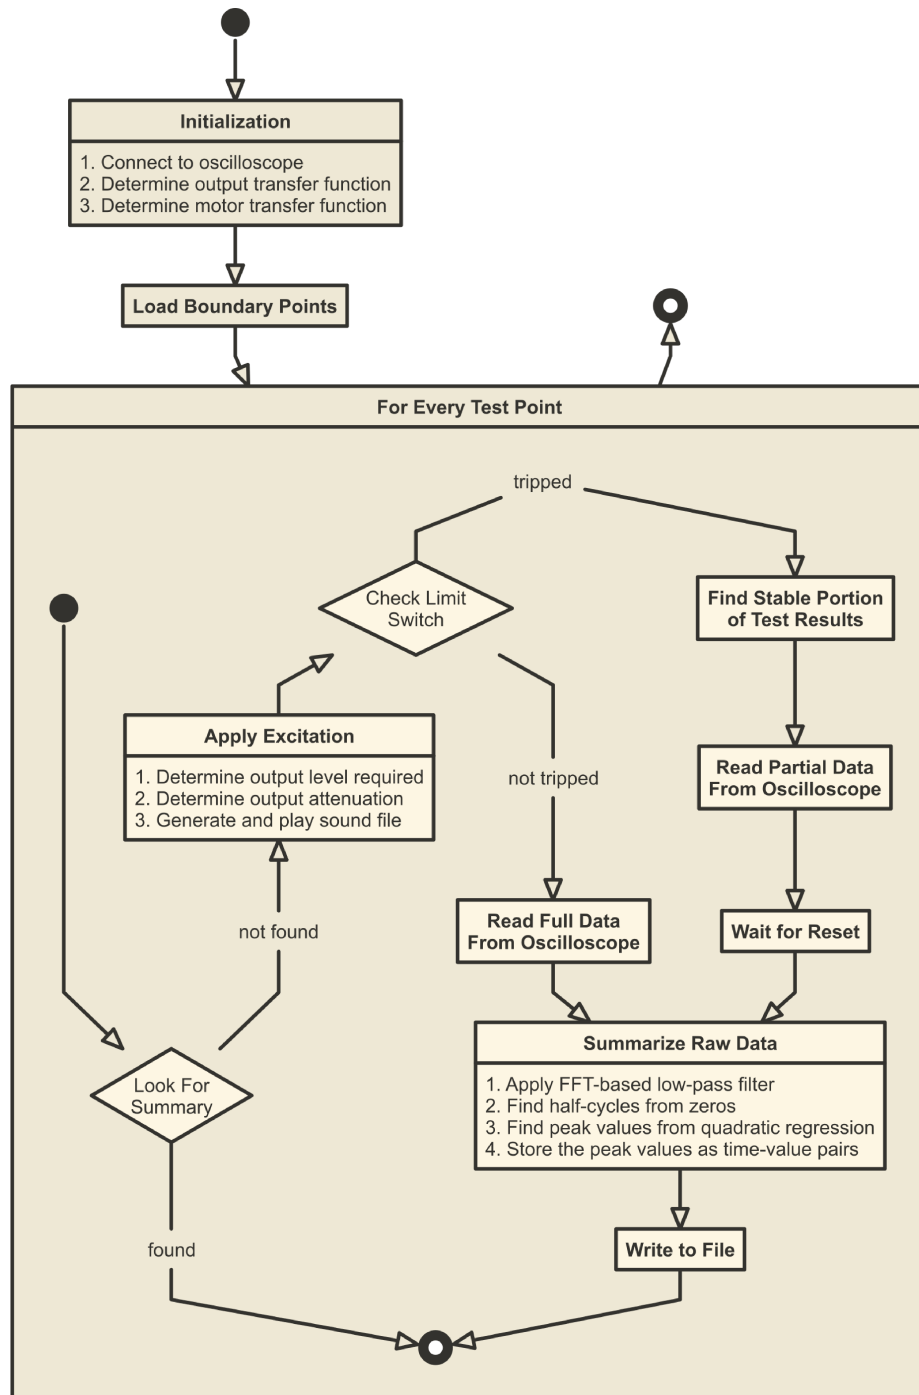


Figure 5.10: Experimental procedure

input signal, so this limitation was considered to be non-critical and can be avoided through the use of a simple amplifier circuit, if strictly necessary.

As a general rule, audio DACs are very precise in terms of both time and amplitude. Because the human ear is so sensitive to frequencies, audio DACs are generally referenced to timing crystals, which typically have relative errors in their central frequency of less than 100ppm (and as low as 1ppm for temperature compensated crystals[27]). While the period of a crystal oscillator does show some variability between cycles, the period of the specimen is so large relative to the period of each crystal oscillation that averaging essentially eliminates the effect of this variability. On the amplitude front, audio DACs typically have a precision of 16-bits per channel. While it is certainly possible to produce measurements and signals with greater than 16-bit precision, doing so typically requires very particular attention to circuit design, wire length and routing, proper grounding, and other such measures to eliminate electromagnetic noise in the associated circuit. As the laboratory in which the experimental apparatus is located contains at least seven poorly shielded desktops in addition to all of the other laboratory equipment and personal laptops, it is not reasonable to expect that such measures can be fully implemented. The 16-bit accuracy of the audio DAC therefore likely represents the practical limit for this application.

5.3.1.1 Excitation Procedure

When an excitation is requested, the automation program first determines the required excitation level. From the audio transfer function and the motor transfer function (discussed in greater detail in section 5.4), the intended force can be converted to the required excitation amplitude, and clipped if necessary.

The program then calculates a sine wave with the intended amplitude, frequency, and duration. This waveform fades in over the first 0.1s and out over the last 0.1s, as the specimen was observed to suffer from significant starting and stopping shocks during excitation. As these shocks tended to dramatically increase the deflection of the specimen, they tended to obscure the exponential response, which is not desirable. The resulting waveform is then written to a `.wav` file and played using the `SimpleAudio` python library, which was called using the `PyCall.jl` Julia library. This decision was made because Julia does not have strong support for media interaction relative to other languages, and as a garbage-collected language the zero-latency required could not be guaranteed inside Julia. The `SimpleAudio` library was selected for its simplicity and convenience rather than for any inherent merit, as it is just a convenient wrapper to a C library.

5.3.2 Detecting Instability and Aborting Tests

Two limit switches were used to detect instability in either direction. This does not represent a true measure of instability but instead detects whether the deflection has increased to a particular extent. Instability is more difficult to detect in real-time, as the definition of instability requires exponential growth, which requires real-time detection of peak deflections and curve-fitting to those peak deflections. As a result, this approach will indicate instability before actual instability occurs, but only where the deflection is large enough to likely result in failure. The circuit designed for this purpose is shown in Figure 5.11.




Figure 5.11: Circuit diagram for the deflection limiter

Like the load cell amplifier circuit, this circuit was designed for simplicity and can be used with essentially any microcontroller and any inexpensive MOSFET. The most significant restriction on the components is that the MOSFET must have a gate threshold voltage that is greater than the maximum negative voltage of the audio signal and less than the difference between the supply voltage and the maximum positive audio signal. If the gate threshold voltage is either too small or too large, the microcontroller will not be able to drive the gate voltage far enough in either direction to block or enable, respectively, the transmission of the audio signal to the power supply. The capacitors exist purely because during experimentation they were found to improve the quality of the transferred signal. Their size is not indicated, but they are indicated as a back-to-back pair of polarized capacitors in the same manner (and for the same reasons) as in the load cell amplifier circuit.

The `detect` pin is configured as an input with the pull-up resistor enabled (if the target microcontroller does not have pull-up resistors, a discrete resistor connected to the positive voltage supply achieves the same results), and the `enable` pin is configured as an output. At startup the enable pin is high and the capacitors are empty, creating a positive gate-source voltage and causing the transistor to conduct. The audio signal is then passed directly to

the output, which sends the excitation signal to the power amplifier. This remains the case unless one of the switches is tripped.

The microcontroller polls continuously to determine whether the `detect` pin has been pulled low. When this happens, it switches the `enable` pin low, which causes the transistor to act like an open circuit, blocking the excitation signal and stopping the test. It then counts upwards to some pre-defined constant, whereupon it switches the `enable` pin high again, and continues to poll the `detect` pin. If the `detect` pin is pulled low at any point during this count, the count resets.

The period for which the excitation is pulled low can be selected arbitrarily by changing the constant to which the microcontroller counts but must be at least long enough that the remainder of the test is blocked and at least long enough that the vibration has sufficient time to decay before the beginning of the next test. As the `enable` pin is connected to the oscilloscope, it is easy enough to find the number corresponding to the desired time by measuring the exact time resulting from a particular number in the oscilloscope. For the experiments performed, a value of 1,150,000 was found to produce a delay of 6.58s, which was deemed adequate.

5.3.3 Retrieving the Data

Again, none of the laboratory equipment available for data collection was in any way programmable with the sole exception of a Rigol DS1054Z 50 MHz oscilloscope, which uses the standard controls for programmable instruments (SCPI) standard. Aside from the obvious advantage of actually being usable for automation, this oscilloscope offers the advantages of taking measurements at 1Gsps, including a programmable gain amplifier (PGA), and having precisely the correct amount of internal noise for oversampling. It also includes the notable disadvantages of offering only four channels and measuring signals at only 8-bit precision. Fortunately, four channels are sufficient for automation (barely), and the low accuracy can be addressed with a combination of amplification and oversampling.

The `Instruments.jl` library, which provides a wrapper to a VISA library (which can be thought of as a low-level driver used for communicating with laboratory instruments), was used to communicate with the oscilloscope over a direct ethernet connection. The exact details of the SCIP communication are not necessary to discuss here, except to indicate that it provides full configuration control for the oscilloscope. In addition to allowing the automation program to ensure that the scope is correctly configured before testing, this

allows the program to do things like dynamically configure the full-scale range of the load cell sensors so that the full 8-bits of accuracy can be used for every measurement. The primary limitation of this approach is its speed, taking approximately 12s to retrieve measurements from a single sensor.

5.4 Data Processing

Before discussing the data processing procedure, it is important to note that there are three types of data that require processing. Each type of data has different processing requirements, which are easiest to discuss separately, but all of them use the same summarization method to reduce the volume and improve the accuracy of the data. Thus, it is easiest to first discuss the summarization procedure, and then discuss the procedure used for each type of data, along with its particular requirements and intended outcome.

5.4.1 Summarizing the Data

The data was summarized by oversampling the data in the frequency domain and recording the maximum value and location of each half-oscillation. The only limitations of this summarization process are that the frequency of the response must be known, and it only produces useful summaries of signals that are well approximated by the extreme values and zeros.

An example of the original and reconstructed curves for a free response is shown in Figure 5.12. Note that the signal reconstruction used here is simply the straight-line connection between peaks and zeros, and as such represents a minimum-effort reconstruction. This was considered acceptable with the simple justification that the zeros are not currently in use in the analysis as they do not directly indicate the amplitude. If the zeros were to be used in calculating the amplitude (for example, by taking the derivative of the signal and using the derivative to estimate the amplitude) a more robust interpolation mechanism is recommended. In any case, the reconstruction represents the signal very accurately, even with a minimum-effort reconstruction.

5.4.1.1 Filtering the Data

Because the lab is electromagnetically noisy, all of the data collected contained some quantity of noise. This was most noticeable in the load cell signal, but was common to all signals and




Figure 5.12: Comparison of the filtered signal and straight-line-connection of signal summary

could not be eliminated through the practical EMI suppression techniques available (using batteries to supply power to the sensors and routing cables through metal channels as far from power lines as possible). As the other steps in processing the data were found to be sensitive to noise, the remaining noise had to be eliminated through filtering.

While several different filters were experimented with, the best performance was using a low-pass filter based on the fast Fourier transform (FFT) of the data. The FFT was computed using the `FFTW.jl` library, which provides a Julia wrapper to a highly optimized C library. The low-pass filtering was then performed by zeroing the entire spectrum above twice the target frequency and reconstructing the signal using the inverse FFT.

5.4.1.2 Oversampling in the Frequency Domain

Because the oscilloscope collects a massive number of relatively low-accuracy points, it is highly desirable to reduce the volume of data and improve the accuracy of that data. This

can be accomplished through oversampling, which is the averaging of multiple measurements to obtain a more accurate estimation of the mean. As a general rule, for every quadrupling of the number of measurements averaged, the resulting accuracy is doubled.

While oversampling is typically performed by taking the mean of a set of measurements, it can be thought of more generally as determining the value of a specific point by filtering a larger number of points to improve accuracy. Using this definition of oversampling, the overall accuracy can be estimated from the number of points used to calculate the value of the filtered signal. In the case of an averaging performed in the frequency domain, the number of points contributing to each term is based on the highest frequency eliminated from the signal. If only the highest frequency found by the FFT is eliminated, the number of points averaged is two, while if the lowest frequency (one cycle over the entire period) is eliminated, the number of points averaged is equal to the number of points in the sample (leaving only the zero-frequency average value). The number of points averaged is therefore easily determined from the low-pass filter performed, since the highest frequency component in the filtered signal is known to be twice the fundamental frequency of the raw signal. For a given sample rate (s) and central frequency (f_0), the increase in accuracy (b_Δ , measured in bits) is,

$$b_\Delta = \log_4 \left(\frac{s}{2f_0} \right) \quad (5.20)$$

Noting that the base measurement accuracy and sample rate of the oscilloscope (in the configuration used) are 8-bits and 250ksps, respectively, the final accuracy (b) of the measurements is very closely approximated by,

$$b = 16.5 - \log_4(f_0) \quad (5.21)$$

For the highest expected frequencies of around 50 Hz with fixed supports and 25 Hz with pinned supports, this suggests a minimum accuracy of about 13.5- to 14.0-bits for any given value. Knowing that the response during each half-cycle is roughly sinusoidal, the peak of each half-oscillation can be found by simply locating the maximum value during that half-cycle.

5.4.2 Processing the Transfer Functions

The objective of the transfer functions is to establish the actual output relative to the theoretical output. The audio transfer function captures the degree of attenuation introduced

by the AC coupling in the audio DAC, which acts like a high-pass filter. The shaker transfer function captures the reduction in force caused by the inertial forces in the shaker armature, which acts like a low-pass filter.

When determining these functions it is necessary to measure the actual output amplitude generated by an input signal of known frequency and amplitude. For the audio transfer function, the audio output was measured by connecting it directly to the oscilloscope, while the shaker transfer function was measured by connecting the audio output to the power amplifier and the shaker directly to the load cell. As the transfer function in both cases is a function in the frequency domain, the excitation was held constant and the frequency varied to determine the effect of the excitation frequency on the amplitude. This frequency sweep was conducted from 200 Hz down to just under 0.2 Hz, with each successive frequency being 90 % of the previous one, giving 67 test frequencies with their corresponding amplitude.

This data was then curve-fit to find a suitable estimation function. Note that the estimation function is not sensitive to the excitation amplitude, but rather indicates the scale of the signal produced from a maximum-output excitation.

5.4.2.1 The Audio Transfer Function

At each frequency, the maximum possible signal was generated by the audio output (around $\pm 1.5\text{V}$ for each of the sound cards tested) for six seconds. The measured voltage was read back from the oscilloscope and summarized as a series of peak values in the same manner as previously described. Because only the amplitude of the resulting signal is needed, the frequency data was then eliminated, leaving one amplitude for each half-period of the excitation. A quantile plot of these values showed that while these samples were very nearly normally distributed, signal noise had an adverse effect on signal integrity when the attenuation was high, leading to an unpredictable skew obscuring the actual behavior.

To eliminate these outliers, Chauvenet's criterion was used. Chauvenet's criterion uses the mean and standard deviation of a sample to identify the upper and lower limits within which all but one of the data points are expected to be found (or, in other words, in the interval $1/2n \leq P \leq 1 - 1/2n$). Because this is a two-sided test, half of this unexplained point is located beyond each boundary, and points evenly spaced at quantiles from a normal distribution will all fall within the suggested limits. This process was repeated until all the points fit within the suggested limits.

Note that Chauvenet's criterion cannot necessarily be used in this manner in all cases, as the detection of outliers is as much art as science in that the appropriate methods depend

a great deal on what is being studied, what is known about the expected behavior, and the overall importance of data integrity. The goal in this case was to find the amplitude at which the probability density function is greatest, and the overall integrity of the data was of no concern. At the same time, it was also observed that there was a tendency for the summary function to produce a small percentage of highly inaccurate values that do not represent the signal amplitude. This is an ideal application for the repeated use of Chauvenet's criterion as there were (unlike with most statistical analyses) no concerns about alterations to the base data or distribution. Once the outliers were eliminated, the median value was taken as the representative amplitude.

Two transfer functions were considered for the audio output. The first was a constant output level (assuming no AC coupling) of the form $V_{out} = C_1 V_{sig}$, and the second was an exponential decay of the form $V_{out} = C_2 - e^{C_3 V_{sig} \nu}$. The first function is easily obtained by averaging the output at every frequency. The second function is more challenging in that finding a solution using regression requires the logarithm of negative numbers, which is not convenient. As audio outputs are intended to cover at least the frequencies in the range of human hearing (20 Hz to 20 kHz), everything above 50 Hz was expected to remain constant regardless of the frequency applied. This allows C_2 to be calculated by taking the mean of the amplitudes measured at frequencies above 50 Hz, leaving only C_3 as an unknown. The value of C_3 was then calculated numerically to six decimal places so that it minimized the residual errors. Comparing the residual errors produced by each function, the function with the lower residual error was selected.

These functions were selected after collecting the attenuation data and were chosen for their ability to model the observed results very closely.

5.4.2.2 The Shaker Transfer Function

The shaker transfer function was obtained in the same manner as the audio transfer function, except that instead of connecting the audio output directly to the oscilloscope, the shaker is connected directly to the load cell, which is then connected to the oscilloscope in the usual manner. This allows the output force to be measured directly rather than through the specimen, which would both introduce additional inertial forces and limit the range of frequencies that can be measured without causing instability.

As with the audio transfer function, the modeling functions were not chosen before the attenuation could be measured and observed. Before this transfer function could be determined the power supply was damaged, preventing further testing. The cause of the damage

remains unclear, and the lead times for repair or replacement were found to be excessive, preventing further testing in the time frame permitted for this thesis. As the transfer function is dependent on both the shaker motor and the power supply, the nature of the power supply has a material impact on the transfer function. For example, an audio amplifier will introduce significant AC coupling and therefore low-frequency attenuation, while a servo drive in the current-proportional mode should produce a constant signal.

5.4.3 Processing the Free Response

The primary obstacle to processing the free response of the specimen is that the frequency is not known and must be determined. Once this is done, the data can be summarized in the usual fashion, and this summary can be used to determine the rate at which the amplitude decays. The only two items of particular interest from the free response are the natural frequency of the system $f_0 = \omega_0/2\pi$, and its associated damping β_0 .

5.4.3.1 Determining the Natural Frequency

As the natural frequency of the system is expected to be approximately 10 Hz, the signal was first filtered at a central frequency of 100 Hz using the method described previously. This eliminates any potentially problematic high-frequency noise, allowing the central frequency to be determined. The central frequency is then estimated by performing an FFT on the six seconds of data following the release of the specimen.

The primary limitation of the FFT in this application is that its accuracy is limited to the nearest multiple of the sample period (± 0.167 Hz in this case). Once the approximate central frequency is known, a more precise estimate can be made using the discrete Fourier transform (DFT), which is dramatically slower than the FFT, but can be performed with arbitrary precision. The natural frequency of the system (f_0) is then estimated using the DFT, testing multiple points over the error range of the DFT, resulting in a proportional increase in precision. Figure 5.13 shows the (normalized) FFT near the central frequency and the corresponding (normalized) DFT obtained in the region near the central frequency. The DFT is very close to the FFT but has some slight differences.

5.4.3.2 Determining the Associated Damping

Once the summary of the free response has been obtained, the damping coefficient associated with the natural frequency can be determined from the amplitude of the signal. The




Figure 5.13: Comparison of the FFT and DFT in the region of the central frequency measured for the pinned specimen

amplitude as a function of time is easily obtained by taking the absolute values of the signal at each point in the summary and eliminating the zeros. Due to asymmetries in the system, the resulting amplitude has small irregularities between points and requires smoothing.

Unlike with other filters, the free response does not benefit from the use of the FFT-based filters, as it forces the beginning and end of the signal to the same value, which seriously affects most of the exponential decay period. Instead, the standard low-pass filter is used, with a central frequency of $1/10^{\text{th}}$ the central frequency. An undesirable side effect of this filter is that it both zeros the initial value and introduces a phase delay, requiring that the signal be reversed before filtering and shifted afterward. Figure 5.14 shows the effect of the filtering after restoring the signal to its original phase and direction. Note that to illustrate more clearly the effect of the filtering, the y-axis in this figure is logarithmic.

Once the signal is smoothed, the exponential decay rate can be determined by calculating the rate of decay between points. Assuming that the decay takes the form $\Delta = e^{-C_1 t}$, the




Figure 5.14: Effect of filtering on the amplitude of the free response of the pinned specimen

exponential decay rate between two points is given by,

$$C_1 = -\frac{\ln(\Delta_1/\Delta_2)}{t_1 - t_2} \quad (5.22)$$

where Δ_1 and Δ_2 are the maximum specimen deflections associated with the first and second points, respectively, and t_1 and t_2 are the corresponding times.

Note that this shows the rate of exponential decay in the section, not β_0 (which is the rate of exponential decay divided by the natural angular velocity of the system):

$$\beta_0 = -\frac{\ln(\Delta_1/\Delta_2)}{\omega_0(t_1 - t_2)} \quad (5.23)$$

The points chosen significantly affect the outcome, as does whether or not the amplitude is smoothed, as shown in Figure 5.15. Because the exponential decay rate is particularly sensitive to noise, both the filtered signal and the wider averaging are helpful in illuminating

the behavior. In addition to the effects of averaging, Figure 5.15 also illustrates that the damping coefficient varies with the maximum deflection of the specimen.




Figure 5.15: Effect of filtering on the amplitude decay rate over time for the free response of the pinned specimen

In the first zone (up to about $t = 4.5$ s) the damping is relatively constant, with $\beta \approx 0.02$. This zone is easily identified as describing the behavior of the column for large deflections. The second zone is shorter (up to around $t = 5.5$ s), during which β_0 reaches approximately twice this value before dropping to a little more than half this value. After the second zone, the damping is once again relatively constant at $\beta_0 \approx 0.01$. Given the variability between regions, the value of β_0 depends on how it is calculated, with the value decreasing as the period over which it is averaged decreases. Ideally, β_0 is calculated over the area that most accurately represents the system being modeled.

From the previous estimates of the minimum and maximum deflections in section 5.2.5.2, the 753 mm pinned specimen will be essentially at rest once the deflection is less than $L/1000 = 0.753$ mm, and instability will be recorded at a deflection of $L/100 = 7.53$ mm. Consulting

Figure 5.14, this suggests that the time-frame during which the amplitude of the free response of the specimen matches that of the behavior to be modeled is approximately $3.8 \text{ s} \leq t \leq 5.2 \text{ s}$, covering approximately the first two zones. The third zone, then, covers the time in which the specimen is essentially at rest, which explains why the measured β_0 is so noisy in this zone.

With the first and third zone identified, the remaining zones appear to be associated with the connection. While in theory all of the connections developed for the apparatus are intended to be purely elastic, in practice there is always some inelasticity in any real connection. If the deflections corresponding to the second zone $\Delta \leq 2.5 \text{ mm} \approx L/300$ are used to determine the equivalent pinned-end angles, the second zone corresponds to angles below $\theta \approx \pi/300 = 0.6 \text{ deg}$. Given the very small rotation at the support in the second zone, the greater damping is very likely a result of the fixed work required to engage and disengage the connection. At large deflections (such as the first zone) this work is not observable because it is very small relative to the energy in the system. At very small deflections (such as the third zone), this work is not performed because the deflection is too small to mobilize the connection. Between these two values, the work done on the connection represents a sizeable proportion of the energy lost during the oscillations.

The most appropriate region to use when calculating the damping coefficient is therefore the first region. Interestingly, this does significantly limit the number of oscillations that can be observed before the damping coefficient begins to increase, and requires the entire trailing portion of the free response to be excluded from the estimation. Fortunately, the damping coefficient calculated from this period is very consistent.

5.4.4 Processing the Experimental Results

The experimental results are primarily complicated by the fact that the shape of the response is unknown, and that part of the response may be unstable. As far as obstacles go, this is much less problematic than those encountered in the previous two sections and is largely solved with the limit switch described earlier. Because the limit switch is connected to the oscilloscope, the stable portion of the response can be obtained easily by simply reading the signal from the limit switch. Once the stable portion is identified, the data can be processed in the same manner as if the entire response is stable. As such, both stable and unstable responses will be discussed together, though only the stable portion of the response is considered.

The primary objective of the experimental process is to identify the stability boundaries. This can be done from stable or unstable results, though the potential problems differ in each case. In practice, the stability regions will be estimated using both the stable and unstable responses.

5.4.4.1 Predicting the Stability Boundaries From Stable Responses

It is difficult to determine the stability boundaries from stable responses because of their tendency to oscillate and because the unstable response does not always appear gradually. While the solution to both of these concerns is quantity, some context is required before discussing the solutions.

In theory, the stability boundaries represent the point where the energy removed through damping (whether due to internal friction or the difference between the excitation frequency and the natural frequency) is precisely equal to the energy added from the excitation. In this idealization, the response either increases or decreases with time, and the amplitude of the oscillations at $t = \infty$ are either zero or infinite. In practice, as the capacity of the column is approached, the oscillations grow in magnitude, and when the column is excited at the stability boundaries the magnitude of the oscillations is infinite. This is most easily pictured by comparing the dynamic response to the buckling behavior of an imperfect column.

When an imperfect column is subjected to an axial load it induces a moment proportional to the column's current deflection. This moment then causes an additional deflection, which increases the moment induced by the axial load, which repeats the process infinitely. This relationship forms a power series, from which the classic equation for the capacity of a beam-column is derived.

In the case of the dynamic column, the axial load is smaller but repeated, with stability hinging on the ability of the column to shed excess energy through internal friction or destructive interference. In this case, the initial excitation causes an oscillation, which sheds some portion of the added energy. Whatever portion of the added energy could not be disposed of then serves to increase the deflection, which increases the energy added by the next excitation. This process continues in the same manner as for the static column, except that where the reinforcing relationship for the static column is between the axial load and the elastic strain energy in the system, the dynamic column also includes the kinetic energy.

In both cases there exists a limiting energy capacity where the minimum-energy resistance mechanism switches from axial to bending, and as the energy level approaches this upper limit, the deflection increases because the column needs to put the energy somewhere.

Expressed mathematically, the energy in either case can be expressed as an infinite series,

$$E_{\infty} = E_0 (1 + E_{\delta} + E_{\delta}^2 + \dots) = \frac{E_0}{1 - E_{\delta}} \quad (5.24)$$


In the most general sense, the change in the energy of the system can be expressed in the form,

$$E_{\infty} = \frac{E_0}{1 - E_{\delta}} \quad (5.25)$$


where E_{∞} is the ultimate energy, E_0 is the initial energy, and E_{δ} is the proportion of the total energy that is preserved between cycles.

In the theoretical idealized model, the initial energy is zero, so this equation always predicts a zero-amplitude oscillation prior to instability. With the imperfect column, however, the oscillations grow before stabilizing at some magnified value. This is well illustrated by the 19.85 Hz and 21.5 Hz curves in Figure 5.16, which shows the reaction of the specimen to a 0.25V signal. This magnification can be used to measure how close the system is to the stability boundaries by using equation 5.25 to solve for E_{δ} . The objective for the stable regions is therefore to use the ratio of the initial and ultimate deflections to estimate how quickly the stability boundary is approached, without ever actually reaching the stability boundary.

Oscillatory Regions The oscillatory behavior near the stability boundaries is well documented in the literature and the 20.0 Hz and 20.5 Hz curves in Figure 5.16 are a good, which show the beginning of the oscillatory behavior. The oscillations in amplitude in these curves is quite brief, as the limit switches were not yet installed when these curves were obtained and the test was stopped prematurely to prevent the possibility of damage to the sensors or specimen, but the behavior was found to be persistent during other tests. This behavior presents a problem in that it is difficult to estimate the degree to which the scale of the oscillations has changed when the amplitude of the oscillations is still changing. The primary counter to this issue is that provided the oscillation of the amplitude can be detected, the oscillations in amplitude can be eliminated. This detection can occur in a single specimen (by using a large number of cycles and a long excitation duration) or regions of tests can be used to identify oscillatory regions, but provided there is a sufficient volume of data the oscillatory behavior can be eliminated, leaving only the average amplitude.



(a) Midspan deflection of the specimen



(b) Magnitude of the oscillations

Figure 5.16: Specimen reaction to 0.25 V pre-gain excitation

Sudden Transitions More problematic than the oscillatory regions are the sudden transitions. This was most clearly observed when approaching the first stability region from the low-frequency side, and can be seen by comparing the difference between the 19.85 Hz and 20.0 Hz curves in Figure 5.16 to the difference between the 20.5 Hz and 21.5 Hz curves. The maximum deflection of the 20.0 Hz response is about $23\times$ that of the 19.85 Hz response, suggesting that the maximum amplitude is increasing over this period at a rate of approximately $150\times$ per Hz, while the 20.5 Hz and 21.5 Hz responses differ by only a factor of approximately 5.5, or $5.5\times$ per Hz as they are separated by exactly 1 Hz. With the exponential response increasing this quickly, it is difficult to obtain enough points on the transition to do any sort of curve fitting. On the other hand, because the transition in these cases is so sudden, it can easily be separated into stable and unstable regions, and so provided a sufficient quantity of points is obtained, the result is largely the same.

5.4.4.2 Predicting the Stability Boundaries From Unstable Responses

The stability boundaries are more easily predicted from the unstable responses in that the unstable responses always exhibit a period of exponential growth. By extrapolating the exponential growth rates, the stability boundaries can be located with relative ease. As with the stable responses, quantity is of great assistance in making these predictions.

5.5 Experimental Error

Every experiment has some difference between the expected and observed outcomes. This error includes anything that could materially affect the measured results, including the physical construction, the environmental conditions, and the sensors used. While some portion of this error is due to known influences and can be corrected for (construction tolerances, variability in material stiffness, etc.), there is always some portion of the error that cannot be identified, and must be determined experimentally. While every effort has been made throughout this thesis to determine the theoretical accuracy and to correct for known errors, it is not possible to estimate the experimental error without actually running the experiments. As a result, the experimental error associated with the apparatus could not be estimated for two reasons.

The first and most obvious reason that the experimental error cannot be estimated at the present time is that the experimental apparatus is missing a power supply. As a result, it is not currently possible to run experiments with the apparatus, which in turn means that

there is no good experimental data with which to estimate the accuracy of the apparatus. While intermediate or proof-of-concept results were obtained, they do not include all of the information required to estimate the experimental error. For example, the responses in Figure 5.16a indicate that an instability region exists in the region of 20 Hz, but because these tests were run before the load cell was properly amplified or isolated from the power lines, and because the output voltage sent to the shaker motor was not measured, it is impossible to tie the results to any analytical or true solution from which to determine the associated error.

The second reason is less obvious but equally significant. During the process of developing the experimental apparatus it became very apparent that the lab space available for the apparatus had a great deal of electromagnetic interference (EMI), and that this introduced significant noise into the measurements. The most obvious sources of this noise were from the nearby power lines, the power amplifier, and the shaker motor itself, which all transfer some noise to both the sensors and the oscilloscope, leading to a loss of accuracy or an increase in the experimental error. As the shaker draws a considerable quantity of power from the amplifier, which draws an even greater quantity of power from the power lines, it is not possible to measure the effect EMI has on the experimental readings without actually powering up the electronics and driving the shaker motor. Further, because the noise is associated with the power drawn through each device, this error is likely to vary with the axial excitation, requiring a more complex estimate than a single measurement.

As a result, while the apparatus is proven in principle, it was not possible to estimate the error in the experimental results both because no final results could be obtained, and because the intermediate or proof-of-concept results cannot be used to determine the experimental error.

5.6 Observed Specimen-Connection Coupling

A significant obstacle to obtaining experimental verification of the elastically restrained connections is that the functional diagrams in chapter 1 assume that the supports have no interaction with the column. This is, of course, untrue, as any support condition will have an associated mass and stiffness which affects the natural frequency and dynamic behavior of the column. This can be neglected in the case of the traditional supports for various reasons, but cannot be ignored in the case of the elastic restraints.

Where no resistance is applied (against axial movement for a roller and rotational movement for a pin), the support undergoes no movement and applies no force, which results in no work performed. Where a support is fixed (lateral fixity for a roller, translational fixity for a pin, and full fixity for a fixed connection), the support develops forces but undergoes no movement, which again results in no work being performed. For a partial restraint (in this case the elastic rotation of the supports), the connection develops force and undergoes movement, which cannot be done without some work being performed on the connection. A summary of these conditions is shown in Table 5.1.

Table 5.1: Work in a connection by degree of fixity

Connection Type	Force	Deflection	Work Performed?
Free	No	No	No
Fixed	Yes	No	No
Elastic	Yes	Yes	Yes

To accurately model the connection, the work performed on the connection must be taken into account. This is most clearly observed by comparing the observed results from a partially restrained connection with those from the pin-pin connection.

For a spring length of $L_k = 59$ mm, the fixity was expected to be $\alpha = 1.392$ ($\rho = 0.273$). Measurements of the deflection at midspan and the rotation at the support for these tests suggest that the actual degree of fixity is approximately two-thirds this stiffness at $\rho = 0.189$ ($\alpha = 1.282$), suggesting that the actual range of stiffnesses provided by the connection is somewhat lower than estimated. The natural frequency in this configuration was measured to be 9.37 Hz, slightly lower than the 10.76 Hz observed in the pin-pin case. While a decrease is expected on account of the increased length of the specimen in the elastically restrained condition, after accounting for both the change in the length of the specimen and the fixity, the expected natural frequency is significantly higher than the observed frequency at 12.58 Hz.

This discrepancy in the natural frequency arises due to a combination of the connection mass and the energy stored in the springs. When the specimen is in the at-rest state, the springs have some non-zero deflection and associated energy storage. When the specimen is in the in-motion state, the connections have some non-zero rotational velocity, which also results in an associated energy storage. If the connections do not store any energy they will have no impact on the natural frequency of the specimen. On the other hand, if

they store energy more effectively than the specimen, they will largely dictate the natural frequency instead of the specimen. The degree of influence the connections have on the natural frequency of the system can be estimated by determining the proportion of the total energy stored in the specimen and the springs.

5.6.1 Relative Bending Energy

The bending energy (or, more formally, the elastic strain energy) associated with the moment in a member is given by,

$$E_b = \frac{1}{EI} \int_0^L M^2 dx = EI \int_0^L (\delta'')^2 dx \quad (5.26)$$

For a given maximum deflection (Δ), the deflection of the specimen is simply the maximum deflection multiplied by the shape function (equation 2.3). Substituting in this deflection gives the bending energy in the specimen:

$$E_{b_s} = \frac{\Delta^2 \alpha^3 \pi^3 E_s I_s (\sin(\alpha\pi) + \alpha\pi)}{2L_s^3 (1 - \cos(\frac{\alpha\pi}{2}))^2} \quad (5.27)$$

The bending energy in the spring is less straightforward to obtain, as it requires some manipulation to express in terms of the maximum deflection of the specimen. The deflection of a cantilever beam is given by,

$$\delta = \frac{P}{6E_k I_k} (2L_k^3 - 3L_k^2 x + x^3) \quad (5.28)$$

where x is measured from the free end of the beam towards the fixed end.

Substituting this equation into equation 5.26 gives the bending energy stored in a cantilever beam. The energy in a single spring (E_{b_k}) is therefore,

$$E_{b_k} = \frac{P^2 L_k^3}{3E_k I_k} \quad (5.29)$$

It is not convenient to express this energy as a function of P because P is not easily compared to Δ . Expressing P as a function of the end displacement expresses the bending

energy in terms than can be easily converted to the connection rotation:

$$E_{b_k} = \frac{3E_k I_k L_p^2}{L_k^3} \theta^2 \quad (5.30)$$

Where θ is the rotation of the specimen at the connection (either $x = 0$ or $x = L_s$).

As θ can be expressed as a function of Δ by taking the first derivative of the deflected shape of the specimen (equation 2.3), substituting in the end rotation of the specimen expresses the bending energy of each spring in the same terms as used for the energy in the specimen. As there are two springs in each connection, doubling this gives the energy associated with each connection:

$$E_{b_c} = \frac{6E_k I_k L_p^2 \Delta^2 \alpha^2 \pi^2 \sin^2\left(\frac{\alpha\pi}{2}\right)}{L_k^3 L_s^2 \left(1 - \cos\left(\frac{\alpha\pi}{2}\right)\right)^2} \quad (5.31)$$

The relative efficacy of the springs is then the ratio of the elastic strain energy in the connections to the elastic strain energy of the specimen, which is easily obtained:

$$\hat{E}_b = \frac{2E_{b_c}}{E_{b_s}} = \frac{E_k I_k L_p^2 L_s}{E_s I_s L_k^3} \frac{24 \sin^2\left(\frac{\alpha\pi}{2}\right)}{\alpha\pi (\sin(\alpha\pi) + \alpha\pi)} \quad (5.32)$$

For the 59 mm connection used here, the four springs will collectively hold approximately 23.00 % as much elastic strain energy as the specimen will.

5.6.2 Relative Kinetic Energy

The specimen used for the pinned supports has a total mass of 464.9 g giving it (after accounting for holes) a relative density of 7.631 kg L⁻¹. Assuming the relative density of each specimen is the same, the 850 mm length of the elastically restrained specimen between supports accounts for 433.6 g of the 663.7 g weight of the final assembly. Therefore, the weld-steel, mounting plates, bolts, and square tubes at each end each have a mass of 115.1 g. Of the connection weight, the bolts and weld steel make up approximately 32.0 g and the square tube makes up approximately 83.1 g. Assuming an average distance from the point of rotation of 10 mm for the bolts and 22.5 mm for the square tube, the moment of inertia of the connection is approximately $I_r = 45.27$ kg mm. The weight of the springs will vary with their length, but the unit weight of the springs is 0.1914 g mm⁻¹, giving the springs measured in this example a mass of 11.29 g each.

As with the elastic strain energy in the previous section, the rotational kinetic energy of the supports can be compared to the kinetic energy of the specimen to estimate the degree to which the weight of the connections may affect the system.

The kinetic energy of the specimen is purely translational and can be obtained from the common definition of kinetic energy. The velocity of the specimen at any point is equal to the deflection, δ , multiplied by the angular frequency, ν . Because the deflection varies along the length of the member, the velocity does too, and the kinetic energy is given by,

$$E_v = \frac{m\nu^2}{2} \int_0^L \delta^2 dx \quad (5.33)$$

Taking δ from the shape function (equation 2.3) gives the energy in the specimen for any maximum deflection Δ :

$$E_{v_s} = \frac{m_s \nu^2 \Delta^2 \left(2 + \cos(\alpha\pi) - \frac{5}{\alpha\pi} \sin(\alpha\pi) \right)}{4 \left(1 - \cos\left(\frac{\alpha\pi}{2}\right) \right)^2} \quad (5.34)$$

The kinetic energy in each support is comprised of a rotational component (in the square tube) and a translational component (in the springs). As these operate on slightly different principles, they will be considered separately.

The translational kinetic energy of the springs can be written as,

$$E_{v_c} = \frac{11m_k \nu^2 P^2 L_k^6}{2520 E_k^2 I_k^2} \quad (5.35)$$

As with the elastic strain, this can be rewritten in terms of the deflection caused by P , which can then be written as a function of the associated angle:

$$E_{v_c} = \frac{11m_k \nu^2 L_p^2}{280} \theta^2 \quad (5.36)$$

Substituting in the end rotation from the shape function gives the kinetic energy in each spring as a function of the maximum deflection of the specimen (Δ):

$$E_{v_c} = \frac{11m_k \nu^2 L_p^2 \Delta^2 \alpha^2 \pi^2 \sin^2\left(\frac{\alpha\pi}{2}\right)}{280 L_s^2 \left(1 - \cos\left(\frac{\alpha\pi}{2}\right) \right)^2} \quad (5.37)$$

The relative kinetic energy stored in the springs is therefore,

$$\hat{E}_v = \frac{4E_{v_c}}{E_{v_s}} = \frac{m_k L_p^2}{m_s L_s^2} \frac{22\alpha^2 \pi^2 \sin^2\left(\frac{\alpha\pi}{2}\right)}{35\left(2 + \cos(\alpha\pi) - \frac{5}{\alpha\pi} \sin(\alpha\pi)\right)} \quad (5.38)$$

For the connection used here, this suggests that the translational kinetic energy stored in the springs is a mere 160ppm of the kinetic energy in the specimen. This is not unexpected, as both the mass and displacement of the springs are much lower than those of the specimen.

The rotational energy differs in that it is derived from the end rotation rather than the deflection along the length of the specimen, and therefore requires no integral:

$$E_{r_c} = \frac{I_r \nu^2 \Delta^2 \alpha^2 \pi^2 \sin^2\left(\frac{\alpha\pi}{2}\right)}{2L_s^2 \left(1 - \cos\left(\frac{\alpha\pi}{2}\right)\right)^2} \quad (5.39)$$

Dividing E_{r_c} by E_{k_s} gives the relative rotational energy storage of the connections:

$$\hat{E}_r = \frac{2E_{r_c}}{E_{v_s}} = \frac{I_r}{m_s L_s^2} \frac{4\alpha^2 \pi^2 \sin^2\left(\frac{\alpha\pi}{2}\right)}{2 + \cos(\alpha\pi) - \frac{5}{\alpha\pi} \sin(\alpha\pi)} \quad (5.40)$$

The pivot and springs collectively hold only 0.2683% of the total kinetic energy in this case, suggesting that despite containing an impressive proportion of the mass, the rotation in the connections is not sufficient to require a significant quantity of energy.

5.6.3 Reducing the Influence of the Connection

To isolate the stiffness behavior, it is helpful to be able to reduce or eliminate the influences that affect the natural frequency of the specimen. As these influences should be related to the strain and kinetic energy stored in the springs and connection, reducing the share of energy stored in these connected items should improve the isolation. Table 5.2 shows the effect increasing each variable has on the connection length and on each type of relative energy storage, when the spring length (L_k) is adjusted to provide a particular degree of fixity. The intent of this table is to illustrate the impact of various design decisions, given that the connection is intended to cover a range of fixities, and therefore must be adjusted to match a particular target fixity. Once the connection has been constructed, this must be done by varying the connection length.

Table 5.2: Effect of changing independent connection properties while holding α constant by varying L_k

Category	Variable	Effect of Variable Increase			
		L_p	\hat{E}_b	\hat{E}_v	\hat{E}_r
Stiffness	I_k	↗	—	↗	↗
	I_s	↘	—	↘	↘
Length	L_p	↗	—	↗	—
	L_s	↗	—	↘	↘
Inertia	m_c	—	—	↗	—
	I_r	—	—	—	↗
	m_s	—	—	↘	↘

Unfortunately, none of the design characteristics of the connection affect the relative bending energy associated with a particular fixity. While equation 5.32 appears to suggest that the relative bending energy should change with any of L_k , L_p , L_s , I_k , or I_s , these terms also appear in α (see equation 5.17) at the same powers. As a result, any change in one of these variables will affect both the relative bending energy and the fixity, but the change in other variables required to correct the fixity will have the effect of exactly countering the change in relative bending energy. The relative bending energy peaks at $L \approx 59.6$ mm, with an expected relative fixity of approximately $\alpha \approx 1.39$. As the fixity approaches the pin-pin and fixed-fixed cases, the relative bending energy approaches zero, suggesting that this behavior is not isolated to this particular connection, but is simply a property of an elastic connection.

As a result of this observation, two possibilities arise. The first is that an elastic connection must be considered together with its specific connections. The second is that an elastic connection will behave identically provided that its fixity can be determined. Whether these are correct will depend on the solution to a more general equation of motion that considers the connections, and is not derived here.

5.7 Summary

An experimental apparatus was designed to allow experimental verification of the dynamic stability of an axially excited column with elastic restraints. This apparatus was constructed, with preliminary experimental results indicating that partial fixity was successfully achieved. The apparatus was then fully automated with the intent of obtaining large volumes of experimental data and determining the variability in the dynamic stability behavior. Due to a combination of equipment failure and supply chain issues, this data could not be collected prior to the publication of this thesis, but preliminary results revealed significant deviations between the theoretical and actual frequencies. On further investigation it was shown that it is not possible to decouple the column and its elastic restraints for a given degree of fixity, meaning that the influence of the elastic restraints at the supports on the total energy stored in the system is a function of the fixity. As dynamic stability is, in essence, an energy storage and transfer problem, this coupling was identified as the most likely candidate for explaining the discrepancy between the theoretical and observed results.

Chapter 6

Conclusions and Future Work

To avoid confounding the two, the conclusions reached as part of this thesis are presented here separately from the potential research opportunities that were identified in the process. The aim in doing this is to clearly delineate between what is known (conclusions), and what is known to be unknown (future work). For the reader seeking an even more succinct summary, the most significant contributions from this thesis can be grouped into the following three points.

1. Derivation and adjustment of the numerical matrix solution, with the corresponding error estimation functions
2. Design and validation of the automated experimental apparatus
3. Identification of the coupling behavior exhibited by an elastically restrained support, and demonstration that this coupling cannot be eliminated

6.1 Conclusions

In the interest of clarity, the conclusions arising from each chapter are presented in the same order, and with the same title as those chapters.

6.1.1 The Analytical Solution

The equation of motion of an elastically restrained column was derived, and Bolotin's method was used to determine the analytical stability boundaries associated with this equation of motion. The analytical solution is the same as that usually obtained through Bolotin's Method, and is a function of the excitation, μ , and frequency, ϕ , parameters. The primary difference in the elastically restrained case is that the frequency parameter is dependent on

the degree of elastic restraint (α). To improve the usability of the stability diagrams produced, the excitation and frequency parameters were restated relative to the Euler buckling load and natural frequency of the unloaded column.

6.1.2 The Numerical Matrix Solution

A numerical method was then derived from the equation of motion to calculate the state transition matrix associated with an arbitrary excitation function, provided that the column is in the at-rest or in-motion conditions. The derivation reduces to the undamped solution when the damping coefficient is equal to zero, but require adjustments as outlined in Chapter 4 to provide usable accuracy in the damped case. It was determined that the third-order small-angle approximations provide the best balance between computational effort and accuracy.

6.1.3 Comparison of the Analytical Solution and the Numerical Matrix Solution

By comparing the results produced by the numerical method to the analytical solutions, the relative error of the numerical method was determined for both the damped and undamped cases. By using an extremely fine discretization in which the error was expected to be extremely small, the relative error in the analytical method was then estimated. As only the undamped numerical solution trended towards infinite accuracy, only the error in the undamped analytical solution could be estimated. Convenient tables were then generated for two different levels of accuracy that can be used to estimate the maximum allowable excitation by stability region and order of the analytical solution.

Modifications to the Numerical Matrix Solution for a Damped Column The numerical solution was found to be a very poor match for the analytical solution when $\beta_0 \neq 0$. It was discovered that this could be corrected by changing the exponential damping term and subtracting the top-right term of the matrix product. These changes do not have any theoretical justification, but improve the accuracy of the numerical solution to usable levels. The correct response and stability behavior are given in Chapter 4 by equations 4.1, and 4.2, respectively. The corresponding equations in Chapter 3 can be used in the undamped case, but as setting $\beta_0 = 0$ for the undamped cases reduces both equations to the same thing, there is no practical reason to use the uncorrected equations.

Relative Error in the Analytical Solution For any stability region, R , excitation level relative to the buckling load of the column, \hat{P} , and order of the analytical solution, o , the relative error in the undamped analytical solution can be estimated as,

$$\epsilon_{ana} \approx \pi^{2R-5oR^{-3/8}-(-1)^R} \hat{P}^{2.7o-0.6R+0.7(-1)^R+2} \quad (6.1)$$

The relative error in the damped analytical solution could not be estimated because the error in the numerical solution does not approach zero as N increases.

Relative Error in the Numerical Matrix Solution The relative error in the numerical matrix solution for an undamped column is a function of the stability region, R , the excitation level relative to the buckling load of the column, \hat{P} , and the number of segments in the discretization, N . Additionally, when the damping at the natural frequency, β_0 , is non-zero, the relative error is also a function of this damping and the excitation frequency relative to the natural frequency of the unloaded column, $\hat{\nu}$:

$$\epsilon_{num} \approx \frac{2^R \hat{P}^{2R}}{N^2} + \frac{R^6}{N^5} + \frac{60\beta_0^2 - 3\beta_0}{8 \hat{\nu}^2} \quad (6.2)$$

6.1.4 Experimental Observations

An experimental apparatus was designed to provide elastically restrained connections of configurable stiffness. This apparatus was then automated, including data collection and processing, to allow for the collection of large quantities of data. Due to equipment failure, this data was not collected, but the automation was sufficiently advanced to allow the collection of preliminary data to verify the functionality of the procedure. Additionally, as certain aspects of the column do not require excitation (such as measuring the fixity or the free response), some data was produced from which a variety of conclusions could be reached. In the process of performing these tasks, two key observations were made that are worth repeating here.

The Effect of Coupling on the Response From the perspective of experimental validation, the most significant observation is that it does not appear to be possible to design an experimental apparatus such that the end connections are partially fixed, but without materially affecting (at least) the natural frequency of the specimen. While the apparatus

presented here can still be used for experimental verification, future attempts at experimental verification would ideally begin from an equation of motion that accounts for and provides an explanation for the effect of coupling on the response. In the absence of such an equation, the coupling will likely limit the accuracy attainable as the effect on the natural frequency will need to be accounted for in other ways.

Variation of the Damping Coefficient The damping coefficient was observed to vary with deflection, though not in the manner predicted by other studies. In this case the deflections in the column were small enough that substantial stress did not develop at any point in the specimen, preventing the varied damping predicted by other studies from being ruled out, but the evidence produced in these experiments either for or against this possibility is entirely insufficient to arrive at any clear conclusion on the subject. In these experiments the damping coefficient was found to vary significantly as the connection mobilized. When the deflection was large the damping coefficient was relatively stable, but the value of the damping coefficient for small deflections was significantly higher, and for very small deflections (vibrations, really) it was found to be significantly lower. As dynamic stability is generally concerned with predicting excessive increases in deflection, the most appropriate damping coefficient should generally be taken only from the portion of the free response in which the deflection is large enough to stabilize the damping coefficient. If the damping coefficient is calculated without sufficient care, it is likely to include the effects of connection plasticity (leading to a higher estimate for the damping coefficient) and the residual damping region (leading to a lower estimate for the damping coefficient), both of which will introduce significant errors into any subsequent calculations.

6.2 Future Work

As future work is, by definition, not a part of the present work, there are no chapter-headings under which to organize these avenues of inquiry and they are presented here en masse.

The most obvious future work that remains to be done on this subject is those portions of the research objectives that could not be accomplished due to extenuating circumstances. In addition to these, as usual, in the process of obtaining these conclusions, a number of potential avenues of investigation were identified.

Experimental Verification of the Stability Regions As the actual experimental verification was not able to be performed, the most obvious and immediate future work to be completed is to collect the data the experimental apparatus was designed to collect. Once this data has been collected, the natural variability associated with the physical behavior can be estimated, which will provide additional insight into the accuracy of the analytical and numerical solutions.

Deriving an Equation of Coupled Motion As mentioned previously, it seems unlikely that great accuracy will be obtained in future experiments unless the natural frequency associated with a particular connection can be estimated analytically. It is therefore important that a proper equation of motion is developed that can provide this.

Variation of the Damping Coefficient with Internal Stress As the specimen used in these experiments does not develop significant internal stress, it is not clear if the varied stress spoken of by others presents as such in an axially-excited column. This could be investigated in greater detail by either replacing the specimen with one that develops greater stresses during the experiment or by adapting the apparatus such that it is capable of applying a static axial excitation in addition to the dynamic excitation. Applying the static excitation was investigated during this thesis, but the mechanisms used introduced very substantial damping to the system and prevented instability from developing. It was determined that a static excitation cannot be obtained by compressing the specimen, due to the change in length during bending, meaning that the most practical solution is almost certainly to replace the shaker motor with something more suitable. The obvious candidate here would be to use a servo motor and ball screw to allow the static load to be applied directly to the specimen with a single connection, but this could not be pursued as part of this thesis due to the unfortunate fact that time and resources are limited.

Investigating Asymmetric Elastic Restraints As structural members are very commonly constructed in cases where the connections at either end are not symmetric, it would be of great benefit to extend the experimental verification of the elastically restrained connections to the case where the connections at either end of the specimen are not symmetric. This was investigated analytically during this thesis but was scrapped when the scale of the work required to complete the investigation became apparent.

The Dynamic Stability of an Axially Excited Column With Elastic Supports

The natural extension of studying partial rotational fixity is to study partial translational fixity. This would allow a more complete understanding of the dynamic stability behavior, and would likewise benefit greatly from experimental verification.

References

- [1] A. L. Kimball and D. E. Lovell, “Internal friction in solids,” *Physical Review*, vol. 30, p. 948,
- [2] D. Lisitano, “Strain proportional damping in Bernoulli-Euler beam theory,” *Mechanical Systems and Signal Processing*, vol. 145, p. 15, 2020.
- [3] V. Fisikopoulos and L. Peñaranda, “Faster geometric algorithms via dynamic determinant computation,” *Computational Geometry*, vol. 54, pp. 1–16, Apr. 2016.
- [4] W.-C. Xie, *Differential Equations for Engineers*. New York: Cambridge University Press, 2010.
- [5] S. Cui, H. K. Cheong, and H. Hao, “Dynamic buckling of simply supported columns under axial slamming,” *Journal of Engineering Mechanics*, vol. 125, no. 5, pp. 513–520, May 1999.
- [6] S. Kenny, F. Taheri, and N. Pegg, “Experimental investigations on the dynamic plastic buckling of a slender beam subject to axial impact,” *International Journal of Impact Engineering*, vol. 27, no. 1, pp. 1–17, Jan. 2002.
- [7] J. R. Gladden, N. Z. Handzy, A. Belmonte, and E. Villermaux, “Dynamic buckling and fragmentation in brittle rods,” *Physical Review Letters*, vol. 94, no. 3, p. 035 503, Jan. 2005.
- [8] H. E. Lindberg, “Impact buckling of a thin bar,” *Journal of Applied Mechanics*, vol. 32, no. 2, pp. 315–322, Jun. 1965.
- [9] G. R. Abrahamson and J. N. Goodier, “Dynamic flexural buckling of rods within an axial plastic compression wave,” *Journal of Applied Mechanics*, vol. 33, no. 2, pp. 241–247, Jun. 1966.
- [10] N. Willems, “Experimental verification of the dynamic stability of a tangentially loaded cantilever column,” *Journal of Applied Mechanics*, vol. 33, no. 2, pp. 460–461, Jun. 1966.

-
- [11] B. J. Ryu, K. Katayama, and Y. Sugiyama, "Dynamic stability of Timoshenko columns subjected to subtangential forces," *Computers & Structures*, vol. 68, no. 5, pp. 499–512, Sep. 1998.
- [12] J. H. Somerset and R. M. Evan-Iwanowski, "Experiments on parametric instability of columns," in *Proceedings, Second Southeastern Conference on Theoretical and Applied Mechanics*, 1964, pp. 503–525.
- [13] T. Iwatsubo, M. Saigo, and Y. Sugiyama, "Parametric instability of clamped-clamped and clamped-simply supported columns under periodic axial load," *Journal of Sound and Vibration*, vol. 30, no. 1, 65–IN2, Sep. 1973.
- [14] R. Ahuja and R. Duffield, "Parametric instability of variable cross-section beams resting on an elastic foundation," *Journal of Sound and Vibration*, vol. 39, no. 2, pp. 159–174, Mar. 1975.
- [15] I. Svensson, "Dynamic buckling of a beam with transverse constraints," *Nonlinear Dynamics*, vol. 11, no. 4, pp. 315–328, Dec. 1996.
- [16] L. Mohanta, "Dynamic stability of a sandwich beam subjected to parametric excitation," MTech Research, Jun. 2006.
- [17] S. C. Mohanty, "Dynamic stability of beams under parametric excitation," Ph.D. dissertation, Sep. 2005.
- [18] S. Shaik, "Dynamic stability of multilayered beam with bolted joints," MTech, 2007.
- [19] V. V. (V. Bolotin, *The dynamic stability of elastic systems*. San Francisco: Holden-Day, 1964.
- [20] W.-C. Xie, *Dynamic Stability of Structures*. Cambridge: Cambridge University Press, 2010.
- [21] O. Giraldo-Londoño and J. D. Aristizábal-Ochoa, "Dynamic stability of slender columns with semi-rigid connections under periodic axial load: Verification and examples," p. 7,
- [22] M. J. Ward, *Basic Floquet Theory*, Oct. 2019.
- [23] J. Deng, "Analytical and numerical investigations on pillar rockbursts induced by triangular blasting waves," *International Journal of Rock Mechanics and Mining Sciences*, vol. 138, p. 104 518, Feb. 2021.
- [24] ASTM, *ASTM A6-09 Standard Specification for General Requirements for Rolled Structural Steel Bars, Plates, Shapes, and Sheet Piling*, 2010.

-
- [25] CSA, *CSA S16-14 Design of Steel Structures*, 2014.
 - [26] CSA, *CSA O86-14 Engineering Design in Wood*, 2014.
 - [27] Anritsu, *Understanding Frequency Accuracy in Crystal Controlled Instruments Application Note*, 2001.