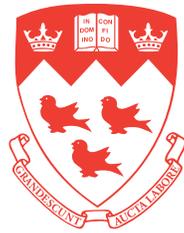


Computational design and simulation of idiophone bars

Douglas Beaton



Computational Acoustic Modeling Laboratory - Music Technology Area
Department of Music Research, Schulich School of Music
McGill University
Montreal, Quebec, Canada

November 2022

A thesis submitted to McGill University in partial
fulfillment of the requirements
of the degree of Doctor of Philosophy

©Douglas Beaton, 2022

Abstract

This thesis develops computational methods for the design and evaluation of idiophone bars, with a focus on marimbas and vibraphones. The methods developed use computer models to enable evaluation of new bar designs and materials, reducing or delaying the need for physical prototypes. The work’s main contributions comprise both design methods and evaluation methods.

Traditional tuning of marimba bars involves carving a “cutaway” or “undercut” on the bottom of the bar to adjust the frequencies of up to three flexural modes of vibration. Other mode types, such as torsional modes, are typically left untuned, with their frequencies left to chance. Makers have complained of timbral problems over a specific range of the marimba keyboard, where these untuned modes may have frequencies close to those of a tuned flexural mode. Anecdotal reports indicate this problem creates a concerning amount of wasted tonewood. Simulation and experimental measurements in this thesis research confirm the proximity of tuned and untuned modal frequencies over this range of notes.

To address problematic untuned modes, this thesis research develops a design methodology to tune idiophone bar flexural and torsional modes concurrently. Tuning is achieved by creating three-dimensional cutaway geometries that vary along both the length and width of the bar (whereas examples in the literature vary only along the length). Bar behaviour is modelled via finite element analysis using three-dimensional solid elements. Cutaway geometry is parametrized in a versatile manner, with an arbitrary number of input parameters. Methods are developed to solve the underdetermined mathematical problem arising when the number of input parameters exceeds the number of tuned modes. Examples of rosewood and aluminum bars tuned using this method are provided. An aluminum prototype bar is fabricated and tested to validate the design methodology.

A second aspect of this thesis research seeks to enable evaluation of proposed bar designs via listening to virtual models of the sound they produce. To achieve this, a hybrid model for idiophone sound radiation is developed. This hybrid model combines a dynamic time history finite element analysis with a modal model, leveraging the advantages of each. The time history model simulates the first ten to fifteen milliseconds of a mallet impact, capturing excitation of the bar and interaction between the bar and mallet head. A modal model

represents the bar's linear response and decay after the mallet impact has concluded. The hybrid model thus constricts the computationally expensive time history model, while alleviating any need for recordings or measurements of prototype bars in fitting parameters of the modal model. All of this is accomplished using the same geometric inputs and definitions as the developed bar design methodology, integrating bar design and evaluation within a single computational framework.

Résumé

Cette thèse développe des méthodes de calcul pour la conception et l'évaluation de barres idiophones, avec un accent mis sur les marimbas et les vibraphones. Les méthodes développées utilisent des modèles informatiques pour permettre l'évaluation de nouvelles conceptions de barres et de nouveaux matériaux, réduisant ou retardant le besoin de prototypes physiques. Les principaux apports de ce travail portent à la fois sur les méthodes de conception et sur les méthodes d'évaluation.

L'accord traditionnel des barres de marimba consiste à tailler un "cutaway" ou un "undercut" au bas de la barre pour ajuster les fréquences jusqu'à trois modes de flexion de vibration. D'autres types de modes, tels que les modes de torsion, ne sont généralement pas réglés, leurs fréquences étant laissées au hasard. Les fabricants se sont plaints de problèmes de timbre sur une plage spécifique du clavier marimba, où ces modes non accordés peuvent avoir des fréquences proches de celles d'un mode de flexion accordé. Des rapports anecdotiques indiquent que ce problème crée un gaspillage non négligeable de bois. La simulation et les mesures expérimentales conduites dans cette recherche de thèse confirment la proximité des fréquences modales accordées et non accordées sur cette gamme de notes.

Pour résoudre les problèmes de modes non accordés, cette recherche de thèse développe une méthodologie de conception pour régler simultanément les modes de flexion et de torsion de la barre idiophone. Le réglage est obtenu en créant des géométries de coupe tridimensionnelles qui varient à la fois sur la longueur et sur la largeur de la barre (alors que les exemples de la littérature ne varient que sur la longueur). Le comportement des barres est modélisé selon une analyse par éléments finis à l'aide d'éléments solides tridimensionnels. La géométrie de coupe est paramétrée de manière polyvalente, avec un nombre arbitraire de paramètres d'entrée. Des méthodes sont développées pour résoudre le problème mathématique sous-déterminé survenant lorsque le nombre de paramètres d'entrée dépasse le nombre de modes accordés. Des exemples de barres en bois de rose et en aluminium accordées selon cette méthode sont fournis. Une barre prototype en aluminium est fabriquée et testée pour valider la méthodologie de conception.

Un deuxième aspect de cette recherche de thèse vise à permettre l'évaluation des conceptions de barres proposées via l'écoute de modèles virtuels du son qu'elles produisent. Pour

ce faire, un modèle hybride de rayonnement sonore idiophone est développé. Ce modèle hybride combine une analyse temporelle dynamique par éléments finis avec un modèle modal, tirant parti des avantages de chacun. Le modèle temporel simule les dix à quinze premières millisecondes d'un impact de maillet, capturant l'excitation de la barre et l'interaction entre la barre et la tête du maillet. Un modèle modal représente la réponse linéaire de la barre et sa décroissance après l'impact du maillet. Le modèle hybride restreint ainsi le coût élevé en calculs du modèle temporel d'éléments finis dynamique, tout en atténuant le besoin d'enregistrements ou de mesures des prototypes de barres lors de l'ajustement des paramètres du modèle modal. Tout cela est accompli en utilisant les mêmes entrées et définitions géométriques que la méthodologie de conception de barres développée, combinant la conception et l'évaluation des barres dans un cadre de calcul unique.

Acknowledgements

First and foremost, I would like to thank my supervisor, Dr. Gary Scavone. His guidance and insights have been an integral part of this journey. Perhaps most appreciated is his ability to help chart a research path, while remaining open to new directions as results unfold.

Thank you to fellow researchers Song Wang and Connor Kemp for helping me find my way around the Computational Acoustics Modeling Laboratory, and teaching me how to operate our measurement equipment. Thanks also to fellow researcher Alberto Acquilino for repeatedly booting my lab computer after countless planned building power outages.

Thank you to Darryl Cameron for ongoing technical support in the Music Technology area. Thanks in particular for helping to build my lab computer and setting up remote access.

Thank you to Dr. Fabrice Marandola for providing valuable insights into percussion performance and marimba quality. Thanks also to Martin Daigle for assisting with access to instruments and even loaning me some mallets.

Thank you to Ramnarine Harihar, Miesam Aghajani and Andy Hofmann for assistance with prototype fabrication, and for answering my endless array of questions.

Thank you to Helene Drouin for helping me navigate the various systems involved in being a graduate student at the Schulich School of Music.

Thank you to Florent Durieux for reviewing my French abstract translation.

Thank you to Drs. Philippe Depalle, Ichiro Fujinaga, Stephen McAdams and Mathias Legrand for serving on my comprehensive exam committee. Thank you to Drs. Péter Rucz, Philippe Depalle and Marcelo Wanderley for serving on my thesis defence committee.

Thanks once again to Dr. Péter Rucz for serving as external thesis reviewer, and to Dr. Philippe Depalle for serving as internal thesis reviewer.

This research was supported by the Vanier Canada Graduate Scholarships, Schulich School of Music Graduate Excellence Award, Engineers Canada–TD Insurance Meloche Monnex Scholarship, and Centre for Interdisciplinary Research in Music Media and Technology Travel Awards. My heartfelt thanks goes to all of these organizations, for alleviating financial concerns and enabling me to focus on my studies.

Finally, thank you to my friends and family for supporting me along a meandering career path through industry and academia. Don't worry: four degrees are surely enough... surely.

Contributions

This thesis, and the research it describes, is original work by the candidate, except for commonly understood and accepted concepts, or where explicit reference to the work of others is made. The dissertation is formatted as a monograph comprised of six chapters and an appendix. It includes content from the following conference and journal publications:

1. D. Beaton and G. Scavone. Measurement-based comparison of marimba bar modal behaviour. In *Proceedings of the International Symposium on Musical Acoustics 2019*, pages 72–77, Detmold, DE, 2019a
2. D. Beaton and G. Scavone. Optimization of marimba bar geometry by 3D finite element analysis. In *Proceedings of the International Symposium on Musical Acoustics 2019*, pages 402–407, Detmold, DE, 2019b
3. D. Beaton and G. Scavone. Tuning idiophone bar torsional modes with three-dimensional cutaway geometries. *The Journal of the Acoustical Society of America*, 149(4):A69, 2021b
4. D. Beaton and G. Scavone. Three-dimensional tuning of idiophone bar modes via finite element analysis. *The Journal of the Acoustical Society of America*, 149(6):3758–3768, 2021a
5. D. Beaton and G. Scavone. Experimental measurements of a prototype vibraphone bar with three-dimensional cutaway geometry. *JASA Express Letters*, 2(8):083201, 2022

The candidate performed all numerical modelling, calculations, coding, experimental measurements and data analysis. The candidate prepared all manuscripts and presentation materials for the conference and journal references cited above.

Supervisor Dr. Gary Scavone reviewed and provided feedback and suggested edits for all manuscripts and presentation materials. He also provided guidance, work space, equipment and software for laboratory measurements.

Contents

1	Introduction and Motivation	1
1.1	Background	1
1.1.1	Marimba History	1
1.1.2	Historical Tuning	2
1.2	Motivation	3
1.2.1	Troublesome Torsional Modes	3
1.2.2	Enabling Design Exploration	4
1.3	Thesis Layout	7
2	Review of Idiophone Computational Modelling Literature	9
2.1	Bar Vibration Modelling	9
2.1.1	Background and History of the Finite Element Method	9
2.1.2	Musical Acoustics Applications of Finite Element Analysis	10
2.1.3	Alternative Modelling Methods	25
2.2	Sound Radiation Modelling	28
2.2.1	Surface Point Source Method	28
2.2.2	Boundary Element Approach	30
2.2.3	Finite Element Approach	30
2.2.4	Scaled Boundary Finite Element Method	33
2.3	Damping Models	33
2.3.1	Rayleigh Damping	34
2.3.2	Generalized Proportional Damping	35
2.4	Tuning Techniques	36
2.4.1	Solving Nonlinear Systems	36
2.4.2	Optimization	37
3	Idiophone Bar Finite Element Models	38
3.1	Programming Packages	38
3.1.1	Pylaster	38
3.1.2	Calculix	39
3.1.3	Code_Aster	39
3.2	Geometry Definition	41
3.2.1	Mesh Definition	41
3.2.2	Cutaway Definition	41
3.2.3	Coordinate Scaling	45
3.3	Material Properties	45
4	Three-Dimensional Tuning of Idiophone Bar Modes	46
4.1	Objective	46
4.2	Method	46
4.2.1	Frequency Targets	46
4.2.2	Tuning Approach	47
4.2.3	Mode Identification	52
4.3	Example Tuned Bars	59
4.3.1	Basic Bars	59
4.3.2	Rosewood Bars	62

4.3.3	Aluminum Bars	66
4.4	Algorithm Performance	68
4.4.1	Tuning Convergence	68
4.4.2	Non-Convergence	70
4.5	Prototype Fabrication	72
4.5.1	Bar Descriptions	72
4.5.2	Model Convergence Test	75
4.5.3	Experimental Modal Analysis and Bar Performance	76
5	Hybrid Modelling of Idiophone Bar Sound Radiation	78
5.1	Objective	78
5.2	Concept	79
5.3	Method	81
5.3.1	Overview	81
5.3.2	Time Domain Modelling	82
5.3.3	Sound Radiation Modelling	91
5.3.4	Mode Selection	92
5.3.5	Numerical Time Integration Compensation	94
5.3.6	Modal Parameter Fitting	99
5.3.7	Signal Crossover	103
5.3.8	Reversing Numerical Time Integration Compensation	105
5.4	Model Parameter Selection	108
5.4.1	Fitting Window Parameters	108
5.5	Example Models	111
5.5.1	Aluminum Bar	111
5.5.2	Rosewood Bar	117
5.6	Fitting Window Parametric Tests	119
5.6.1	Aluminum Test	122
5.6.2	Rosewood Test	124
6	Conclusion and Future Directions	126
6.1	Contributions	127
6.2	Future Directions	128
	References	130
A	Bar Measurements	142
A.1	Marimba Bar Measurements	142
A.1.1	Motivation	142
A.1.2	Scope	143
A.1.3	Method	143
A.1.4	Results	144
A.1.5	Discussion	144
A.2	Additional Bar Measurements and Reanalysis	147
A.2.1	Motivation	147
A.2.2	Scope	147
A.2.3	Method	147
A.2.4	Results	148
A.2.5	Discussion	150

List of Figures

1.1	Photo of the bottom of a Bergerault F#3 marimba bar. The central area where material has been removed is the bar's <i>cutaway</i> . Also visible are two holes drilled across the bar's width to house support cables.	4
1.2	Vibration mode types and naming convention in this thesis. Undeformed bar geometry is outlined in black. Free-free boundary conditions are employed. Mode types: (V) - vertical-flexural mode, (A) - axial mode, (L) - lateral-flexural mode, (T) - torsional mode. Note that the bar cutaway in this example results in coupling between deformation types. Vertical deformation is present in the axial mode, while torsional deformation occurs in the lateral-flexural mode. Figure taken from Beaton and Scavone (2021a).	5
2.1	Marimba bar finite element mesh from Bork et al. (1999). Three distinct mesh densities can be seen along the bar's length.	11
2.2	Wooden bar finite element meshes from Bretos et al. (1999). Three geometry types were considered: a) constant cross-section, b) parabolic undercut and c) rectangular undercut.	12
2.3	Optimized geometries for an aluminum vibraphone bar tuned to 880 Hz from Henrique and Antunes (2003). Geometries were optimized via (a) simulated annealing using element thicknesses as inputs, and (b) sequential quadratic programming using 10 cosinusoidal shape functions. Bar length was a variable in each case (note differing x-axis scales between panels).	14
2.4	Optimized geometries for an aluminum vibraphone bar tuned to 440 Hz from Henrique and Antunes (2003). Flexural mode frequencies were tuned to ratios 1:3:5:7:9. Geometry was optimized via sequential quadratic programming using 10 cosinusoidal shape functions. Bar length was fixed at 350 mm.	15
2.5	Vibraphone bars tuned to musical note A4 and fabricated via CNC mill per Kirkland and Moradi (2016).	16
2.6	Example finite element mesh of an aluminum bar model from Bestle (2017).	17
2.7	Aluminum bar models from Bestle et al. (2017). Each bar is tuned to a fundamental frequency of 440 Hz with differing frequencies of higher-order partials.	18
2.8	Fabricated aluminum vibraphone bar by Bestle (2017). Cutaway geometry was optimized to tune the frequency ratios of the first three vertical-flexural modes. A CNC mill was used to fabricate the precise bar geometry.	19
2.9	Schematic of xylophone sound radiation simulation domain from Rucz et al. (2021). (a) Top view. (b) Side view.	20
2.10	Side view of idiophone bar with parameters defining the simplified undercuts from Soares et al. (2021a). This example has two distinct steps, though generally the model may use any number of steps.	21
2.11	Prototype bars with stepped undercuts fabricated by Soares et al. (2021b). Frequency ratios are indicated for tuned modes: (V1:V2:V3:V4 T1:T2).	21

2.12	Axisymmetric finite element mesh of an English church bell from Perrin et al. (1983). Solid lines indicate undeformed mesh. Dashed lines show modelled results for the six-meridian inextensional radial ring driven mode known as the <i>Tierce</i>	22
2.13	Finite element model results from McLachlan et al. (2012). A) Residual radial stresses in polar coordinates after dimple formation. B) Residual tangential stresses in polar coordinates after dimple formation. C) Effect of adding a 20 mm diameter dimple of various depths on the frequency of the 0,1 and 2,0 modes. Results are shown for both a linear finite element analysis (FEA) model with residual stresses, and a physical prototype (P'TYPE).	24
2.14	Variables in the calculation of sound radiation from a single surface element by Bestle et al. (2017).	29
2.15	Example mesh from a two-dimensional model of an open elliptic amphitheatre from Bermúdez et al. (2008). Amphitheatre boundaries are displayed in white. Light grey triangular elements represent the air. Darker grey square elements define a perfectly matched layer to model radiative boundary conditions.	32
2.16	Conceptual sketch of a two-dimensional unbounded domain meshed with finite elements and infinite elements. Figure from Givoli (2008).	33
2.17	Model of an ancient Chinese bell from Liu et al. (2019a). Both the bell and a portion of the surrounding air are shown. Radiative boundary conditions are imposed at the outer surface of the modelled air.	34
2.18	Plot of damping ratio, ζ , vs. angular frequency, ω , for Rayleigh damping. Contributions of mass proportional and stiffness proportional damping are shown separately.	35
3.1	These panels depict the stages of creating a meshed bar model: (a) a structured mesh of uniform density is generated over a unit cube in natural coordinates (u, v, w) , this example has $(80, 7, 3)$ elements in the (length, width, thickness) directions; (b) cutaway shape is defined by scaling w-coordinates of the nodes according to a thickness function $t(u, v)$; (c) final bar geometry is generated by scaling nodal positions from natural coordinates (u, v, w) to model coordinates (x, y, z) via Equation 3.1.	42
3.2	Example of a bar cutaway defined by cross-section interpolation. Cross-sections are outlined in blue. The red line shows an example of thickness interpolated between cross-sections along the length of the bar. This figure is borrowed from Beaton and Scavone (2019b).	43
3.3	Example of a bar cutaway defined by surface interpolation. This example demonstrates the use of control point mirroring to enforce double symmetry. The blue dots shown in panel (a) serve as the control points for bar tuning. Red dots are mirrored versions of the blue control points. Green dots are mirrored versions of both the blue and red points. Panel (b) shows a more complex geometry that can be attained by adjusting the blue control points and mirroring twice. This example uses Clough-Tocher interpolation, thus the control points lie directly on the cutaway surface.	44
4.1	Example of identifying a vertical-flexural mode via shape recognition based on node groups with low displacement. Figure from Beaton and Scavone (2021a).	54
4.2	Similar but not degenerate mode shapes. These modes are from a rosewood bar model with vertical-flexural modes tuned to the typical 1:4:10 ratios. Coupled lateral and torsional deformations make identifying the modes exceedingly difficult without further information. (a) Mode L2, $f = 2260$ Hz. (b) Mode T3, $f = 2759$ Hz. Figure from Beaton and Scavone (2021a).	56

4.3	Corner nodes used to identify mode shapes, as shown in Soares et al. (2021b). Visible on the bar’s side is the step-wise cutaway definition employed in their paper.	58
4.4	Modal frequency ratios from simulated bars over the lower two octaves of a five-octave marimba. Outer bar dimensions are based on measurements of a Yamaha instrument. This figure is adapted from Beaton and Scavone (2019b).	60
4.5	Example rosewood bar tuned to the typical 1:4:10 ratios using a thin cutaway as its initial input geometry. (a) Initial geometry. (b) Tuned geometry. . . .	61
4.6	Example rosewood bar tuned to the typical 1:4:10 ratios using a moderate cutaway as its initial input geometry. (a) Initial geometry. (b) Tuned geometry.	61
4.7	Example rosewood bar tuned to the typical 1:4:10 ratios using a uniform bar as its initial input geometry. (a) Initial geometry. (b) Tuned geometry. . . .	61
4.8	Tuned bar geometries and frequency ratios for models that tune progressively more modes. Model 1(a) begins by tuning only mode V1. Model 1(b) tunes modes V1 and V2. These additions continue to model 1(f), which tunes all six of the modes plotted above. Table 4.1 outlines the tuning targets of each submodel. Figure borrowed from Beaton and Scavone (2021a).	63
4.9	Tuned geometries for and frequency ratios for various rosewood bar models of the musical note F3. Tuning ratios are indicated for each bar with “-” denoting an untuned mode. Table 4.2 outlines the tuning targets of each example model. Figure borrowed from Beaton and Scavone (2021a).	65
4.10	Tuned geometries for and frequency ratios for various aluminum bar models of the musical note C4. Tuning ratios are indicated for each bar with “-” denoting an untuned mode. Table 4.3 outlines the tuning targets of each example model. Figure borrowed from Beaton and Scavone (2021a).	67
4.11	Cutaway geometries and resulting modal frequencies for the iterations required to tune model 1(e). Iteration 0 is the same initial geometry shown in Figure 4.5(a).	69
4.12	Cutaway geometries and resulting modal frequencies for the iterations required to tune the model for fabrication bar Prototype 2. The initial geometry in Iteration 0 was selected to eliminate sharp edges along the sides of the cutaway.	71
4.13	Images of Prototype 1: (a) bar model meshed with tetrahedral elements, (b) photo of fabricated bar. Figure borrowed from Beaton and Scavone (2022). . . .	73
4.14	Images of Prototype 2: (a) bar model meshed with 20-node hexahedral elements, (b) photo of fabricated bar. Figure borrowed from Beaton and Scavone (2022).	75
5.1	Example output waveform from the hybrid sound radiation model. (a) Full result for a model with a simulation time of 2 seconds. (b) Detailed view of the first 50 milliseconds. The three distinct segments of the model are colour-coded with: blue - time history excitation, orange - crossover, and green - Prony tail.	80
5.2	Mallet impulse force curves. Two curve fits based on data from Bork (1990) and Bestle (2017) are compared with the mallet impact force simulated using contact mechanics. The three curves have been offset along the time axis to roughly align their initial rise in force.	84
5.3	Initial geometry of a mallet impact model for use in Code_Aster. The bar is meshed with 720 20-node hexahedral elements. The mallet head is meshed with 1106 10-node tetrahedral elements.	85
5.4	A view from below an Adams five-octave marimba showing the support cable passing through several bars and resting on cable posts. Rubber wraps are visible that prevent the bars from striking the metal cable posts. The bar F#3 from Figure 5.5 is visible toward the right of the photo.	86

5.5	An F#3 bar from an Adams five-octave marimba used in experimental modal measurements (see Appendix A). Oversized holes for the support cable are visible near either end of the cutaway.	87
5.6	Damping models employed in the dynamic time domain finite element models in this work. Damping ratios vary according to the Rayleigh damping model in Equation 5.3 with coefficients per Table 5.2.	89
5.7	Definition of geometric parameters used in sound radiation modelling by the Surface Point Source Method. Figure adapted from Bestle et al. (2017). . . .	92
5.8	Example frequency compression functions showing the numerical frequencies resulting from different time integration methods. The red line shows the effect of HHT on frequency with $\alpha_h = -0.1$, while the blue line represents Newmark's Method with $\beta_w = 0.25$ and $\gamma_w = 0.5$. Both examples use an input damping ratio, $\zeta = 0.001$ for all frequencies and a sampling rate of 48 kHz.	99
5.9	Examples of total numerical damping ratios resulting from HHT numerical time integration with different values of α_h . The input damping ratios are those from a Rayleigh model using input parameters for aluminum from Table 5.2. Results shown assume a sampling rate of 48 kHz.	100
5.10	Mixing function $m(t)$ from Equation 5.31 along with $1 - m(t)$. The two lines illustrate how $p_t(t)$ and $p_p(t)$ are scaled over the transition zone before being summed to create $p_h(t)$	105
5.11	Hybrid radiation aluminum bar example model meshed geometry. Red circles represent support locations in the finite element model. Origin of the global coordinate system is shown at the corner of the bar's playing surface. The model is comprised of 5,760 8-node hexahedral elements.	113
5.12	Waveforms of radiated sound at the virtual microphone position for the aluminum bar example model. (a) Results produced using finite element dynamic time history analysis throughout the simulation. (b) Results from the hybrid sound radiation model described in this chapter, before applying the damping and frequency corrections of Section 5.3.8.	114
5.13	Frequency content of radiated sound waves shown in Figure 5.12 for the example aluminum bar time history and hybrid models. Hybrid sound radiation model results are shown before applying the damping and frequency corrections of Section 5.3.8.	114
5.14	Frequency content of the aluminum hybrid radiation model, with and without reversing the effects of numerical time integration.	116
5.15	Hybrid radiation rosewood bar example model meshed geometry. Red circles represent support locations in the finite element model. Origin of the global coordinate system is shown at the corner of the bar's playing surface. The model is comprised of 5,760 8-node hexahedral elements.	118
5.16	Waveforms of radiated sound at the virtual microphone position for the rosewood bar example model. (a) Results produced using finite element dynamic time history analysis throughout the simulation. (b) Results from the hybrid sound radiation model described in this chapter, before applying the damping and frequency corrections of Section 5.3.8.	120
5.17	Frequency content of radiated sound waves shown in Figure 5.16 for the example rosewood bar time history and hybrid models. Hybrid sound radiation model results are shown before applying the damping and frequency corrections of Section 5.3.8.	120
5.18	Frequency content of the rosewood hybrid radiation model, with and without removing non-participating modes. Results are shown before applying the damping and frequency corrections of Section 5.3.8.	122

5.19	Fit quality as a function of fitting window start time and duration for the aluminum bar example model in Section 5.5.1. Colour map indicates the fit quality, defined by Equation 5.39, between time history model results and hybrid radiation models with various fitting window parameters. The yellow circle marks the expected suitable parameter values shown in Table 5.5. . . .	123
5.20	Fit quality as a function of fitting window start time and duration for the rosewood bar example model in Section 5.5.2. Colour map indicates the fit quality, defined by Equation 5.39, between time history model results and hybrid radiation models with various fitting window parameters. Range of the colour map is set equal to that of Figure 5.19. The yellow circle marks the expected suitable parameter values shown in Table 5.10.	124
A.1	Measurement setup for experimental modal analysis of marimba bars. Components: 1) Polytec PDV-100 Laser Doppler Vibrometer, 2) Brüel & Kjær Type 8203 Impact Hammer, 3) test bar, 4) foam supports, 5) pendulum, 6) National Instruments USB-4431 Signal Acquisition Board, 7) laptop with MATLAB digital acquisition platform. Figure taken from the poster presented at ISMA 2019.	145
A.2	Measurement results for the 15 measured marimba bars. Bar images are shown approximately to scale. Figure taken from the poster presented at ISMA 2019.	146
A.3	Measured modal frequencies and damping ratios for the rosewood F#3 marimba bars from Figure A.2. The fitted curve shown is that of a Rayleigh damping model with parameters given in Table 5.2.	149
A.4	Modal frequencies and damping ratios measured for the aluminum stock bar. The fitted curve shown is that of a Rayleigh damping model with parameters given in Table 5.2.	149

List of Tables

3.1	Material properties of aluminum and rosewood used in this work. Symbols include: density, ρ , Young's modulus, E , Poisson's ratio, ν , and shear modulus, G . Rosewood properties are adopted from Bork et al. (1999). Aluminum properties are those of the standard 6061 alloy (Battelle Memorial Institute, 2020). This table is borrowed from Beaton and Scavone (2021a).	45
4.1	Target tuning ratios of all model 1 submodels in Figure 4.8. Table adapted from Beaton and Scavone (2021a).	62
4.2	Target tuning ratios of example rosewood bar models in Figure 4.9. Table adapted from Beaton and Scavone (2021a).	64
4.3	Target tuning ratios of example aluminum bar models in Figure 4.10. Table adapted from Beaton and Scavone (2021a).	66
4.4	Prototype bar dimensions.	72
4.5	Prototype bar modal tuning ratios. Untuned modes are indicated by '-'.	72
4.6	Numbers of elements used for Prototype 2 convergence test models. Model X-, Y- and Z-axes correspond to the bar's length, width and thickness directions, respectively.	76
4.7	Measured and modelled modal frequencies for the tuned modes of Prototype 2. Modelled results were produced from Model 6 in Table 4.6 using both Pylaster and Calculix. In all cases modelled frequencies in the two programs agreed to the first decimal place. Measured frequencies were determined using the LSCE method (see Appendix A). Frequency differences are expressed in cents, relative to the modelled values. Measured frequency ratios are calculated relative to the measured frequency of mode V1. Table borrowed from Beaton and Scavone (2022).	76
5.1	Mallet impact force curve parameters.	83
5.2	Rayleigh damping coefficients used for time domain modelling.	89
5.3	Hybrid radiation aluminum bar example model dimensions.	112
5.4	Hybrid radiation aluminum bar example model modal tuning ratios.	112
5.5	Hybrid radiation aluminum bar example model parameters.	112
5.6	Hybrid radiation example model mallet impact and virtual microphone positions. The same positions are used for both the aluminum and rosewood example modes. However, the two example models have different bar dimensions (see Tables 5.3 and 5.8), thus the relative positions of the mallet impact and virtual microphone differ between the models.	113
5.7	Aluminum bar example model, mode frequencies and normalized effective modal mass in the z-direction. Values reported are from modal analysis in Code_Aster.	115
5.8	Hybrid radiation rosewood bar example model dimensions.	117
5.9	Hybrid radiation rosewood bar example model modal tuning ratios.	117
5.10	Hybrid radiation rosewood bar example model parameters.	117

5.11	Rosewood bar example model, mode frequencies and normalized effective modal mass in the z-direction. Values reported are from modal analysis in Code_Aster.	121
A.1	Material properties of aluminum used for bar fabrication compared with published standard values. Density, ρ , was measured for the stock material, as was Young's modulus, E . Poisson's ratio, ν , is taken as the standard published value for both cases. Shear modulus, G , is calculated from E and ν for isotropic materials. Standard aluminum properties are those of the 6061 alloy (Battelle Memorial Institute, 2020).	148

List of Symbols

General

- A** matrices are indicated using boldface capital letters
- A**⁺ ‘+’ superscript indicates the *pseudoinverse* of a matrix
- b** column vectors are indicated using boldface lowercase letters
- \dot{x} overdots indicate derivatives of a quantity with respect to time

Mode Labels

- A1, A2 ... axial mode 1, axial mode 2, ...
- L1, L2 ... lateral flexural mode 1, lateral flexural mode 2, ...
- T1, T2 ... torsional mode 1, torsional mode 2, ...
- V1, V2 ... vertical flexural mode 1, vertical flexural mode 2, ...

Review of Idiophone Computational Modelling Literature (Chapter 2)

Symbols in Chapter 2 are adopted from the reviewed literature. They are defined in, and applicable only to, their respective subsections.

Finite Element Models (Chapter 3)

- ν Poisson’s ratio for an isotropic material
- ρ material density (kg/m³)

- E elastic modulus (GPa) of an isotropic material
- G shear modulus (GPa) of an isotropic material
- L overall length (m) of a bar
- T end thickness (m) of a bar (i.e. thickness outside the cutaway)

$t(u, v)$	thickness function defining the cutaway thickness at natural coordinates (u, v)
u	natural coordinate in the longitudinal (length) direction of the bar, $u \in [0, 1]$
v	natural coordinate in the transverse (width) direction of the bar, $v \in [0, 1]$
W	overall width (m) of a bar
w	natural coordinate in the vertical (thickness) direction of the bar, $w \in [0, 1]$
x	model coordinate (m) in the longitudinal (length) direction of the bar
y	model coordinate (m) in the transverse (width) direction of the bar
z	model coordinate (m) in the vertical (thickness) direction of the bar

Three-Dimensional Tuning of Idiophone Bar Modes (Chapter 4)

ζ	natural coordinate for evaluating finite element shape functions
η	natural coordinate for evaluating finite element shape functions
λ_j	j^{th} eigenvalue of a finite element modal analysis
ξ	natural coordinate for evaluating finite element shape functions
ρ	material density (kg/m^3)
ϕ_j	j^{th} eigenvector of a finite element modal analysis
$\psi(\mathbf{s})$	mode shape displacement vector at point s
$\psi_x(s)$	x-component of mode shape displacement vector at point s
$\psi_y(s)$	y-component of mode shape displacement vector at point s
$\psi_z(s)$	z-component of mode shape displacement vector at point s
\mathbf{A}_e	finite element Jacobian matrix containing partial derivatives of model coordinates with respect to natural coordinates of an element
$b_n(\mathbf{x})$	binary variable used to activate or deactivate the n^{th} input parameter in the vector \mathbf{x}
$\mathbf{D}(\mathbf{x})$	diagonal matrix of binary values used to activate or deactivate input parameters in the vector \mathbf{x} when tuning cutaway geometry
$f_{\text{model},k}(\mathbf{x})$	frequency (Hz) of the k^{th} tuning mode in a model with geometry input vector \mathbf{x}
$f_{\text{target},k}$	target frequency (Hz) of the k^{th} tuning mode
$\mathbf{g}(\mathbf{x})$	vector of tuning error functions for bar geometry based on input parameter vector \mathbf{x}
$g_k(\mathbf{x})$	tuning error function for the k^{th} mode based on input parameter vector \mathbf{x}

i	index variable for iteration number
j	index variable for mode number
$\mathbf{J}(\mathbf{x})$	Jacobian matrix containing partial derivatives of tuning error functions with respect to input parameters in the vector \mathbf{x}
\mathbf{K}	system global stiffness matrix in a finite element model
K	number of modes to be tuned in a given bar
k	index variable for tuning mode number
\mathbf{M}	system global mass matrix in a finite element model
\mathbf{M}_e	element mass matrix in a finite element model
\mathbf{N}	finite element shape function matrix
N	number of input parameters defining a bar's cutaway geometry
n	index variable for individual geometry input parameters
s_1, s_2	corner node positions used to identify mode shapes
T_{max}	maximum allowable value of a geometry input variable
T_{min}	minimum allowable value of a geometry input variable
\mathbf{x}	vector of input design parameters for tuning bar geometry
x_n	the n^{th} input variable defining a bar's geometry

Hybrid Sound Radiation Model (Chapter 5)

α	Rayleigh damping model mass proportional coefficient
α_h	Control parameter for Hilber-Hughes-Taylor numerical time integration
α_n	decay coefficient of mode n in a Prony model
β	Rayleigh damping model stiffness proportional coefficient
β_w	Control parameter in Newmark's Method for numerical time integration
γ_w	Control parameter in Newmark's Method for numerical time integration
ΔT	time step (s) of a digital signal (inverse of sampling frequency)
ζ	damping ratio (proportion of critical damping)
$\bar{\zeta}$	total numerical damping ratio apparent in a system modelled using numerical integration
ζ_n	damping ratio of mode n in a Prony model
Θ_n	angle (radians) between the normal vector of surface element n and a vector from that element to the listener position
λ	eigenvalue of the amplification matrix \mathbf{A}_m
μ	average term in modified log-normal function used to model mallet impact force

ρ	density (kg/m ³) of bar material
ρ_0	density (kg/m ³) of air
σ	standard deviation term in modified log-normal function used to model mallet impact force
ϕ_n	mode shape vector of the n^{th} mode
φ_n	phase angle (radians) of mode n in a Prony model
Ω	incremental angle used in calculating amplification matrix \mathbf{A}_m
A	scaling amplitude factor in modified log-normal function used to model mallet impact force; real part of an amplification matrix principal eigenvalue
A_1	first invariant of amplification matrix \mathbf{A}_m (equal to half the trace of \mathbf{A}_m)
A_2	second invariant of amplification matrix \mathbf{A}_m (equal to the sum of the principal minors of \mathbf{A}_m)
A_3	third invariant of amplification matrix \mathbf{A}_m (the determinant of \mathbf{A}_m)
\mathbf{A}_m	amplification matrix for evaluating numerical integration methods
A_{31}	component of amplification matrix \mathbf{A}_m at position 3,1
A_{32}	component of amplification matrix \mathbf{A}_m at position 3,2
A_{33}	component of amplification matrix \mathbf{A}_m at position 3,3
A_n	amplitude of mode n in a Prony model
B	imaginary part of an amplification matrix principal eigenvalue
b_{fw}	beginning time (s) of hybrid radiation model fitting window
b_{tz}	beginning time (s) of signal crossover transition zone
\mathbf{C}	finite element system damping matrix
c_a	speed (m/s) of sound in air
c_{si}	speed (m/s) of sound in an isotropic material
D	constant used to calculate amplification matrix \mathbf{A}_m
d_{fw}	duration (s) of the hybrid radiation model fitting window
d_{sim}	total duration (s) of a finite element dynamic time history simulation
E_T	elastic modulus (Pa) of an orthotropic bar material in the thickness direction
e_{fw}	end time (s) of hybrid radiation model fitting window
e_{tz}	end time (s) of signal crossover transition zone
f	frequency (Hz)
\bar{f}	numerical frequency (Hz) apparent in a system modelled using numerical integration

\mathbf{f}_h	vector of frequencies (Hz) from a Fourier transform of a hybrid radiation model
\mathbf{f}_h^*	non-uniformly spaced vector of frequencies (Hz) from a Fourier transform of a hybrid radiation model
$\bar{\mathbf{f}}_h$	vector of numerical frequencies (Hz) from a Fourier transform of a hybrid radiation model
$f_m(t)$	mallet impact force (N) at time t
f_n	frequency (Hz) of mode n in a Prony model
f_0	fundamental frequency (Hz) of a bar
G	shear modulus (Pa) of an isotropic material
G_{LT}	shear modulus (Pa) of an orthotropic bar material between the length and thickness directions
G_{WT}	shear modulus (Pa) of an orthotropic bar material between the width and thickness directions
\mathbf{h}	vector of Prony system residues
h_n	n^{th} complex system residue of a Prony model
i	index variable
j	$\sqrt{-1}$
\mathbf{K}	finite element system stiffness matrix
K	total number of points in a discrete time signal
k	index variable for signal discrete time instances
L	length (m) of a bar
\mathbf{M}	finite element system mass matrix
$m(t)$	signal crossover mixing parameter at time t
$m_{\text{eff},n,z}$	effective modal mass of mode n in the global z-direction
$\hat{m}_{\text{eff},n,z}$	normalized effective modal mass of mode n in the global z-direction
$m_{\text{total},z}$	total unrestrained mass of a model in the global z-direction
N	total number of modes used in a Prony model
n	index variable for summation
$\hat{\mathbf{n}}_n$	unit vector normal to surface element n
$p(t)$	sound pressure at time t
$p_n(t)$	sound pressure radiated by surface element n at time t
$p_h(t)$	sound pressure calculated by the hybrid sound radiation model at time t
$p_p(t)$	sound pressure calculated by a Prony model at time t
$p_t(t)$	sound pressure calculated by a finite element dynamic time history model at time t

$\mathbf{r}_{c,i}$	coordinate (m) vector of the i^{th} corner of a bar's playing surface
\mathbf{r}_ℓ	coordinate (m) vector of the listener position
r_n	distance (m) from surface element n to the listener position
S_n	area (m^2) of surface element n
T	thickness (m) of a bar at its end (away from cutaway)
t	time (s; or as noted)
t_{lo}	time (s) at which mallet head lifts off the bar's playing surface following an impact
t_0	reference time in modified log-normal function used to model mallet impact force
t_{pb}	time (s) required for waves to propagate through the bar
t_{pa}	time (s) required for waves to propagate through the air
u	displacement (m) of a single-degree-of-freedom system
\mathbf{u}_z	displacement vector of a finite element model, with all displacements in the global z -direction set equal to one, and all other values set to zero
$\tilde{v}_n(t)$	normal velocity (m/s) of surface element n at time t
$\bar{\mathbf{v}}_n(t)$	velocity (m/s) vector in three dimensions for surface element n
W	width (m) of a bar
\mathbf{X}_h	Fourier transform of hybrid radiation model output signal
$\bar{\mathbf{X}}_h$	Fourier transform of hybrid radiation model constructed with numerical modal frequencies
\mathbf{x}	vector of discrete time signal values
\mathbf{x}_h	hybrid radiation model output signal
$\bar{\mathbf{x}}_h$	hybrid radiation model output signal constructed with numerical modal frequencies $\{\bar{f}_n\}$
\mathbf{Z}	Vandermonde matrix of Prony system poles
z_n	n^{th} complex system pole of a Prony model

Chapter 1

Introduction and Motivation

1.1 Background

1.1.1 Marimba History

Origins of the marimba can be traced back a very long time to what today may be described as “primitive” instruments. In fact, one of the world’s oldest known instruments is a lithophone comprised of stone bars tuned to a type of pentatonic scale (Adcock, 2017). *The New York Times* (1954) described the instrument as a “Stone-Age ‘Marimba’ ”, and reported that prehistorians believe its construction may date back more than 4,000 or 5,000 years. Tool markings on the stones provide evidence of purposeful tuning (Adcock, 2017), as does their pentatonic scale. Clearly the origins of bar percussion instruments are ancient indeed.

MacCallum (1969) provides an excellent history of the marimba in more recent times. Published in 1969, his book provides possible origins of the marimba with wooden bars. There are, however, many such possible origins, and the true provenance of the instrument remains mired in obscurity.

MacCallum later describes the instrument’s proliferation in Guatemala, beginning circa 1894. It was around this time the modern layout, with two rows of bars arranged to resemble the black and white keys of a piano, first appeared. MacCallum indicates the largest Guatemalan marimbas as spanning six and a half octaves from note F1 to B7. Resonators were commonly made of soft pine or cedar, with a shape that combines a pyramidal frustum on top and inverted pyramid on the bottom. This flaring of the resonators allows for shorter lengths than uniform tubes of the same pitch.

Per MacCallum, marimba manufacturing enjoyed a golden age in the United States from 1910 to 1920. The J. C. Deagan Company was a prominent contributor in this era, offering marimbas anywhere from three octaves to six octaves in size and available in half-octave

increments. Deagan also created and produced the “nabimba” - essentially an instrument comprised of the lower few octaves of a marimba. They even produced a “marimbaphone” —a type of marimba wherein the bars and resonators could be rotated 90 degrees such that the bar ends pointed up and the resonator tubes were parallel to the floor. In this configuration the bars were played by bowing their ends (which were curved for this purpose).

Around 1920, manufacturers in the United States ceased production of marimbas with notes below C3, and four-octave marimbas because the standard. More exotic instruments, like the nabimba, also disappeared from catalogues. This situation remained the norm until Musser introduced the four-and-a-third octave marimba, beginning with note A2, in 1960. At the time of MacCallum’s writing, Musser was the only manufacturer producing marimbas larger than four octaves. Bass marimbas, with notes beginning at C2, were available upon special order.

Fast forward to today and marimba size availability has expanded once more. Five-octave marimbas are commonly available from major manufacturers, including Adams, Bergerault, Malletch, Marimba One, Musser and Yamaha. Such instruments will typically include notes C2 through C7. It is not unusual to see marimbas of this size of service in an orchestra. Presently, the Schulich School of Music has three marimbas of this size.

1.1.2 Historical Tuning

According to MacCallum (1969), prior to the 1920’s, only the fundamental mode and frequency was considered when tuning marimba bars. Such bars were then at the mercy of the “clank note” (most prominent overtone). Should this clank note fall near, but not directly on, a higher octave of the fundamental, bar timbre could be rendered intolerable. MacCallum points out that striking a bar directly in its centre would minimize the clank, implying that the clank is indeed a product of the second vertical bending mode. Limiting mallet strikes to the centre of each bar is impractical and would be overly burdensome for players. Quality manufacturers would surely have repurposed or discarded bars with unacceptable clank notes.

MacCallum (1969) goes on to describe “octave tuning”, which allows the player to strike anywhere on the bar and produce a satisfactory sound.

In the case of marimba bars modern tuning brings down the errant first two overtones to notes of the same lettered name of the fundamental note. Tuning the third overtone is often done, but is difficult and is a luxury. It is usual to tune the bars of the xylophone (which are supposed to be rather thick) so that the most prominent overtone is a fifth lettered note above the fundamental.

This description departs in some ways from modern marimba tuning practice. Today, the second partial of a marimba bar is typically tuned two octaves higher than the fundamental, which would indeed give it the same lettered note, thereby matching MacCallum’s described practice. However, these days the third partial is commonly tuned to a factor of 10 times the fundamental, an interval of roughly three octaves and a major third. Thus this second overtone would not share the same lettered note as the fundamental. Achieving MacCallum’s description would require tuning this overtone to either eight times the fundamental, for a three-octave increase, or 16 times the fundamental, for a four-octave increase. Tuning the second overtone down to a ratio of eight times the fundamental may be achievable, with concentrated removal of material at strategic locations. Such tuning would likely be difficult, and may result in a fragile bar. Tuning the third partial up to a ratio of 16 times the fundamental is surely impractical, given the proportion of bar materials that affect both of these vibration modes. Claims of tuning a marimba bar’s fourth partial (third overtone) have only appeared anecdotally while working on this thesis. Given the variation in measured overtone frequencies notable in Appendix A, any such claims of tuning the third overtone seem dubious at best. If MacCallum has inadvertently used the term “overtone” the way the term “partial” is used today, his description of bar tuning would make more sense as a logical improvement over tuning only the fundamental. One can only speculate whether this is, indeed, the case.

Published not long after MacCallum’s book, Moore’s thesis (Moore, 1970) indicates tuning ratios of 1:4:10 for the first three partials of a marimba bar. It can be supposed then, with some confidence, that the tuning ratios seen today in marimba bars first appeared somewhere between 1920 and 1970, and were motivated by a desire to improve bar timbre and avoid the so-called “clank tones”.

This phenomenon of untuned “clank tones” producing problematic bar timbre persists to this day, though the vibration modes responsible and range of occurrence on the keyboard have changed. Addressing this issue forms part of the motivation for this work, as described in the next section.

1.2 Motivation

1.2.1 Troublesome Torsional Modes

The bars of marimbas, xylophones, vibraphones and similar bar percussion instruments are tuned by carving an “undercut” or “cutaway” on the bottom of the bar. Figure 1.1 shows an example of such a cutaway on a Bergerault F \sharp 3 marimba bar. Shaping a bar in this

manner produces desirable timbre while providing a smooth playing surface on the bar’s topside. Typically, up to three flexural modes of vibration are tuned to produce harmonic relationships between their modal frequencies. For marimbas and vibraphones, the first three flexural modes are most commonly tuned to have their frequencies align in ratios of 1:4:10, while 1:3:6 may be observed in xylophones (Moore, 1970).



Figure 1.1: Photo of the bottom of a Bergerault F#3 marimba bar. The central area where material has been removed is the bar’s *cutaway*. Also visible are two holes drilled across the bar’s width to house support cables.

Bars also exhibit other mode types, including flexure in both vertical and lateral directions, as well as torsional and axial modes. Figure 1.2, borrowed from Beaton and Scavone (2021a), shows examples of each mode type along with the names used to describe them here. Other mode types that do not significantly affect bar timbre (e.g. axial modes along a vertical axis) are omitted. These other mode types will have frequencies well above our range of interest and radiate a negligible amount of sound.

Makers have complained that, over a specific range of the marimba keyboard, the frequencies of untuned torsional modes may fall close enough to those of tuned flexural modes to pollute bar timbre (Stevens, 2015), sending valuable tonewood to the reject bin. Solving this problem forms part of the motivation for this work. Specifically, it is desirable to remedy the problematic torsional modes solely by shaping a bar’s cutaway, and without adding concentrated mass or secondary materials to the bar. This approach is preferred as adding materials would incur the additional costs of these materials and require some form of (likely adhesive) connection, which could become a point of failure.

1.2.2 Enabling Design Exploration

New Materials

Modern day professional quality marimbas and xylophones are invariably made with rosewood bars. MacCallum (1969) mentions a number of possible materials for making marimba bars, including Mexican ironwood, Brazilian Pernambuco wood, “chico zapote” (sapodilla), “encino roble” (white oak), hormigo, which was the preferred choice for Guatemalan marimbas, and “bois d’arc” (bodark) a wood native to the southern United States. MacCallum agrees with Moore (1970), though, that rosewood (specifically Honduran rosewood; *Dalber-*

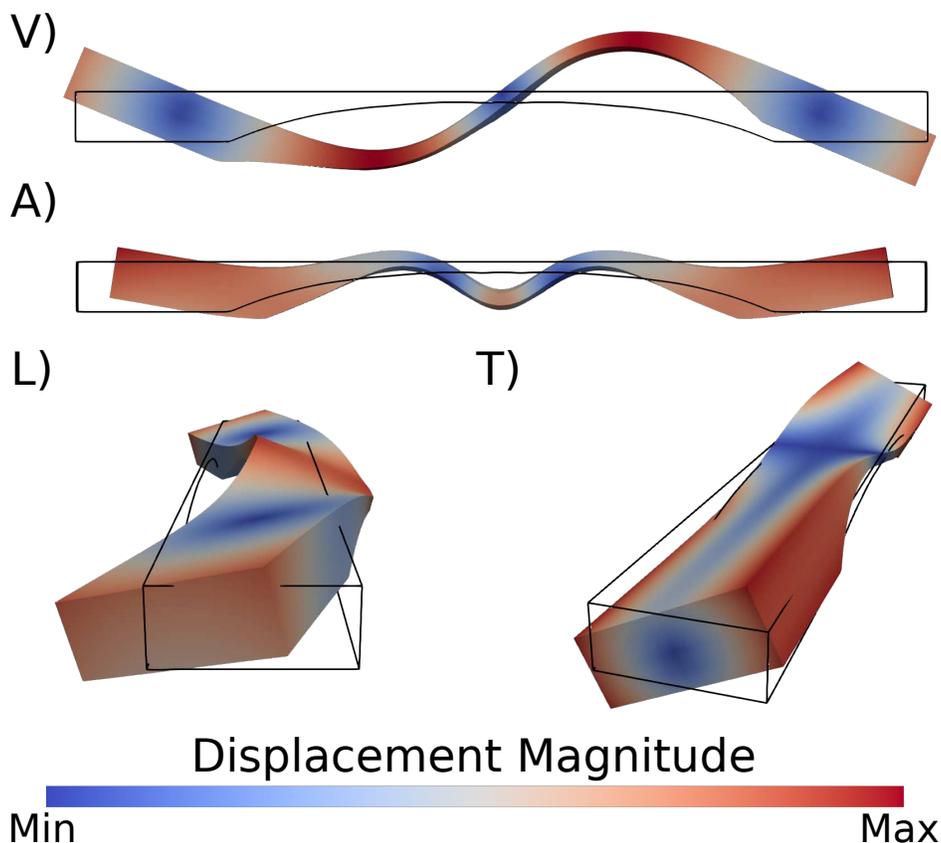


Figure 1.2: Vibration mode types and naming convention in this thesis. Undeformed bar geometry is outlined in black. Free-free boundary conditions are employed. Mode types: (V) - vertical-flexural mode, (A) - axial mode, (L) - lateral-flexural mode, (T) - torsional mode. Note that the bar cutaway in this example results in coupling between deformation types. Vertical deformation is present in the axial mode, while torsional deformation occurs in the lateral-flexural mode. Figure taken from [Beaton and Scavone \(2021a\)](#).

gia stevensonii) was the only material used for the manufacture of quality marimba and xylophone bars in the United States at the time of their respective publications. This trend continues to today, where a search for new or used marimbas will yield wholly rosewood results for quality instruments.

Such uniformity in marimba construction is imperilled somewhat by the vulnerability of rosewood species around the world. While some species of rosewood (e.g. Brazilian rosewood) are more endangered than others, all are in a vulnerable state. The *Convention on International Trade of Endangered Species of Wild Fauna and Flora* (CITES) controls international trade of specific wild animals and plants between participating countries (of which there are currently 184 listed on the CITES website). Brazilian rosewood has been listed in Appendix I of CITES since 1992, placing it in the most restrictive category of trade for the

most endangered species. In 2017, all other species of rosewood were added to Appendix II of CITES for species under threat, leading to restrictions on global trade. In 2019 an exception was added to allow the free trade of non-Brazilian species of rosewood used in musical instruments. While this exception makes it easier to bring instruments across international borders, clearly there would be advantages to finding a viable alternative to rosewood.

African padauk can sometimes be found in more inexpensive marimbas or xylophones, possibly intended for practice or instruction. Padauk has even been described as “diet rosewood” in the context of marimba and xylophone materials (Dooley, 2015). Though useful for cheaper instruments, the rarity of padauk marimbas reflects a general consensus on its inferiority to rosewood.

Synthetic materials have found application in musical instruments for some time now. For example, RainSong guitars are built almost entirely from carbon fibre, while Phoenix Pianos uses carbon fibre to produce soundboards and other components (e.g. pin blocks, hammer shanks). Synthetic materials also appear in mallet percussion instruments. Presently Musser produces synthetic marimba bars using a proprietary material named Keylon. Yamaha also produces a series of synthetic bars using their own proprietary material called Acoustalon. These synthetic materials have much lower damping than rosewood, and produce a timbre quite distinct from any marimba with wooden bars.

Given the situation described above, there exists a motivation to explore new potential materials for mallet percussion instruments such as marimbas and xylophones. This raises the obvious question of which materials may be best suited to the task. While some work has been done to assess the potential of different wood species (Holz, 1996; Aramaki et al., 2007; Brémaud, 2012), it would be advantageous to also explore new synthetic or other natural materials without having to first construct prototype bars.

Alternate Tuning Ratios

Addressing the problematic torsional modes described above raises some interesting questions. If the proximity of torsional mode frequencies to tuned mode frequencies is a problem, by how much must the modes be separated to fix this problem? With the right frequency ratios, could the torsional modes be made useful rather than problematic? If so, what would be the best frequency ratios for tuning some torsional modes along with the flexural modes?

Bork and Meyer (1985) investigated listeners’ preferred frequency ratio for the third partial of a xylophone sound. Their study used synthesized tones comprised of three partials, each with independently controlled frequency, peak amplitude and decay time. Listeners were found to prefer a frequency ratio of about 9.79 for the third partial in the xylophone sound. This value falls midway between three octaves plus a minor third, and three octaves plus a

major third, above the fundamental frequency. Such results show that listeners' preferences may not necessarily fall on harmonic ratios. This study also demonstrates an advantage of using synthesized tones for listener evaluation, as creating prototype bars for each tuning ratio evaluated would be costly, time consuming and wasteful of materials.

The relative amplitudes of each partial in [Bork and Meyer \(1985\)](#), as well as their evolution over time, are based on measurements of existing rosewood bars. In evaluating new bar designs, it would be most advantageous if measurements on physical prototypes were not a required input of the analysis.

Sound Radiation from Virtual Bars

Enabling the evaluation of new materials and tuning ratios by creating virtual models of test bars is thus the second motivation of this work. It is desirable that such evaluations can be carried out based solely on material parameters and proposed bar geometry, avoiding any need to first create a physical prototype. Such an approach would serve to greatly reduce the time and material cost of creating physical prototypes. In fact, if sufficient data is available for a candidate material, it could be evaluated without procuring any samples at all.

It is also desirable that an analyst could evaluate a proposed bar design by listening to its virtual representation. Developing quantitative measures of instrument quality remains a stubbornly elusive goal in musical acoustics. Rather than attempting to compare bar designs via any such quantitative measure, the model should instead output the sound radiated from a bar after a mallet strike. This will allow the analyst (and others) to listen to candidate bar designs and use their own subjective opinions for comparison. Any such radiated sound outputs should thus be as realistic as possible, with less importance placed on computational efficiency.

1.3 Thesis Layout

The subsequent chapters of this thesis are organized as follows:

Chapter 2 provides a review of published literature on computational modelling techniques for idiophone bars. It considers approaches to model vibrational behaviour, as well as sound radiation. Damping models are discussed, as they play an important role in sound production. Techniques to solve the mathematical problem of bar tuning are also considered.

Chapter 3 describes the finite element models of idiophone bars used in this work. The modelling approach described in this chapter is common to work appearing in chapters [4](#)

and 5.

Chapter 4 outlines a method developed to tune idiophone bars using three-dimensional cutaway geometries. The tuning problem is defined mathematically, and the solution method is described. Examples of aluminum and rosewood bars are presented and the solving algorithm's performance is discussed. The chapter concludes with an overview of fabricated prototype bars.

Chapter 5 presents a developed hybrid sound radiation model for idiophone bars. The modelling steps are described in detail. Suggested approaches to compute parameters specific to the model are provided. Example models are presented and tests are performed to evaluate the suggested model parameters.

Chapter 6 concludes the thesis, giving a summary of the work and recommending potential directions for future research.

Appendix A provides information on experimental measurements of marimba bars as well as stock materials. These measurements investigated the behaviour of professional-quality instruments and determined material properties for bar models.

Chapter 2

Review of Idiophone Computational Modelling Literature

2.1 Bar Vibration Modelling

Modelling the vibrational properties of idiophone bars is of primary importance in this work. Specifically, any models created must determine the modal frequencies at which the bar will vibrate for a given cutaway profile. While numerous options exist for modelling a bar's natural frequencies, the Finite Element Method (FEM) was selected to model bar vibration in this research. Its selection was based on numerous factors, including widespread availability of FEM packages, numerous example applications in the literature, and previous experience and familiarity with the method. The following sections provide a literature review of FEM in a context of musical acoustics and, more specifically, applications to idiophones. Alternative methods that were considered are also briefly discussed.

2.1.1 Background and History of the Finite Element Method

The Finite Element Method is a numerical method for solving differential equations. [Turner et al. \(1956\)](#) give the first thorough overview of the method as it is known today, including use of the term “element” in the manner now common. The concept of subdividing a region for analysis also appears in earlier works by [Hrennikoff \(1940\)](#) and [Courant \(1943\)](#). The popularity of FEM increased as computers grew faster, thereby enabling larger problems to be modelled within practical time limits. Today FEM is regularly employed to solve problems of engineering and design.

FEM as applied to structural analysis (i.e. solving the equations of elasticity) can be generally divided into two categories based on the types of elements used: solid elements and structural elements. Each node in a solid element will have a number of displacement degrees-

of-freedom (DOFs) equal to the dimension of the problem (e.g. for a three-dimensional problem, each node in a solid element will have three displacement DOFs). By contrast, structural elements include rotational DOFs in addition to displacement DOFs. Structural elements will also have dimensions less than or equal to the dimension of the problem domain. For example, it is common to see one-dimensional beam elements and two-dimensional shell elements modelling structural problems in three-dimensional space. For three-dimensional problem domains, each node of a structural element will have three displacement DOFs as well as three rotation DOFs, for a total of six DOFs per node. Structural elements are generally intended to represent large volumes of material compared to solid elements. In doing so they reduce the total number of DOFs in a problem, thereby lowering its computational demand. The following section contains examples of FEM applied to problems of musical acoustics. Both solid and structural elements can be found in these examples.

2.1.2 Musical Acoustics Applications of Finite Element Analysis

The popularity of FEM in engineering and design also extends to research in musical acoustics. Many examples of FEM applied to the analysis of musical instruments are available in the literature. This section will outline published applications of FEM to percussion instruments (idiophones), stringed instruments (chordophones) and wind instruments (aerophones). The papers discussed herein have been selected to exemplify various model types and applications. Truly, a detailed and comprehensive review of FEM applications in musical acoustics could fill volumes. The selection below is intended to be representative, though certainly not exhaustive.

Bork et al. - 1999

Bork et al. (1999) applied FEM to the modal analysis of a marimba bar and compared results with experimental measurements. Their model represented a real marimba bar of note C3 using 8-node three-dimensional solid elements. Figure 2.1 shows the mesh used in their model. Mesh density was varied along the length of the bar to reduce overall DOFs and computational costs. The mesh shown used 2880 elements. Modal frequencies between 100 Hz and 8000 Hz were considered (the fundamental frequency for the C3 bar was around 130 Hz). Their model produced 24 modes with frequencies within these limits. An average runtime of 10 hours was reported for their modal analyses, using a software package called *CASTEM* on HP workstations. Similar models built today using the program *CalculiX* (**Dhondt**) and a desktop computer will run in a few seconds.

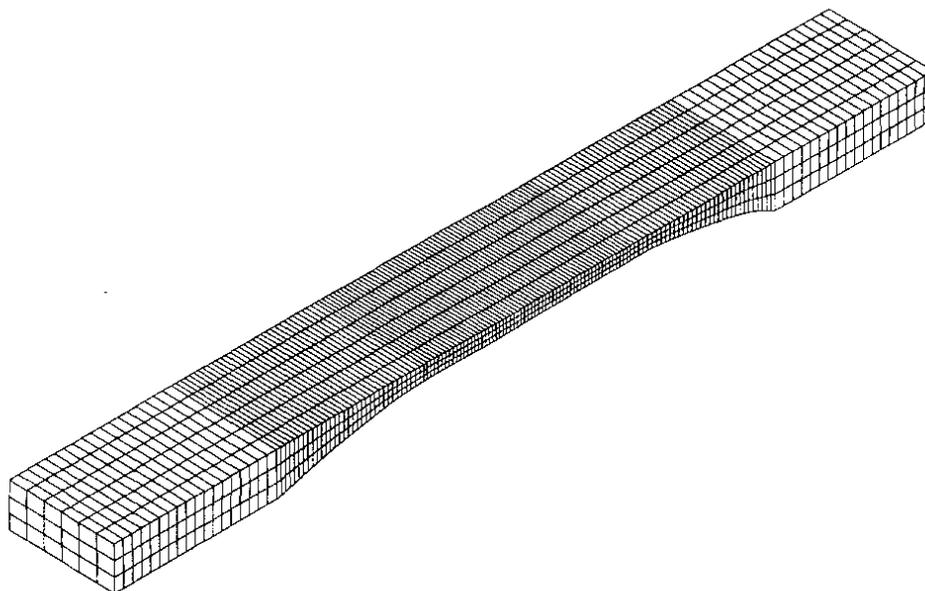


Figure 2.1: Marimba bar finite element mesh from [Bork et al. \(1999\)](#). Three distinct mesh densities can be seen along the bar's length.

[Bork et al. \(1999\)](#) defined their geometry by measuring a rosewood marimba bar at discrete points along its length and used an averaged value for the cutaway thickness. Experimental modal analysis was performed to measure mode shapes and frequencies. The modal frequencies predicted by FEM and those measured experimentally agreed to within 1% below 500 Hz and within 4% below 4000 Hz. Sensitivity analyses were reported for some of their results. They found the bar's fundamental frequency is sensitive to changes in bar thickness, with a 0.5 mm change over the length of the cutaway resulting in an 8% change in fundamental frequency. They found less sensitivity to Poisson's ratio, noting that changes of $\pm 50\%$ to Poisson's ratio resulted in changes of only $\pm 5\%$ for modal frequencies under 5 kHz. Importantly, [Bork et al. \(1999\)](#) noted that assuming isotropic material behaviour may produce reasonable estimates of the first few vertical-flexural modes, but modelling orthotropic material behaviour is essential to capture accurate behaviour of torsional modes.

Bretos et al. - 1999

[Bretos et al. \(1999\)](#) applied FEM to the modal analysis of wooden bars for musical instruments. Their aim was to study the influence of elastic material parameters and undercut shapes on modal frequencies. Three types of bar geometry were considered: constant cross-section, parabolic undercut and rectangular undercut. Figure 2.2 displays the meshes used for each geometry type, along with the control parameters h_u and L_u . Each mesh used 150 three-dimensional solid elements with a total of 668 nodes. The outer dimensions of the

bar were 330 mm x 34 mm x 18 mm, yielding an aspect ratio exemplary of xylophone bars. Material properties were modelled after a tropical wood commonly known as Azobé.

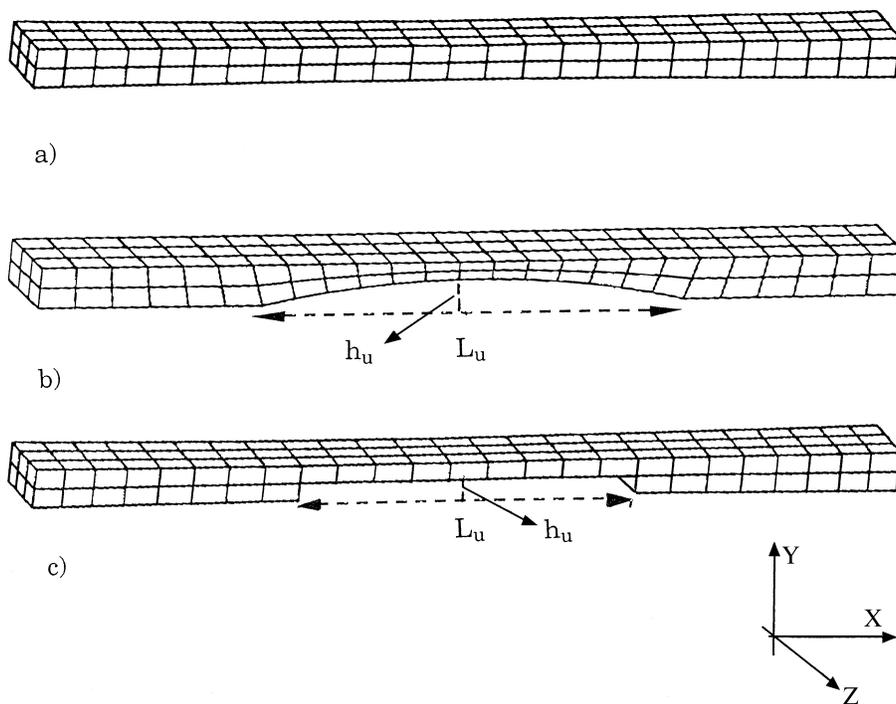


Figure 2.2: Wooden bar finite element meshes from [Bretos et al. \(1999\)](#). Three geometry types were considered: a) constant cross-section, b) parabolic undercut and c) rectangular undercut.

In analyzing the sensitivities of modal frequencies, [Bretos et al. \(1999\)](#) identify four material parameters as having significant influence: Young’s modulus in the principal (longitudinal) direction, E_1 , shear moduli along the principal axis, G_{12} and G_{13} , and density, ρ . They indicate that elastic parameters in the secondary axes (E_2 , E_3 and G_{23}) do not affect any of the natural eigenmodes. This is true for the modes considered in the paper. Elastic parameters in the secondary axes will, in fact, affect some modes of vibration; however, for any typical xylophone or marimba bar geometry, such modes will be much higher in frequency than the tuned modes and unlikely to noticeably affect the bar’s timbre.

Part of their parametric study considered how different undercuts affect the ratio of frequencies between the first two vertical flexural modes, f_2/f_1 . To reduce the risk of damage from mallet strikes they impose a limit on the undercut depth, h_u , such that $h_u \leq 0.7h$, where h is the overall bar thickness. Their study considered undercut lengths, L_u , up to $0.9L$, where L is the overall bar length. Their results show that over these undercut length and depth ranges a frequency ratio of $f_2/f_1 = 3$ (as is common in xylophone bars) is always

attainable. Their results also show that a frequency ratio of $f_2/f_1 = 4$ (more common for marimba bars) is attainable provided that $L_u \leq 0.75L$. Their investigation did not report on undercut influence on the third vertical bending mode frequency, f_3 .

In addition to the numerical sensitivity and parametric analyses, bars were fabricated for vibration testing and comparison with model results. One bar of constant cross-section was tested, along with a bar with parabolic undercut, tuned to note E4, and a bar with rectangular undercut, tuned to C4. The experimental setup used a forced vibration measurement technique described by [Bretos et al. \(1997\)](#). Two different excitation points were selected: one intended to excite many modes and another intended to excite only vertical flexural modes. For all three bars tested [Bretos et al. \(1999\)](#) found the measured spectral peaks aligned well with modal frequencies from the finite element models.

Henrique and Antunes - 2003

[Henrique and Antunes \(2003\)](#) applied one-dimensional structural finite elements in optimizing bar geometry for mallet percussion instruments. Their model utilized Timoshenko thick-beam elements which account for shear deformations and rotary inertia. This is in contrast to classical Euler-Bernoulli beam elements, which assume a beam's cross-section is small compared to its length. Previous experimental work has demonstrated Timoshenko beam theory to be more accurate in capturing the dynamic behaviour of percussion instrument bars than classical Euler-Bernoulli beam theory ([Orduña-Bustamante, 1991](#)). Additional examples of bar vibration modelled using one-dimensional beam elements are available in the literature ([Akay et al., 1983](#); [Petrolito and Legge, 1997](#)), including work by the same authors ([Henrique et al., 2001](#)).

[Henrique and Antunes \(2003\)](#) tested two optimization methods for tuning bar geometry: a stochastic method called simulated annealing, and a deterministic method called sequential quadratic programming. One example provided in their paper is the tuning of an aluminum vibraphone bar to a fundamental frequency of 880 Hz with frequency ratios 1:4:10 between the first three flexural modes. Their objective function calculated errors between the three tuned modes and their target frequencies, along with a parameter to quantify surface roughness. The bar was divided into 64 beam elements. Symmetry was enforced by mirroring the bar's geometry about its midpoint, leaving 32 element thicknesses to be defined. Bar length was either included as an input parameter or imposed as a fixed value.

Stochastic optimization via simulated annealing considered each of the 32 beam element thicknesses as separate inputs. The procedure tuned each of these inputs independently, relying on the objective function's roughness parameter to ensure practical bar geometries. [Figure 2.3\(a\)](#) displays the resulting bar geometry from simulated annealing for the 880 Hz

vibraphone bar. Bar length was included as an input variable in this example.

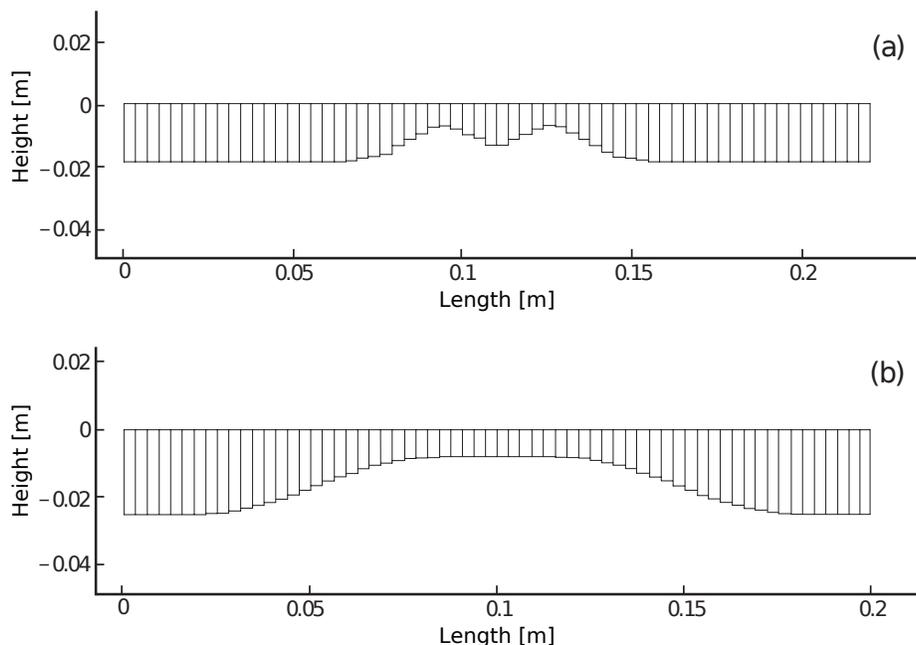


Figure 2.3: Optimized geometries for an aluminum vibraphone bar tuned to 880 Hz from [Henrique and Antunes \(2003\)](#). Geometries were optimized via (a) simulated annealing using element thicknesses as inputs, and (b) sequential quadratic programming using 10 cosinusoidal shape functions. Bar length was a variable in each case (note differing x-axis scales between panels).

For deterministic optimization using sequential quadratic programming, [Henrique and Antunes \(2003\)](#) employed orthotropic shape functions to define a bar’s thickness profile. A bar’s thickness at a given point along its length, $H(x)$, was determined by summing the weighted contributions of S orthogonal shape functions, $\psi(x)$. Equation 2.1 illustrates this calculation with weighting coefficients, A_s , and mean value, A_0 :

$$H(x) = A_0 + \sum_{s=1}^S A_s \psi_s(x). \quad (2.1)$$

Thus, for a given set of shape functions, $\psi_s(x)$, and bar length, L , bar geometry is entirely defined by the coefficients A_s , with $s \in [0, S]$. This approach has the advantage of separating the mesh and geometry definitions, thereby allowing for mesh refinement independent of the input parameters. It also reduces the problem complexity by limiting the number of inputs to the selected value of S . Figure 2.3(b) shows results from deterministic optimization of the example 880 Hz vibraphone bar using 10 sinusoidal shape functions.

[Henrique and Antunes \(2003\)](#) also investigated nonstandard frequency ratios leading to

bar geometries not commonly seen in practice. Figure 2.4 shows the resulting geometry of deterministic optimization with frequency ratios 1:3:5:7:9 and an imposed bar length of 350 mm.

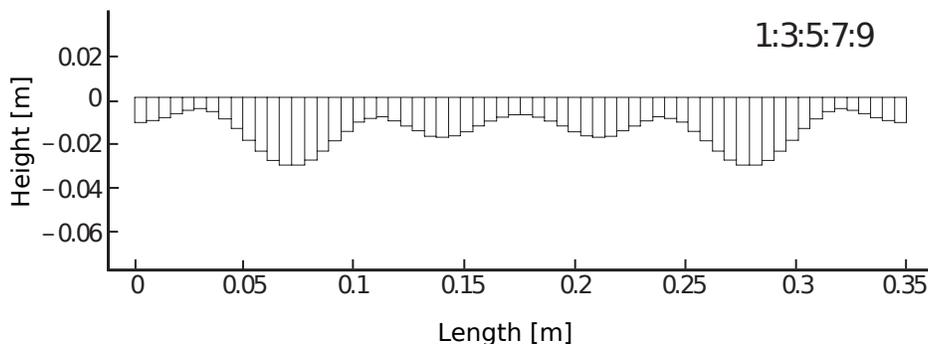


Figure 2.4: Optimized geometries for an aluminum vibraphone bar tuned to 440 Hz from [Henrique and Antunes \(2003\)](#). Flexural mode frequencies were tuned to ratios 1:3:5:7:9. Geometry was optimized via sequential quadratic programming using 10 cosinusoidal shape functions. Bar length was fixed at 350 mm.

In the second half of their study, [Henrique and Antunes \(2003\)](#) utilized their optimized bar geometries for sound synthesis via nonlinear dynamical modelling. Their model employed a modal decomposition of the bar’s response with averaged damping values from previous experimental results. Nonlinear interaction between the bar and an impacting mallet was incorporated. Mallet impact force was modelled using the Hertz model, similar to [Chaigne and Doutaut \(1997\)](#). This synthesis model was used to investigate the effects of varying strike location and the timbre of bars with unconventional modal frequency ratios.

Kirkland and Moradi - 2016

[Kirkland and Moradi \(2016\)](#) used three-dimensional solid finite elements to investigate the tuning of vibraphone bars. Up to 10,000 tetrahedral elements were used in their models, depending on bar length. They reported selecting tetrahedral elements for their improved accuracy in dynamic systems; however, they do not give a baseline for comparison or provide any substantiation. The geometry of their model is interesting in that it uses a spline curve to define the bar’s cutaway, and also accounts for the holes drilled for cable supports (which they term “node holes”).

Their goal was to find the geometry of a vibraphone bar that minimizes its mass while tuning its first three vertical-flexural modes of vibration to frequency ratios of 1:4:10. A tolerance target of 0.25% was established for the modal frequencies. Their intention was to enable fabrication of aluminum vibraphone bars on a computer numerical control (CNC)

mill, without any need for subsequent fine tuning.

[Kirkland and Moradi \(2016\)](#) used sequential quadratic programming (SQP) to tune their bar geometry. The vertical (thickness direction) coordinates of spline control points were used as inputs. The bar's width, end thickness and longitudinal control point coordinates were held constant. Bar length was initially considered as an input, but it was found that this resulted in a nonlinear distribution of bar lengths over the vibraphone keyboard, which would be impractical. They state that five spline control point coordinates were optimized by the tuning algorithm; however, Table I and Equation 6 of their paper seem to imply that two of these values may have been held constant, leaving only three input variables.

[Kirkland and Moradi \(2016\)](#) performed analyses by first optimizing a bar without node holes to within a frequency tolerance of 1.0%; this model was used to determine node hole locations. Subsequent models added these node holes and the analysis was repeated with frequencies tuned to within 0.25%.



Figure 2.5: Vibraphone bars tuned to musical note A4 and fabricated via CNC mill per [Kirkland and Moradi \(2016\)](#).

Results of their tuning process were demonstrated by machining two copies of a bar tuned to note A4 (442 Hz), as shown in Figure 2.5. They then mounted the sample bars on string passing through the node holes (similar to the condition on a vibraphone) and recorded the results of a mallet strike using a microphone. The recorded sounds were fed into the program *Audacity* for spectral analysis. The authors report identical measured frequencies

for the two fabricated bars to the nearest 0.01 Hz. However, they also state that their measured frequencies are accurate to within ± 1.35 Hz. Thus the -2 cents frequency error reported for their first transverse mode (mode V1 in the nomenclature of this thesis) could be anywhere from -7.5 cents to 3 cents. While performing experimental modal analysis on the fabricated bars could have yielded more precise modal frequencies, these results nevertheless demonstrate the potential for CNC milling to accurately tune aluminum bars.

Despite some minor criticisms here, [Kirkland and Moradi \(2016\)](#) present interesting work. Their study is instructive in its use of SQP to treat the modal frequencies as constraints.

Bestle et al. - 2017

In work detailed in both a paper ([Bestle et al., 2017](#)) and PhD thesis ([Bestle, 2017](#)), Bestle and his co-authors used 20-node three-dimensional solid elements to model and tune aluminum vibraphone bars. Figure 2.6 shows an example mesh from one of Bestle’s models. This model included 12,267 elements with 59,257 nodes and a total of 177,771 degrees-of-freedom. Bestle also modelled sound radiation from the vibraphone bars using surface nodes of the finite element mesh. Details of the sound radiation model are discussed in Section 2.2.1.

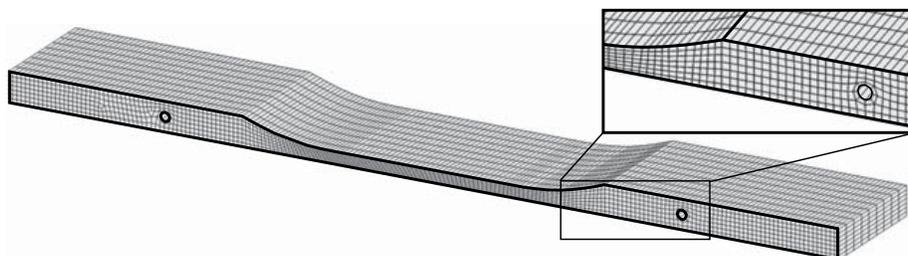


Figure 2.6: Example finite element mesh of an aluminum bar model from [Bestle \(2017\)](#).

[Bestle et al. \(2017\)](#) show an example of adjusting bar geometry to change the frequency ratios of higher-order partials. They provide an example of a flat bar with a fundamental frequency of 440 Hz. Two subsequent models with the same outer bar dimensions, but material removed at strategic locations, are tuned to have the same fundamental frequency, but different frequencies for other partials. Figure 2.7 shows the geometry of these three example bars.

Similar to [Kirkland \(2014\)](#), [Bestle \(2017\)](#) used spline curves to define the cutaway of an aluminum vibraphone bar. A cutaway was optimized for a bar tuned to musical note F3 using seven input parameters: one defining the length of the cutaway, and six defining depths of cut at points along the cutaway. Symmetry was enforced by mirroring bar geometry about

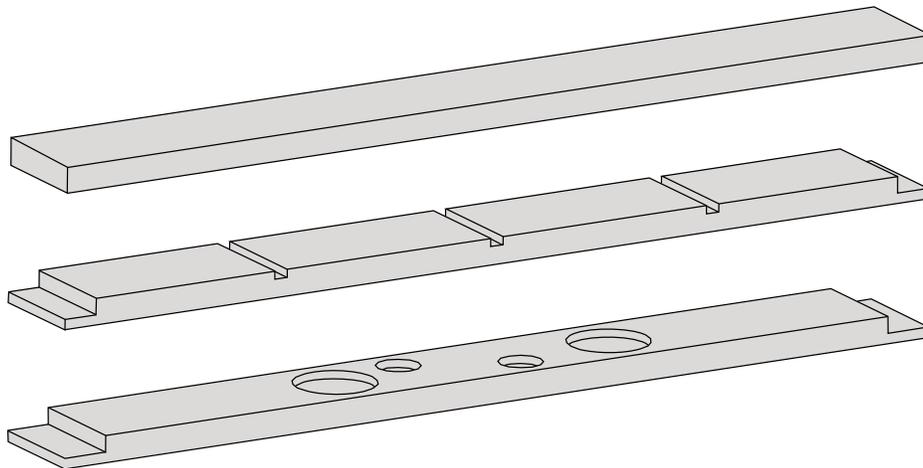


Figure 2.7: Aluminum bar models from Bestle et al. (2017). Each bar is tuned to a fundamental frequency of 440 Hz with differing frequencies of higher-order partials.

its centreline. Bar thickness was uniform across its width at any given point along the longitudinal axis. Equation 2.2 shows the tuning objective function used:

$$g(x) = \sum_{i=1}^5 w_i \frac{|f_i^{\text{FEM}} - f_i^{\text{target}}|}{f_i^{\text{target}}}, \quad (2.2)$$

where x is the vector of geometry inputs, w_i are weighting coefficients for the various modes, f_i^{FEM} are the modelled modal frequencies and f_i^{target} are the corresponding target frequencies. A vector $w = [5, 0, 3, 0, 1]^T$ was used to assign importance factors to the first three vertical flexural modes while ignoring the other (torsional and lateral flexural) modes during tuning. Note that this formulation assumes a constant mode ordering. A gradient-based optimization algorithm was used to minimize the objective function in Equation 2.2, but details of this algorithm were not reported.

Bestle (2017) goes on to model, optimize and fabricate two aluminum bars on a CNC mill. The first bar has what is described as a “classic” cutaway, which includes curved transition zones on each end with a centre portion of uniform thickness (similar to that shown in Figure 2.6). The second bar, shown in Figure 2.8, uses the spline formulation detailed earlier. Prototypes of the bars were fabricated and experimental modal analysis performed on each. Previous measurements on a Yamaha vibraphone bar indicated frequency ratios of 1:3.99:10.76 between the first three vertical flexural modes. The fabricated, optimized “classic” cutaway produced ratios of 1:4.03:10.29, while the fabricated, optimized cutaway with spline definition achieved a ratio of 1:4.01:10.01.

Both works (Bestle et al., 2017; Bestle, 2017) show the suitability of FEM for accurately

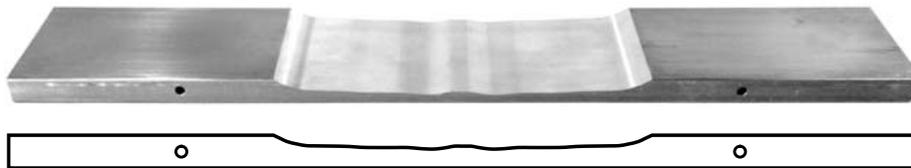


Figure 2.8: Fabricated aluminum vibraphone bar by [Bestle \(2017\)](#). Cutaway geometry was optimized to tune the frequency ratios of the first three vertical-flexural modes. A CNC mill was used to fabricate the precise bar geometry.

capturing the modal behaviour of bar models for subsequent fabrication via CNC mill. These works are also interesting for their proposed approach to model sound radiation from a struck bar (see Section [2.2.1](#)).

Rucz et al. - 2021

In a paper published while this thesis research was ongoing, [Rucz et al. \(2021\)](#) demonstrated a coupled vibroacoustic model of a xylophone. Using finite elements, their model considered a xylophone bar under mallet impact, an adjacent resonator, several nearby bars (whose presence affects the resonator), and a hemispherical portion of the radiation domain above the bar. Figure [2.9](#) shows a schematic of the simulation domain included in their model. Areas labelled Ω_a model acoustical wave propagation via the Helmholtz equation. The orange-coloured areas represent five xylophone bars (one struck and four passive) governed by the equation of motion. The acoustic and mechanical fields are coupled in the model. Mallet impact was modelled using the approach outlined by [Chaigne and Doutaut \(1997\)](#), which applies the excitation forces on a single finite element node and accounts for interaction between the bar and mallet at that node.

Both the sound radiation domain, Ω_a , and xylophone bars were meshed using three-dimensional solid finite elements. [Rucz et al. \(2021\)](#) report using hexahedral elements for the xylophone bars and tetrahedral elements for the acoustic domain, with interpolation employed between these incompatible meshes. Infinite elements are defined around the exterior of the unbounded acoustic domain, Γ_∞ . They report roughly 17,500 to 25,000 nodes in the sound bar meshes and roughly 100,000 nodes with around 400,000 elements in the acoustic domain mesh.

To reduce computational demand, [Rucz et al. \(2021\)](#) use a modal basis to reduce the size of the coupled system matrices. They then employ Newmark’s method with coefficients for average acceleration to evaluate the modal system’s response to mallet impact in the time domain.

[Rucz et al. \(2021\)](#) report evaluating eigenmodes of the acoustic domain to be the most

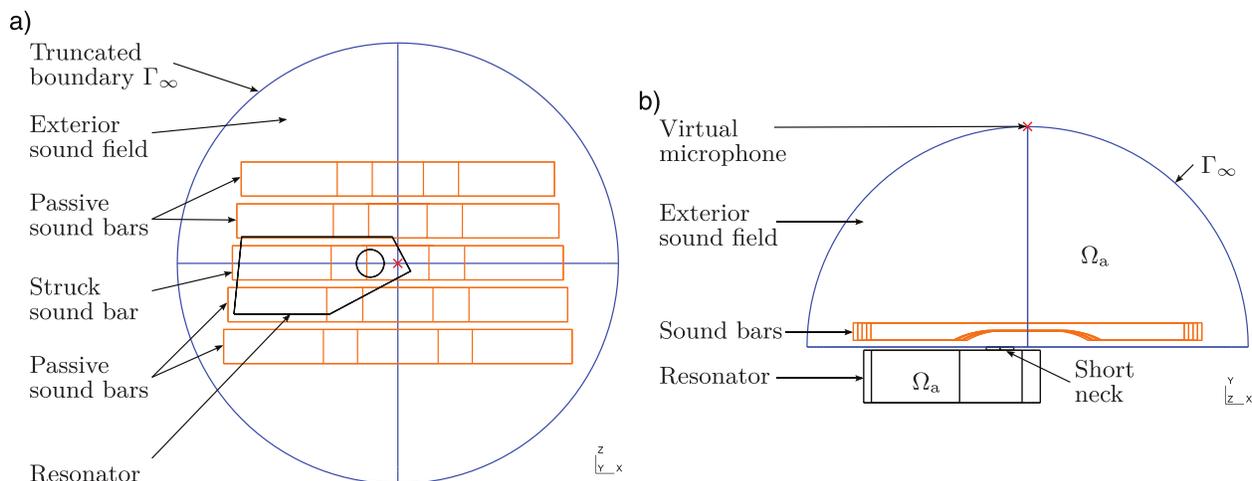


Figure 2.9: Schematic of xylophone sound radiation simulation domain from Rucz et al. (2021). (a) Top view. (b) Side view.

time-consuming step in their analysis, requiring around 40 to 60 minutes for a given example model. Other preprocessing steps are reported as taking about 15 minutes, while time domain simulation of a 10-second long bar system response required around three minutes after completing the other steps.

Rucz et al. (2021) perform several measurements to validate their model, including investigations of rosewood mechanical properties and resonator cavity resonance properties. They also consider energy distribution between mode types depending on mallet strike location. Example sound files are provided demonstrating some very realistic sounding mallet strike results from their model.

Soares et al. - 2021

Two papers using a simplified definition of bar undercuts were published while this thesis research was ongoing (Soares et al., 2021a,b). The second paper (Soares et al., 2021b), which discusses tuning of idiophone bar torsional modes, appeared in the literature after the methods developed in Chapter 4 were completed and published (Beaton and Scavone, 2021a).

Figure 2.10 shows the parameters used to define stepped undercut profiles used in the two papers. The stepped profile serves to limit the number of input variables while producing shapes that are simple to manufacture (Soares et al., 2021a).

In the first paper, Soares et al. (2021a) employ these simplified undercuts to define one-dimensional finite element models of idiophone bars using Timoshenko beams elements. The flexural modes of these models were then tuned using an evolutionary algorithm. Vari-

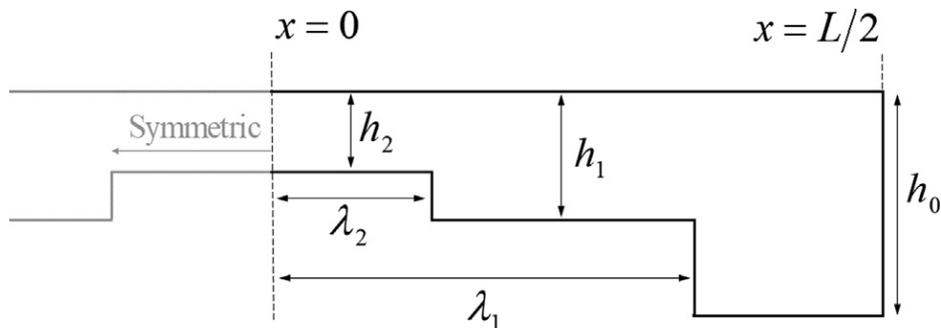


Figure 2.10: Side view of idiophone bar with parameters defining the simplified undercuts from Soares et al. (2021a). This example has two distinct steps, though generally the model may use any number of steps.

ous non-standard tuning ratios were considered. Some interesting examples were presented comparing simplified undercut shapes from the evolutionary algorithm to shapes found by Henrique and Antunes (2003) using cosinusoidal shape functions (including the bar shown in Figure 2.4).

The second paper (Soares et al., 2021b) applies this same simplified undercut profile to three-dimensional finite element bar models, and tunes both flexural and torsional modes. A surrogate optimization algorithm is employed to reduce the number of function evaluations (FE modal analyses) required in a global optimization scheme. Several example models are shown, tuning the first three vertical-flexural modes to the usual 1:4:10 frequency ratios, while also tuning up to two torsional modes.

Soares et al. (2021b) validate their methods by fabricating three prototype bars with these stepped undercuts, designed using their optimization approach. Figure 2.11 shows the three prototype bars along with the frequency ratios of tuned vertical-flexural and torsional modes. Performance of these bars is discussed in Section 4.5.3.

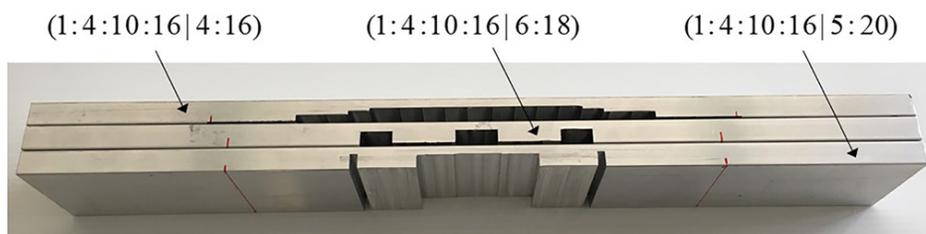


Figure 2.11: Prototype bars with stepped undercuts fabricated by Soares et al. (2021b). Frequency ratios are indicated for tuned modes: (V1:V2:V3:V4|T1:T2).

Soares et al. (2021b) also provide sound recordings of their prototype bars being struck with a typical vibraphone mallet, with resonator tubes positioned beneath each bar. Spectrograms of the resulting sounds are provided. The paper concludes with a discussion of

modal damping and the timbral features of the different recordings.

Perrin et al. - 1983

Perrin et al. (1983) performed modal analysis on an English church bell using an axisymmetric finite element model. Axisymmetric finite elements are designed to analyze axisymmetric structures (also called solids of revolution) using a two-dimensional mesh of one slice of their profile. This approach has the advantage of drastically reducing the number of DOFs, thereby reducing computational complexity compared to a three-dimensional model of the full structure. The fact that this paper was published in 1983 is evidence of this advantage. Full three-dimensional finite element bell models became more common in the 1990s (Kim-mich et al., 1992) and thereafter (Zhang et al., 2007), as the computing power available to researchers increased.

Perrin et al. (1983) compared their results with experimental measurements of a 56.6 cm tall, 70.2 cm diameter, 214 kg English church bell tuned to the musical note D5. The bell was first excited acoustically using four loudspeakers and an oscillator. This setup was used to identify all apparent partials up to 9.3 kHz. Next the bell was excited magnetically by attaching a small disc and employing a magnetic transducer. Two accelerometers were then used to survey the bell's response when driven at a specific frequency. One accelerometer was connected to the bell near the drive point, or diametrically opposite it, as a phase reference. The second accelerometer was held against the bell by hand and used to explore and map nodal lines of the active mode.

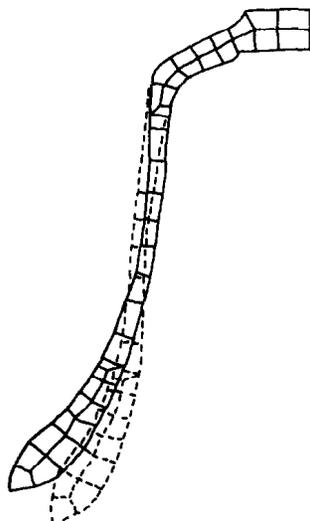


Figure 2.12: Axisymmetric finite element mesh of an English church bell from Perrin et al. (1983). Solid lines indicate undeformed mesh. Dashed lines show modelled results for the six-meridian inextensional radial ring driven mode known as the *Tierce*.

Figure 2.12 shows the axisymmetric finite element mesh employed by Perrin et al. (1983). Forty elements can be counted in the figure, demonstrating the simplicity of this mesh compared to many seen today. The use of this relatively small mesh enabled the modelling of this 214 kg bell at a time when computational power was comparatively limited. By contrast, a 2007 study modelling wave propagation in chime bells used between one million and ten million three-dimensional solid finite elements per model (Zhang et al., 2007).

Also shown in Figure 2.12 is the displaced shape of a mode known as the *Tierce* – one of the commonly tuned modes of a bell. Perrin et al. (1983) identified 134 vibrational modes up to 9.3 kHz. They used the mode shapes computed by the finite element model to match those from experimental measurements and compared frequency results. In some cases mode shapes calculated by the model, but missing from measurement data, prompted the authors to perform additional measurements using driving locations informed by model results. Good agreement was found between modelled and experimental results for lower modal frequencies, with larger deviations seen as frequency increased. This trend, also observed by Bork et al. (1999), is common in modal analysis using finite elements. Considering the relative coarseness of the mesh in Figure 2.12, discrepancies between modelled and measured results for high modal frequencies are not surprising.

McLachlan et al. - 2012

McLachlan et al. (2012) used explicit finite element analysis to model tuning the natural modes of a gong with a punch. Their model included nonlinear geometry (large displacements) and material behaviour (plastic deformation and residual stresses) to capture the effects of cold-forging small dimples in the gong. These sophisticated models were validated by fabricating physical prototypes.

McLachlan et al. (2012) used second-order three-dimensional solid tetrahedral elements in their gong model. The model was created using simulation software *Simulia Abaqus* (Abaqus). Mild steel was selected to model the gong and produce the physical prototype. The steel's elastic and elasto-plastic behaviour was determined via measurements on test specimens. An 11.895 kg drop hammer attached to a 1 m pivot arm was used to produce dimples in the physical prototypes. The hammer's velocity immediately before impact with the gong was estimated using a high speed video camera. This value was then used as the initial hammer velocity in the finite element model. After forming dimples, a scanning laser Doppler vibrometer and shaker were used to measure mode shapes of the prototype gongs.

Panels (A) and (B) in Figure 2.13 display model results for residual stresses after dimple formation. Panel (C) of Figure 2.13 plots the change in frequency for two vibrational modes as a function of dimple depth at the centre of the gong. The plot gives results from both

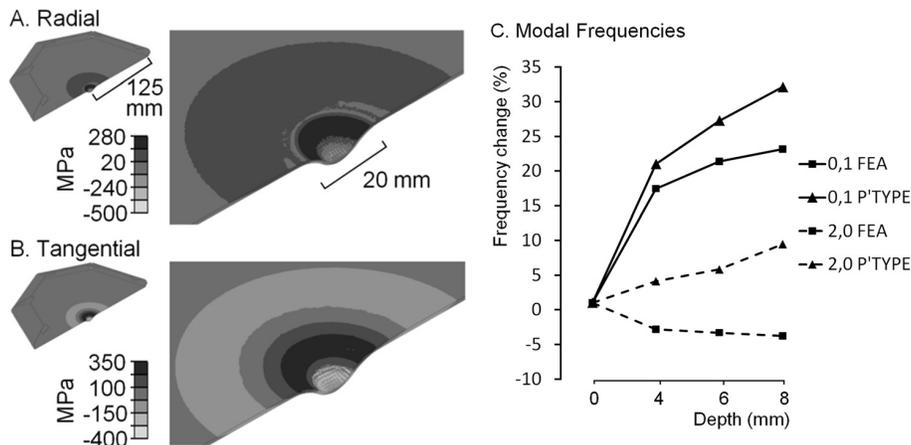


Figure 2.13: Finite element model results from [McLachlan et al. \(2012\)](#). A) Residual radial stresses in polar coordinates after dimple formation. B) Residual tangential stresses in polar coordinates after dimple formation. C) Effect of adding a 20 mm diameter dimple of various depths on the frequency of the 0,1 and 2,0 modes. Results are shown for both a linear finite element analysis (FEA) model with residual stresses, and a physical prototype (P'TYPE).

the finite element model and measured physical prototype. Significant differences can be seen in the change in modal frequencies between the modelled and measured results in [Figure 2.13](#). Strain rate dependencies for plastic deformation, which were not included in the finite element model, were identified as a likely source of these discrepancies. Note that their paper used a m,n naming convention when referring to the gong's vibrational modes, where m is the number of meridial nodal lines and n is the number of nodal rings.

Interestingly, [McLachlan et al. \(2012\)](#) report that models run with the dimpled gong geometry, but with residual stresses removed, showed no appreciable frequency changes in the 0,1 and 2,0 modes. Thus they attribute the change in frequencies of these modes solely to the presence of residual stresses.

[McLachlan et al. \(2012\)](#) go on to attempt tuning the first five modes of a prototype gong using four dimples evenly spaced at half the gong's radius. They provide results showing the effect of increasing dimple depth on the five modal frequencies of interest; however, they do not achieve the desired harmonic relationships using dimples alone. Instead, they report that the harmonic relationships between modes were eventually achieved by stressing a broad region about the centre of the gong surface to the steel's yield point without forming a dimple and then allowing it to spring back. Limited details are provided as to why this extra step was effective in achieving the desired modal frequencies. They conclude that their method of creating residual stresses via cold forming small dimples can reliably tune gong modes to harmonic relationships. The validity of this result is obscured, however, by the required additional step of yielding a broad area of the centre of the gong.

Additional Idiophone Modelling Applications

The papers outlined above represent a sample of finite element analysis applied to musical acoustics research on idiophones. Numerous other applications are available in the literature.

[Legge and Petrolito \(2007\)](#) combined FEM with constrained optimization to tune the shape of xylophone and marimba bars. [Laukkanen and Worland \(2011\)](#) used FEM to model sequential removal of material from a vibraphone-like aluminum bar, and compared their results with measurements of a machined bar.

[McLachlan \(1997\)](#) performed a parametric study of geometry changes and their effects on gong vibrational modes using FEM. [Perrin et al. \(2014\)](#) modelled the vibrational modes of a gamelan gong using FEM and compared results with measurements by electron speckle pattern interferometry (ESPI).

[Perrin et al. \(1995\)](#) employed FEM in analyzing the vibrational modes of several types of bells. A later study compared FEM results with ESPI measurements of an Indian elephant bell's vibrational modes ([Perrin et al., 2012](#)). Several studies ([Kimmich et al., 1992](#); [McLachlan et al., 2003](#); [Özakça and Gögüş, 2004](#); [Legge and Petrolito, 2007](#)) have used FEM to optimize bell geometries. [Zhang et al. \(2007\)](#) created a very large FEM model to study wave propagation in Chinese chime bells with a two-tone acoustic characteristic. [Carvalho et al. \(2017\)](#) used FEM to study tuning bells by attaching masses at strategic locations. Their study was motivated by a desire to improve the tuning of existing historic bells. Other studies have used FEM to recreate or reverse-engineer the sounds of historic bells ([Debut et al., 2016, 2019](#)).

2.1.3 Alternative Modelling Methods

The examples outlined above clearly illustrate the prevalence of FEM in musical acoustics research. One practical advantage of this ubiquity is the vast assortment of FEM packages, both commercial and open-source, available to researchers. However, numerous alternative methods to FEM do exist. Depending on the intended application, some may hold advantages over FEM and may therefore be more popular. Outlined below are various alternative methods with discussion of their application.

Finite Difference Method

Another very commonly applied method of solving partial differential equations in musical acoustics is the Finite Difference Method (FDM). This method is often taught in courses on musical acoustics, with the advantage that its formulation is easier to understand and apply to a new differential equation than that of FEM. Specifically, FDM can be applied

directly to a differential equation by substituting FDM approximations of derivatives into the equation. By contrast, FEM requires first converting the differential equation into integral form, a process requiring more sophisticated mathematics. This simplicity in formulation makes applying FDM to a new differential equation easier than applying FEM.

One disadvantage of FDM compared to FEM is that it can be challenging to define complex geometries using FDM. FDM is most easily applied on a uniform mesh (whether in space or time) and grows more complicated when nonuniform spacing is included. By contrast, once FEM elements have been defined for a given set of governing equations, it is easy to apply them to model complex geometries. Another disadvantage of FDM is that care must be taken to avoid numerical instabilities, which can arise from ill-suited discretizations.

FEM and FDM can also be combined to leverage the advantages of both. For example, a structural dynamic time history model will commonly use a mix of these two methods. An FEM formulation will commonly be applied to solve the structural equations, while an FDM formulation will handle the progression in time.

Examples of FDM applications in musical acoustics include 1D models of xylophone bar vibration (Chaigne and Doutaut, 1997), as well as piano string vibration (Chaigne and Askenfelt, 1994) and nonlinear cymbal models (Bilbao, 2010).

Boundary Element Method

The Boundary Element Method (BEM) is another numerical approach to solving differential equations with applications in musical acoustics. BEM differs from FEM in that its formulation is based on boundary integral equations of a problem rather than differential equations. Solving a problem using BEM requires a knowledge of the problem's *fundamental solution* in an infinite domain. The results of a BEM analysis for a given point in the domain can be conceptualized as the summed influence of the fundamental solution acting from every point on the boundary to the point of interest within the domain.

A key advantage of BEM in musical acoustics applications is that only the boundary is discretized. For exterior problems modelling sound radiation from a vibrating body, this eliminates the complication of establishing infinite or non-reflecting boundary conditions (see further discussion in Section 2.2.3). By contrast, FEM and FDM require discretizing the entire domain. As a result, BEM models will typically require far fewer degrees-of-freedom than FEM or FDM, leading to smaller system matrices. Defining meshes only on the boundary can have practical advantages as well, in that geometry data from animation software or 3D scanners could be more readily imported into a BEM model with less need for further meshing.

One disadvantage of BEM is that it will often lead to non-symmetric fully populated system

matrices. This differs from FEM, which typically produces sparsely populated symmetric matrices that can benefit from efficient solving algorithms. As a result, some of the advantage of having fewer degrees-of-freedom in BEM is offset by its fully populated matrices. Another drawback of BEM is that it is not well suited for nonlinear problems or domains with multiple materials. From a practical standpoint, BEM also suffers from a shortage of available software packages compared to the much more popular FEM. The mathematics underlying BEM is sophisticated and BEM is less commonly taught at universities compared to FDM and FEM. For this reason, applying BEM to a new differential equation may be quite challenging, thereby presenting an obstacle for the analyst. A related disadvantage is that fundamental solutions may not be available for all problems.

Examples of BEM applications in musical acoustics include modelling sound radiation from marimba bars (Bestle et al., 2017), optimizing bell geometry (Yu and Kwak, 2011), and simulating guitar sound radiation (Tahvanainen et al., 2019).

Other Methods

There are many other numerical methods available for solving differential equations. A few examples with applications in musical acoustics are listed below.

1. *Meshfree Methods* are a class of numerical methods which, as the name implies, do not include the types of meshes seen in FEM or FDM. Examples of meshfree methods include *Smoothed Particle Hydrodynamics* (SPH), the *Element-Free Galerkin Method* (EFG) and numerous others. Specific applications of meshfree methods to musical acoustics are difficult to find in the literature. Examples of general applications which could be extended to musical acoustics do exist, such as vibration analysis of a pre-stressed membrane (Unnikrishnan et al., 2020).
2. *Isogeometric Analysis* is a relatively recent approach to computational modelling that uses information (e.g. splines) from computer aided design (CAD) packages to define the geometry of modelled objects rather than relying on traditional FEM meshes. As is the case with meshfree methods, examples of general applications to acoustics are available in the literature (Manh et al., 2011; Kolman et al., 2015; Wu et al., 2015; Dinachandra and Raju, 2017), though specific applications to musical acoustics are difficult to find.
3. The *Finite Volume Method* (FVM) is another numerical method that solves differential equations by subdividing the problem domain. Similar to FEM, this method requires meshing the entire problem domain and can more easily handle complex geometries

than FDM. FVM is more commonly applied to problems of computational fluid dynamics (CFD), and is thus better suited to modelling aerophones than idiophones or chordophones. Examples of FVM applied to musical acoustics include modelling gas flow inside a bassoon (Richter, 2012).

4. The *Scaled Boundary Finite Element Method* (SBM) is another relatively recent numerical method for solving differential equations. It is similar to, and compatible with, FEM, but formulates elements differently. Example musical acoustics applications of SBM in the literature are rare, but include modelling sound radiation from an ancient Chinese bell (Liu et al., 2019a). Example applications to general acoustics in the literature are more easily found (Liu et al., 2016, 2019b).

The methods listed above are some of the more well-known approaches for solving differential equations with applications in musical acoustics. Other methods exist but pale in terms of popularity. For brevity, the discussion of alternative methods to FEM will be limited to those outlined above.

2.2 Sound Radiation Modelling

In addition to modelling a bar’s vibrational behaviour, this research also considered sound radiation. The choice of method to model sound radiation was considered after selecting FEM to model bar vibration. Thus, methods of modelling sound radiation that are compatible or easily interoperable with FEM were strongly preferred. Ultimately, a method that is herein called the *Surface Point Source Method* (SPS) was selected. This recent method was developed by Bestle et al. (2017) and originally applied to model vibraphone bars. It was selected for use in this research based on compatibility with FEM, relative ease of implementation, and mesh-free formulation (i.e. no meshing of the propagation domain). The following sections provide a general review of SPS, along with discussion of other methods considered for sound radiation modelling. A description of the bar modelling used by Bestle et al. (2017) appears in Section 2.1.2. Mathematical details of the method are discussed in Section 5.3.3.

2.2.1 Surface Point Source Method

Bestle et al. (2017) approximated the sound radiated from a vibraphone bar (see Section 2.1.2 for description of the bar model) by treating each element on the discretized surface as a point source emitting a spherical wave. The resulting wavefront encountering a listener

is then calculated as the summation of component waves from each of the surface point sources. This approach thus represents an application of Huygens' principle that each point on a complex wavefront is itself a source of spherical wavelets that sum to form the complex wave.

The approach taken by [Bestle et al. \(2017\)](#) decouples the structural dynamics problem from the sound radiation problem, ignoring any influence of the surrounding air on the bar's motion. The wave front emanating from a single surface element to a given listener position is calculated as:

$$S_f(t) = \frac{\rho_0 A_f}{4\pi r_f} \dot{v}_f(t) \cos(\Theta_f), \quad (2.3)$$

where S_f is the sound pressure from element f , ρ_0 is the density of air, A_f is the surface element area, r_f is the distance from the element centroid to the listener position, \tilde{v}_f is the surface element normal velocity and Θ_f is the angle between the surface element normal and a vector from the element centroid to the listener position. The sound pressure at a listener position at time t is calculated by summing the contributions of each surface element with appropriate time delays:

$$S(t) = \sum_f S_f \left(t - \frac{r_f}{c} \right), \quad (2.4)$$

where c is the speed of sound in air. Thus calculating the resulting sound pressure requires assembling a buffer of contributions from each element. Since the ratio $\frac{r_f}{c}$ will not necessarily be a multiple of the selected time step, contributions from each element must be interpolated to avoid unwanted artifacts in the resulting sound.

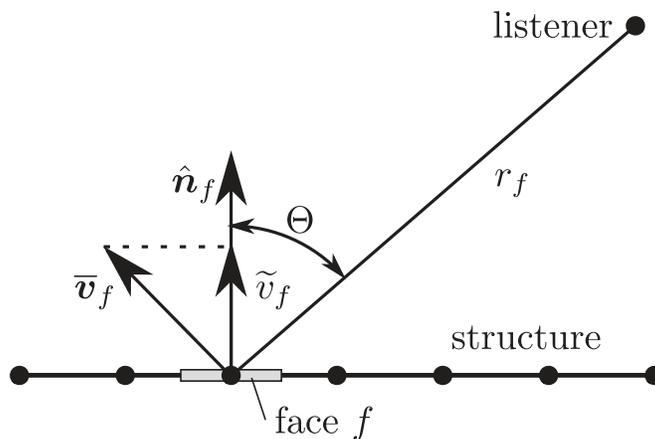


Figure 2.14: Variables in the calculation of sound radiation from a single surface element by [Bestle et al. \(2017\)](#).

Figure 2.14 illustrates some of the parameters in Equations 2.3 and 2.4. Since sound radiation is calculated solely based on surface motion and listener position, this method

avoids any need to discretize the propagation domain. [Bestle et al. \(2017\)](#) also describe a method to account for relative motion of the listener and sound source, thereby modelling the Doppler effect with their approach.

2.2.2 Boundary Element Approach

In modelling sound radiation, the Boundary Element Method (BEM) has the significant advantage that it does not require meshing of the propagation domain. Instead, only the surfaces of modelled objects need to be meshed. Models in this research consider idiophones in an anechoic environment, with no sound waves returning from outer walls or any other surfaces. This is equivalent to modelling the idiophones in an unbounded domain. Thus, when using BEM to model wave propagation, only the surface of the idiophone itself requires meshing. This saves considerable time and modelling effort compared with the FEM approaches described below.

Surface meshes used in BEM require careful consideration. Researchers have sometimes applied rules of thumb when meshing BEM models for acoustics, specifying a minimum number of elements per wavelength, though example models are more often for interior problems ([Marburg, 2008, 2002](#)). BEM can be prone to numerical issues, such as instability and fictitious modes, making its application more challenging than one might hope. [Tahvanainen et al. \(2019\)](#) coupled a FEM structural model of a guitar with a BEM sound radiation model. Their work used a specific type of BEM called the *Fast Multipole Boundary Element Method* ([Nishimura, 2002](#)), with improved computation speed.

2.2.3 Finite Element Approach

In contrast to BEM, adopting a finite element approach for modelling sound radiation from an idiophone requires discretizing the propagation domain. Any required mesh must be limited to a practical number of elements. The number of required elements will be proportional to the volume of the modelled domain. Thus it is not feasible to define a domain large enough that waves reflecting from the boundary would not reach the listener position during a simulation. The number of elements and DOFs required for such a domain would surely be computationally prohibitive. Instead, applying FEM to the propagation domain requires defining some form of computationally non-reflective boundary condition. In other words: a boundary that permits waves *outgoing* from the domain to pass without allowing/producing any *returning* reflected waves. While simple to state, this condition has proven difficult to achieve in practice, leading to a great deal of research efforts over the years ([Givoli, 2008](#)).

Different approaches to defining such conditions have been explored in the literature. Some

of these approaches resemble traditional boundary conditions that are defined over a surface in a three-dimensional model, while others occupy three-dimensional space. So long as an approach performs well computationally, it does not matter whether additional volume is required beyond the propagation domain. The literature on approaches to absorbing boundary conditions is extensive, with a good survey provided by [Givoli \(2008\)](#). Several approaches detailed in the literature are outlined below. The papers referenced herein are only a sample of those available.

Absorbing Boundary Conditions

The first boundary conditions considered here resemble more traditional boundaries which are defined over a surface in a three-dimensional model, and do not require additional elements or nodes. Two early approaches to absorbing boundary conditions considered by [Givoli \(2008\)](#) to be “classical” are those by [Engquist and Majda \(1979\)](#) and [Bayliss and Turkel \(1980\)](#).

These methods are considered *low-order* as their formulation in higher orders leads to numerical difficulties in implementation. A review of methods designed for implementing higher-order absorbing boundary conditions is provided by [Givoli \(2004\)](#).

Absorbing Layers

[Givoli \(2008\)](#) credits [Israeli and Orszag \(1981\)](#) as having devised the absorbing layer. As the name implies, absorbing layers create an effectively absorbing boundary condition using a layer of finite thickness around the domain. In the original formulation, this was accomplished by adding damping terms to the governing equation within the absorbing layer ([Israeli and Orszag, 1981](#)).

Perfectly Matched Layer (PML) is a popular form of absorbing layer first developed in the 1990s by [Berenger \(1994\)](#). PMLs are designed to have no reflection at the boundary to the modelled domain, with the solution decaying exponentially within the PML ([Givoli, 2008](#)). [Figure 2.15](#) shows an example of a two-dimensional FEM model with a PML defined around the exterior. The finite width of the PML is clearly demonstrated in the figure.

Since their initial development PMLs have gained popularity in modelling radiative boundaries. [Bermúdez et al. \(2008\)](#) provide a survey of developments in PMLs up to 2008. More recent applications include [Kaltenbacher et al. \(2013\)](#) and [Vermeil de Conchard et al. \(2019\)](#). PMLs can also be applied with other computational methods, including FDM ([Chobeau and Savioja, 2016](#)) and Meshfree Methods (MFM) ([Yang et al., 2019](#)).

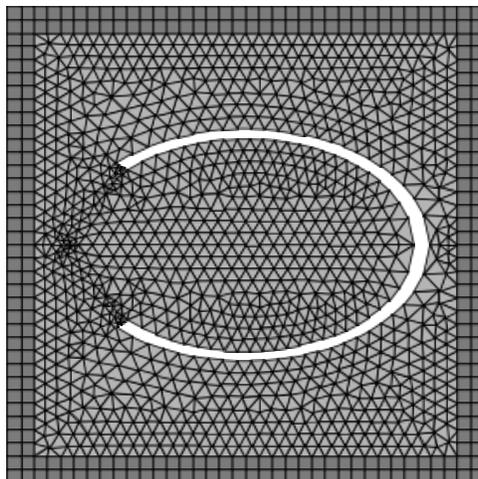


Figure 2.15: Example mesh from a two-dimensional model of an open elliptic amphitheatre from [Bermúdez et al. \(2008\)](#). Amphitheatre boundaries are displayed in white. Light grey triangular elements represent the air. Darker grey square elements define a perfectly matched layer to model radiative boundary conditions.

Infinite Elements

As the names implies, infinite elements are designed to represent portions of a domain that is unbounded in a given direction. Figure 2.16 shows a conceptual sketch of their application. Early works on infinite elements were first published in the 1970s ([Bettess, 1977](#); [Bettess and Zienkiewicz, 1977](#)). [Givoli \(2008\)](#) credits [Burnett \(1994\)](#) and [Astley \(1996\)](#) with establishing important milestones in the treatment of waves in unbounded domains.

Infinite elements are regularly applied to problems of acoustic radiation. Recent applications in the literature include examples of modal analysis with infinite elements ([Peters et al., 2013](#); [Moheit and Marburg, 2017](#)), most notably applied to model sound radiation from xylophone bars ([Rucz et al., 2021](#)).

Similar to PMLs, infinite elements have been coupled with other methods, including MFM ([Wu et al., 2019](#)).

Other Absorbing Boundaries

Further examples of absorbing boundary conditions include continued fraction absorbing boundaries ([Guddati and Lim, 2006](#)) and reciprocal absorbing boundaries ([Nguyen and Tasoulas, 2018](#)).

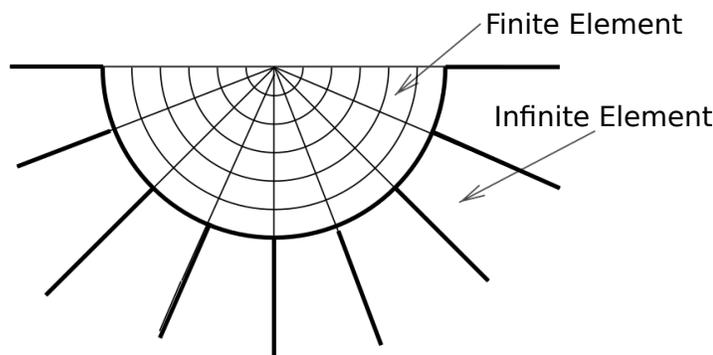


Figure 2.16: Conceptual sketch of a two-dimensional unbounded domain meshed with finite elements and infinite elements. Figure from [Givoli \(2008\)](#).

2.2.4 Scaled Boundary Finite Element Method

The Scaled Boundary Finite Element Method (SBM) is a relatively recent numerical method, having been initially conceived in the 1990s. It was originally called the “consistent infinitesimal finite-element method” and was formulated for dynamic analysis of unbounded domains ([Wolf and Song, 1995](#)). The method has since been expanded to include numerous problem types including structural dynamics ([Song, 2009](#)).

SBM has the same advantage as BEM, in that it requires discretizing only the boundaries of an unbounded domain ([Song, 2018](#)). This leads to a reduction in spatial dimensions by one and eliminates any need for the absorbing boundary conditions required with FEM.

SBM also provides the option to model a portion of the propagation domain, while still providing radiative exterior boundary conditions. [Liu et al. \(2019a\)](#) provide an interesting example of such a case. Their model considered an ancient Chinese bell excited at a single point by a smooth pulse of finite duration. Figure 2.17 displays portions of the meshed bell and modelled propagation domain.

2.3 Damping Models

Appropriate modelling of damping is a critical aspect of producing realistic bar sounds. Relative decay rates of the various vibration modes will have a large effect on a bar’s timbre. Damping of the bar’s response, in particular, will dominate the overall energy loss compared with any loss in the propagation medium. Unfortunately, computing a damping matrix is not nearly as straightforward as forming mass or stiffness matrices, as the physical principles underlying damping are not as well understood ([Adhikari, 2006](#)).

This research employed Rayleigh damping to model the bar’s response to impulsive excita-

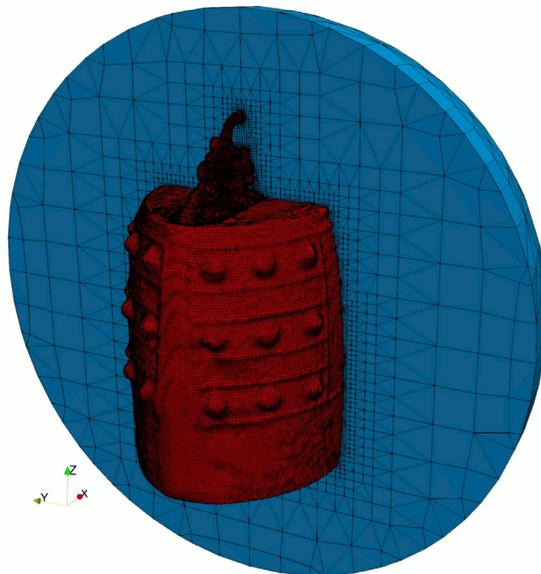


Figure 2.17: Model of an ancient Chinese bell from Liu et al. (2019a). Both the bell and a portion of the surrounding air are shown. Radiative boundary conditions are imposed at the outer surface of the modelled air.

tion in the time domain. This choice was based entirely on availability, as Rayleigh damping is often the only damping model available for dynamic time history modelling in FEM software. Details of the damping models used in this research are provided in Section 5.3.2.

The following sections provide a description of Rayleigh damping, as well as an alternative model that could be considered in certain scenarios. Damping is a complicated phenomenon attracting substantial research efforts. The approaches listed below are but a sample of those available in the literature.

2.3.1 Rayleigh Damping

Rayleigh damping was first outlined in the book *Theory of Sound* (Rayleigh, 1945), originally published in 1877, and remains popular today. Rayleigh damping computes the damping matrix, \mathbf{C} , as a weighted combination of the mass matrix, \mathbf{M} , and stiffness matrix, \mathbf{K} :

$$\mathbf{C} = \alpha\mathbf{M} + \beta\mathbf{K}, \quad (2.5)$$

where α and β are coefficients that determine how damping will vary with frequency. This variation with frequency can be calculated for Rayleigh damping as:

$$\zeta(\omega) = \frac{1}{2} \left(\frac{\alpha}{\omega} + \beta\omega \right), \quad (2.6)$$

where ζ is the damping ratio, expressing damping as a portion of critical damping, and ω is an input radial frequency. Looking at the terms inside the parenthesis of Equation 2.6 the $\frac{\alpha}{\omega}$ term represents the contribution of mass proportional damping, while the $\beta\omega$ term comes from stiffness proportional damping. Figure 2.18 gives an example of the damping coefficient from Equation 2.6 as a function of frequency, along with the individual mass proportional and stiffness proportional components. With the two parameters, α and β , available in Equations 2.5 and 2.6, Rayleigh damping can generally be calibrated to match damping values at two specific frequencies. The remaining damping values will follow the general shape shown in Figure 2.18. This model has remained popular in FEM codes as it is easy to scale and combine the global mass and stiffness matrices which are already required in the computation.

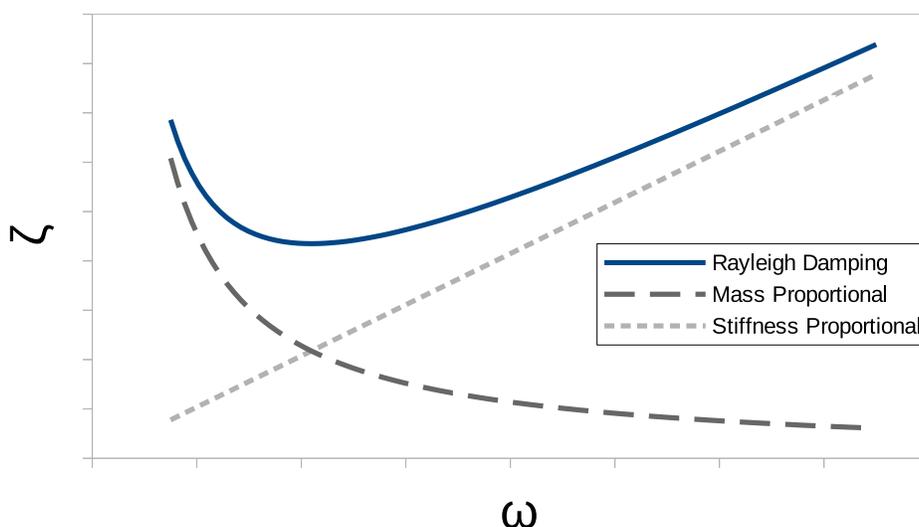


Figure 2.18: Plot of damping ratio, ζ , vs. angular frequency, ω , for Rayleigh damping. Contributions of mass proportional and stiffness proportional damping are shown separately.

2.3.2 Generalized Proportional Damping

Adhikari (2006) provides a versatile method for constructing a damping matrix called *Generalized Proportional Damping* (GPD). The method requires an input of any arbitrary function that computes damping ratio, ζ , as a function of radial frequency, ω . Such a function could be formed, for example, by fitting a polynomial curve to experimentally measured damping ratios at various frequencies. To compute the damping matrix, each instance of the radial frequency term, ω , in the input function is replaced by $\sqrt{\mathbf{M}^{-1}\mathbf{K}}$, where \mathbf{M} is the mass matrix and \mathbf{K} is the stiffness matrix. Any constant terms in the input function are multiplied by

the identity matrix, \mathbf{I} . With these matrix substitutions made, the resulting function is multiplied by $2\mathbf{M}\sqrt{\mathbf{M}^{-1}\mathbf{K}}$ and evaluated at $\sqrt{\mathbf{M}^{-1}\mathbf{K}}$ to produce the damping matrix. Thus, the method allows a damping matrix to be computed from the global mass and stiffness matrices, along with any function defining damping ratio as a function of radial frequency.

While this approach sounds very promising in its versatility, computational costs may become high depending on the complexity of the function selected to compute damping ratio as a function of frequency (matrix functions may be required). GPD was not available in any of the FEM packages considered for this thesis. Its use would have necessitated implementing all solving routines from scratch.

2.4 Tuning Techniques

The problem of tuning idiophone bar cutaway geometry can alternatively be formulated as minimizing an error function, or solving a system of nonlinear equations. Depending on the formulation, the problem can be solved using either optimization techniques, or methods of solving simultaneous nonlinear equations. Numerous tuning techniques were explored in this research, including both problem formulations. Ultimately the nonlinear system formulation was selected, as described in Section 4.2.2. The sections below reference sources which informed the solving method described in Section 4.2.2. Other methods considered or explored are also discussed.

2.4.1 Solving Nonlinear Systems

Most relevant to the bar tuning problem in this research are underdetermined systems of nonlinear equations in which each input variable has both a minimum and maximum bound (sometimes called simple bounds or box constraints in the literature). [Macconi et al. \(2009\)](#) outline a method for solving such systems. They describe their approach as:

... a Gauss-Newton method embedded into a trust-region strategy [taking] advantage of the so-called Coleman and Li affine-scaling matrix ...

The Coleman and Li affine-scaling matrix ([Coleman and Li, 1996](#)) appears regularly in the literature on solving such systems.

Levenberg-Marquardt methods, such as those described by [Kanzow et al. \(2004\)](#), provide an alternative approach of solving underdetermined systems of nonlinear equations subject to constraints.

2.4.2 Optimization

Optimization techniques tested in this research were well-known methods described in most textbooks on the subject. These included *gradient descent* and *sequential quadratic programming*, as outlined by Nocedal and Wright (2006).

Genetic algorithms (Goldberg, 1989) were considered early in this research work, with results reported in a conference paper Beaton and Scavone (2019b). While the literature on genetic algorithms is vast, several convenient surveys of the method and similar methods in a structural engineering context form a good starting point (Kicinger et al., 2005; Hare et al., 2013; Munk et al., 2015).

Chapter 3

Idiophone Bar Finite Element Models

This chapter gives a general overview of the idiophone bar finite element models created in this work. Similar models were used in both the tuning of torsional modes (Chapter 4) and hybrid modelling of sound radiation (Chapter 5). Different finite element packages were used to solve the models, depending on the task under consideration. Finite element programs were also changed over the course of this research, as modelling needs evolved. The governing equations, and their approximation using finite element methods, will not be repeated here, as such information is readily available in standard texts (Cook et al., 2001; Zienkiewicz et al., 2013).

The following sections outline some salient features of the models and analysis software used for various tasks and at various times in this work. To avoid ambiguity, modelling aspects specific to the hybrid sound radiation model (e.g. damping models, boundary conditions, mallet properties) are discussed in Chapter 5.

3.1 Programming Packages

Three programming packages were used to model idiophone bars in this work, as described below.

3.1.1 Pylaster

Pylaster is an FEA package written by the author in Python. This software was adapted from Matlab code used in modelling composite materials for the author's first master's thesis (Beaton, 2010). The package has also been used to automate analysis tasks in bridge engineering (Radojevic and Kirkwood, 2017). In this work *Pylaster* was extended to include features such as 20-node 3D solid elements and an eigenvalue solver for performing modal analyses.

Pylaster served as the central modelling program in this work. All bar models were originally defined in Pylaster. In cases where other packages were used, Pylaster models were exported to input file formats compatible with those other packages. Results were analyzed in Python scripts written specifically for this work, with output from the various analysis packages either stored in text-based output files or databases.

3.1.2 Calculix

Calculix (sometimes written in the stylized form “CalculiX”) is an open-source finite element analysis package originally written for Linux environments (Dhondt). Input file format in Calculix is similar to, and largely compatible with, the popular commercial FE program *Abaqus* (currently offered by Dassault Systemes). Calculix is largely written in FORTRAN, and thus benefits from the speed of this compiled language.

Calculix was primarily used around the beginning of this research, before a solver for eigenvalue problems was implemented in Pylaster. Bar tuning algorithms were programmed to work with either Calculix or Pylaster. Ultimately, Pylaster was selected for bar tuning as it allowed for easier implementation of semi-analytical modal frequency derivatives with respect to input variables. This resulted in faster run times for tuning models. Presently, at the end of each bar tuning analysis a single Calculix model is produced using the final geometry for record keeping purposes.

Calculix was also considered for dynamic time history models of mallet impacts. However, producing output at the desired time steps proved easier with the dynamic time-stepping capabilities of *Code_Aster* than those of Calculix.

3.1.3 Code_Aster

Code_Aster is an open-source software package for finite element analysis of structural and thermomechanical problems (Électricité de France). The program is developed and maintained by Électricité de France, a multinational electric utility company headquartered in France and operating in many countries around the world. The name “Code_Aster” is an acronym for the French description of the program’s capabilities: “Analyses des Structures et Thermomécanique pour des Études et des Recherches”. Overall, *Code_Aster* is a very large program with many available functionalities. Its core code is written in FORTRAN, while many modules are written in Python. Python code can be included in *Code_Aster* input files and will be interpreted by the program when processing those input files.

Code_Aster was selected for dynamic time history analysis of mallet impact models in this work. It was chosen based on its vast array of capabilities, compatibility with Python,

and flexibility in dynamically changing time steps in time history models. Two rationales ultimately lead to the decision to use Code_Aster:

1. Code_Aster was able to produce time history output at consistent time steps (desired to maintain a uniform sampling rate), while also dynamically altering the time step when required for model convergence. This did not seem possible with Calculix.
2. It was thought that learning and using Code_Aster would be quicker than implementing the necessary functionality (large displacements, large strains, contact mechanics) in Pylaster.

Whether or not rationale 2 was indeed true is debatable. Getting the mallet impact models to successfully run in Code_Aster required an enormous amount of time. Such slow progress with the program was a product of numerous factors. These factors are repeated below for any future reader who may consider using Code_Aster for similar work:

- Commands and error messages written in a foreign language.
- Documentation originally written in a foreign language and machine-translated to English.
- Memory (RAM) issues - Code_Aster would keep many results in memory, even if they were not requested as output, eventually exhausting memory limits. As a result, analyses had to be broken down into smaller sequential pieces (e.g. a model simulating 2 seconds of bar response to a mallet impact had to be run as a sequence of separate models, each with a length of 0.15 to 0.25 seconds).
- Disk space issues - Code_Aster would store vast amounts of data for each model run, even after attempts to instruct the program to limit stored data. This data would be duplicated in each sequential model described above. A simulation of 2 seconds, sampled at a full 48 kHz, may produce about 11 GB of necessary output, while producing hundreds of GB of unrequested temporary files. Available hard drive space on the Linux machine used to run Code_Aster was exhausted on multiple occasions. Only a single full model could be run at a time, with results transferred to other storage before running another full model.
- Time history analyses simulating one to two seconds of bar response were performed for comparison with hybrid radiation model results. These analyses are very computationally demanding, and would regularly take upwards of three or four days to run. Some of the problems described above would occur several days into an analysis. Thus troubleshooting a single model could take a very long time in itself.

3.2 Geometry Definition

All bar models were defined in Pylaster using a standard procedure to generate desired mesh density and cutaway geometry. Details are outlined in the following sections.

3.2.1 Mesh Definition

Bar models in this work have a structured finite element mesh comprised of 3D solid elements. Depending on the application, either 20-node or 8-node hexahedral elements were used. The numbers of elements in the length, width and thickness directions were specified as input parameters. Given these parameters, a Python script generated a mesh of the desired element type with uniform density over a unit cube. This cube was oriented in a natural coordinate system, (u, v, w) , such that $u, v, w \in [0, 1]$. An example of a meshed cube is shown in Figure 3.1 with 80, 7 and 3 elements in the length (u), width (v) and thickness (w) directions, respectively.

Most analyses in this work used a mesh of 8-node hexahedral elements. An import exception were the models employed in bar fabrication (including convergence checks of these models) which used 20-node elements.

3.2.2 Cutaway Definition

Cutaway geometry was defined by a thickness function, $t(u, v)$, that gives the cutaway thickness at natural coordinates (u, v) . Two different types of thickness functions were explored over the course of this work: cross-section interpolation and surface interpolation. Each interpolation type is described in the following sections.

Cross-Section Interpolation

Structural engineers commonly think of beams in terms of their cross-sections (i.e. the shape created if one were to cut a “slice” of the beam, perpendicular to its longitudinal axis). Idiophone bars are generally beam-like in their shapes and aspect ratios, thus it is natural to conceptualize their geometry as varying between some number of fixed cross-sections defined along their length.

Figure 3.2 gives an example of bar cutaway geometry defined by *cross-section interpolation*. For ease of understanding, the bar in Figure 3.2 is shown in model coordinates, though in practice the interpolation occurs in natural coordinates. Each cross-section in the figure is outlined in blue. The red line demonstrates interpolating bar thickness along the bar’s

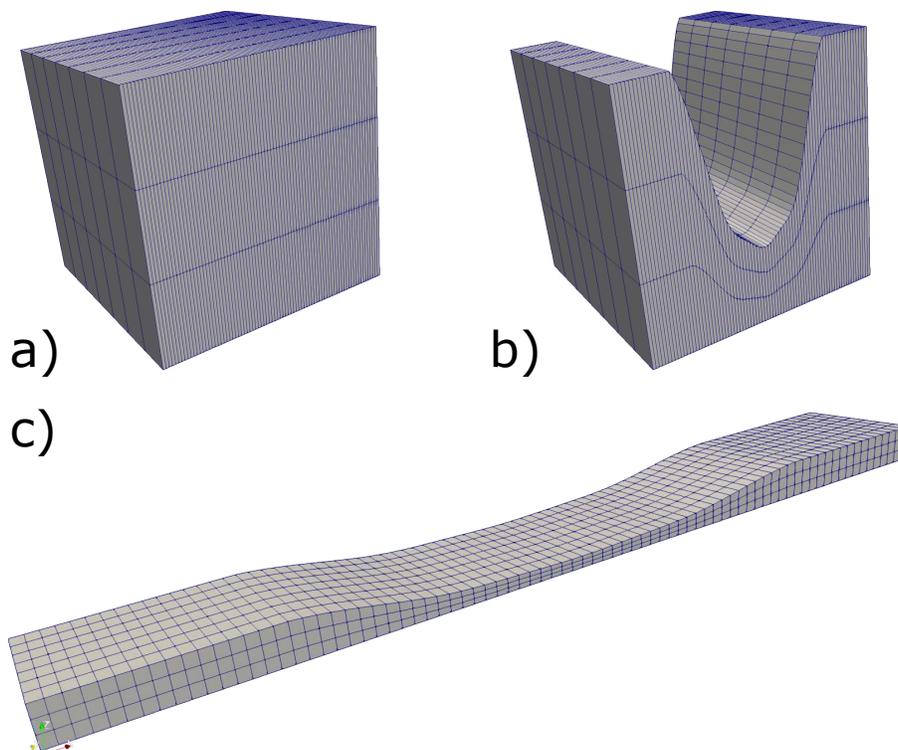


Figure 3.1: These panels depict the stages of creating a meshed bar model: (a) a structured mesh of uniform density is generated over a unit cube in natural coordinates (u, v, w) , this example has $(80, 7, 3)$ elements in the (length, width, thickness) directions; (b) cutaway shape is defined by scaling w -coordinates of the nodes according to a thickness function $t(u, v)$; (c) final bar geometry is generated by scaling nodal positions from natural coordinates (u, v, w) to model coordinates (x, y, z) via Equation 3.1.

length (x - or u -coordinate) for a fixed position along its width (y - or v -coordinate). This line, which is a contour of the function $t(u, v)$ for a constant value v , is defined using cubic spline interpolation between the cross-sections. The thickness of each cross-section is evaluated at the specified v -coordinate. These thickness values and cross-section locations are then used to interpolate thickness along the u -axis. Cubic interpolation was provided by the Python function *interp1d* in the *scipy.interpolate* module.

Cross-section interpolation was the first type of cutaway definition considered in this work. It was used mainly to tune vertical bending modes only, demonstrating how frequency ratios of untuned torsional modes may evolve along the keyboard (Beaton and Scavone, 2019b).

Ultimately, this interpolation type provided limited flexibility in tuning torsional modes, as achieving a wide range of modal tuning ratios would require an even wider range of possible input cross-sections (each of which having one or more tunable parameters). Thus, after producing valuable early results, cross-section interpolation was abandoned in favour of the

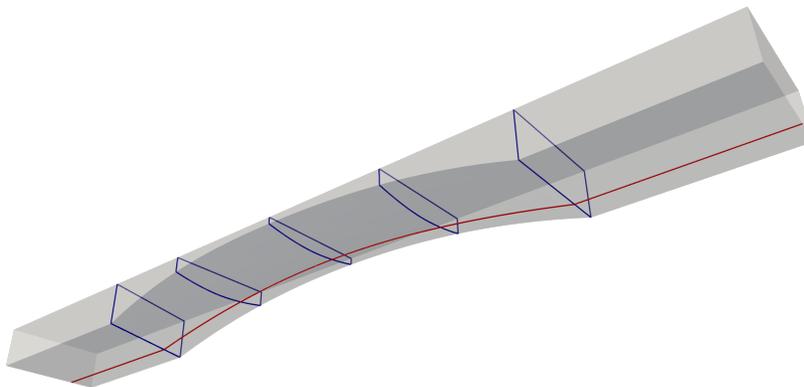


Figure 3.2: Example of a bar cutaway defined by cross-section interpolation. Cross-sections are outlined in blue. The red line shows an example of thickness interpolated between cross-sections along the length of the bar. This figure is borrowed from [Beaton and Scavone \(2019b\)](#).

more versatile surface interpolation described in the next section.

Surface Interpolation

To provide more flexibility in defining sophisticated cutaway geometry (and thus a larger range of achievable tuning ratios) *surface interpolation* was employed. In this method, a grid of *control points* was defined over the cutaway’s surface. Each point served to control the cutaway’s thickness locally, with the remaining geometry determined using a 2D interpolation scheme. Figure 3.3 gives an example layout of control points on a cutaway surface. Panel (a) of the figure begins with a basic cutaway geometry with control points shown in three colours. The blue and red points each cover one quarter of the cutaway surface, while the green points cover half. If desired, these groups can be used to enforce cutaway symmetry, and reduce the number of inputs used to control geometry. All models considered in this work took advantage of this symmetry. The blue control points of Figure 3.3 were used as inputs. These were mirrored across the width of the beam to produce the red control points. Both red and blue points were then mirrored to produce the green control points. Figure 3.3 (b) gives an example of a more complex geometry that can be produced using this approach. The sophisticated cutaway geometry of panel (b) was produced using only the blue control points shown and mirroring as described. This example shows the versatility of surface interpolation, as producing a similar geometry using cross-section interpolation would clearly be difficult.

Two methods of 2D interpolation between control points were explored. The first method

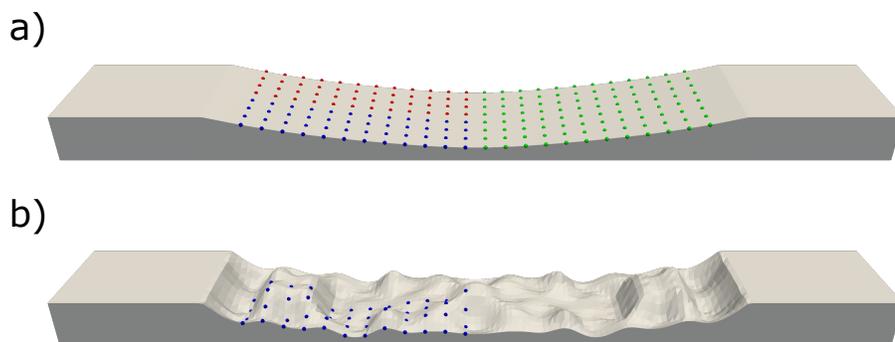


Figure 3.3: Example of a bar cutaway defined by surface interpolation. This example demonstrates the use of control point mirroring to enforce double symmetry. The blue dots shown in panel (a) serve as the control points for bar tuning. Red dots are mirrored versions of the blue control points. Green dots are mirrored versions of both the blue and red points. Panel (b) shows a more complex geometry that can be attained by adjusting the blue control points and mirroring twice. This example uses Clough-Tocher interpolation, thus the control points lie directly on the cutaway surface.

considered was that provided by the Python function *griddata* in the *scipy.interpolate* module. This function implements a Clough-Tocher interpolation scheme, a description of which may be found in Farin (1986). This method was used to produce cutaway geometries in Beaton and Scavone (2021a), as well as Figure 3.3. With this interpolation scheme, control points lie directly on the cutaway surface, making their manipulation an intuitive exercise. One drawback of this approach was that, in certain scenarios, the interpolated surface between control points could extend above or below the maximum and minimum control point coordinates. Thus limits on minimum bar thickness over the cutaway, which were enforced via limits on control point coordinates, could sometimes be violated using this interpolation.

The second 2D interpolation method explored was non-uniform rational B-spline (NURBS) surfaces (Piegl and Tiller, 1995). Switching to this interpolation was necessary for compatibility with computer aided manufacturing software used to fabricate prototype bars (see Section 4.5). The most salient difference between NURBS and Clough-Tocher interpolation is that with NURBS, control points may not (and often *will* not) be positioned directly on the cutaway surface. Instead, the control points form a “control polygon”, which roughly approximates the surface itself. One benefit of this structure is that the NURBS surface cannot extend beyond the extrema of its control polygon. Thus using NURBS surfaces to define cutaway geometry will also ensure strict adherence to thickness limits enforced by limits on control point coordinates.

Examples of cutaway geometries defined using NURBS are provided in Section 4.5.

Table 3.1: Material properties of aluminum and rosewood used in this work. Symbols include: density, ρ , Young’s modulus, E , Poisson’s ratio, ν , and shear modulus, G . Rosewood properties are adopted from [Bork et al. \(1999\)](#). Aluminum properties are those of the standard 6061 alloy ([Battelle Memorial Institute, 2020](#)). This table is borrowed from [Beaton and Scavone \(2021a\)](#).

Material	ρ [kg/m ³]	E [GPa]	ν	G [GPa]
Aluminum	2700	68.9	0.33	25.9
Rosewood	1080			
x-axis		23.0		
y-axis		2.3		
z-axis		1.15		
xy-plane			0.30	3.0
yz-plane			0.60	1.0
xz-plane			0.45	3.0

3.2.3 Coordinate Scaling

The final step in creating a meshed bar model was to convert the meshed cube, with its applied cutaway, from natural coordinates (u, v, w) to model coordinates (x, y, z) . By virtue of the natural coordinates being defined over a range $[0, 1]$, converting to model coordinates was a simple matter of scaling:

$$\begin{Bmatrix} x \\ y \\ z \end{Bmatrix} = \begin{bmatrix} L & 0 & 0 \\ 0 & W & 0 \\ 0 & 0 & T \end{bmatrix} \begin{Bmatrix} u \\ v \\ w \end{Bmatrix}, \quad (3.1)$$

where L is the overall length of the bar, W is the width of the bar, and T is the end thickness (i.e. the thickness everywhere outside the cutaway).

3.3 Material Properties

Aluminum and rosewood idiophone bars were considered in this work, reflecting the materials typically used for vibraphones and marimbas, respectively. Table 3.1 outlines the properties of these materials generally used for bar models in this work. An important exception is the models used for bar fabrication. In that case, material properties were based on measurements of the stock aluminum used for fabrication. This adjustment was made to maximize tuning accuracy of the fabricated bars. Further info on bar measurements and stock bar properties is provided in Appendix A.

Chapter 4

Three-Dimensional Tuning of Idiophone Bar Modes

4.1 Objective

As described in Section 1.2.1, this work seeks to remedy the problem of untuned torsional modes polluting bar timbre. It is preferable to accomplish this goal solely by shaping bar cutaways. Other potential approaches, such as adding concentrated mass or reinforcing materials, would complicate the design and could introduce new points of failure (e.g. adhesives).

Untuned torsional modes become problematic when their modal frequencies come close to those of tuned flexural modes. To alleviate this issue, the method developed in this chapter shapes bar cutaways to tune both flexural and torsional modes concurrently.

The method described in this chapter was published by [Beaton and Scavone \(2021a\)](#) in a special issue of the *Journal of the Acoustical Society of America* (JASA). This article will be widely cited in the sections that follow.

4.2 Method

4.2.1 Frequency Targets

Alleviating timbral problems caused by untuned torsional modes can be accomplished by separating their modal frequencies from those of the tuned flexural modes. This begs the question: “by how much must these frequencies be separated to produce suitable bar timbre?” The answer to this question is both perceptual, and subjective, as presently there is no objective measure of “suitable bar timbre”.

While examples of studies on bar timbre may be found in the literature (Bork and Meyer, 1985; McAdams et al., 2004; Aramaki et al., 2007), perceptual investigation of the effects of torsional modes with frequencies similar to flexural modes was not considered in this work. Instead, this chapter takes as a given that an appropriate frequency target can be set for each mode to be tuned. In practice, one can always choose an arbitrarily large difference between flexural mode frequencies and torsional mode frequencies that would avoid timbral issues. In this chapter modes will be tuned to integer multiples of a bar’s fundamental frequency. Thus any two tuned modes will either be set to have the same frequency, or differ by some multiple of the fundamental.

Setting two modes to have equal frequency may produce interesting theoretical results, as will be shown. In practice, however, such bars would be difficult to fabricate with the precision required to produce effectively equal frequencies in the two modes. More likely the two frequencies would be close, but unequal, yielding beating effects.

Frequency targets in this work are thus set in three steps:

1. Select the modes to be tuned along with tuning ratios for their frequencies.
2. Set the bar’s fundamental frequency based on the desired musical note and tuning standard (e.g. A440, A442, etc.).
3. Multiply tuning ratios by the fundamental frequency to produce target frequencies for each tuned mode.

Throughout this chapter the terms “target frequencies” and “target frequency ratios” or “target tuning ratios” may be used interchangeably; note that the latter two refer to the same values, and can be converted to/from the first by multiplying/dividing by the fundamental frequency.

4.2.2 Tuning Approach

Note that the nomenclature and symbols adopted in this chapter will largely match that of Beaton and Scavone (2021a). For brevity, only direct quotes from that paper will be cited; equations will not be cited.

Problem Formulation

The problem of finding a cutaway geometry that tunes modal frequencies to the desired targets will be formulated as follows. Let \mathbf{x} be a vector of input parameters that define the bar cutaway. For the cutaway definition types described in Section 3.2.2, \mathbf{x} will either be

a vector of values defining cross-sectional properties at specified positions, or a vector of thickness values defining the cutaway surface at discrete points. Let N be the total number of input parameters defining the cutaway geometry. For the n^{th} input variable, x_n , minimum and maximum limits are established such that:

$$T_{min} \leq x_n \leq T_{max}, n \in [1, N], \quad (4.1)$$

where T_{min} and T_{max} are the minimum and maximum allowable values of input parameter x_n . Equation 4.1 assumes all input parameters in \mathbf{x} are of the same type and subject to the same limits. Individual limits for each input parameter are easily accommodated, with no difference in the remaining steps of this procedure. Limits to the input variables may be assigned for practical purposes (e.g. bar cutaway thickness should not exceed the thickness of available stock materials), strength purposes (e.g. exceeding some minimum thickness may produce an easily damaged bar) or other purposes.

Let K be the total number of modes tuned by the desired output bar geometry. For the k^{th} mode, the *tuning error function*, $g_k(\mathbf{x})$, is defined as:

$$g_k(\mathbf{x}) = \frac{f_{model,k}(\mathbf{x}) - f_{target,k}}{f_{target,k}}, \quad (4.2)$$

where $f_{model,k}(\mathbf{x})$ is the frequency of the k^{th} tuning mode in a model with geometry input vector \mathbf{x} , and $f_{target,k}$ is the target frequency of tuning mode k . The problem of tuning a bar's geometry to produce the desired modal frequencies is then formulated as:

$$\text{set } \mathbf{g}(\mathbf{x}) = \mathbf{0}, \quad (4.3)$$

where $\mathbf{g}(\mathbf{x})$ is a vector of the K tuning error functions described by Equation 4.2. The problem defined by Equation 4.3 is subject to the limits on input variables in Equation 4.1.

Tuning Algorithm

In this approach it is assumed that $K \leq N$, so there will be at least as many input parameters as there are tuning modes. The example bar shown in Figure 3.3 uses 48 points to define the bar's cutaway. Typical marimba or vibraphone bars will tune up to three modes. In this work the maximum number of modes tuned is six. Thus, in the problems considered herein, $K \leq 6 \ll 48 = N$. Any time $K < N$ the system in Equation 4.3 becomes underdetermined and multiple solutions may exist.

While it is possible to solve Equation 4.3 using optimization methods to minimize its norm,

herein the system is solved directly, using an iterative, gradient-based approach. At any given iteration, the next estimated input variable vector to solve Equation 4.3 is computed as:

$$\mathbf{x}_{i+1} = \mathbf{x}_i - [\mathbf{J}(\mathbf{x}_i)\mathbf{D}(\mathbf{x}_i)]^+ \mathbf{g}(\mathbf{x}_i), \quad (4.4)$$

$$\mathbf{J}(\mathbf{x}) = \begin{bmatrix} \frac{\partial g_1}{\partial x_1} & \frac{\partial g_1}{\partial x_2} & \dots & \frac{\partial g_1}{\partial x_N} \\ \frac{\partial g_2}{\partial x_1} & \frac{\partial g_2}{\partial x_2} & \dots & \frac{\partial g_2}{\partial x_N} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial g_K}{\partial x_1} & \frac{\partial g_K}{\partial x_2} & \dots & \frac{\partial g_K}{\partial x_N} \end{bmatrix}, \quad (4.5)$$

$$\mathbf{D}(\mathbf{x}) = \begin{bmatrix} b_1(\mathbf{x}) & & & \\ & b_2(\mathbf{x}) & & \\ & & \ddots & \\ & & & b_N(\mathbf{x}) \end{bmatrix}, \quad (4.6)$$

$$b_n(\mathbf{x}) = \begin{cases} 0 & \text{if } [\mathbf{J}^+(\mathbf{x})\mathbf{g}(\mathbf{x})]_n > 0, x_n = T_{min}, \\ & \text{or } [\mathbf{J}^+(\mathbf{x})\mathbf{g}(\mathbf{x})]_n < 0, x_n = T_{max}, \\ 1 & \text{otherwise} \end{cases} \quad (4.7)$$

where $\mathbf{J}(\mathbf{x})$ is a Jacobian matrix of partial derivatives of the tuning error functions with respect to the geometry input parameters, and $\mathbf{D}(\mathbf{x})$ is a diagonal matrix of binary terms defined by Equation 4.7. This matrix is similar in concept to the scaling matrix used by [Macconi et al. \(2009\)](#), though its details differ. Note that the superscript “+” after a matrix denotes the Moore-Penrose inverse (or *pseudoinverse*) ([Ben-Israel, 2002](#); [Penrose, 1955](#)) of that matrix. Given some initial input vector, \mathbf{x}_0 , Equation 4.4 is applied iteratively until all terms in $\mathbf{g}(\mathbf{x})$ are within defined tolerances of zero.

Matrix $\mathbf{D}(\mathbf{x})$ is designed to aid in convergence of Equation 4.4. Any values in vector \mathbf{x}_{i+1} that are pushed beyond the limits T_{min} or T_{max} will be set equal to the limit they have surpassed. Convergence will slow as values in \mathbf{x}_{i+1} reach one of these limits ([Beaton and Scavone, 2021a](#)). Matrix $\mathbf{D}(\mathbf{x})$ helps to mitigate any reduction in convergence rate by “deactivating” those inputs that have hit their limits, thereby forcing greater reliance on inputs that still have room for adjustment.

Jacobian Matrix

Matrix $\mathbf{J}(\mathbf{x})$ will be a dense matrix with dimensions $K \times N$. Populating this matrix will be the most time consuming operation in each evaluation of Equation 4.4.

The partial derivative terms $\left\{ \frac{\partial g_k}{\partial x_n} \right\}$ can be computed numerically or analytically. If these terms are evaluated numerically, using first-order finite different methods, the function $\mathbf{g}(\mathbf{x})$ will need to be computed $N + 1$ times per iteration. Computing $\mathbf{g}(\mathbf{x})$ requires performing a finite element modal analysis to determine tuning mode frequencies for a given vector of input parameters. Depending on the size of the finite element model, performing this modal analysis $N + 1$ times could require considerable time, particularly for $N = 48$, as was used in this work.

Computing time can be reduced, to some extent, by instead populating $\mathbf{J}(\mathbf{x})$ analytically. To do so, consider a single term within $\mathbf{J}(\mathbf{x})$:

$$\frac{\partial g_k}{\partial x_n} = \frac{\partial}{\partial x_n} \left(\frac{f_{model,k}(\mathbf{x}) - f_{target,k}}{f_{target,k}} \right) = \frac{1}{f_{target,k}} \frac{\partial f_{model,k}(\mathbf{x})}{\partial x_n}. \quad (4.8)$$

For brevity, explicit dependence on vector \mathbf{x} will be omitted from the following equations for populating matrix $\mathbf{J}(\mathbf{x})$. Frequency $f_{model,k}$ is calculated from the equation:

$$f_{model,k} = \frac{\sqrt{\lambda_k}}{2\pi}, \quad (4.9)$$

where λ_k , corresponding to tuning mode k , is an eigenvalue that solves the equation:

$$(\mathbf{K} - \lambda_j \mathbf{M}) \boldsymbol{\phi}_j = 0, \quad (4.10)$$

where \mathbf{K} is the system stiffness matrix of a finite element model of the idiophone bar, \mathbf{M} is the system mass matrix, λ_j is the j^{th} eigenvalue of the system and $\boldsymbol{\phi}_j$ is the j^{th} eigenvector. Note that k and j need not be equal, as the K tuned modes will not necessarily be the first K bar modes. Determining which mode number, j , corresponds to which tuning mode number, k , is discussed in Section 4.2.3 on mode identification.

Taking the partial derivative of Equation 4.9 with respect to input parameter x_n yields:

$$\frac{\partial f_{model,k}}{\partial x_n} = \frac{1}{4\pi\sqrt{\lambda_k}} \frac{\partial \lambda_k}{\partial x_n}. \quad (4.11)$$

The partial derivative of eigenvalue λ_k with respect to input parameter x_n , in Equation 4.11

can be evaluated as (for derivation see [Lund \(1994\)](#) or [Yang and Peng \(2020\)](#)):

$$\frac{\partial \lambda_k}{\partial x_n} = \boldsymbol{\phi}_j^T \left(\frac{\partial \mathbf{K}}{\partial x_n} - \lambda_j \frac{\partial \mathbf{M}}{\partial x_n} \right) \boldsymbol{\phi}_j. \quad (4.12)$$

A detailed discussion of evaluating the partial derivative of the system stiffness matrix, \mathbf{K} , with respect to an input parameter, x_n , is provided by [Wang et al. \(1985\)](#). For the finite element models in this work, the system mass matrix, \mathbf{M} , is computed as the sum of element mass matrices, \mathbf{M}_e . Thus the partial derivative of the system mass matrix, \mathbf{M} , with respect to x_n , is equal to the sum of the partial derivatives of the element mass matrices, \mathbf{M}_e , with respect to x_n . An element mass matrix is computed as:

$$\mathbf{M}_e = \rho \int_{-1}^1 \int_{-1}^1 \int_{-1}^1 \mathbf{N}^T \mathbf{N} \det \mathbf{A}_e d\xi d\eta d\zeta, \quad (4.13)$$

where \mathbf{M}_e is the element mass matrix, ρ is the material density, \mathbf{N} is a matrix of finite element shape functions, \mathbf{A}_e is a Jacobian matrix containing partial derivatives of model coordinates with respect to natural coordinates ξ , η and ζ . Here the symbol \mathbf{A}_e is used for the finite element Jacobian matrix to avoid confusion with matrix $\mathbf{J}(\mathbf{x})$ from Equation 4.4.

The shape function matrix \mathbf{N} in Equation 4.13 is a function only of the natural element coordinates, ξ , η and ζ , and is independent of any input parameter, x_n . Thus the derivative of matrix \mathbf{M}_e with respect to x_n is:

$$\frac{\partial \mathbf{M}_e}{\partial x_n} = \rho \int_{-1}^1 \int_{-1}^1 \int_{-1}^1 \mathbf{N}^T \mathbf{N} \frac{\partial \det \mathbf{A}_e}{\partial x_n} d\xi d\eta d\zeta. \quad (4.14)$$

Equation 4.14 applies for cases where the material density, ρ , is not a function of any input parameter. Calculation of the term $\frac{\partial \det \mathbf{A}_e}{\partial x_n}$ is discussed by [Wang et al. \(1985\)](#). Its calculation involves computing partial derivatives of finite element nodal coordinates with respect to input parameters. It is in the calculation of this term, which also appears in the partial derivative of element stiffness matrices, that the effect of input parameters on cutaway geometry (nodal coordinates) comes into play.

In this work, partial derivatives of nodal coordinates with respect to input parameters were evaluated numerically using finite different techniques. From there, all terms in Equation 4.12 could be evaluated, followed by Equation 4.11, and finally Equation 4.8, which populates matrix $\mathbf{J}(\mathbf{x})$. This semi-analytical approach, while somewhat extensive in its implementation, did save some computation time compared to populating matrix $\mathbf{J}(\mathbf{x})$ completely numerically.

Initial Geometry Selection

When the system in Equation 4.3 is underdetermined (i.e. when $K < N$), the pseudoinverse in Equation 4.4 has the desirable property of returning the next estimated trial solution, \mathbf{x}_{i+1} , that is nearest the current trial solution, \mathbf{x}_i . This can be shown by rearranging Equation 4.4 to the following:

$$[\mathbf{J}(\mathbf{x}_i)\mathbf{D}(\mathbf{x}_i)](\mathbf{x}_i - \mathbf{x}_{i+1}) = \mathbf{g}(\mathbf{x}_i), \quad (4.15)$$

which has the familiar general form:

$$\mathbf{A}\mathbf{x} = \mathbf{b}. \quad (4.16)$$

It can be shown that if the general system in Equation 4.16 is underdetermined (with multiple solutions possible), then applying the pseudoinverse:

$$\mathbf{x} = \mathbf{A}^+\mathbf{b}, \quad (4.17)$$

yields the solution, \mathbf{x} , with minimum norm (Ben-Israel and Greville, 2003). Applying this result to Equation 4.15 implies that Equation 4.4 produces the new estimate, \mathbf{x}_{i+1} , that minimizes $\|\mathbf{x}_i - \mathbf{x}_{i+1}\|_2$, making it the estimated new trial solution that lies closest to the current trial solution, \mathbf{x}_i .

The overall effect of this property is that the tuning algorithm will generally seek suitable bar geometries (those that tune modes as desired) similar to whatever bar geometry is present in the current iteration. As a result, the initial input vector, \mathbf{x}_0 , selected by the analyst will affect the tuned bar geometry produced by the algorithm. For this reason, it is recommended that the analyst select a desirable bar shape (by whatever criterion they should choose to define “desirable”) for the initial geometry, as the final geometry produced by the algorithm will tend to be similar to this initial geometry (insofar as is feasible, while also tuning the modes). Examples of the impact of initial bar geometry are provided in Section 4.3.1.

4.2.3 Mode Identification

The ability to identify modes in the finite element results is an essential aspect of successfully tuning bar cutaway geometry. Finite element modal analyses will return mode shapes and frequencies sorted in ascending order of frequency. Thus the third mode reported will always be the one with the third-lowest frequency, and so on. Depending on the initial geometry provided, it is entirely possible for modes to reorder over the course of tuning bar geometry.

For example, the tuning algorithm in this chapter may be provided an initial geometry

where mode V3 is the 5th mode overall (i.e. it has the 5th lowest modal frequency in the model). However, in the final geometry tuned by the algorithm, mode V3 may now be the 6th mode overall. For this reason, it is insufficient to instruct the algorithm to tune the 5th mode to a specified frequency. Instead, tuning objectives must be defined for modes based on their mode label (V1, T3, etc.), and the algorithm must identify a mode's label when analyzing finite element model results.

Discussion of these methods is also provided in Section IV of [Beaton and Scavone \(2021a\)](#).

Shape Recognition

Emulating the manner in which an analyst would manually identify modes, an algorithm was implemented to determine mode type based on the shape formed by nodal displacements. This algorithm functioned very well for bars of uniform thickness, but could break down for bars with long, deep cutaways. The modes of such bars can exhibit coupling between deformation types. For example, the axial mode in [Figure 1.2](#) includes significant vertical deformation, while the lateral mode in that same figure includes noticeable torsional deformation.

The following paragraphs discuss how the mode types shown in [Figure 1.2](#) were identified via shape recognition in this work. The x-, y- and z-axes used to define model coordinates are shown in [Figure 4.1](#).

Axial modes were the simplest mode type to identify in this work, as only the first axial mode (A1) appeared over the frequency range of interest. If a mode's maximum absolute nodal displacement in the x-direction was larger than the maximum absolute displacements in the y- and z-directions, the mode was identified as axial ([Beaton and Scavone, 2021a](#)).

Flexural modes are of primary interest in this work, as vertical-flexural modes are tuned in every model. Their identification is described succinctly by [Beaton and Scavone \(2021a\)](#):

Flexural modes were identified by locating and assessing areas with low displacement magnitude (i.e., the norm of the three displacement components); these are the dark blue areas shown in [[Figure 4.1](#)] and throughout the paper. First, all nodes with displacement magnitude below a specified threshold were marked (pink dots, [[Figure 4.1](#)]). These nodes were then separated into contiguous groups, based on element connectivity. Then, a line of best fit was calculated for each group (green lines, [[Figure 4.1](#)]). Identification of flexural modes was based on the number and orientation of these lines. Two or more groups, each

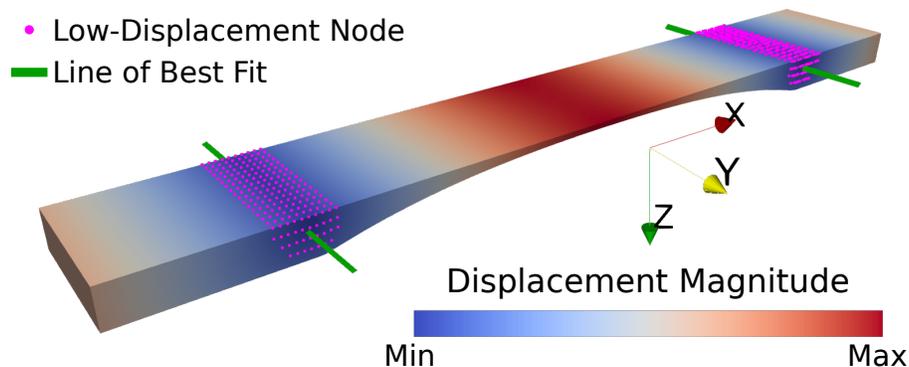


Figure 4.1: Example of identifying a vertical-flexural mode via shape recognition based on node groups with low displacement. Figure from [Beaton and Scavone \(2021a\)](#).

with a best-fit line parallel to the y axis indicates a vertical-flexural mode. Similarly, two or more groups, all with best-fit lines parallel to the z -axis indicates a lateral-flexural mode. For flexural modes, the mode number (e.g., third vertical mode, second lateral mode) is one less than the number of low-displacement groups.

Torsional modes were also of great interest in this work, as the method described in this chapter was developed to enable their tuning concurrent with other mode types. Again, from [Beaton and Scavone \(2021a\)](#):

For flat and near-flat bars, over the frequency ranges considered, it was possible to identify torsional modes largely via process of elimination. After ruling out axial and flexural modes, any remaining modes were identified as torsional. For bars with significant cutaways, where torsional deformation may be present in lateral modes, a check was performed on the rotation of bar ends. If they were found to rotate to a greater extent about the z -axis than the x axis, the mode was instead labelled as lateral-flexural. This method of identifying torsional modes by process of elimination was feasible for this work as other mode types, outlined below, appeared only in frequency ranges well above any modes that were tuned.

Modes not considered include axial modes along the y - and z -axes, torsional modes about those axes, and flexural modes in the y - z plane. In this work, frequencies of these mode types occur well above the range of interest for bar tuning, which was generally below 6 kHz. As such, these mode types did not appear in finite element results in this work, rendering their identification unnecessary.

Modal Assurance Criterion

A popular tool for the identification of mode shapes is the *Modal Assurance Criterion* (MAC) (Allemang, 2003). This criterion essentially takes the dot product between two mode shapes, using the result as a measure of similarity between the two. Compared to shape recognition, the MAC has the advantage of being much faster to compute, with the drawback that its use requires a set of known reference modes for comparison.

In this work the MAC was used in conjunction with shape recognition to leverage the advantages of each. Within a given bar tuning analysis, mode shapes of the initial bar geometry (i.e. the 0th iteration) were identified via shape recognition. This was generally successful as initial bar geometries tended to be relatively simple. In subsequent iterations, the MAC was used to identify modes, with mode shapes identified in the previous iteration used for reference.

This combination of shape recognition and MAC was generally sufficient to identify modes during bar tuning. However, caveats could occur in two scenarios:

Degenerate modes occur when two modes have very similar frequencies. In the literature this is also known as the problem of repeated eigenvalues. In such cases, the mode shape of each degenerate mode will appear as a weighted combination of the mode shapes from the two modes with similar frequencies.

Mixing two mode shapes in this manner renders identification by shape recognition infeasible. It can even be difficult for an analyst to identify such modes manually. The MAC also struggles in such cases as there will be rows and columns in the MAC matrix with two similar values, thus producing ambiguity (Beaton and Scavone, 2021a).

Degenerate modes are most likely to appear when tuning two modes to have the same frequency ratio. However, they can also occur if the iterative tuning process causes two modes to reorder, and they end up with similar frequencies during an intermediate tuning step.

In this work, degenerate modes were identified by mapping from the bar geometry that caused them to a geometry with known modes. This process is described later in this section.

Similar modes can also cause ambiguity when attempting to automatically identify mode shapes. This can occur, in particular, for bars with deep cutaways. In such cases, coupling between deformation types can produce modes with very similar shapes. Figure 4.2 shows an example of very similar mode shapes that are not degenerate and, in fact, have significantly different modal frequencies. While difficult to identify using shape recognition or the MAC,

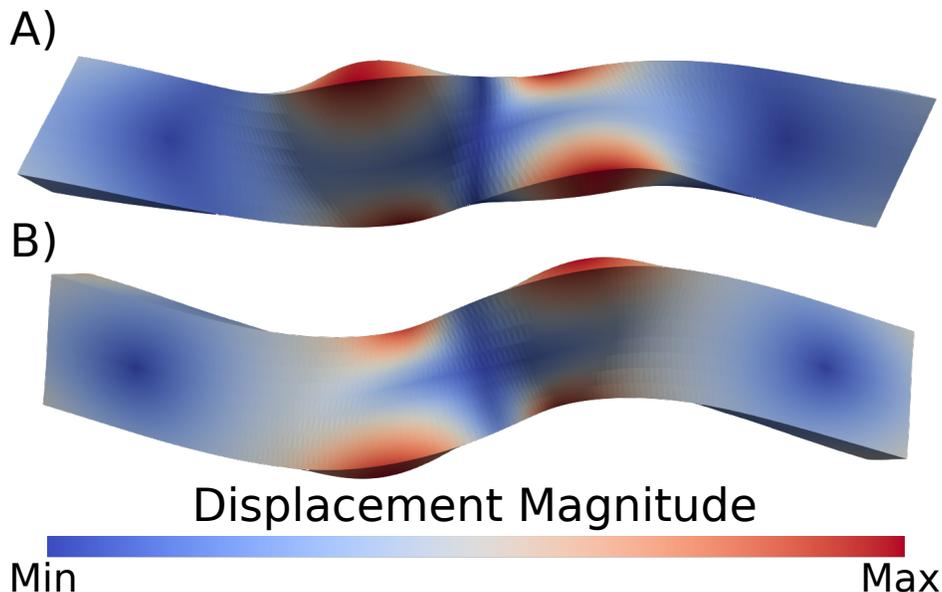


Figure 4.2: Similar but not degenerate mode shapes. These modes are from a rosewood bar model with vertical-flexural modes tuned to the typical 1:4:10 ratios. Coupled lateral and torsional deformations make identifying the modes exceedingly difficult without further information. (a) Mode L2, $f = 2260$ Hz. (b) Mode T3, $f = 2759$ Hz. Figure from [Beaton and Scavone \(2021a\)](#).

similar modes like these are easily determined by mapping to known modes, as described below.

Mapping Known Modes

The procedure for identifying degenerate modes by mapping from a bar model with known modes is described concisely by [Beaton and Scavone \(2021a\)](#):

In cases where degenerate modes were encountered a third, robust yet labor intensive mode identification approach was employed. In this method, the current geometry, with degenerate modes, is mapped to a (typically simpler) reference model with known modes. A mapping variable, m , is defined to compute input vectors for intermediate models between the reference model and the current model. These intermediate input vectors are calculated as:

$$\mathbf{x}_{\text{int}} = \mathbf{x}_{\text{ref}} + m (\mathbf{x}_{\text{deg}} - \mathbf{x}_{\text{ref}}), \quad (4.18)$$

where \mathbf{x}_{int} is the intermediate model input vector, \mathbf{x}_{ref} is the reference model input vector, m is the mapping variable, and \mathbf{x}_{deg} is the degenerate model input vector. From here the method works by using the known modes at $m = 0$ to identify

modes in models with progressively larger values of m using MAC. The degenerate modes can then be identified in one of two ways. Both options involve generating and evaluating intermediate models with values of m progressively closer to 1. MAC is used to identify modes in the intermediate models, using the model with the next-highest value of m as reference (beginning with the known model at $m = 0$). In the first option, this process continues with progressively larger values of m until an intermediate model is produced that enables identifying the degenerate modes (at $m = 1$) using MAC.

If identifying the degenerate modes in this manner using MAC proves problematic, an alternative is to generate and evaluate enough intermediate models to enable approximating the frequency of a given mode (e.g., mode V3) as a function of m . These functions can then be used to predict the frequency of a given mode type at $m = 1$. Additional intermediate models are generated until the resulting frequency predictions align with the degenerate mode frequencies sufficiently well to identify the modes.

Corner Displacements

Several months after the method described in this chapter was published (Beaton and Scavone, 2021a), another paper on the tuning of torsional modes appeared, authored by Soares et al. (2021b). While the methods used to define and tune bar geometry in that paper are markedly different than those used here, the need to automatically identify mode types is shared.

The method used by Soares et al. (2021b) to identify mode shapes considers only the displacement of nodes at the two corners of one end of a bar. They describe their method as follows:

In most cases, the same mode type classification used in uniform beams can be made with regard to undercut beams. However, in some particular geometries, this formal classification might become ambiguous as some mode shapes pertain to combinations of different modal families. This typically occurs when two modes of distinct families are close in frequency. In these cases, care must be taken to correctly identify the modal family associated with each mode. Here, a set of simple and reliable criteria was used, based on the evaluation of the 3-D mode shape $\boldsymbol{\psi}(s) = [\psi_x(s), \psi_y(s), \psi_z(s)]$ at two corner points s_1 and s_2 , as illustrated in [Figure 4.3].

The modal family identification process goes as follows:

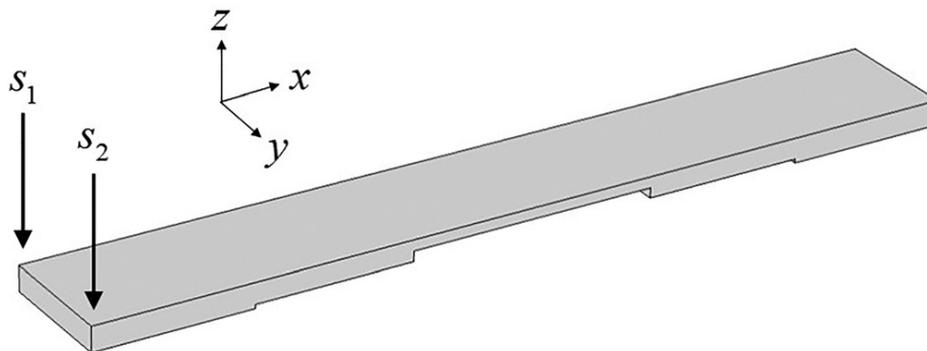


Figure 4.3: Corner nodes used to identify mode shapes, as shown in Soares et al. (2021b). Visible on the bar’s side is the step-wise cutaway definition employed in their paper.

- Lateral-bending modes and axial modes are identified directly if any of the conditions are true: $\max(|\boldsymbol{\psi}(s)|) = |\psi_y(s_1)|$ (lateral) or $\max(|\boldsymbol{\psi}(s)|) = |\psi_x(s_1)|$ (axial).
- If none of the above are true, i.e., $\max(|\boldsymbol{\psi}(s)|) = |\psi_z(s_1)|$, the mode can be either vertical-bending or torsional. Then the signs of $\psi_z(s_1)$ and $\psi_z(s_2)$ provide a disambiguation, i.e., if $\text{sign}(\psi_z(s_1)) = \text{sign}(\psi_z(s_2))$, the mode is vertical-bending; otherwise, if $\text{sign}(\psi_z(s_1)) \neq \text{sign}(\psi_z(s_2))$, the mode is torsional.

So far, the authors have seen no mis-identification using these criteria, mainly because, even in the troubling scenario of coupled modal families, there tends to be a clear domination of one modal family over the other.

The corner displacements method of mode identification described by Soares et al. (2021b) is elegantly simple, easy to implement and quick to compute. While not tested rigorously in this work, this method presents some clear advantages over shape recognition, MAC and mapping to known modes. That said, only aluminum bars were considered by Soares et al. (2021b). This method should be tested before use with rosewood bars, as their orthotropic material properties and tendency for long, deep cutaways may create substantial coupling between lateral and torsional deformation.

All of the mode identification methods outlined above were applied to bar models with free-free boundary conditions. Changing these boundary conditions will have a large impact on mode shapes and frequencies. New methods of mode identification may likely be required for bar models with different boundary conditions.

4.3 Example Tuned Bars

The following subsections provide examples of bar geometry tuned over the course of this research using the methods developed in this chapter. Material and geometry definition types will be specified for each. The majority of examples covered have appeared previously through conference presentations (Beaton and Scavone, 2019b, 2021b) or a journal paper (Beaton and Scavone, 2021a).

4.3.1 Basic Bars

Examples contained in this subsection tune only vertical-flexural modes (V1, V2, etc.) to illustrate various points.

Untuned Modes in Lower Marimba Octaves

One of the earliest studies performed in this work examined the variation of untuned modal frequency ratios over the lower two octaves of a five-octave marimba (note C2 to note C4) (Beaton and Scavone, 2019b). Each bar was modelled using the cross-section interpolation method described in Section 3.2.2. Three rectangular cross-sections were defined and mirrored lengthwise to enforce cutaway symmetry. Using the depths of these three cross-sections as inputs, modes V1, V2 and V3 were tuned to their typical 1:4:10 ratios so as to produce results representative of conventional bar tuning. Outer bar dimensions were based on measurements of a five-octave Yamaha marimba. Figure 4.4 plots the resulting frequency ratios for the first nine modes of each bar from these simulations.

The results of Figure 4.4 confirm anecdotal reports from makers that the frequency of mode T2 can be close to that of mode V3 over a range from approximately note E3 to note B3. Figure 4.4 also shows a relatively uniform frequency ratio for mode T1, placing it consistently between modes V1 and V2, without being near either. Changes in bar width, which occur only at select points along the keyboard, are noticeable between notes E2 and F2, and notes F3 and F♯3. Interestingly, mode L1 does come close to mode V2 at note C♯2, though the author is not aware of any complaints of timbral issues caused in this case. This may be related to the mode shapes of L1 and V2. Coupling between torsional and lateral deformation does allow mode L1 to be excited when striking a bar's playing surface. However, areas of the playing surface that, if struck, would excite mode L1 do not line up well with areas that would excite mode V2. Thus the two modes are less likely to be excited at the same time. The opposite is true of modes T2 and V3, between which the areas causing maximal excitation coincide.

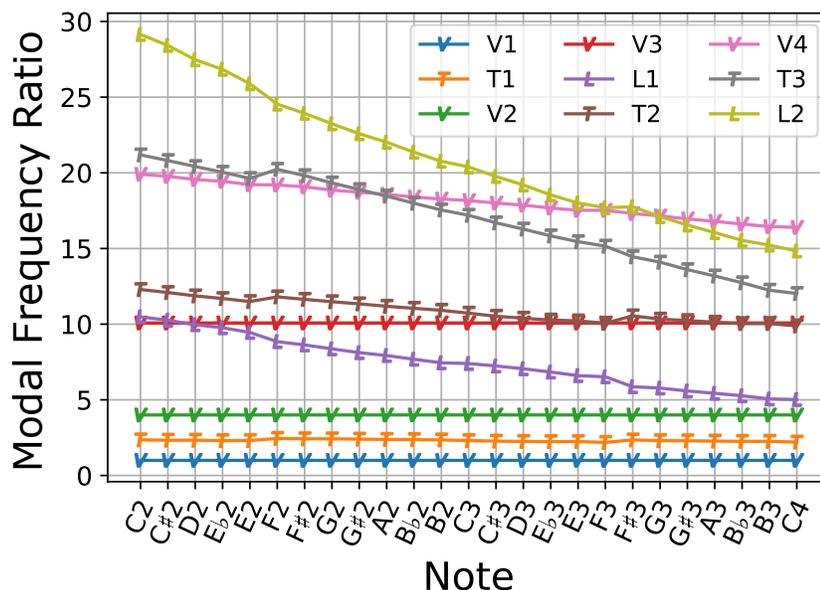


Figure 4.4: Modal frequency ratios from simulated bars over the lower two octaves of a five-octave marimba. Outer bar dimensions are based on measurements of a Yamaha instrument. This figure is adapted from [Beaton and Scavone \(2019b\)](#).

Impacts of Initial Geometry

Importance of initial geometry selection is discussed as part of the tuning methodology described in Section 4.2.2. Figures 4.5 to 4.7 demonstrate the impact of this choice. Each of the three figures show the initial and final geometry of a marimba bar tuned to the typical 1:4:10 ratios in flexural modes V1 to V3. All three bars have the same outer dimensions, material properties and number of input parameters. The only differences between the three are the choice of initial geometry. Figure 4.5(a) begins with a very thin cutaway in its initial geometry, resulting in a similarly thin cutaway for its tuned geometry in Figure 4.5(b). Figure 4.6(a) starts with a more moderately thick initial geometry and ends with a bar that is somewhat hollowed out in Figure 4.6(b), with thicker portions at various location and along the sides. In sharp contrast to these two, Figure 4.7(a) begins with a bar of uniform thickness, and outputs a bar with deep “hinge-like” segments that control the relevant modal frequencies in Figure 4.7(b). These hinged areas provide control over the vertical modes while keeping large portions of the cutaway quite thick, similar to the initial geometry. Looking at Figures 4.5 through 4.7 the importance of selecting a desirable initial geometry is clear.

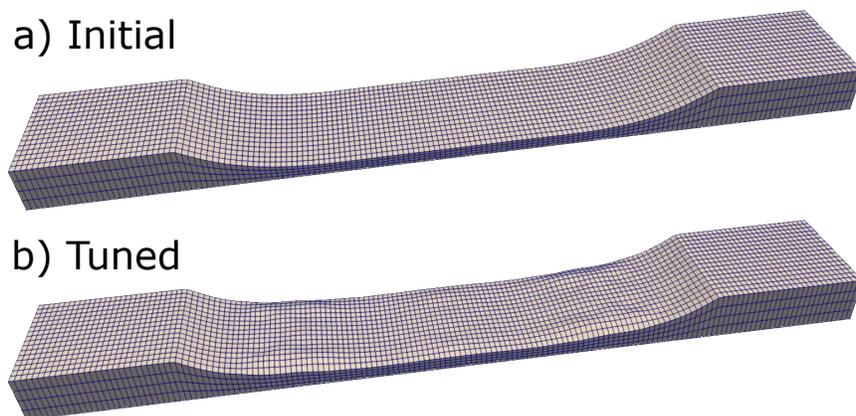


Figure 4.5: Example rosewood bar tuned to the typical 1:4:10 ratios using a **thin cutaway** as its initial input geometry. (a) Initial geometry. (b) Tuned geometry.



Figure 4.6: Example rosewood bar tuned to the typical 1:4:10 ratios using a **moderate cutaway** as its initial input geometry. (a) Initial geometry. (b) Tuned geometry.

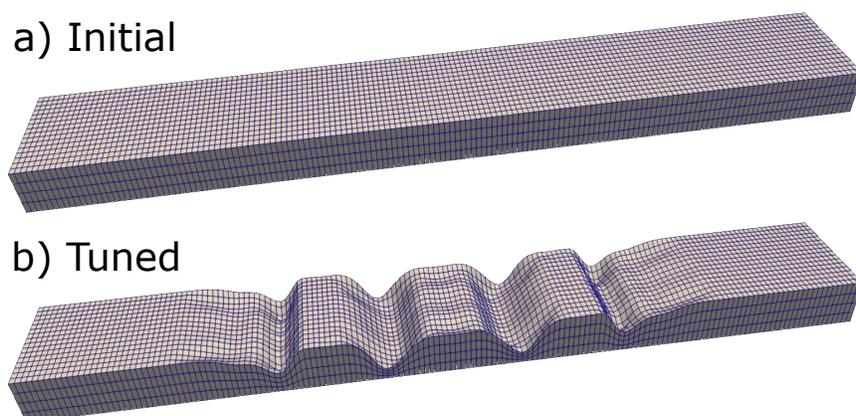


Figure 4.7: Example rosewood bar tuned to the typical 1:4:10 ratios using a **uniform bar** as its initial input geometry. (a) Initial geometry. (b) Tuned geometry.

Table 4.1: Target tuning ratios of all model 1 submodels in Figure 4.8. Table adapted from [Beaton and Scavone \(2021a\)](#).

Material	Note	Target Tuning Ratios					Model	
		V1	V2	V3	T1	T2	L1	No.
Rosewood	F3	1	-	-	-	-	-	1(a)
		1	4	-	-	-	-	1(b)
		1	4	10	-	-	-	1(c)
		1	4	10	2	-	-	1(d)
		1	4	10	2	8	-	1(e)
		1	4	10	2	8	6	1(f)

4.3.2 Rosewood Bars

This section provides example tuned rosewood bars, most of which were previously shown by [Beaton and Scavone \(2021a\)](#). These example bars use the Clough-Tocher interpolation to compute cutaway geometry, as discussed in Section 3.2.2.

Effects of Tuning Additional Modes

Figure 4.8 shows the effects of tuning progressively more modes in a rosewood example bar. Each of the submodels in Figure 4.8 used a thin, smooth initial geometry similar to that in Figure 4.5(a). Model 1(a), shown on the far left of Figure 4.8, tunes only mode V1, leaving the other reported modes untuned. The tuned geometry of model 1(a) looks very similar to the smooth, thin, initial geometry, indicating very little change was necessary to tune just a single mode. Model 1(b) tunes modes V1 and V2, while model 1(c) tunes V1, V2 and V3. In both cases, the final geometry is still quite smooth and rather similar to the initial geometry. Notably, the untuned frequency ratio of mode T2 moves from around 8 in model 1(a) to about 10 in model 1(c), placing it in close proximity to mode V3.

Model 1(d) in Figure 4.8 adds mode T1 to the tuned modes in the progression of submodels, and produces the first significant change in the tuned bar geometry. These changes become more pronounced and more interesting in model 1(e), which tunes mode T2 along with the others from previous submodels. Two “teeth” are noticeable on either side of the bar around its midpoint. These features become elongated in model 1(f), which tunes mode L1 along with the others. This set of submodels clearly demonstrates the versatility of tuning torsional and lateral-flexural modes using the surface interpolation described in Section 3.2.2.

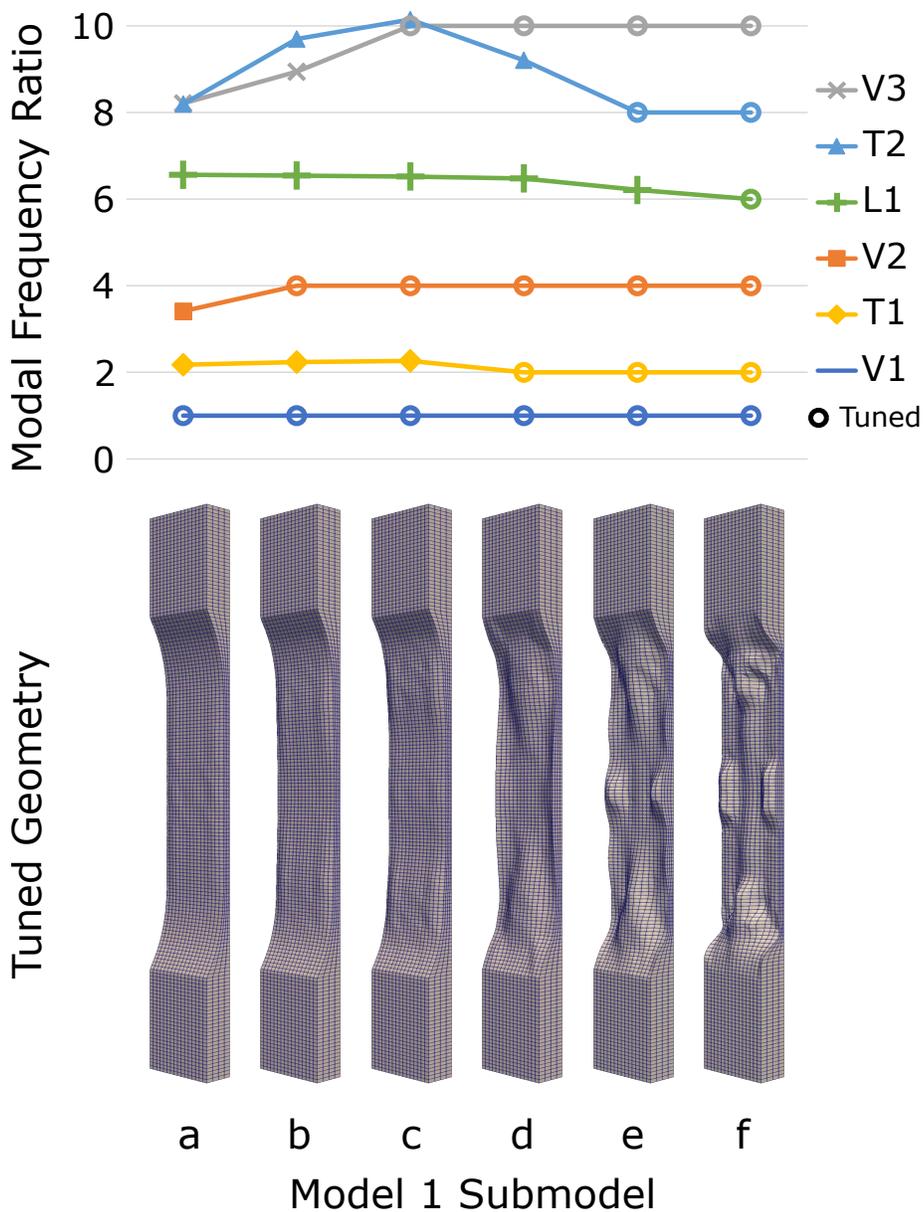


Figure 4.8: Tuned bar geometries and frequency ratios for models that tune progressively more modes. Model 1(a) begins by tuning only mode V1. Model 1(b) tunes modes V1 and V2. These additions continue to model 1(f), which tunes all six of the modes plotted above. Table 4.1 outlines the tuning targets of each submodel. Figure borrowed from [Beaton and Scavone \(2021a\)](#).

Table 4.2: Target tuning ratios of example rosewood bar models in Figure 4.9. Table adapted from [Beaton and Scavone \(2021a\)](#).

Material	Note	Target Tuning Ratios					Model No.	
		V1	V2	V3	T1	T2		L1
Rosewood	F3	1	4	10	2	8	-	1(e)
		1	4	10	-	11	-	2
		1	4	8	2	8	-	3
		1	5	9	3	7	-	4
		1	4	10	2	10	6	5

Example Tuned Bars

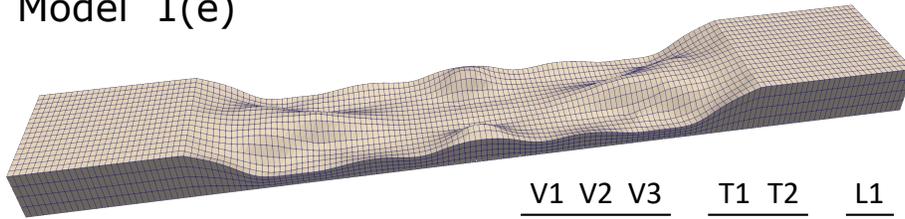
Figure 4.9 provides various examples of tuned rosewood bar geometries. All bars in the figure are tuned to the note F3. Each bar in Figure 4.9 has outer dimensions 406 mm × 58 mm × 24 mm based on measurements of an existing marimba. A minimum thickness of 6 mm is maintained throughout these models. Tuning ratios are shown in Figure 4.9 and summarized in Table 4.2.

The example models in Figure 4.9 tune anywhere from four to six modes, with cutaway geometries ranging from slightly sophisticated to decidedly impractical. Models 1(e) and 2 show interesting cutaway geometries of reasonable complexity. In some ways, the two models show opposing features in their cutaways. Model 1(e) tunes mode T2 down to a frequency ratio of 8, compared to its value of roughly 10 when left untuned in model 1(c) (see Figure 4.8). By contrast, Model 2 tunes mode T2 up to a ratio of 11. The opposing effects these models have on mode T2 are reflected in the material distribution over their cutaways. Near either end of model 1(e)’s cutaway, material is concentrated near the centreline of the bar, and thinned at the sides. Model 2 shows the opposite pattern, with material concentrated toward the sides of the bar and thinned near the centreline. These trends are reversed at the midpoints of the two bars. Material is distributed near the sides of the bar at the midpoint of model 1(e) (these are the “teeth-like” structures mentioned previously), whereas model 2 concentrates material along the centreline of the bar at its midpoint.

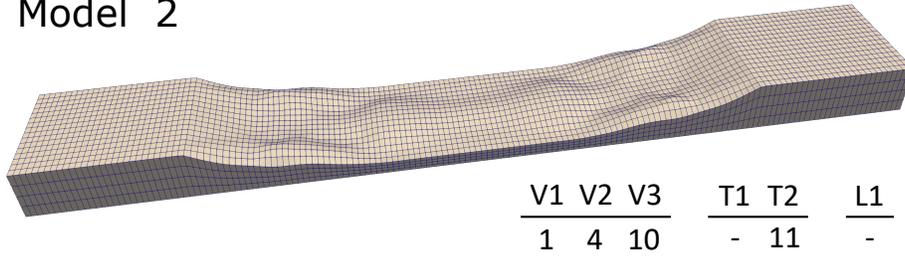
Model 3 has similar tuning ratios to model 1(e), with the exception that mode V3 is lowered from a ratio of 10 to a ratio of 8. This change is sufficient to cause significant changes in cutaway geometry. Most notable are the deep hinge-like regions at the ends of model 3’s cutaway. These hinges are positioned at a point that will have considerable impact on mode V3 and substantially less impact on modes V1 and V2. While models 1(e) and 2 have rather modestly complex cutaway geometries, model 3 is on the impractical side, with rapidly changing slopes and sharps edges.

Model 4 has perhaps the most complex cutaway in Figure 4.9. This very interesting and

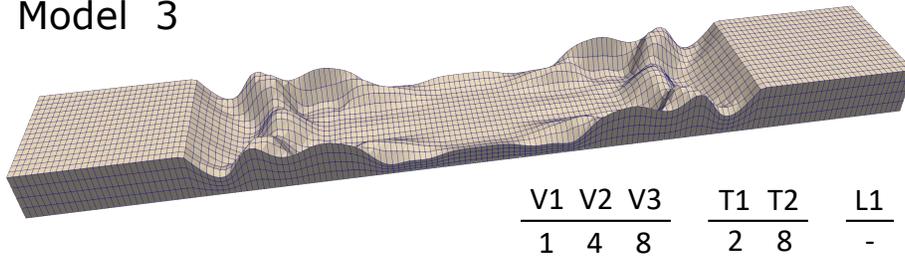
Model 1(e)



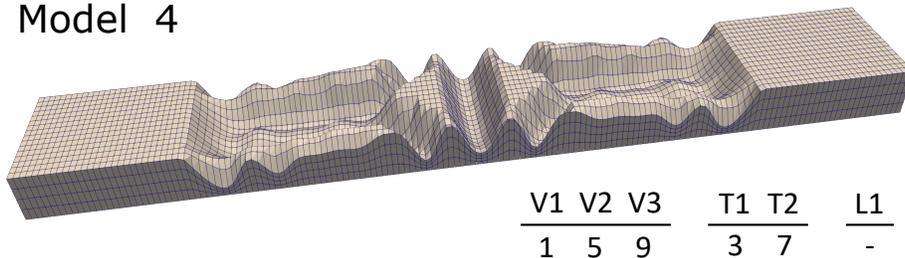
Model 2



Model 3



Model 4



Model 5

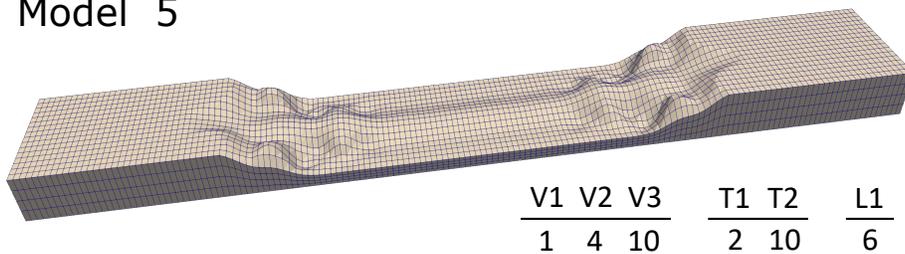


Figure 4.9: Tuned geometries for and frequency ratios for various rosewood bar models of the musical note F3. Tuning ratios are indicated for each bar with “-” denoting an untuned mode. Table 4.2 outlines the tuning targets of each example model. Figure borrowed from [Beaton and Scavone \(2021a\)](#).

Table 4.3: Target tuning ratios of example aluminum bar models in Figure 4.10. Table adapted from [Beaton and Scavone \(2021a\)](#).

Material	Note	Target Tuning Ratios						Model No.
		V1	V2	V3	T1	T2	L1	
Aluminum	C4	1	3	9	-	7	-	6
		1	4	8	3	10	-	7
		1	4	10	3	12	8	8
		1	5	11	-	-	-	9

entirely impractical geometry is a result of the atypical tuning ratios in the model, in which modes V2, V3, T1 and T2 are all tuned to odd harmonics. Notable in the cutaway are very thick areas along the edges, which serve to increase mode T1 to a frequency ratio of 3. A deep hinge at the bar’s midpoint primarily affects modes V1 and V3. Most interestingly, a set of slanted hinge-like areas on either side of the midpoint serve to lower mode T2 to a tuning ratio of 7. While interesting as theoretical results, models 3 and 4 would both be very difficult to fabricate, even on a CNC mill.

Model 5 is the lone example that includes mode L1 amongst its tuning targets. This mode is primarily controlled by the long stretch of cutaway in which material is thinned on the bar edges and thickened slightly along a sort of central “spine”. The resulting cutaway in this model is less complex than models 3 and 4, but does contain some intricate details at either end, which could prove challenging to fabricate.

In terms of practicality, models 3 and 5 have an additional caveat not present in the other example models. In each of models 3 and 5 modes V3 and T2 are tuned to the same frequency ratio. While again interesting as a theoretical result, in practice any small misalignment in the frequencies of these modes could produce an unintended beating effect, which is likely undesirable.

4.3.3 Aluminum Bars

Figure 4.10 provides various examples of tuned aluminum bar geometries. All bars in the figure are tuned to the note C4. Each bar in Figure 4.10 has outer dimensions 333 mm × 57 mm × 13 mm based on measurements of an existing vibraphone. A minimum thickness of 5 mm is maintained throughout these models. Tuning ratios are shown in Figure 4.10 and summarized in Table 4.3.

The example models in Figure 4.10 again tune anywhere from four to six modes, with varying levels of cutaway complexity. With an end thickness of 13 mm and minimum thickness of 5 mm the aluminum bar examples in Figure 4.10 provided less distributable material with which to shape their cutaways compared to the rosewood bar examples in Figure 4.9.

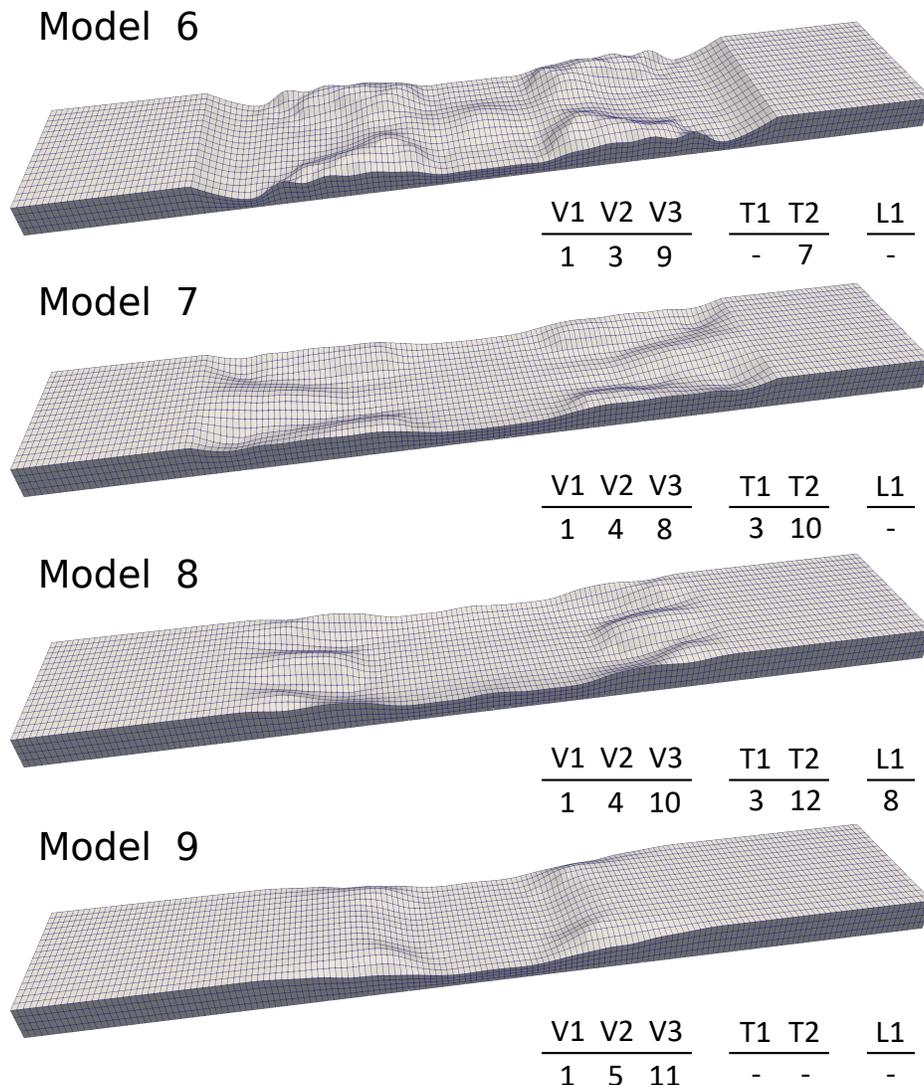


Figure 4.10: Tuned geometries for and frequency ratios for various aluminum bar models of the musical note C4. Tuning ratios are indicated for each bar with “-” denoting an untuned mode. Table 4.3 outlines the tuning targets of each example model. Figure borrowed from [Beaton and Scavone \(2021a\)](#).

Model 6 has perhaps the most interesting cutaway geometry in Figure 4.10. As is seen in other models, two hinge-like areas are present at the ends of the cutaway, where they will have greatest effect on mode V3. Also of note is the interesting X-shaped pattern created by thickened ridges on either side of the bar’s midpoint. These features presumably affect mode T2, as it is the only torsional mode tuned in the model.

Model 7 creates a more practical cutaway geometry than model 6. With only slight complexity at the ends of its cutaway, and a smooth, flat area around its midpoint, this example bar should be relatively easy to fabricate with a CNC mill.

Model 8 manages to tune six modes, including mode L1, while maintaining a relatively simply cutaway geometry. The ends of its cutaway even bear some similarity to model 7. Despite its unsophisticated geometry, some sharp edges may render this example slightly challenging to fabricate.

Tuning the typical V1, V2, V3 modes to the atypical ratios 1:5:11, model 9 produces the simplest geometry of any example in Figure 4.10. The straightforward shape with its gradual slope and rounded edges should be an excellent candidate for fabrication.

Overall, the relatively simplicity of cutaway shapes in Figure 4.10 compared to those in Figure 4.9 is a product of the aluminum bar examples have less material available to distribute over their cutaways. This results from the minimum 5 mm thickness of the aluminum examples representing 38% of their end thickness, leaving 62% of material distributable over the cutaway. By contrast, the 6 mm minimum thickness of the rosewood bar examples represents 25% of the end thickness of those bars, allowing 75% of material to be distributed over their cutaways.

4.4 Algorithm Performance

4.4.1 Tuning Convergence

The bar tuning method outlined in this chapter centres around the iterative application of Equation 4.4. Given an initial cutaway geometry, each application of Equation 4.4 seeks a similar geometry that produces the desired frequency ratios in specified modes. Figure 4.11 shows the iterations involved in tuning model 1(e) beginning from the initial geometry shown in Figure 4.5(a). A tolerance of one cent was enforced for each of the five modes tuned in Figure 4.11 and in all examples throughout this chapter.

It is apparent from Figure 4.11 that the frequencies of individual modes do not converge monotonically. Modes V1 and V2 clearly overshoot their target values before later converging. The little difference apparent between iterations 4 and 5 is a product of the rather stringent one cent tuning tolerance that was used. In iteration 4 only mode V2 is tuned within this tight tolerance.

The tuning example in Figure 4.11 shows relatively fast convergence compared to other models in this research work. This model benefits from a strategy of having many of the thickness inputs begin near one of their limits (in this example, many values in \mathbf{x}_0 are set equal to T_{min}). In such scenarios the individual values in each iterative input vector \mathbf{x}_i tend to move away from the boundary where they started. In these cases $\mathbf{D}(\mathbf{x})$ in Equation 4.6 becomes an identity matrix and the entirety of matrix $\mathbf{J}(\mathbf{x})$ is active in Equation 4.4, which

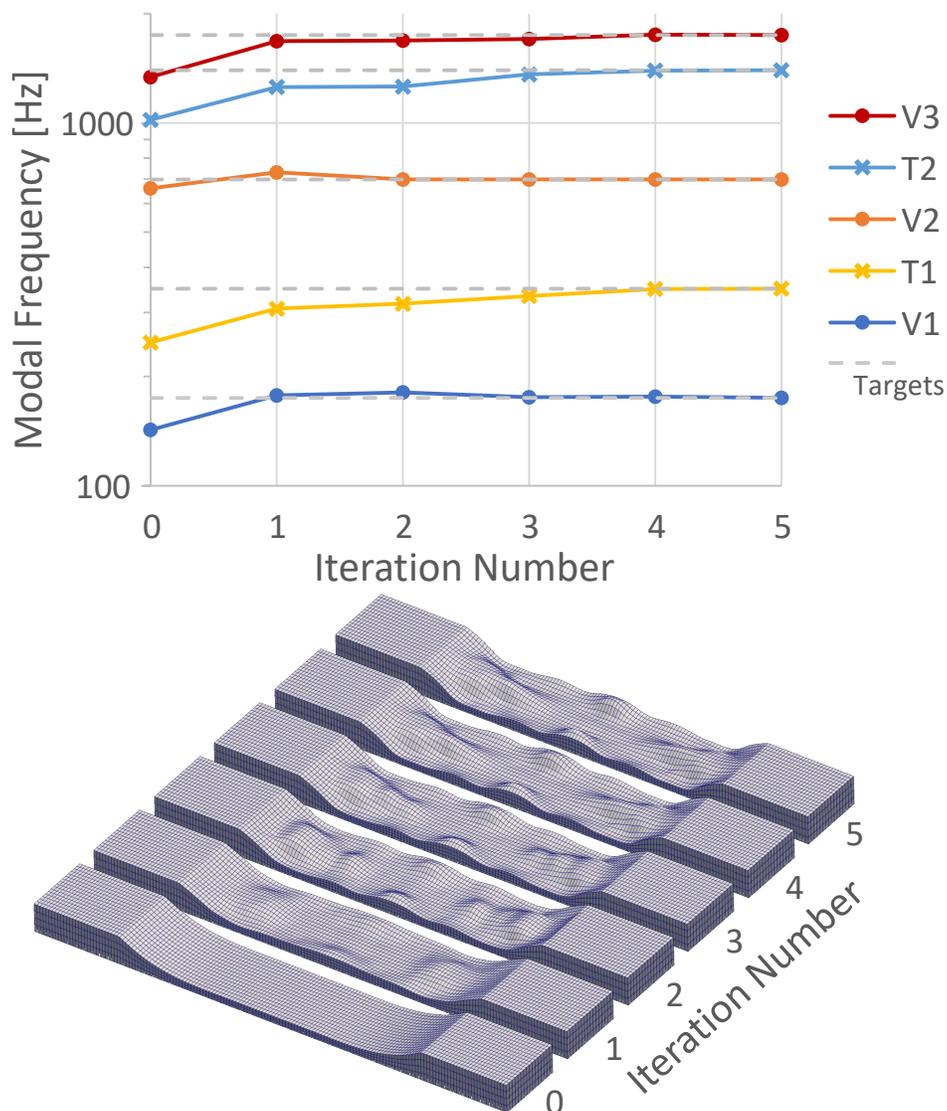


Figure 4.11: Cutaway geometries and resulting modal frequencies for the iterations required to tune model 1(e). Iteration 0 is the same initial geometry shown in Figure 4.5(a).

can lead to faster convergence.

Convergence of the example in Figure 4.11 is also aided by the fact that modal frequencies in iteration 0 are not too distant from their target values. The modes stay well separated throughout the iterations, implying that at no point were degenerate modes present.

The absence of any of these conditions can slow down convergence. Figure 4.12 shows algorithm convergence for convergence run 3 of Prototype 2 in Section 4.5, if that model were run from its initial geometry (in practice, convergence runs begin with the final input vector from the previous convergence model run).

The example in Figure 4.12 required slightly longer to converge than the example shown

in Figure 4.11. The two examples begin with very different initial geometries. The initial geometry in Figure 4.12 was designed to eliminate sharp corners along the cutaway edges. This was achieved by creating an initial geometry that is thicker near the bar’s centreline and thinned toward the edges. While successful in creating a tuned geometry free from sharp edges, modal frequencies from this initial geometry were farther from their target values than those in Figure 4.11, thus necessitating some additional iterations. Again the effects of a stringent tuning tolerance are seen in iterations 4 through 7 of Figure 4.12, where small changes in bar geometry and modal frequencies bring the modes within one cent of their target frequencies.

Figure 4.12 also gives an example of mode reordering, with modes T1 and V2 reversing order between iterations 1 and 2.

4.4.2 Non-Convergence

Several scenarios also exist that could result in the algorithm not converging. If the analyst requests a set of tuning ratios that are infeasible for the input geometries and constraints, the model will not converge and will instead stop after a maximum number of iterations. In other, relatively rare cases, the algorithm can enter an “orbit” while trying to tune a feasible set of tuning ratios. In this scenario, the estimated solution vectors, \mathbf{x}_i , can enter an orbit, wherein they alternate between some fixed number of two or more repeated similar vectors. When this occurs, the algorithm is said to be orbiting about some fixed point that may be a solution, or may be a local minimum. Such orbits are a known caveat of gradient-based solving methods. Over hundreds of model runs in this research, orbits were observed to occur a handful of times.

Troubleshooting

In the event of non-convergence, numerous troubleshooting techniques can be explored before concluding that a desired set of tuning ratios may be infeasible. Troubleshooting approaches employed in this work include:

1. Rerunning the tuning algorithm with a different initial geometry.
2. Enforcing a maximum allowable change in input vector for each iteration of Equation 4.4.
3. Running the tuning algorithm using one less mode tuned. For a given set of modal tuning objectives to be feasible, any subset of those objectives must also be feasible (and should be quicker to investigate).

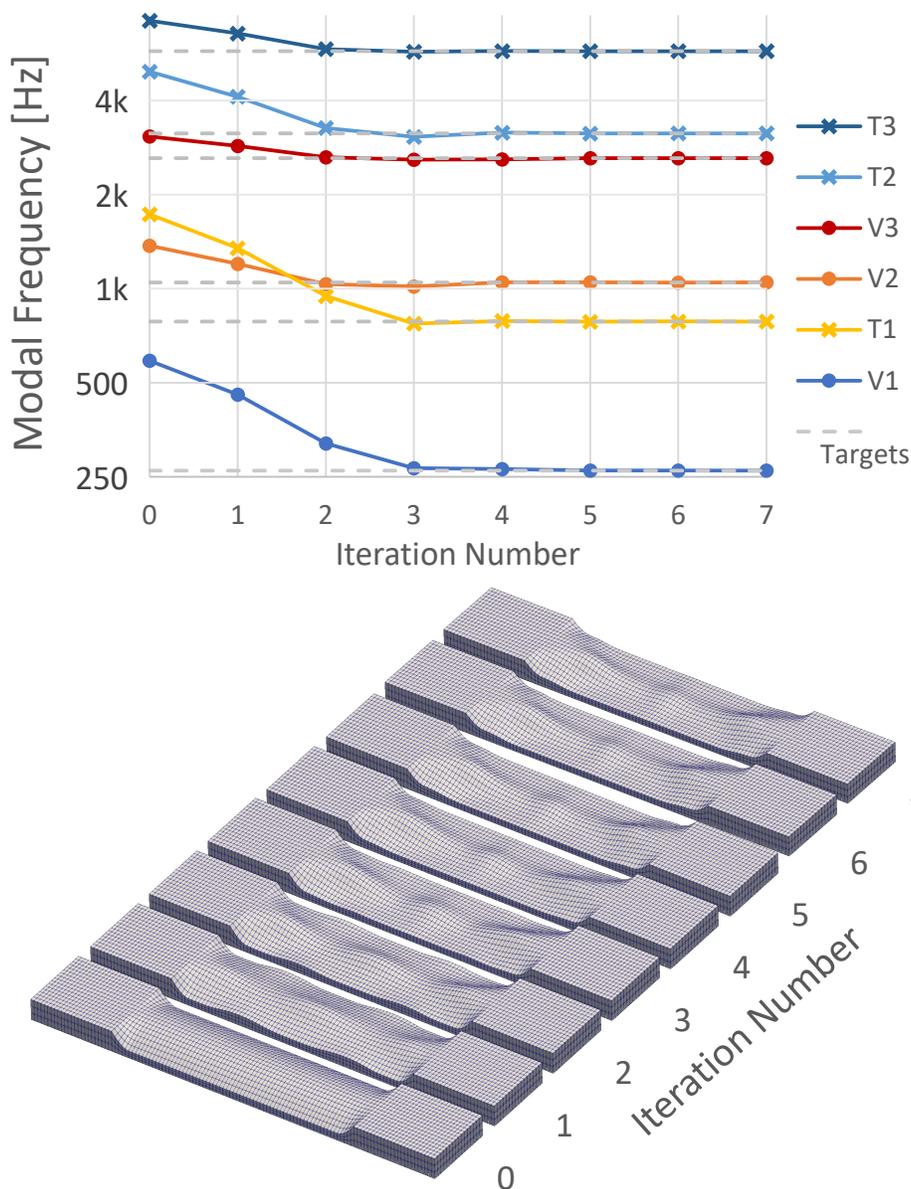


Figure 4.12: Cutaway geometries and resulting modal frequencies for the iterations required to tune the model for fabrication bar Prototype 2. The initial geometry in Iteration 0 was selected to eliminate sharp edges along the sides of the cutaway.

The methods above are by no means an exhaustive list of troubleshooting techniques for gradient-based tuning methods. For example, one interesting potential approach could be to create an algorithm to detect orbiting. In the event that orbiting is detected, the tuning algorithm could be instructed to set the next input vector to the average of those input vectors present in the orbit.

Table 4.4: Prototype bar dimensions.

Bar	Length [mm]	Width [mm]	End Thickness [mm]
Prototype 1	400	63.5	19
Prototype 2	365	63.5	18

Table 4.5: Prototype bar modal tuning ratios. Untuned modes are indicated by ‘-’.

Bar	Tuning Ratio					
	V1	V2	V3	T1	T2	T3
Prototype 1	1	4	10	4	14	-
Prototype 2	1	4	10	3	12	22

4.5 Prototype Fabrication

Two aluminum prototype bars were fabricated in this work. The isotropic, manufactured nature of aluminum makes it an easier starting point than a natural, orthotropic material like wood. Compared to rosewood, aluminum was also more readily available, and familiar to the fabrication shop technicians. The number of prototype bars created was ultimately governed by cost. The bars were produced on a computer numerical control (CNC) mill at McGill’s Faculty of Engineering Workshop. Each prototype had a price tag of around \$400, with the bulk of this cost paying for shop technician services (only shop technicians are permitted to operate the CNC mill).

The second prototype bar was produced several months after the first, taking advantage of lessons learned from that initial bar fabrication. This resulted in vastly improved performance of the second bar compared to the first. The following sections describe the modelling, fabrication and performance of these bars, with an emphasis on the second, more successful, prototype. Discussion of prototype bar fabrication can also be found in an article ([Beaton and Scavone, 2022](#)) published as this thesis was under review.

4.5.1 Bar Descriptions

Each prototype was tuned to the musical note C4. Bar outer dimensions are given in Table 4.4 while Table 4.5 outlines the tuning ratios used for each. Stock aluminum material of the popular 6061 alloy ([Battelle Memorial Institute, 2020](#)) was selected based on price, availability and ease of machining. Stock material was supplied in 2½-inch by ¾-inch (63.5 mm by 19 mm) bars.

Prototype 1

In selecting dimensions for the first prototype bar, an effort was made to reduce the amount of required CNC milling, thereby limiting the bar’s cost. The overall width and end thickness of the bar were set to match those of the stock material. With a 19 mm thickness, the bar would be significantly thicker than the 13 mm thick bars measured on a Yamaha vibraphone. For this reason, no attempt was made to match the bar’s other dimensions to those of a typical vibraphone bar. The bar’s length was set at 400 mm to allow either end of the stock material to be cut square to the sides.

With the bar’s outer dimensions set, dozens of analyses were performed. These analyses explored different combinations of tuning targets (musical note and modal tuning ratios) and initial input geometries. The musical notes considered were: F3, G3, A3, B3 and C4. Ultimately, note C4 was selected along with the tuning ratios shown in Table 4.5. This combination was chosen as its resulting bar geometry was deemed most preferable (in terms of being practical for fabrication and aesthetically interesting).

Figure 4.13 shows the geometry of the first prototype bar, with panel (a) giving a screenshot of the modelled geometry and panel (b) a photo of the fabricated bar. Initial bar geometry was selected so as to avoid any sharp edges along the sides of the cutaway. This is most obvious near the ends of the cutaway, where material is concentrated along the bar’s central axis while sloping down toward its sides.

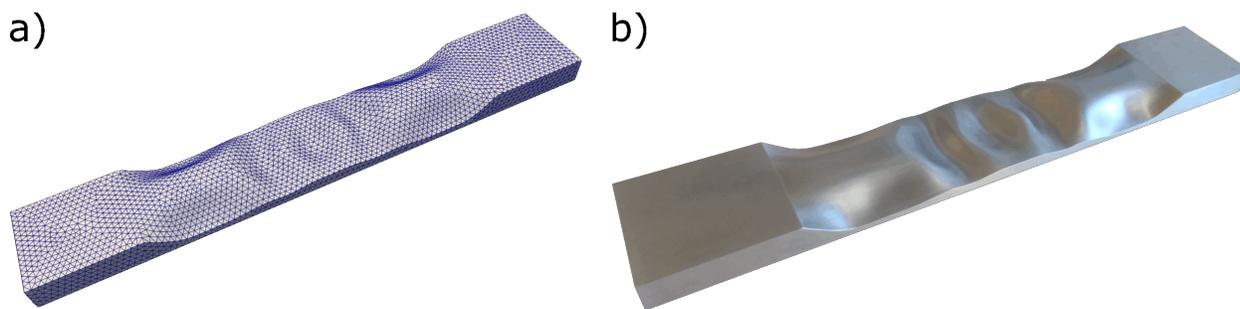


Figure 4.13: Images of Prototype 1: (a) bar model meshed with tetrahedral elements, (b) photo of fabricated bar. Figure borrowed from [Beaton and Scavone \(2022\)](#).

Though its shape is aesthetically interesting, in terms of vibrational properties Prototype 1 performed quite poorly. The modal frequencies of Prototype 1, measured by experimental modal analysis (see Appendix A), departed significantly from those predicted by the computer models. The primary source of these discrepancies were significant differences between the modelled bar geometry (i.e. its intended geometry) and the actual fabricated geometry.

The stock aluminum bar material purchased for prototyping was not a stress-relieved variety. Instead the aluminum stock pieces used for bar fabrication contained residual stresses

imparted by their manufacture. In the stock bar's original state, internal forces from these residual stresses are in equilibrium, and the stock material remains quite straight. Removal of bar material during CNC milling disturbs the equilibrium of these internal forces, causing the bar to warp. The milling procedure itself can also impart residual stresses (Beaton and Scavone, 2022), even on stress-relieved materials, making some degree of warping unavoidable.

The warped bar no longer has a flat surface upon which it can be placed in the CNC mill. To compensate, shop technicians must remove additional material to flatten the warped bar before additional rounds of milling. The author was not aware of this necessity prior to fabrication. As shown in Table 4.4, the designed end thickness of Prototype 1 was equal to the stock bar's thickness of 19 mm. With no excess thickness to spare, the removal of material to flatten the bar after warping resulted in Prototype 1 being anywhere from 0.5 mm to 0.8 mm thinner than expected at any given point.

This unanticipated thinning of the bar resulted in a fundamental frequency roughly 10% lower than intended. Results from Prototype 1 are omitted from the remainder of this section, as the observed discrepancies in modal frequencies were clearly a result of fabrication issues and not a product of the tuning procedure described in this chapter.

Prototype 2

After assessing performance of Prototype 1, and diagnosing the source of geometry deviations, a second prototype was designed. The goal in designing Prototype 2 was to minimize any errors in final bar geometry caused by the fabrication process. To achieve this goal, Prototype 2 was designed in close collaboration with the fabrication shop technicians, so as to accurately produce the intended geometry in the final product.

The most important adjustment made to the second prototype was to reduce the overall thickness of the final design. As shown in Table 4.4, Prototype 2 had an end thickness of 18 mm rather than the 19 mm used in Prototype 1. Given that the stock material was 19 mm thick, this reduction in bar thickness left 1 mm of buffer material that could be removed to compensate for warping caused by the milling process. Bar width was kept at 63.5 mm, matching that of the stock material, as advised by the fabrication technicians. Musical note C4 was maintained, matching Prototype 1.

With the bar note, width and end thickness selected, over 150 bar models were tuned, testing different combinations of bar length, tuning ratios, initial geometries, cutaway lengths and minimum thickness limits. The aim was to produce a bar with the following characteristics:

1. Smooth geometry (no sharps edges).

2. Six tuned modes (one more than Prototype 1).
3. Relatively small amount of material removed for cutaway (reduces warping).
4. Interesting cutaway shape.

Items 1 and 2 were mandatory. Item 3 was considered preferable while Item 4 was optional. Figure 4.14 shows the geometry selected for Prototype 2, along with a photo of the fabricated bar.

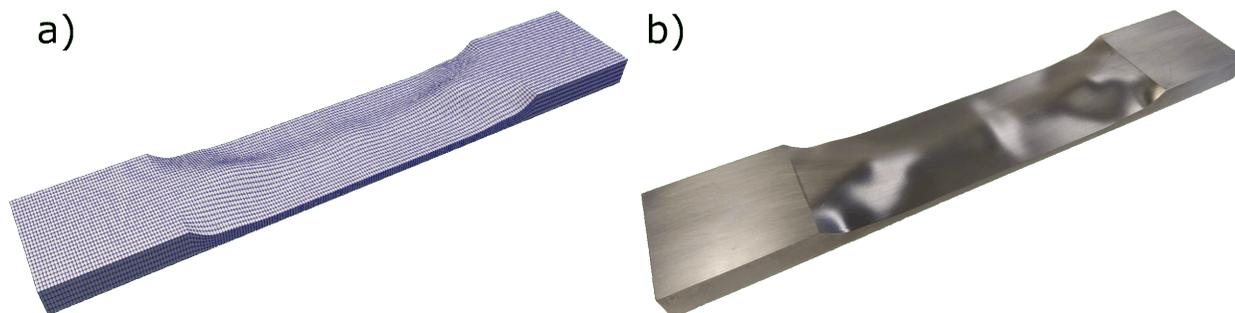


Figure 4.14: Images of Prototype 2: (a) bar model meshed with 20-node hexahedral elements, (b) photo of fabricated bar. Figure borrowed from [Beaton and Scavone \(2022\)](#).

The fabricated geometry of Prototype 2 was far more accurate than Prototype 1, leading to greatly improved bar performance. The same modelling and tuning procedures were employed for both prototypes. Thus the improved performance of Prototype 2 confirms that poor performance from Prototype 1 was a product of fabrication issues. To provide commentary indicative of the performance possible using tuning methods described in this chapter, the following sections consider only Prototype 2.

4.5.2 Model Convergence Test

Prior to fabrication, each of the tuned bar models selected for prototyping was tested for convergence. Model runs used for convergence tests were comprised of 20-node hexahedral elements. Table 4.6 shows the number of elements used in each convergence model run. Mesh densities were progressively increased until the control point positions, which govern the NURBS surface defining the cutaway, did not change between runs. Thus the control point positions from run 5 of Table 4.6, are the same as those from run 6. These control point positions produced the modal tuning ratios shown in Table 4.5, with each modal frequency tuned to within one cent of its target value. The final model, with a total of 31,200 elements, is shown in Figure 4.14(a).

Table 4.6: Numbers of elements used for Prototype 2 convergence test models. Model X-, Y- and Z-axes correspond to the bar’s length, width and thickness directions, respectively.

Run Number	Elements Per Axis			Total
	X	Y	Z	
1	40	8	2	640
2	80	12	2	1,920
3	120	16	3	5,760
4	160	20	4	12,800
5	200	24	5	24,000
6	240	26	5	31,200

Table 4.7: Measured and modelled modal frequencies for the tuned modes of Prototype 2. Modelled results were produced from Model 6 in Table 4.6 using both Pylaster and Calculix. In all cases modelled frequencies in the two programs agreed to the first decimal place. Measured frequencies were determined using the LSCE method (see Appendix A). Frequency differences are expressed in cents, relative to the modelled values. Measured frequency ratios are calculated relative to the measured frequency of mode V1. Table borrowed from [Beaton and Scavone \(2022\)](#).

Mode	Frequencies			Frequency Ratios		
	Modelled [Hz]	Measured [Hz]	Difference [cents]	Target	Measured	Difference [cents]
V1	261.7	261.1	-4.1	1	1	-
T1	785.2	772.1	-29.1	3	2.96	-24.9
V2	1046.8	1047.2	0.7	4	4.01	4.6
V3	2616.3	2632.5	10.6	10	10.08	14.2
T2	3139.5	3108.7	-17.4	12	11.91	-13.6
T3	5755.8	5757.1	0.5	22	22.05	3.9

4.5.3 Experimental Modal Analysis and Bar Performance

Assessment of the fabricated bars’ performance was carried out via experimental modal analysis. The bar under test was supported on two small foam blocks positioned below the nodes of its fundamental mode. This support condition was designed to emulate how support cables are located on a vibraphone or marimba. Experimental modal analysis was then performed according to the same procedure described in Appendix A for rosewood bars.

Table 4.7 outlines the modal frequencies measured for Prototype 2 compared with those predicted by finite element models used for bar tuning. The average absolute error for the modal frequencies shown in Table 4.7 works out to 10.4 cents. Average absolute error for the frequency ratios is 12.2 cents. Measured frequency ratios in Table 4.7 are calculated relative to the *measured* frequency of mode V1; thus the reported differences in frequency ratios depart from the reported differences in frequencies.

Results shown in Table 4.7 are most readily compared with those of Soares et al. (2021b). Their work, published several months after Beaton and Scavone (2021a), though likely performed concurrently, also tuned idiophone bar torsional modes. The two works take decidedly different approaches to parameterizing and shaping bar cutaways. As described earlier in the chapter, this work tuned bar cutaways in a three-dimensional manner, allowing bar shapes to vary in all three dimensions. By contrast, Soares et al. (2021b) used what a previous paper by the same authors (Soares et al., 2021a) called “simplified undercuts”. Their approach formed bar cutaways that varied in a step-wise manner, with geometry remaining consistent across the width of the bar.

Soares et al. (2021b) included measurement results for three fabricated bars with different tuning ratios. Each of their bars tuned four vertical bending modes and two torsional modes. Looking at Table 4.7 the average absolute tuning ratio error of 12.2 cents is comparable to the results of Table I from Soares et al. (2021b), which report average absolute errors of 9.0, 10.9 and 32.9 cents for the three bars fabricated in that work. The -4.1 cent error in tuning Prototype 2’s fundamental frequency is notably better than fundamental frequency errors of -26, -12 and -38 cents for the three bars of Soares et al. (2021b).

Overall the measured frequency results in Table 4.7 indicate very satisfactory performance from Prototype 2, thereby demonstrating the potential for the three-dimensional bar tuning methodology developed in this chapter.

Chapter 5

Hybrid Modelling of Idiophone Bar Sound Radiation

5.1 Objective

As outlined in Section 1.2.2, this work seeks to enable the evaluation of potential new bar geometries and materials in a virtual manner, minimizing the need for physical prototypes. Evaluations are carried out using finite element models (FEM), as described in Chapter 3. An important aspect of this objective is that FE models of new bar designs should function without the need for **any** physical prototype to be built first. The only necessary inputs must be bar geometry and material properties. The required material properties are:

- Elastic modulus/moduli
- Poisson's ratio(s)
- Shear modulus/moduli
- Density
- Damping ratio (as a function of frequency)

These properties can be measured from stock material or taken from any available published sources. For sound radiation models that include the impact of a mallet, material properties must be established for both the bar and mallet. The mallet material is an important consideration, as it will act as a low-pass filter on the radiated sound, with harder mallets exciting more higher partials.

Time domain models of instrument excitation and sound radiation provide the benefit of great flexibility, but suffer from high computational costs. For percussion instruments, time

domain models can capture interaction between the mallet head and instrument. Doing so can account for the mallet head’s mass and the effects it has during impact with the bar. Including this interaction thus leads to more accurate simulation of the impact excitation and the radiated attack transient. [Chaigne and Doutaut \(1997\)](#) included mallet/bar interaction in their simulations of xylophone bar vibration. [Henrique and Antunes \(2003\)](#) later employed this same approach in their work. Both sets of authors considered one-dimensional bar models.

Capturing torsional mode behaviour requires three-dimensional bar models with total degrees-of-freedom orders of magnitude larger than models from these previous works ([Chaigne and Doutaut, 1997](#); [Henrique and Antunes, 2003](#)). Even with today’s computers the increased computational cost of tracking so many nodes can be prohibitive. For this reason, time history FEM is rarely employed in dynamic structural models with large numbers of nodes.

The objective of this chapter is thus to create a model that can simulate sound radiation from an impacted bar, using only geometry and material property inputs, and accounting for bar/mallet interaction, while reducing computation time in the FE model.

5.2 Concept

To accomplish the objective outlined above, this chapter proposes a *Hybrid Sound Radiation Model*. The concept is centred around a hybridization of finite element dynamic time history analysis, and modal models with decaying sinusoids (also known as *Prony models*). The hybrid model produced by this concept will output a waveform that tracks sound pressure at a given position relative to the instrument. This position, termed the “listening position” is equivalent to placing a virtual microphone at the corresponding location inside the model, and outputting the signal it records.

The output waveform will consist of three segments:

- (A) *Time history excitation*: A first segment produced by modelling sound radiation using the results of a **finite element dynamic time history model**. This time history model tracks the bar’s motion in response to a given excitation (mallet strike) and includes any nonlinear effects, such as interaction between the mallet head and bar. Any appropriate method may be used to calculate sound radiation at the listener position based on the simulated motion of the bar.
- (B) *Crossover*: A second segment that serves to transition between the first and third segments.

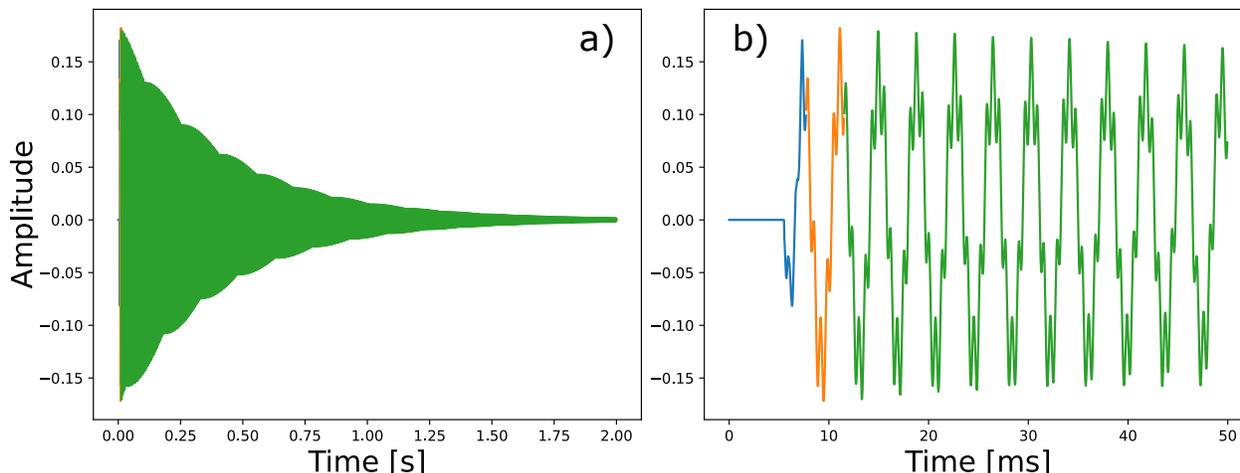


Figure 5.1: Example output waveform from the hybrid sound radiation model. (a) Full result for a model with a simulation time of 2 seconds. (b) Detailed view of the first 50 milliseconds. The three distinct segments of the model are colour-coded with: **blue** - time history excitation, **orange** - crossover, and **green** - Prony tail.

(C) *Prony tail*: A third segment produced by simulating sound pressure at the listener position as a **Prony model**, with a set number of decaying sinusoids. Parameters for the modes of this model will be determined using output from the first segment, and also prior knowledge of the bar behaviour from the damping model input to the time history model, and a separate finite element modal analysis.

Of the three segments outlined above, segments A and B are by far the most computationally expensive, as both include results from the time history analysis. Segment C, by contrast, is much less computationally demanding, as once parameters of the Prony model are determined, its output signal can be calculated very rapidly. The primary benefit on the proposed hybrid model is to minimize the required length of segments A and B in favour of the much less computationally demanding segment C. Figure 5.1 gives an example output waveform produced using this hybrid model approach. Segments A, B and C in Figure 5.1 are coloured blue, orange and green, respectively. It is immediately apparent how well the method is able to limit the extent of segments A and B in favour of the computationally advantageous segment C.

If the output signal in Figure 5.1 were generated using only finite element time history analysis (i.e. if the signal was comprised entirely of segment A), the required run time on the desktop computer used in this work would be measured in days. By contrast, the hybrid model output of Figure 5.1 was generated in less than an hour, the bulk of which time was spent producing segments A and B.

All of this begs the question “why not use a Prony model for the entire response?” The answer here is twofold: first, the time history model is capable of capturing nonlinear phenomena, such as interaction between the mallet head and bar surface, whereas Prony models rely on the principle of modal superposition, which itself assumes system linearity. Second, the parameters of a Prony model are determined by fitting the generated output to some input data. In this method, the input data comes from the time history analysis, along with a separate modal analysis and input damping model. Absent these initial simulation results to tune the Prony model parameters, it is common to use recorded data as input for parameter fitting (see, for example, [Aramaki et al. \(2007\)](#)). While using recorded data as input is a perfectly viable approach, the goal of this work is to enable evaluating new materials and geometries without the need for physical prototypes (which would be required to produce recorded inputs). By using only material property and geometry inputs, the proposed hybrid model approach accomplishes this goal. It does so while leveraging the Prony model’s efficiency to greatly reduce computational cost compared to a full time history model.

The following sections outline the steps involved in the hybrid sound radiation model, and provide example results and discussion.

5.3 Method

The objective for this chapter is accomplished by employing a hybrid approach to modelling bar impact and sound radiation in the time domain. This approach combines a nonlinear dynamic time history FE model (in which the equations of motion are solved using numerical time integration) with a modal model comprised of exponentially decaying sinusoids. Sound radiation is simulated by modelling the bar’s surface as a collection of point sources based on its finite element mesh ([Bestle et al., 2017](#)). Details of the method and example results are outlined in the following sections.

5.3.1 Overview

The hybrid radiation model is constructed via the following steps:

1. Perform a modal analysis of the bar FE model including the boundary conditions that will be used to model mallet impact and sound radiation.
2. Run a time history analysis of the mallet impact, including the complete time in which the bar and mallet make contact and a short duration thereafter.

3. Simulate sound radiation to a given listener position by treating the bar’s playing surface as a collection of point sources and inputting surface motion output from Step 2.
4. Select the modes to use in the modal model. Look up frequencies for each mode from the results of Step 1. Calculate damping ratios for each mode using the damping model selected for Step 2.
5. Adjust the modal frequencies and damping ratios from Step 4 to account for the effects of numerical time integration in the time history model.
6. Using the adjusted modal frequencies and damping ratios, fit the amplitude and phase of each mode to a portion of the radiated sound signal using the second step of Prony’s Method.
7. With all modal parameters fit, evaluate the decaying sinusoid for each mode over the desired range of time. Sum these decaying sinusoids to create the modal model.
8. Combine the initial sound radiation from Step 3 with the modal model from Step 7 by crossfading between the two over a specified overlap.
9. Perform a frequency warping on the radiated sound signal from Step 8 to convert the adjusted frequencies from Step 5 back to those expected in the modal analysis results from Step 1.

These steps are explained in further detail in the following sections.

5.3.2 Time Domain Modelling

Impact of the mallet head on the idiophone bar, interaction between the bar and mallet, and the subsequent bar motion are all modelled in the time domain using the FE model described in Chapter 3. All models including contact mechanics were run in Code_Aster, while some preliminary models were run in Pylaster and Calculix. Numerous boundary conditions were considered, though it proved difficult to create boundaries that could adequately restrain the model while maintaining modal frequencies from bar models with free-free boundaries. Instead, modal frequencies were determined by a modal analysis using the same boundary conditions applied for the impact model. Support locations were generally positioned near areas where the fundamental bar mode has minimal displacement (the “nodes” of the fundamental mode) so as to have minimal effect on the bar’s fundamental frequency. However, restraining these nodes would invariably alter frequencies of the other modes. To produce a bar with some harmonically tuned modes for the hybrid radiation model, the bar tuning

Table 5.1: Mallet impact force curve parameters.

Parameter	Bork Curve	Bestle Curve
A	44.77	76.62
μ	0.5266	0.5899
σ	0.1200	0.1150
t_0	-0.8715	-1.148

methods described in Chapter 4 were applied to a bar with the boundary conditions used in the time domain mallet impact model.

Mallet Impact as a Time-Varying Nodal Load

Some preliminary models defining mallet impact as a time-varying nodal load were run in Pylaster and Calculix. These models ran successfully, though they did not include interaction between the bar and mallet, and thus the effect of the mallet’s mass was lost. Applied forces were spread equally over a number of nodes within a small area on the bar’s playing surface. The total applied force was varied over time according to one of two excitation functions. These functions, shown in Figure 5.2 are based on curve fits from data in Figure 2 from [Bork \(1990\)](#), and Figure 4.14 from [Bestle \(2017\)](#). A modified log-normal function, used to fit each curve, takes the form:

$$f_m(t) = \frac{A}{(t - t_0)\sigma\sqrt{2\pi}} \exp\left(-\frac{\log(t - t_0) - \mu}{2\sigma^2}\right); \quad t > t_0, \quad (5.1)$$

where t is time, expressed in milliseconds, and $f_m(t)$ gives the mallet impact force in Newtons. The remaining terms have no physical interpretation and inconsistent units, as Equation 5.1 is here used empirically to fit the mallet impact force curves from the two cited papers. The log-normal distribution was taken as a basis for this function given its asymmetric pulse-like shape. Equation 5.1 is not applied as a typical log-normal distribution, as it is not modelling a random variable and has been modified by adding the term t_0 . Parameters of the fitted functions from Figure 5.2 are given in Table 5.1. Note that the fitted log-normal force curves in Figure 5.2 were used only to test explicit and implicit time integration, as described in the following section. Example models in Section 5.5 used contact mechanics to simulate mallet impact. Figure 5.2 also shows an example mallet impact force from these models. Further discussion is provided in the impact modelling subsection below.

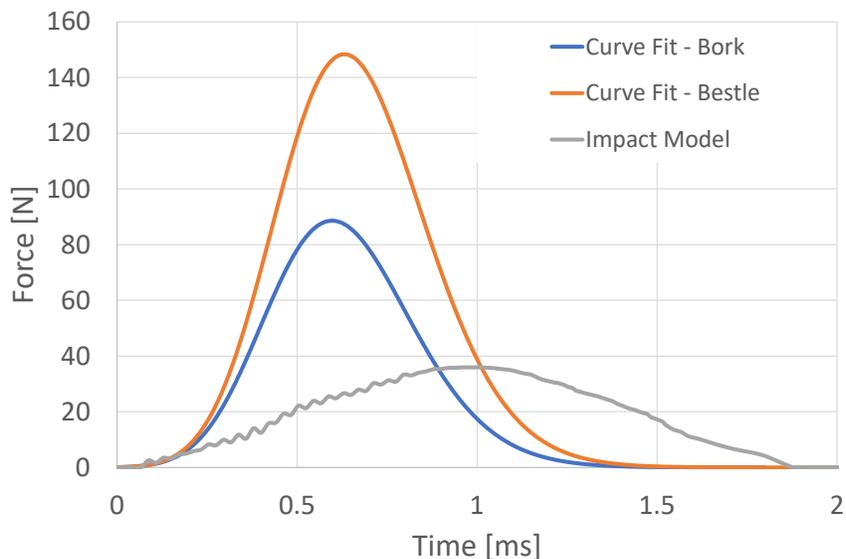


Figure 5.2: Mallet impulse force curves. Two curve fits based on data from [Bork \(1990\)](#) and [Bestle \(2017\)](#) are compared with the mallet impact force simulated using contact mechanics. The three curves have been offset along the time axis to roughly align their initial rise in force.

Explicit vs. Implicit Time Integration

Models that applied the mallet impact as a time-varying nodal load were used to test running the analysis with implicit or explicit time integration schemes. It was found that both integration schemes produced similar results in terms of nodal displacements. However, the time steps required for stability of the explicit integration scheme were so small, and the resulting number of required time steps so large, that implicit integration proved to be a much faster approach. Later it was found that frequency warping corrections, as outlined in Section 5.3.8, were more accurate for implicit time integration, further reinforcing its selection for this work.

Impact Modelling

To provide greatest flexibility in modelling mallet impacts, FE models were created including a model of the mallet head, along with its initial velocity. These models capture the impact between the mallet head and idiophone bar via finite element contact modelling. Initial models were run in Calculix, taking advantage of its capacity for contact detection. For analyses such as these, Calculix has the ability to dynamically adjust the time step size based on the current state of the model. Doing so can save considerable time for those

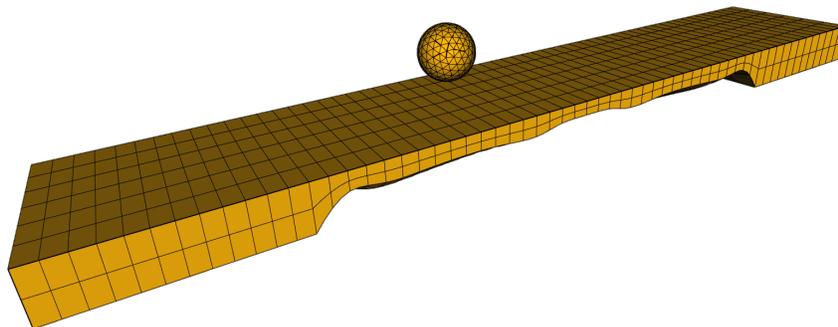


Figure 5.3: Initial geometry of a mallet impact model for use in Code_Aster. The bar is meshed with 720 20-node hexahedral elements. The mallet head is meshed with 1106 10-node tetrahedral elements.

portions of the analysis where no contact occurs. However, it was found that using dynamic time step adjustments in Calculix made it difficult to produce results at times corresponding to the desired audio sampling rate. The dynamically adjusted time steps varied considerably and resulting surface displacements would need to be interpolated between points with non-uniform distribution in time.

A decision was made to try running the dynamic impact analyses in Code_Aster, as it has a reputation for being quite powerful, and is largely based on Python, which could make for easier interfacing with Pylaster. Though the learning curve with Code_Aster was onerous, ultimately, it did prove capable of both dynamically altering the time step during impact analyses, while still producing output at the time steps desired for audio sampling.

Figure 5.3 shows the initial geometry of a mallet impact model. The mallet head is modelled as a simple sphere and given initial position close to the bar's playing surface to reduce the number of time steps required for pre-impact travel. The mallet head material was defined with arbitrary material properties loosely approximating rubber, as finding reliable mechanical properties for rubber in the literature proved very difficult.

Figure 5.2 shows an example mallet impact force simulated via contact mechanics. This example is from an aluminum bar impact model with initial mallet head velocity of 1.75 metres per second. The resulting force curve has a lower peak than the fitted curves in Figure 5.2, indicating a softer mallet strike. Some fluctuations are noticeable during the initial rise in force, which may stem from interaction between the mallet head and bar. Adjusting the mallet head initial velocity or elastic modulus were each found to affect the peak force, impact duration and observed fluctuations. Softer mallet impacts were found to



Figure 5.4: A view from below an Adams five-octave marimba showing the support cable passing through several bars and resting on cable posts. Rubber wraps are visible that prevent the bars from striking the metal cable posts. The bar $F\sharp 3$ from Figure 5.5 is visible toward the right of the photo.

reduce model run times slightly, as models using smaller initial velocities or softer elastic moduli permitted larger dynamic time steps and required fewer iterations during the impact simulation. No attempt was made to match the simulated mallet impulse to either of the fitted curves in Figure 5.2. The relatively soft mallet impact, with its beneficial effect on run time, was deemed suitable to test and demonstrate the hybrid sound radiation model of this chapter.

Boundary Conditions

Numerous boundary conditions were explored for the idiophone bars in the mallet impact models. Modern marimba and vibraphone bars are commonly supported on cables that pass through holes drilled near the nodes of the fundamental (first flexural) vibration mode. Positioning the supports in this manner allows movement of the first mode with minimal restriction. Figure 5.4 gives a view of bars being supported by a cable on an Adams five-octave marimba. Figure 5.5 gives a detailed view of one of these bars after being removed from the instrument. The two support holes are clearly visible at either end of the cutaway in Figure 5.5. Note that the support holes are not parallel. The support hole nearest the centre of the instrument runs perpendicular to the end of the bar. The hole nearer the edge of the instrument is drilled at an angle so that a single cable may pass through all support holes in a straight line as the bars change in length along the keyboard.



Figure 5.5: An F \sharp 3 bar from an Adams five-octave marimba used in experimental modal measurements (see Appendix A). Oversized holes for the support cable are visible near either end of the cutaway.

The support holes drilled through the bar are significantly larger than the cable diameter. This makes it easier to pass the support cable through the hole and also serves to reduce any restriction on the bar's movement after it is struck. The amount of play in the cable support holes is enough that, on one Yamaha marimba, a bar could be lifted about 2 to 3 mm before encountering resistance from the support cable. The opposite effect was noted when pushing down on the bar. In that direction the support cable, already engaged to support the bar's self weight, stiffens considerably with less than 1 mm movement. This can be attributed to the help of cable support posts positioned within a few millimetres on either side of the bar. As the bar moves downward, the angle of the support cable changes quickly, creating a large stiffening effect.

These support conditions are ideal for marimbas and vibraphones. When striking a bar with a mallet, the support cable stiffens quickly, ensuring a large portion of the impact energy will serve to deform the bar (though some will be translated into the bar "bouncing" up and down on the cable). After the mallet strike, the amount of play in the support cable holes limits any restriction on bar vibration, particularly that of the fundamental mode.

Reproducing these precise boundary conditions in an FE model is challenging. In the upward direction there should be a gap, with no resistance, followed by a gradual, yet nonlinear increase in resistance (with excessive upward movement, the cable would lift off the cable supports and begin to lift adjacent bars). In the downward direction there should be an immediate nonlinear stiffening of resistance from the cable. The current version of Code_Aster does not provide an easy mechanism by which to define such a complex nonlinear boundary condition. Several attempts were made to approximate these boundaries, including: linear spring supports, bilinear spring supports and linear springs supports with an initial gap. The desired effect of such support conditions was to allow a bar model that was tuned using free-free boundary conditions to be used in the impact model without altering its modal behaviour. None of the FE boundary conditions considered were successful in doing so.

Linear spring supports, with very low stiffness, had little effect on the bar modes. However, when low-stiffness springs are used in the impact model, the energy imparted on the bar by

the mallet head translates primarily into bar motion instead of bar deformation. In other words, with low-stiffness springs the bar would simply bounce on the springs instead of deforming and vibrating. Increasing the spring stiffness sufficiently to cause bar deformation altered the modal behaviour.

Bilinear springs, with initial low stiffness followed by higher stiffness, created artifacts in the radiated bar sound, as did using spring with initial gaps.

The purpose of the bar impact model was to serve as an input to demonstrate the hybrid radiation model. While a higher degree of realism in this model is preferable it is not strictly necessary for its purpose. The decision was made to use fully fixed boundary conditions on specified support nodes. Figure 5.11 shows the restrained boundary nodes in the aluminum bar model in Section 5.5.1. These nodes are positioned at the ends of the bar cutaway, roughly below where support cable holes would be drilled. Only nodes on the sides of the bar are restrained. On an existing instrument, the support cable passing through the bar would contact the bar only at the ends of the hole, as there is invariably some sag along the length of the cable within the bar. Therefore support nodes in the FE model are positioned at the bar sides only.

Given that changes to bar modal behaviour imparted by the boundary conditions are unavoidable, and with a desire to observe tuned modes in the mallet impact and hybrid radiation models, the example bars described in Section 5.5 were tuned using the methods described in Chapter 4 and the boundary conditions described here.

Damping Model

The most ubiquitous damping model for time history analyses in commercial and open-source finite element software is Rayleigh damping (Rayleigh, 1945). Originally published between 1877 and 1878 in Lord Rayleigh’s two-volume *Theory of Sound*, the model has proliferated and endured. Its popularity is likely due to its simplicity, being a two-parameter model, and ease of implementation. Using this model, the global damping matrix for a finite element system is calculated as

$$\mathbf{C} = \alpha\mathbf{M} + \beta\mathbf{K}, \quad (5.2)$$

where \mathbf{C} is the system damping matrix, \mathbf{M} , is the system mass matrix, \mathbf{K} , is the system stiffness matrix, and α and β are the Rayleigh damping coefficients. Note that, while α and β in Equation 5.2 are typically arranged in the order shown, with α scaling the mass matrix and β scaling the stiffness matrix, this is not always the case. Importantly for the current work, these two parameters are reversed in Code_Aster.

Keeping with the notation used in Equation 5.2, damping ratio, ζ , for Rayleigh damping

Table 5.2: Rayleigh damping coefficients used for time domain modelling.

Material	α	β
Rosewood	13.877	$4.754 \cdot 10^{-7}$
Aluminum	4.4521	$8.185 \cdot 10^{-9}$

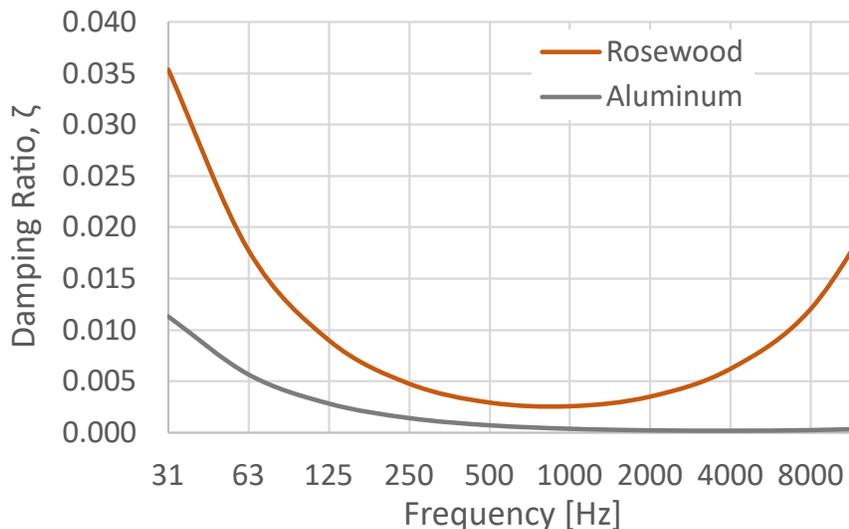


Figure 5.6: Damping models employed in the dynamic time domain finite element models in this work. Damping ratios vary according to the Rayleigh damping model in Equation 5.3 with coefficients per Table 5.2.

can be calculated as a function of frequency:

$$\zeta(f) = \frac{1}{2} \left(\frac{\alpha}{2\pi f} + \beta 2\pi f \right), \quad (5.3)$$

where $\zeta(f)$ is the damping ratio for a given frequency f (in Hz), and α and β are the Rayleigh damping coefficients from Equation 5.2. Thus Rayleigh damping is controlled by two parameters. Depending on the intended modelling application, an analyst can use these two parameters to precisely set the damping ratios for two different frequencies, or use them to tune a line of best fit over a number of desired (or measured) damping ratios at various frequencies.

Table 5.2 gives the Rayleigh damping coefficients used in this work to model rosewood and aluminum bars in the time domain. These values were determined from measurements of rosewood marimba bars from professional instruments, and the aluminum stock material used in fabricating prototype bars. More information on bar measurements is provided in Appendix A.

Figure 5.6 plots the damping ratio curves from Equation 5.3 for the parameter values in Table 5.2. The curve shapes in Figure 5.6 conform to expectations upon examining Equation 5.3. For low frequencies, mass-proportional damping controlled by the α parameter dominates the behaviour. As frequency increases, this mass-proportional terms dies off and the stiffness-proportional component, controlled by the β parameter takes over, with damping ratio approaching a linear increase with frequency. Therein lies one of the limitation of Rayleigh damping. With only two parameters available to tune behaviour, damping ratios can only be controlled within a limited frequency range. The β term in particular exerts a large effect, as a significant value of β will cause damping to rise rapidly in high frequencies. By contrast, as is seen for aluminum, a low value of β will keep damping low over a large range of frequencies. Finding middle ground between these extremes can be challenging with Rayleigh damping. Some compromise will often be required.

Alternative damping models were also considered for this work. The most promising model, *Generalized proportional damping* (GPD) (Adhikari, 2006), allows damping ratio to be defined as a general function of frequency (using a limited but versatile set of mathematical functions and operations), and then uses the input damping ratio curve to determine how the system damping matrix is calculated from the system mass and stiffness matrices. Though not tested in this work, conceptually this model provides great potential for defining more varied damping ratio curves than those shown in Figure 5.6 for Rayleigh damping. A drawback, however, is that calculation of the global damping matrix will be considerably more complicated than with Rayleigh damping. The more complex the input damping ratio curve used in GPD, the more time consuming it may be to calculate the system damping matrix. It may also be necessary to approximate the system damping matrix, if the input damping curve contains things like trigonometric or exponential functions. Thus, the additional computational demand of GPD will be exacerbated in situations where the global mass and stiffness matrices are changing. So this damping model may result in increased run times for nonlinear models. Nevertheless, GPD may well provide intriguing new options, particularly in situations where some portion of the system has consistent mass and stiffness matrices.

Unfortunately, at the time of this work, GPD was not available as a damping option in Calculix, Code_Aster, or any other available software (to the best of the author’s knowledge). Implementing GPD in Pylaster was considered, and partially completed. To perform a complete time domain mallet strike model with GPD in Pylaster would also have necessitated implementing contact mechanics in the code. Adding contact mechanics would have constituted a very large investment in implementation time for something that is ultimately not an integral part of the hybrid radiation model (which would function equally well with

Rayleigh damping or GPD). Thus the decision was made to use Code_Aster, and its available Rayleigh damping model, for the mallet strike time domain model.

5.3.3 Sound Radiation Modelling

In principle, the hybrid radiation model developed here can be employed with any sound radiation model that is compatible with the finite element bar models described in Chapter 3. Several methods of sound radiation modelling were considered, include the *Boundary Element Method* (BEM) (Ali and Rajakumar, 2004; Atalla and Sgard, 2015) and *Scaled Boundary Finite Element Method* (SBM) (Song, 2018; Liu et al., 2019a). Ultimately, the method adopted here was that of Bestle et al. (2017). This method was selected for its relative simplicity and ease of implementation. While the method may have some limitations, it was found to produce realistic sounds, and was deemed sufficient for demonstrating and evaluating the hybrid radiation model of this chapter.

The method described by Bestle et al., herein termed the *Surface Point Source Method* (SPS), is based on Huygen’s principle that a complex wavefront can be described as the combined behaviour of a set of simpler wavefronts (Bestle et al., 2017). Employing this principle, SPS calculates sound radiation by considering the surface motion of sound sources. A surface is subdivided into a number of smaller component elements, each with an associated position, tributary area and described motion over time. Each element is then treated as a monopole source. The resulting radiated sound thus becomes the summed contributions from each of these monopole sources.

In this work, the source object is an idiophone bar. The relevant sound radiating surface will be the bar’s playing surface (upon which the mallet is struck). One constraint of this approach is that the position at which radiated sound is modelled (the “listener position” or “virtual microphone position”) must be such that a line drawn from this position to any point on the playing surface will be unobstructed. In practical terms this simply means the listener position must be facing the bar’s playing surface.

Using this method, radiated sound at a given listener position is calculated as:

$$p(t) = \sum_n p_n \left(t - \frac{r_n}{c_a} \right), \quad (5.4)$$

where $p(t)$ is the sound pressure at the listener position at time t , $p_n(t)$ is the sound pressure at the listener position contributed by surface element n , r_n is the distance from the centroid of surface element n to the listener position, and c_a is the speed of sound in air. Note that the operand $t - \frac{r_n}{c}$ will generally not be an integer multiple of the sampling time step. Thus

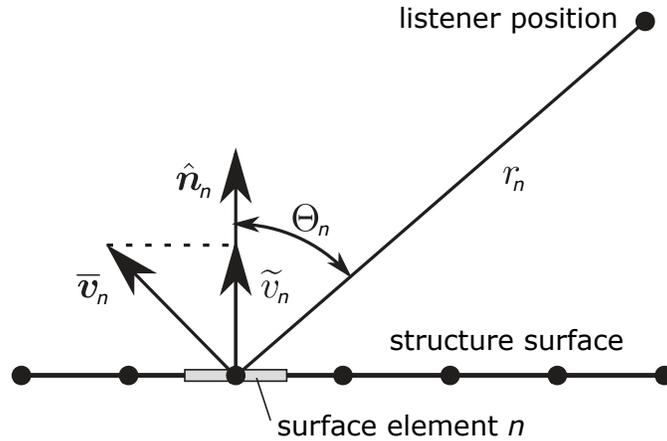


Figure 5.7: Definition of geometric parameters used in sound radiation modelling by the Surface Point Source Method. Figure adapted from Bestle et al. (2017).

each time series $p_n(t)$ will need to be interpolated when evaluating Equation 5.4. Linear interpolation was used in evaluating these time series in this work.

Contributions from the individual surface elements is calculated as:

$$p_n(t) = \frac{\rho_0 S_n}{4\pi r_n} \dot{\tilde{v}}_n(t) \cos(\Theta_n), \quad (5.5)$$

where ρ_0 is the density of air, S_n is the area of surface element n , $\dot{\tilde{v}}_n(t)$ is the first derivative of surface normal velocity with respect to time and Θ_n is the angle between the surface normal and vector from the element centroid to the listener position. Surface normal velocity, $\tilde{v}_n(t)$, is calculated as:

$$\tilde{v}_n(t) = \bar{\mathbf{v}}_n^T(t) \hat{\mathbf{n}}_n, \quad (5.6)$$

where $\bar{\mathbf{v}}_n(t)$ is a column vector of velocities in three dimensions for surface element n at time t , and $\hat{\mathbf{n}}_n$ is a unit vector normal to surface element n . These parameters, along with others from Equations 5.4 and 5.5 are shown in Figure 5.7, which is adapted from Bestle et al. (2017).

5.3.4 Mode Selection

An important factor in the hybrid radiation model is not just how many modes from the finite element modal analysis are included in the Prony model, but *which* modes are included. Some modes present in the modal analysis results will not radiate sound efficiently toward the listener position. However, if the frequency of one of these modes is near that of a mode that does radiate sound efficiently, both modes may appear in the hybrid model response

(when only one *should*, in fact, appear).

This situation can be remedied by selecting only modes that radiate significant sound toward the listener position. A method for selecting such modes is outlined below.

Normalized effective modal mass

A useful modal parameter commonly reported by finite element software is the *effective modal mass* of each mode in a given direction (global axis). These parameters describe the amount of mass that is active in a given mode in a given direction. The effective modal mass is computed as:

$$m_{\text{eff},n,z} = \frac{(\boldsymbol{\phi}_n^T \mathbf{M} \mathbf{u}_z)^2}{\boldsymbol{\phi}_n^T \mathbf{M} \boldsymbol{\phi}_n}, \quad (5.7)$$

where $m_{\text{eff},n,z}$ is the effective modal mass of mode n in the global z -direction, $\boldsymbol{\phi}_n$ is the shape of mode n , \mathbf{M} is the global mass matrix, and \mathbf{u}_z is a vector wherein each entry corresponding to a displacement degree-of-freedom in the global z -direction is set equal to one, while all other entries are set to zero. In this work mode shapes were mass-normalized, and thus:

$$\boldsymbol{\phi}_n^T \mathbf{M} \boldsymbol{\phi}_n = 1. \quad (5.8)$$

Typically of greater interest than effective modal mass is the *normalized effective modal mass*. This is the effective modal mass in a given direction, divided by the total effective mass in that same direction. It is computed as:

$$\hat{m}_{\text{eff},n,z} = \frac{m_{\text{eff},n,z}}{m_{\text{total},z}}, \quad (5.9)$$

where $\hat{m}_{\text{eff},n,z}$ is the normalized effective modal mass of mode n in the global z -direction, and $m_{\text{total},z}$ is the total unrestrained mass of the model in the z -direction. In practice, $m_{\text{total},z}$ will be very nearly equal to the total mass of the bar, subtracting only a very small amount of mass that is tributary to finite element nodes with boundary conditions applied in the z -direction.

The normalized effective modal mass is commonly used to assess whether a given set of modes have sufficiently captured the total mass of a structure, active in a given direction. For example, an engineering standard may specify that a response spectrum analysis, which employs modal decomposition, must include sufficient modes to capture at least 90% of the total mass of a structure in a given direction.

Of greater interest in this work, is employing normalized effective modal mass to eliminate inconsequential modes. Any modes that do not contribute significant mass in the direction

of the listener position should be discarded. Removing such modes helps to avoid any interference they may cause in fitting parameters to modes that do contain significant active mass in the relevant direction.

In the bar models of this work, the global z -axis is oriented perpendicular to each bar's playing surface. Thus bar motion in the z -direction is perpendicular to the surface that radiates sound to the listener position, as described in Section 5.3.3. Normalized effective modal mass in the z -direction is therefore the relevant measure for including or excluding a given mode in the hybrid sound radiation model.

In this work, the following approach was employed in the selection of modes for the hybrid sound radiation model:

1. Identify a frequency range of interest and calculate modes until their frequencies surpass this range.
2. Calculate normalized effective modal mass in the z -direction for each mode. Eliminate any modes with normalized effective mass less than 10^{-10} in the z -direction.
3. Select all remaining modes in the frequency range of interest for use in the hybrid model.

Note that the threshold of 10^{-10} for normalized effective modal mass should be considered a suggested starting point; this threshold has not been exhaustively explored. However, the normalized effective modal masses of inconsequential modes in this work had magnitudes in range of 10^{-26} or smaller (see Tables 5.7 and 5.11). Thus, establishing a threshold for modal inclusion was hardly ambiguous. Figure 5.18 in Section 5.5.2 illustrates the effect of selecting only modes with significant normalized effective modal mass.

This technique for selecting modes was largely successful, though certainly not perfect. Exploration of other approaches to mode selection, perhaps based on some metric of sound radiation strength, is an avenue for potential future research.

5.3.5 Numerical Time Integration Compensation

The numerical time integration method used in the time history analysis will affect the spectral content and damping behaviour of the bar's response in that model. The modes predicted by the initial modal analysis will be present in the bar's time history response, but their modal frequencies will differ from those in the modal analysis results. Damping behaviour of these modes in the time history model may also differ from that predicted by the input damping model. To ensure accurate fitting of parameters for the Prony model, modal frequencies and damping parameters must first be adjusted to account for the effects

of numerical time integration. This section details how to adjust modal frequencies and damping parameters for two common types of numerical time integration.

Two popular methods of numerical time integration commonly used for finite element time history analyses are *Newmark's Method* (Newmark, 1959) and the *Hilber-Hughes-Taylor Method* (Hilber et al., 1977). Conveniently, the Hilber-Hughes-Taylor Method (HHT) reduces to Newmark's Method when certain values of input parameters are specified. This section will outline numerical time integration compensation for HHT, noting that the same adjustments will work with Newmark's Method using appropriate parameter values. In principle, other numerical time integration schemes would also work with the hybrid sound radiation model, provided their effects on frequency and damping of each mode can be predicted for compensation prior to fitting parameters of the Prony model.

Integration Method Parameters

Newmark's method is controlled by two dimensionless input parameters, β_w and γ_w (herein given the subscript 'w' to distinguish from similar variables in this chapter). These parameters control how much the acceleration at the end of a time step enters into the calculation of velocity (influenced by γ_w) and displacement (influenced by β_w) at the end of that same time step. Higher values of either parameter produce greater influence of acceleration at the end of the time step. Setting these parameters will determine the method's stability and imposed numerical errors (Newmark, 1959). A common approach when applying Newmark's method is to set $\beta_w = 0.25$ and $\gamma_w = 0.5$. Using these specific parameters is termed the *average acceleration method* or *midpoint rule*, which is unconditionally stable (Cook et al., 2001). Other values of β_w and γ_w may also be used, with differing effects on stability and accuracy.

HHT expands on Newmark's Method by adding another parameter, α_h , and using it to determine β_w and γ_w . Specifically, for unconditional stability with HHT, these parameters are set to:

$$\beta_w = \frac{1}{4}(1 - \alpha_h)^2; \quad \gamma_w = \frac{1}{2}(1 - 2\alpha_h); \quad \alpha_h \in \left[-\frac{1}{3}, 0\right]. \quad (5.10)$$

With β_w and γ_w determined by α_h in Equation 5.10, HHT becomes a single parameter method. A common choice, which is also the default option in Code_Aster, is to set $\alpha_h = -0.1$. It can also be observed from Equation 5.10 that setting $\alpha_h = 0$ will produce $\beta_w = 0.25$ and $\gamma_w = 0.5$, in which case HHT reduces to the average acceleration method.

Numerical Damping and Frequency Compression

Hilber et al. (1977) provide discussion of stability, numerical damping and frequency effects

for various cases with no input damping. For cases with input damping, such as those models used in this work, an overview of HHT's effects on frequencies and damping is provided by Hilber (1976). Note that the literature tends to speak in terms of *period elongation* when describing spectral changes imparted by time integration methods. Discussing period, rather than frequency, is common in structural engineering where frequencies of interest are typically less than 1 Hz. In this work it is more convenient to speak in terms of frequency. Thus, what the literature terms *period elongation*, herein will be called *frequency compression*.

The relevant steps to predict frequency and damping behaviour in the time history model are outlined below. For a more thorough discussion, refer to Hilber (1976), from whom the following method is sourced.

The effects a numerical time integration scheme will have on frequencies and damping behaviour can be assessed by analyzing an *amplification matrix*. This amplification matrix can be used to recursively calculate displacement, velocity and acceleration of a dynamic single degree-of-freedom system. For such a system undergoing free vibration (absent any applied force), the amplification matrix can be applied as follows:

$$\mathbf{u}_{i+1} = \mathbf{A}_m \mathbf{u}_i, \quad (5.11)$$

$$\mathbf{u}_i = \begin{Bmatrix} u_i \\ \dot{u}_i \Delta T \\ \ddot{u}_i \Delta T^2 \end{Bmatrix}, \quad (5.12)$$

where \mathbf{A}_m is the amplification matrix, u_i , \dot{u}_i and \ddot{u}_i are displacement, velocity and acceleration at time step i , respectively, and ΔT is the time step (inverse of sampling frequency).

Looking at Equation 5.11 it is clear that the system's behaviour is calculated via repeated application of the amplification matrix, \mathbf{A}_m . Thus this matrix contains all relevant information about how the system will respond over time. Of particular interest are the eigenvalues of the amplification matrix.

For HHT with non-zero input damping, Hilber (1976) gives the following amplification matrix:

$$\mathbf{A}_m = \begin{bmatrix} 1 + \beta_w A_{31} & 1 + \beta_w A_{32} & \frac{1}{2} + \beta_w (A_{33} - 1) \\ \gamma_w A_{31} & 1 + \gamma_w A_{32} & 1 + \gamma_w (A_{33} - 1) \\ A_{31} & A_{32} & A_{33} \end{bmatrix}, \quad (5.13)$$

$$\left. \begin{aligned}
 A_{31} &= -\frac{\Omega^2}{D} \\
 A_{32} &= -\frac{1}{D} [2\zeta\Omega + (1 + \alpha_h)\Omega^2] \\
 A_{33} &= 1 - \frac{1}{D} \left[1 + 2\zeta\Omega + (1 + \alpha_h)\frac{\Omega^2}{2} \right] \\
 D &= 1 + 2\gamma_w\zeta\Omega + (1 + \alpha_h)\beta_w\Omega^2 \\
 \Omega &= 2\pi f\Delta T
 \end{aligned} \right\} \quad (5.14)$$

where f and ζ are the frequency and damping ratio for the single-degree-of-freedom system under consideration; other terms as previously defined or defined in Equation 5.14.

The eigenvalues of \mathbf{A}_m can be calculated using the following matrix invariants:

$$\left. \begin{aligned}
 A_1 &= 1 - \frac{1}{D} \left[\zeta\Omega + \frac{1}{2}\Omega^2 \left((1 + \alpha_h)(\gamma_w + \frac{1}{2}) - \alpha_h\beta_w \right) \right] \\
 A_2 &= -\frac{1}{D} \left[2\zeta\Omega + \Omega^2 \left(\frac{1}{2}\gamma_w + 2\alpha_h(\gamma_w - \beta_w) \right) \right] \\
 A_3 &= \frac{1}{D}\alpha_w\Omega^2 \left(\beta_w - \gamma_w + \frac{1}{2} \right)
 \end{aligned} \right\} \quad (5.15)$$

where A_1 is half the trace of \mathbf{A}_m , A_2 is the sum of the principal minors of \mathbf{A}_m , and A_3 is the determinant of \mathbf{A}_m . Using these invariants, the eigenvalues of \mathbf{A}_m can be found by solving the characteristic equation:

$$\lambda^3 - 2A_1\lambda^2 + A_2\lambda - A_3 = 0, \quad (5.16)$$

where λ is an eigenvalue of \mathbf{A}_m . Hilber (1976) shows that within the range of α_h values in Equation 5.10, the eigenvalues of \mathbf{A}_m , which satisfy Equation 5.16, will occur as two principal eigenvalues (appearing as complex conjugate pairs) and a third spurious eigenvalue. Let the complex-valued principal eigenvalues take the form:

$$\lambda_{1,2} = A \pm Bj. \quad (5.17)$$

The frequency and damping ratio with which the single-degree-of-freedom system will respond in the time history model can then be calculated as:

$$\bar{f} = \frac{\bar{\Omega}}{2\pi\Delta T}, \quad (5.18)$$

$$\bar{\zeta} = \frac{-\ln(A^2 + B^2)}{2\bar{\Omega}}, \quad (5.19)$$

$$\bar{\Omega} = \arctan\left(\frac{B}{A}\right), \quad (5.20)$$

where \bar{f} is the *numerical frequency* and $\bar{\zeta}$ is the *total numerical damping*. These are the frequency and damping ratio with which a single-degree-of-freedom system, with natural frequency f and input damping ζ , will respond in a time history model using HHT. Hilber (1976) showed that this same procedure can be applied to a linear, time invariant multi-degree-of-freedom system with numerous modes. For each mode in the system, with modal frequency f and damping ratio ζ , Equations 5.15 through 5.20 can be applied to determine that mode's numerical frequency \bar{f} and total numerical damping ratio $\bar{\zeta}$.

Section 5.3.6 outlines the procedure for fitting modal parameters in the hybrid sound radiation model developed in this chapter. One of the benefits of this hybrid radiation model is that modal frequencies and damping ratios are known in advance, leaving only amplitudes and phase angles to be determined by fitting. To ensure the hybrid radiation model produces the same system response that would otherwise have been produced by a complete time history simulation, the known frequencies and damping ratios used for parameter fitting must be the numerical frequencies and total numerical damping ratios calculated from Equations 5.18 and 5.19.

Figure 5.8 plots numerical frequency \bar{f} as a function of modal frequency f for a system modelled via HHT with $\alpha_h = -0.1$, a sampling frequency of 48 kHz and a damping ratio of $\zeta = 0.001$ at all frequencies. Note that the input damping ratio need not be uniform for this method to apply. The function shown in Figure 5.8, which takes a modal frequency as input and returns the corresponding numerical frequency resulting from a given time integration method, is herein called the *frequency compression function*. Also shown in Figure 5.8 is the frequency compression function for Newmark's Method with parameters corresponding to the average acceleration method.

Figure 5.9 plots total numerical damping, $\bar{\zeta}$, as a function of frequency taking a Rayleigh damping model as inputs. The specific damping model used as input is that for aluminum with Rayleigh input parameters as shown in Table 5.2. It is clear from Figure 5.9 that damping ratio is affected much more in higher frequencies. Note that the average acceleration method is non-dissipative, in that it does not affect damping ratios.

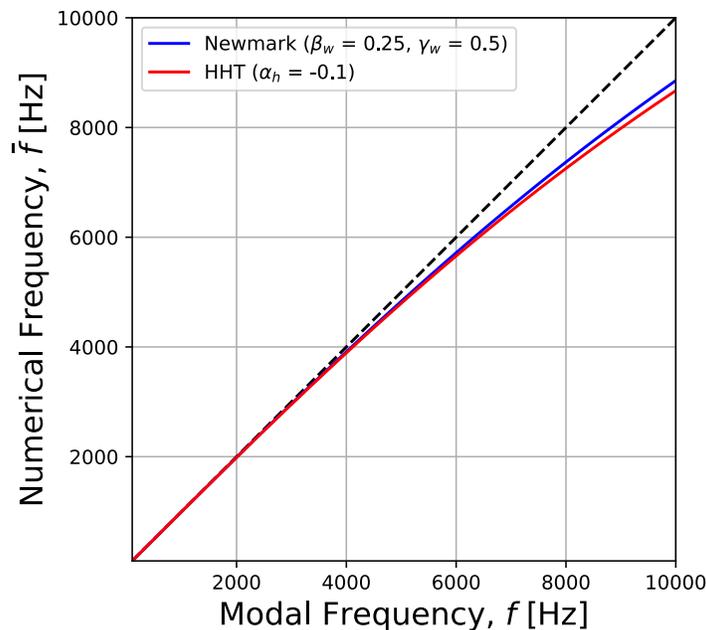


Figure 5.8: Example frequency compression functions showing the numerical frequencies resulting from different time integration methods. The red line shows the effect of HHT on frequency with $\alpha_h = -0.1$, while the blue line represents Newmark’s Method with $\beta_w = 0.25$ and $\gamma_w = 0.5$. Both examples use an input damping ratio, $\zeta = 0.001$ for all frequencies and a sampling rate of 48 kHz.

5.3.6 Modal Parameter Fitting

One of the main efficiencies of this hybrid sound radiation model is in the fitting of modal parameters. The first portion of the hybrid model is calculated using the sound radiation methods of Section 5.3.3 with displacement results from the dynamic time history analysis of Section 5.3.2. The second portion of the hybrid model is comprised of a sum of decaying sinusoids, which is sometimes called a *Prony model* after the classic work by [de Prony \(1795\)](#). Using notation convenient for describing parameters, this sum of decaying sinusoids can be written as

$$p_p(t) = \sum_{n=1}^N A_n \exp(-2\pi f_n \zeta_n t) \cos(2\pi f_n t + \varphi_n), \quad (5.21)$$

where $p_p(t)$ is the radiated sound pressure at time t , N is the number of vibration modes considered, and $\{A_n, \zeta_n, f_n, \varphi_n\}$ are the amplitude, damping ratio, frequency (in Hz) and phase angle (in radians) of the n^{th} mode, respectively. For convenience, Equation 5.21 uses the damping ratio, ζ_n , which expresses damping in a given mode as a ratio of *critical damping* at that frequency. This convention is adopted for easy compatibility with the damping models in Section 5.3.2.

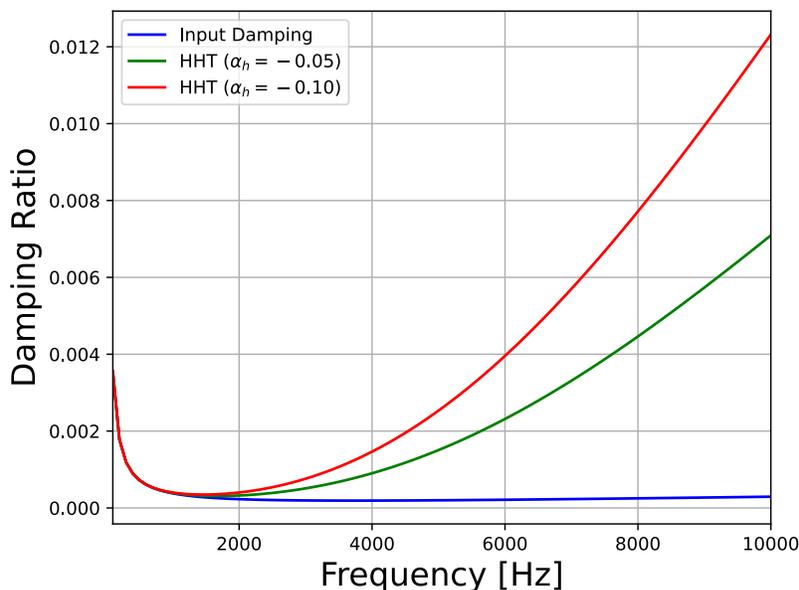


Figure 5.9: Examples of total numerical damping ratios resulting from HHT numerical time integration with different values of α_h . The input damping ratios are those from a Rayleigh model using input parameters for aluminum from Table 5.2. Results shown assume a sampling rate of 48 kHz.

A primary benefit of using the Prony model of Equation 5.21 in this hybrid radiation model, is that the modal frequencies, $\{f_n\}$, and damping ratios, $\{\zeta_n\}$, are already known. Modal frequencies are determined by running a modal analysis on the same model (including the same boundary conditions) that is used for the time domain modelling. Note that these modal frequencies will need to be adjusted, as discussed in Section 5.3.8. Once the adjusted modal frequencies are known, the damping ratio of each mode can be calculated from Equation 5.3, or any other damping model used for the time history modelling.

Traditional Prony analysis, which fits the parameters of Equation 5.21 to an observed time series (or similar measured data) for a given number of modes, N , is performed in two parts. The first, more complicated part, fits frequency and damping parameters for each of the modes. The second, simpler part, fits amplitudes and phase angles for each mode, and is dependent on the frequency and damping parameters already determined. With both frequency and damping parameters known in advance, the hybrid radiation model has the advantage of requiring only the second stage of Prony analysis, skipping the first, more complicated part.

The complete procedure of Prony analysis can be found in numerous texts, including Marple (1987). The second step, once modal frequencies and decay parameters are known,

proceeds as follows. For a prony model with N damped sine waves, Equation 5.21 can be rewritten in discrete time as:

$$x[k] = \sum_{n=1}^N h_n z_n^{k-1}, \quad (5.22)$$

$$h_n = A_n \exp(j\varphi_n), \quad (5.23)$$

$$z_n = \exp[(\alpha_n + j2\pi f_n)\Delta T], \quad (5.24)$$

$$\alpha_n = -2\pi f_n \zeta_n, \quad (5.25)$$

where $x[k]$ is the signal at discrete time index k , h_n is the n^{th} residue of the system - a complex number containing both the amplitude A_n and phase angle φ_n of the n^{th} mode, z_n is the n^{th} pole of the system - a complex number containing the damping coefficient α_n and frequency f_n of the n^{th} mode, while ΔT is the time step between samples (inverse of the sampling frequency) and $j = \sqrt{-1}$.

Equation 5.22 can be rewritten in matrix form as:

$$\mathbf{Z}\mathbf{h} = \mathbf{x}, \quad (5.26)$$

where:

$$\mathbf{Z} = \begin{bmatrix} 1 & 1 & \dots & 1 \\ z_1^1 & z_2^1 & \dots & z_N^1 \\ z_1^2 & z_2^2 & \dots & z_N^2 \\ \vdots & \vdots & & \vdots \\ z_1^{K-1} & z_2^{K-1} & \dots & z_N^{K-1} \end{bmatrix}, \mathbf{h} = \begin{Bmatrix} h_1 \\ h_2 \\ \vdots \\ h_N \end{Bmatrix}, \mathbf{x} = \begin{Bmatrix} x[1] \\ x[2] \\ \vdots \\ x[K] \end{Bmatrix}. \quad (5.27)$$

Once the system's poles are known, matrix \mathbf{Z} can be populated. Then, this matrix can be used along with the input signal vector \mathbf{x} to solve Equation 5.26 for the vector of unknown residues \mathbf{h} . Traditional Prony analysis determines the system poles by finding the roots of a characteristic equation. In this work, the poles are known in advance from a finite element modal analysis and an input damping model, so matrix \mathbf{Z} can be populated immediately. Therefore the length, K , of signal \mathbf{x} , from the finite element time domain model, need only be large enough to provide suitable results for the residue vector \mathbf{h} . Reducing the required length of vector \mathbf{x} , from the computationally expensive finite element time domain model, is the major efficiency of this hybrid radiation model approach.

A few important practicalities of solving Equation 5.26 warrant mentioning here. First, the poles of Equation 5.24 occur in complex conjugate pairs. So, for each modal frequency selected from the finite element modal analysis results, two complex conjugate poles must be

defined by inputting one with negative frequency. If desired, the amplitudes of these pairs of complex conjugate poles may later be combined into a single pole with positive frequency.

A second practicality of solving Equation 5.26, is that frequencies taken from the finite element modal analysis, and used to populate matrix \mathbf{Z} , must be adjusted to account for any effects of numerical time integration in the finite element time domain model. After fitting the amplitudes and phase angles for each mode, frequencies for the entire hybrid model can be readjusted to undo these effects. This procedure is explained in detail in Section 5.3.8. For the moment, it will suffice to say that frequencies taken from the finite element modal analysis must be adjusted before being used to define the poles in Equation 5.24. Any such adjustments will be compensated later in the process.

A third important practicality is that the matrix \mathbf{Z} will have dimension $K \times N$, where generally $K > N$. So the system described by Equation 5.26 will be overdetermined. This system must therefore be solved using an appropriate approximate method, such as least squares. In this work, the vector of residues was determined using the equation:

$$\mathbf{h} = \mathbf{Z}^+ \mathbf{x}, \quad (5.28)$$

where \mathbf{Z}^+ is the Moore-Penrose *pseudoinverse* of matrix \mathbf{Z} . Note that calculating the pseudoinverse of a large matrix can produce significant computational demand. Should this become problematic, more efficient methods of solving Equation 5.26 may be employed, potentially taking advantage of the Vandermonde structure of matrix \mathbf{Z} (see, for example, Section 11.5 of Marple (1987)). In this work, with the signal lengths considered for fitting modal amplitudes and phase angles, Equation 5.28 performed adequately.

Fitting Window

An important aspect of the modal parameter fitting described in this section is the input signal used to fit the parameters. This input signal appears as vector \mathbf{x} in Equations 5.26 and 5.27. Note that it is **not** the entire output of sound radiation calculated from the finite element time history model that is used to fit parameters for the Prony model. Instead, a specific portion of the time history results are used for modal fitting. This selected portion of time history sound radiation output signal is that which occurs over a specified length of time called a *fitting window*.

The fitting window is defined by two parameters: the *fitting window start time*, and the *fitting window duration* (which could as easily be called *length*). Each of these parameters are measured in modelled time. So, for example, a fitting window could be defined with a start time of 8 milliseconds, and a duration of 10 milliseconds. In such a case, the portion of

signal used to fit modal parameters would be the segment occurring between 8 milliseconds and 18 milliseconds in simulation time.

In this chapter the fitting window start time and duration are given the variables b_{fw} and e_{fw} , respectively. Together, these parameters will determine the accuracy of the hybrid sound radiation model, as well as its computational efficiency. Generally speaking, a later start time and longer duration should lead to great accuracy. However, increasing either of these variables will necessitate a longer finite element time history model, thereby increasing computational demand.

Best performance of the hybrid model will occur with the earliest and shortest fitting window that can still produce suitably accurate results. Selection of fitting window parameters is discussed in Section 5.4.

5.3.7 Signal Crossover

Once the initial time domain model has been run, and modal parameters for the Prony model of Equation 5.21 have been determined, the time domain results and Prony model are combined to create a single output signal. The beginning of this output signal will be comprised solely of time domain model results. The end, or “tail”, of the output signal will consist solely of the Prony model. In most cases, this tail will constitute the majority of the output signal.

Connecting these two portions of the signal has potential to produce artifacts in the output signal. Specifically, if a significant mode in the time history response is not included in the Prony tail, the sudden removal of this mode could create artifacts in the hybrid radiation model output. To avoid the potential of such artifacts, a smooth transition from the initial time domain model results to the Prony tail will be employed. Between the two segments there will be a *transition zone* where the two components are mixed at varying ratios. These three zones are illustrated in Figure 5.1.

Transition Zone

The transition zone, where the time domain results and Prony tail mix, or “crossover”, must occur over a time span in which the two component signals are defined. Parameters for the Prony tail are determined over the length of the fitting window described in Section 5.3.6. Thus, the Prony tail is undefined for any time $t < b_{fw}$. So b_{fw} represents the earliest time at which signal crossover can begin.

The hybrid radiation model will be maximally efficient when the required length of the finite element time domain simulation is minimized (as this simulation is by far the most

computationally expensive component). To facilitate the parameter fitting described in Section 5.3.6, the time domain model must extend to the end of the fitting window. So the time domain simulation must extend to at least the end time of the fitting window, e_{fw} . Thus, to avoid any unnecessary extension of the time domain model, the transition zone should not extend beyond e_{fw} .

Ideally, the Prony model will contain a sufficient number of modes to capture all relevant (perceptible) spectral content in the finite element time history model. However, the modelling procedure should allow for the possibility that the time history results may include some spectral content not present in the Prony tail. Any such content will be attenuated over the transition zone as the output signal changes from purely time history results to purely Prony model. Thus, to avoid the potential for artifacts in the hybrid model output, it is best to provide a gradual transition between the two component signals, avoiding any abrupt attenuation of excess spectral content.

The preceding paragraphs establish limits on the earliest time at which the transition zone may begin, and the latest time at which it can end. A longer transition zone would lend itself well to a smooth/gradual transition. Therefore, to achieve the smoothest transition possible between the two component signals, the transition zone will be set to its maximum allowable length. With b_{fw} the earliest allowable transition start time, and e_{fw} the latest allowable transition end time, the transition zone is thus set equal to the fitting window described in Section 5.3.6.

Mixing Function

Content of the hybrid signal over the transition zone will be a mix of sound radiation calculated from the finite element time history model results, and the Prony tail with fitted modal parameters. The overall hybrid signal is thus calculated as:

$$p_h(t) = \begin{cases} p_t(t) & t < b_{tz} \\ m(t)p_t(t) + (1 - m(t))p_p(t) & b_{tz} \leq t \leq e_{tz} \\ p_p(t) & t > e_{tz} \end{cases}, \quad (5.29)$$

where $p_h(t)$ is the hybrid sound radiation model output at time t , $p_t(t)$ is radiated sound at time t calculated from the time history model, $p_p(t)$ is radiated sound at time t calculated from the Prony model, $m(t)$ is a mixing function at time t , b_{tz} is the time at which the signal transition begins, and e_{tz} is the time at which the signal transition ends.

As described in the previous section, it is desirable to have a smooth transition between

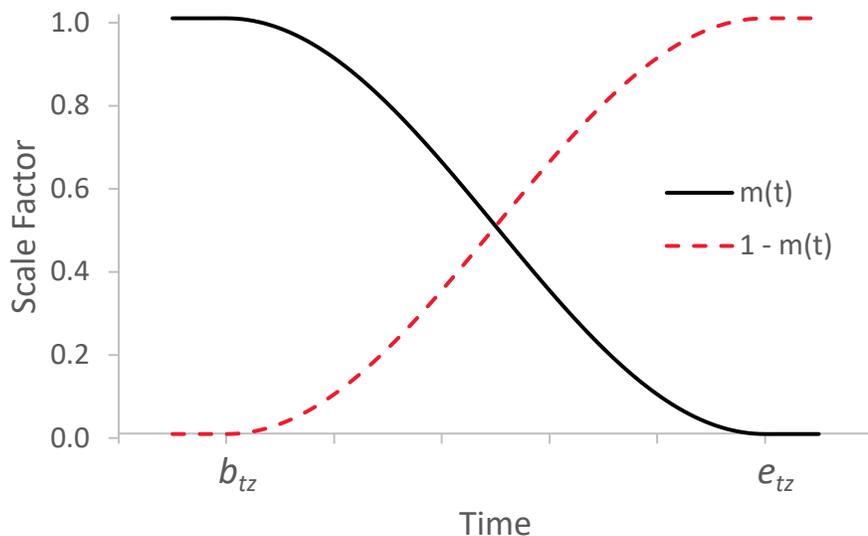


Figure 5.10: Mixing function $m(t)$ from Equation 5.31 along with $1 - m(t)$. The two lines illustrate how $p_t(t)$ and $p_p(t)$ are scaled over the transition zone before being summed to create $p_h(t)$.

the time history results and Prony results. Thus the transition time limits are set as:

$$b_{tz} = b_{fw}; e_{tz} = e_{fw}. \quad (5.30)$$

To further ensure a smooth transition between the two component signals, the mixing function, $m(t)$, in Equation 5.29 is defined such that it is continuously differentiable with zero slope at each end of the transition zone. This is achieved using a half-wavelength of a cosine wave:

$$m(t) = \begin{cases} 1 & t < b_{tz} \\ \frac{1}{2} \left(1 + \cos \left(\pi \frac{(t-b_{tz})}{(e_{tz}-b_{tz})} \right) \right) & b_{tz} \leq t \leq e_{tz} \\ 0 & t > e_{tz} \end{cases}. \quad (5.31)$$

Figure 5.10 plots $m(t)$ and $1 - m(t)$ over the transition zone to demonstrate how $p_t(t)$ and $p_p(t)$ are scaled over this time.

5.3.8 Reversing Numerical Time Integration Compensation

Section 5.3.5 discussed corrections required to use predicted modal frequencies and input damping parameters to fit the Prony tail to output from the time history model. These corrections were necessary as the time integration method used in the time history model

will alter the frequencies and damping behaviour of each mode.

For completeness, the final step in the hybrid radiation model is to undo the adjustments made in Section 5.3.5. Specifically, it is desirable to remove any numerical damping and/or frequency deviations imparted by the time integration method. Of the time integration methods considered in this work, the Hilber-Hughes-Taylor Method (HHT) will impose both numerical damping and frequency changes, while Newmark’s Method will affect frequencies only, leaving input damping unchanged.

Damping Adjustment

Using the hybrid model approach, reversing the effects of numerical damping is quite straightforward. At this stage in the analysis, for the n^{th} mode in the Prony tail, the following computational steps have been taken:

1. Modal frequency, f_n , was predicted by a finite element modal analysis
2. Input damping ratio, ζ_n , was computed from the (Rayleigh) damping model for the input modal frequency, f_n (Section 5.3.2).
3. Numerical frequency, \bar{f}_n , and total numerical damping ratio, $\bar{\zeta}_n$, were computed based on the time integration method used for the time history model (Section 5.3.5).
4. Modal amplitude, A_n , and phase angle, φ_n , were fit to the time history results (Section 5.3.6).

To reverse the effect of time integration on modal damping the Prony tail is simply constructed using input modal damping ratios, ζ_n , from Step 2, rather than the total numerical damping ratios of Step 3. It is possible for the total numerical damping for a given mode to be significantly larger than the corresponding input damping. In such a case, switching from the total numerical damping ratio, back to the input damping ratio, may create significantly more contribution from that mode in the hybrid model output. Care must then be taken to avoid integration methods with excessive numerical damping. If numerical damping is so large as to effectively remove a mode from the time history response, then switching back to the input damping value may have an effect similar to amplifying noise.

This approach to reversing numerical damping effects does not remove such effects from the time history excitation segment of the hybrid model response. However, given the short duration of this segment, and given that the system is under excitation for much of it, the effect of leaving numerical damping in the time history excitation was deemed negligible.

Frequency Adjustment

In contrast to numerical damping effects, reversing frequency compression is not as simple as changing a mode's numerical frequency, \bar{f}_n , back to its original frequency, f_n . As shown in Figure 5.8, \bar{f}_n can differ significantly from f_n . These differences become even larger at smaller sampling rates. Rapidly switching modal frequencies over the crossover segment is expected to have a non-negligible effect on timbre, particularly for modes where these frequencies differ significantly. Instead, reversing frequency compression must be accomplished in a manner that affects all segments of the hybrid model.

Reversing frequency compression is accomplished via the follow steps:

1. Create a hybrid model output signal, $\bar{\mathbf{x}}_h$, using the parameters $(A_n, \zeta_n, \bar{f}_n, \varphi_n)$ for each mode (note that numerical frequency, \bar{f}_n , remains).
2. Zero-pad the hybrid model signal from Step 1 as necessary to produce the desired frequency resolution in the following steps. If necessary, low-pass filter this signal to remove any artifacts created by zero-padding.
3. Compute a Fourier transform, $\bar{\mathbf{X}}_h$, of the zero-padded hybrid model $\bar{\mathbf{x}}_h$ from Step 2.
4. Compute the corresponding vector of frequencies, $\bar{\mathbf{f}}_h$, based on the sampling frequency used in creating the hybrid model signal. Values in this vector correspond to the numerical frequencies (\bar{f}) affected by numerical time integration.
5. Compute a vector of original frequencies, \mathbf{f}_h^* , by taking numerical frequencies in $\bar{\mathbf{f}}_h$, and computing the inverse of the frequency compression function (Figure 5.8) for the applicable time integration method. Note that frequencies in \mathbf{f}_h^* will be non-uniformly spaced.

The frequency content $\bar{\mathbf{X}}_h$, and frequency vector \mathbf{f}_h^* now represent a non-uniformly sampled Fourier transform of the desired output signal.

6. Let \mathbf{f}_h be a vector of uniformly spaced frequencies corresponding to a Fourier transform of the desired output signal, \mathbf{x}_h , with applicable sampling frequency. Compute the frequency content, \mathbf{X}_h , of the desired output signal by interpolating $\bar{\mathbf{X}}_h$ (sampled at non-uniformly spaced frequencies \mathbf{f}_h^*), at the uniformly spaced frequencies \mathbf{f}_h .
7. Finally, compute the inverse Fourier transform of \mathbf{X}_h to produce the hybrid model output signal \mathbf{x}_h . Trim as necessary to remove padded zeros.

The output signal \mathbf{x}_h produced in Step 7 will be the final output of the hybrid sound radiation model. Its frequency content will correspond to the original input modal frequencies absent the effects of numerical time integration. This applies to all three segments of the model, not just the Prony tail.

In cases where finding an analytical inverse to the frequency compression function in Figure 5.8 would be difficult (i.e. HHT), Step 5 is performed numerically.

Any reasonable one-dimensional interpolation scheme may be selected for Step 6. In this work, linear interpolation was employed along with significant zero-padding for fine frequency resolution.

5.4 Model Parameter Selection

5.4.1 Fitting Window Parameters

The ability to set any required parameters of the hybrid radiation model in advance is an important aspect of creating a versatile and readily applicable method. While things like material properties and boundary conditions are important parameters of the bar model, these parameters would also apply if employing other approaches to sound radiation (e.g. boundary element methods). Some parameters, while not necessary for *all* possible methods, do appear in *some* other methods. For example, modal models appear in many different applications in musical acoustics and beyond. Therefore, selection of the number of modes to include in the hybrid radiation model may be informed by experience with modal models in general.

Most importantly, the parameters that apply uniquely to the hybrid radiation model are those that define the fitting window described in Section 5.3.6. Specifically, they are the fitting window *start time*, and fitting window *duration*. Setting these parameters will impact the accuracy of the sound radiation simulation, and also determine the required length of the time domain modelling. By inspection, it is evident that a later fitting window start time and longer duration should result in increased accuracy. At the same time, the total required simulation time for the finite element time history model will be the sum of these two parameters. Beyond the end of the fitting window, the hybrid model is fully defined by the Prony tail, and thus no further finite element results are required. So the sum of the fitting window start time and duration will determine the necessary length of finite element time domain modelling, which comprises the vast majority of required computational effort.

With these two parameters controlling both output accuracy and computational effort, selecting suitable values for each in advance will be imperative in creating a useful method.

Procedures for setting each parameter are outlined in the following sections. Parametric tests of these parameters are described in Section 5.6.

Fitting Window Start Time

Two considerations will be key in selecting an appropriate start time for the fitting window. They are:

- Has all nonlinear behaviour of the system dissipated?
- Have all sound waves produced by nonlinear behaviour passed the listener/microphone position?

For the bar models considered in this work, the significant nonlinear effects stem from interaction between the mallet head and the bar during the mallet strike. An obvious minimum criterion for the fitting window start time thus becomes the time at which the mallet ceases to be in contact with the bar (herein called the “lift-off” time). After this time, there should be no further nonlinear interaction in the system.

In addition to the lift-off time, this work also considered the maximum time required for sound waves to propagate through the bar material. This can be viewed as the time required for the “information” that the mallet has lifted off the bar to permeate the system. In the models considered here, the time required for wave propagation is small compared to the lift-off time, but is included nevertheless. The goal of this section is to find a suitably accurate yet reasonably efficient combination of window start time and duration. This need not necessarily be the absolute minimum suitable combination of the two.

Finally, after considering mallet lift-off and time required for waves to traverse the bar, the necessary time for sound waves to propagate from the bar to the listener/microphone position must be included.

The fitting window start time is thus calculated by the simple formula,

$$b_{fw} = t_{lo} + t_{pb} + t_{pa}, \quad (5.32)$$

where b_{fw} is the fitting window start time, t_{lo} is the mallet lift-off time, t_{pb} is wave propagation time within the bar and t_{pa} is wave propagation time from the bar to the microphone position.

The lift-off time, t_{lo} can be determined from reports of the number of elements in contact in the finite element time history model. After this number has peaked and begun to drop, the lift-off time will be the point at which the number of elements in contact returns to zero. To be thorough, this number can continue to be checked throughout the fitting window duration

to confirm there were no secondary impacts from the bar rebounding and contacting the mallet again (in which case the lift-off time can be extended accordingly).

The wave propagation time, t_{pb} is taken as the time required for a shear wave to twice travel the longest linear distance in the bar. In this work effects of the bar cutaway are ignored and the longest linear distance is simply considered the bar's diagonal length in three-dimensions. For an isotropic material, shear waves will propagate with equal speed in all directions. Thus, t_{pb} can be calculated for a bar of isotropic material as:

$$t_{pb} = \frac{2\sqrt{L^2 + W^2 + T^2}}{c_{si}}, \quad (5.33)$$

$$c_{si} = \sqrt{\frac{G}{\rho}}, \quad (5.34)$$

where $\{L, W, T\}$ are the length, width and end thickness of the bar, respectively, c_{si} is the shear wave speed in an isotropic material, G is the shear modulus and ρ the material density. For an orthotropic material there will be six different shear moduli, in pairs of two, corresponding to each possible ordered combination of different principal material directions. For an orthotropic material, Equations 5.33 and 5.34 are combined and rearranged to naively approximate t_{pb} as:

$$t_{pb} \approx 2\sqrt{\rho \left(\frac{L^2}{G_{LT}} + \frac{W^2}{G_{WT}} + \frac{T^2}{E_T} \right)}, \quad (5.35)$$

where G_{LT} and G_{WT} are shear moduli between the subscripted directions and E_T is the elastic modulus in the thickness-direction, and $\{L, W, T\}$ indicate the length, width and thickness directions, respectively.

Sound propagation in air is taken simply as the longest distance from any corner of the bar's playing surface to the virtual microphone position, divided by the speed of sound in air. Thus:

$$t_{pa} = \frac{1}{c_a} \max_{i=1}^4 \|\mathbf{r}_{c,i} - \mathbf{r}_\ell\|_2, \quad (5.36)$$

where c_a is the speed of sound in air, $\mathbf{r}_{c,i}$ is the coordinate vector of the i^{th} playing surface corner of the bar, and \mathbf{r}_ℓ is the coordinate vector of the microphone position.

Fitting Window Duration

The fitting window duration is selected based on a much simpler criterion. It is expected that one full period of the lowest modal frequency in the tuned bar should provide a sufficient

fitting window duration. Thus the fitting window duration is expressed as:

$$d_{fw} = \frac{1}{f_0}, \quad (5.37)$$

where d_{fw} is the fitting window duration and f_0 is the fundamental frequency of the bar.

Required Simulation Time

The required time over which the mallet strike and subsequent bar response should be simulated thus becomes the sum of the fitting window start time and duration:

$$d_{sim} = b_{fw} + d_{fw}, \quad (5.38)$$

where d_{sim} is the required duration of the time history simulation, and other terms are as defined above. Looking at Equation 5.38 it is clear that adding or subtracting time from either the fitting window start time or duration will have equal impact on the overall simulation time needed.

The following sections provide details of the example aluminum and rosewood bar models used to evaluate the hybrid radiation model. Section 5.6 outlines parametric tests performed on both the aluminum and rosewood example models to test performance of the fitting window start time and duration criteria suggested above.

5.5 Example Models

Two example models were developed and tested to demonstrate the hybrid radiation model. The models varied in their bar materials, geometry, tuning ratios, sampling frequencies and fitting window parameters. Each model was tuned with boundary conditions as described in Section 5.3.2. These same boundary conditions were applied in the time domain hybrid sound radiation models. Tuning the bar with the same boundary conditions as the hybrid radiation model serves to ensure the bar radiates with the same frequencies to which it was tuned.

5.5.1 Aluminum Bar

Model description

The first example bar used to test the hybrid radiation model is an aluminum bar with outer dimensions matching those of Prototype 2 described in Section 4.5. While the outer

Table 5.3: Hybrid radiation aluminum bar example model dimensions.

Dimension	Length	Width	End Thickness
Value (mm)	365	63.5	18

Table 5.4: Hybrid radiation aluminum bar example model modal tuning ratios.

Mode	V1	V2	V3	T3
Tuning Ratio	1	4	6	18

bar dimensions (repeated in Table 5.3) match those of the prototype bar, the tuning of this model differs. This example bar was tuned with the same boundary conditions used in the mallet impact simulation, so as to produce expected modal frequencies in that simulation. Table 5.4 outlines the tuning ratios used to shape the cutaway of the example bar. Hybrid radiation model parameters used in this example are outlined in Table 5.5, while Table 5.6 gives the mallet strike and virtual microphone positions. The tuned bar geometry is plotted in Figure 5.11.

To assess performance of the hybrid radiation model, the finite element dynamic time history model of mallet impact and subsequent bar response was run for a simulation time of two seconds. By comparison, only about 12 milliseconds of finite element simulation time was required to produce the hybrid model, which was then used to generate two seconds worth of radiated sound. Thus, for a radiated sound two seconds long, the hybrid radiation model requires only 0.6% of the dynamic time history simulation time, which comprises the vast majority of required computational effort.

Model results and discussion

Table 5.7 shows the modal frequencies and normalized effective modal mass for the first 20 modes of the aluminum bar model in Figure 5.11. Of the 20 modes calculated, nine had significant effective modal mass in the z-direction and were selected for inclusion in the hybrid radiation model. Interestingly, mode 12 in Table 5.7, corresponding to mode T3 of the model, did not exhibit significant effective modal mass. Thus, despite being tuned by the bar geometry shown in Figure 5.11, this mode does not contribute significantly to sound

Table 5.5: Hybrid radiation aluminum bar example model parameters.

Parameter	Value
Fitting Window Start Time (ms)	7.73
Fitting Window Duration (ms)	3.82
Number of Modes Included	10
Sampling Frequency (kHz)	48

Table 5.6: Hybrid radiation example model mallet impact and virtual microphone positions. The same positions are used for both the aluminum and rosewood example modes. However, the two example models have different bar dimensions (see Tables 5.3 and 5.8), thus the relative positions of the mallet impact and virtual microphone differ between the models.

Parameter	X [mm]	Y [mm]	Z [mm]
Mallet strike centroid	150	20	0
Virtual microphone position	185	40	-1500

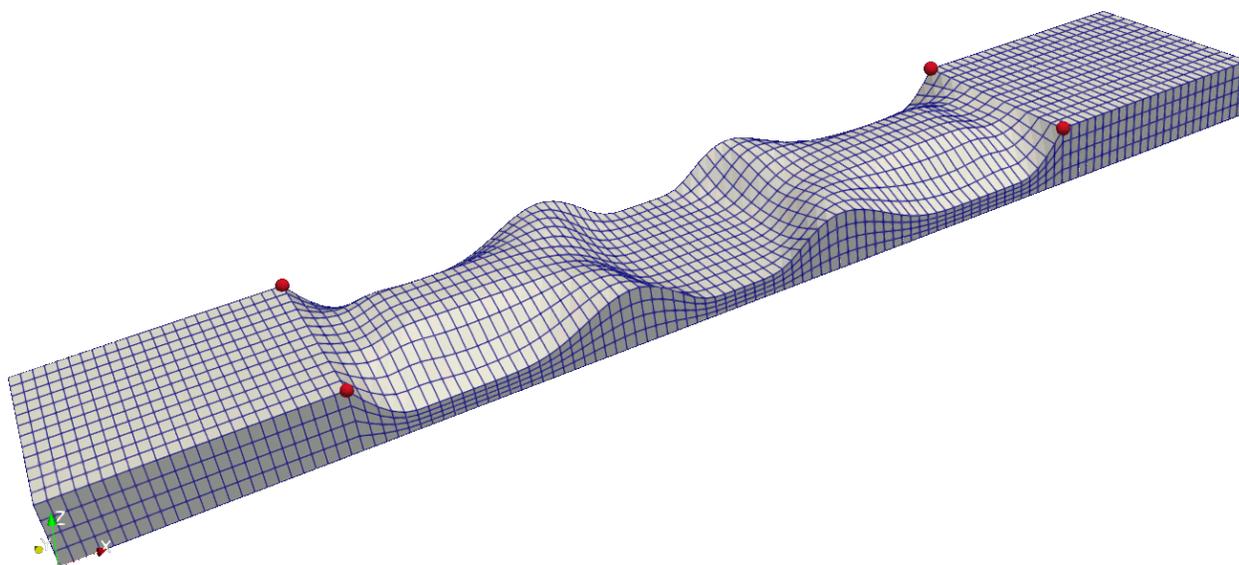


Figure 5.11: Hybrid radiation aluminum bar example model meshed geometry. Red circles represent support locations in the finite element model. Origin of the global coordinate system is shown at the corner of the bar’s playing surface. The model is comprised of 5,760 8-node hexahedral elements.

radiation in this model.

Figure 5.12 shows the resulting sound radiation captured at the virtual microphone position from the finite element time history and hybrid radiation models. To facilitate a fair comparison, the hybrid radiation results in this plot are shown *before* applying the frequency warping described in Section 5.3.8. This is necessary as the finite element time history results do not contain any such warping.

Viewed in the time domain, the waveforms in Figure 5.12 are clearly indistinguishable. Figure 5.12 was split into two panels to avoid having one waveform obscure the other when plotted together. Note also that this same example hybrid radiation model is shown in Figure 5.1, wherein the three distinct segments of the hybrid model are shown. Figure 5.1 gives a sense of just how much the required length of time history modelling can be reduced by employing this hybrid modelling approach.

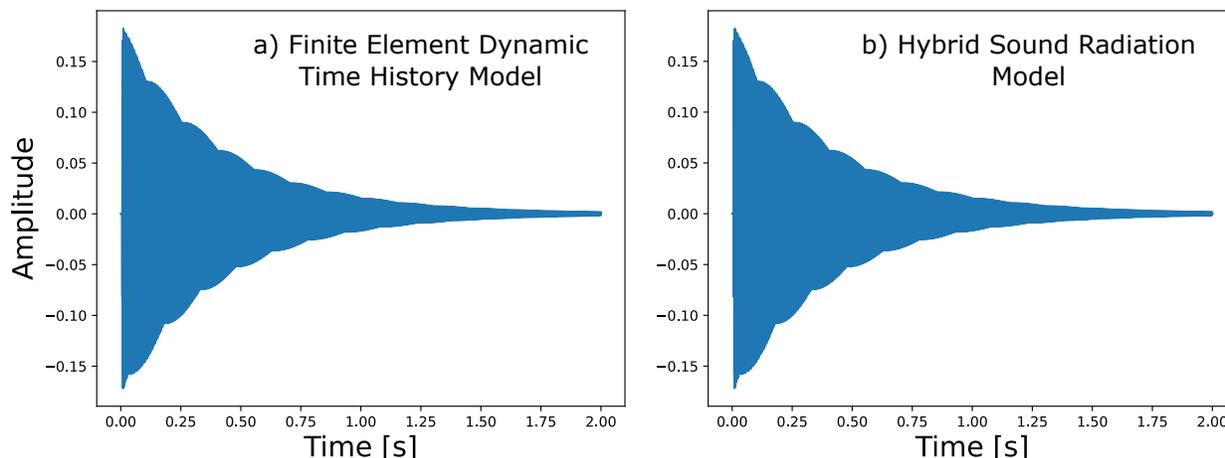


Figure 5.12: Waveforms of radiated sound at the virtual microphone position for the aluminum bar example model. (a) Results produced using finite element dynamic time history analysis throughout the simulation[†]. (b) Results from the hybrid sound radiation model described in this chapter, before applying the damping and frequency corrections of Section 5.3.8[‡].

[†]http://www.music.mcgill.ca/caml/lib/exe/fetch.php?media=publications:AL_TH.wav

[‡]http://www.music.mcgill.ca/caml/lib/exe/fetch.php?media=publications:AL_HY_NoCorr.wav

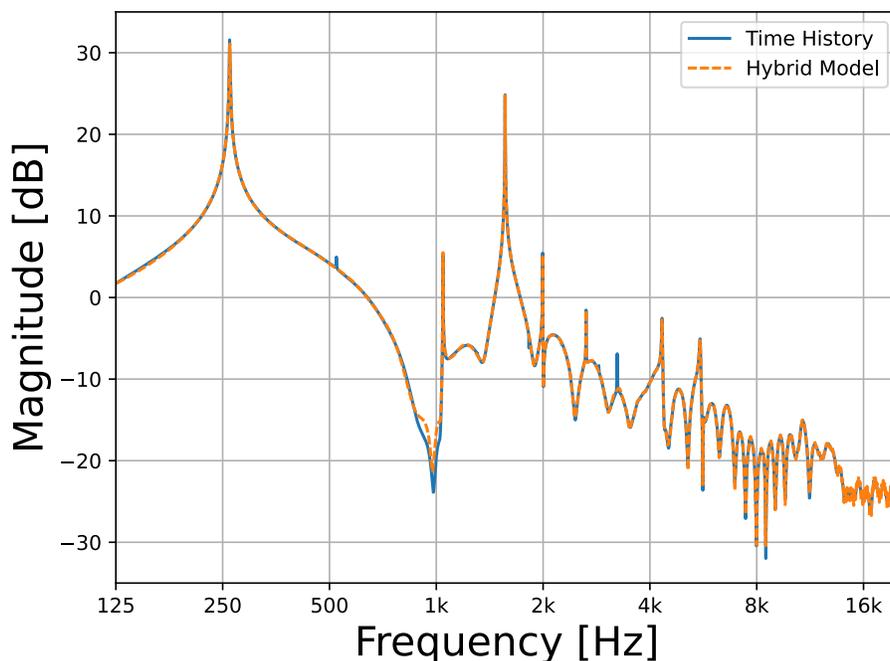


Figure 5.13: Frequency content of radiated sound waves shown in Figure 5.12 for the example aluminum bar time history and hybrid models. Hybrid sound radiation model results are shown before applying the damping and frequency corrections of Section 5.3.8.

Table 5.7: Aluminum bar example model, mode frequencies and normalized effective modal mass in the z-direction. Values reported are from modal analysis in Code_Aster.

Mode #	Frequency [Hz]	Normalized Effective Modal Mass (z-axis)	Include in model?
1	261.6	0.012686	✓
2	1046.5	0.001878	✓
3	1569.8	0.732680	✓
4	1943.9	9.66e-26	
5	2008.3	0.024274	✓
6	2582.6	2.04e-27	
7	2677.1	6.78e-10	✓
8	2914.2	1.53e-28	
9	3296.4	9.94e-27	
10	4120.5	2.26e-26	
11	4479.8	0.177019	✓
12	4709.3	6.74e-27	
13	4829.6	4.35e-29	
14	5527.9	3.33e-28	
15	5862.3	0.030865	✓
16	6000.3	0.002999	✓
17	8514.3	0.013583	✓
18	8749.2	1.43e-28	
19	9565.7	0.000262	✓
20	10902.4	3.46e-30	

Figure 5.13 shows the frequency content of the sound waves from Figure 5.12. Observable differences between the two spectra are generally very minor. A small peak in the time history results spectrum appears just over 500 Hz; this is curious as, per Table 5.7, the model does not contain a mode near that frequency. The one noticeable difference between the two spectra occurs around 3300 Hz. A spectral peak occurs near this frequency in the time history results of Figure 5.13, while this peak does not appear in the hybrid model results. Per Table 5.7 the corresponding mode at 3296 Hz has very low normalized effective modal mass and was thus not included in the hybrid model. Interestingly, of the ten modes excluded due to low modal mass, only this mode is visibly absent in Figure 5.13.

Figure 5.14 demonstrates the impact of reversing numerical time integration effects. The blue curve in Figure 5.14 plots the hybrid model spectrum for the bar in Figure 5.11 before reversing time integration effects (this is the same curve shown in Figure 5.13), while the orange curve plots the same data after reversing these effects. Two predictable outcomes are readily apparent in Figure 5.14, the first of which are frequency effects. Comparing peaks between the orange and blue spectra in Figure 5.14, it is evident that peaks in the

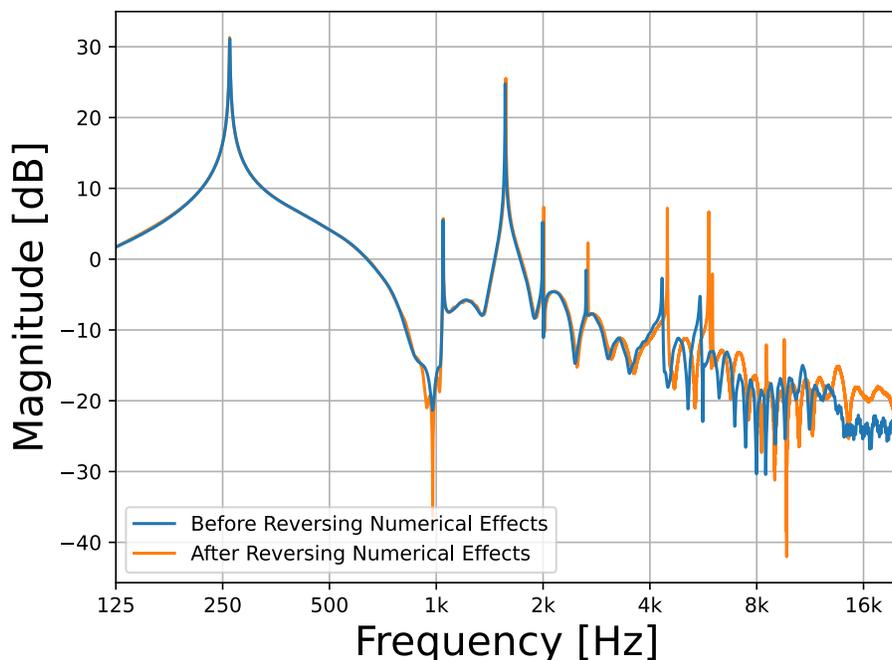


Figure 5.14: Frequency content of the aluminum hybrid radiation model, with[†] and without[‡] reversing the effects of numerical time integration.

[†]http://www.music.mcgill.ca/caml/lib/exe/fetch.php?media=publications:AL_HY_Corr.wav

[‡]http://www.music.mcgill.ca/caml/lib/exe/fetch.php?media=publications:AL_HY_NoCorr.wav

orange spectrum are shifted slightly right compared to those in the blue spectrum. In other words, the frequencies of these spectral peaks have increased. This effect is more noticeable in higher frequencies, with obvious changes in peaks above 4 kHz and negligible effect on the fundamental frequency near 250 Hz. Such results are as expected, given the frequency compression functions in Figure 5.8, which show larger differences in numerical frequency at higher modal frequencies. Notwithstanding amplitude differences due to damping corrections, the orange spectrum in Figure 5.14 has a similar shape as the blue spectrum, but stretching to the right as frequency increases. This verifies that the frequency adjustment method described in Section 5.3.8 is functioning well.

The second notable changes in Figure 5.14 are the increased magnitudes of peaks in the orange spectrum. These increases in spectral peak magnitudes at each modal frequency are a direct result of the damping adjustments described in Section 5.3.8. As shown in Figure 5.9, the HHT method used in this example model increases damping more at higher frequencies. Reversing this effect reduces damping more for modes at higher frequencies, leading to greater increases in their spectral peak magnitudes.

Importantly, there are also examples of peaks in the orange spectrum of Figure 5.14 that

Table 5.8: Hybrid radiation rosewood bar example model dimensions.

Dimension	Length	Width	End Thickness
Value (mm)	406	58	24

Table 5.9: Hybrid radiation rosewood bar example model modal tuning ratios.

Mode	V1	V2	V3	T1	T2
Tuning Ratio	1	3	4	7	10

do not appear to have a corresponding peak in the blue spectrum. Specifically, these are the second, lower peak around 6 kHz and the two peaks above 8 kHz. Clearly adjusting these modes' damping factors to reverse the effects of numerical time integration has increased their prominence in the spectrum. Cases such as these should be handled with caution, as increasing the amplitude of a previously unnoticeable mode is similar in nature to amplifying noise. In such a scenario, the analyst may compare output spectra from the hybrid model before and after reversing time integration effects, and decide if any newly present modes should, in fact, be removed.

5.5.2 Rosewood Bar

Model description

The second example bar used to test the hybrid radiation model is a rosewood bar with outer dimensions matching an A3 bar measured on a Yamaha 5-octave marimba. While the outer bar dimensions (repeated in Table 5.8) match those of the measured A3 bar, the tuning of this model differs. This example bar was tuned with the same boundary conditions used in the mallet impact simulation, so as to produce expected modal frequencies in that simulation. The bar was tuned to note F3 rather than A3 and has tuning ratios outlined in Table 5.9. Hybrid radiation model parameters used in this example are outlined in Table 5.10. Tuned bar geometry is plotted in Figure 5.15.

To assess performance of the hybrid radiation model, the finite element dynamic time history model of mallet impact and subsequent bar response was run for a simulation time

Table 5.10: Hybrid radiation rosewood bar example model parameters.

Parameter	Value
Fitting Window Start Time (ms)	8.03
Fitting Window Duration (ms)	5.73
Number of Modes Used	12
Sampling Frequency (kHz)	48

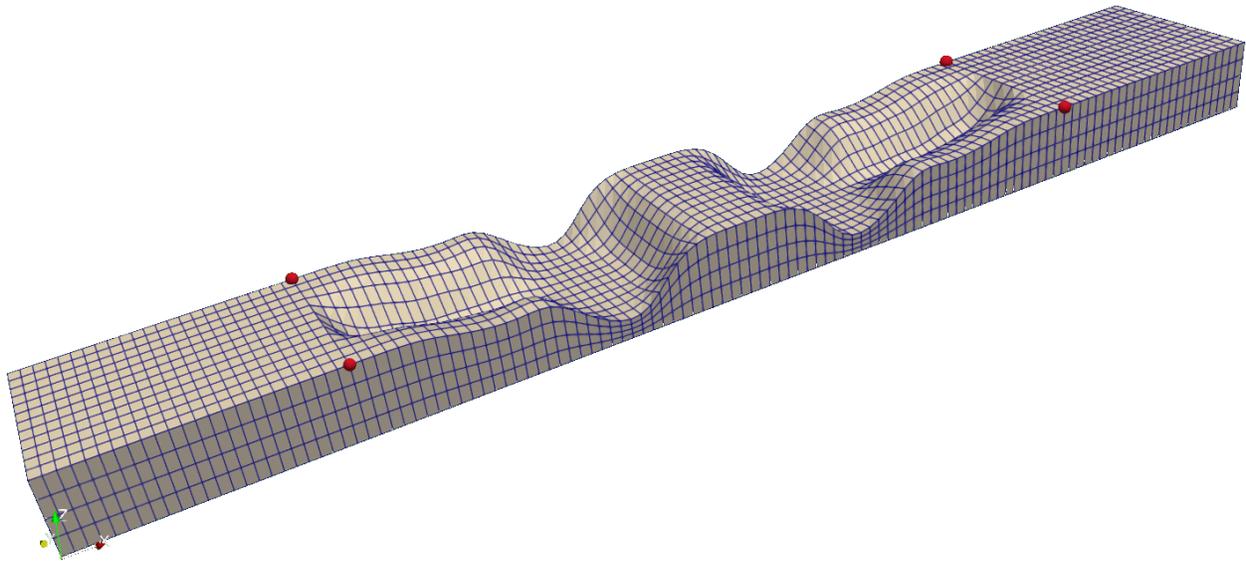


Figure 5.15: Hybrid radiation rosewood bar example model meshed geometry. Red circles represent support locations in the finite element model. Origin of the global coordinate system is shown at the corner of the bar’s playing surface. The model is comprised of 5,760 8-node hexahedral elements.

of one second. By comparison, only about 14 milliseconds of finite element simulation time was required to produce the hybrid model, which was then used to generate a full second worth of radiated sound. Thus, for a radiated sound one second long, the hybrid radiation model requires only 1.4% of the dynamic time history simulation time, which comprises the vast majority of required computational effort.

Model results and discussion

Table 5.11 shows the modal frequencies and normalized effective modal mass for the first 24 modes of the rosewood bar model in Figure 5.15. Of the 24 modes calculated, 12 had significant effective modal mass in the z -direction and were selected for inclusion in the hybrid radiation model.

Notably, modes 7 and 10 in Table 5.11, corresponding to modes T1 and T2 of the model, do not exhibit significant effective modal mass. Thus, as was seen in the aluminum bar example, once again these tuned torsional modes do not contribute significantly to sound radiation in this model. This is almost certainly attributable to the rigid boundary conditions (shown in Figure 5.15) applied to support the bar during the mallet strike. With rigid supports positioned on the outer edges of the bar, it is not surprising that the torsional modes exhibit little effective modal mass. On a real marimba, there is substantial play between the cable supporting a bar and the hole through which it passes (see Figures 5.4 and 5.5). Thus actual

boundary conditions would be much less restrictive to torsional motion. See Section 5.3.2 for further discussion of boundary conditions.

Figure 5.16 shows the resulting sound radiation captured at the virtual microphone position in the finite element time history and hybrid radiation models. To facilitate a fair comparison, the hybrid radiation results in the figure are shown *before* applying the frequency adjustment described in Section 5.3.8. This is necessary as the finite element time history results do not contain any such adjustment.

Figure 5.18 shows the possible effect of including superfluous modal frequencies in the hybrid model. The blue curve in Figure 5.18 is the same spectrum shown for the hybrid model in Figure 5.17, which uses only those 12 modes indicated in Table 5.11 for inclusion in the model. The red curve in Figure 5.18 shows a hybrid model of the rosewood bar that includes *all* 24 modes listed in Table 5.11. While the two curves overlap for the majority of frequencies, clear differences are apparent near 1 kHz. Around this frequency, five spectral peaks are visible in the red curve that do not appear in the blue curve. These peaks correspond to five modes in Table 5.11 with negligible normalized effective modal mass. Looking at the time history model spectrum in Figure 5.17 it is clear that these additional peaks do not appear in results from the time history model, and thus do not belong in the hybrid model. While only five of the twelve modes with low modal mass in Table 5.11 produce noticeable peaks in the red spectrum of Figure 5.18, their inclusion is enough to create an audible change in timbre. This example thus demonstrates the importance of selecting appropriate modes for use in the hybrid model.

5.6 Fitting Window Parametric Tests

Parametric tests of fitting window start time and duration were performed to assess the suitability of values suggested for these parameters in Section 5.4. Tests were performed for both the aluminum example model in Section 5.5.1 and rosewood example from Section 5.5.2. To focus on changing fitting window parameters, and eliminate other choices from consideration, parametric tests were run using *all* modes shown in Tables 5.7 and 5.11, not just those identified as significant. Results for the aluminum model are plotted in Figure 5.19, while Figure 5.20 gives similar results for the rosewood model.

To assess the performance of different fitting window parameters, a quantitative measure of output quality was required. Quality was measured by quantifying the similarity between hybrid models produced using different fitting window parameters, and the signals produced using time history analysis. Signal similarity, and thereby fit quality, was calculated using a

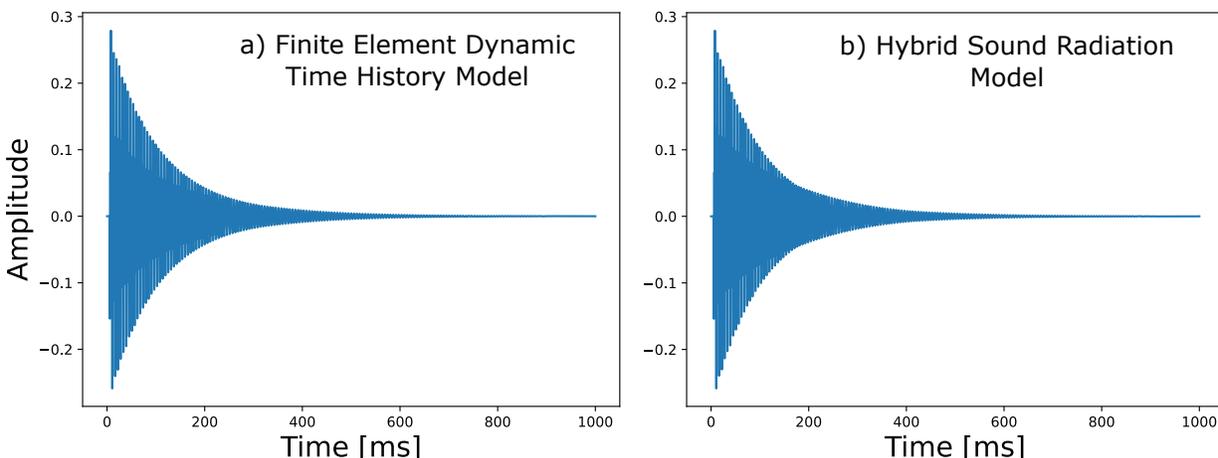


Figure 5.16: Waveforms of radiated sound at the virtual microphone position for the rosewood bar example model. (a) Results produced using finite element dynamic time history analysis throughout the simulation[†]. (b) Results from the hybrid sound radiation model described in this chapter, before applying the damping and frequency corrections of Section 5.3.8[‡].

[†]http://www.music.mcgill.ca/caml/lib/exe/fetch.php?media=publications:RW_TH.wav

[‡]http://www.music.mcgill.ca/caml/lib/exe/fetch.php?media=publications:RW_HY_NoCorr.wav

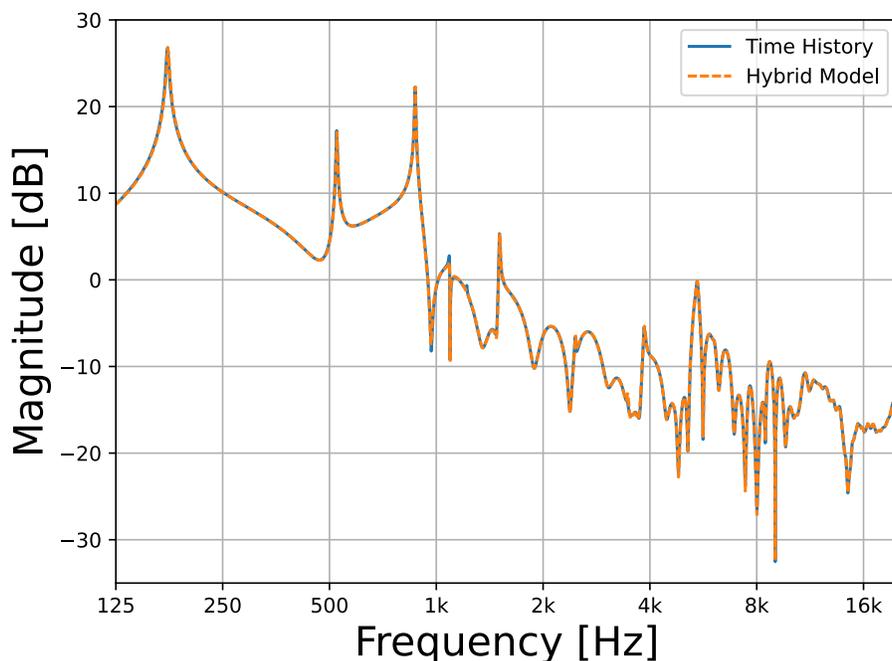


Figure 5.17: Frequency content of radiated sound waves shown in Figure 5.16 for the example rosewood bar time history and hybrid models. Hybrid sound radiation model results are shown before applying the damping and frequency corrections of Section 5.3.8.

Table 5.11: Rosewood bar example model, mode frequencies and normalized effective modal mass in the z-direction. Values reported are from modal analysis in Code_Aster.

Mode #	Frequency [Hz]	Normalized Effective Modal Mass (z-axis)	Include in model?
1	174.7	0.002688	✓
2	523.9	0.000126	✓
3	873.0	0.969546	✓
4	950.9	6.57E-28	
5	1006.8	4.10E-28	
6	1092.7	0.000262	✓
7	1222.3	6.41E-29	
8	1305.5	5.23E-30	
9	1511.6	0.000382	✓
10	1746.1	1.01E-30	
11	2008.6	4.77E-30	
12	2491.2	0.025398	✓
13	2985.9	2.07E-30	
14	3531.9	0.000003	✓
15	3694.4	5.61E-31	
16	3948.6	0.000109	✓
17	4003.7	3.84E-30	
18	4928.2	2.97E-31	
19	5614.5	0.000075	✓
20	5731.7	0.000664	✓
21	5746.9	0.000135	✓
22	6469.0	7.62E-08	✓
23	6682.4	3.02E-33	
24	7104.9	4.08E-33	

dot product in the time domain:

$$\text{Fit Quality} = \frac{\mathbf{x}_1^T \mathbf{x}_2}{\|\mathbf{x}_1\|_2 \|\mathbf{x}_2\|_2}, \quad (5.39)$$

where \mathbf{x}_1 and \mathbf{x}_2 are the two signals to be compared.

In producing Figures 5.19 and 5.20, signal comparison was limited to the Prony tail segment of the hybrid model (the green segment in Figure 5.1) and corresponding portion of the time history results. The time history excitation segment of the hybrid model (the blue segment in Figure 5.1) is taken directly from the time history results. Thus, including this initial portion would have unfairly increased the fit quality rating of all hybrid models. The crossover segment (orange in Figure 5.1) is a mixture of time history and modal results; its inclusion would also have increased the fit quality of all hybrid models, though to a lesser

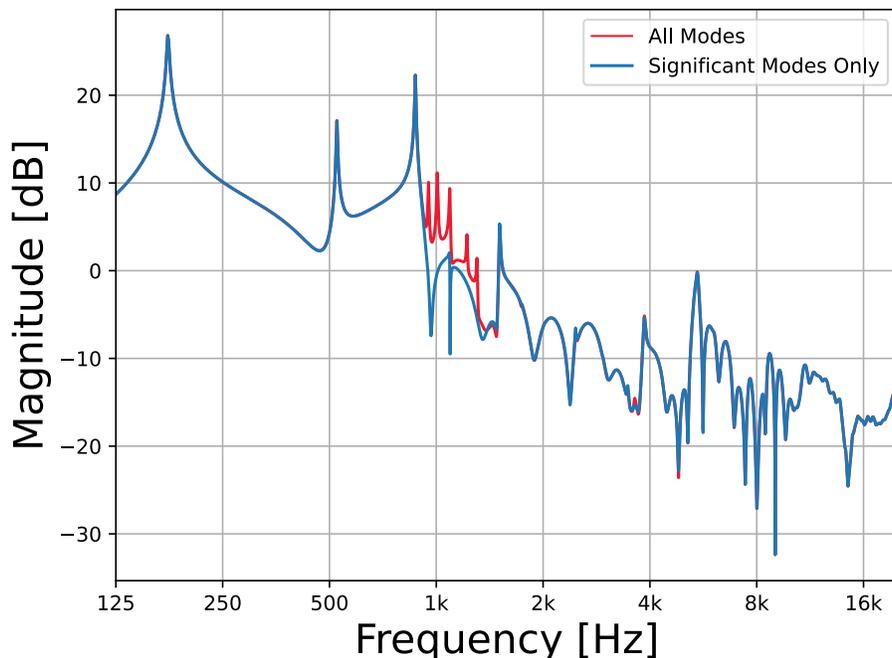


Figure 5.18: Frequency content of the rosewood hybrid radiation model, with[†] and without[‡] removing non-participating modes. Results are shown before applying the damping and frequency corrections of Section 5.3.8.

[†]http://www.music.mcgill.ca/caml/lib/exe/fetch.php?media=publications:RW_HY_NoCorr.wav

[‡]http://www.music.mcgill.ca/caml/lib/exe/fetch.php?media=publications:RW_HY_NPM.wav

extent than the initial segment. For these reasons, comparing only the third segment was deemed the most representative approach.

5.6.1 Aluminum Test

Figure 5.19 plots the resulting fit quality for hybrid radiation models of the aluminum example bar in Section 5.5.1, using different fitting window parameters. The yellow circle in Figure 5.19 represents the combination of fitting window start time and duration predicted by Equations 5.32 and 5.37, and shown in Table 5.5 for the aluminum model.

Looking at Figure 5.19 the large red area shows a plateau, where further increases in fitting window duration or start time do not significantly increase the fit quality of the Prony tail. By contrast, the blue area near the bottom of the plot indicates very poor fit quality. Looking at this area it is evident that a fitting window duration below roughly two milliseconds would be a very poor choice, regardless of the fitting window start time. The yellow circle in Figure 5.19, which represents suggested fitting window parameters, occurs within the red plateau, somewhat near its lower edge. This is an excellent result, as these

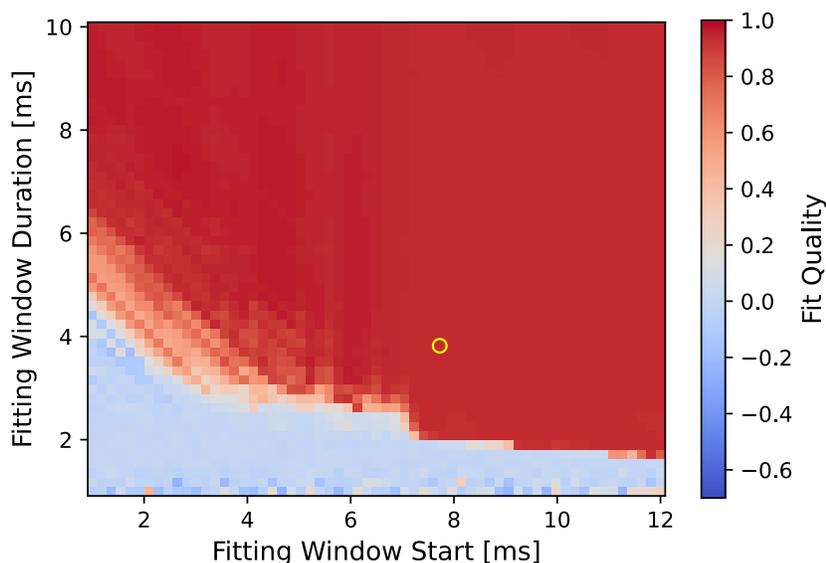


Figure 5.19: Fit quality as a function of fitting window start time and duration for the aluminum bar example model in Section 5.5.1. Colour map indicates the fit quality, defined by Equation 5.39, between time history model results and hybrid radiation models with various fitting window parameters. The yellow circle marks the expected suitable parameter values shown in Table 5.5.

suggested fitting window parameters are intended to produce suitable fit quality while being relatively efficient (i.e. keeping the total of fitting window start time plus duration low, to reduce the time history model’s length).

The transition from the region of poor fit to the red plateau of Figure 5.19 varies from left to right. To the left there is a more gradual slope between the two regions. Lines of equal fit quality in this region are oriented at a -45° angle in the plot, corresponding to positions with equal totals of fitting window start time plus fitting window duration. In this region of the plot, the fitting window start time is too early, occurring before the suggested time. Thus, increasing either start time or duration shows a similar improvement in this region.

On the lower right of Figure 5.19, beyond the suggested fitting window start time, the transition from poor fit to good fit is more abrupt. The sudden transition is oriented nearly parallel to the horizontal axis, indicating that further increases in fitting window start time have little-to-no benefit in this region. Notably, this sudden transition that is parallel to the horizontal axis begins near the suggested fitting window start time.

In terms of fitting window duration, the suggested value represented by the yellow circle in Figure 5.19 is clearly within the quality plateau, as one would hope. There is some room to reduce the fitting window duration while remaining in the plateau zone, which would

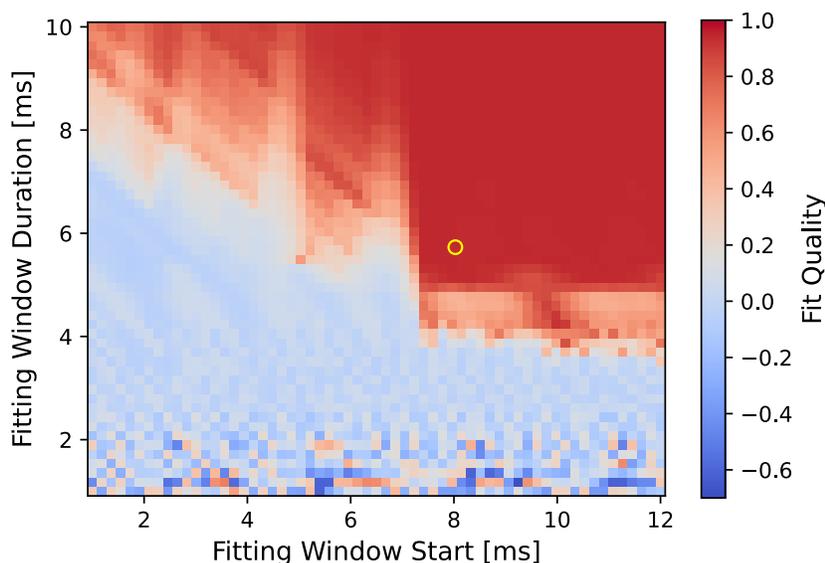


Figure 5.20: Fit quality as a function of fitting window start time and duration for the rosewood bar example model in Section 5.5.2. Colour map indicates the fit quality, defined by Equation 5.39, between time history model results and hybrid radiation models with various fitting window parameters. Range of the colour map is set equal to that of Figure 5.19. The yellow circle marks the expected suitable parameter values shown in Table 5.10.

increase efficiency by reducing the required length of the time history model. However, there is no guarantee that a shorter fitting window duration will always suffice, as is evident in the rosewood example of Figure 5.20.

5.6.2 Rosewood Test

Figure 5.20 plots the parametric test results for fit quality of the rosewood example bar from Section 5.5.2, using different fitting window parameters. The yellow circle in Figure 5.20 represents the combination of fitting window start time and duration predicted by Equations 5.32 and 5.37, and shown in Table 5.10 for the rosewood model.

Figure 5.20 again shows a large red plateau in fit quality, similar to that in Figure 5.19. In contrast to Figure 5.19, however, the plateau in Figure 5.20 has an almost square shape. It is very promising to see the recommended parameters, represented by the yellow circle, near the lower-left corner of the squarish plateau. This implies both the suggested start time and duration of the fitting window are performing very efficiently.

Similar to the aluminum example, once again the left side of the plot shows strips of approximately equal fit quality oriented at a -45° angle. To the right of the suggested

parameters, the transition from the blue region to the red region is not quite so abrupt as was seen in the aluminum example. However, this transition is once again oriented parallel to the horizontal axis, indicating that increasing the fitting window start time beyond the suggest value adds little in the way of fit quality.

The suggested fitting window duration in Figure 5.20 is slightly above a relatively gradual transition to an area of poor fit. In contrast to the aluminum results, where it appeared that window duration could be reduced somewhat, Figure 5.20 shows little room for any reduction in window duration without risking an accompanying loss of fit quality.

While by no means an exhaustive search, the parametric test results in Figures 5.19 and 5.20 both show suitable performance of the fitting window parameter values recommended in Section 5.4.1 when applied to the example models of Section 5.5. This provides some confidence in the recommendations of Section 5.4.1.

Chapter 6

Conclusion and Future Directions

This thesis set out to enable the design and evaluation of idiophone bars in a virtual environment, reducing the need for physical bar prototypes. The methods and models developed have achieved this goal. Using the finite element model definition of Chapter 3, the tuning method developed in Chapter 4 enables the design of a bar cutaway to achieve a desired set of tuning ratios for specified modes of vibration. The developed tuning approach solves the problem of timbre pollution from torsional modes by allowing designers to concurrently tune flexural and torsional modes.

The hybrid sound radiation model developed in Chapter 5 allows analysts and makers to listen to their bar designs in a virtual manner. This model uses the same computational framework as the tuning method in Chapter 4, facilitating evaluation of new bar designs. Both the tuning method and hybrid sound radiation model rely solely on bar geometry and material properties. Depending on the material selected, these properties can be taken from published sources, or measured from stock material. Thus bar designs can be devised and evaluated in a virtual manner, before producing physical prototypes. If sufficient material property information is available in published sources, new materials could feasibly be tested before even purchasing material samples.

The developed hybrid model combines finite element dynamic time history modelling with modal models, leveraging the benefits of each. The time history model enables simulation of complex physical interactions during excitation of the instrument. The modal model provides computationally efficient calculation of instrument response after the complex behaviour has subsided. The hybrid model formulation combines these components in a manner that minimizes the extent of the computationally expensive time history model.

6.1 Contributions

The research undertaken in this thesis has:

1. Confirmed makers' anecdotal reports of torsional modes polluting bar timbre through experimental measurements of marimba bars (Appendix A).
2. Explored the range of bars susceptible to torsional mode problems via simulations (Section 4.3.1).
3. Established a versatile model for defining three-dimensional bar cutaway geometry (Section 3.2.2).
4. Developed methods of automatically identifying mode shapes to enable unsupervised bar tuning algorithms (Section 4.2.3).
5. Produced a method of solving the underdetermined system of nonlinear equations defining the bar tuning problem, in a manner that incorporates analyst input via initial geometry selection (Section 4.2.2).
6. Demonstrated the concurrent tuning of flexural and torsional modes using three-dimensional cutaway geometries (Sections 4.3.2 and 4.3.3).
7. Designed, fabricated and tested a prototype bar to validate the developed modelling and tuning methods (Section 4.5).
8. Developed a hybrid model of sound radiation from idiophones within the same computational framework established for bar tuning (Chapter 5).
9. Investigated parameters controlling the hybrid model (Section 5.6) and suggested values for modelling idiophone bars (Section 5.4).
10. Demonstrated the hybrid radiation model through examples of aluminum and rosewood bars (Section 5.5).

Items 6 and 7 are particularly notable, as they represent the first example of three-dimensional bar cutaway design in the literature on musical acoustics. Item 8 is noteworthy in its application of finite element dynamic time history modelling - a method often dismissed as computationally prohibitive in musical acoustics.

6.2 Future Directions

The computational design and evaluation framework developed for idiophones in this work provides a variety of potential avenues for further exploration. Some interesting possible routes for inquiry include:

- **Three-dimensional cutaway fabrication for wooden bars.** High quality marimba bars are invariably manufactured from rosewood. While examples of fabricated aluminum bars may be found in this work, and the musical acoustics literature (Laukkanen and Worland, 2011; Kirkland and Moradi, 2016; Bestle, 2017; Soares et al., 2021b), examples of fabricated wooden bars have thus far been less common, and moderately successful (Mingming, 2011). Fabricating sophisticated bar cutaway designs from a natural material like wood is an interesting and challenging area of potential future research.
- **Exploration of preferable damping behaviour and new material design.** Marimba bars made from synthetic materials are currently available on the market (e.g. Musser’s *Kelon*, or Yamaha’s *Acoustalon*). These instruments offer some advantages, such as being less affected by changes in humidity. However, the damping properties of these synthetics produce a sound distinctly different from rosewood. The methods developed in this thesis offer the opportunity to evaluate new potential materials for bar manufacture. One could even use listening tests to find desirable material properties for idiophone bars and then attempt to design a material with the desired properties.
- **Addition of a resonator to the radiation model.** As shown by Rucz et al. (2021), the addition of a resonator to the hybrid sound radiation model would enable listening to virtual bar designs in an arrangement more reflective of actual playing conditions. Incorporating resonators in a manner that avoids meshing their full volume would do much to aid computational efficiency.
- **Alternative approaches to mode selection.** While mode selection using normalized effective modal mass (see Section 5.3.4) was largely successful, it is by no means perfect. Exploring alternative methods of mode selection could result in a more universally reliable approach. Consideration of sound radiation strength, or spectral content of the initial finite element time domain model results, may help inform such methods.
- **Realistic boundary conditions for marimba and vibraphone models.** As discussed in Section 5.3.2, the support conditions for a marimba bar positioned on an

instrument are challenging to define in a finite element model. The support cables passing through a bar offer varying degrees of resistance depending on the distance and direction of travel. Further investigation of boundary conditions could enable sound radiation from bars designed with free-free boundary conditions to be modelled without altering their modal behaviour. Defining nonlinear nodal spring supports in the finite element model is one potential approach. Experimental measurements of bar displacements on an existing instrument, using varying levels of applied force, could facilitate defining these nonlinear spring supports.

References

- Abaqus. Abaqus unified FEA - SIMULIA by Dassault Systèmes. <https://www.3ds.com/products-services/simulia/products/abaqus/>. Accessed: 2020 Apr 28.
- M. Adcock. The ancient stone instruments of Vietnam. *Time and Mind*, 10(1):23–37, 2017.
- S. Adhikari. Damping modelling using generalized proportional damping. *Journal of Sound and Vibration*, 293(1):156–170, 2006.
- A. Akay, M. Bengisu, and M. Latcha. Transient acoustic radiation from impacted beam-like structures. *Journal of Sound and Vibration*, 91(1):135–145, 1983.
- A. Ali and C. Rajakumar. *The Boundary Element Method - Applications in Sound and Vibration*. Taylor & Francis Group, Boca Raton, FL, 2004.
- R. J. Allemang. The modal assurance criterion - Twenty years of use and abuse. *Sound and Vibration*, 37(8):14–21, 2003.
- M. Aramaki, H. Baillères, L. Brancheriau, R. Kronland-Martinet, and S. Ystad. Sound quality assessment of wood for xylophone bars. *The Journal of the Acoustical Society of America*, 121(4):2407–2420, 2007.
- R. J. Astley. Transient wave envelope elements for wave problems. *Journal of Sound and Vibration*, 192(1):245–261, 1996.
- N. Atalla and F. Sgard. *Finite Element and Boundary Methods in Structural Acoustics and Vibration*. CRC Press, Boca Raton, FL, 2015.
- P. Avitabile. *Modal Testing: A Practitioner’s Guide*. John Wiley & Sons, Chichester, UK, 2018.
- Battelle Memorial Institute. *Metallic Materials Properties Development and Standardization (MMPDS-15)*. Battelle Memorial Institute, Columbus, OH, 2020.

- A. Bayliss and E. Turkel. Radiation boundary conditions for wave-like equations. *Communications on Pure and Applied Mathematics*, 33(6):707–725, 1980.
- D. Beaton. Implementation and Evaluation of a Coupled Thermal-Structural Analysis Module for Laminated Composites in an Open-Source Finite Element Code. Master’s thesis, The University of British Columbia, Vancouver, BC, 2010.
- D. Beaton and G. Scavone. Measurement-based comparison of marimba bar modal behaviour. In *Proceedings of the International Symposium on Musical Acoustics 2019*, pages 72–77, Detmold, DE, 2019a.
- D. Beaton and G. Scavone. Optimization of marimba bar geometry by 3D finite element analysis. In *Proceedings of the International Symposium on Musical Acoustics 2019*, pages 402–407, Detmold, DE, 2019b.
- D. Beaton and G. Scavone. Three-dimensional tuning of idiophone bar modes via finite element analysis. *The Journal of the Acoustical Society of America*, 149(6):3758–3768, 2021a.
- D. Beaton and G. Scavone. Tuning idiophone bar torsional modes with three-dimensional cutaway geometries. *The Journal of the Acoustical Society of America*, 149(4):A69, 2021b.
- D. Beaton and G. Scavone. Experimental measurements of a prototype vibraphone bar with three-dimensional cutaway geometry. *JASA Express Letters*, 2(8):083201, 2022.
- A. Ben-Israel. The Moore of the Moore-Penrose inverse. *The Electronic Journal of Linear Algebra*, 9:150–157, 2002.
- A. Ben-Israel and T. N. E. Greville. *Generalized Inverses: Theory and Applications*. Number 15 in CMS books in mathematics. Springer, New York, NY, second edition, 2003.
- J.-P. Berenger. A perfectly matched layer for the absorption of electromagnetic waves. *Journal of Computational Physics*, 114(2):185–200, 1994.
- A. Bermúdez, L. Hervella-Nieto, A. Prieto, and R. Rodríguez. Perfectly Matched Layers. In S. Marburg and B. Nolte, editors, *Computational Acoustics of Noise Propagation in Fluids - Finite and Boundary Element Methods*, pages 167–196. Springer-Verlag, Berlin, DE, 2008.
- P. Bestle. *Zur Klangoptimierung von Idiophonen mittels Schallsynthesierung (For sound optimization of idiophones by means of sound synthesis)*. Number 52 in Schriften aus

- dem Institut für Technische und Numerische Mechanik der Universität Stuttgart. Shaker Verlag, Aachen, DE, 2017.
- P. Bestle, P. Eberhard, and M. Hanss. Musical instruments - Sound synthesis of virtual idiophones. *Journal of Sound and Vibration*, 395:187–200, 2017.
- P. Bettess. Infinite elements. *International Journal for Numerical Methods in Engineering*, 11(1):53–64, 1977.
- P. Bettess and O. C. Zienkiewicz. Diffraction and refraction of surface waves using finite and infinite elements. *International Journal for Numerical Methods in Engineering*, 11(8):1271–1290, 1977.
- S. Bilbao. Percussion synthesis based on models of nonlinear shell vibration. *IEEE Transactions on Audio, Speech, and Language Processing*, 18(4):872–880, 2010.
- I. Bork. Measuring the acoustical properties of mallets. *Applied Acoustics*, 30(2-3):207–218, 1990.
- I. Bork and J. Meyer. On the tonal evaluation of xylophones. *Percussive Notes Research Edition*, 23(6):48–57, 1985.
- I. Bork, A. Chaigne, L.-C. Trebuchet, M. Kosfelder, and D. Pillot. Comparison between modal analysis and finite element modelling of a marimba bar. *Acta Acustica united with Acustica*, 85(2):258–266, 1999.
- I. Brémaud. Acoustical properties of wood in string instruments soundboards and tuned idiophones: Biological and cultural diversity. *The Journal of the Acoustical Society of America*, 131(1):807–818, 2012.
- J. Bretos, C. Santamaría, and J. Moral. Frequencies, input admittances and bandwidths of the natural bending eigenmodes in xylophone bars. *Journal of Sound and Vibration*, 203(1):1–9, 1997.
- J. Bretos, C. Santamaría, and J. A. Moral. Finite element analysis and experimental measurements of natural eigenmodes and random responses of wooden bars used in musical instruments. *Applied Acoustics*, 56(3):141–156, 1999.
- D. L. Brown, R. J. Allemang, R. Zimmerman, and M. Mergeay. Parameter estimation techniques for modal analysis. *SAE Transactions*, 88(1):828–846, 1979. 1979 Automotive Engineering Congress and Exposition Technical Paper 790221.

- D. S. Burnett. A three-dimensional acoustic infinite element based on a prolate spheroidal multipole expansion. *The Journal of the Acoustical Society of America*, 96(5):2798–2816, 1994.
- M. Carvalho, V. Debut, and J. Antunes. Development of a non-destructive bell-tuning technique through optimized structural modifications. *Procedia Engineering*, 199:330–335, 2017.
- A. Chaigne and A. Askenfelt. Numerical simulations of piano strings. I. A physical model for a struck string using finite difference methods. *The Journal of the Acoustical Society of America*, 95(2):1112–1118, 1994.
- A. Chaigne and V. Doutaut. Numerical simulations of xylophones. I. Time-domain modeling of the vibrating bars. *The Journal of the Acoustical Society of America*, 101(1):539–557, 1997.
- P. Chobeau and L. Savioja. Memory cost of absorbing conditions for the finite-difference time-domain method. *The Journal of the Acoustical Society of America*, 140(1):EL119–EL124, 2016.
- T. F. Coleman and Y. Li. An interior trust region approach for nonlinear minimization subject to bounds. *SIAM Journal on Optimization*, 6(2):418–445, 1996.
- R. D. Cook, D. S. Malkus, M. E. Plesha, and R. J. Witt. *Concepts and Applications of Finite Element Analysis*. John Wiley & Sons, Hoboken, NJ, fourth edition, 2001.
- R. Courant. Variational methods for the solution of problems of equilibrium and vibrations. *Bulletin of the American Mathematical Society*, 49(1):1–23, 1943.
- G. C. F. M. R. de Prony. Essai expérimental et analytique: sur les lois de la dilatabilité de fluides élastique et sur celles de la force expansive de la vapeur de l’alkool, à différentes températures. (Experimental and analytical test: on the laws of the dilatibility of elastic fluids and on those of force expansion of water vapor and alcohol vapor, at different temperatures). *Journal de l’École Polytechnique*, 1(2):24–76, 1795.
- V. Debut, M. Carvalho, E. Figueiredo, J. Antunes, and R. Silva. The sound of bronze: Virtual resurrection of a broken medieval bell. *Journal of Cultural Heritage*, 19:544–554, 2016.
- V. Debut, M. Carvalho, F. Soares, and J. Antunes. Reverse Engineering Techniques for Investigating the Vibro-Acoustics of Historical Bells. In T. Fakhfakh, C. Karra, S. Bouaziz,

- F. Chaari, and M. Haddar, editors, *Advances in Acoustics and Vibration II*, volume 13 of *Applied Condition Monitoring*, pages 218–226. Springer Nature, Cham, CH, 2019.
- G. Dhondt. CalculiX: A Three-Dimensional Structural Finite Element Program. <http://www.calculix.de>. Accessed: 2022 May 30.
- M. Dinachandra and S. Raju. Isogeometric analysis for acoustic fluid-structure interaction problems. *International Journal of Mechanical Sciences*, 131-132:8–25, 2017.
- M. Dooley. Rosewood vs. Synthetic vs. Padauk Marimba & Xylophone Bars. <https://www.lonestarpercussion.com/Rosewood-vs-Synthetic-vs-Padauk-Marimba-Bars.html>, 2015. Accessed: 2022 February 1.
- Électricité de France. Finite Element Code_Aster, Analysis of Structures and Thermomechanics for Studies and Research. <https://www.code-aster.org>. Accessed: 2022 May 30.
- B. Engquist and A. Majda. Radiation boundary conditions for acoustic and elastic wave calculations. *Communications on Pure and Applied Mathematics*, 32(3):313–357, 1979.
- G. Farin. Triangular Bernstein-Bézier patches. *Computer Aided Geometric Design*, 3(2):83–127, 1986.
- D. Givoli. High-order local non-reflecting boundary conditions: a review. *Wave Motion*, 39(4):319–326, 2004.
- D. Givoli. Computational Absorbing Boundaries. In S. Marburg and B. Nolte, editors, *Computational Acoustics of Noise Propagation in Fluids - Finite and Boundary Element Methods*, pages 145–166. Springer, Berlin, DE, 2008.
- D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning*. Addison-Wesley Publishing Company, Reading, MA, 1989.
- M. N. Guddati and K.-W. Lim. Continued fraction absorbing boundary conditions for convex polygonal domains. *International Journal for Numerical Methods in Engineering*, 66(6):949–977, 2006.
- W. Hare, J. Nutini, and S. Tesfamariam. A survey of non-gradient optimization methods in structural engineering. *Advances in Engineering Software*, 59:19–28, 2013.
- L. Henrique and J. Antunes. Optimal design and physical modelling of mallet percussion instruments. *Acta Acustica united with Acustica*, 89(6):948–963, 2003.

- L. Henrique, J. Antunes, and J. de Carvalho. Design of musical instruments using global optimization techniques. In *Proceedings of the International Symposium on Musical Acoustics*, pages 545–548, Perugia, IT, 2001.
- H. M. Hilber. *Analysis and Design of Numerical Integration Methods in Structural Dynamics*. PhD Thesis, University of California at Berkeley, Berkeley, CA, 1976.
- H. M. Hilber, T. J. R. Hughes, and R. L. Taylor. Improved numerical dissipation for time integration algorithms in structural dynamics. *Earthquake Engineering & Structural Dynamics*, 5(3):283–292, 1977.
- D. Holz. Acoustically important properties of xylophone-bar materials: Can tropical woods be replaced by European species? *Acta Acustica united with Acustica*, 82(6):878–884, 1996.
- A. Hrennikoff. *Plane Stress and Bending of Plates by Method of Articulated Framework*. PhD Thesis, Massachusetts Institute of Technology, Cambridge, MA, 1940.
- M. Israeli and S. A. Orszag. Approximation of radiation boundary conditions. *Journal of Computational Physics*, 41(1):115–135, 1981.
- B. Kaltenbacher, M. Kaltenbacher, and I. Sim. A modified and stable version of a perfectly matched layer technique for the 3-d second order wave equation in time domain with an application to aeroacoustics. *Journal of Computational Physics*, 235:407–422, 2013.
- C. Kanzow, N. Yamashita, and M. Fukushima. Levenberg-Marquardt methods with strong local convergence properties for solving nonlinear equations with convex constraints. *Journal of Computational and Applied Mathematics*, 172(2):375–397, 2004.
- R. Kicinger, T. Arciszewski, and K. D. Jong. Evolutionary computation and structural design: A survey of the state-of-the-art. *Computers & Structures*, 83(23-24):1943–1978, 2005.
- S. Kimmich, R. Reitingner, and E. Ramm. Integration of different numerical techniques in shape optimization. *Structural Optimization*, 4(3-4):149–155, 1992.
- W. B. Kirkland. *Topographical Optimization of Structures for use in Musical Instruments and Other Applications*. Master’s thesis, The University of Alabama at Birmingham, Birmingham, AL, 2014.

- W. B. Kirkland and L. Moradi. Topographical optimization of structures for use in vibraphones and other applications. In *Proceedings of the Society for Design and Process Science 21st International Conference*, pages 69–77, Orlando, FL, 2016.
- R. Kolman, S. Sorokin, B. Bastl, J. Kopačka, and J. Plešek. Isogeometric analysis of free vibration of simple shaped elastic samples. *The Journal of the Acoustical Society of America*, 137(4):2089–2100, 2015.
- E. M. Laukkanen and R. Worland. Acoustical effect of progressive undercutting of percussive aluminum bars. *Proceedings of Meetings on Acoustics*, 14(1):025003, 2011.
- K. A. Legge and J. Petrolito. Designing idiophones with tuned overtones. *Acoustics Australia*, 35(2):47–50, 2007.
- L. Liu, C. Song, W. Gao, and C. Birk. Scaled Boundary Finite Element Method for Exterior Acoustic-Structure Interaction Analyses. In H. Hao and C. Zhang, editors, *Mechanics of Structures and Materials: Advancements and Challenges*, volume 1, pages 299–304. Taylor & Francis Group, London, UK, 2016.
- L. Liu, J. Zhang, C. Song, C. Birk, and W. Gao. An automatic approach for the acoustic analysis of three-dimensional bounded and unbounded domains by scaled boundary finite element method. *International Journal of Mechanical Sciences*, 151:563–581, 2019a.
- L. Liu, J. Zhang, C. Song, C. Birk, A. A. Saputra, and W. Gao. Automatic three-dimensional acoustic-structure interaction analysis using the scaled boundary finite element method. *Journal of Computational Physics*, 395:432–460, 2019b.
- E. Lund. *Finite Element Based Design Sensitivity Analysis and Optimization*. PhD Thesis, Aalborg University, Aalborg, DK, 1994.
- F. K. MacCallum. *The Book of the Marimba*. Carlton Press, New York, NY, 1969.
- M. Macconi, B. Morini, and M. Porcelli. A Gauss-Newton method for solving bound-constrained underdetermined nonlinear systems. *Optimization Methods & Software*, 24(2):219–235, 2009.
- N. D. Manh, A. Evgrafov, A. R. Gersborg, and J. Gravesen. Isogeometric shape optimization of vibrating membranes. *Computer Methods in Applied Mechanics and Engineering*, 200(13-16):1343–1353, 2011.
- S. Marburg. Six boundary elements per wavelength: Is that enough? *Journal of Computational Acoustics*, 10(1):25–51, 2002.

- S. Marburg. Discretization Requirements: How many Elements per Wavelength are Necessary? In S. Marburg and B. Nolte, editors, *Computational Acoustics of Noise Propagation in Fluids - Finite and Boundary Element Methods*, pages 309–332. Springer-Verlag, Berlin, DE, 2008.
- S. L. Marple. *Digital Spectral Analysis: With Applications*. Prentice-Hall Signal Processing Series. Prentice-Hall, Englewood Cliffs, NJ, 1987.
- S. McAdams, A. Chaigne, and V. Roussarie. The psychomechanics of simulated sound sources: Material properties of impacted bars. *The Journal of the Acoustical Society of America*, 115(3):1306–1320, 2004.
- N. McLachlan. Finite element analysis and gong acoustics. *Acoustics Australia*, 25(3):103–107, 1997.
- N. McLachlan, B. K. Nigjeh, and A. Hasell. The design of bells with harmonic overtones. *The Journal of the Acoustical Society of America*, 114(1):505–511, 2003.
- N. McLachlan, R. Adams, and C. Burvill. Tuning natural modes of vibration by prestress in the design of a harmonic gong. *The Journal of the Acoustical Society of America*, 131(1):926–934, 2012.
- Z. Mingming. Automatic Multi-Modal Tuning of Idiophone Bars. Master’s thesis, Curtin University, Perth, AU, 2011.
- L. Moheit and S. Marburg. Infinite elements and their influence on normal and radiation modes in exterior acoustics. *Journal of Computational Acoustics*, 25(4):1650020, 2017.
- J. L. Moore. *Acoustics of Bar Percussion Instruments*. PhD Thesis, The Ohio State University, Columbus, OH, 1970.
- D. J. Munk, G. A. Vio, and G. P. Steven. Topology and shape optimization methods using evolutionary algorithms: a review. *Structural and Multidisciplinary Optimization*, 52(3):613–631, 2015.
- N. M. Newmark. A method of computation for structural dynamics. *Journal of the Engineering Mechanics Division*, 85(3):67–94, 1959.
- C. T. Nguyen and J. L. Tassoulas. Application of reciprocal absorbing boundary condition to transient analysis of acoustic wave propagation. *Computer Methods in Applied Mechanics and Engineering*, 329:55–74, 2018.

- N. Nishimura. Fast multipole accelerated boundary integral equation methods. *Applied Mechanics Reviews*, 55(4):299–324, 2002.
- J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer Science+Business Media, New York, NY, second edition, 2006.
- F. Orduña-Bustamante. Nonuniform beams with harmonically related overtones for use in percussion instruments. *The Journal of the Acoustical Society of America*, 90(6):2935–2941, 1991.
- M. Özakça and M. Göğüş. Structural analysis and optimization of bells using finite elements. *Journal of New Music Research*, 33(1):61–69, 2004.
- R. Penrose. A generalized inverse for matrices. *Mathematical Proceedings of the Cambridge Philosophical Society*, 51(3):406–413, 1955.
- R. Perrin, T. Charnley, and J. dePont. Normal modes of the modern English church bell. *Journal of Sound and Vibration*, 90(1):29–49, 1983.
- R. Perrin, T. Charnley, and G. Swallowe. On the tuning of church and carillon bells. *Applied Acoustics*, 46(1):83–101, 1995.
- R. Perrin, L. Chalmers, D. P. Elford, G. M. Swallowe, and T. R. Moore. Normal modes of the Indian elephant bell. *The Journal of the Acoustical Society of America*, 131(3):2288–2294, 2012.
- R. Perrin, D. P. Elford, L. Chalmers, G. M. Swallowe, T. R. Moore, S. Hamdan, and B. J. Halkon. Normal modes of a small gamelan gong. *The Journal of the Acoustical Society of America*, 136(4):1942–1950, 2014.
- H. Peters, N. Kessissoglou, and S. Marburg. Modal decomposition of exterior acoustic-structure interaction. *The Journal of the Acoustical Society of America*, 133(5):2668–2677, 2013.
- J. Petrolito and K. A. Legge. Optimal undercuts for the tuning of percussive beams. *The Journal of the Acoustical Society of America*, 102(4):2432–2437, 1997.
- L. Piegl and W. Tiller. *The NURBS Book*. Monographs in Visual Communications. Springer-Verlag, Berlin, DE, 1995.

- D. Radojevic and K. F. Kirkwood. The new old bridge: Renovation of the Macdonald suspension bridge. *Structural Engineering International*, 27(1):32–37, 2017.
- J. W. S. Rayleigh. *Theory of Sound*. Dover Publications, New York, NY, Dover reprint edition, 1945.
- A. Richter. Numerical investigations of the gas flow inside the bassoon. *Journal of Fluids Engineering*, 134(4):041104, 2012.
- P. Rucz, M. A. Ulveczki, J. Angster, and A. Miklós. Simulation of mallet percussion instruments by a coupled modal vibroacoustic finite element model. *The Journal of the Acoustical Society of America*, 149(5):3200–3212, 2021.
- C. Shih, Y. Tsuei, R. Allemang, and D. Brown. Complex Mode Indication Function and its applications to spatial domain parameter estimation. *Mechanical Systems and Signal Processing*, 2(4):367–377, 1988.
- F. Soares, J. Antunes, and V. Debut. Multi-modal tuning of vibrating bars with simplified undercuts using an evolutionary optimization algorithm. *Applied Acoustics*, 173:107704, 2021a.
- F. Soares, J. Antunes, and V. Debut. Tuning of bending and torsional modes of bars used in mallet percussion instruments. *The Journal of the Acoustical Society of America*, 150(4):2757–2769, 2021b.
- C. Song. The scaled boundary finite element method in structural dynamics. *International Journal for Numerical Methods in Engineering*, 77(8):1139–1171, 2009.
- C. Song. *The Scaled Boundary Finite Element Method: Introduction to Theory and Implementation*. John Wiley & Sons, Hoboken, NJ, 2018.
- L. H. Stevens. The acoustic complexity of keyboard percussion instruments. Part I: What we know and what we wish we knew. *The Journal of the Acoustical Society of America*, 138(3):1912, 2015.
- H. Tahvanainen, H. Matsuda, and R. Shinoda. Numerical simulation of the acoustic guitar for virtual prototyping. In *Proceedings of the International Symposium on Music Acoustics 2019*, pages 380–386, Detmold, DE, 2019.
- The New York Times. Stone-Age ‘Marimba’ is Discovered in Vietnam: Rock Plates Pitched in Java Scale May Be 5,000 Years Old. *The New York Times*, 1954.

- M. J. Turner, R. W. Clough, H. C. Martin, and L. J. Topp. Stiffness and deflection analysis of complex structures. *Journal of the Aeronautical Sciences*, 23(9):805–823, 1956.
- K. R. Unnikrishnan, I. R. Praveen Krishna, and C. O. Arun. Free Vibration Analysis of Prestressed Membrane Using Element Free Galerkin Method. In S. C. Satapathy, K. S. Raju, K. Molugaram, A. Krishnaiah, and G. A. Tsihrintzis, editors, *International Conference on Emerging Trends in Engineering (ICETE)*, volume 2 of *Learning and Analytics in Intelligent Systems*, pages 542–550. Springer Nature Switzerland, Cham, CH, 2020.
- A. Vermeil de Conchard, H. Mao, and R. Rumpfer. A perfectly matched layer formulation adapted for fast frequency sweeps of exterior acoustics finite element models. *Journal of Computational Physics*, 398:108878, 2019.
- S.-Y. Wang, Y. Sun, and R. Gallagher. Sensitivity analysis in shape optimization of continuum structures. *Computers & Structures*, 20(5):855–867, 1985.
- J. P. Wolf and C. Song. Consistent infinitesimal finite-element cell method: in-plane motion. *Computer Methods in Applied Mechanics and Engineering*, 123(1-4):355–370, 1995.
- H. Wu, W. Ye, and W. Jiang. Isogeometric finite element analysis of interior acoustic problems. *Applied Acoustics*, 100:63–73, 2015.
- S. Wu, Y. Xiang, J. Yao, and S. Wang. An element-free Galerkin coupled with improved infinite element method for exterior acoustic problem. *Journal of Theoretical and Computational Acoustics*, 27(02):1850021, 2019.
- J. Yang, X. Zhang, G. R. Liu, and W. Zhang. A compact perfectly matched layer algorithm for acoustic simulations in the time domain with smoothed particle hydrodynamic method. *The Journal of the Acoustical Society of America*, 145(1):204–214, 2019.
- Q. Yang and X. Peng. An exact method for calculating the eigenvector sensitivities. *Applied Sciences*, 10(7):2577, 2020.
- Y. Yu and B. M. Kwak. Design sensitivity analysis of acoustical damping and its application to design of musical bells. *Structural and Multidisciplinary Optimization*, 44(3):421–430, 2011.
- H. Zhang, Z. Yan, C. Yang, and Y. Shi. Study on excitation of the two-tone acoustic characteristic of the chime bell of Marquis Yi of Zeng by finite element method. *Chinese Science Bulletin*, 52(16):2167–2173, 2007.

O. C. Zienkiewicz, R. L. Taylor, and J. Z. Zhu. *The Finite Element Method: Its Basis and Fundamentals*. Elsevier, Butterworth-Heinemann, Amsterdam, seventh edition, 2013.

Appendix A

Bar Measurements

This appendix outlines experimental modal analyses performed on numerous bars over the course of this work. It is subdivided into two main components: early measurements performed on marimba bars, and later measurements performed on aluminum bars, along with reanalysis of marimba bar data.

A.1 Marimba Bar Measurements

This section expands on measurement results reported by [Beaton and Scavone \(2019a\)](#). It contains further results that were presented in a poster at the International Symposium on Musical Acoustics (ISMA) 2019, but did not appear in the corresponding conference paper ([Beaton and Scavone, 2019a](#)).

A.1.1 Motivation

Near the outset of this work experimental modal analysis was performed on a selection of marimba bars from different instruments. These measurements assessed how modal tuning ratios varied between different instruments and different manufacturers. Of primary interest were the questions:

1. How accurately do manufacturer's achieve the typical 1:4:10 tuning ratio in their marimba bars?
2. How do bar tuning ratios vary between different instruments and different manufacturers?
3. How do the frequency ratios of untuned modes compare with those of tuned modes?

4. Do the frequency ratios of modes V3 and T2 come close to each other, as reported anecdotally by makers and expected by preliminary simulations?

Question 4 was of particular interest in these measurements. Early simulation results (see Figure 4.4) over the lower two octaves of a marimba indicated that modes V3 and T2 could have similar frequencies from roughly note E3 through note B3. These results matched comments made by makers complaining of timbral problems likely caused by torsional modes. These bar measurements were thus intended to confirm whether untuned mode T2 could have similar frequency to tuned mode V3 within this range of notes.

A.1.2 Scope

To facilitate a fair comparison between different instruments and makers, all bars measured were of the same musical note. The note F \sharp 3 was selected as it lies within the range E3 to B3. Coincidentally, this note represents a position at which bar widths were observed to change on numerous instruments (with note F3 and below having larger bar width than note F \sharp 3 and above). From a practical perspective, F \sharp 3 was also the most readily accessible bar in the range E3 to B3, as accidentals are more easily removed from a marimba.

Measurements were performed on as many bars of the note F \sharp 3 as could be accessed at the time. Ultimately, 15 bars were measured, of which two were synthetic and the remainder made of rosewood. Six manufacturers were represented in the measurement sample. Two bars had been retuned at some point after their manufacture; a third bar appears to have been retuned, but this is not known with certainty.

A.1.3 Method

Measurements

Figure A.1 shows the measurement setup used to perform experimental modal analysis. Bars were supported on small foam blocks positioned below the holes where cables would support the bars on an instrument (these holes are typically located at the nodes of the bar's fundamental mode). The Polytec laser Doppler vibrometer (LDV) was positioned below the bar, with its laser aimed upward, hitting the bar's bottom surface near a corner.

With the LDV aimed at the bar's bottom surface, the Brüel & Kjær (B&K) impact hammer was positioned to strike the bar on its playing surface. From there, the impact hammer was used to excite the bar over a grid of strike locations with the LDV recording the bar's response to each strike. A minimum of four measurements were recorded for averaging at each strike position.

Data processing

Measured data from the impact hammer and LDV was fed into a MATLAB script for processing. Modal frequencies were determined using the *Complex Mode Indication Function* (CMIF) (Shih et al., 1988; Avitabile, 2018). Results using this method are shown in the following section for all measured marimba bars.

A.1.4 Results

Figure A.2 displays measured frequency ratio results for the 15 bars considered, along with geometric data and images of each bar. This figure was taken from the poster presented at ISMA 2019 Beaton and Scavone (2019a). Bar images are shown approximately to scale, providing visual indication of the variability in bar geometry. Mode shape images were produced by MATLAB, and show only vertical displacements.

A.1.5 Discussion

Immediately noticeable in Figure A.2 is the variability in bar size. Bar lengths ranged from 373 mm up to 456 mm. Bar widths, by contrast, were quite consistent. All bars were between 56 and 58 mm wide, save for two exceptions around 50 mm. End thicknesses were similarly consistent, with all rosewood bars between 22 and 25 mm thick at their ends.

Looking at the frequency ratios, mode V2 is consistently the standard ratio of 4 times the fundamental, with hardly any deviation observed in any bar. More variability is apparent in the tuning of mode V3, with some bars noticeably off the typical target of 10. The two largest deviations, however, belong to bar 8, a synthetic bar, and bar 9, which is by far the oldest bar, being estimated as 60 years old or more.

Of the untuned modes, mode T1 seems to hover around a ratio of 2 for all of the bars, with relatively small variability compared to the other untuned modes. Mode L1 shows a much wider range of observed frequency ratios, ranging from less than 5 to more than 7, though the two highest values occur in the synthetic bars. Most importantly, the frequency ratio of mode T2 was indeed observed to be close to that of mode V3. In some cases, such as bars 1, 2, 5 and 10, the two modes had very similar ratios. In the extreme case, the two modes had frequencies so similar in were so close in bar 15 that CMIF identified them as a single degenerate mode.

Thus the answer to Question 4 is a resounding “yes”: the frequency ratios of modes T2 and V3 can certainly come close to each other for these bars of note F \sharp 3. Similar behaviour can be expected from bars positioned nearby on the keyboard.

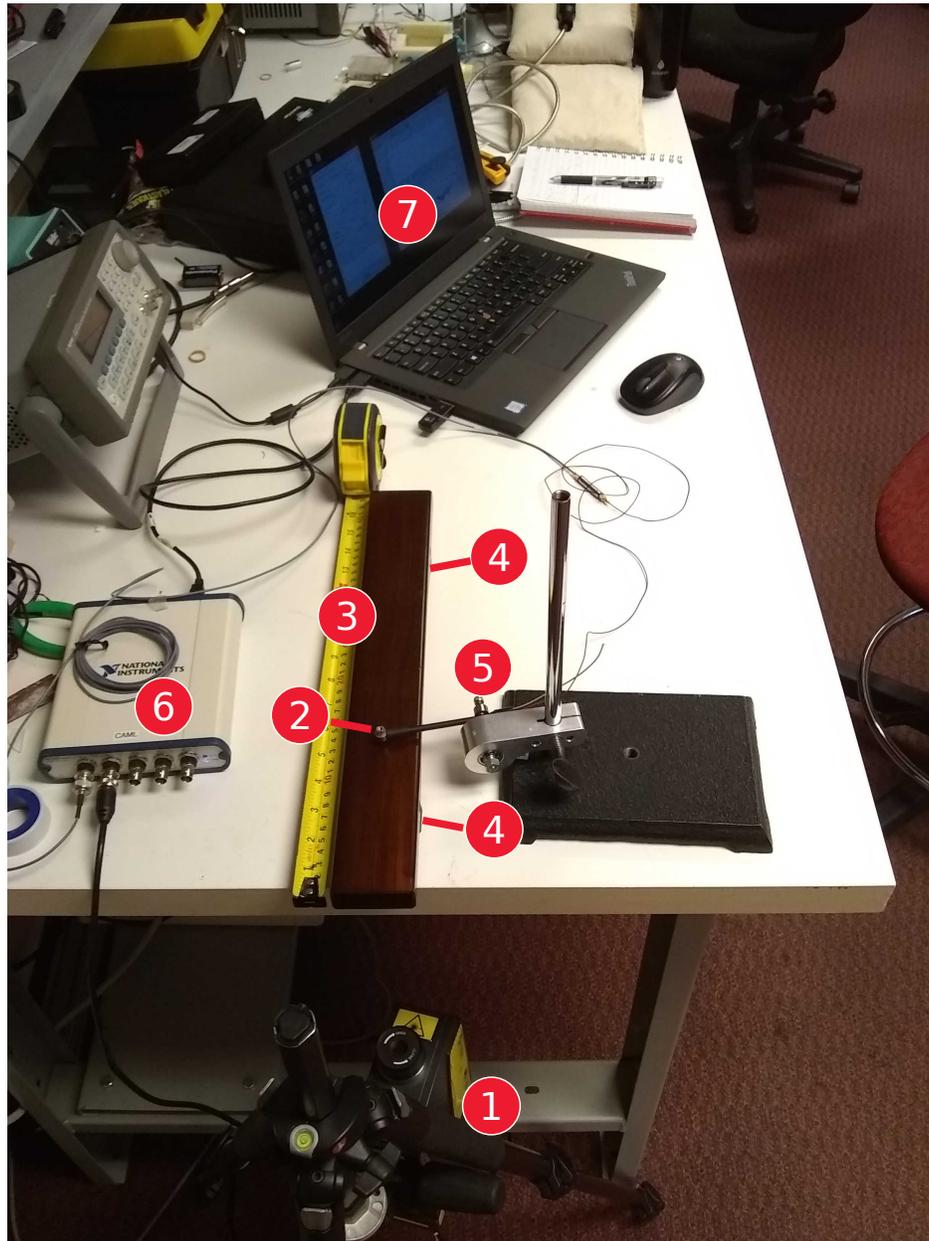
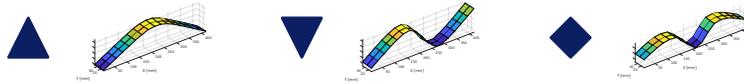


Figure A.1: Measurement setup for experimental modal analysis of marimba bars. Components: 1) Polytec PDV-100 Laser Doppler Vibrometer, 2) Brüel & Kjær Type 8203 Impact Hammer, 3) test bar, 4) foam supports, 5) pendulum, 6) National Instruments USB-4431 Signal Acquisition Board, 7) laptop with MATLAB digital acquisition platform. Figure taken from the poster presented at ISMA 2019.

Measured Mode Shape Legend

(Vertical displacement component only)

Tuned Vertical Modes



Target Values



Lateral and Torsional Modes



Degenerate Mode

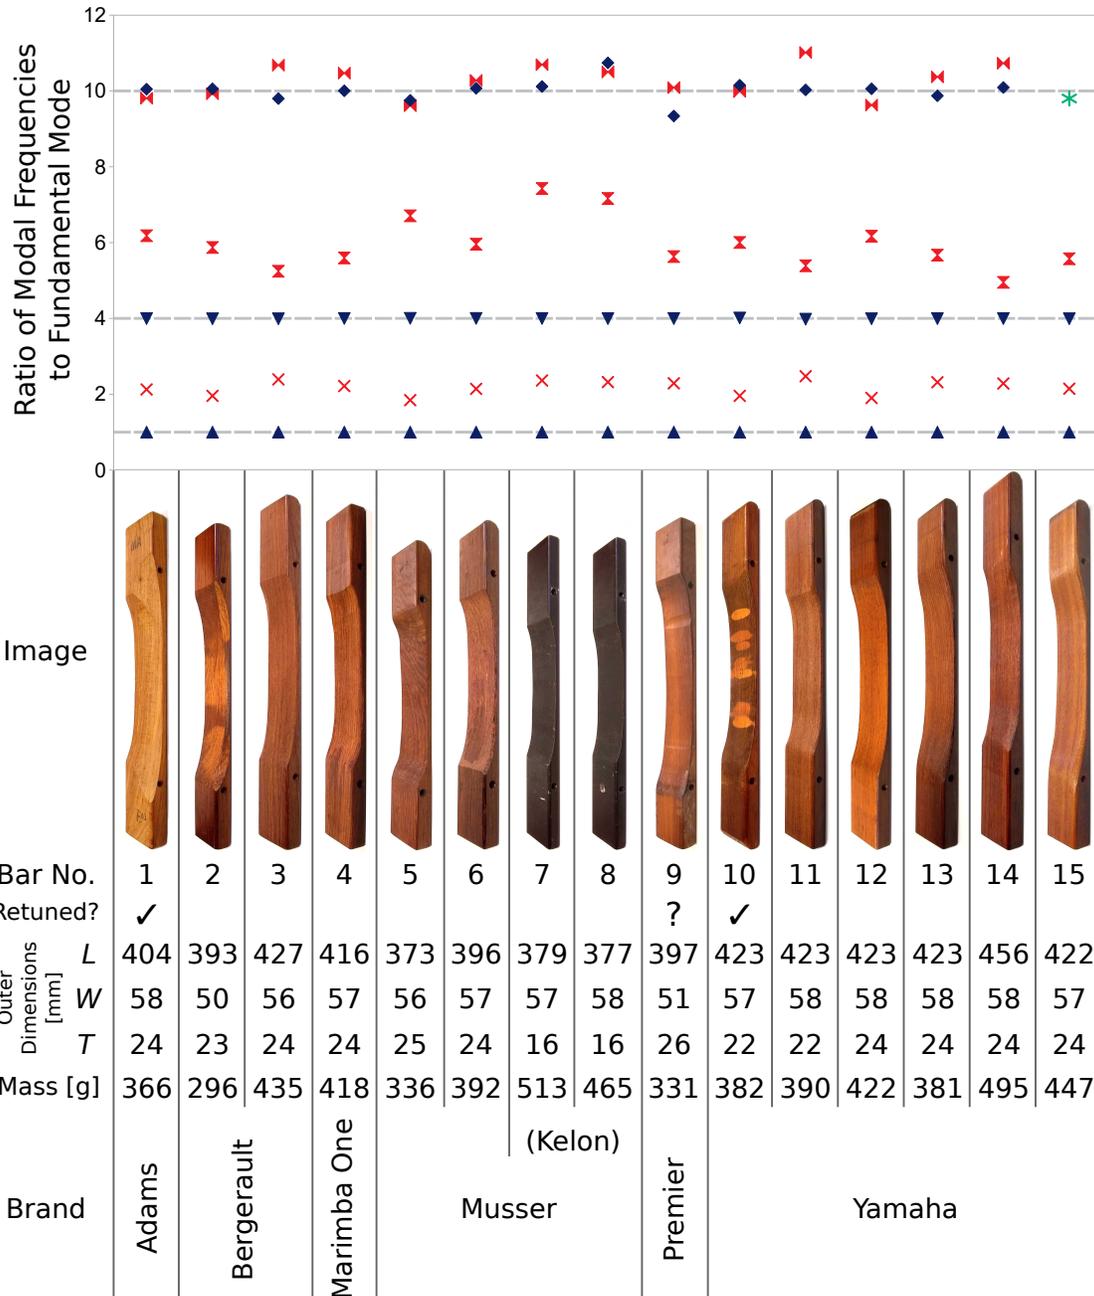
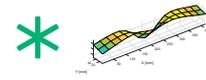


Figure A.2: Measurement results for the 15 measured marimba bars. Bar images are shown approximately to scale. Figure taken from the poster presented at ISMA 2019.

A.2 Additional Bar Measurements and Reanalysis

This section covers bar measurements made in this work, unrelated to those reported by [Beaton and Scavone \(2019a\)](#). It also discusses further analysis of measurement data from Section [A.1](#), performed after ISMA 2019.

A.2.1 Motivation

Additional bar measurements were motivated by the following:

1. To measure material properties of the aluminum stock material to be used for bar fabrication (see Section [4.5](#)), rather than relying on standard published values.
2. To assess the performance of fabricated bar prototypes.
3. To provide damping parameters inputs for the hybrid radiation model in Chapter [5](#).

The decision was made to determine damping parameters from bar measurements after a literature search yielded no suitable results.

A.2.2 Scope

Aluminum stock material used for bar fabrication was purchased in four pieces. These pieces were originally a single, long bar before being cut by the supplier. As all four stock bars were cut from the same original piece, and since aluminum is a relatively consistent manufactured material, measurements were performed on just one of the four stock bars.

Both prototype bars described in Section [4.5](#) were measured to assess their performance.

Data from all of the rosewood bar measurements from Section [A.1](#) was reanalyzed to determine damping parameters.

A.2.3 Method

Measurements were performed in the same manner described in Section [A.1.3](#). For the stock aluminum bar, a finite element model estimated the location of nodes of the fundamental mode. Foam support blocks were placed under these locations for the measurements.

Measurements of stock bar mass and dimensions were taken and used to estimate material density. Elastic properties of the aluminum stock bar were determined by converging the finite element model of the stock bar, using the estimated density, and adjusting the input elastic modulus until the predicted fundamental frequency matched that from measurements.

Table A.1: Material properties of aluminum used for bar fabrication compared with published standard values. Density, ρ , was measured for the stock material, as was Young’s modulus, E . Poisson’s ratio, ν , is taken as the standard published value for both cases. Shear modulus, G , is calculated from E and ν for isotropic materials. Standard aluminum properties are those of the 6061 alloy (Battelle Memorial Institute, 2020).

Property	ρ [kg/m ³]	E [GPa]	ν	G [GPa]
Measured Value	2693	69.39	0.33	26.09
Standard - 6061 Alloy	2700	68.90	0.33	25.90

Poisson’s ratio was set to the standard published value (Battelle Memorial Institute, 2020) in this procedure.

To determine damping properties for both the aluminum stock bar and previous marimba bar measurements, data was (re)analyzed using the *Least Squares Complex Exponential* (LSCE) method (Brown et al., 1979; Avitabile, 2018). This method represents frequency response functions as the sum of decaying sinusoids (a Prony model), and thus provides damping terms for individual modes.

The damping terms and associated modal frequencies calculated by the LSCE method were used as data describing damping as a function of frequency for each material. Damping models used in Chapter 5 were determined by fitting Rayleigh damping models to each material data set in a least squares manner.

A.2.4 Results

Aluminum material properties for the measured stock material are shown in Table A.1. Density and elastic modulus are based on bar measurements. Poisson’s ratio is taken as the standard published value, while shear modulus is calculated from the measured elastic modulus and published Poisson’s ratio.

Modal frequencies for fabricated bar Prototype 2 are provided in Table 4.7, along with their expected values from the converged finite element model.

Figure A.3 shows measured damping ratios and associated modal frequencies, along with the best-fit Rayleigh damping model for rosewood. Similar data is shown in Figure A.4 for aluminum. Corresponding Rayleigh damping parameters for aluminum and rosewood are shown in Table 5.2.

A.2.5 Discussion

Looking at Table A.1 the measured values of density, ρ , and elastic modulus, E , are quite close to the published standard values for the 6061 alloy. This is as expected for a manufactured material like aluminum, and provides confidence in the measured values.

A few interesting observations can be made from Figure A.3. Notably, damping ratios for modes V1, V2 and V3 are generally consistent. Mode V1 values are clustered around 186 Hz, while mode V2 data points are clustered, rather tightly, around 744 Hz. Mode V3 data forms a group around 1840 Hz, with considerably more spread in frequency values (as is expected based on Figure A.2). The fitted Rayleigh damping model shows generally good agreement with these mode types and most data in Figure A.3. The sole exception is the group of data around 300 to 500 Hz, which represent mode T1. Damping ratios for these modes lie noticeably above the Rayleigh curve. Whether this was caused by some aspect of the measurements (perhaps the foam blocks supporting the bar), or is simply an aspect of damping not adequately captured by a Rayleigh model, is difficult to say.

Looking at Figure A.4 the fitted Rayleigh damping model shows generally good agreement with measured damping ratios for the aluminum stock bar. The sole exception to this may be the data point around 5400 Hz with a damping ratio around 0.0008. This point represents mode T3 of the stock bar. Noting the different vertical axis scales in Figures A.3 and A.4, the aluminum damping ratios are observed to be an order of magnitude smaller than those for rosewood. This is evident in Figure 5.6 where the two damping models are plotted together. Given how long a struck aluminum bar's sound can sustain, compared to that of a rosewood bar, the much lower damping of aluminum conforms to expectations.